

Cyber-Physical Security for Additive Manufacturing Systems

Logan Sturm

Dissertation submitted to the faculty of the Virginia Polytechnic Institute and State University in
partial fulfillment for the degree of

Doctor of Philosophy
In
Mechanical Engineering

Christopher B. Williams (Chair)
Jaime A. Camelio
Pablo A. Tarazaga
Xiaoyu (Rayne) Zheng

October 9th, 2020
Blacksburg, VA

Keywords: Additive Manufacturing, Cyber-physical Security, Side-channels, In Situ Monitoring

Copyright

Cyber-Physical Security for Additive Manufacturing Systems

Logan Sturm

ACADEMIC ABSTRACT

Additive manufacturing (AM) is a growing section of the advanced manufacturing field and is being used to fabricate an increasing number of critical components, from aerospace components to medical implants. At the same time, cyber-physical attacks targeting manufacturing systems have continued to rise. For this reason, there is a need to research new techniques and methods to ensure the integrity of parts fabricated on AM systems. This work seeks to address this need by first performing a detailed analysis of vulnerabilities in the AM process chain and how these attack vectors could be used to execute malicious part sabotage attacks. This work demonstrated the ability of an internal void attack on the .STL file to reduce the yield load of a tensile specimen by 14% while escaping detection by operators.

To mitigate these vulnerabilities, a new impedance-based approach for in situ monitoring of AM systems was created. Two techniques for implementing this approach were investigated, direct embedding of sensors in AM parts, and the use of an instrumented fixture as a build plate. The ability to detect changes in material as small as 1.38% of the printed volume (53.8 mm³) on a material jetting system was demonstrated.

For metal laser powder bed fusion systems, a new method was created for representing side-channel meltpool emissions. This method reduces the quantity of data while remaining sensitive enough to detect changes to the toolpath and process parameters caused by malicious attacks. To enable the SCMS to validate part quality during fabrication required a way to receive baseline part quality information across an air-gap. To accomplish this a new process noise tolerant method of cyber-physical hashing for continuous data sets was presented. This method was coupled with new techniques for the storage, transmission, and reconstructing of the baseline quality data was implemented using stacks of “ghost” QR codes stored in the toolpath to transmit information through the laser position.

A technique for storing and transmitting quality information in the toolpath files of parts using acoustic emissions was investigated. The ATTACH (additive toolpath transmission of acoustic cyber-physical hash) method used speed modulation of infill roads in a material extrusion system to generate acoustic tones containing quality information about the part. These modulations were able to be inserted without affecting the build time or requiring additional material and did not affect the quality of the part that contained them.

Finally, a framework for the design and implementation of a SCMS for protecting AM systems against malicious cyber-physical part sabotage attacks was created. The IDEAS (Identify, Define, Establish, Aggregate, Secure) framework provides a detailed reference for engineers to use to secure AM systems by leveraging the previous work in vulnerability assessment, creation of new side-channel monitoring techniques, concisely representing quality data, and securely transmitting information to air-gapped systems through physical emissions.

Cyber-Physical Security for Additive Manufacturing Systems

Logan Sturm

GENERAL ABSTRACT

Additive manufacturing (AM), more widely known as 3D printing, is a growing field of manufacturing where parts are fabricated by building layers of material on top of each other. This layer-based approach allows the production of parts with complex shapes that cannot be made using more traditional approaches such as machining. This capability allows for great freedom in designing parts, but also means that defects can be created inside of parts during fabrication. This work investigates ways that an adversary might seek to sabotage AM parts through a cyber-physical attack.

To prevent attacks seeking to sabotage AM parts several new approaches for security are presented. The first approach uses tiny vibrations to detect changes to part shape or material by attaching a small sensor either directly to the parts or to the surface that they are built on. Because an attack that sabotages an AM system (3D printer) could also affect the systems used to detect part defects these systems should be digitally separated from each other. By using a series of QR codes fabricated by the AM system along with the parts, information can be sent from the AM system to the monitoring system through its sensors. This prevents a cyber-attack from jumping from the AM system to the monitoring system. By temporarily turning off the laser power and tracking the movements of the guiding mirrors the QR code information can be sent to the monitoring system without having to actually print the QR code. The information stored in the QR code is compared to the emission generated when fabricating the parts and is used to detect if an attack has occurred since that would change the emissions from the part, but not from the QR code.

Another approach for sending information from the AM system using physical emissions is by using sounds generated during part fabrication. Using a desktop scale 3D printer, the speed of certain movements was increased or decreased. The change in speed causes the sound emitted from the printer to change, while not affecting the actual quality of the print. By using a series of tones, similar to Morse code, information can be sent from the printer. Research was performed on the best settings to use to transmit the information as well as how to automatically receive and decode the information using a microphone.

The final step in this work is a framework that serves as a guide for designing and implementing monitoring systems that can detect sabotage attacks on AM parts. The framework covers how to evaluate a system for potential vulnerabilities and how to use this information to choose sensors and data processing techniques to reduce the risk of cyber-physical attacks.

Table of Contents

1. Introduction and Motivation	1
1.1. Vulnerabilities of CPS	1
1.2. Introduction to Additive Manufacturing (AM) as a Cyber-Physical System.....	3
1.3. Cyber-physical security	4
1.4. Need for Quality Control and Security in AM.....	4
1.5. Research Gap	5
1.6. Research Questions	6
1.7. Roadmap	9
2. Literature review	10
2.1. AM Research	10
3. Cyber-physical vulnerabilities in additive manufacturing systems: A case study attack on the .STL file with human subjects.....	12
3.1. Introduction	12
3.1.1. Attacks on Cyber-Physical Systems.....	12
3.1.2. Cyber-Attacks on Manufacturing.....	13
3.1.2. Context: Cyber-Physical Vulnerabilities in Additive Manufacturing.....	13
3.2. Cyber-Vulnerabilities in the Additive Manufacturing Process.....	14
3.2.1. CAD Model	15
3.2.2. STL/.AMF File	15
3.2.3. Toolpath File	16
3.2.4. Physical Machine.....	17
3.2.5. Summary	18
3.3. Case Study: Cyber-Physical Attack on AM Systems via Altering .STL Files	19
3.3.1. Types of .STL Attacks.....	19
3.3.2. Void Attack Considerations.....	22
3.3.3. Attack Embodiment	23
3.4. Effects of .STL Void Attack	24
3.4.1. Effect of Voids on Part Strength.....	25
3.4.2. Human Subjects Experimentation	26
3.5. Discussion and Recommendations	29
3.6. Closure	31

4. In-situ Monitoring of Additive Manufacturing Processes via Impedance-based Measurements	34
4.1. Introduction	34
4.1.1 In situ monitoring of AM	35
4.1.2 Impedance-based monitoring	36
4.2 EXPERIMENTAL METHODS	38
4.2.1 Method Overview	38
4.2.2 Material Jetting Process	39
4.2.3 Materials	39
4.2.4. Test specimen design	40
4.2.5 Impedance measurements and analysis	41
4.3 RESULTS	42
4.3.1 Embedded Piezos	43
4.3.2 Fixture-based Piezos	45
4.4 DISCUSSION	47
4.5 CONCLUSIONS	48
5. In-situ Detection of Build Tampering in Metal Additive Manufacturing Using a Cyber-physical Hash ..	54
5.1. Cyber-physical security in additive manufacturing and the use of side-channel monitoring	55
5.2. Side-channel monitoring system and transmission approach	58
5.2.1 Toolpath representation of data	59
5.2.2. Frequency representation and comparison of toolpath data	60
5.2.3. Build validation using cyber-physical hashing	62
5.2.4. Distribution of intensity data	64
5.2.5. QR code extraction	64
5.3. Experimental	66
5.3.1. Laser powder bed fusion and side-channel monitoring system setup	66
5.3.2. Build setup	67
5.4. Results and Discussion	69
5.4.1. Toolpath Comparison	69
5.4.2. Intensity Comparison	75
5.4.3. Mapping frequency representation into predefined ranges	77
5.4.4. QR Code Extraction	78
5.5. Conclusions	81

5.6. References	82
6. ATTACH: Additive Toolpath Transmission of an Acoustic Cyber-Physical Hash	85
6.1. Cyber vulnerabilities and process monitoring in additive manufacturing systems.....	85
6.1.1. Ensuring quality and security of cyber-physical manufacturing systems	85
6.1.2. In situ monitoring via side-channel measurement systems in AM.....	86
6.1.3. In situ transmission of quality data.....	87
6.2. Acoustically Transmitted Cyber-Physical Hash	89
6.2.1. Velocity modulation as a transmission method	89
6.2.2. Selecting toolpath roads to modulate	90
6.2.3. Insertion of modulation into the toolpath.....	93
6.2.4. Error correction.....	94
6.2.5. Acoustic Processing.....	94
6.3. Experimental methods.....	95
6.3.1. Physical System	96
6.3.2. Tone parameters effect on transmissibility	96
6.3.3. Demonstrate the ability to successfully transmit a message stored in the toolpath to the SCMS	97
6.3.4. Frequency Modulation Effect on Part Strength	98
6.4. Results and Discussion	98
6.4.1. Effect of parameters on transmissibility.....	98
6.4.2. Message Transmission	101
6.4.3. Effect on tensile strength.....	102
6.5. Conclusions	103
6.6. References	105
7. IDEAS (Identify, Define, Establish, Aggregate, Secure): A cyber-physical framework for securing additive manufacturing systems using physical side-channels	111
7.1. The need for side-channel monitoring to secure additive manufacturing systems	111
7.2. Identifying Attack Vectors.....	116
7.2.1. Identifying critical part properties	117
7.2.2. Establishing tolerances for part properties	117
7.2.3. Identify high-level process overview (mass and energy inputs).....	117

7.2.4. Identify process/machine specific implementation and correlation to geometry (extrinsic) and parameter (intrinsic) effects on part properties.....	120
7.2.5. Quantify relationships between process settings and part properties to determine maximum allowable variation.....	121
7.2.6. Define attack detection threshold(s)	122
7.3. Defining Side-Channels	126
7.3.1. Monitoring approaches.....	127
7.3.2. Sensor requirements.....	128
7.4. Establish Baselines	130
7.4.1. Direct comparison.....	132
7.4.2. Machine learning	132
7.4.3. Physics-based modeling/Digital twin	132
7.5. Aggregate Datasets	133
7.5.1. Downsampling	133
7.5.2. Compression	133
7.5.3. Hashing.....	133
7.5.4. Control Charting.....	134
7.5.5. Data Fusion	134
7.6. Secure Transmission	135
7.6.1. Security Approaches	136
7.7. Conclusions	140
7.8. References	141
8. Conclusions and broader impacts.....	152
8.1. Summary of research.....	152
8.1.1 Objective 1.	153
8.1.2. Objective 2	154
8.1.3. Objective 3	155
8.1.4. Objective 4	156
8.1.5. Objective 5	157
8.2. Limitations and future work	157
Publications.....	159
Research Contributions.....	160

Broader impacts..... 161
References 162

List of Figures

Figure 1.1. Additive Manufacturing Process Chain.....	4
Figure 3.1. Additive Manufacturing Process Chain.....	14
Figure 3. 2. Configuration data for a Material Jetting AM system intercepted over WiFi using network protocol analyzer software.....	18
Figure 3.3. .STL Attacks on an ASTM D638 tensile specimen. A) an unaffected dogbone; B) a scaled down dogbone; C) an indentation; D) a protrusion; E) a vertex moved inward; F) a vertex moved outward.....	21
Figure 3.4. Process flow of void placement algorithm; 1) Void is fully encapsulated inside of the part; 2) The void is positioned in a location that is more likely to cause a failure; 3) The void is scaled either up or down to increase the chance of failure or decrease the odds of detection.....	24
Figure 3.5. Cross sectional slice of a dogbone infected with a void.....	25
Figure 3.6. Von Mises Stress of a dogbone infected with a void.....	25
Figure 3.7. On Left: Uninfected dogbones breaking at the gauge section. On Right: Infected dogbones breaking at the void location within the specimen neck.....	25
Figure 3.8. Load and strain data of parts with and without voids.....	26
Figure 4.1. A) Example of embedded sensor method B) Example of fixture-based sensor method.....	39
Figure 4.2. Test parts A) Control sample “A”, B) Triangular prism cavity “B” (lines indicate defect dimensions).....	40
Figure 4.3. A comparison of multiple signatures taken at the same layer from the same part using an embedded sensor.....	43
Figure 4.4. A) Difference comparison for embedded piezos between baseline, control samples, and defect samples across each layer. The defect is introduced in layer 151 (as indicated by the red line) and detectable in all samples at layer 210 (red oval). The defect size when detected is 95.6 mm ³ (2.28% of printed volume). B) The part signatures at layer 210. The red oval indicates the area where the defect is evident.....	44
Figure 4.5. A) Difference comparison for embedded piezos between baseline, control samples, and defect samples across each layer. The defect is introduced in layer 150 (as indicated by the red line) and detectable in all samples at layer 245 (red oval). The defect size when detected is 53.8 mm ³ (1.38% of printed volume). B) The part signatures at layer 195. The red oval indicates the area where the defect is evident. White areas indicate frequency ranges that were used in analysis.....	46
Figure 5.1. Cyber-physical hash overview: Using an air-gapped SCMS to validate AM parts by incorporating a data package in the toolpath or model file to send information using physical emissions during the fabrication process.....	57

Figure 5.2. Examples of different cases for the SCMS where the part/data package are modified/unmodified. The manufacturer needs to ensure that an attacker cannot alter the data package to match a modified part or the system could reject or pass the part.....59

Figure 5.3. Example of how infill frequency representations change with different toolpaths. The as-built scanlines, sequentially ordered scanlines, and the corresponding frequency representations of the scanlines are shown for three different toolpaths (normal, offset, and gap).....61

Figure 5.4. Example of monitoring system for metal powder bed fusion. Emissions from the melt pool are captured by a pair of photodiodes and synchronized with the position of the laser read off of the galvos in the scanner.....67

Figure 5.5. External geometry change of tensile test specimen compared to control specimen. The curvature at the neck of one side of the specimen has been reduced from a radius of 6 to a radius of 4.5.....68

Figure 5.6. QR code stack for transmitting the data package to the SCMS through physical emissions from the build. QR codes pictured are spaced at intervals of each five layers with a total of six QR codes per stack.....68

Figure 5.7. First test build layout consisting of 11 parts along with the process parameters used for each part.....69

Figure 5.8. Frequency representation of laser toolpath changes throughout build 1, demonstrating unique patterns for different parts.....70

Figure 5.9. Part frequency signatures isolated and overlaid for tensile test specimens for later 68. The blue dotted lines show the control samples and the red dashed line shows the sample with altered geometry. The location where the curvature is altered can be clearly seen as a slightly lower height in the curve.....71

Figure 5.10. Part frequency signatures of different tabs overlaid to demonstrate the signature variations caused by process parameter and geometry changes. The signature representation allows for a visual inspection of part variation as well as the ability to identify the type of defect based.....71

Figure 5.11. Frequency comparison of toolpaths for tab test parts during fabrication. Green/yellow indicate the greatest difference from the control values. The first 33 layers are support material and have identical process parameter and toolpath settings (overriding part specific settings). The layers where the void occurs are highlighted with a red outline.....73

Figure 5.12. Frequency comparison of control tabs for Build 2. Tab 2 has the same geometry, but a different toolpath on layers 1-9 which was not part of the designed experiment, but was detected by the proposed method.....74

Figure 5.13. Example of different toolpath for Tab 2 (taken from layer 3 of build 2). Dark blue colors indicate low intensity points while brighter green/yellow points indicate high intensity points.....74

Figure 5.14. Example of different toolpath signature for Tab 2 (taken from layer 3 of Build 2). Tabs 1 and 3 have two scan blocks (corresponding to their two signature peaks). Tab 2 has three scan blocks (corresponding to three signature peaks).....75

Figure 5.15. Intensity comparison of each part in build 1. Green/yellow values indicate larger changes in average intensity. Some slight build location dependency can be seen from layers 1-33, when the settings for the parts are identical. Once the process parameters are changed, the increased power creates a noticeable difference. The increase in scanning speed and changes to the toolpath from the void defect also causes a detectable change in the average intensity.....76

Figure 5.16. Intensity comparison from Build 2. The third test specimen exhibits significantly higher average intensity from between layer 34 and layer 66. It is unclear what changed in the system to cause this difference.....77

Figure 5.17. An example of mapping the control points of the frequency representation of the toolpath into predefined ranges and how combinations of adjacent bins can be used to increase the robustness of the system against process noise. Expanded views show the location of control points for different tabs and how they map into the predefined ranges. Green in the target range, yellow indicates an edge (one bin adjacent on a single axis) and orange indicates a corner (one bin adjacent on two axis). Red values show control points that are at least two bins away from the target value.....78

Figure 5.18. Printed QR code stack in print bed (A) and after fabrication (B). The poor contrast between the light and dark cells in the QR code make it difficult to extract the stored information using a camera.....79

Figure 5.19. Scan line representation of QR code, a) jump lines included b) jump lines removed and scan lines filtered.....79

Figure 5.20. Extracting QR codes from scan lines using overlapping gaussian distributions. A) Threshold value set too low, resulting in over detection B) threshold value set acceptably C) threshold value set too high, resulting in under detection.....80

Figure 5.21. Extracted “ghost” QR code containing the hex string: f8016e158dfe9280829f052c2018fd9280

Figure 6.1. Use of a data package stored in the model file/toolpath of a part to transmit quality information to an air-gapped SCMS.....88

Figure 6.2. Example of frequency modulation of toolpath infill.....90

Figure 6.3. Histograms comparing the length of 45° infill roads in different part geometries. A simple 50mm cube has a flat distribution of road lengths. A tensile test specimen (165mm x 19mm x 3.2mm) has a spike corresponding to the road length in the gauge section. The Stanford Bunny has an uneven distribution of road lengths, with a larger cluster of long roads in the body of the bunny (50mm x 38.7mm x 50mm).....91

Figure 6.4. Theoretical maximum amount of transmittable data for infill modulation of an ASTM D638-14 Type 1 tensile test specimen using a single level of modulation. As tone time or feed rate increase, the number of roads long enough to hold data decrease.....92

Figure 6.5. Modulated toolpath comparison between 0.5s tone (a) and 1.0s tone (b) for type 1 ASTM tensile test specimen with a 45°. The shorter tone extends into the roads in the gauge section of the tensile test specimen, while the 1.0s tone is limited to the longer roads in the neck and grip sections.....93

Figure 6.6. Modulated toolpath of a tensile test specimen. Light blue represents standard infill, dark blue is the slow modulation speed, yellow is the fast modulation speed, and red is the contour/travel road speed. Note that contour/travel movements have been reduced in this image for better color scaling.94

Figure 6.7. Message representation in ASCII and reed-solomon encoded binary with 5-bytes of error correction. The total length of the encoded message is 20 bytes (160 bits).....97

Figure 6.8. Overlaid spectrograms of modulated toolpath data for 0.5s, 1.0s, and 1.5s tones with 20Hz modulation.....99

Figure 6.9. Overlaid spectrograms of modulated toolpath data for 1.0s tones, 20Hz modulation.....100

Figure 6.10. Overlaid spectrograms of modulated toolpath data for 1.0s tones, 100Hz modulation. Note: due to the large modulation the frequencies overlap, so they have been separated for clarity.....100

Figure 6.11. Comparison of tone transmission between odd layers (A) and even layers (B). When both stepper motors are moving in the same signed direction the tones are louder. When the stepper motors are moving in opposite signed directions the resulting tones are quieter.....101

Figure 6.12. Example of message extraction from acoustic emissions. The colored area in the middle is the spectrogram representation of the low modulation. The orange bars at the bottom are the corresponding spectral centroid results (shifted down for visibility) and the ones and zeros are the message bits being interpreted. “01100010100011000100110000001100” = “F1”102

Figure 6.13. Tensile test specimens before and after testing. (A) modulated specimens, (B) control specimens. There is a minor visual difference between the modulated samples and the control samples.

Figure 7.1. Cyber-physical hash overview: Using an air-gapped SCMS to validate AM parts by incorporating a data package in the toolpath or model file to send information using physical emissions during the fabrication process [22].....113

Figure 7.2. Different domains of additive manufacturing based on the part geometries and process parameters used. Fixed geometries and settings are easier to establish baselines for while unique parts with in-situ adjustments are much harder to validate.....114

Figure 7.3. Overview of IDEAS framework.....116

Figure 7.4. Fishbone diagram of the FDM process and how different process parameters affect various part properties [53].....121

Figure 7.5. Fishbone diagram of a laser powder bed fusion system.[54]121

Figure 7.6. Cross sectional slice of a dogbone infected with a void.....125

Figure 7.7. Examples of different types of geometric and parameter changes for AM parts fabricated on one or more machines. Establishing a baseline needs to start by comparing fixed settings and geometries and then build up to more complex changes and interactions.....131

Figure 7.8. Ways of transmitting build information to a SCMS [22].....135

Figure 7.9. Flowchart for selecting side-channels for use with a data package, stored in the toolpath, that contains part quality information for validating the build against sabotage attacks.....136

Figure 7.10. Examples of different cases for the SCMS where the part/data package are modified/unmodified. The manufacturer needs to ensure that an attacker cannot alter the data package to match a modified part or the system could reject or pass the part.....137

List of Tables

Table 3.1. Group caliper measurements of tensile test specimens.....	28
Table 4.1. Part layers where measurements were taken. For each layer, the volume of model material in each part is shown along with the size of the defect, and the percentage of the printed material that the defect represents.....	42
Table 6.1. Test Parameters used and corresponding road distance of tones. Combinations of feedrates and tone times that cause the message length to exceed ~24mm were too long to fit within the test specimen and were excluded.....	97
Table 6.2. Tensile test results for control and modulated ABS parts fabricated using desktop additive manufacturing systems.....	103
Table 7.1. Overview of AM process types and the physical material and energy inputs into the system along with their dimensionality (i.e. voxel, pass, layer, or volume).....	118
Table 7.2. Example attacks and their corresponding risk evaluation.....	125

1. Introduction and Motivation

“Cyber-physical systems (CPS) are smart systems that include engineered interacting networks of physical and computational components” [1]. The coupling of these two domains is integral in the designed function of these systems and often provides new modes of user interaction and enhanced functionality. Cyber-physical systems are present in a wide range of areas, such as manufacturing, transportation, and the smart grid. Other growing areas for CPS are in medical devices and consumer electronics where “smart” devices like the Nest thermostat are becoming increasingly popular [2].

One of the most important CPS is the electric grid, where power plants, relay stations, and transmission lines comprise the physical component and the digital communication and supervisory control and data acquisition (SCADA) system comprise the cyber component [3]. Another CPS is cars (particularly self-driving). Modern cars increasingly incorporate more cyber components for sensors, control, and user feedback. These cyber features control a wide range of physical components such as fuel injectors, airbag controllers, and climate control [4]. Advanced features such as driver assistance tools and self-driving cars are only increasing this trend towards cyber-physical coupling.

Manufacturing continued to expand its use of CPS. At a low level, numerical controllers are now widely used for things like temperature controllers, computer-numerical-control (CNC) machining, robotic arms, and additive manufacturing (AM). At this level, they provide the advantage of precise and repeatable control, without the need for constant manual input. This allows products to be created more quickly and more repeatable than traditional approaches. At a higher-level, CPS are used in manufacturing to coordinate multiple systems to work together efficiently. The concept of “Industry 4.0” is built around the idea of low-level CPS being networked together to create a much larger CPS that is able to efficiently schedule work and shift loads around a network. With this large-scale inclusion, manufacturers will be better able to utilize equipment and to adapt to changes in load or demand. The competitive benefits of CPS will continue to drive further adoption and integration of these products into manufacturing.

The benefits of CPS systems also come with additional risks. Because physical systems are now coupled with digital ones any digital attack has the potential to affect the physical world. This means that physical security, such as key access, is no longer sufficient to protect physical assets and production. Manufacturers must now contend with bad actors that may be located hundreds or thousands of miles away [5–8].

1.1. Vulnerabilities of CPS

The coupling of cyber and physical systems to create CPS creates significant benefits; however, it also creates a number of new vulnerabilities. With the growth of the Internet of Things (IoT) the number of CPS systems on networks continues to increase [9]. Concurrently, cyber-attacks have become more prevalent in recent years, increasing in maliciousness and decreasing in visibility [9–11]. This poses a significant issue, as cyber-attacks on cyber-physical systems could result in damage to the machines themselves or the humans who interact with them.

A prominent example of a cyber-physical attack was the Stuxnet worm that targeted Iranian centrifuges used for refining uranium. In this attack, the worm was able to infect the software system and affect the physical hardware, causing damage to the centrifuges. By sending false data back to the operators, Stuxnet was able to make it appear as though the centrifuges were operating correctly, while it caused them to damage themselves. The ability of Stuxnet both to cause damage to physical systems and to hide itself illustrates the ability of a cyber-physical attack to disrupt manufacturing systems and the need for physical methods of detection[12].

Another example of a cyber-physical attack is the hijacking of insulin pumps. In this case a hacker is able to connect to a Bluetooth enabled insulin pump to control the dose of insulin given to the wearer. By increasing or decreasing the dosage of insulin, it is possible to cause serious injury or even death in the user. The currently security system for these pumps is insufficient to prevent a cyber-physical attack that could have potentially lethal consequences[13].

The previously mentioned examples demonstrate the ability for cyber-attacks to cross over into the physical world. Attacks on CPS are even more alarming when considering the ever-increasing amount of networked devices that are being connected to machines in the manufacturing world. A cyber-attack on these machines could cause injury to plant workers and damage to the machine itself. This has already been demonstrated at a German foundry wherein a cyberattack on a blast furnace's control systems prevented a safe shut down and resulted in "massive damage" [14,15]. Perhaps even more insidiously, an attack could be designed to cause a process to produce faulty parts that might find their way into end-user products[16]. For example, an attack could be designed to affect the production of a jet turbine part such that it would pass inspection but fail during operation and cause significant damage.

An even more recent example of a cyber-physical attack that had a real-world effect on consumers was an attack in December 2015 on the power system in the Ukraine. This attack affected approximately 250,000 households with power blackouts. The attack used BlackEnergy malware along with KillDisk to gain remote access to the power grid and cut off power [17]. For one power producer, Kyivoblenergo, this attack led to the temporary disconnection of thirty different substations, affecting approximately 80,000 customers. In addition to the remote access component of the attack a simultaneous distributed denial of service (DDoS) attack was launched on the phone systems to attempt to prevent customer reports of power outages [18].

With the rise in both the number of CPS connected to networks and in malicious cyber-attacks, there is a clear need for research to understand the vulnerabilities of cyber-physical systems. Recent work on the current status of cyber-physical systems has identified security as a major issue [19]. Additionally, there is a need to identify the skills that are necessary to effectively operate a cyber-physical manufacturing environment [20], which includes skills pertaining to the identification and prevention of attacks on manufacturing systems.

Related research at Virginia Tech has shown that attacks on subtractive manufacturing can have a demonstrable effect on the part being produced [21,22]. In this study students were asked to fabricate

a tensile test specimen as part of a manufacturing lab. Unknown to the students, the toolpath of the tensile test specimen was modified to reduce the cross-sectional area of the specimen by 19% corresponding to a 19% reduction in performance. This research then evaluated the students' response to the attack. [21,22]. In this work the authors demonstrated a toolpath attack by manually modifying and replacing machining toolpath (e.g., GCODE). This work was limited to modifying the external geometry of a part in a way that affected the part strength, but was detectable using common inspection techniques. As such, the authors have investigated cyber-physical vulnerabilities in manufacturing systems. The authors have demonstrated such attacks on subtractive manufacturing, and have found that they can have a demonstrable effect on the part being produced. The National Defense Industrial Association (NDIA) has also published a white paper investigating the needs and challenges of manufacturing security [23].

1.2. Introduction to Additive Manufacturing (AM) as a Cyber-Physical System

Additive manufacturing (AM) refers to a variety of manufacturing processes that allow for the fabrication of an arbitrary geometry through the selective placement and/or fusion of material in a layer-wise fashion. AM offers several advantages over traditional manufacturing. First, AM does not require retooling or refixturing to fabricate new parts and designs. The ability of AM to rapidly and inexpensively transition from design to fabrication is a key advantage. A second advantage of AM is the ability to easily fabricate complex geometries that traditional manufacturing cannot fabricate[24]. AM also allows for multiple parts to be fabricated as a single assembly, reducing the number of components and reducing the need for costly and time-consuming assembly processes.

The advantages of AM are achieved through it being a strongly coupled cyber-physical system. The first step in the AM process chain (shown in Figure 1.1) is the design of the model geometry; this can be done through a variety of processes such as using CAD software or 3D scanning techniques. Once the model has been created it is converted to the .STL file, the standard file format for AM systems. This .STL file is then fed through machine specific CAM software to generate a toolpath file that is sent to the machine. The AM system then fabricates the model and the completed part is removed, post processed if necessary, and inspected, often using the original model as a digital reference. The digital transfer of the model throughout the AM process is essential to achieving the flexibility and complexity that AM systems are capable of producing.

The physical components of an AM system consist of the material that the part is fabricated out of, the components of the machine that is fabricating the part, any sensors that are included or connected to the system, and the part itself. The cyber components consist of the digital model file, the .STL file, the toolpath file, and any internal computers or controllers in the system. AM systems are able to affect not only part geometry but also its material properties by altering the printing process parameters or toolpath. This leads to a strong coupling between the cyber and the physical in AM systems[25,26].

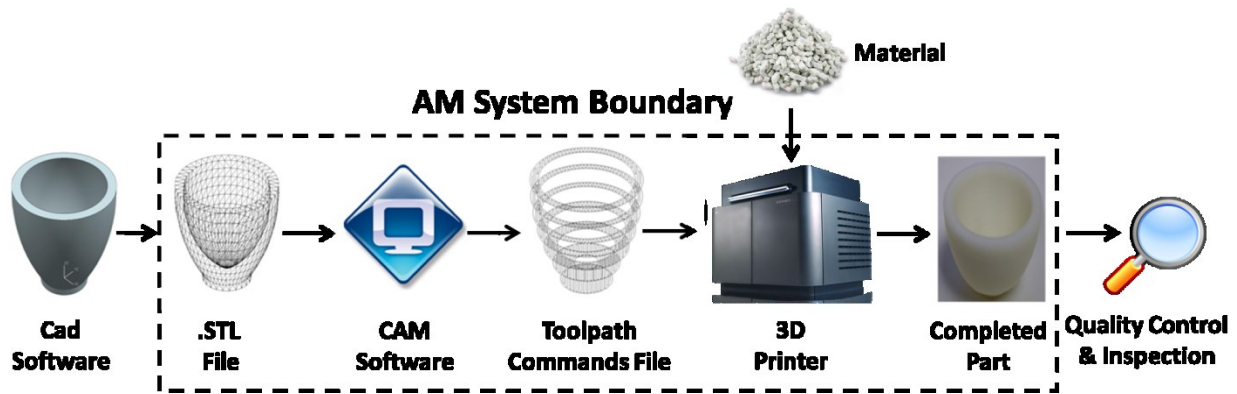


Figure 1.1. Additive Manufacturing Process Chain

1.3. Cyber-physical security

As a cyber-physical system, AM needs security against malicious attacks. Traditionally security has focused on the two domains discretely, physical systems being protected by physical security measure such as locks and building access while cyber systems are protected by cyber security measures such as firewalls and antivirus software. With the development of CPS, the potential for cross-domain attacks emerges. An attack in the cyber domain may cause a physical impact such as a robotic arm damaging physical hardware or an electronically secured door being unlocked. Attacks in the physical domain can also have effects in the cyber domain such as a thermal sensor being manipulated to cause the control software to push the system out of the correct temperature range.

Because of the close integration that exists in cyber-physical systems it is necessary to consider the security of these systems holistically. Focusing on the two domains as separate entities leaves the potential for unexpected cross-domain attacks. In addition to this, by looking at security as a whole it is possible to more efficiently identify and protect attacks. Cyber-attacks may leave physical traces and physical attacks may leave cyber traces. By looking at both domains together the chancing of detecting and preventing attacks increases. As security approaches start to focus on the physical domain in addition to the cyber domain, they begin to overlap with the role of quality control systems.

1.4. Need for Quality Control and Security in AM

The maturation and growing adoption of additive manufacturing (AM) as a means for making end-use products has led to an increased need for part traceability and process monitoring [27]. Unlike prototypes or fixtures, end-use products must often meet rigorous performance standards and be certified before they can be put into use. In particular, the aerospace industry has begun to implement AM technologies for fabricating high value, low volume parts. Some examples of AM parts being produced and tested include direct printed metals (e.g., Rolls Royce Trent-XWB bearing [28]) and polymers (e.g., FAA-approved ULTEM 9085 aircraft air duct created by Stratasys and Orbis [29]).

In a process such as CNC machining it is possible for an attack to alter the toolpath and affect the geometry and surface finish of a part. In AM however, the layer-based deposition of materials

means that it is possible for an attack not only to affect the external geometry and surface finish, but also to cause internal defects in parts. Unlike external defects that can be detected through measurement systems or visually, internal defects are much harder to detect once they have been enclosed in the part, requiring more advanced and expensive detection techniques such as x-ray computed tomography (CT) scans. In addition to geometric defects, AM also presents the unique opportunity to selectively alter the material properties inside the part. In AM the fabrication process affects both the material properties and geometry. Effects such as the localized heating and cooling of sections of a part are dependent not only on the overall geometry of the part, but also on the toolpath and parameters that are used to fabricate that geometry. By altering the toolpath or process parameters during fabrication it is possible to locally introduce defects such as weaker material properties in an AM part.

An additional feature of AM is the ability to fabricate complex geometries. While these designs can be functionally superior, the introduction of complexity such as lattice structures can make it hard to use traditional quality control measures to validate the part. Traditional quality control relies on the abilities to make dimensional measurements of important features, using tools such as calipers, gauges, or surface scanning to validate features. Parts, such as lattice structures, that contain features that are inaccessible to these tools are difficult to validate properly without resorting to more expensive and time-consuming methods such as CT scanning.

The process differences between AM and traditional mean that traditional quality control and security measures are insufficient for preventing and detecting attacks on AM systems. To address this, new quality control techniques are needed that can detect the defects unique to AM systems. While existing inspections such as CT scanning are able to detect some of these defects, it is both expensive and unable to detect all the types of possible defects that can be introduced in AM systems. Additionally, the post process nature of most inspection techniques means that even defects that are detected can lead to costly wastes of time and materials since a defect that occurs early in a part is not detected until after the entire part has been fabricated. While traditional quality control focuses on post-process inspection of parts, in-situ monitoring is desirable in AM for two reasons: to detect defects before they become internal to the part, and to save cost by detecting defects before the entire build has been completed.

1.5. Research Gap

While prior work has been done in both quality control and security, little work has focused on their intersection in AM. While existing work has looked at vulnerabilities in CPS such as the electric grid, little work has been done to investigate the potential for targeted attacks to cause bad parts. In AM, the dependence of the (physical) part properties on the (cyber) process parameters and toolpath makes it particularly important to understand the complex cyber-physical interactions. In order to develop effective mitigation techniques, it is vital to understand the vectors through which an attack could be carried out and to identify the approaches through which an attack might attempt to compromise part quality.

Once attack surfaces have been identified, it is still necessary to identify systems that are able to detect the physical effects of an attack. While significant work is being done to develop new monitoring systems for AM, there is still a need for non-destructive techniques that can validate not only the geometry, but also the material properties of a part. An ideal monitoring technique for AM would be able to monitor material properties and geometry layer-by-layer. This technique could be used to qualify parts and would be inexpensive and non-invasive. While current monitoring techniques are able to achieve some of these qualities, they are limited in others (e.g. optical layer-by-layer imaging cannot interrogate material properties throughout the part, x-ray CT scanning can volumetrically image a part, but is expensive).

The ability to detect the physical effects of an attack is irrelevant if the monitoring system being used to do so is compromised. While a substantial amount of work has looked at developing systems that can detect defects in AM systems, little to none has looked at integrating these sensors into the AM workflow in a secure way. Many approaches directly integrate the sensors into the printer software (often with the goal of closed-loop feedback). In this case compromising either the monitoring or the printer can result in the entire system being compromised. To more thoroughly secure the process against deliberate attacks, more work is needed to separate these two systems while still being able to communicate information.

1.6. Research Questions

The goal of this research is to identify unique vulnerabilities in AM systems and determine their effectiveness against human subjects and to research novel techniques for detecting these attacks on AM systems. This involved research on the use of impedance sensors, side-channel monitoring systems (SCMS), cyber-physical hashing, new data processing techniques, and new ways of representing and transmitting information through part toolpaths and physical emissions. To achieve this, five main thrusts are presented:

- 1) An examination of vulnerabilities in the AM process chain,
- 2) The investigation of a new in situ monitoring and testing technique for detecting material property or geometry changes,
- 3) The identification of new techniques for processing and analyzing data from side-channels,
- 4) The use of physical side-channels to securely transmit information to an air-gapped SCMS,
- 5) The creation of a framework for designing side-channel monitoring systems to mitigate the risk of cyber-physical attacks on AM systems.

The first thrust is achieved through a detailed analysis of the AM process chain, and a detailed case study looking at the .STL file as an attack vector (Chapter 3). The second thrust is achieved through the extending impedance-based monitoring techniques to validate AM parts. Specifically, the use of piezo-ceramic sensors for the in-situ detection of internal defects in AM parts is explored (Chapter 4). The third thrust is achieved through the use of a visually readable frequency response representation for

identifying defect types in a metal powder bed fusion (MPBF) system (Chapter 5). The fourth thrust is demonstrated on two representative systems, (MPBF) (Chapter 5) (through the use of QR code stacks) and material extrusion (Chapter 6) (through the use of toolpath modulation). The final thrust leverages the findings from the previous findings to provide a comprehensive framework for designing a SCMS and a detailed discussion of the considerations involved (Chapters 7).

Overall Goal	
To improve the security of AM systems by understanding the potential nature of attacks and by developing new techniques for in-situ monitoring to detect build defects as they occur.	
Objective 1	Research Question 1
To understand the vulnerabilities of AM systems	How can one execute an automated void attack with solely the information provided in a .STL file?
Objective 2	Research Question 2
To develop a non-destructive side-channel monitoring technique that can detect changes to material properties and geometry	How can one use impedance-based monitoring in-situ in AM systems to detect internal part defects?
Objective 3	Research Question 3
To utilize an air-gapped side-channel measurement system on powder bed fusion that can detect changes to process parameters and toolpath	How can one detect a geometry or process parameter attack on a metal powder bed fusion (MPBF) system using side-channel measurements?
Objective 4	Research Question 4
To research techniques for securely communicating a physical-hash to a side-channel measurement system	What factors are most important when communicating an acoustic physical hash to a side-channel measurement system?
Objective 5	IDEAS Framework
To create a framework for designing a side-channel monitoring system to mitigate cyber-physical attacks on AM systems	Identify vulnerabilities, define side-channels, establish baselines, aggregate data-sets, secure transmission

Hypothesis 1	Hypothesis 3
Triangle size and aspect ratio can be used to select areas and generate voids within these areas that are likely to reduce part strength.	Monitoring laser position and intensity can be used to detect both geometry and process parameter attacks on a (MPBF) system as well as to transmit quality information to a SCMS.
Sub Question 1.1	Sub Question 3.1
How do inserted voids affect part strength?	How do toolpath changes affect the laser velocity/intensity profile?
Sub Question 1.2	Sub Question 3.2
Can an automatically inserted void affect part strength enough to cause a failure?	How do process parameter changes affect the laser velocity/intensity profile?
Sub Question 1.3	Sub Question 3.3
Can operators detect a void attack?	How can laser galvo data be integrated into a physical hash for metal powder bed fusion?
Hypothesis 2	Hypothesis 4
Fabricating parts with attached piezo sensors will allow the in situ detection of internal cavities in AM parts.	Tone length, feed rate, and modulation amplitude will affect the transmissibility of the physical hash to the SCMS.
Sub Question 2.1	Sub Research Question 4.1
What is the sensitivity of embedded impedance monitoring to internal defects?	How much information does the hash need to hold?
Sub Question 2.2	Sub Research Question 4.2
What is the sensitivity of fixture-based impedance monitoring to internal defects?	How frequently does the hash need to be transmitted?
	Sub Research Question 4.3
	How sensitive to process factors is the transmission of the hash?

IDEAS Framework
Identify, Define, Establish, Aggregate, Secure
Identify
How can attack vectors be identified?
Define
How does a manufacture select appropriate side-channels?
Establish
What approach is used for determining the side-channel baseline?
Aggregate
How can side-channel data size be aggregated to reduce size and improve performance?
Secure
What steps can be taken to secure the transmission of quality information to the side-channel monitoring system?

1.7. Roadmap

In Chapter 1, an introduction to additive manufacturing as a cyber-physical system was presented along with a brief introduction to the concept of cyber-physical security. The need for quality control and security in AM was also presented along with an overall research goal and a list of research questions.

Motivated by the primary research goal, Chapter 2 will present a literature review of the research that has been done in the AM space on attacks and vulnerabilities of AM systems as well as some proposed mitigation techniques.

Next, Chapter 3 details the investigation of the AM process chain for attack vectors and vulnerabilities and looks specifically at a vector unique to AM systems, the internal void attack. This section examines the process an attack would go through in designing and executing an attack at looks at the response of operators to an attack occurring.

Chapter 4 presents the use of impedance-based monitoring as a novel method for the in-situ detection of build defects in AM systems. This section evaluates the technique for application to AM and compares the use of embedded sensors vs. fixture-based sensors for detecting internal defects.

Chapter 5 expands previous work on the cyber-physical hash (CPH) by incorporating it into a metal PBF system. This section examines how changes to geometry and process parameters affect the side-channel signatures on a metal PBF system and presents an approach for reducing the volume of side-channel data while maintaining detection ability. This section also presents an approach for transmitting information in a QR code without require the physical fabrication of the code.

Chapter 6 further explores the secure transmission of quality data to SCMSs through physical emissions. In this section, data is incorporated directly into the toolpath parts fabricated on a material extrusion system and transmitted through acoustic emissions using a “toolpath modulation” technique. The work investigates how modulation parameters affects data storage, transmissibility, and part quality.

Chapter 7 leverages the understanding gained from the work in the previous chapters to create the IDEAS (Identify, Define, Establish, Aggregate, Secure) framework for designing a SCMS to secure AM systems against malicious cyber-physical sabotage attacks. This section covers the full development chain from identifying attacks, to selecting side-channels, collecting and processing side-channel data, and securely transmitting quality information to an air-gapped side-channel monitoring system (SCMS). The goal is to provide a reference for researchers and manufacturers looking to develop SCMSs for AM security.

Chapter 8 provides a summary of the research accomplishments from this work as well as a discussion on the limitations and topics of future work in the area. Lists of publications produced directly or in relationship to this work, research contributions, and broader impacts of this work are also presented.

2. Literature review

At the inception of this research, little to no work had been done looking specifically at security in AM systems. Since the inception of this work, significant interest in security for AM systems has arisen. In section 1.2 and 1.3 an overview of some cyber-physical attacks that have occurred or been investigated and shown to be vulnerable was given. Section 1.5 highlighted some of the work that has been done on in-situ monitoring of AM. In this section the review is focused on the growing body of security research focused on AM. Most, if not all of this work has been developed since the presentation of the original work in this dissertation (Chapter 3)[30], showing how this work has helped to promote research in a new area.

2.1. AM Research

Prior to the initial research presented in this dissertation in 2014 little to no research had looked at the potential for attacking AM systems for the purpose of causing damage to physical systems [30]. While some early work considers the security of AM from an intellectual property standpoint [31], security is not considered to be of importance, “security and privacy issues are not a concern because 3-D printing enables in-house production.”[31]. Since 2014 a growing body of work on security in AM has emerged.

One of the first additional works of research in AM security was the dr0wned demonstration of a void attack [32], very similar to the research presented in 2014. In this attack a manually designed void was placed in a propeller that was fabricated on an AM system. The propeller was mounted to a quadcopter and caused a failure during operation. This attack shows one example of the real world affects an attack, such as the one presented in Chapter 3, could have. Work by Zeltmann, demonstrated the effects that alterations to the part toolpath could have on the part strength [33]. In this study the fill pattern of a tensile test specimen, fabricated on an extrusion-based AM system, was altered to affect the mechanical properties of the test specimen. A void attack was also carried out where square defects were placed inside of a tensile test specimen and ultrasonic imaging was used to attempt to detect the defects. In the intellectual property space attacks based on side-channel measurements have also been demonstrated [34]. By using a microphone to listen to the sounds made by the stepper motors in an AM system it was possible to recreate the toolpath with an error as small as 5.87%. While not directly affecting part quality, this type of attack demonstrates how tightly the cyber and physical domains are coupled in AM systems.

Other work has looked at the broader picture, evaluating the security of the process chain as a whole for metal AM [35]. Deloitte has looked in at the overall security of AM, discussing the reliance on digital files, digital connectivity, and the role that it plays in the supply chain [36]. In this work the scale and scope of AM manufacturing are investigated and the potential threat actors (such as criminals, nation states, insiders, etc.) and impacts (such as IP theft, infrastructure destruction are, reputation damage, etc.) are discussed. This work proposes broad solutions to security problems, such as conducting risk assessments and designing protections into the system from the start.

To address the security issues that arise with AM, a variety of approaches has been investigated. One proposed method for preventing attacks in extrusion based AM systems is through the use of

Carefully designed defects in the model file [37]. In this approach, one or more defects are inserted into the model file during the design process. Because the AM process converts model files from a surface-based format (the .STL file) to a toolpath file there are by nature translation errors that occur. The proposed method relies on these errors prevent the theft of intellectual property (IP). The defects are designed so that when printed with the correct process parameters the translation error eliminates (or significantly reduces) the defect. If the part is printed with the incorrect process parameters the defect becomes significant and the final part is non-operable. This requires an attacker to steal both the model file and the process parameters in order to fabricate a high-quality part.

One of the earliest works looking at anti-counterfeiting was through the use of quantum dots to create unclonable features[38]. A similar approach to prevent counterfeiting in AM parts is to embed tags inside of fabricated parts [39] In this approach a tag is placed inside of the part using pockets of support material or air. Under certain lighting conditions this tag is not visible, however when lit and observed from the appropriate angle the tag becomes visible. By storing data in the tag, the manufacturer is able to validate that the part is genuine, similar to a watermark. An attacker than attempted to reverse engineer the product through a method such as 3D scanning the surface would be unable to reproduce the tag and any counterfeit part could be identified by the fact that it either didn't contain a tag or the tag was invalid. More recent work in this space has looked at embedding anti-counterfeiting tags, such as QR codes, into metal AM parts using multiple material fabrication[40].

Another approach is to investigate a system of side-channel measurements to detect defects when they occur[41]. These methods use techniques such as power monitoring[42] , acoustic emissions[43–45], magnetic response, thermal characteristics [46], and motion sensing to detect when a build goes outside of normal operating parameters, indicating a failure. Other approaches have used image analysis techniques to monitor builds for the presence of attacks [47] . Further, by combining a variety of side-channel measurements it is possible to develop a model to detect both build failures and the type of failure that has occurred[47,48] .

3. Cyber-physical vulnerabilities in additive manufacturing systems: A case study attack on the .STL file with human subjects

Coauthors: Dr. Christopher Williams, Dr. Jamie Camelio, Dr. Jules White, Robert Parker

Abstract

One of the key advantages of additive manufacturing (AM) is its digital thread, which allows for rapid communication, iteration, and sharing of a design model and its corresponding physical representation. While this enables a more efficient design process, it also presents opportunities for cyber-attacks to impact the physical world. In this paper the authors examine potential attack vectors along the Additive Manufacturing process chain. Specifically, the effects of cyber-physical attacks, and potential means for detecting them, are explored. To explore the potential implications of such an attack, a case study was conducted to evaluate the ability of human subjects to detect and diagnose a cyber-physical attack on the STL file of a test specimen. Based on the results of this study, recommendations are presented for preventing and detecting cyber-physical attacks on AM processes.

Keywords: Additive manufacturing, Cyber-physical security, Advanced manufacturing

3.1. Introduction

3.1.1. Attacks on Cyber-Physical Systems

Cyber-physical systems (CPS) are systems that integrate physical hardware with software systems, often with the use of a network. With the growth of the Internet of Things (IoT) the number of CPS systems on networks continues to increase [1]. Concurrently, cyber-attacks have become more prevalent in recent years, increasing in maliciousness and decreasing in visibility [1–3]. This poses a significant issue, as cyber-attacks on cyber-physical systems could result in damage to the machines themselves or the humans who interact with them.

A prominent example of a cyber-physical attack was the Stuxnet worm that targeted Iranian centrifuges used for refining uranium. In this attack the worm was able to infect the software system and affect the physical hardware, causing damage to the centrifuges. By sending false data back to the operators, Stuxnet was able to make it appear as though the centrifuges were operating correctly, while it caused them to damage themselves. The ability of Stuxnet both to cause damage to physical systems and to hide itself illustrates the ability of a cyber-physical attack to disrupt manufacturing systems and the need for physical methods of detection[4].

Another example of a cyber-physical attack is the hijacking of insulin pumps. In this case a hacker is able to connect to a Bluetooth enabled insulin pump to control the dose of insulin given to the wearer. By increasing or decreasing the dosage of insulin, it is possible to cause serious injury or even death in the user. The currently security system for these pumps is insufficient to prevent a cyber-physical attack that could have potentially lethal consequences[5].

3.1.2. Cyber-Attacks on Manufacturing

The previously mentioned examples demonstrate the ability for cyber-attacks to cross over into the physical world. Attacks on CPS are even more alarming when considering the ever-increasing amount of networked devices that are being connected to machines in the manufacturing world. A cyber-attack on these machines could cause injury to plant workers and damage to the machine itself. This has already been demonstrated at a German foundry wherein a cyberattack on a blast furnace's control systems prevented a safe shut down and resulted in "massive damage" [6,7]. Perhaps even more insidiously, an attack could be designed to cause a process to produce faulty parts that might find their way into end-user products[8]. For example, an attack could be designed to affect the production of a jet turbine part such that it would pass inspection but fail during operation and cause significant damage.

With the rise in both the number of CPS connected to networks and in malicious cyber-attacks, there is a clear need for research to understand the vulnerabilities of cyber-physical systems. Recent work on the current status of cyber-physical systems has identified security as a major issue [9]. Additionally, there is a need to identify the skills that are necessary to effectively operate a cyber-physical manufacturing environment [10], which the authors believe includes skills pertaining to the identification and prevention of attacks on manufacturing systems.

As such, the authors have investigated cyber-physical vulnerabilities in manufacturing systems. The authors have demonstrated such attacks on subtractive manufacturing, and have found that they can have a demonstrable effect on the part being produced [11,12]. In the previous work the authors demonstrated a toolpath attack by manually modifying and replacing machining toolpath (e.g., GCODE). This work was limited to modifying the external geometry of a part in a way that affected the part strength, but was detectable using common inspection techniques.

3.1.2. Context: Cyber-Physical Vulnerabilities in Additive Manufacturing

In this paper, the authors scope their research solely on Additive Manufacturing (AM, commonly referred to as "3D Printing") systems. The process chain of these networked machines has unique vulnerabilities that warrant a detailed investigation. Specifically, due to (i) their layer-wise fabrication approach and (ii) the processes' impact on parts' final material properties due to process-induced transformation of raw material to finished good (as opposed to subtractive machining, where the material properties are defined by the feed stock), there are several cyber-physical vulnerabilities that are unique to AM. For example, voids can be placed inside of a part and the material properties of internal layers can be changed without affecting the exterior layers, which makes detection with traditional part inspection techniques difficult (as discussed in Section 3). Because of the potential damage from a cyber-physical attack, there is a need to investigate AM systems to determine what vulnerabilities exist and how to prevent and mitigate the threat of cyber-attacks. Recently, there has been increasing interest in evaluating these threats to AM systems, with the effects of toolpath modifications and embedded defects being simulated and evaluated through physical testing [13]. The key differences between additive and subtractive manufacturing are the features that make AM

unique: the ability to place defects internally and to change the material properties through either geometry or processing parameters.

An overview of the cyber-physical attack vulnerabilities of the AM process chain is presented in Section 2. A case study of an AM cyber-physical attack, in which the .STL file structure is altered, is presented in Section 3. The resulting effectiveness of this attack is reported in Section 4 via part testing and experimentation with human subjects. Finally, in Section 5, the results of the attack are analyzed to identify ways of preventing and mitigating future attacks.

3.2. Cyber-Vulnerabilities in the Additive Manufacturing Process

To be able to prevent a cyber-attack, one must first understand the vulnerabilities and weaknesses of the system. To do this, it is necessary to follow a cyber-attack through the process chain, from conception to simulated deployment. In this section, the AM process chain is examined for potential vulnerabilities to cyber-attacks. Previous work by Bridges [14] has given a high level process view by highlighted points in the process chain where the theft or corruption of a design could occur; however, in this paper the authors seek to provide additional detailed insight by discussing specific ways such an attack could occur as well as providing commentary on the difficulty of attacking each step in the process chain and the value associated with each.

The digital nature of the AM process chain, shown in Figure 1, provides an opportunity for a cyber-attack to cross into the physical world. AM systems are often connected to internal networks and may even have internet connectivity. This can allow for useful features, such as remote diagnosis troubleshooting, but also opens up the potential for an attack to compromise the systems remotely. There are four main steps on the process chain where an attack could take place: the CAD model, the .STL file, the toolpath file, and the physical machine itself. Two other areas where attacks could occur are (i) in the feed material (e.g., by altering a material database) and (ii) in the process monitoring and quality control system (e.g. in-situ monitoring of build volume temperature, dimension checks of finished parts, etc.). The primary differences from the authors’ previous work in subtractive manufacturing [11] are the features that make AM unique: the ability to place defects internally and to change the material properties through either geometry or processing parameters.

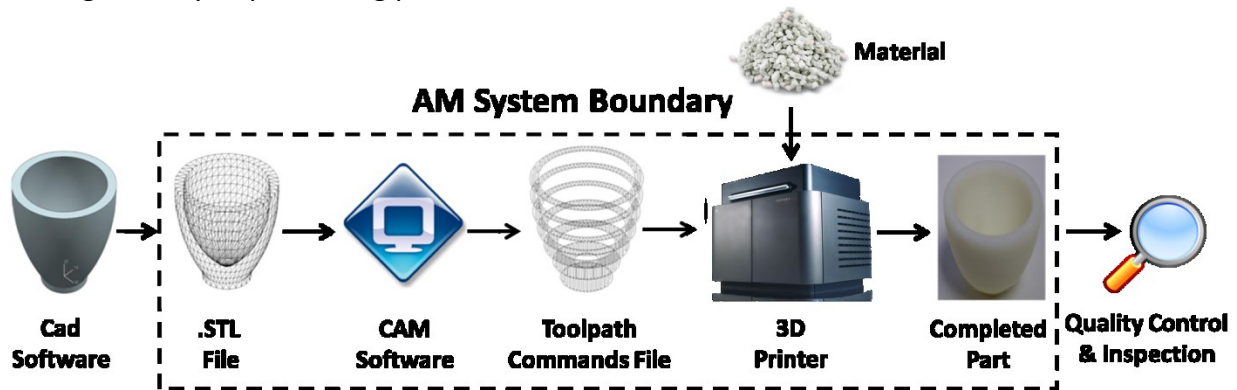


Figure 3.1. Additive Manufacturing Process Chain

3.2.1. CAD Model

The computer-aided design (CAD) model is the first step in the additive manufacturing chain and is common both to additive and subtractive manufacturing. This step in the process is the most valuable in terms of information, as it contains all of a part's geometric data. If connected to a product lifecycle management (PLM) software suite, the model could also include information related to simulated performance (e.g., results from finite element analysis, computational fluid dynamics, multi-physics simulations, etc.), failure modes, and the associated parameters of the part's intended use.

Attacks at this point in the process may focus on stealing or corrupting files. One example of this is the ACAD/Medre.A worm, which was designed to infect and steal AutoCAD drawings [15]. Another example of this is the CryptoLocker malware, which encrypts a user's files and then demands a ransom for the encryption key to unlock these files [16]. Because CAD files are the basis from which other drawings and files are generated, the theft or loss of these files can be costly. Furthermore, a competitor who stole these files would have access to all of the design steps that went into making the CAD model and could easily reverse engineer or modify the part.

The complex nature and proprietary format of most CAD files makes it more difficult for an attack to directly alter the part file; however, one could be designed to do so. For example, a CAD file for a crankshaft could be altered to reduce the area of the load bearing member, resulting in premature failure. Any corruption at this phase would propagate through the entire process chain, resulting in a part that is "bad" from start to finish. However, because parts may still be edited during the CAD phase, and the use of revision management in PLM software, the chances of detecting an attack are increased.

3.2.2. STL/.AMF File

The second step in the process chain is to convert the CAD model into a .STL file, the current defacto and open-source file standard used in AM. Upon conversion, most of the model data is lost, and only the surface information of the part remains. The constructive solid geometry (CSG) representation of how the geometry was created is lost in the translation from CAD to .STL; so, it becomes harder to reverse engineer the information needed to create a similar part. Despite this loss of information, a theft of a .STL file is still costly as it (i) contains all of the information needed to fabricate the geometry of the part (which could result in the production of counterfeit copies) and (ii) the surface geometry data can be attacked to nefariously change part geometry.

The .STL representation of the solid model is no longer represented by complex mathematical equations, but by a simple surface geometry composed of triangular elements, called facets. Each facet is defined by three vertices, specified by a set of x , y , and z coordinates, and a normal vector. Each closed set of facets comprises a shell, and multiple parts can be represented inside a single .STL file by including multiple shells. As with the CAD file, the .STL representation of a part could be altered to reduce the finished part's performance. This can easily be accomplished by altering the coordinates of the facets' vertices in the .STL file, which would change the part's exterior dimensions. Furthermore, additional vertex coordinates can

be appended to the STL file to create new features. Perhaps most dangerous of all, facets could be appended to the file in such a manner as to affect the interior of the part. One can imagine malware that took advantage of AM's unique layer-by-layer fabrication process to fabricate parts that look and feel strong on their exterior, but have an embedded weakness (e.g., crack propagation site).

The open nature of the .STL file makes it easy to modify. The lack of volumetric information makes it more difficult for an attacker to determine where the most dangerous location to cause a defect is than if they were attacking the CAD file directly. The file is generated fairly early in the process making it valuable to attack. Finally, the lack of process parameters means that only geometric data can be attacked at this point. Additional details of the vulnerabilities of a file that contains a faceted representation of solid geometry are presented in Section 3 and demonstrated as a case study in Section 4. It is noted that both the .STL and .AMF file (the current ISO/ASTM standard [17]) are open-source file formats that contain vertices of facets, and thus are both susceptible to the same types of attacks. The .AMF file contains more model information, which arguably makes it vulnerable to a longer set of attacks than the .STL. For this reason the focus of this research will be on the .STL file, with the understanding that these techniques are also valid on the .AMF file.

3.2.3. Toolpath File

Upon receiving a .STL file, each additive process converts the model into layers and generates a toolpath from these layers. Conceptually similar to GCODE, this toolpath file contains the commands for the controllers that move the AM systems' coordinate axes and deposition mechanisms (e.g., extrusion federate, laser power, inkjet pulses, etc.). Potential attacks on the toolpath could include rewriting the file to feature maleficent instructions to (i) to place/remove material in the wrong location, (ii) to cause layers to be placed too close/far away from each other, (iii) to alter deposition timing/patterning so as to affect interlayer adhesion, and (iv) to damage the part/machine by driving the tool into the part/machine. As demonstrated by Zeltmann, the alteration of the toolpath direction can have significant effect of the mechanical strength of a part [13].

Toolpath attacks offer the most freedom in what an attack can achieve because any operation the machine could normally use when creating a part can be altered in the toolpath. A toolpath file could be intercepted by a virus on the machine's computer or when the file is sent from one computer to the machine. For example, with a wireless-enabled AM system the information might be sent over a network where it could be intercepted. While subtractive manufacturing is also vulnerable to toolpath attacks, AM systems are more vulnerable due to the greater freedom inherent in the process. While an attack on a subtractive process will primarily affect the surface where it may be detectable, toolpaths attacks on AM systems can affect the internal properties or geometry of a part (as discussed in Section 3.1).

Theft and alteration of the toolpath file is more challenging than that of the .STL file, since it toolpath files are proprietary and specific to individual machines (even within the same product family). While it could be reverse engineered or used on another machine of the same type to produce the part, it requires more effort to do so compared to that of a .STL file. Attacking the

model through the toolpath does allow any of the changes that can be made in the .STL file to be made to the part; however, the implementation of such changes is more difficult in the toolpath.

3.2.4. Physical Machine

The final component in the AM process chain is the machine itself. This comprises of two parts, the physical hardware and the internal microcontroller or embedded computer. This section will focus on the controller for the printer and how it can be attacked to affect part quality or damage the physical hardware. This stage of the process is vulnerable to attacks that are similar to the aforementioned manufacturing cyber-physical attacks (e.g., Stuxnet and German foundry, Section 1.2), in which a worm alters the firmware of the system's controller. AM systems come in a variety of configurations. Some systems have a simple internal microcontroller and require an external computer to prepare the build for them. Some systems have embedded computers that allow for more advanced functionality. In both cases the systems may or may not be connected to a network. Without network access an attack would need to be brought over using some type of physical media such as a USB drive. This is feasible due to the need to transfer part files to the machine in order to fabricate them. In the case of a networked system (which is becoming a more common requirement for advanced AM systems to enable OEMs to ensure their machines are not being augmented by the machine owner), an attacker could gain virtual access to the system. Many manufacturing systems are locked down once a final configuration is determined, often preventing those systems from being able to receive updates and security patches that address known vulnerabilities. Interestingly, as most AM systems have both open USB ports and open IP ports in the PC firewall to enable off-site machine maintenance and troubleshooting, such a sophisticated attack vector is not needed.

Using this as an attack vector, the machine's process parameters (e.g., nozzle temperature, laser intensity, etc.) could be altered to affect the material properties of the part. For example, inkjet nozzles might be turned off at one point to prevent material from being placed, or turned on at another point to place excess material. Temperatures of extrusion nozzles could be altered to change the mechanical properties of the extruder or to clog the nozzle all together. Laser intensity could be altered to provide lower energy density to the bed to alter the final part's mechanical properties. An example of such a vulnerability is shown in Figure 2, which contains configuration data for a layer being printed on a Material Jetting AM system. The system's configuration file was intercepted in plaintext using an open-source network protocol analyzer (Wireshark[12]) over the laboratory's intranet. By intercepting and altering this configuration file, it would be possible to remotely alter vital process parameters that could alter mechanical properties of the printed part (e.g., by altering UV intensity), and/or harm the machine itself (e.g., by increasing limits on operating temperature). This type of attack can be difficult to detect without some type of independent monitoring.

There does exist at this step some potential for model data to be extracted and reconstructed using side channel measurement techniques such as acoustic measurements. An example of this is work done by Chhetri et al., where a device was used to record the sounds of the motors of an extrusion machine and to reconstruct the part geometry from the calculated

movement of the printer[18]. Other examples of this type of attack could include the use of video footage (gathered from a smartphone or security camera) to reconstruct the part geometry. These attacks require sophisticated knowledge of the process being attacked, may be subject to interference such as factory floor noise, and would not work for energy-based systems (where material formation does not create a detectable acoustic signal).

Compared to the prior vulnerabilities discussed, the physical machine has some security advantages that would make it more challenging for an attacker since (i) it requires access to the machine (either physical or virtual) and (ii) configuration files are typically proprietary, machine-specific (i.e., while one machine might be particularly vulnerable to an attack, another machine could be well secured and hard to infect). Such an attack would also require in-depth knowledge of the impacts of process parameters on the printed part. While these challenges might deflect a cursory attempt at an attack, a targeted attack could certainly overcome them. It is also noted that this vulnerability is focused on the machine only; no model data is available to be stolen or altered in this aspect of the AM process chain.

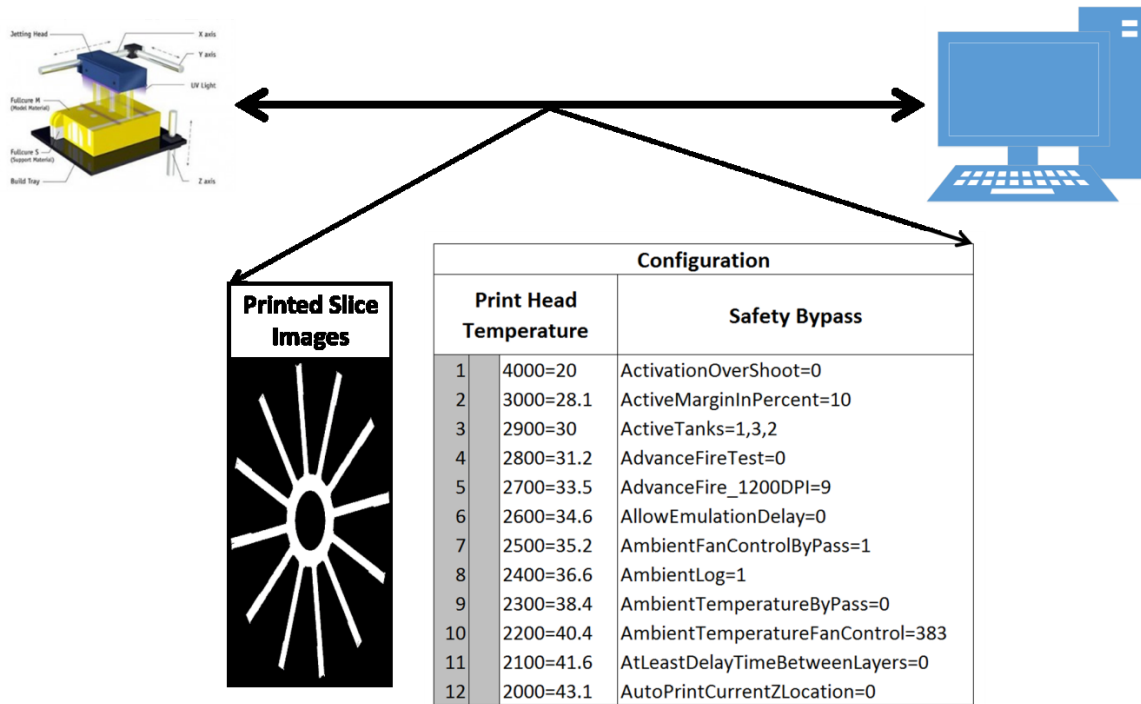


Figure 3.2. Configuration data for a Material Jetting AM system intercepted over WiFi using network protocol analyzer software.

3.2.5. Summary

Based on the above discussion, the authors believe that the .STL file is currently the most vulnerable step in the AM process chain. The .STL is the one potential attack that does not require specific modification for every individual AM machine, since it is an attack on the

standardized file format that every AM machine takes as input. As .STL file creation occurs at the beginning of the process chain, and is general to every AM machine, a focused attack on the file could have severe implications and could cripple an AM production line. Currently, the standard .STL file is not encrypted or obfuscated in any way, which makes it easy to modify. Modifying the .STL does not require significant technical knowledge of the process or the implementation. Finally, the .STL file is often transferred widely increasing the number of locations where it may be attacked. As previously mentioned, AM allows for defects to be placed that are internal to the part being attacked. This distinction from subtractive manufacturing makes AM more vulnerable to an attack at this step than subtractive manufacturing would be.

Of the other attack vectors, attacking the physical machine has the greatest potential to introduce defects that could escape detection. Since the integrity of the model file remains intact, many cyber based solutions will be unable to detect this attack. By altering machine process parameters, changes to the material properties can be introduced that can be difficult to detect in the physical realm. The toolpath is the next most dangerous attack vector. Once an operator has validated the toolpath (if possible) the build is assumed to be correct. An attack that occurs after this point would have to be detected in the physical realm, either during printing or in final inspection. Further, some processes include process parameters in the toolpath file. This allows for some physical machine type attacks to occur at the toolpath step. Finally, the CAD file is the least vulnerable from a part quality standpoint. This is due to the fact that it is early in the process chain and many steps exist where a change might be detected. It is still a crucial step to protect to prevent the theft of intellectual property.

3.3. Case Study: Cyber-Physical Attack on AM Systems via Altering .STL Files

To gain a better understanding of existing vulnerabilities, to determine if these vulnerabilities are significant, to understand the circumstances that allow attacks to occur, and to develop better methods for preventing cyber-physical attacks from occurring in the future, the authors explored the effects of a cyber-physical attack on AM systems. Specifically, following the discussion of cyber-physical vulnerabilities along the AM process chain (Section 2), the authors decided to explore the potential impacts of the vulnerabilities of the .STL file. A description of the specific vulnerabilities in a faceted representation of solid geometry is presented in this section. To prevent external replication of the attack conducted in this case study, details of the final attack algorithm, code, and implementation have been omitted.

3.3.1. Types of .STL Attacks

As described in Section 2.2, while a .STL file only has surface data in the form of a list of triangle vertices, there are several different types of file attacks that can affect the final part geometry. Such attacks could happen in the form of malware, wherein .STL files are intercepted and automatically altered without the operator's knowledge. The effects of these attacks on sample part geometry (an ASTM D638 tensile test specimen) are illustrated in Figure 3.

- *Corruption/Encryption* – A traditional cyber-attack where the file is damaged or encrypted, this renders it inaccessible to the user. A corruption or encryption attack is a straightforward attempt to damage or extort the owner of the file. By rendering the file unusable it makes it very evident that an attack has taken place. However, if backups are not available, this type of attack can cause a lot of damage before it is detected. (Figure 3a)
- *Scaling* – The part is scaled up or down in one or more axes by simply applying a scaling factor to the vertex coordinates (Figure 3b). Such an alteration would result in a changed form that may affect the fit or strength of the part. For example, a tensile test specimen that is 10% thinner may not be noticeably different to the eye, but will cause a measurable change to performance. This type of attack could be checking the digital dimensions in AM pre-processing software and/or by measuring the external dimensions of the completed part.
- *Indents/Protrusions* – Small protrusions or indents may be added to a part to affect the fit, surface finish, or strength of the part (Figures 3c and 3d). This can be done by changing the coordinates of a single vertex, in this case making it identical to a vertex movement attack, or by adding and modifying facets and vertexes. For example, the inside of a printed duct could be altered to contain indents that negatively affect the flow. While such attacks would be visibly detectable, they may be placed in locations where they are difficult to see or measure.
- *Vertex Movement* – One or more vertices in the part is altered, resulting in a changed form that may affect the fit or strength of the part (Figures 3e and 3f). This has the advantage of manipulating the shape of the file without changing its size. Unlike scaling, vertex manipulation allows almost any section of a part to be sized both up and down, as well as warped or altered. Measurements can detect this attack, but by altering areas that are difficult to measure, the presence of this attack can be hidden.
- *Voids* – In this attack, a set of triangular facets (and corresponding vertices) could be appended to the .STL file that define a new shell. This shell could be placed within the extents of the existing model and, if its normal vectors were flipped relative to the parts' exterior surfaces, it would effectively define a “negative” space inside of the model. Because such voids are completely enclosed, they are undetectable by dimensional measurements and may be difficult or impossible to find visually. The use of supporting material in many processes also renders the void difficult to detect by mass measurement, since the void is filled with a structurally deficient, but similarly dense, material. Small differences in weight already occur due to process variation. The presence of a void would weaken a part and, if placed in a loadbearing location, may cause premature part failure. This attack does increase the file size slightly, but the small change is unlikely to be noticed (Section 3.2).




Figure 3.3. .STL Attacks on an ASTM D638 tensile specimen. A) an unaffected dogbone; B) a scaled down dogbone; C) an indentation; D) a protrusion; E) a vertex moved inward; F) a vertex moved outward.

None of the attack types listed would violate the governing rules of .STL files (e.g., watertight surfaces, non-manifold geometries, Euler-Poincare characteristic), and would thus not automatically raise errors during slicing or toolpath generation. The first four types of attacks affect the exterior surface of the printed part, and are thus shared by both AM and subtractive machining (via GCODE alterations).

The void vulnerability, however, is unique to AM. Due to its layer-based fabrication method, the system has access to the entire volume of the part during fabrication, not just the exterior surface as in traditional subtractive machining. Furthermore, unlike a protrusion, indentation, or other external change, a void cannot be physically measured once it is enclosed inside of the part. Because voids are not readily visible when looking at a model on a computer screen, and are not measurable via traditional quality control processes (e.g., dimensional checks or visual inspection methods) once the part has been made, they present a higher likelihood of passing undetected. In some AM processes such as extrusion, a visual representation of the toolpath may be generated and could be used by an operator to detect the presence of a void. In other systems, such as material jetting or powder bed fusion, there may be no toolpath preview available for the operator to validate. The lack of a toolpath validation step makes these systems particularly vulnerable to attacks on the model file that occur after the model has been validated. The authors' prior work demonstrated the potential of a void attack [19]. Later research by Belikovetsky, et al. has further validated the potential of this work by causing catastrophic failure on a propeller fabricated on an AM system [20].

3.3.2. Void Attack Considerations

Given the potential for a bad actor to maliciously alter a .STL file to include an internal void in a part, the authors explored the feasibility of such an attack. Specifically, the authors sought to answer the following questions: *Can a software automatically execute a void attack with solely the information provided in a .STL file? Could a void be automatically sized and placed appropriately to cause part failure? Could the void attack be detected?* To answer these series of questions, the authors began by exploring considerations for implementing such an attack:

- *File Size* – When altering a .STL file a change in file size may make it easier to detect an attack. Some attacks, such as scaling or moving a vertex, do not change the file size. Other attacks, such as adding voids or other features, will change the file size depending on the complexity of the feature added. It was found that a tetrahedron void can be added by inserting four additional facets to the .STL file, which results in a file size increase of only 200 bytes. This size increase is negligible in all but the simplest files, as .STL file sizes are typically on the order of mega- and gigabytes.
- *Location* – The location of a void in a part has a significant impact on the effect of the attack on a part. A void is not a threat if it does not result in any reduction in part strength. A well-placed void is a concern, as it can lead to a critical part failure.
- *Void Shape* – The shape of the void affects the amount of information needed to represent it in a .STL file. The vertices of the simplest solid body, a four-faceted tetrahedron, is sufficient for initiating a crack.
- *Void Size* – Void size is important for affecting the mechanical strength of the part. A void that is too small will not be fabricated, due to process resolution limits. Large voids increase the chance of part failure, but may be easier to detect during fabrication. As demonstrated by Zeltmann, defects that are too small may not significantly reduce the strength of a part[13],
- *Void Number* – The number of voids placed can be varied based on the desired effect. A larger void might be made less noticeable by replacing it with several smaller aligned voids. The addition of more voids will add more to the file size and may make the change more noticeable.
- *Full Enclosure* – To be the most effective, the void must be fully enclosed in the part model. If a void is not fully enclosed inside of a single shell, it will cause printing and model slicing errors.
- *Malware Run Time* – The time it takes for the attacking software to analyze and place a void is directly related to how easy it is to detect. A several second delay caused by the attack algorithm would likely draw suspicion.
- *Silent* – To be effective, such an attack should not set off any warnings. For example if the void contained a facet with flipped normal, AM pre-processing software would display a warning to the user alerting them that there was a problem with the file.

In summary, the most dangerous void attack would be one where the void was located in a structurally important location and sized such that it was difficult to detect while still causing a critical structural failure in the part. Such an attack would run quickly, add a minimal amount of

data to the part file, and would not cause any warning messages when inspecting or printing the part.

3.3.3. Attack Embodiment

To assess the potential impact of a .STL void attack, the authors created a piece of software that featured three key pieces of functionality: (i) the ability to ensure that the void is located completely inside the part, (ii) the ability to automatically determine a location for the void, and (iii) the ability to automatically scale the void based on the part.

3.3.3.1. Void Encapsulation

The ability to place a void completely inside of the part is a key part of a void attack, both for the functionality of the void and for avoiding detection. If a void is placed completely outside of the existing shell of the part, it will serve no function and is not a concern. Similarly, a void that extends partially outside of a part will cause file processing errors. The only dangerous void attacks are those where the void is completely enclosed. The task of placing a void inside of an arbitrary .STL file poses several challenges due to the limited information contained in the file. First, .STL files may contain multiple shells. A void may be fully enclosed inside of the part, but still intersect multiple shells causing an error. Second, a .STL file can contain almost any geometry (e.g., convex or concave). Because of this, techniques for determining collisions based on convex geometries cannot be used. Other approaches, such as determining the center of mass, are met with similar difficulties, as the center of mass may be outside of the object.

In analyzing this problem, a two-step method was found that could be used by an attacker to ensure that a void is placed completely inside of a part. The first step determines the number of closed shells and attempts to locate a likely failure location before placing a small void. This initial void location is based on the existing coordinates found in the facets in the file. The second step uses ray tracing to determine the rough size of the location and to scale the void up to a larger size in order to increase the probability of causing a failure, since very small defects may not cause a meaningful reduction in part strength [13]. The method also allows for multiple smaller voids to be used instead of a single large void and for the maximum void size to be limited.

3.3.3.2. Void Location

Based on the previous discussion, the authors were able to successfully demonstrate void encapsulation. Next, the authors explored the feasibility of strategically locating voids within a part to maximize its impact on the part's structural integrity. Because .STL files are not solid models, it is difficult to attempt any type of structural analysis of the part. Additionally, since the parts infected could vary significantly, and the purpose of each part is not known, it is hard to place a void in any meaningful way. However, one could write an algorithm that places voids in locations that are likely to have stress concentrations (e.g., areas with holes or sharp changes in curvature). These areas can be identified within the limited information provided by a .STL file by analyzing the density of the mesh. Similar to FEA, more complex areas will be represented by a denser mesh than flat areas. By placing voids near locations where the mesh is dense, the likelihood of placing them somewhere that will cause a critical failure increases

significantly. Once the location of the densest mesh has been found, the surface facet is used as a reference to create the void. The reference is offset from the surface by a small amount and then used as the basis for the void. This method allows for voids to be placed in multiple locations that are likely to cause failure in order to increase the overall chance of causing a part failure.

3.3.3.2. Void Scaling

Another feature that makes a void attack more dangerous is the ability to scale the size of the void to ensure that it is printable by the AM system. An attacker could increase the chance of failure by making the void larger, while still keeping it small enough to avoid detection. An algorithm can be written that automatically scales the void geometry appropriately by referencing the dimensional extents of the part geometry. By combining scaling techniques with the previously discovered method for ensuring the void is completely enclosed, the chance that the void placed by the software would cause a failure was increased.

The resulting attack software automatically (i) analyzes a .STL file, (ii) places a tetrahedron-shaped void completely inside the part (iii) near areas where the mesh is most dense, (iv) scales the void appropriately, and (v) verifies that the void is fully enclosed inside of a single shell. The base process is illustrated in Figure 4. Adding the void adds a total of four new facets to the file and creates a new negative volume shell into an existing shell. For a binary .STL file the total change in file size is 200 bytes since each facet requires 50 bytes to represent. Though this might be a discernable change in a small, simple part, it will be nearly invisible in a larger more complex part that has a file size of several megabytes.




Figure 3.4. Process flow of void placement algorithm; 1) Void is fully encapsulated inside of the part; 2) The void is positioned in a location that is more likely to cause a failure; 3) The void is scaled either up or down to increase the chance of failure or decrease the odds of detection.

3.4. Effects of .STL Void Attack

To ascertain the potential impact of this specific attack, the authors sought to answer two research questions: (i) *Can an automatically inserted void affect part strength enough to cause a failure?*, and (ii) *Can a void attack be detected by operators?* These questions were answered via two experiments: first, the authors evaluated the effect of a “printed void” on the mechanical strength of a printed specimen; second, the authors evaluated the ability of AM operators to notice the attack via a human subjects experiment.

3.4.1. Effect of Voids on Part Strength

The goal of the attack was to evaluate the impact of a void placed inside of a part for measurable degradation of part quality. An ASTM Standard D638-10 tensile test specimen was used as a sample part as it provided a straightforward way of measuring the quality of a part with and without a void. Figure 5 shows a cross-sectional slice of a part with a void placed inside of it using the author's software. The void placed was a tetrahedron with a volume of 3.3mm^3 and a surface area of 16.9mm^2 . The longest edge was 4.1mm and the shortest edge was 2.3mm . To quantify the effects of the void placement, finite element analysis was run to determine if the void was likely to cause a failure. Unsurprisingly, the Von Mises stress of the part is highest at the location of the void (Figure 6).

In order to verify this result for parts manufactured using an AM process, several tensile specimens with and without voids were fabricated via Powder Bed Fusion AM system (DTM Sinterstation 2500 Plus) using Nylon 12 powder. Upon testing, all of the specimens containing voids fractured at the void location, while the dogbones without voids failed normally, as shown in Figure 7. The average reduction in yield load was 14%, from 1085N to 930N, and the strain at failure was reduced from 10.4% to 5.8% as shown in Figure 8.




Figure 3.5. Cross sectional slice of a dogbone infected with a void.



Figure 3.6. Von Mises Stress of a dogbone infected with a void.




Figure 3.7. On Left: Uninfected dogbones breaking at the gauge section. On Right: Infected dogbones breaking at the void location within the specimen neck.




Figure 3.8. Load and strain data of parts with and without voids.

In addition, the tensile specimens containing voids were tested to see if the voids were detectable through common inspection techniques. Physical dimensions of the part revealed no differences, since the void was fully enclosed. Mass measurements of the parts did not reveal any differences since the void was small and filled with loose nylon powder. For spherical particles, packed nylon powder would have a theoretical density of 0.5 that of solid nylon. In practice, this difference is somewhat smaller due to the non-uniform particles and parts that may have slightly reduced density. The void size was 3.3 mm^3 while the model size was 8440 mm^3 . With a density difference of 50%, the total percentage change to the part mass would be 0.02% (mass difference of $\sim 2 \text{ mg}$). Normal process variation such as differences in the cleaning process results in larger changes to the model mass/density. As a result, the void is undetectable through mass measurements. In material jetting and material extrusion, the support material has a very similar density to the model material, within the range of a few percentage points. Voids in these parts also cannot be detected by weighing.

3.4.2. Human Subjects Experimentation

3.4.2.1. Participants and Context

A case study was performed to determine the feasibility of a cyber-attack on a simple AM system and to evaluate the ability of AM operators to detect an attack. Participants were senior-level undergraduate and graduate students enrolled in an Additive Manufacturing course at a large, ABET accredited land-grant university. These mechanical engineering students had completed (i) a large portion of their required degree coursework, (ii) a comprehensive capstone design project, and (iii) a class project that required them to create a unique product designed specifically for fabrication via additive manufacturing. In addition, many of the

student participants had previously participated in industrial internships or co-ops. Given these prior experiences, and their completion of the course that covered all additive manufacturing processes, materials, process physics, and AM-specific design methodologies [21] the study participants are considered to have intermediate to advanced knowledge and understanding of AM. To mimic today's AM workforce, the participants were not given any specific information or training in the possibilities and implications of cyber-physical attacks on AM systems prior to the exercise.

In this experiment, the participants were challenged to design, manufacture, and destructively evaluate a tensile test specimen. The exercise was offered to the students as an extra-credit opportunity that would provide them a chance to learn about setting up builds on a Material Jetting AM system, and to learn about quality control methods by testing the parts on a tensile test machine. However, unbeknownst to them, their submitted STL files were attacked with the authors' algorithm (Section 3.3) to include internal voids (described in Section 4.2.2). This participant deception was approved by the authors' Institutional Review Board (Virginia Tech IRB #13-959). After the completion of the study, all participants received full disclosure of the deception and true purpose of the study. The participants were then given the opportunity to consent or withdraw their results from the study. All participants consented to the anonymous release of the data collected from the exercise.

3.4.2.2. Experimental Approach

A tensile test specimen was chosen as the target geometry due its established standard for analysis and because it was a simple (and professionally relevant) part for students to design. The students were asked to complete the following steps in the assignment:

1. Create an ASTM Standard D638-10 tensile test specimen using CAD software.
2. Export the part as a .STL file.
3. Open the .STL file on the lab computer and use an AM pre-processing software (Netfabb) to verify that the .STL model is valid and ready for fabrication.
4. Load the part onto the build tray using Objet Studio and print the part on a Connex 350 machine.
5. Use calipers to measure the completed part to verify that it matched the designed model.
6. Perform a tensile test on the part and evaluate the results in a report.

Unknown to the students, the computer used in Step 3 was infected by the authors' .STL attack software. Upon transferring their part file to the computer, a second piece of software running in the background (referred to as the wrapper) detected the presence of the file and executed the .STL attack software, which automatically appended void facets to the .STL file (and thus into the part). This attack was executed while normal anti-virus software was running (Microsoft Forefront) and did not trigger any alerts in the AV software.

3.4.2.3. Results

The experiment was performed on five groups of students (total number of students = 21). One group was given a part that contained a void, but its location would not cause a premature failure. The remaining four groups were given parts with voids that would cause a premature failure at an unexpected location.

From observations of each team’s participation during the exercise, and analysis of their resulting group reports, none of the students detected the presence of the void during the copying of their .STL file and the inspection in Netfabb. The void was able to pass the standard Netfabb warning test as the void was a fully enclosed separate shell with a negative volume. A more thorough investigation using the repair tool would have revealed the presence of a second shell. While in this study the void was inserted before the part was verified, the attack could also be designed to affect parts after the inspection stage instead.

While printing the parts, two of the five groups remained to watch the entire printing process. Only one of these two groups noticed the presence of the void during printing and remarked that there was a “divot” in their part. Upon completion of the print, none of the groups were able to detect the presence of the voids. The group that had previously noticed the “divots” during printing assumed that the divots they had seen had been filled in. When measuring with calipers, all of the groups reported that their parts were within a reasonable margin of their original design dimensions.

Table 3.1. Group caliper measurements of tensile test specimens.

	Width (mm)	Thickness (mm)
Design	13.00 (0.00)	3.20 (0.00)
Group 1	13.22 (0.03)	3.19 (0.01)
Group 2	13.12 (0.02)	3.19 (0.01)
Group 3	13.09 (0.00)	3.19 (0.00)
Group 4	13.13 (0.03)	3.22 (0.05)
Group 5	13.21 (0.07)	3.19 (0.05)

Upon breaking the part via an Instron tensile testing machine, all four teams with a performance affecting void recognized that their part failed prematurely. Of these, two teams did not notice the voids and attributed the failure to the anisotropic nature of additively manufactured parts. Two teams detected the presence of the void on the fracture surface; one team correctly identifying it as the “divot” that they had seen earlier. Both of these teams concluded that the void was most probably due to problems with the AM system itself (e.g., temporarily clogged jetting nozzles). In all cases, similar to the victims of the Stuxnet attack, the premature failure of the part was attributed to problems with the machine and not with a cyber-attack. The team with the parts containing voids that would not cause premature failure during tensile testing was unaware that their parts had been attack and made no mention of any defects, voids, or abnormal behavior in their report. This result demonstrates (i) the feasibility of a .STL “void” attack, (ii) the potential of a void attack to cause premature part

failure, (iii) the challenges in detecting a void attack, and perhaps most importantly, (iv) the need for educating future engineers about the potential for cyber-physical attacks on advanced manufacturing systems.

3.5. Discussion and Recommendations

The results of this study show that more work is needed to protect AM systems against cyber-physical attacks. While the focus of this work was on the vulnerabilities of the universal AM file format, it is not done to suggest that we should move to a more closed, proprietary format. Not only would such an effort eliminate the primary reason behind the recent, widespread proliferation of desktop-scale AM systems within the hobbyist/maker community, it would also fail to address the root of the problem. Any security system can be broken eventually, and replacing/encrypting the .STL file would do little more than provide a false sense of security and delay the inevitable hack. The authors argue that a better way of addressing the threat of cyber-physical attacks is to develop advanced monitoring systems and quality-control procedures that take into account cyber-physical vulnerabilities, and to better train the workforce to detect these attacks.

Improved Software Checks – While AM pre-processing software are capable of detecting shells with negative volumes, this ability alone is not sufficient protection from a void .STL attack. While an operator checking for extra shells might be able to detect an attack that occurs before the model is inspected, a previously validated model could just as easily be attacked and would escape detection. Another common check is for very small shells and features since they are unlikely to be important features, and are often unneeded artifacts left over from another process like 3D scanning. While testing for these shells increases the chance of detecting a void based cyber-attack additional validation is needed further down the chain after the part is inspected in order to ensure that no additional defects have been added. The incorporation of toolpath process checks is also important in the detection of cyber-attacks. Being able to visually validate the toolpath before sending it to a printer is an important tool in providing an extra check to detect the presence of an attack. While some AM systems include this feature, it should be extended to all AM systems.

Hashing/Secure Signing/Blockchain – Hashing is a technique commonly used in security to ensure the validity of a file. The file is run into the hashing function, which generates a string of character called a hash. The hash is then posted along with the file. When a user downloads a file they can run it through the same hashing function and compare the resulting hash with the posted hash. If the two hashes match, the file can be assumed to be identical to the original. What makes hashes effective is their ability to convert a large file into a simple string that can be easily shared. Any small change in the file generates a large change in the hash. The simple adding of a single character to a text file will completely change the hash that is generated. At the most basic level hashes can be (and are) used to ensure that the .STL file received has not been tampered with. While this does increase security it adds some additional work to the process. Additionally, a file could potentially be attacked after it has been received and hashed, or before it was hashed to begin with. This problem could be addressed in part by including the hash function at the time of file creation, within the CAD software itself, and generating a hash

function at the last step in the process, where the .STL file is loaded in to the printer software to be converted to a toolpath. Another approach that has recently been shown is the use of a block chain to prevent and detect tampering with files[22]. This type of implementation could potentially stop an attack such as the .STL one discussed by the authors. A second benefit would be the potential to track the point where such an attack had occurred. Other attacks such as those affecting process parameters might still be vulnerable. One area of further research is the creation of a physical hash, a hash that incorporates elements from both the physical and the cyber side to give added resilience against cyber-attacks.

Improved Process Monitoring – Improved process monitoring is an ongoing goal and area of research for AM in an effort to better control and improve their performance. Some solutions have already been explored, such as using optical sensors to provide closed loop control for AM (e.g., layerwise laser melting [23]). Indirect measuring, or “side channel” measurements, such as measuring the temperature of the melt pool to determine the laser power instead of simply asking the machine can be effective at detecting cyber-attacks effecting the machine parameters, since, as Stuxnet demonstrated, a clever attack can cause a system to report false data. While feedback systems provide valuable information for process control, used alone they are insufficient to detect cyber-attacks. If the .STL file for a part is infected with a void a closed loop system will notice the void, compare it to the part file, and conclude that it is an intended design feature. For this reason more research needs to be done to establish statistical process control (SPC) for AM systems that works to detect the effects of cyber-attacks. By combining side channel measurements and using SPC to establish baseline operating parameters, systems can be made more robust and also more resistant to cyber-attacks.

Operator Training – The final area that needs improvement is in education and training. The vast majority of malicious software is installed or transmitted unwittingly by uneducated users. Educated workers are more likely to detect and prevent a cyber-attack. As demonstrated in the human subject study (Section 4.2), future engineers are not aware of the threat of potential cyber-attacks. Even when confronted with an obvious problem, no group was able to identify the issue as a cyber-attack. Informing students and workers about the potential risks of cyber-attacks can prevent these attacks from occurring and help diagnose them faster when they do occur. The greatest need is for operator awareness; however, a methodology for identifying cyber-physical attacks and differentiating those from traditional manufacturing errors is needed to properly educate operators about the risks of cyber-physical attacks on manufacturing.

Of these areas the authors believe that the most efficient component to address is the improvement of software checks and tools. Specifically, checks that validate the geometry and toolpath of the model and that allow operators to validate the toolpath before fabricating. Next, the authors’ believe that existing tools for security from the cyber domain such as hashing or blockchains should be integrated without requiring significant additional development. Such tools could fix many of the cyber vulnerabilities that currently exist in AM system. The use of these cyber-security tools alone; however, is insufficient to address the *cyber-physical* vulnerabilities that exist in AM. Operating manufacturing systems are often unable to accept software upgrades due to software compatibility (e.g., operating system, process drivers, etc.) issues and concerns about process downtime; for such systems, it may not be possible to

implement these types of protections. In systems where these protections can be implemented they cannot prevent all attacks. An attack on the process parameters of a machine will not affect the integrity of the model file, but it will affect the final part quality. To complement these tools, the authors suggest expanding current research on post- and in-situ process monitoring to provide cyber-physical security as side-channel measurements. While these tools are currently being developed for quality control issues due to normal process variations, the extension to consider attacks is needed to help detect vulnerabilities, such as a change in process parameters. Finally, all of the above methods rely on improved training of the workforce to be aware of the potential of cyber-physical attacks and to be able to diagnose and prevent such attacks. Due to the prevalence of cyber-attacks grounded in social engineering (e.g., spear-phishing attacks), an AM operator's lack of diligence in cyber-security can ultimately undermine any efforts to secure the system.

3.6. Closure

With the increasing number of manufacturing systems connected to networks, more work needs to be done to ensure the safety of these systems. Additive manufacturing systems in particular have unique vulnerabilities presented by the ability to affect the internal layers without affecting the exterior. An overview of the AM process chain (Section 2) showed that the .STL file was a vulnerable attack vector due to its universality and ease of editing. Further investigation into the .STL file revealed a method by which a void could be automatically placed inside a part, while avoiding detection by common process checks (Section 3). This void placement was demonstrated to cause a 14% reduction in yield load in a tensile test specimen. A human-subjects experiment, wherein a void was inserted into a test specimen via an automated attack algorithm, demonstrated that such STL attacks could be implemented without being detected by machine operators, common STL validation software, or the installed virus checking software installed on the process's computers (Section 4). Based on the results of this study, it appears that a real threat from cyber-attacks exists and that further research needs to be done on how to mitigate such attacks. The inclusion of software checks, hashing, process monitoring, and worker training are proposed as methods of reducing these threats. Future work includes the development of physical hashing techniques and of improved side channel process monitoring and control

Acknowledgements

This work was supported by the National Science Foundation Grant #1446804: "CPS: Synergy: Collaborative Research: Cyber-Physical Approaches to Advanced Manufacturing Security." Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

Portions of this work (graduate student funding) were also provided through internal funding at Virginia Tech.

References

- [1] Evans D. The Internet of Things: How the Next Evolution of the Internet is Changing Everything. CISCO White Pap 2011;1:1–11.
- [2] Bayuk J, Cavit D, Guerrino E, Mahony J, McDowell B, Nelson W, et al. Malware Risks and Mitigation Report. BITS Financ Serv Roundtable, Washington, DC 2011.
- [3] Watin-Augouard M. Prospective Analysis on Trends in Cybercrime from 2011 to 2020. Natl Gendarm 2011.
- [4] Falliere N, Murchu LO, Chien E. W32.Stuxnet Dossier 2011;4:1–69.
- [5] Li C, Raghunathan A, Jha NK. Hijacking an insulin pump: Security attacks and defenses for a diabetes therapy system. e-Health Netw. Appl. Serv. (Healthcom), 2011 13th IEEE Int. Conf., 2011, p. 150–6. doi:10.1109/HEALTH.2011.6026732.
- [6] Zetter K (Wired). A Cyberattack Has Caused Confirmed Physical Damage for the Second Time Ever | WIRED 2015. <http://www.wired.com/2015/01/german-steel-mill-hack-destruction/> (accessed January 7, 2016).
- [7] BSI B für S in der I. Die Lage der in Deutschland 2011. Informationstechnik 2011.
- [8] National Defence Industrial Association's. Cybersecurity for Advanced Manufacturing. White Pap 2014.
- [9] Wang L, Törngren M, Onori M. Current status and advancement of cyber-physical systems in manufacturing. J Manuf Syst 2015;37:517–27. doi:10.1016/j.jmsy.2015.04.008.
- [10] Dworschak B, Zaiser H. Competences for cyber-physical systems in manufacturing – first findings and scenarios. Procedia CIRP 2014;25:345–50. doi:10.1016/j.procir.2014.10.048.
- [11] Wells LJ, Camelio JA, Williams CB, White J. Cyber-physical security challenges in manufacturing systems. Manuf Lett 2013;2:74–7.
- [12] Turner H, Amos B, White J, Camelio J, Williams C. Bad parts: Are our manufacturing systems at risk of silent cyber-attacks. IEEE Secur Priv 2015;40–7. doi:10.1109/MSP.2015.60.
- [13] Zeltmann SE, Gupta N, Tsoutsos NG, Maniatakos M, Rajendran J, Karri R. Manufacturing and Security Challenges in 3D Printing. JOM 2016;68:1872–81. doi:10.1007/s11837-016-1937-7.
- [14] Bridges SM, Hall STAT, Graves SJ, Hall STAT, Keiser K, Hall STAT, et al. Cyber Security for Additive Manufacturing. Proc. 10th Annu. Cyber Inf. Secur. Res. Conf., New York, NY, USA: ACM; 2015, p. 14:1--14:3. doi:10.1145/2746266.2746280.
- [15] Eset Corporation. ACAD/Medre.A. White Pap 2012.
- [16] Mustaca S. Are your IT professionals prepared for the challenges to come? Comput Fraud Secur 2014;2014:18–20. doi:10.1016/S1361-3723(14)70472-5.

- [17] ASTM F2915-12. Standard Specification for Additive Manufacturing File Format (AFM) Version 1.1. ASTM Int 2012;2013:1–15.
- [18] Chhetri SR, Abdullah M, Faruque A. Side-Channels of Cyber-Physical Systems : Case Study in Additive Manufacturing. IEEE Des Test 2017;PP:1–1.
doi:10.1109/MDAT.2017.2682225.
- [19] Sturm LD, Williams CB, Camelio JA, White J, Parker R. Cyber-Physical Vulnerabilities in Additive Manufacturing Systems. Solid Free Fabr Symp 2013:951–63.
- [20] Belikovetsky S, Yampolskiy M. dr0wned – Cyber-Physical Attack with Additive Manufacturing n.d.
- [21] Williams CB, Seepersad CC. Design for Additive Manufacturing Curriculum: A Problem and Project Based Approach. Solid Free Fabr Symp 2012:81–92.
- [22] Trouton S, Vitale M, Killmeyer J. 3D opportunity for blockchain. Deloitte Univ Press 2016.
- [23] Craeghs T, Bechmann F, Berumen S, Kruth J-P. Feedback control of Layerwise Laser Melting using optical sensors. Phys Procedia 2010;5:505–14.
doi:10.1016/j.phpro.2010.08.078.

4. In-situ Monitoring of Additive Manufacturing Processes via Impedance-based Measurements

Coauthors: Dr. Mohammed Albakri, Dr. Pablo Tarazaga, Dr. Christopher Williams,

Abstract

In this paper, the authors explore the use of impedance-based monitoring techniques for in-situ detection of additive manufacturing build defects. By physically coupling a piezoceramic (PZT) sensor to the part being fabricated, the measured electrical impedance of the PZT can be directly linked to the mechanical impedance of the part. It is hypothesized that one can detect build defects in geometry or material properties in-situ by comparing the signatures collected during printing of parts with that of a defect-free control sample. In this paper, the authors explore the layer-to-layer sensitivity for both PZT sensors embedded into printed parts and for a fixture-based PZT sensor. For this work, this concept is evaluated in context of material jetting. A set of control samples is created and used to establish a baseline signature. (e.g., internal voids) are fabricated and their layer-to-layer signatures are compared to a control sample. Using this technique, the authors demonstrate an ability to track print progress and detect defects as they occur. For embedded sensors the defects were detectable at 2.28% of the part volume (95.6 mm³) and by fixture-based sensors when it affected 1.38% of the part volume.

Keywords

Additive Manufacturing, Nondestructive Testing, In-situ Monitoring, Material Jetting, impedance-based monitoring, Photopolymer systems

4.1. Introduction

The maturation and growing adoption of additive manufacturing (AM) as a means for making end-use products has led to an increased need for part traceability and process monitoring [1]. Unlike prototypes or fixtures, end-use products must often meet rigorous performance standards and be certified before they can be put into use. In particular, the aerospace industry has begun to implement AM technologies for fabricating high value, low volume parts. Some examples of AM parts being produced and tested include direct printed metals (e.g., Rolls Royce Trent-XWB bearing [2]) and polymers (e.g., FAA-approved ULTEM 9085 aircraft air duct created by Stratasys and Orbis [3]). While traditional quality control focuses on post-process inspection of parts, in situ monitoring is desirable in AM for two reasons: to detect defects before they become internal to the part, and to save cost by detecting defects before the entire build has been completed.

4.1.1 In situ monitoring of AM

As AM has become more mature, there has been an increasing amount of research into in situ monitoring techniques for all AM processes. In particular, there has been a strong focus on the performance of metal AM processes such as powder bed fusion (PBF) and directed energy deposition (DED). In situ monitoring approaches fall into three categories: surface (monitoring the top layer), volumetric (monitoring through some depth into a part), and indirect (monitoring the machine performance instead of the part). In the first case, the primary focus is with monitoring the formation of a layer either through optical technology or by monitoring the thermal characteristics (e.g. meltpool temperature or extrudate temperature). These monitoring techniques are able to identify defects that occur on the surface of the part, such as geometry changes or small voids, but they lack the ability to directly measure the material properties of the part[4]. Volumetric approaches such as acoustic emission testing, ultrasonic testing, impedance-based testing, eddy current testing, or x-ray computed tomography (CT) can detect internal defects below the surface of a part and may be able to detect changes in the material properties of a part. Indirect monitoring may be able to predict that a part will contain a defect, but it is unable to directly detect defects in the part.

One of the most popular approaches for in-situ monitoring in AM systems is surface monitoring using image based sensors. For powder bed fusion (PBF) and directed energy deposition (DED). A variety of research approaches have been presented [5–7] including optical (using high-speed cameras) [8–11], monitoring of the laser [12], and infrared[13]. Some optical methods focus on monitoring the entire layer while other in-situ systems focus on monitoring the melt pool or melt plume instream of the entire build area since this allows for increased resolution [14]. Another method is through the use of an illumination laser in combination with a high-speed camera [15]. Similar approaches have been used for fused filament fabrication, with optical monitoring [16] and laser scanning [17]. For stereolithography, interference monitoring has been used to detect curing and monitor geometry [18].

The body of research for material jetting is significantly smaller than that for metal or extrusion based systems. For polymer systems, the use of IR LEDs and detectors to monitor spatial location of droplets has been demonstrated [19]. For liquid metal jetting, a CCD camera and strobing LED have been used to monitor droplet shape for quality assurance purposes [20]. In both cases, these approaches are limited to monitoring single or small numbers of nozzles and are currently not suited for the large jetting arrays used in many commercial systems. These systems also do not monitor the material properties of the process. While optical techniques limit inspection to the surface geometry of the part at each layer, when used in-situ, surface techniques are able to monitor cross-sections of the part throughout the print. In this way, in-situ surface techniques are able to approximate volumetric inspection of the part. The limitation of this approach is that surface inspection is unable to capture any changes that occur below the

layer being monitored. Surface monitoring cannot directly detect thermal or other effects that propagate down through many layers.

With machine behavior, the focus is on detecting when the machine is operating outside of normal parameters. In traditional manufacturing processes, like CNC milling, the material properties are fixed and only the geometry and surface finish need to be monitored. In AM both the material properties and the geometry can be affected by the process parameters and toolpath used to fabricate the geometry can affect the internal material properties of the part. Monitoring process parameters can detect potential build failures, but does nothing to directly monitor the material properties of the part[21].

Other approaches exist to monitor the thermal characteristics of the system using thermocouples and pyrometers [7,22,23]. Ultrasonic [24–26] and laser ultrasonic [4] testing has also been investigated as a means for detecting surface defects and they have some capability to detect volumetric properties such as thickness and internal defects. Acoustic emission (AE) testing is another vibration-based approach that detects energy released from irreversible changes that occur on a micro scale in a part when it is loaded (e.g. crack growth) [27,28]. Thermography can detect internal defects close to the surface by monitoring the heating and cooling of an object, as defects cause irregularities compared to the bulk material properties [27]. More exotic methods of in-situ inspection have also been considered, such as neutron diffraction and X-ray backscattering [29]. These methods can detect defects volumetrically, but require expensive radiation sources. Some work outside of direct-metal AM has been conducted using a variety of sensors and integrating their results together. This work has demonstrated the use of data fusion techniques to leverage signals from thermocouples, accelerometers, an IR temperature sensor, and a borescope to monitor and detect build defects in an material extrusion process [30,31].

An ideal monitoring technique for AM would be able to monitor material properties and geometry layer-by-layer. This technique could be used to qualify parts and would be inexpensive and non-invasive. While current monitoring techniques are able to achieve some of these qualities, they are limited in others (e.g. optical layer-by-layer imaging cannot interrogate material properties throughout the part, x-ray CT scanning can volumetrically image a part, but is expensive [32]).

4.1.2 Impedance-based monitoring

Based on the previous work that has been done on in-situ monitoring, the authors believe that there is still a need for additional in-situ monitoring techniques that are able to directly interrogate both the material properties and the geometry of the entire part as it is being fabricated. One technique that has been used successfully in structural health monitoring (SHM) applications is impedance-based monitoring. This technique utilizes piezoelectric materials,

specifically lead zirconate titanate (PZT) wafers, as collocated sensors and actuators to simultaneously excite the structure and measure its response[33,34].The fundamental basis for impedance based SHM is that the presence of damage will alter the mass, stiffness, or dampening characteristics of the structure, which in turn reflects on its measured dynamic response. Impedance based SHM has been shown to be a promising, non-intrusive, cost-effective, and sensitive solution for real-time damage assessment[35].

Ultrasonic testing and acoustic emission testing share some similarities with impedance based monitoring (e.g. all are vibration-based approaches) however, there are several important differences. Acoustic emission occurs when a material undergoes irreversible changes in its internal structure during mechanical loading. This limits detection to dynamic defects (i.e. a defect can only be detected once and if missed, higher loading is need). While AE can detect very small cracks (25 um) it requires a reference signature and is best suited for continuous testing. Ultrasonic testing is able to detect internal defects by propagating waves through a part and monitoring their reflections. This allows for accurate localization of defects within a part; however it also means that the waves need a clear path to be able to travel along. In a complex part like a lattice structure, waves may be unable to reach a location hidden from “line-of-sight” and therefore any defects in that location will be undetectable [36]. Impedance based monitoring is steady state instead of dynamic and monitors the impedance response of the entire part at a given frequency [37,38]. This allows it to repeatedly identify the same defect (unlike AE) and to detect defects that are hidden from “line-of-sight” since they affect the overall response of the part.

In an AM context, each part design has a distinct mass, stiffness, and dampening response that can be detected through impedance-based monitoring. The introduction of a defect into a part will cause a corresponding change in the response. By physically attaching a PZT sensor to a part, it is possible to couple the mechanical response of the system to the electrical response of the sensor. In a manufacturing context, this means that a simple electrical signal can be used to evaluate the mechanical properties of a part without the need for destructive testing. The authors’ previous work has established that piezo-ceramic sensors can be used as a post process NDE tool to detect defects in fully fabricated AM parts[39]. An initial series of defect free control parts, are fabricated and measured using this technique to establish a baseline response signature. These control parts can be validated as needed with methods such as CT scanning and destructive testing. Once a baseline signature has been established, the signature of each subsequent part is compared to the baseline to determine if a defect has occurred. In post process inspection, this technique detected defects as small as 8 mm³ in polymer AM parts fabricated by material jetting and material extrusion [39].

The overall goal of this paper is to assess the feasibility of using electromechanical impedance measurements for in-situ non-destructive evaluation of AM parts. For in-situ impedance-based monitoring a similar approach is proposed. After mounting, sensors monitor several control parts at select layers during fabrication. These control parts are validated and a baseline signature is generated for each layer. As new parts are fabricated their signatures are compared to the baseline signature at each monitored layer. If the variation in the part signature exceeds the established variation threshold it indicates the presence of a build defect.

There are three primary goals for this research. The first goal is to assess the feasibility of using impedance-based monitoring in-situ. The second is to compare approaches for mounting the sensors in-situ, specifically by comparing embedded sensors to fixture based sensors. The third is to determine the size of the defects that are detectable when using impedance based measuring in-situ on plastic AM parts, specifically VeroWhite parts made using material jetting.

4.2 EXPERIMENTAL METHODS

4.2.1 Method Overview

The first goal of this study is to evaluate the suitability of using impedance-based monitoring in-situ on an AM system. To investigate this, the authors' designed (i) a test part that a sensor could be embedded into during fabrication, and (ii) a fixture with a mounted sensor on which parts could be fabricated. After embedding the sensor, multiple measurements are taken to ensure that a consistent baseline can be obtained.

The second goal is to compare two methods of in-situ monitoring, embedded sensors and fixture-based sensors. In the first case, the sensor is embedded in a small cavity in the part during fabrication (Figure 1 A). This method permanently attaches the sensor to the part and requires a separate sensor for each test part. In the second case, a sensor is mounted to a fixture and the parts are directly fabricated on the fixture (Figure 1 B). After fabrication, the part is removed and the fixture is reused. In this way a single sensor (or set of sensors) can be used to measure a large number of parts. To compare these two methods, two identical sets of parts are fabricated: one using the embedded technique and the other using the fixture technique and the sensitivity of both methods are compared.




Figure 4.1. A) Example of embedded sensor method B) Example of fixture based sensor method

The final goal is to determine the resolution of defects that in-situ sensing can detect. To accomplish this, a defect with an increasing size is fabricated. During fabrication, measurements are taken at set intervals and compared to the baseline signature. When the deviation from the baseline becomes significant, the layer number is noted and used to calculate the size of the defect based on the model. Since the defect is increasing in size, the first layer that the defect is detected on indicates the size of the smallest detectable defect.

4.2.2 Material Jetting Process

Material Jetting (MJ) is an AM process that uses an array of nozzles to selectively deposit material (usually a resin) that is cured by a broad area energy source (e.g. a UV lamp). These systems can offer high resolution (droplet and layer sizes < 100 microns) and are used for fabricating dental and medical models [40]. Material jetting systems also have the ability to fabricate models using multiple materials, which allows for performance testing of composite (hard/soft) designs [41,42]. In the context of this study, material jetting was chosen due to accessibility of the system for embedding sensors and for the ability to selectively change the material/stiffness of a part while maintaining the overall geometry and mass. While material jetting was selected, the monitoring process as used should also be applicable to extrusion systems, and potentially stereolithography systems.

4.2.3 Materials

All parts were printed using a Stratasys Connex 350 with a resolution of 300 x 600 DPI in the x-y plane with 0.03 mm layers [43]. The material used was VeroWhite [44] (a hard acrylate based photopolymer), and SUP705 support material with a matte finish. VeroWhite was chosen because it is a standard material used for fabricating parts on this system. While the specific results, such as sensitivity, cannot be directly transferred to other materials (e.g. stiffer materials may be more sensitive, softer materials may be less), the general approach and methods should

be applicable to different types of materials as demonstrated in the authors' previous work [39]. The piezoelectric material was 0.1905 mm thick lead zirconate titanate (PZT) wafers[45] cut into 20 mm x 8 mm pieces. Cyanoacrylate was used to bond the sensors to the part and the fixture. Copper tape with a conductive adhesive and flux core solder were used to attach electrodes to the PZTs. The fixture consisted of 1/8" acrylic sheet and 1/16" stainless steel cut into a 2" x 3" section and mounted on top of a 3D printed stand.

4.2.4. Test specimen design

For the embedded sensing, two test parts were designed, (i) a control (Figure 2A, "A") and (ii) a triangular prism cavity (Figure 2B, "B"). The triangular cavity feature was chosen so that the defect would start small and grow increasingly larger with each layer. The defect simulates the effect of a void being placed in the model file. Both parts also contained a small cavity at the base for embedding the PZT sensor/actuator. The piezos are embedded after 2.5 mm (layer# 83) of the parts have been printed. The defects begin after 4.5 mm (layer# 151). The thickness of the wall containing the defects is 1 mm. Part "B" has a right triangular cross-section with dimensions shown in figure 2B. The total volume of the defect is 2665 mm³ (26.7% of total volume). Due to the ramp shape, the size of the fabricated defect increases with each layer and the smallest increase in defect size in a layer is 0.027 mm³ (0.00027% of total volume).




Figure 4.2. Test parts A) Control sample "A", B) Triangular prism cavity "B" (lines indicate defect dimensions)

4.2.4.1 Embedded Measurements

Because the defect is enclosed, it will be filled with a supporting material that has similar density, but significantly different stiffness. The mass and geometry of the part remain the same. The defect does simulate how a void in the model file would be fabricated (e.g. through a malicious cyber-attack [46]), and the substitution of model material for support material weakens a part in a similar way that an empty void would.

Three copies of each test part were printed simultaneously, arrayed in a random layout. The print was paused at the completion of the small cavity (2.5 mm, layer# 83) and support material was removed. The sensors were bonded to the parts and allowed to set for one hour. After resuming printing, measurements were taken at the layers shown in Table 1. To avoid any

interference in the signal caused by the machine operating, the printer was paused before each measurement and resumed after all parts had been measured (~5s per part).

4.2.4.2 Fixture-based Measurements

For the fixture-based sensing, a PZT sensor was prepared similarly to the embedded sensors, but was mounted to a steel sheet using cyanoacrylate. Steel was chosen as the fixture material because initial experiments with acrylic demonstrated that the stiffer material transferred vibrations better, resulting in better signatures. After the sensors were mounted to the sheet, the piece was placed on top of a 3D printed fixture aligned to the top left corner of the build area. Parts were printed one at a time and a single measurement was taken at each layer shown in Table 1 without pausing the printer. A total of three control samples and three defective parts were printed on the fixture for a total of six prints. To ensure good transfer of vibration, the standard pedestal of support material was removed for the fixture prints and they were built directly on the fixture surface. The defect starts in layer 151, which is 4.23mm into the print.

4.2.5 Impedance measurements and analysis

The response of the sensor was measured using a Keysight E4990A impedance analyzer over a frequency range of 10 to 100 kHz. Each embed part was measured (~5s per measurement) three times at the end of each of the layers shown in Table 1 and the mean of the measurements was used for comparison. The fixture-based parts were measured a single time after the layer while the print was running. To quantify the difference between signatures, the damage metric defined in Equation 1 was used:

$$RMSD = \sqrt{\sum \frac{(Z_D - Z_{BL})^2}{Z_{BL}^2}} \quad (1)$$

Where Z_D is the impedance at a given frequency of the part being measured and Z_{BL} is the impedance at a given frequency of the baseline (established by averaging together the control samples). This damage metric is explained more in detail in the authors' previous work[39]. After establishing a baseline signature and variance with the control samples new specimens are compared to the baseline and if the damage metric is greater than the sample variance it indicates the presence of a defect.

Table 4.1. Part layers where measurements were taken. For each layer, the volume of model material in each part is shown along with the size of the defect, and the percentage of the printed material that the defect represents.

Part Height		Model Material		Defect Size
mm	Layer	Control	Defect	
4.20	140	2793.0	2793.0	0.00%
4.35	145	2892.8	2892.8	0.00%
4.50	150	2992.5	2992.5	0.00%
4.65	155	3092.3	3091.6	0.02%
4.80	160	3192.0	3189.3	0.08%
4.95	165	3291.8	3285.8	0.18%
5.10	170	3391.5	3380.9	0.31%
5.25	175	3491.3	3474.7	0.48%
5.40	180	3591.0	3567.1	0.67%
5.55	185	3690.8	3658.2	0.88%
5.70	190	3790.5	3748.0	1.12%
5.85	195	3890.3	3836.5	1.38%
6.00	200	3990.0	3923.6	1.66%
6.30	210	4189.5	4093.9	2.28%
6.60	220	4389.0	4258.9	2.96%
6.90	230	4588.5	4418.6	3.70%
7.20	240	4788.0	4573.0	4.49%
7.50	250	4987.5	4722.0	5.32%

4.3 RESULTS

3.1 Embedded piezo repeatability

Figure 3 shows an example of an impedance signature taken at (layer 150#, embedded part A1). As shown in figure 3. multiple measurements of a single part at a set layer had very little deviation. The small variations that do occur can be attributed to normal noise experienced by the analyzer when performing impedance measurements. This lack of deviation shows that there is little to no machine interference or sources of random error occurring in the measurements, outside of what is normally expected from the monitoring equipment.

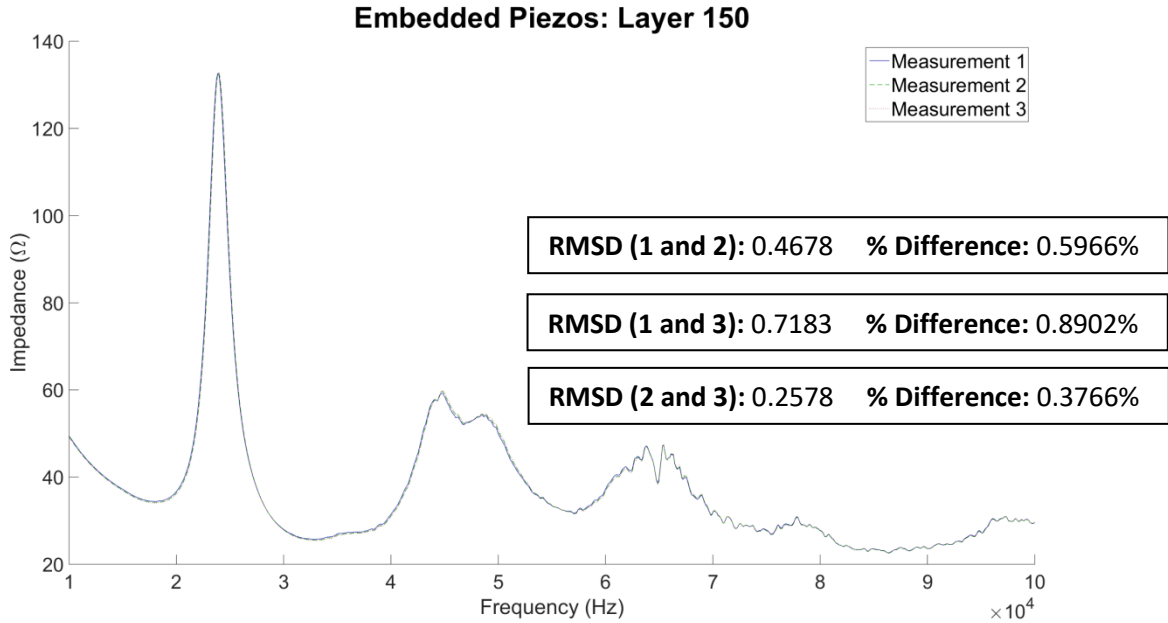


Figure 4.3. A comparison of multiple signatures taken at the same layer from the same part using an embedded sensor.

4.3.1 Embedded Piezos

Figure 4a. provides a visualization of the damage metric (eqn. 1) calculated from comparing control to defect specimens each part at each layer. The y-axis indicates the sample and the x-axis indicates the layer number. Thus each cell in the array represents the comparison of a signature, like the one shown in figure 3., to the baseline signature using the RMSD (equation 1) damage metric. In this representation, dark purple (as shown on the baseline row) indicates little to no difference from the baseline signature at that layer. Brighter colors indicate increasingly large deviations from the baseline signature. The difference is deemed significant when the damage metric indicates a greater change in the defective parts than the variance in the control parts.

The red line indicates the layer at which the defect starts to be fabricated (layer# 151) and the red oval indicates the point where the damage metric reports a significant change (i.e. the RMSD is greater than that of the controls). Using the signature from the embedded piezos, the damage metric was able to detect the triangular prism cavity at layer 210 as shown in Figure 4a. The defect size when detected is 95.6 mm³ (2.28% of printed volume)(Table 1). Figure 4b. shows the signature for the controls, baseline, and defective parts at layer 245.

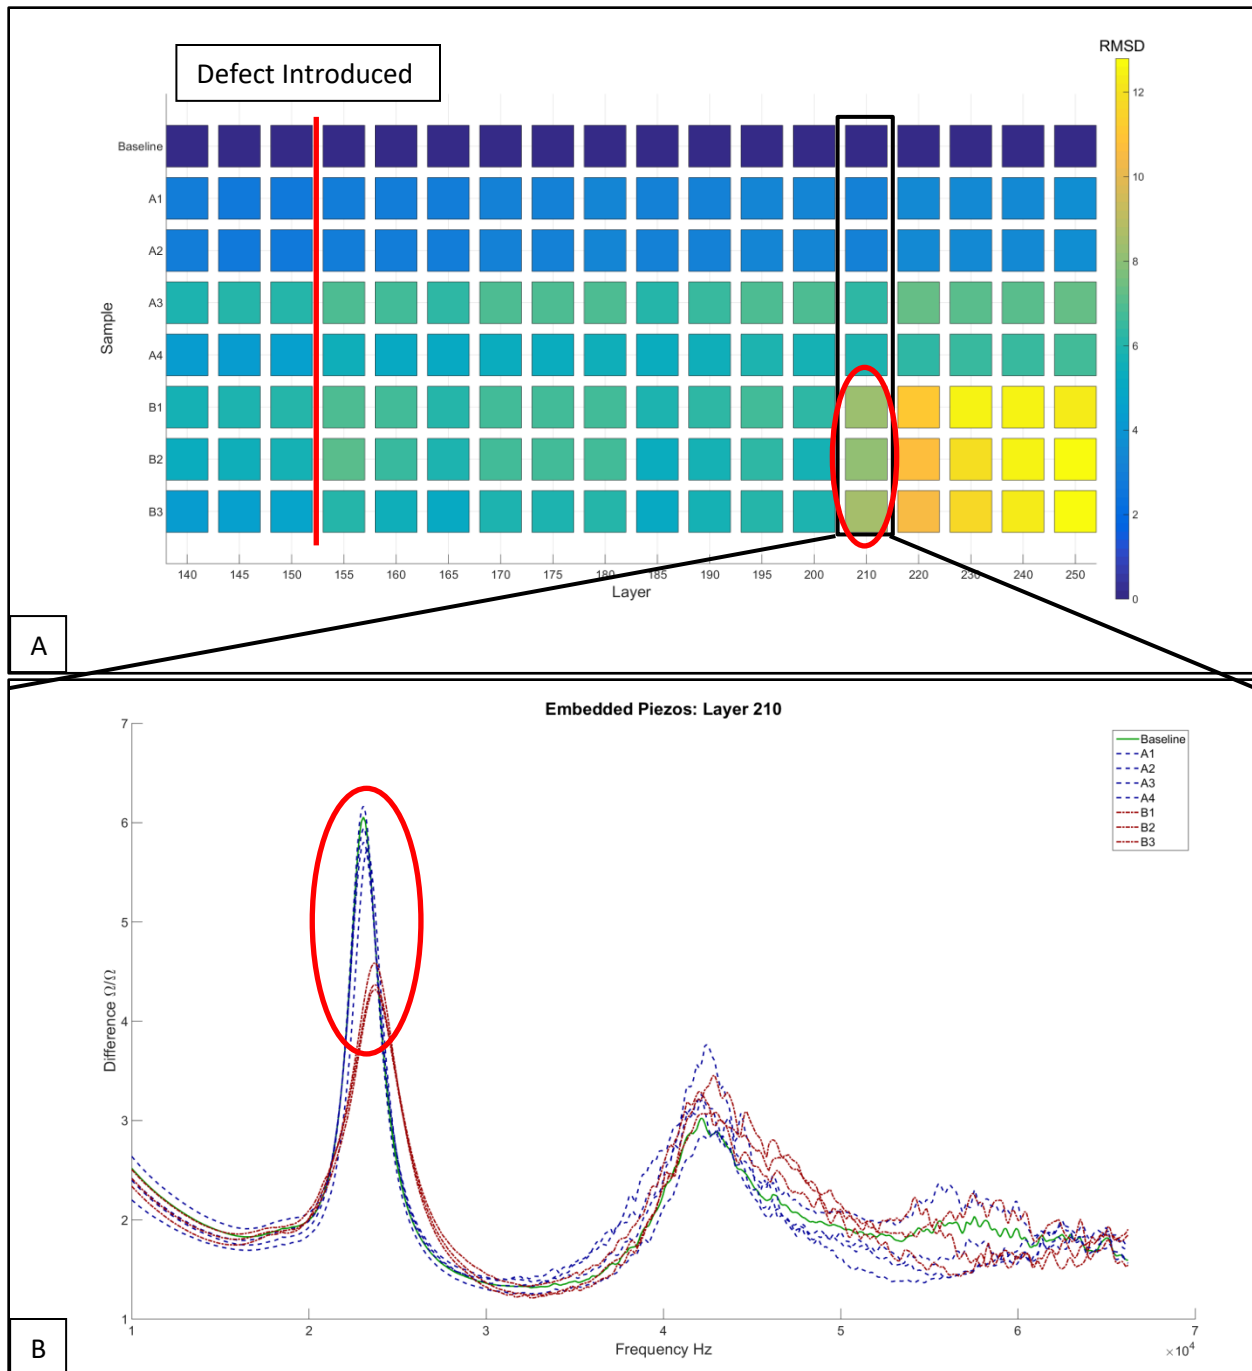


Figure 4.4. A) Difference comparison for embedded piezos between baseline, control samples, and defect samples across each layer. The defect is introduced in layer 151 (as indicated by the red line) and detectable in all samples at layer 210 (red oval). The defect size when detected is 95.6 mm³ (2.28% of printed volume). B) The part signatures at layer 210. The red oval indicates the area where the defect is evident.

4.3.2 Fixture-based Piezos

When using the damage metric with the steel fixture, it was able to detect the defect at layer 195 (Figure 5a.) This corresponds to a defect of size 53.8 mm^3 (1.38% of printed volume). Mounting the piezo to a steel fixture results in significantly more distinct peaks (due to the higher stiffness of the material), the variation between measurements is also significantly reduced. The increased mass of the system significantly reduces the magnitude of the differences. For the best results it is necessary to identify the peaks where the defect is manifested and to compare the curves in those areas, otherwise the noise of small sharp peaks may reduce the sensitivity. By visual comparison of the signatures, it was determined that the defect was manifested in the signature primarily in the following four frequency ranges, 30-33 kHz, 37-40 kHz, 41-44 kHz, and 53-56 kHz. This subset of frequencies was used to calculate the damage metric, which showed a significant difference between the baseline and the defective parts.

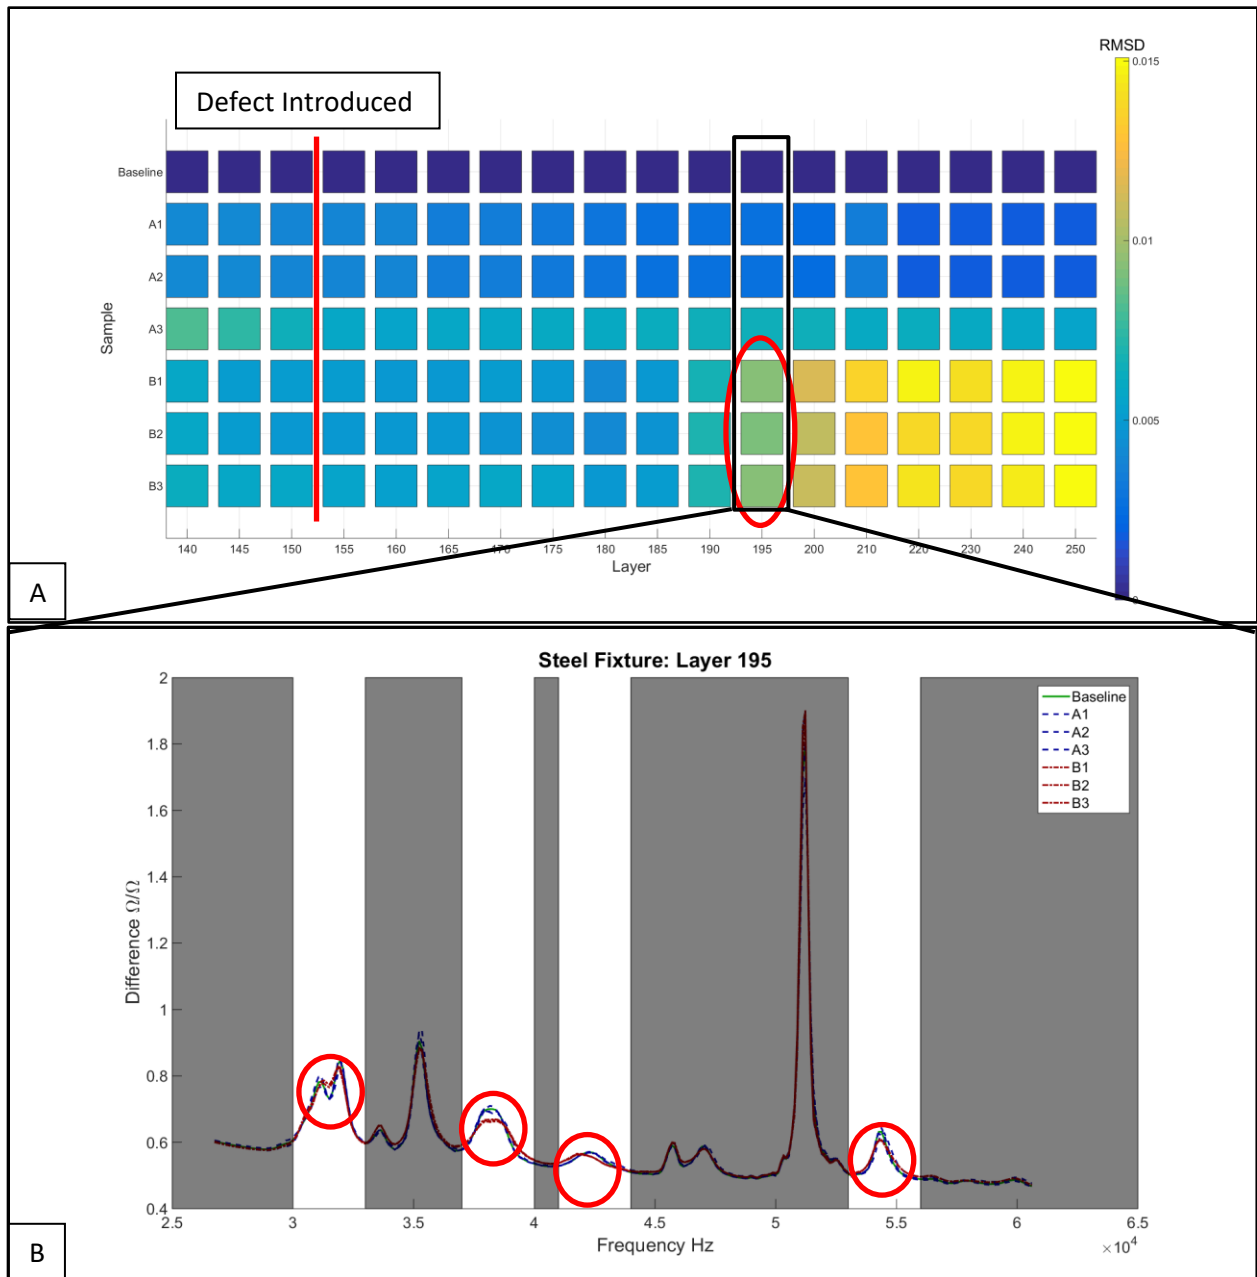


Figure 4.5. A) Difference comparison for embedded piezos between baseline, control samples, and defect samples across each layer. The defect is introduced in layer 150 (as indicated by the red line) and detectable in all samples at layer 245 (red oval). The defect size when detected is 53.8 mm^3 (1.38% of printed volume). B) The part signatures at layer 195. The red oval indicates the area where the defect is evident. White areas indicate frequency ranges that were used in analysis.

4.4 DISCUSSION

By embedding piezos into printed parts it was possible to detect changes in parts' stiffness introduced by creating cavities containing support material. These defects were detectable as small as 95.6 mm^3 (2.28% of printed volume). Natural variation between each sensor, and differences in how the piezos are mounted, can cause variation in sensor readings. These natural variations introduce noise that effectively reduces the sensitivity of the technique. While smaller print defects might be detectable, the sensor variation results in the resulting signal being undistinguishable from that of the control specimen. Another drawback of the embedding method is that it requires modification of the model design to accommodate the sensor and for the build to be paused to allow the sensor to be embedded.

The fixture-based method addresses these issues by using the same piezo sensor across multiple parts. This eliminates the variation that occurs between sensors and keeps the mounting constant. There still is some variation based on the adhesion of the part to the fixture. This reduction in variation allowed for the defect to be detected 15 layers earlier, at layer 195 instead of layer 210. The use of a fixture adds additional mass to the system, which has the potential to reduce the sensitivity of the system. The signatures showed more peaks in the fixture than for the embedded piezos. This is due to the stiffer steel fixture having more resonance at higher frequencies than the model material. The presences of these peaks present more areas where changes might be detected. The material of the fixture is quite important in transferring the vibrations from the piezo to the part. When using an acrylic fixture, the dampening effect was significant enough to render defects undetectable from baseline variation. A steel fixture has significantly greater stiffness, which allows for better transfer of vibrations from the sensor to the part. Using a steel fixture, it was possible to detect a 53.8 mm^3 defect (1.38% of printed volume). The steel fixture has the ability to detect smaller defects than the embedded piezos. A drawback of the fixture method is the need for the part to be located in the same position each time. Small variation in location will cause changes in the signature. This variation introduces noise that can make the detection of very small defects impossible.

The authors' previous work using impedance-based monitoring as a post processing inspection technique were able to detect internal "voids" as small as 8 mm^3 (0.083% of volume) for the VeroWhite material on the Connex. While the current in-situ work only able to detect defects as small as 53.8 mm^3 , the prior work shows that improved resolution should be possible. While this is large compared to the size of droplets ($\sim 100 \text{ }\mu\text{m}$), the nature of material jetting means that small geometric defects will be filled by fluid flow, while large defects will be visible on the surface. Changes to the material properties will be more likely to affect entire layers resulting in much larger affected areas (a single layer of the test part is 19.95 mm^3). Part of for

less sensitivity can be attributed to the fact that these techniques are highly dependent on the stiffness of the material being used. The support structure used in Material Jetting has a high amount of dampening, which reduces the sensitivity of the sensors. In post process application, the surrounding support material is removed before the sensors are attached, which may help increase the sensitivity. For the parts with embedded sensors another factor that may reduce the sensitivity is the difficulty of cleaning the embedding location of support material in-situ. Any residual support material would interfere with the attachment of the sensor to the part and could cause variation that would reduce the sensitivity.

The authors believe that the sensitivity of the in-situ measurements could be brought close to that of the post process application with further refinements of the process. It should be noted that these sensitivity results are specific to the material being used, in this study VeroWhite. The use of stiffer materials, such as those found in direct metal AM processes, can significantly improve the ability of the technique to detect small defects. This was shown in the difference between the acrylic and steel fixtures.

Another opportunity for improving sensitivity is in the calculation of the damage metric. While the current technique is capable of detecting differences between signatures, it gives more weight to RMS magnitude differences than changes to the shape of the curve. These shape changes, while potentially small in magnitude, are indicators of significant changes in the response of the part. In parts fabricated on the fixture, the changes in the signatures are happening in a small subset of the frequency range tested. By reducing the range measured, similar to the post-process measurements, and increasing the resolution of the measurements, it should be possible to detect smaller defects without increasing the time needed to acquire the measurements.

4.5 CONCLUSIONS

The layer-by-layer fabrication process of AM systems makes a volumetric evaluation of part quality important. With existing in-situ techniques it can be difficult to directly measure the quality of a part. Because impedance monitoring is linked to the mass, stiffness, and damping of a part, it is able to detect both geometric changes and material property changes volumetrically throughout a part. This ability allows impedance-based monitoring to be used as a side-channel technique (an indirect measurement that can be correlated to the desired properties). The study successfully demonstrated the ability of impedance-based techniques to be used as an in-situ monitoring approach for AM. The study demonstrated that both embedded and fixture based approaches were feasible in a Material Jetting system, with the fixture based method being more sensitive to internal defects. Specifically, internal defects could be detected by embedded piezos when they affected 2.28% of the part volume (95.6 mm^3) and by steel fixture-based piezos when they affected 1.38% of the part volume (53.8 mm^3). While this was demonstrated for simple parts

the authors' previous work shows that this approach can also be used on complexed geometries [39]. By inspection it was possible to determine the frequency ranges where the defect was manifested in the signature. By narrowing the sampled frequency range it is possible to increase the resolution without increasing the time required. Increasing the sampling resolution should make it possible to increase the resolution of the detection method. The improved performance by increasing the stiffness of the fixture indicates that higher resolution detection should be possible in parts fabricated out of stiffer materials, such as metals.

In future work the authors hope to expand the application of this technique to other AM processes, to examine additional types of defects, and to improve the sensitivity by increasing the size/number of the sensor(s), by refining the damage metric to be able to more precisely determine when variation is occurring the signature, and by refining the frequency range being measured.

This material is based upon work supported by the National Science Foundation under Grant No. CMMI-1436365 and Grant No. CMMI-1635356.

Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

References

- [1] NIST, Measurement Science Roadmap for Metal-based Additive Manufacturing, 2013. http://www.nist.gov/el/isd/upload/NISTAdd_Mfg_Report_FINAL-2.pdf.
- [2] D. Howie, High powered Trent XWB-97, (2015) 12–15. <http://www.rolls-royce.com/media/insights/simon-burr.aspx> (accessed January 31, 2017).
- [3] Stratasys, FAA-Approved Air Duct for 'Flying Eye Hospital' Produced in Just Days, (2015). <http://blog.stratasys.com/2015/03/05/3d-printed-air-duct-flying-eye-hospital/> (accessed January 31, 2017).
- [4] S.K. Everton, M. Hirsch, P. Stravroulakis, R.K. Leach, A.T. Clare, Review of in-situ process monitoring and in-situ metrology for metal additive manufacturing, *Mater. Des.* 95 (2016) 431–445. doi:10.1016/j.matdes.2016.01.099.
- [5] G. Tapia, A. Elwany, A Review on Process Monitoring and Control in Metal-Based Additive Manufacturing, *J. Manuf. Sci. Eng.* 136 (2014) 60801–60810. doi:10.1115/1.4028540.
- [6] E.W. Reutzler, A.R. Nassar, A survey of sensing and control systems for machine and

- process monitoring of directed-energy, metal-based additive manufacturing, *Rapid Prototyp. J.* 21 (2015) 159–167. doi:10.1108/RPJ-12-2014-0177.
- [7] Z.Y. Chua, I.H. Ahn, S.K. Moon, Process monitoring and inspection systems in metal additive manufacturing: Status and applications, *Int. J. Precis. Eng. Manuf. - Green Technol.* 4 (2017) 235–245. doi:10.1007/s40684-017-0029-7.
- [8] S. Berumen, F. Bechmann, S. Lindner, J.-P. Kruth, T. Craeghs, Quality control of laser- and powder bed-based Additive Manufacturing (AM) technologies, *Phys. Procedia.* 5 (2010) 617–622. doi:10.1016/j.phpro.2010.08.089.
- [9] T. Craeghs, S. Clijsters, J.-P. Kruth, F. Bechmann, M.-C. Ebert, Detection of process failures in Layerwise Laser Melting with optical process monitoring, *PhD.* 39 (2012) 753–759. doi:10.1016/j.phpro.2012.10.097.
- [10] J. Schwerdtfeger, R.F. Singer, C. Körner, In situ flaw detection by IR-imaging during electron beam melting, *Rapid Prototyp. J.* 18 (2012) 259–263. doi:10.1108/13552541211231572.
- [11] R.B. Dinwiddie, R.R. Dehoff, P.D. Lloyd, L.E. Lowe, J.B. Ulrich, Thermographic in-situ process monitoring of the electron-beam melting technology used in additive manufacturing, *Proc. SPIE.* 8705 (2013) 87050K-87050K–9. doi:10.1117/12.2018412.
- [12] A.G. Demir, C. De Giorgi, B. Previtali, Design and implementation of a multi-sensor coaxial monitoring system with correction strategies for selective laser melting of a maraging steel, *J. Manuf. Sci. Eng.* 140 (2017) 1–14. doi:10.1115/1.4038568.
- [13] M. Grasso, A.G. Demir, B. Previtali, B.M. Colosimo, In situ monitoring of selective laser melting of zinc powder via infrared imaging of the process plume, *Robot. Comput. Integr. Manuf.* 49 (2018) 229–239. doi:10.1016/j.rcim.2017.07.001.
- [14] M. Khanzadeh, S. Chowdhury, M.A. Tschopp, H.R. Doude, M. Marufuzzaman, L. Bian, In-situ monitoring of melt pool images for porosity prediction in directed energy deposition processes, *IISE Trans.* 0 (2018) 1–19. doi:10.1080/24725854.2017.1417656.
- [15] P. Lott, H. Schleifenbaum, W. Meiners, K. Wissenbach, C. Hinke, J. Bültmann, J. Bültmann, Design of an Optical system for the In Situ Process Monitoring of Selective Laser Melting (SLM), *Phys. Procedia.* 12 (2011) 683–690. doi:10.1016/j.phpro.2011.03.085.
- [16] S. Nuchitprasitchai, M. Roggemann, J.M. Pearce, Factors effecting real-time optical monitoring of fused filament 3D printing, *Prog. Addit. Manuf.* 2 (2017) 133–149. doi:10.1007/s40964-017-0027-x.
- [17] M. Faes, F. Vogeler, K. Coppens, H. Valkenaers, E. Ferraris, W. Abbeloos, T. Goedeme, Process Monitoring of Extrusion Based 3D Printing via Laser Scanning, *Proc. PMI 2014 Conf.* (2014). doi:10.13140/2.1.5175.0081.
- [18] X. Zhao, D.W. Rosen, Real-time interferometric monitoring and measuring of photopolymerization based stereolithographic additive manufacturing process: Sensor

- model and algorithm, *Meas. Sci. Technol.* 28 (2017). doi:10.1088/0957-0233/28/1/015001.
- [19] A. Wang, T. Wang, C. Zhou, W. Xu, LuBan: Low-Cost and In-Situ Droplet Micro-Sensing for Inkjet 3D Printing Quality Assurance, *Proc. 15th ACM Conf. Embed. Netw. Sens. Syst.* (2017) 27:1--27:14. doi:10.1145/3131672.3131686.
- [20] T. Wang, T.H. Kwok, C. Zhou, S. Vader, In-situ droplet inspection and closed-loop control system using machine learning for liquid metal jet printing, *J. Manuf. Syst.* 47 (2018) 83–92. doi:10.1016/j.jmsy.2018.04.003.
- [21] S.B. Moore, J. Gatlin, S. Belikovetsky, M. Yampolskiy, W.E. King, Y. Elovici, Power Consumption-based Detection of Sabotage Attacks in Additive Manufacturing, (2017) 1–19. <http://arxiv.org/abs/1709.01822>.
- [22] T. Furumoto, T. Ueda, M.R. Alkahari, A. Hosokawa, Investigation of laser consolidation process for metal powder by two-color pyrometer and high-speed video camera, *CIRP Ann. - Manuf. Technol.* 62 (2013) 223–226. doi:10.1016/j.cirp.2013.03.032.
- [23] T. Furumoto, T. Ueda, N. Kobayashi, A. Yassin, A. Hosokawa, S. Abe, Study on laser consolidation of metal powder with Yb: fiber laser—Evaluation of line consolidation structure, *J. Mater. Process. Technol.* 209 (2009) 5973–5980. doi:10.1016/j.jmatprotec.2009.07.017.
- [24] H. Rieder, A. Dillhöfer, M. Spies, J. Bamberg, T. Hess, M.S. Bamberg, H. Rieder, A. Dillhöfer, M. Spies, J. Bamberg, T. Hess, Ultrasonic online monitoring of additive manufacturing processes based on selective laser melting, *AIP Conf. Proc.* 1650 (2015) 184–191. doi:10.1063/1.4914609.
- [25] H. Rieder, M. Spies, J. Bamberg, B. Henkel, On- and offline ultrasonic characterization of components built by SLM additive manufacturing, *AIP Conf. Proc.* 1706 (2016). doi:10.1063/1.4940605.
- [26] S. Kenderian, O. Esquivel, K.R. Olson, E.C. Johnson, A general overview of some Nondestructive Evaluation (NDE) techniques for materials characterization, *Opt. Mater. Struct. Technol. IV.* 7425 (2009) 742506. doi:10.1117/12.826906.
- [27] Q.Y. Lu, C.H. Wong, Additive manufacturing process monitoring and control by non-destructive testing techniques: challenges and in-process monitoring, *Virtual Phys. Prototyp.* 13 (2018) 39–48. doi:10.1080/17452759.2017.1351201.
- [28] S.A. Shevchik, C. Kenel, C. Leinenbach, K. Wasmer, Acoustic emission for in situ quality monitoring in additive manufacturing using spectral convolutional neural networks, *Addit. Manuf.* 21 (2018) 598–604. doi:10.1016/j.addma.2017.11.012.
- [29] T. Watkins, H. Bilheux, K. An, A. Payzant, R. Dehoff, C. Duty, W. Peter, C. Blue, C. Brice, Neutron Characterization for Additive Manufacturing, 171 (2013) 23–27. <https://ntrs.nasa.gov/archive/nasa/casi.ntrs.nasa.gov/20140005932.pdf>.
- [30] P.K. Rao, J. (Peter) Liu, D. Roberson, Z. (James) Kong, C. Williams, Online Real-Time

- Quality Monitoring in Additive Manufacturing Processes Using Heterogeneous Sensors, *J. Manuf. Sci. Eng.* 137 (2015) 61007–61012. doi:10.1115/1.4029823.
- [31] Z. Li, Z. Zhang, J. Shi, D. Wu, Prediction of surface roughness in extrusion-based additive manufacturing with machine learning, *Robot. Comput. Integr. Manuf.* 57 (2019) 488–495. doi:10.1016/j.rcim.2019.01.004.
- [32] A. Thompson, I. Maskery, R.K. Leach, X-ray computed tomography for additive manufacturing: A review, *Meas. Sci. Technol.* 27 (2016). doi:10.1088/0957-0233/27/7/072001.
- [33] C. Liang, F.P. Sun, C.A. Rogers, Coupled Electro-Mechanical Analysis of Adaptive Material Systems-Determination of the Actuator Power Consumption and System Energy Transfer, *J. Intell. Mater. Syst. Struct.* 8 (1997) 335–343. <https://doi.org/10.1177/1045389X9700800406>.
- [34] V. Giurgiutiu, A.N. Zagari, Characterization of piezoelectric wafer active sensors, *J. Intell. Mater. Syst. Struct.* 11 (2001) 959–976. doi:10.1106/A1HU-23JD-M5AU-ENGW.
- [35] G. Park, H. Sohn, C.R. Farrar, D.J. Inman, Overview of piezoelectric impedance-based health monitoring and path forward, *Shock Vib. Dig.* 35 (2003) 451–463. doi:10.1177/05831024030356001.
- [36] Z. Bai, S. Chen, L. Jia, Z. Zeng, Phased array ultrasonic signal compressive detection in low-pressure turbine disc, *NDT E Int.* 89 (2017) 1–13. doi:10.1016/j.ndteint.2017.03.002.
- [37] D.M. Peairs, P.A. Tarazaga, D.J. Inman, Frequency Range Selection for Impedance-Based Structural Health Monitoring, *J. Vib. Acoust.* 129 (2007) 701. doi:10.1115/1.2775506.
- [38] J.P. Nokes, G.L. Cloud, The application of interferometric techniques to the nondestructive inspection of fiber-reinforced materials, *Exp. Mech.* 33 (1993) 314–319. doi:10.1007/BF02322147.
- [39] M.I. Albakri, L.D. Sturm, C.B. Williams, P.A. Tarazaga, Impedance-based non-destructive evaluation of additively manufactured parts, *Rapid Prototyp. J.* 23 (2017) 589–601. doi:10.1108/RPJ-03-2016-0046.
- [40] M.I. Mohammed, B. Cadd, G. Peart, I. Gibson, Augmented patient-specific facial prosthesis production using medical imaging modelling and 3D printing technologies for improved patient outcomes, *Virtual Phys. Prototyp.* 13 (2018) 164–176. doi:10.1080/17452759.2018.1446122.
- [41] V. Dikshit, A.P. Nagalingam, Y.L. Yap, S.L. Sing, W.Y. Yeong, J. Wei, Crack monitoring and failure investigation on inkjet printed sandwich structures under quasi-static indentation test, *Mater. Des.* 137 (2018) 140–151. doi:10.1016/j.matdes.2017.10.014.
- [42] V. Dikshit, A.P. Nagalingam, Y.L. Yap, S.L. Sing, W.Y. Yeong, J. Wei, Investigation of quasi-static indentation response of inkjet printed sandwich structures under various indenter geometries, *Materials (Basel)*. 10 (2017). doi:10.3390/ma10030290.

- [43] Stratasys Ltd., Objet350 and Objet500 Connex3 Spec Sheet, (2017). http://usglobalimages.stratasys.com/Main/Files/Machine_Spec_Sheets/PSS_PJ_Connex3.pdf?v=635836085717699464.
- [44] Stratasys Ltd., Polyjet Materials data sheet, (2015) 3. global72.stratasys.com/~media/Main/Files/Material_Spec_Sheets/MSS_PJ_PJMaterialsDataSheet.ashx.
- [45] P.S. Inc., PIEZOELECTRIC & MATERIAL PROPERTIES OF PSI-5A4E SINGLE SHEETS, (2011) 26. <http://www.piezo.com/catalog8.pdf> files/Cat8.26.pdf.
- [46] L.D. Sturm, C.B. Williams, J.A. Camelio, J. White, R. Parker, Cyber-physical vulnerabilities in additive manufacturing systems: A case study attack on the.STL file with human subjects, *J. Manuf. Syst.* 44 (2017) 154–164. doi:10.1016/j.jmsy.2017.05.007.

5. In-situ Detection of Build Tampering in Metal Additive Manufacturing Using a Cyber-physical Hash

Coauthors: Dr. Christopher Williams

Abstract:

The proliferation of high-value, end-use parts fabricated using additive manufacturing continues to increase. Concurrently, cyber-attacks have begun to more actively target cyber-physical systems such as digital manufacturing facilities. In order to ensure the functionality and quality of these critical components against malicious part sabotage, new techniques are needed for detecting the physical effects of sabotage attacks. In this paper the authors present an approach for using an air-gapped side-channel monitoring system (SCMS) to protect a metal powder bed fusion process against part sabotage attacks. While this overall approach is generalizable to many different types of side-channels, the goal of this paper is to demonstrate an embodiment of the approach using melt pool data. Quality specifications (data package), stored in the part toolpath, are detected by the SCMS and compared against the in situ monitoring to validate the part. These specifications need to be converted into a compact representation because the use of physical emissions limits the data transfer rate. The SCMS has two key requirements i) the ability to detect significant changes to the toolpath and process parameters and ii) to be able to extract the stored data package during the build. This is accomplished by using a frequency representation of the fill pattern to condense the data, hashing the result into predefined ranges, and storing the data package in the part file in a series of QR codes. The SCMS then monitors the laser position (by tracking the galvos) and the emissions from the build area (using photodiodes) to validate the parts and extract the data package from the QR codes.

Keywords: Cyber-physical security, additive manufacturing, laser powder bed fusion, cyber-physical hash, defect detection

5.1. Cyber-physical security in additive manufacturing and the use of side-channel monitoring

Additive manufacturing (AM) has continued to expand and mature as a growing area of manufacturing. The AM metal powders market alone was over \$360 Million in 2019 and is predicted to grow to over \$970 Million by 2026 [1]. This growth has been driven by high-value, end-use parts such as engine fuel nozzles [2], medical implants [3], aircraft parts [4,5], and spacecraft components by companies and organizations such as SpaceX [6] and NASA [7]. These high-value applications require exceptional performance and reliability, resulting in an increased focus on in-situ process monitoring, nondestructive testing, quality control measures, traceability, and provenance in order to ensure performance. This has led to a large quantity of data being generated and digitally transferred throughout the AM process chain, starting with the design and simulation of parts and progressing through the generation of toolpaths, determining optimal process parameters, and collecting and analyzing process monitoring data. This large digital thread is essential to achieve the gains in efficiency and performance of modern manufacturing, labeled “Industry 4.0”, but creates an opportunity for bad actors to target these systems using cyber-physical attacks to cause the failure of physical parts in critical applications.

The most prevalent example of a cyber-physical attack on manufacturing is the use of the STUXNET worm that attacked and disabled the PLCs of centrifuges being used in the Iranian nuclear program [8,9]. Another more recent example is that of a German foundry in which its blast furnace’s digital controller was overloaded due to a cyber-attack [10]. In both cases, significant physical damage was caused as the result of a cyber-attack. In addition to these real world examples, previous research has shown that advanced manufacturing has significant vulnerabilities, both in the manufacturing systems [11–14] as well as the quality systems that monitor them [15,16].

When compared to traditional manufacturing approaches, AM presents additional difficulties for detecting and mitigating attacks. First, the complex geometries AM is often used for can be difficult to measure or have areas that are inaccessible for measurement by traditional quality approach. Second, the layered addition of material means that voids or localized changes to material properties can be completely enclosed inside the part where they cannot be accessed by conventional measurement tools. Within this complexity, researchers have shown a variety of potential vulnerabilities that exist within AM systems [17–20]. Attack vectors include altering the .STL file (to alter the geometry or create internal voids) [18,21–23], changes to the toolpath [17,19], modification of process parameters [19,20,24], and theft of intellectual property (IP) [25–30]. In the context of potential cyber threats, it is necessary to consider approaches for implementing cyber-physical security into AM systems and the quality monitoring associated with them.

One method for increasing the security of an AM system is through in situ process monitoring. However, since monitoring systems can also be targeted by attacks, [15,16] to be secure, these monitoring systems need to be digitally separated from the AM system and preferably air-gapped altogether. For this reason, a SCMS, which takes measurements of either system behavior (e.g. vibrations, temperatures, acoustic emissions, etc.) or direct measurements of system performance (e.g.

tool position), is preferred. These measurements are separate from the system's embedded controllers and provide a second layer of validation. If the SCMS detects deviations from expected values, it may indicate an attack. A second advantage of being a standalone system is that the hardware and software of the SCMS can be more easily upgraded and updated without causing interruptions to the production AM system.

While air-gapping the SCMS is beneficial from a security standpoint, it poses a challenge from a communication standpoint, since without a digital connection between the two systems, the SCMS does not have prior knowledge of the parts being fabricated or of the build specific parameters. While some of this information can be preloaded or manually transferred to the SCMS, manual entry of data is both tedious and error prone which makes this approach undesirable for production use. Transfer of data on physical media is also possible, but the use of optical disks or removable drives present a potential vector for compromising the SCMS and should be kept to a minimum as much as possible. One method for avoiding these challenges is to transfer information to the SCMS physically through the emissions from the AM system during fabrication.

In the authors' previous work the use of a cyber-physical hash as a method for communicating build information through physical emissions to a SCMS was demonstrated [31]. The basis of the approach, illustrated in Figure 1, is to record side-channel information of printing a known good set of parts, condense this quality data into a compact representation (referred herein as a data package) that is then stored in either the model file or the toolpath, and to receive this information using the SCMS during fabrication. This information is then compared against the side-channel emissions of the current printing process to validate the build. Since the data package stored in the toolpath/model file is transmitted using physical emissions that can be read by the SCMS, the SCMS is able to remain air-gapped. The previous work demonstrated this concept using a desktop-scale fused filament fabrication (FFF) AM system, in which temperature and layer-timing information were converted into a hash string that was then encoded into a QR code that was included in the original part's model file. The QR code was printed by the system using two different colored materials and read using an embedded camera on the SCMS. Monitored data from the build was then converted into a hash string and compared against the read QR hashed string, in order to validate the build. While this was a useful test case, many high value industrial AM systems with functionally critical AM parts are fabricated on direct metal AM systems, where the monitoring system requirements are significantly different from material extrusion (ME).

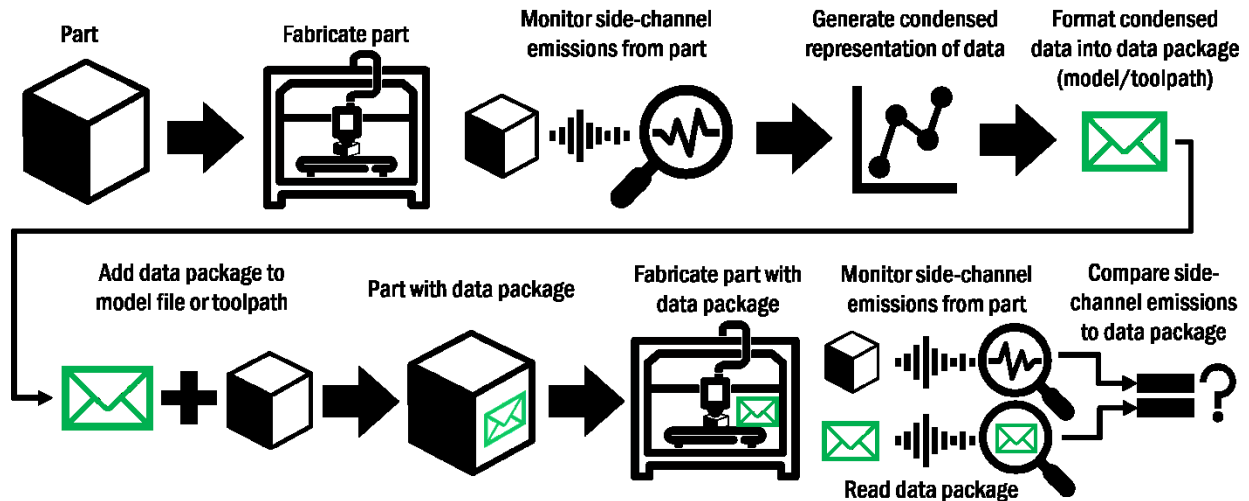


Figure 5.1. Cyber-physical hash overview: Using an air-gapped SCMS to validate AM parts by incorporating a data package in the toolpath or model file to send information using physical emissions during the fabrication process.

The increased speed and resolution of LPBF compared to a ME system create new challenges and opportunities for implementing an air-gapped SCMS. LPBF requires significantly higher resolution and respond time from the sensors. This increases the amount of monitoring data exponentially, a single layer in a LPBF system can contain millions of individual data points, far more than can be directly sent to the monitoring system. A benefit of this increased resolution and speed is that it has the potential to allow for significantly a larger data package to be stored and transmitted to the SCMS. In this paper, the authors seek to expand the ideas presented in the previous cyber-physical hash work [31] by exploring demonstrating a new embodiment of the concept that is able to address the significant differences between the two system types.

At its core, any embodiment of the cyber-physical hash approach will require four core requirements to be met:

- 1) Ability to monitor the geometry and parameters of parts during fabrication
- 2) Ability to store information in the toolpath or model file in a way that can be physically communicated to a SCMS
- 3) Ability to condense quality information into a data package that fits in toolpath/part file
- 4) Ability to extract information package back from the physical emissions

In this paper, the authors will address examine how these requirements can be satisfied on a metal LPBF system by introducing new methods for comparing toolpaths and process parameters between parts, techniques for reducing the quantity of raw data into something that can be stored in the data package, an improved toolpath representation of the data package, an hashing approach that can accommodate process variation, and techniques for extracting the data packing using the available side-channels.

5.2. Side-channel monitoring system and transmission approach

The choice of side-channels greatly affects the affects the embodiment of the cyber-physical hash approach. Sensors selection is paired with the data density that the monitoring system and data package need to accommodate. The authors' previous work on the IDEAS framework provides a detailed approach for the selection of side-channels. In this paper, the work will focus on the use of a galvo side-channel to track laser position and photodiodes averaging the spectral emissions from the build area to approximate meltpool emissions. Raw meltpool data cannot be stored in the toolpath because it is infeasible to use physical emissions to transmit gigabytes of data, both in time and toolpath complexity. The data package transmitted needs to contain a condensed representation of the quality data that can be used to validate the build, while still being small enough to be physically stored and transmitted.

The cyber-physical hash concept presented in the authors' previous work [32] allows an arbitrary amount of data to be condensed into a finite length string that is stored in the toolpath as a data package. If the part is altered, either through changes to the part itself or through attacking the AM system, without making the corresponding changes to the data package, the SCMS will reject the build/part as being invalid once the discrepancy has been detected. If the data package is altered without changing the part, the SCMS will similarly reject the build/part. Only if the part and the data package have matching modifications will the build/part pass the SCMS as shown in Figure 2. By securing the information stored in the data package (e.g., using encryption or hashing) the manufacturer can make it significantly more difficult for an attacker to compromise the system. While part/toolpath data can also be encrypted for additional security, the toolpath information must eventually be decrypted by the machine in order for the machine to physically fabricate the part. If the machine becomes compromised, the attacker may be able to alter the toolpath after it has been decrypted. With the data package, the key difference is that the information remains encrypted even during fabrication and is not decrypted until it reaches the air-gapped SCMS. While the physical toolpath of the data could be seen by an attacker, the information contained is never decrypted on the printer in a way that could allow an attacker to access its contents.

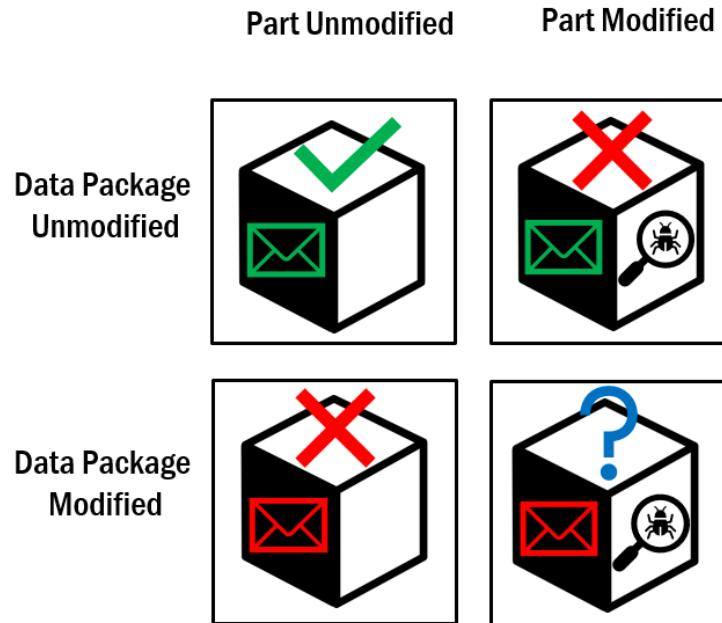


Figure 5.2. Examples of different cases for the SCMS where the part/data package are modified/unmodified. The manufacturer needs to ensure that an attacker cannot alter the data package to match a modified part or the system could reject or pass the part.

5.2.1 Toolpath representation of data

After choosing the SCMS, the next step is to determine how information can be stored in the toolpath or model file, and how that information will be extracted by the SCMS. Since the SCMS that was chosen used both a camera and a photodiode setup, it was decided to use a QR code as the method for transmitting the information. The reason for choosing the QR code is that it could be included in the model file directly without requiring direct access to modify the toolpath files and had a higher data density than a barcode or similar model file-based approach. The QR code could also be interpreted by either the camera or the photodiode monitoring system. To demonstrate this proof of concept the data package (hash string of frequency representation control points and the average photodiode intensity), was stored in a series of QR codes stacked in the Z-axis of the build. During fabrication, this QR code data would be extracted either using an embedded high-resolution camera or laser scan lines from the galvos movement and emissions from the build area.

The main drawback of using a fabricated QR code to transmit information to the SCMS is that it increases the build time, energy, and material consumption. An ideal transmission method would not require any additional material and would add as little time as possible to the build. Given the high resolution of the laser galvo tracking on the SCMS and direct access to the toolpath it would be possible to implement more sophisticated methods of transmitting data such as those demonstrated for data exfiltration such as (Fansmitter [33], Brightness [34], ATTACH [X], etc.) that could transfer data at a higher rate and adding less time/material to the build. Unfortunately, without direct access to the toolpath, it is not possible to implement these types of approaches.

To circumvent these challenges, the process instead scans the QR code with the laser power set to 0 W and the scan speed increased to 1 m/s. While no physical QR code is fabricated, the SCMS still tracks the movement of the galvos and can use this toolpath information to reconstruct the geometry of the QR code and extract the information stored inside. This “ghost” QR code enables data transmission without consuming material and also reduces the transmission time due to the higher scan speeds that can be used when not melting/solidifying material. Since the operator can adjust the power and scan speed settings when setting up a build, this method can be implemented without requiring direct access to the toolpath.

5.2.2. Frequency representation and comparison of toolpath data

Due to the high spatial and temporal resolution required to monitor the melt pool, the amount of data generated by the SCMS is quite large (millions of data points per layer). This large quantity of data makes it unfeasible to directly compare each discretely measured data point so it is necessary to design a method for reducing the amount of data. To accomplish this, a frequency representation of the fill lines in the toolpath was used. This approach first determines the change in galvo position from the previous point and uses this to approximate the speed of the laser. After this, the system uses a known bounding area for each part and then performs the following steps to convert it to a frequency representation and identify any deviations from the expected values:

- 1) Determine the X-Y bounding box for each part on the build
- 2) Use change in position between measured points to approximate speed
- 3) Filter the data using speed/intensity to extract the fill lines (max speed, min speed, intensity threshold)
- 4) Apply a rolling average filter (window size $k=2$) to smooth small jitters/noise
- 5) Determine locations where different fill lines are (change in velocity)
- 6) Calculate the number of points within each fill line (duration) to get a frequency representation of the part
- 7) Apply a rolling average filter ($k=2$) to smooth small jitters/noise
- 8) (Optional) Remove any very steep lines ($\Delta y > 250$) on the signature
- 9) Find difference (e.g. root mean squared (RMS)) between the frequency representation of the toolpath of the control signature and the frequency representation of each part
- 10) Plot a bar graph for each layer showing RMS difference of each part between layers

The advantage of this approach is that it can significantly reduce the volume of data in a build. Instead of storing and transmitting hundreds of thousands to millions of data points related to galvo position within each layer, the amount of data can be reduced by several orders of magnitude. The resulting signature also provides a more human readable way to visualize differences in the part toolpath, as illustrated in Figure 3, thus making it easier to detect and locate attacks. When using the cyber-physical hashing method, steps 9 and 10 are omitted and replaced by the mapping of control points to predefined ranges as discussed in Section 2.6.

A primary feature of this signature approach is that it is highly sensitive to the toolpath itself, not just the part area being scanned. If a small void is inserted into a part, it will split scanlines and result in a significant change to the signature as illustrated in Figure 3c. The same geometry being scanned in

two different patterns will result in different signatures (Figure 3a/b). Since part properties and quality can be affected not just by geometry, but also by toolpath driven physics, such as cooling rate, these can have a significant effect on final part quality. This is beneficial in detecting sabotage attacks that alter the toolpath because even if the geometry of two parts is the same, small changes to the toolpath pattern will be detectable.

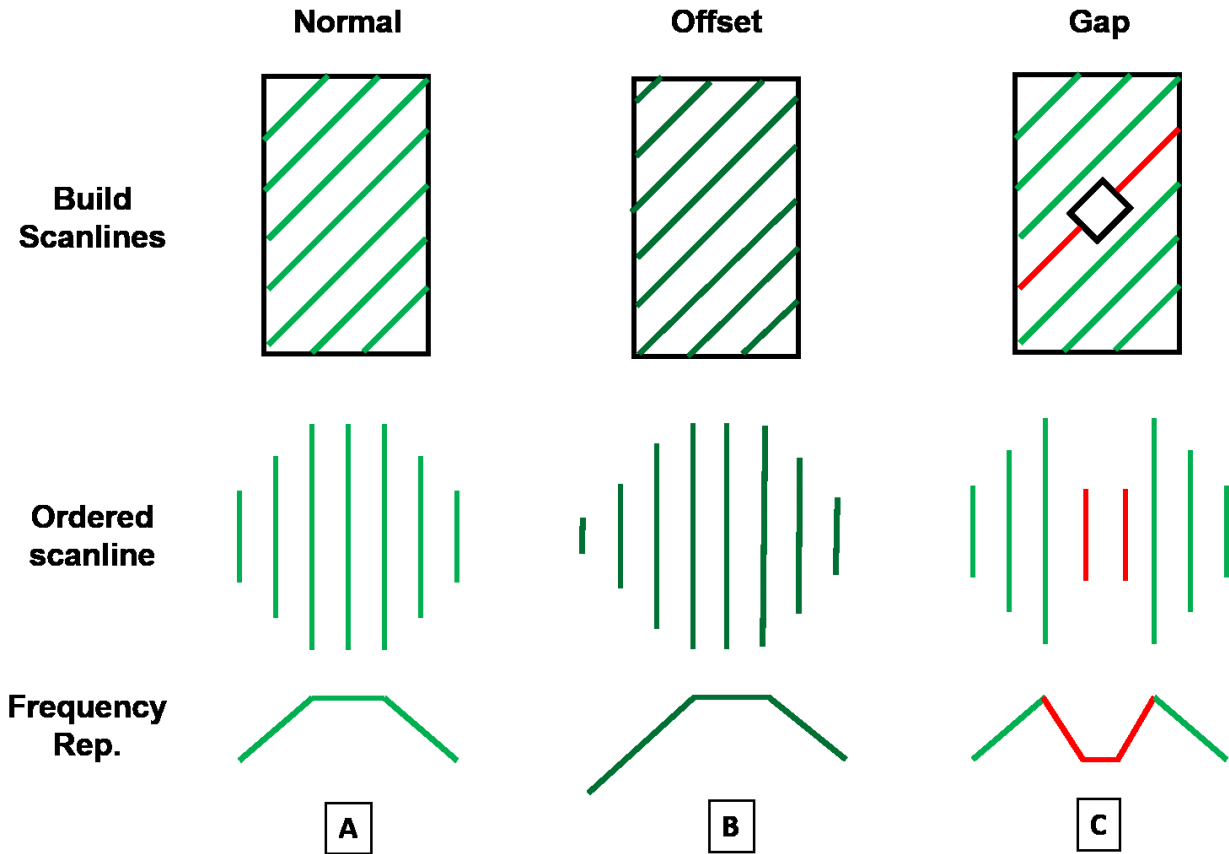


Figure 5.3. Example of how infill frequency representations change with different toolpaths The as-built scanlines, sequentially ordered scanlines, and the corresponding frequency representations of the scan lines are shown for three different toolpaths (normal, offset, and gap).

A drawback of the frequency representation is that small changes in the toolpath that may not affect the functionality or properties of a part may result in significantly different signatures that may falsely indicate an error or attack. An example of this is the addition of a few small scan lines at the beginning or end of a rastered set, slightly offsetting the series of scan lines as shown in Figure 2b. This may not have a notable effect on the final part, but would shift the signature curve, misaligning it and resulting in a large sum error over the total curve. Depending on the application and system set up this could be either an advantage or a drawback to this approach. A way of further reducing the amount of data required to represent the meltpool data is through the use of control points. Control points are the locations in the frequency representation where the slope changes by more than a threshold value (in this study a value of 3). By only storing the endpoints of line segments the amount of data can often be significantly reduced and intermediated values can be interpolated from the control points.

5.2.3. Build validation using cyber-physical hashing

With a side-channel selected and a toolpath representation for transmitting the data defined, it is necessary to determine the method that the SCMS will use to compare the in situ measured data against the data package embedded in the toolpath/QR code. Some approaches will require more information to be communicated to the SCMS (more data stored in the toolpath) while others will require more data to be preloaded to the SCMS. In section 3 a direct root mean squared (RMS) comparison of the frequency representations throughout the build is shown to better illustrate how changes to the toolpath and process parameters affect the side-channel emissions from the build.

5.2.3.1. Cyber-physical hash and predefined ranges

To implement the cyber-physical hash on a LPBF system a series of predefined ranges for photodiode intensity and frequency representation control point position are determined and preloaded onto the SCMS. Next the control part(s) baseline photodiode intensity and control points of the frequency representation is mapped into these ranges using a selected tolerance value. The ranges from the baseline are then combined with a key and cryptographically hashed. The resulting hash string is then stored in the toolpath data package. During fabrication the SCMS takes measurements via the embedded air-gapped sensors. It then maps the monitored side-channels into the predefined ranges and generates a hashed string from the resulting ranges. The hash string is compared against the hash string stored in the data package. If the hashes do not match, the system has been attacked and the operator is alerted. The ranges are a necessary feature for this approach because without them any small change in monitored values (unavoidable in physical manufacturing) would result in hash values that would not match, resulting in falsely rejected parts.

The advantage of this approach is that the hash string is a fixed length, which can greatly decrease the amount of data, and it is irreversible, which prevents an attacker from reversing the quality information from the hash. While this approach does require some information to be digitally transferred to the SCMS (i.e., the ranges used and keys to hash them with), this information only needs to be set up initially and does not require regular updates for new parts like the identifier approach. More details on the use of cyber-physical hashing using predefined ranges can be found in the authors' previous work [31].

5.2.3.2. Implementing predefined ranges over a continue domain

While predefined ranges work well for parameters that have discrete set points with minimal variation around those set points (such as set laser power) they work poorly for parameters that consistently vary across a continue range (such as tool position). If the measured point is in the center of a bin the method will not present an issue. However, any continuous line from one bin to another will have to cross the boundary between bins. In the case of a measurement on the edge of two bins, small

process variations will cause the measurement to bounce between the two bins. If this bin is used in a hash function it means that the system will generate false detection positives with small perturbations. To be able to use predefined ranges with the frequency representation of the toolpath data a technique is needed to compensate for these edge cases. Increasing the size of the bins can reduce the number of edges cases, but will not eliminate the problem completely.

One approach to solving this problem is to calculate the hash not only for the bin the data point falls into, but also for each adjacent bin. Using a modern system and GPU acceleration it is possible to calculate more than 5 billion simple MD5 hashes every second[35], with computational capability providing an upper limit to the number of bin combinations that can reasonable be calculated by the SCMS. For the frequency representation each bin has 8 neighboring bins. A naïve approach would require testing 9 bins for each hashed point. Because each combination of bins must be tested this approach scales at a rate of 9^n where n is the number of points that can be tested. When $n = 10$, this approach already requires close to 5 billion hashes to be calculated.

This first step that can be taken to improve this approach is to reduced the number of adjacent bins. Unless there is a large amount of variation in the process, data points will only vary between ranges on one side. This reduces the number of adjacent bins for each point to either 4 (if near a corner) or 4 (if near an edge). This results in a value of $n = 16$ (for corners) or $n = 32$ (for edges), which is a significant improvement over the naïve approach. Another step that can be taken to reduce the number of points combinations that need to be calculated is to ignore points that fall in the center of a range and to only test those that fall near an edge or corner (and are likely to be shifted by small process variation). By calculating adjacent bins for values that appear closes to edges/corners first, the approach can increase the likelihood of detecting the correction combination early, requiring few combinations to be calculated. If the number of adjacent bins that needs to be calculated exceeds a reasonable amount the system determines that the part is outside of acceptable parameters and sends an alert. If a matching hash is achieved, the number of points that needed to be mapped to adjacent bins can be used as a measure of the quality of the match. Large numbers of adjacent bins may indicate that an attack has occurred. If the system is consistently generating too many false alerts, it may be necessary to increase the size of the predefined ranges, decreasing the amount of measured points that are mapped to the incorrect range due to process noise. A summary of the steps in the approach is as follows:

- 1) Place measured points into predefined ranges
- 2) Calculate hash for resulting ranges
- 3) If the hash values do not match, select the point that is closest to the edge of a predefined range and generated a new hash using the nearest adjacent range
- 4) Repeat this process for each point in order of the smallest distance from the edge of range to the point and calculate every combination of ranges for these points until either a matching hash is found, all adjacent ranges have been checked, or a maximum value of computations is reached
- 5) If a match is found, report the number of control points that had to be mapped into adjacent ranges and the distance of these points from the adjacent range. This serves as a

measure of the variation in the signature. A large number of adjacent ranges may indicate that an attack has occurred and should raise a warning for further inspection.

- 6) If a match is not found, reject the part and raise an alert.

5.2.4. Distribution of intensity data

While the frequency representation of the scan pattern is able to detect changes to part geometry, toolpath, and laser speed, it is unable to detect changes to the laser power which may affect the material properties of the part. To achieve this, it is necessary to measure the intensity values from the photodiode. While the photodiode intensity values do not directly measure the laser power, they provide a correlated response that can be used for validation purposes. By mapping the intensity values to the same positional data points as in the frequency representation, the laser power can be validated. A straightforward approach for doing this is to determine the mean and variance of the intensity measured by the photodiode for each part on each layer. During fabrication, the SCMS can calculate the mean/variance of the intensity and map it into a range to compare against the control values. This allows for the detection of large-scale changes, such as altering the laser power for an entire layer or part, and may also provide some ability to detect smaller changes (if they increase the overall variance significantly). If more precise/localized detection of intensity changes is needed, additional rules/checks can be implemented such as a localized rolling average mean/variance or a control chart-based detection of out of range values (e.g., if 10 of 12 consecutive values fall outside of a defined limit). In contrast to toolpath, intensity will have a significant amount of natural variation from part to part and build to build. While creating tighter rule sets will allow the detection of smaller changes by an attacker, the natural variation means that tighter rules will also increase the chances of having a false detection event, where the difference is due not to an attack, but to natural process variation.

5.2.5. QR code extraction

The most straightforward approach to reading a QR code is to use a camera to extract the regions of interest. This works well for high contrast systems such as a screen, black-and-white printing, or distinctly colored multi-material printing, however does not work well in systems with poor contrast. In the case of LPBF both the melted powder and the unmelted powder are similarly colored, making it difficult to correctly identify black/white areas in the QR code. Due to the small nature of the cells and the nature of the process, the physical “pixels” of the QR code may be irregularly formed, further increasing the difficulty of accurately extracting the data.

Due to the potential limitations of camera-based reading of a QR code fabricated on a LPBF process, a method was designed for extracting the QR code information directly from the photodiode/galvo data. The nature of the LPBF process means that the toolpath is not a direct representation of the fabricated geometry (). In metal LPBF the part geometry is formed by using multiple overlapping scan lines to transfer sufficient heat into the powder to melt it in the desired area. In addition to the powered scan lines there are also travel movements where the laser is turned off and moving from the end of one scan line to the beginning of another.

To extract the information stored in the QR code it is necessary to recreate the intended geometry from the toolpath data. There are a variety of approaches that can be used to accomplish this task, two of which are detailed below.

5.2.5.1. Extracting fabricated QR code from scan lines using summed intensity

The first approach utilizes the intensity measurements from the photodiodes to generate an overlapping intensity map. These areas are then attributed to either a black or white cell depending on the amount of energy detected in the area. The core approach is as follows:

- 1) Filter out values below minimum intensity threshold (eliminate travel movements)
- 2) Grow each point using a gaussian intensity function
- 3) Create a grid to map the intensity values
- 4) Sum the resulting intensity distributions over the grid
- 5) Threshold the grid values to determine white/black

Because the travel movements are done while the laser is off, they can be filtered out by removing any scan data under a minimum intensity threshold. This removes data points that are not directly contributing to the fabricated QR code. Performing this thresholding results in a remaining set of scanlines that much more closely resembles a QR code. By taking these resulting scanlines and applying a gaussian distribution to each point converts the point data obtained by the monitoring system, to a series of small overlapping areas. A grid is generated and the intensity in each cell of the grid is summed together to give a value for that cell. Each cell is then assigned either a white or black coloration based on a threshold value, resulting in a final QR code.

5.2.5.2. Extracting ghost QR code from scan lines using scan line averaging

A key requirement of the previous approach is the use of intensity thresholding to eliminate jump lines and to determine black/white areas. With the laser power set to zero, it is no longer possible to use the intensity in this way since all scan lines have the same (near zero) intensity value. For this reason, a second approach for extracting the QR code information is needed. This approach is as follows:

- 1) Filter out lines points faster than a speed threshold
- 2) Extract the horizontal and vertical scan lines from the QR code
- 3) Remove lines below a minimum length threshold (eliminates noise)
- 4) Calculate the center point of each scan line
- 5) Average the center points of each scan line together to determine the QR code cell

Jump lines typically move at the maximum speed possible for the galvos. By setting the scan speed of the ghost QR code slightly slower than this maximum, speed thresholding can be used to eliminate jump lines. Once the jump lines have been eliminated, the next step is to extract only the horizontal and vertical (relative to the QR code) scan lines from the QR code. These lines contain the border information for each cell in the QR code and excludes additional scan lines that may overlap multiple cells and may make identifying cells more difficult. Any horizontal and vertical scan lines below a minimum length threshold are removed as these may be from the edges of other scan lines or other

types of noise. The center point of the remaining scan lines is then calculated and the average of each set of scan lines is used to determine the center of each QR code cell. Effectively this using the scanned borders of each cell to determine that cell's position. While scan lines of nearby cells may overlap, keeping track of the time order of each scan line allows for appropriately grouping cell borders.

5.3. Experimental

The goal of the experimental work was to validate the proposed SCMS technique for verifying LPBF build quality. This was accomplished with the following steps i) collect a control dataset from the SCMS, ii) use the frequency representation approach to reduce the volume of the collected data, iii) hash the quality information into predefined ranges, iv) store the information in a data package in the model file that was loaded onto a build with other parts, v) use the SCMS to read the stored data, vi) detect modifications to parts by comparing stored data and in situ measurements. Specifically, the experimental work investigated the following red team attacks:

- 1) Internal geometry changes (voids)
- 2) External geometry changes (altered curvature)
- 3) Altered toolpath with identical geometry
- 4) Increased/decreased laser speed
- 5) Increased/decreased laser power

In addition, the red team attacks, two data package representations were used

- 1) Physically printed QR code stack
- 2) Ghost" QR code stack, "fabricated" with increased scan speed and laser power $P = 0 \text{ W}$.

To accomplish these goals, two builds were fabricated, detailed in Section 3.2. This first build was used to test the ability of the proposed approach to detect red team attacks on part quality. The second build was used to validate the ability to process the toolpath data to recover the QR code and read the hash value stored within.

5.3.1. Laser powder bed fusion and side-channel monitoring system setup

For the experimental work, a 3D systems ProX DMP 320 LPBF system was used that featured an array of embedded side-channel sensors including a pair of photodiodes and a high-resolution camera. The details of the system are the same as those described by Coeck et. al., where a pair of photodiodes capture all light emitted from the build chamber at a rate of 50kHz and then synchronize this information with the laser position from the galvo controller [36] as shown in Figure 4. While this approach is unable to isolate the emissions from the meltpool, it does capture the total conditions from the build chamber and should be sufficient for detecting a malicious attack on the process, while being easier and less expensive to implement than a co-axial scanning system.

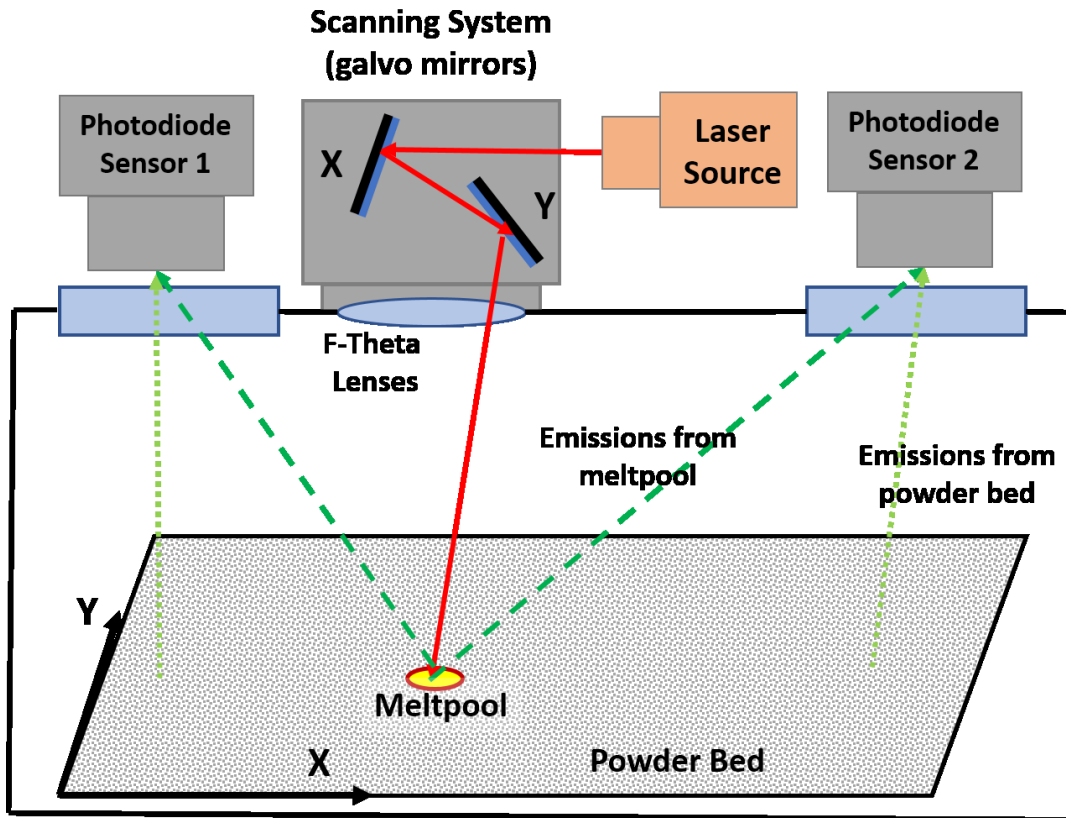


Figure 5.4. Example of monitoring system for metal powder bed fusion. Emissions from the meltpool are captured by a pair of photodiodes and synchronized with the position of the laser read off of the galvos in the scanner.

5.3.2. Build setup

Build 1 consisted of tensile test specimens (ASTM E8/E8M -16a, Plate Type Subsize Specimen) (Figure 5) and small rectangular tabs (25 mm x 10 mm x 4 mm). The second build contained a second set of test tabs and two series of QR code stacks (Figure 6, one that was physically fabricated and one “ghost” QR code stack that had the laser power turned off and the scan speed increased). These builds contained both control parts and parts with either modified geometries (Figure 5) or modified process parameters (speed or laser power). The specific contents of each build are detailed as follows:

Build 1 (11 Test Parts, Figure 5):

- 2 Tensile Test Specimens - Control (Figure 3a)
- 1 Tensile Test Specimen - External Geometry Defect (smaller neck curvature) (Figure 3b)
- 3 Tab - Control
- 1 Tab - Internal geometry defect (different sized cylindrical voids) (Figure X)
- 1 Tab - Laser Power Increase (+150W, +50%)
- 1 Tab - Laser Power Decrease (-6W, -2%)
- 1 Tab - Scan Speed Increase (450mm/s, +50%)
- 1 Tab - Scan Speed Decrease (-18mm/s, -2%)

Build 2 (4 Test Parts):

- 3 Control Tabs
- 3 QR Code Stacks – Printed (Figure 4)
- 1 QR Code Stack – Laser turned off

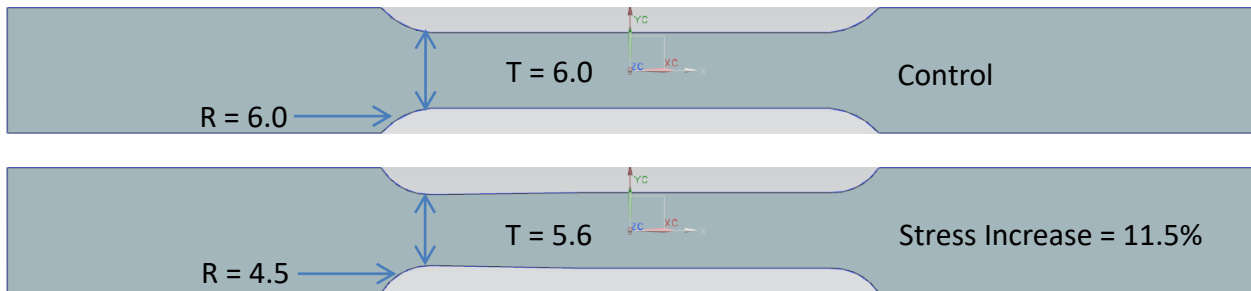


Figure 5.5. External geometry change of tensile test specimen compared to control specimen. The curvature at the neck of one side of the specimen has been reduced from a radius of 6 to a radius of 4.5.




Figure 5.6. QR code stack for transmitting the data package to the SCMS through physical emissions from the build. QR codes pictured are spaced at intervals of each five layers with a total of six QR codes per stack.

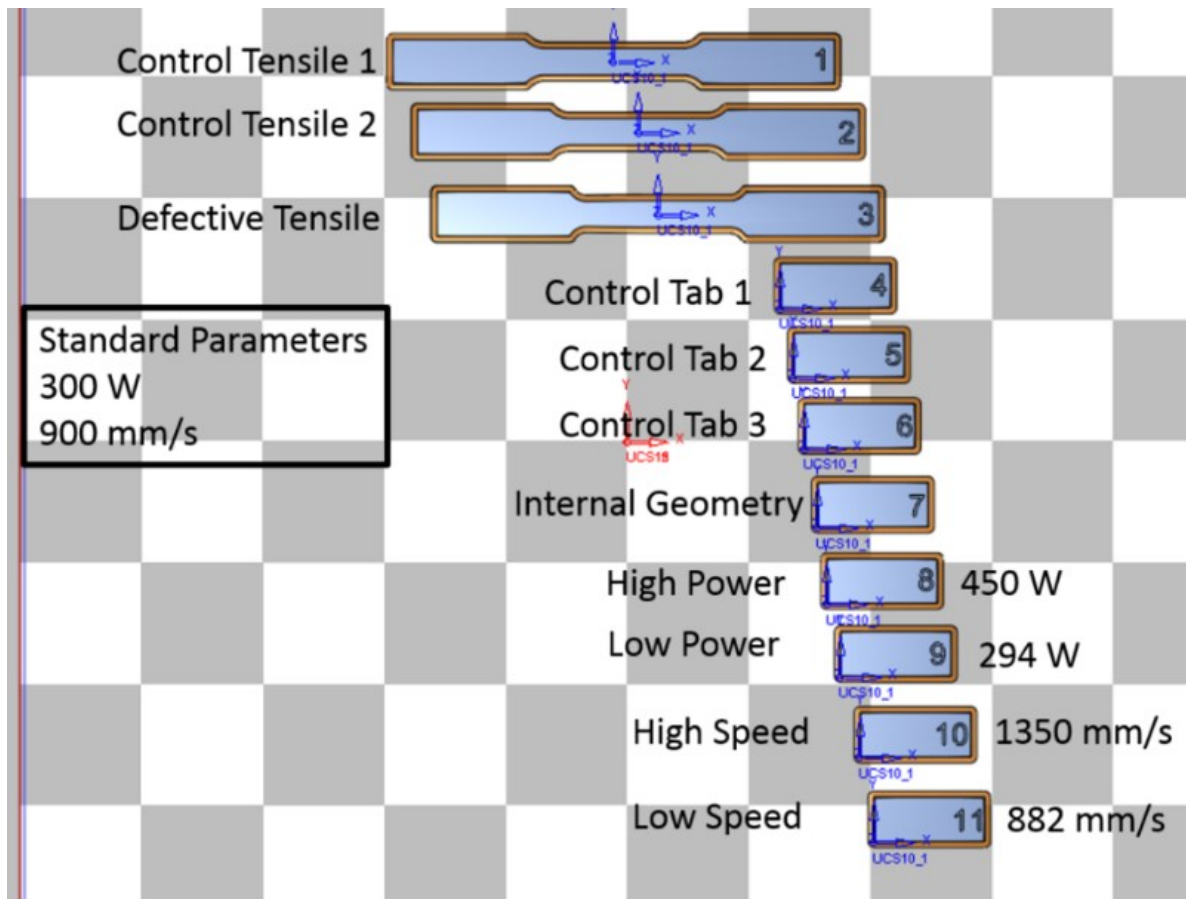


Figure 5.7. First test build layout consisting of 11 parts along with the process parameters used for each part.

5.4. Results and Discussion

After fabricating both sets of parts an analysis was done on the laser galvo positions to generate a frequency representation of the toolpath of each part on every layer. The frequency representations of the altered parts were then compared against the frequency representation of the control part in each build. To validate laser power, the average and variance of the photodiode intensity for each layer of each part was determined and compared against the average of the control samples. Section 4.1 details the results of the toolpath comparison (geometry and laser speed attacks), Section 4.2 details the results of the intensity comparison, and Section 4.3 details the resolving of the QR codes from the SCMS data.

5.4.1. Toolpath Comparison

As discussed in Section 2.2, a frequency representation of the infill scanlines was used to condense the large number of data points into a smaller data set. Figure 8 shows the frequency signature for layer 68 of Build 1 in the time order that they appear during the build. The sections of the overall signature corresponding to each part are highlighted. The tensile specimens contain more scan

vectors and have correspondingly longer signatures representation that is distinct from that of the tabs. The frequency representation of position data for the tensile specimen with the external geometry change can be seen in Figure 9.

The various printed tabs exhibit a variety of signatures depending on the types of modifications made to the base toolpath as shown in Figure 10:

- The control tabs (shown in green) overlap with each other and with the tabs that have altered laser power
- The tabs with increased (“+”) and decreased (“-”) laser power (shown in blue) overlap with each other and with the control tabs, since laser power does not affect the toolpath
- The tab with the 50% speed increase (shown with a dotted red line) appears significantly below the other signatures due to having a reduced number of data points in each scan line over the same time period
- The tab with the 2% speed decrease (shown with a dashed red line) appears slightly above the other signatures due to having an increased number of data points in each scan line over the same time period
- The tab containing the internal void (shown with a solid black line) is initially similar to the control tabs and then develops a jagged, up/down series of jumps due to the splitting of single long scan lines into multiple short scan lines around the locations of the internal voids (as shown in Figure 3). On layers that do not contain these internal defects, the signature is the same as those of the controls




Figure 5.8. Frequency representation of laser toolpath changes throughout build 1, demonstrating unique patterns for different parts.




Figure 5.9. Part frequency signatures isolated and overlaid for tensile test specimens for later 68. The blue dotted lines show the control samples and the red dashed line shows the sample with altered geometry. The location where the curvature is altered can be clearly seen as a slightly lower height in the curve.




Figure 5.10. Part frequency signatures of different tabs overlaid to demonstrate the signature variations caused by process parameter and geometry changes. The signature representation allows for a visual inspection of part variation as well as the ability to identify the type of defect based.

Figure 11 shows the RMS difference of each tab from the control tab for each layer in build 1. In this build the first 33 layers are supporting material layers that override any scan speed or laser power settings. Since all of the tabs have the same geometry, this results in similar signatures for all parts and little total difference between all the signatures. From layer 34 to 66, the process parameter changes take effect, but the internal void is not currently present. The 50% increase in speed is immediately detectable, while the 2% decrease in speed is not significant enough to cause detection in the current approach. The void defect occurs between layers 67 and 78, and results in a detectable change in signature from the control. Due to an error in the data collection, layer 47 is omitted from the data set, which is why there is no variation on this layer in the figure.

Figure 12 shows the signature comparison between the test tabs in the second build. While the geometry and process parameter settings were the same, the second tab has a different scan pattern, as shown in Figure 13. This difference in scan pattern results in a different frequency signature (Figure 14) which causes a significant deviation compared to the other controls. From layer 10 onward, the toolpath of the second tab returns to matching the toolpaths of the other two tabs as can be seen from the lack of difference detected in the remaining layers.




Figure 5.11. Frequency comparison of toolpaths for tab test parts during fabrication. Green/yellow indicate the greatest difference from the control values. The first 33 layer are support material and have identical process parameter and toolpath settings (overriding part specific settings). The layers where the void occurs are highlighted with a red outline.




Figure 5.12. Frequency comparison of control tabs for Build 2. Tab 2 has the same geometry, but a different toolpath on layers 1-9 which was not part of the designed experiment, but was detected by the proposed method.

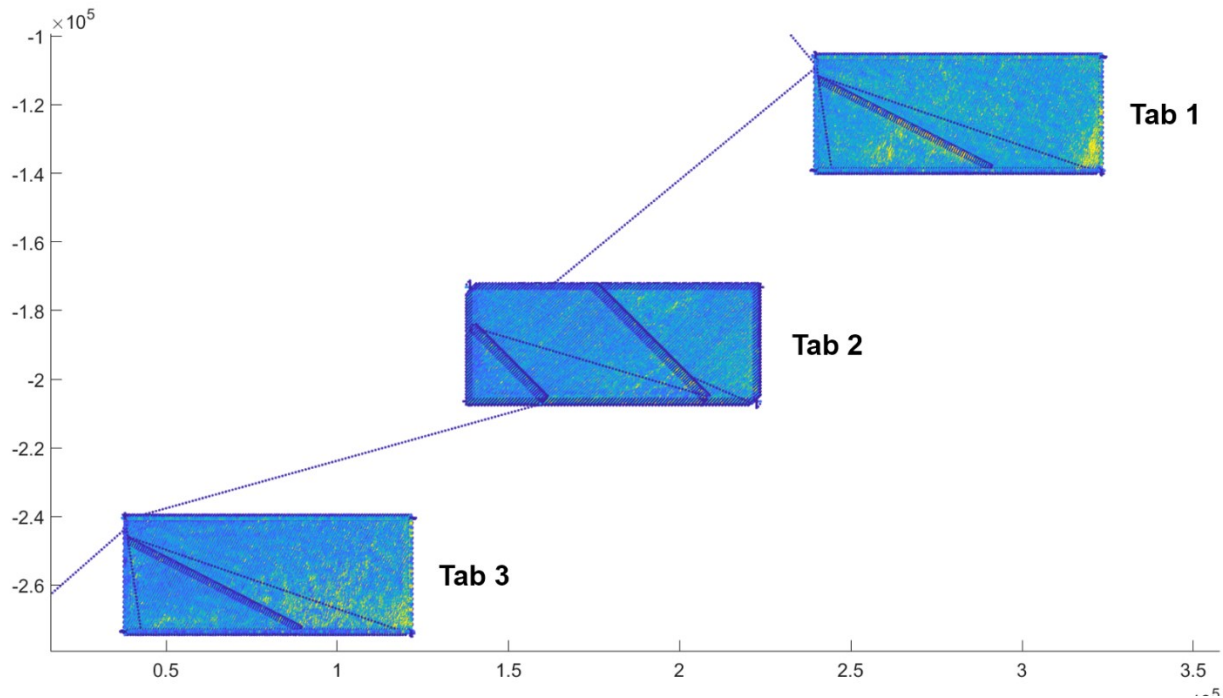


Figure 5.13. Example of different toolpath for Tab 2 (taken from layer 3 of build 2). Dark blue colors indicate low intensity points while brighter green/yellow points indicate high intensity points.

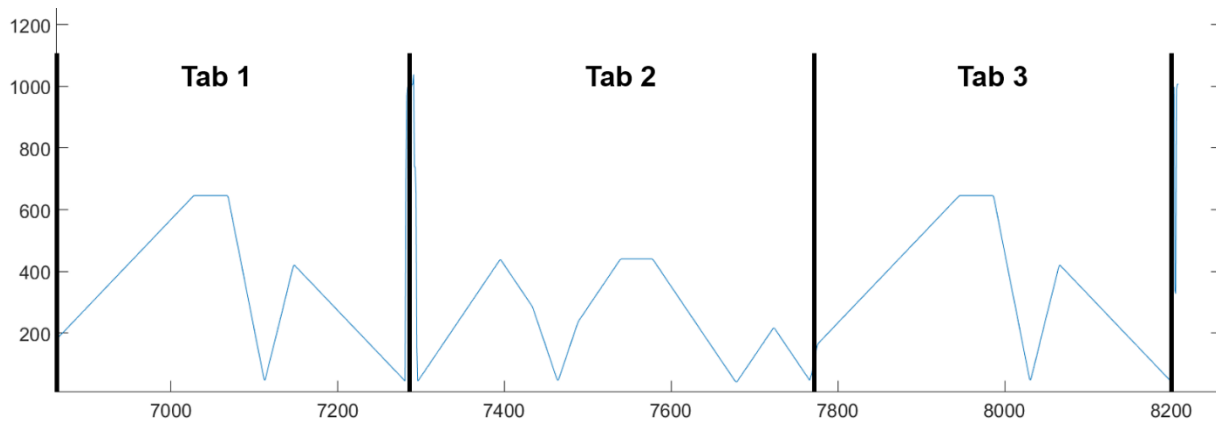


Figure 5.14. Example of different toolpath signature for Tab 2 (taken from layer 3 of Build 2). Tabs 1 and 3 have two scan blocks (corresponding to their two signature peaks). Tab 2 has three scan blocks (corresponding to three signature peaks).

5.4.2. Intensity Comparison

To evaluate the ability of the photodiode to detect changes in process parameter, the average intensity for each part was calculated and compared against the average for the control specimens in the build as shown in Figure 15. In Build 1, a slight, but noticeable trend could be seen during the first 33 support layers (where process parameters and geometry were identical). Test tabs located later in the build had slightly higher average intensity values. It is unclear whether this was due to the position on the build tray affecting the photodiode responsiveness to the meltpool emission or if the increase was due to residual heating of the build area from previous parts in the layer. After the process parameters were changed the tab with increased laser power was able to be discerned from the control signals. Unexpectedly, the tab with the increased laser power showed a greater variation from the average intensity than the tab with increased laser power. The tab with a 2% decrease in laser power had an average photodiode intensity that was indistinguishable from the control tabs.

In Build 2 the control tabs exhibit similar average intensity values during the first 33 layers (Figure 16). However, from layer 34-36 the intensity of the third control tab was noticeable different from the other two controls. While this difference can be seen in the raw meltpool data it is unclear what the physical cause of this increase was. Further study is needed to identify the causes of intensity variation between parts within a build.




Figure 5.15. Intensity comparison of each part in build 1. Green/yellow values indicate larger changes in average intensity. Some slight build location dependency can be seen from layers 1-33, when the settings for the parts are identical. Once the process parameters are changed, the increased power creates a noticeable difference. The increase in scanning speed and changes to the toolpath from the void defect also causes a detectable change in the average intensity.

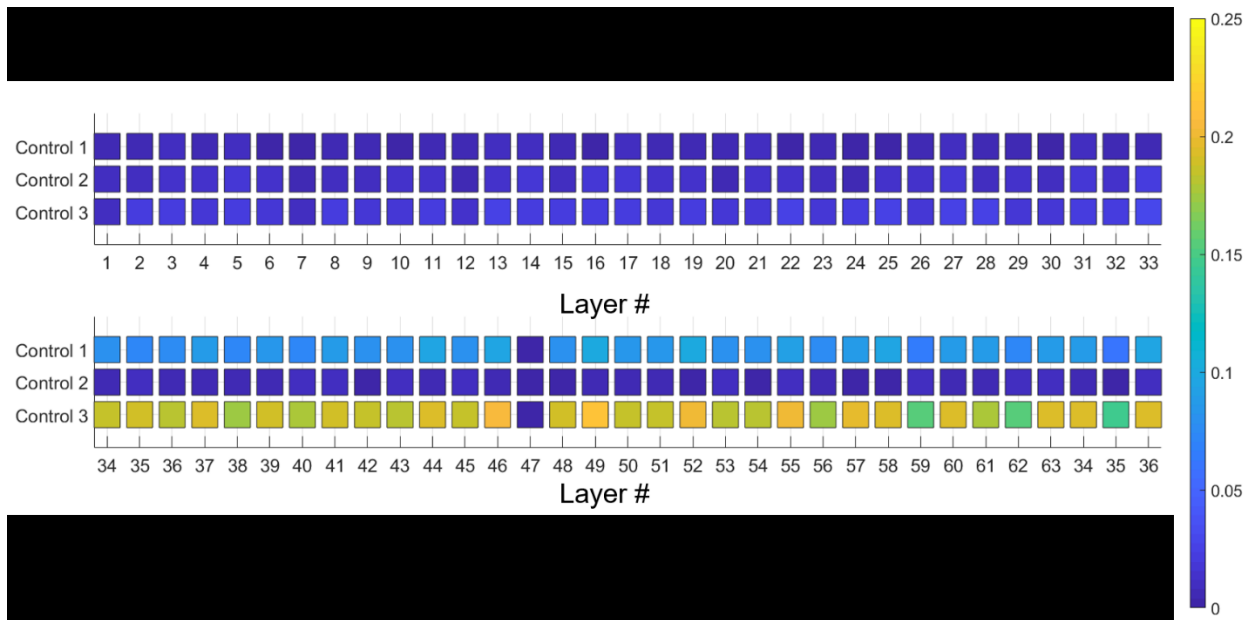


Figure 5.16. Intensity comparison from Build 2. The third test specimen exhibits significantly higher average intensity from between layer 34 and layer 66. It is unclear what changed in the system to cause this difference.

5.4.3. Mapping frequency representation into predefined ranges

Using approach presented in Section 2.32 the frequency representation of the toolpath was mapped into a series of predefined ranges. It was observed that numerous control points fell close to the edges of the predefined ranges. In some cases, the control tabs had points that mapped into ranges that were adjacent to each other as shown in Figure 17. The green bins indicate the target ranges for the hash string, yellow bins are ranges that are one bin adjacent in a single direction and orange bins are ranges that are one bin adjacent in two directions. Red bins indicate ranges that are not directly adjacent to the target range. This approach was able to easily detect the 50% increase in scan speed and the inclusion of internal voids. The 2% speed decrease remained within the adjacent bins, but might be discernable by the increased number of adjacent ranges needed to achieve a match. The results show that this approach can be used to increase the robustness of the physical hashing approach against process noise.

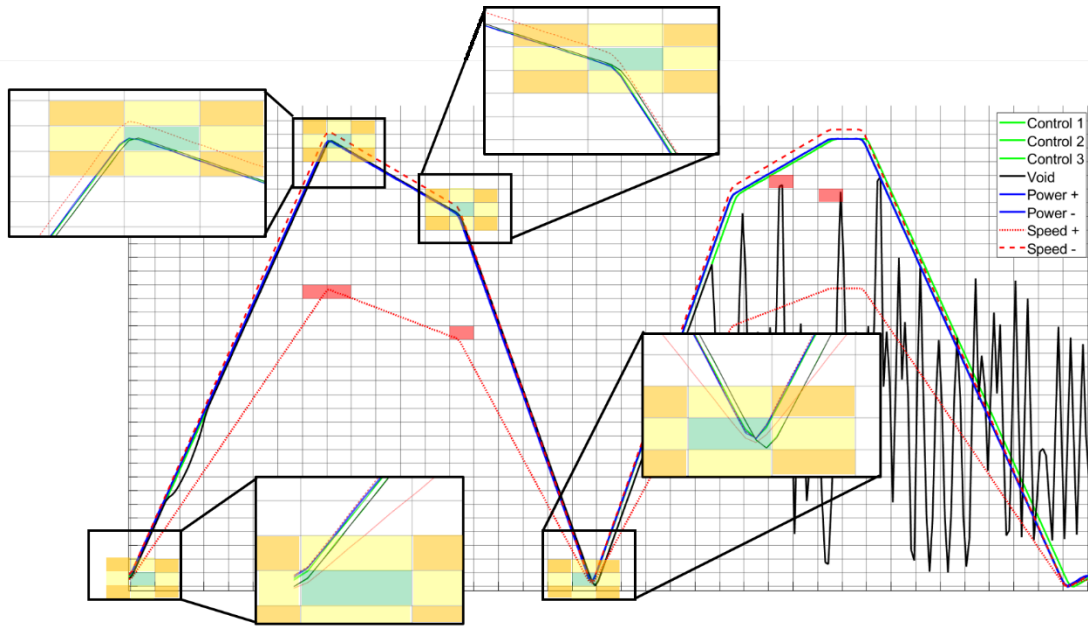


Figure 5.17. An example of mapping the control points of the frequency representation of the toolpath into predefined ranges and how combinations of adjacent bins can be used to increase the robustness of the system against process noise. Expanded views show the location of control points for different tabs and how they map into the predefined ranges. Green in the target range, yellow indicates an edge (one bin adjacent on a single axis) and orange indicates a corner (one bin adjacent on two axis). Red values show control points that are at least two bins away from the target value

5.4.4. QR Code Extraction

Initially, images from the camera were used to attempt to read information from the QR code; however, at the resolution used the fabricated QR code was not reliably able to be extracted from the image. The primary reason for this was that there was not good contrast between the solidified powder and the remaining unsolidified powder around it, as shown in Figure X. Due to similar colors and rough features, it was difficult to accurately extract the light/dark sections of the QR code. While further refinement could have been made to the image processing algorithm it was instead decided to attempt to extract the QR code from the scan paths gathered using the photodiode/galvo monitoring system.




Figure 5.18. Printed QR code stack in print bed (A) and after fabrication (B). The poor contrast between the light and dark cells in the QR code make it difficult to extract the stored information using a camera.

It is not possible to extract the QR code from the scan data without performing some initial processing due to the overlapping nature of scan and jump lines in the PBF process, as shown in Figure 16a. For the fabricated QR code it is possible to use the intensity of the laser to filter out a significant amount of extra noise in the scan lines and to keep the sections that correspond to the desired QR code geometry. This is illustrated in Figure 16b. By applying the gaussian grid method discussed in Section 2.6.1 these remaining scan points can be converted into a readable QR code. This approach is sensitive to the input parameters of the processing algorithm. If the threshold is set too low, over detection of white areas will occur (Figure 17a). If the threshold is set too high, under detection of white areas will occur, resulting in white cells being lost (Figure 17c). Figure 17b shows an example of an acceptable setting for the threshold. While this method works for extracting most QR code cells from the toolpath data, it may over or under detect some cells, so additional error correction may need to be included in the QR code as a precaution.

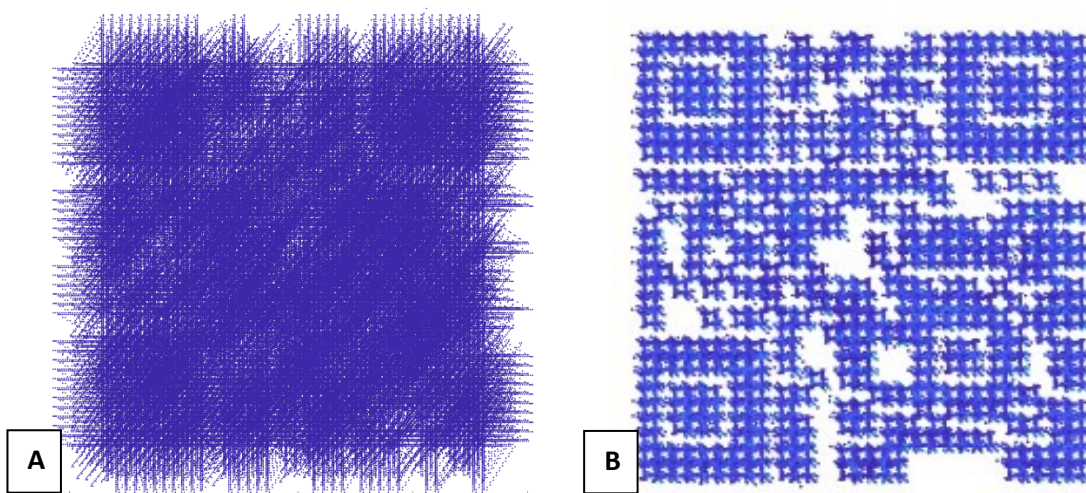


Figure 5.19. Scan line representation of QR code, a) jump lines included b) jump lines removed and scan lines filtered




Figure 5.20. Extracting QR codes from scan lines using overlapping gaussian distributions. A) Threshold value set too low, resulting in over detection B) threshold value set acceptably C) threshold value set too high, resulting in under detection.

Because the “ghost QR codes are not physically fabricated by the AM system, all the intensities are the same and it is not possible to filter out scanlines using intensity. While speed can still be used to eliminate the jump lines, there are still a large number of overlapping datapoints that cannot directly be used to extract the QR code. By using the time domain information (ordering) of the scanlines, the method discussed in Section 2.6.2 can be used to extract the QR code. This approach results in a clean QR code as shown in Figure 18.




Figure 5.21. Extracted “ghost” QR code containing the hex string: f8016e158dfe9280829f052c2018fd92

5.5. Conclusions

In this paper the authors explored the use of an air-gapped SCMS to validate an AM build against cyber-physical sabotage attacks. Discussion was provided on the selection of side-channels as well as on techniques (frequency representation, image representation) for reducing the amount of data that needs to be sent to the SCMS. This work, tested on a 3D systems ProX DMP 320 and using DMP Monitoring sensors, demonstrated the ability to use a cyber-physical hash to transmit information to a SCMS through physical emissions by using both a physical and a “ghost” QR code to store and communicate information. Monitoring the laser path and intensity was able to detect changes to part geometry (both internal and external changes), changes in infill pattern, changes in laser speed, and changes in laser intensity. By using a frequency-based analysis of the infill scan pattern it is possible to reduce that data volume several orders of magnitude (millions to thousands) while still maintaining the ability to detect small changes in the geometry or toolpath. Mean and variance information about the intensity was able to detect changes in laser power, as well as speed differences. Some toolpath changes (internal voids) also generated a detectable change in these values. The work demonstrated the ability to transmit information to an air-gapped SCMS by incorporating data into the model file, through the use of a QR code stack. Two methods for extracting and processing QR codes from the side-channels were presented. The first, used intensity values to filter out non-scanning jump lines and then uses a gaussian based method together with a grid to find the energy overlap of the scan lines and the resulting fabricated QR code. To detect the unprinted “ghost” QR code that contained laser movements without turning on the laser, a second was used which extracted the scan lines from the QR code cell borders and averaged them together in order to determine the center point of each cell. The use of an unprinted QR code stack as a data transmission vector demonstrates an advancement from the authors’ previous work, increasing the amount of data that can be transmitted, while eliminating material consumption, and reducing the “fabrication” time for each QR code by increasing the scanning speed. In conjunction with the use of frequency representation and intensity monitoring to reduce the data volume this data transmission approach demonstrates how an air-gapped SCMS for cyber-physical security could be implemented on a metal PBF system, a method of using physical sensors to detect cyber-attacks. Further work could consider additional techniques for reducing data volumes and comparing side-channel data to detect attacks to create a higher resolution and more robust detection algorithm. On the data package and transmission side, investigation could be done into reducing the size of the QR code cells while still maintain transmissibility and in using direct toolpath access to develop additional, higher bandwidth techniques for transmitting information to the SCMS.

Acknowledgements:

The authors would like to acknowledge the help and contributions of 3D Systems in fabricating parts, data collection, and technical expertise.

This material is based upon work supported by the National Science Foundation under Grant No. CMMI-1436365 and Grant No. CMMI-1635356.

Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

5.6. References

- [1] K. Ahuja, S. Singh, Additive Manufacturing with Metal Powders Market Statistics – 2026, 2019. <https://www.gminsights.com/industry-analysis/additive-manufacturing-with-metal-powders-market> (accessed August 14, 2020).
- [2] New manufacturing milestone: 30,000 additive fuel nozzles | GE Additive, (n.d.). <https://www.ge.com/additive/stories/new-manufacturing-milestone-30000-additive-fuel-nozzles> (accessed August 14, 2020).
- [3] M. Javaid, A. Haleem, Additive manufacturing applications in medical cases: A literature based review, *Alexandria J. Med.* 54 (2018) 411–422. <https://doi.org/10.1016/j.ajme.2017.09.003>.
- [4] Making 3D-printed parts for Boeing 787s | Aerospace America, (n.d.). <https://aerospaceamerica.aiaa.org/departments/making-3d-printed-parts-for-boeing-787s/> (accessed August 14, 2020).
- [5] First titanium 3D-printed part installed into serial production aircraft - Commercial Aircraft - Airbus, (n.d.). <https://www.airbus.com/newsroom/press-releases/en/2017/09/first-titanium-3d-printed-part-installed-into-serial-production-.html> (accessed August 14, 2020).
- [6] H. Post, SpaceX Launches 3D-Printed Part to Space, Creates Printed Engine Chamber, SpaceX. (2014).
- [7] A. Bandyopadhyay, S. Bose, Additive Manufacturing: Pioneering Affordable Aerospace Manufacturing National, (2015).
- [8] N. Falliere, L.O. Murchu, E. Chien, *W32.Stuxnet Dossier*, 4 (2011) 1–69.
- [9] M. Holloway, *Stuxnet Worm Attack on Iranian Nuclear Facilities*, (2015).
- [10] K. (Wired) Zetter, A Cyberattack Has Caused Confirmed Physical Damage for the Second Time Ever | WIRED, (2015). <http://www.wired.com/2015/01/german-steel-mill-hack-destruction/> (accessed January 7, 2016).
- [11] Z. Desmit, A.E. Elhabashy, L.J. Wells, J.A. Camelio, Cyber-physical Vulnerability Assessment in Manufacturing Systems, *Procedia Manuf.* 5 (2016) 1060–1074. <https://doi.org/10.1016/j.promfg.2016.08.075>.
- [12] Z. DeSmit, A.E. Elhabashy, L.J. Wells, J.A. Camelio, An approach to cyber-physical vulnerability assessment for intelligent manufacturing systems, *J. Manuf. Syst.* 43 (2017) 339–351. <https://doi.org/10.1016/j.jmsy.2017.03.004>.
- [13] H. Turner, B. Amos, J. White, J. Camelio, C. Williams, Bad parts: Are our manufacturing systems at risk of silent cyber-attacks, *IEEE Secur. Priv.* (2015) 40–47. <https://doi.org/10.1109/MSP.2015.60>.
- [14] L.J. Wells, J.A. Camelio, C.B. Williams, J. White, Cyber-physical security challenges in manufacturing systems, *Manuf. Lett.* 2 (2014) 74–77. <https://doi.org/10.1016/j.mfglet.2014.01.005>.
- [15] A.E. Elhabashy, J.A. Camelio, L.J. Wells, W.H. Woodall, V. Tech, Misuse of Quality Control Tools in Manufacturing Abstract ID : 2226, (2018) 2–4.
- [16] A.E. Elhabashy, L.J. Wells, J.A. Camelio, W.H. Woodall, A cyber-physical attack taxonomy for production systems: a quality control perspective, *J. Intell. Manuf.* (2018) 1–16.

<https://doi.org/10.1007/s10845-018-1408-9>.

- [17] L.D. Sturm, C.B. Williams, J.A. Camelio, J. White, R. Parker, Cyber-physical vulnerabilities in additive manufacturing systems: A case study attack on the.STL file with human subjects, *J. Manuf. Syst.* 44 (2017). <https://doi.org/10.1016/j.jmsy.2017.05.007>.
- [18] L.D. Sturm, C.B. Williams, J.A. Camelio, J. White, R. Parker, Cyber-Physical Vulnerabilities in Additive Manufacturing Systems, *Solid Free. Fabr. Symp.* (2014) 951–963.
- [19] M. Yampolskiy, W.E. King, J. Gatlin, S. Belikovetsky, A. Brown, A. Skjellum, Y. Elovici, Security of additive manufacturing: Attack taxonomy and survey, *Addit. Manuf.* 21 (2018) 431–457. <https://doi.org/10.1016/j.addma.2018.03.015>.
- [20] M. Yampolskiy, L. Schutzle, U. Vaidya, A. Yasinsac, Security Challenges of Additive Manufacturing with Metals and Alloys, in: M. Rice, S. Shenoi (Eds.), *Crit. Infrastruct. Prot. IX 9th IFIP 11.10 Int. Conf. ICCIP 2015, Arlington, VA, USA, March 16-18, 2015, Revis. Sel. Pap., Springer International Publishing, Cham, 2015: pp. 169–183*.
- [21] L.D.L.D. Sturm, C.B.C.B. Williams, J.A.J.A. Camelio, J. White, R. Parker, Cyber-physical vulnerabilities in additive manufacturing systems: A case study attack on the.STL file with human subjects, *J. Manuf. Syst.* 44 (2017) 154–164. <https://doi.org/10.1016/j.jmsy.2017.05.007>.
- [22] S.E. Zeltmann, N. Gupta, N.G. Tsoutsos, M. Maniatakos, J. Rajendran, R. Karri, Manufacturing and Security Challenges in 3D Printing, *JOM.* 68 (2016) 1872–1881. <https://doi.org/10.1007/s11837-016-1937-7>.
- [23] S. Belikovetsky, M. Yampolskiy, J. Toh, Y. Elovici, drOwned - Cyber-Physical Attack with Additive Manufacturing, *CoRR. abs/1609.0* (2016). <http://arxiv.org/abs/1609.00133>.
- [24] A. Slaughter, M. Yampolskiy, M. Matthews, W.E. King, G. Guss, Y. Elovici, How to Ensure Bad Quality in Metal Additive Manufacturing, *Proc. 12th Int. Conf. Availability, Reliab. Secur. - ARES '17.* (2017) 1–10. <https://doi.org/10.1145/3098954.3107011>.
- [25] M.A. Al Faruque, S.R. Chhetri, A. Canedo, J. Wan, Forensics of thermal side-channel in additive manufacturing systems, *CECS Tech. Report# 16-01.* (2016).
- [26] M.A. Al Faruque, S.R. Chhetri, A. Canedo, J. Wan, Acoustic Side-Channel Attacks on Additive Manufacturing Systems, *2016 ACM/IEEE 7th Int. Conf. Cyber-Physical Syst. ICCPS 2016 - Proc.* (2016) 1–10. <https://doi.org/10.1109/ICCPS.2016.7479068>.
- [27] S.R. Chhetri, A. Canedo, M.A. Al Faruque, Confidentiality Breach Through Acoustic Side-Channel in Cyber-Physical Additive Manufacturing Systems, *ACM Trans. Cyber-Physical Syst.* 2 (2017) 1–25. <https://doi.org/10.1145/3078622>.
- [28] C. Song, F. Lin, Z. Ba, K. Ren, C. Zhou, W. Xu, My Smartphone Knows What You Print: Exploring Smartphone-based Side-channel Attacks Against 3D Printers, in: *Proc. 2016 ACM SIGSAC Conf. Comput. Commun. Secur., ACM, New York, NY, USA, 2016: pp. 895–907*. <https://doi.org/10.1145/2976749.2978300>.
- [29] T. Mativo, C. Fritz, I. Fidan, Cyber acoustic analysis of additively manufactured objects, *Int. J. Adv. Manuf. Technol.* 96 (2018) 581–586. <https://doi.org/10.1007/s00170-018-1603-z>.
- [30] A. Hojjati, A. Adhikari, K. Struckmann, E.J. Chou, T.N.T. Nguyen, K. Madan, M.S. Winslett, C.A.

- Gunter, W.P. King, Leave your phone at the door: Side channels that reveal factory floor secrets BT - 23rd ACM Conference on Computer and Communications Security, CCS 2016, October 24, 2016 - October 28, 2016, 24-28-Octo (2016) 883–894.
<https://doi.org/10.1145/2976749.2978323>.
- [31] J. Brandman, L. Sturm, J. White, C. Williams, A physical hash for preventing and detecting cyber-physical attacks in additive manufacturing systems, *J. Manuf. Syst.* 56 (2020) 202–212.
<https://doi.org/10.1016/j.jmsy.2020.05.014>.
- [32] J. Brandman, A Physical Hash for Preventing and Detecting Cyber-Physical Attacks in Additive Manufacturing Systems, Virginia Tech, 2017.
- [33] M. Guri, Y. Solewicz, A. Daidakulov, Y. Elovici, Fansmitter: Acoustic Data Exfiltration from (Speakerless) Air-Gapped Computers, (2016). <http://arxiv.org/abs/1606.05915>.
- [34] M. Guri, D. Bykhovsky, Y. Elovici, Brightness: Leaking Sensitive Data from Air-Gapped Workstations via Screen Brightness, 2019 12th C. Conf. Cybersecurity Privacy, C. 2019. (2019).
<https://doi.org/10.1109/CMI48017.2019.8962137>.
- [35] Hashcat Benchmark - 2x RTX 2080 Ti · GitHub, (n.d.).
<https://gist.github.com/shahril96/168975e2bc8de49492180ae4f1721294> (accessed September 25, 2020).
- [36] S. Coeck, M. Bisht, J. Plas, F. Verbist, Prediction of lack of fusion porosity in selective laser melting based on melt pool monitoring data, *Addit. Manuf.* 25 (2019) 347–356.
<https://doi.org/10.1016/j.addma.2018.11.015>.

6. ATTACH: Additive Toolpath Transmission of an Acoustic Cyber-Physical Hash

Co Authors: Nathan Raeker-Jordan, Christopher Williams

Abstract:

One of the key challenges facing industrial adoption of additive manufacturing is the need to validate both part quality and security of associated digital data. One way to achieve this is through the use of an air-gapped side-channel monitoring system (SCMS), a system that provides independent verification of process actions by monitoring physical emissions. However, this air-gap poses a challenge for communicating information about the part, since standard digital methods of transmitting part information cannot be used. While methods exist for tracking physical parts and for connecting the digital chain, there still remains a need for approaches that are able to bridge the *cyber-physical gap* between the digital data and the physical properties. To solve this problem, the authors present the Additive Toolpath Transmission of an Acoustic Cyber-Physical Hash (ATTACH) framework as an approach for embedding part quality/side-channel quality measurement information into the digital toolpath of a material extrusion system, that is transmitted physically through physical emissions from the printer during the fabrication process to an air-gapped SCMS. This paper demonstrates two approaches for embedding data into the toolpath g-code (between layers and through infill modulation) and receiving this data using a microphone as a side-channel monitoring system. The paper demonstrates the ability to embed 5263 bits of data in a tensile test specimen (165mm x 19mm x 3.2mm) with no significant effect on finished part properties.

Keywords:

Additive Manufacturing, Cyber-physical Security, Side-channel monitoring, Acoustic, Extrusion

6.1. Cyber vulnerabilities and process monitoring in additive manufacturing systems

6.1.1. Ensuring quality and security of cyber-physical manufacturing systems

The rise of digital manufacturing and a push to adopt Industry 4.0 practices has resulted in an ever-increasing amount of information being generated during the manufacturing process. A single part now can have multiple design files, machine specific toolpaths and process parameters, process monitoring data collected during fabrication, and post process inspection data. For additive manufacturing (AM) in particular, the design freedom and distributed nature of part production creates a strong need for security, traceability, and provenance in the manufacturing process. The Global Brand Counterfeiting Report for 2018-2020 estimates that the total amount of counterfeiting globally will reach 1.82 Trillion USD by the year 2020[1]. Traceability and provenance are important for detecting and preventing these counterfeits from entering the supply chain. Recently, researchers have raised growing concerns about the security of manufacturing systems, demonstrating vulnerabilities with the ability to compromise machines and fabricated parts[2–11]. The possibility of malicious attacks adds another layer of difficulty to the challenge of ensuring that good parts make it into the supply chain to the end

user. Manufacturers need to be able to identify the source of design files, that appropriate quality standards are met, and that customers are able to validate parts as genuine and not counterfeit.

To address these needs, a variety of solutions, both physical and digital, have been proposed. On the digital side, traditional product lifecycle management (PLM) has been used. More recently, blockchain has been proposed as a possible solution for tracing parts in the AM supply chain [12–18]. The use of the blockchain has the potential to allow a fabricated part to be traced back to the original design file. However, in order to use a tool like blockchain to establish provenance throughout the entire AM supply chain, it is necessary to include an identifier within the physical part [19–21] through some form of watermarking, serial number, or physically unclonable function (PUF). This allows the physical part to be connected to the digital data. A wide variety of watermarking/PUF approaches have been suggested for AM systems [22], such as randomly placing quantum dots inside a part [23], chemical fingerprinting [18,24,25], microstructure control [26], inherent machine/process variation [27,28], “smart” inks [29], acoustic barcodes [30], and embedded features [31–41] which can be read in a variety of ways including terahertz imaging [38], thermal imaging [41–43], and optical measurements [31–33,37,44].

While improved digital security and watermarking approaches can improve the security of the AM process, there is still a need to create a secure and traceable thread that is connected through the entire AM process chain. In particular, the difficulty is in linking the cyber (digital design data) with the physical (process monitoring, parts). While approaches such as (blockchain) secure the digital thread they are unable to directly connect to the physical outputs from a manufacturing system. Similarly, they often lack a connection to either in process monitoring data, the original design file, or both. These methods need to be tied to strong quality monitoring in order to ensure that good parts are received by the end user.

6.1.2. In situ monitoring via side-channel measurement systems in AM

The capability of AM systems to produce complex parts with features that range in both scale and shape poses a significant challenge for traditional post process inspection methods. Features may be inaccessible to techniques such as computer measurement machines (CMM) or structured light scanning [45]. Even volumetric scanning methods such as computerized tomography (CT) scanning may be unable to ensure quality in all parts, due to the unique ability of the AM process to affect the internal microstructure of part. To address the limitations of traditional quality control, AM systems are increasingly moving towards in situ process monitoring [46–51]. These systems are able to monitor fabrication and can detect defects as they occur. While these approaches represent a significant improvement in process monitoring and quality assurance, the connectedness and digital nature of the process creates opportunities for a bad actor to maliciously affect the process.

One approach for mitigating these cyber-vulnerabilities is through the use of in situ air-gapped side-channel monitoring systems (SCMS). Side-channels are indirect, physical emissions from a system that correspond to the desired properties of the part being fabricated or of the process being run. Typically, side-channels have been used as a way to exfiltrate data from an air-gapped system using thermal [52,53], acoustic [54–63], or visual emissions [64]. These methods have demonstrated the

ability of various systems to transmit data at rates of up to 5-10 bps using physical means, without requiring a digital connection. Using embedded microphones, researchers have demonstrated the ability to reconstruct part toolpaths using acoustic emissions from the system, as shown by Song et al., where a smartphone was used as an IP attack to collect acoustic side-channels and reconstruct part geometries[59]. More recently work has been done to use side-channels as a method for detecting attacks. Belikovetsky et al. demonstrated how an acoustic side-channel could be used as a method for detect attacks on a part file [54]. While these methods have shown promise in using side-channels to detect defects in AM parts they either require prior knowledge of the part being fabricated or that the side-channel be directly, digitally connected to the AM system. If the former case it is necessary to repeated and manually transfer new data to the SCMS, in the latter case the direct connection opens the SCMS to potentially being compromised if the AM system is attacked, neither is an ideal situation. To solve these problems, the authors propose a method for storing information in the part toolpath that can be transmitted through physical emissions to an air-gapped SCMS.

6.1.3. In situ transmission of quality data

In this work the authors seek to leverage these types of physical methods of transmitting information to couple digital design data with the physical emissions of the system, allowing the designing of a part to communicate information regarding the design intent to an air-gapped side-channel monitoring system being used for quality assurance. By automating this process, it removes the need for manual setup of the quality system and reduces the potential for operator errors. Figure 1 shows how quality information can be stored in a part file/toolpath. First, a known good part is fabricated while a SCMS monitors the emissions. These emissions are then processed and aggregated into a condensed data package. This data package is then stored in the model file or toolpath. The new part with the data package is then fabricated while being monitored by the SCMS. The SCMS reads the data package and compares the known good quality values to the emissions from the current part. If the part is outside of expected ranges the SCMS alerts that an attack or defect has occurred. Previous work has demonstrated the ability to implement this cyber-physical hash approach using a camera as the SCMS and a printed QR code as the data package [65].

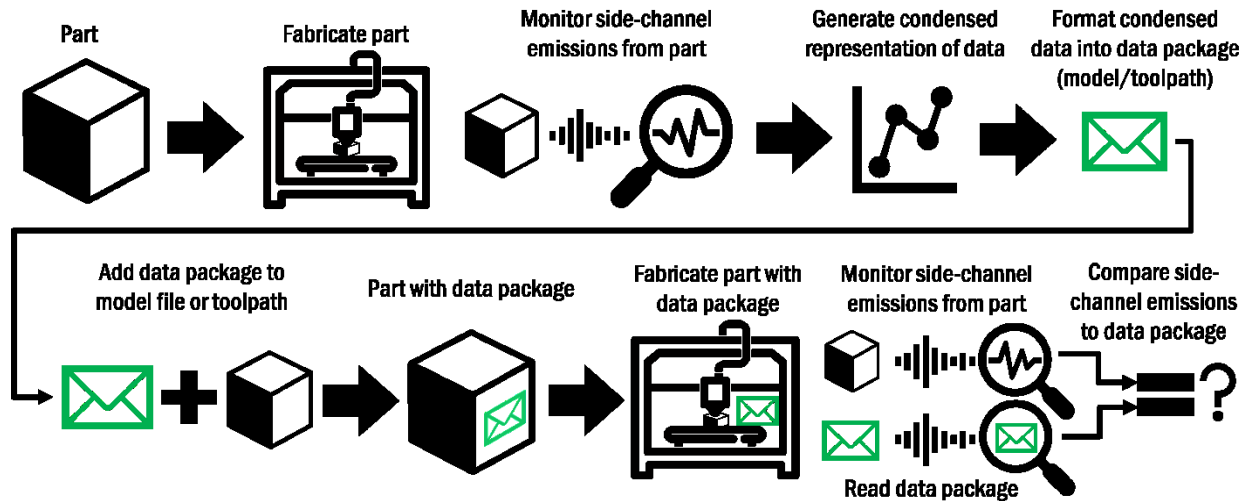


Figure 6.1. Use of a data package stored in the model file/toolpath of a part to transmit quality information to an air-gapped SCMS.

The ATTACH approach allows data to be stored on top of a standard fused filament fabrication (FFF) toolpath through slight modulation of the feed rate. These modulations can be detected by an air-gapped side-channel monitoring system and can be used to communicate part specific information to the monitoring system without having a digital connection that could be subject to being compromised by an attacker. This increases the overall robustness and security of the fabrication process by increasing the difficulty for the attacker to compromise the entire system, since both the toolpath, the machine, and the monitoring system would need to be independently compromised for an attack to occur and escape detection.

A second benefit of the ATTACH approach is that it allows the monitoring system to be independent of the machine doing the fabrication. No changes need to be made to the firmware of the AM system and physical integration of sensors can be kept to a bare minimum (e.g. a microphone outside of the AM system). The only change that needs to be made is to incorporate the modulated data into the toolpath file. This allows the ATTACH approach to be implemented on existing systems, without requiring those systems to be upgraded or changed from their current functional settings. The monitoring system can be easily moved from one machine to another, the only change that needs to be made is to adjust the acoustic processing to match that of the system that it will be used on. This means that the monitoring system can easily and independently be upgraded in order to take advantage of the latest in software security, something that physical manufacturing systems may not always be able to do due to hardware or process limitations.

In this paper, the authors present an approach for adding an additional layer of data on top of an existing toolpath through the use of velocity modulation. This modulation allows information to be communicated to an air-gapped side-channel monitoring system without compromising part quality or increasing fabrication time.

6.2. Acoustically Transmitted Cyber-Physical Hash

To incorporate metadata into the toolpath, it is necessary to choose a modification that will result in a detectable physical signal. In the case of an additive manufacturing system that uses stepper motors for x-y movement, one such signal is the acoustic noise generated by the motors, which directly corresponds to the speed at which the motors are operating. While not the only possible side-channel, this acoustic method has several advantages, i) fast response time, ii) monitored using an inexpensive microphone, iii) not sensitive to changes in lighting, iv) minimally affected by the overall part geometry (compared to optical or penetrating techniques), v) minimally invasive to the printer. Acoustic emissions also have the advantage of being tied only to the movement and do not require modification of other process parameters to convey information. One drawback of this method is sensitivity to external noise, which could be a concern in a noisy manufacturing environment. Other signal vectors are either slow (thermal), intrusive (direct position monitoring), or require a more expensive and sensitive instrumentation setup (camera vision). Previous literature has also shown acoustic monitoring to be a promising method of process monitoring [66]. While similar transmission methods could be developed for other side-channels, this work chooses to focus on the acoustic side-channel as the transmission approach.

6.2.1. Velocity modulation as a transmission method

Normal printing operations on an FFF printer generate acoustic emissions directly related to the power radiating from the stepper motors and the natural vibration frequencies of the system[72]. When using stepper motors, the acoustic frequencies that can be generated by an additive manufacturing system directly correspond to the number of steps per second that the stepper motors can generate as shown in Equation 1.

$$Feed\ rate * \frac{Revolutions}{Feed} * \frac{Steps}{Revolution} = Tone\ (Hz) \quad (1)$$

By increasing/decreasing the feed rate it is possible to increase/decrease the frequency of the tone generated as shown in Figure 2. Changing the feed rate of a build can affect the material properties of the fabricated part (due to part under/overfill or increase/decreased cooling time between roads/layers) and will affect the overall build time. However, by using small, short modulations any potential effects on material properties can be kept to a minimum. To address the issue of build time, modulations can be sent in low-high (0) or high-low (1) pairs. By combining two oppositely modulated tones, the total time can be kept constant and the value of the bit being transmitted can be determined by the ordering of the modulation.




Figure 6.2. Example of frequency modulation of toolpath infill.

6.2.2. Selecting toolpath roads to modulate

A standard g-code file for an FFF printer consists of lines of linear movement commands, referred to as roads, to set points at set feed rates. These movements fall into three broad categories (that usually occur at different rates) i) contours, ii) infill, iii) travel. Contours follow the outline of the part and are usually done at the slowest speed to ensure the best external geometry tolerance. Fill lines place material that is enclosed inside the part and are usually a fixed pattern at a faster rate since there is less need for accuracy. Travel lines are movements of the print head where no extrusion takes place. These moves do not affect the quality of the part and are generally performed at the fastest possible speed in order to minimize fabrication time. While many types of infill patterns exist, for the purpose of this study only a 45° solid infill pattern is considered as this represents a standard setting when fabricating the type of functional load bearing parts that would require robust quality monitoring.

The modulation approach has three core components:

- The feed rate of the base g-code,
- The amount of modulation for the tone
- The length of the tone.

For the purpose of inserting data using frequency modulation, the base frequency is fixed to the carrier frequency (the base toolpath). Of the three movement categories, the two best candidates for adding modulation are travel movements (which will not affect build quality) and fill lines (where movement rate is consistent and accuracy are not as critical). While travel lines have the advantage of not affecting extruded material, they usually occur at a much faster feed rate than the rest of the toolpath. Because the movements are faster, the length of the tones are limited to being much shorter. The capacity for information in a part is directly related to the length and number of eligible roads in that part. Since larger parts generally have more and longer roads, they will also have a larger capacity for information.

A simple histogram of road lengths can be used as a quick metric to determine the potential data capacity and corresponding tone settings as shown in Figure 3.

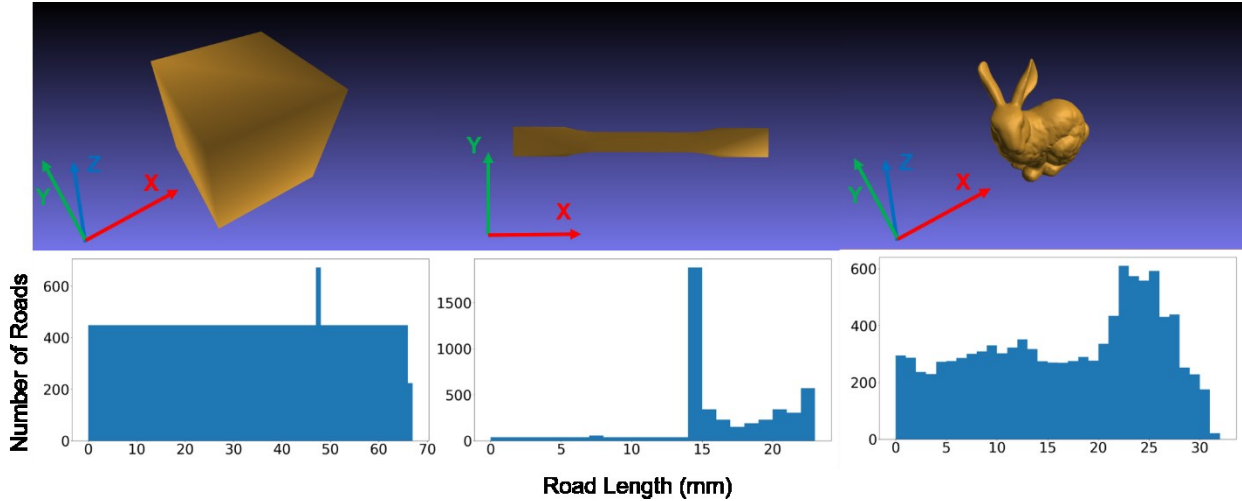


Figure 6.3. Histograms comparing the length of 45° infill roads in different part geometries. A simple 50mm cube has a flat distribution of road lengths. A tensile test specimen (165mm x 19mm x 3.2mm) has a spike corresponding to the road length in the gauge section. The Stanford Bunny has an uneven distribution of road lengths, with a larger cluster of long roads in the body of the bunny (50mm x 38.7mm x 50mm).

Since each road in a build has a fixed length, feed rate, and acceleration, their interaction determines the maximum possible tone length that can be inserted in a given road as shown in Equation 2:

$$t = \left(\left(\frac{d}{v} - \frac{v}{a} \right) - 2n \frac{\Delta v}{a} \right) / n \tag{2}$$

Where d is the length of the road, v is the feed rate of the road, a is the acceleration rate of the printer, Δv is the velocity modulation amount, and n is the number of modulations. The first part of the equation represents the time that printer is moving at the set feed rate, and the second part represents the time required to accelerate/decelerate from the modulated frequency. In systems where the value of acceleration is significantly higher than the velocity, a simplified approximation can be used (Equation 3).

$$t = d / (v * n). \tag{3}$$

Since the length of roads in a part vary, so do the locations where a tone of a given length can be inserted. Shorter tones will be able to be inserted in shorter roads and multiple short tones may be inserted in one sufficiently long road. The number and length of these roads determines one half of the potential amount of data that can be stored. The other half of the data equation is the rate at which the

tone can be sent and interpreted. Faster tones will allow for more data to be inserted; however, the minimum tone length is limited by the change in velocity and the acceleration speed of the system. The other limiting factor is the accuracy at which the side-channel monitoring system can resolve the tones. As tones become faster, it becomes more difficult for the monitoring system to detect the change. For a chosen tone length, the total amount of data that can be stored in a toolpath is shown in Equation 4:

$$\#Bits = (bits\ per\ tone) * \sum \frac{road\ length}{tone\ length} \tag{4}$$

Where the total number of bits that can be stored is the number of tone levels multiplied by the sum of the integer amount of tones that can fit in a road. This value is geometry and rate dependent. Figure 4 shows how this varies in an ASTM D638-14 Type 1 tensile test specimen, with varying tone lengths.




Figure 6.4. Theoretical maximum amount of transmittable data for infill modulation of an ASTM D638-14 Type 1 tensile test specimen using a single level of modulation. As tone time or feed rate increase, the number of roads long enough to hold data decrease.

The feed rate is set by the parameters of the roads that that are chosen, in this case the fill lines. This imposes an upper boundary on both the modulation and the tone length. When selecting the modulation, the primary considerations are being able to discern the modulated tones from standard printing tones and to avoid causing a physical effect on the part. The modulation needs to be sufficiently different from the base feed rate so that the signals do not overlap. While increasing modulation amount does increase separation from the base feed rate, larger modulations will require more time to accelerate to the modulated speed, shortening the length of the tone, as shown in Equation 2. Large modulations are also more likely to cause a physical effect on the part since large changes in velocity

could lead to uneven extrusion. Since the acceleration rate of FFF systems is usually high relative to the feed rate, this does not present a large issue, except in the case of very large modulations. Another factor to consider in selecting the modulation amount is the frequency bands that this will place the modulation in. The acoustic energy emitted at a given frequency varies based on the physical system. Some frequencies will generate loud tones (resonance) while others will generate quiet tones (dampening). The modulation frequencies should be selected such that both the high and low modulations will result in a sufficiently loud tone that can be detected. Finally, when selecting a modulation amplitude, the value should be selected to avoid overlapping with common printing movements to avoid the potential for false positives. Common printing movements can be identified and avoided by generating an unmodulated toolpath, identifying unique feed rates/frequencies in the toolpath, and then counting the number of occurrences at each frequency

6.2.3. Insertion of modulation into the toolpath

The process for inserting modulation into infill roads in the toolpath is as follows:

- 1) Select roads at the desired feed rate
- 2) Select roads at a 45° angle
- 3) Select roads that are long enough to contain at least one tone (Equation 2,3)
- 4) Calculate the number of tones that can be inserted in each road

Each road is then split into one or more pairs of modulated movements when one half has the feed rate increased and the other half has the feed rate decreased. The length of these modulated segments is adjusted so that the time remains the same (i.e., the fast segment is longer and the slow segment is shorter). The ordering of the fast/slow pairs indicates the bit value. A padding value can also be used for segments that are long enough to contain multiple tones to maintain some separation between the tones. These new roads replace the existing roads in the toolpath. Figure 5 shows an example of a modulated toolpath using two different tone lengths. In Figure 5a, multiple tones are able to be inserted into the longer roads. In Figure 5b the longer tone time means that only a single tone can be inserted in the roads and some roads can no longer contain a tone. Figure 6 shows an enlarged version of Figure 5a.




Figure 6.5. Modulated toolpath comparison between 0.5s tone (a) and 1.0s tone (b) for type 1 ASTM tensile test specimen with a 45°. The shorter tone extends into the roads in the gauge section of the tensile test specimen, while the 1.0s tone is limited to the longer roads in the neck and grip sections.




Figure 6.6. Modulated toolpath of a tensile test specimen. Light blue represents standard infill, dark blue is the slow modulation speed, yellow is the fast modulation speed, and red is the contour/travel road speed. Note that contour/travel movements have been reduced in this image for better color scaling.

6.2.4. Error correction

Due to the noisy nature of acoustic emissions and the potential for background interference it is prudent to include some degree of error correction in messages transmitted using this method. A variety of error correction methods exist, one commonly used approach is Reed-Solomon error correction. Reed-Solomon codes are used in applications such as barcodes, mobile communications, television, and high-speed modems [73]. For a given number of parity bytes (e.g. 32) the decoder can detect and correct up to half as many bytes (e.g. 16) of errors in the message. Depending on the environment more or fewer parity bytes may be included for error correction. While adding additional parity bytes increases the robustness of the system by allowing for additional error correction, this will also reduce the amount of data that can be stored in the toolpath.

6.2.5. Acoustic Processing

As the modulated toolpath is run the embedded microphone detects the frequency shifts between fast/slow movements in the roads. The side-channel monitoring system records these fast/slow acoustic emissions from the build. It is then necessary to process the acoustic data in order to extract the information being transmitted. A general procedure includes:

- 1) Downsample the data to reduce computational cost:
- 2) Notch filter the raw data around the tone frequencies:
- 3) Perform a series of windowed Fast Fourier Transforms (FFT) into the frequency domain
- 4) Threshold the data to reduce noise
- 5) Calculate the spectral centroid of the signal at each time stamp
- 6) Determine if spectrum was above intensity threshold (contained a tone)
- 7) Use the spectral centroid to assign a tonal bin at each timestamp
- 8) Use tone ordering low/high (0), high/low (1) to assign a bit value to the tone

The downsampling reduces both the required memory and the computational cost of processing the acoustic signal. Downsampling will also reduce the resolution of the data, so it is necessary to select a value that is a good compromise between resolution and data volume. In this study all data was downsampled by a factor of five. The notch filters are used to exclude any emissions that are outside of the tone frequencies. This decreases the chances of the system picking up build emissions other than the tones and makes it easier to extract the data from the tones. A third order Butterworth filter with a band size of 4Hz (centered on the tone frequencies) was used in this study. The windowed FFTs create time dependent frequency bins for extracting the tonal information. The ideal windowing size varies somewhat based on the length of the tone being sent, shorter tones will require shorter windows, but longer windows will tend to be less noisy. The ideal threshold level also varies depending on the inherent loudness of the modulated bands, but a threshold of -40dB was used for visualization purposes in the spectrograms shown. The spectral centroid is the barycenter of the spectrum. It is a weighted average of the frequencies in a time window [74]. The spectral intensity threshold is used to determine if the timestep contains a tone. If the sum total of energy contained in the timestep is above the threshold it is assumed to contain a tone if it is below the threshold it is assumed to not contain a tone. If a timestamp is determined to contain a tone it is assigned to either a high or low bin. Since each transmitted bit is a pair of low/high or high/low tones, only one set of the tones (high or low) is needed to extract the data. Once the tones have been extracted, a bit value is assigned to each pair. The first bit is determined from the first tone pair and a sign bit is used to flip between 0 and 1. If only half of the tones are used (e.g. the low half of the tones) the change in sign can be determined by the spacing between the tones. A half tone gap (low-high, low-high) indicates that the next message bit is the same as the previous one while a zero (high-low, low-high) or two (low-high, high-low) tone gap indicates that the next message bit is opposite of the current one.

6.3. Experimental methods

There were three primary goals of the experimental work:

- Investigate the effects of modulation parameters on transmissibility
- Demonstrate the ability to successfully transmit a message stored in the toolpath to the SCMS
- Investigate the effect(s) of modulating the toolpath on the final part properties

To investigate the effects of modulation parameters multiple builds were fabricated using a variety of different feed rates, tone lengths, and modulation amplitudes. After identifying an optimal set of parameters for transmissibility a test message was transmitted to the SCMS to demonstrate the viability

of the approach. Finally, a set parts with a worst-case scenario (largest effect on the toolpath) was fabricated and compared against a control set using tensile testing.

6.3.1. Physical System

The physical system that was used chosen to validate this approach was a Lulzbot Taz6 with a single extruder[75]. This system was chosen due to its widespread use and direct control of the toolpath. For the recording device, an inexpensive (\$28) FiFine K669B microphone with a frequency response range of 20-20kHz was used[76]. The system used represents easily accessible and affordable equipment that could easily be implemented on an existing system. If more fidelity is required for an application, additional microphones or more expensive equipment could be used.

The maximum travel speed of the Lulzbot Taz6 is 200 mm/s (firmware limited) and there are 0.5025 revolutions per millimeter of feed, corresponding to a maximum motor speed of 100.5 revolutions per second. The Nema 17 motors (MS17HD6P4150) used in the Lulzbot Taz6 have 200 steps/revolution making the highest frequency tone that can be generated on the system 20.1 kHz, which is just above the response range of the microphone used. Since modulation will be inserted below the maximum travel speed the tones fall within the response range. The maximum acceleration of the Taz6 is 500mm/s² meaning that the time to accelerate from a stationary position to maximum speed is 400ms and the acceleration time to switch directions at maximum speed is 800ms. The fastest feed rate used in the study was 22.5 mm/s, which requires 45ms of acceleration and the largest tone modulation was 100Hz, which requires 2ms of acceleration. Because these values are small compared to the length of the tones used, the simplified form of equation 1 can be used to approximate tone lengths.

6.3.2. Tone parameters effect on transmissibility

To determine the effect of different parameters on the transmissibility of the data, the base feed rate, tone time, and tone modulation were varied, as shown in Table 1. The part that was used for the base toolpath was an ASTM D638-14 Type 1 tensile test specimen. Due to geometric constraints, some combinations were not possible (e.g., a 1.5s tone at 22.5mm/s feed rate would be approximately 33.75mm long, which was longer than the longest road in the test specimen). A total of twenty-four builds were fabricated with a series of tones inserted into the first layer. For repeatability, each build was recorded and the saved audio was then processed. This allowed the audio processing on the monitoring system to be tuned while avoiding any build-to-build variations and reducing the number of builds required to tune the system.

Table 6.1. Test Parameters used and corresponding road distance of tones. Combinations of feedrates and tone times that cause the message length to exceed ~24mm were too long to fit within the test specimen and were excluded.

<i>Parameter</i>	<i>Test Value</i>
Infill Density	Solid
Infill Pattern	45°
Feed Rates (mm/s)	12.5, 15.0, 17.5, 20.0, 22.5
Tone Time (s)	0.5, 1.0, 1.5
Modulation (Hz)	20, 100

6.3.3. Demonstrate the ability to successfully transmit a message stored in the toolpath to the SCMS

In order to test the ability to store, transmit, and receive information a simple test message string containing feed rate and extruder temperature was used. The 10-byte string shown in Figure 7 used reed-solomon code with 5- bytes of error correction, allowing up to half of the message to be corrected in case of errors. The total length of the encoded message was 20bytes (160-bit pairs of tones). The message was inserted into a tensile test specimen with an infill speed of 20mm/s (feed rate 1200) using 20Hz modulation and 1s tones.

Message = "F1200 T225"

0110001010001100	0100110000001100
0000110000000100	0010101001001100
0100110010101100	1100000101000010
0110101111100101	1010001000011101
0011000011101111	1101000101000011

Figure 6.7. Message representation in ASCII and reed-solomon encoded binary with 5-bytes of error correction. The total length of the encoded message is 20 bytes (160 bits).

For the acoustic processing, the same downsampling and notch filter parameters as detailed in Section 2.5 were used. A window size of 1323 (0.15*Sample Rate) was used for the FFT, a threshold of -30dB was used to further reduce external noise and a value of 0.025 for the spectrum intensity threshold was experimentally determined and used to identify time windows containing tones.

6.3.4. Frequency Modulation Effect on Part Strength

To determine the effects of modulating infill speed on part strength, a two sets of D638-14 Type 1 tensile test specimens were fabricated, each containing five samples. The first set was fabricated as a control using a fixed infill speed of 15mm/s. The second set was fabricated using infill modulation with the following settings:

- Tone time: 500ms
- Padding time: 200ms
- Modulation: 100Hz

Toolpaths were generated using Cura (Lulzbot Edition 3.6.20) and fabricated on a Lulzbot Taz6 single extruder printer. The message length was set to 5263 bits (the maximum possible with these settings as determined by the technique described in Section 2) to ensure that the modulation occurred throughout the height of the part, instead of being limited to a few layers, so that any negative effects of the modulation would be maximized. The higher modulation amount was chosen, since larger modulation would be expected to have a greater effect on the part. Finally, the tone time of 500ms and feed rate were chosen in order to ensure that modulation would occur within the gauge section of the part (Figure 5a and Figure 6). For longer tones or faster feed rates, the modulation may become limited by the cross-sectional infill area, causing the modulation to occur only in the wider sections of the part (Figure 4). While this would be beneficial to real parts if the modulation reduced the local material properties, it also would cause any effect to be undetectable when testing the sample, since failure would occur in the gauge section. The inclusion of the modulation in the gauge section represents the worst-case scenario and will display the greatest effect on part performance. Test specimens were fabricated in the XY, orientation (i.e. the bar length oriented along the x-axis). Because the tones are inserted as high-low and low-high pairs within a single road, the overall road time and layer time remain the same. Consequently, there should be no effect on the interlayer strength of the part. As a result, it was unnecessary to fabricate test specimens in orientations outside of the x-y plane, since any effect of modulation on part strength would be primarily limited to the x-y plane.

6.4. Results and Discussion

6.4.1. Effect of parameters on transmissibility

Twenty-four tensile process specimens were printed using the parameters shown in Table 1 to validate the effectiveness of different process parameters on the tone transmission. All the tone time/feed rate combinations tested were detectible; however, as shown in Figure 8, the leading and trailing edges of the tone become slightly more rounded, when feed rate increases, which can be attributed in part to the time spent accelerating to the desired frequency and also due to the effects of the window size in the FFT. This effect is more pronounced on the shorter tones, as they have less time to reach and stabilize at the desired frequency/speed.




Figure 6.8. Overlaid spectrograms of modulated toolpath data for 0.5s, 1.0s, and 1.5s tones with 20Hz modulation.

Figure 9 shows an overlaid comparison of different feed rates using a 1.0s tone with 20Hz of modulation. As the feed rate increases, so does the corresponding tone distance, which results in fewer roads long enough to contain tones. In the 12.5mm/s case, all of the tones are able to be inserted into sequential roads; however, as the feed rate increases, these tones become split over increasingly large gaps as previously shown in Figure 5. While all of the tones are visible, certain frequency bands generate louder and/or more equal tones making them better for transmission. The loudest tones were detected in the upper bands of 22.5 mm/s and 17.5 mm/s, while the quietest was below 22.5 mm/s. The feed rates of 20.0 mm/s and 15.0 mm/s had the most balanced level of noise, 20mm/s had the best combination of balance and loudness.




Figure 6.9. Overlaid spectrograms of modulated toolpath data for 1.0s tones, 20Hz modulation.

Figure 10 shows an overlaid comparison of different feed rates using a 1.0s tone with 100Hz of modulation. As in Figure 9, while all the tones are detectable, some combinations produce more ideal results. The 20.0 mm/s range produces the best combination of loudness and balance, followed by the 15.0 mm/s and 12.5 mm/s ranges. The 17.5 mm/s tones are imbalanced with the top tone being significantly louder, which is also true for the 22.5 mm/s tone.

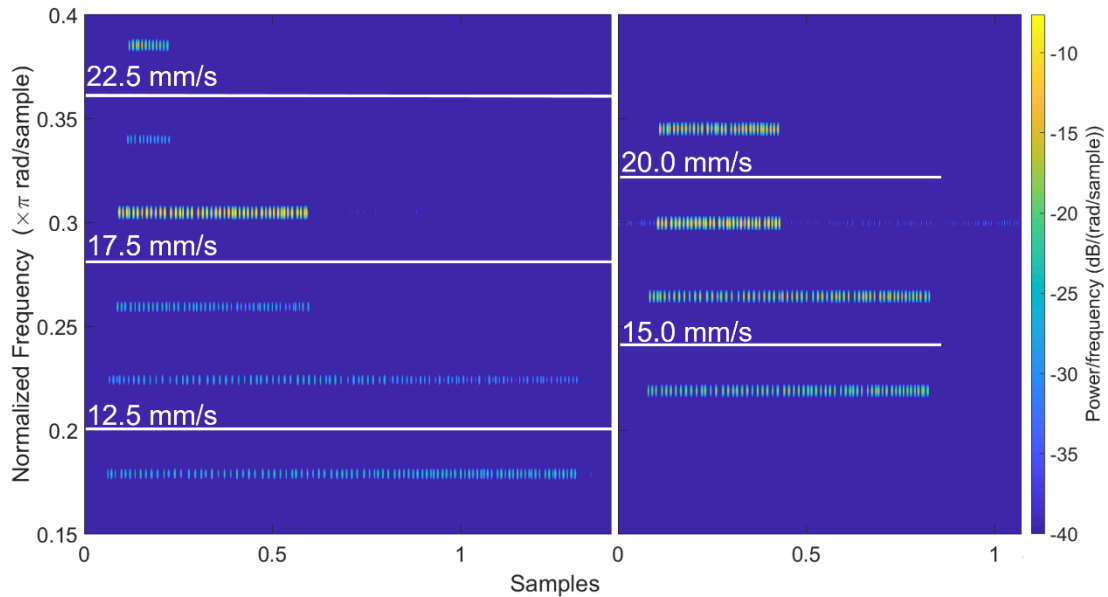


Figure 6.10. Overlaid spectrograms of modulated toolpath data for 1.0s tones, 100Hz modulation. Note: due to the large modulation the frequencies overlap, so they have been separated for clarity.

These results show that the choice of modulation amount is important for achieving tones that are easy to detect and extract. Because the system will emit more noise at certain frequencies, it is important to select modulation(s) that will place the tones in bands that are loud and balanced. Failing to do so may result in a weak signal which can be difficult to accurately filter and extract from the system noise and thus result in poor transmission. It is important to note that the optimal feed rates may differ between systems. While two similar systems (e.g., two Lulzbot Taz6 printers) will likely have similar results, physical characteristics of the system, such as the stepper motors or rigidity of the system, may result in different optimal frequencies. When two systems differ significantly (e.g., different stepper motors, different construction, etc.) the optimal modulation parameters will also be likely to differ. Before implementing toolpath modulation, each system should be evaluated to determine the ideal parameters. Once parameters have been established they can continue to be used for different parts on the same system.

6.4.2. Message Transmission

During testing it was discovered that the volume of the transmitted tones varied based on the layer. Odd layers generated louder tones while even layers generated quieter tones. The reason for this layer-to-layer variation was the alternating direction of the infill pattern as shown in Figure 11. On odd layers both stepper motors move in the same signed direction (i.e. both positive or both negative). On even layers the stepper motors move in opposite signed directions (i.e. one positive and one negative). The result is that when the stepper motors have the same signed movement the tones are louder and when they have opposite signed movements the tones are quieter. This is likely due to the specific resonance modes of the authors' Taz6 printer causing either more constructive or more destructive vibration in the printer from the two stepper motors. While the potential for these effects should be considered when implementing this approach, it is important to note that this affect may not happen on other models of printers or even on other Taz6 printers depending on the specific physical properties of those systems.

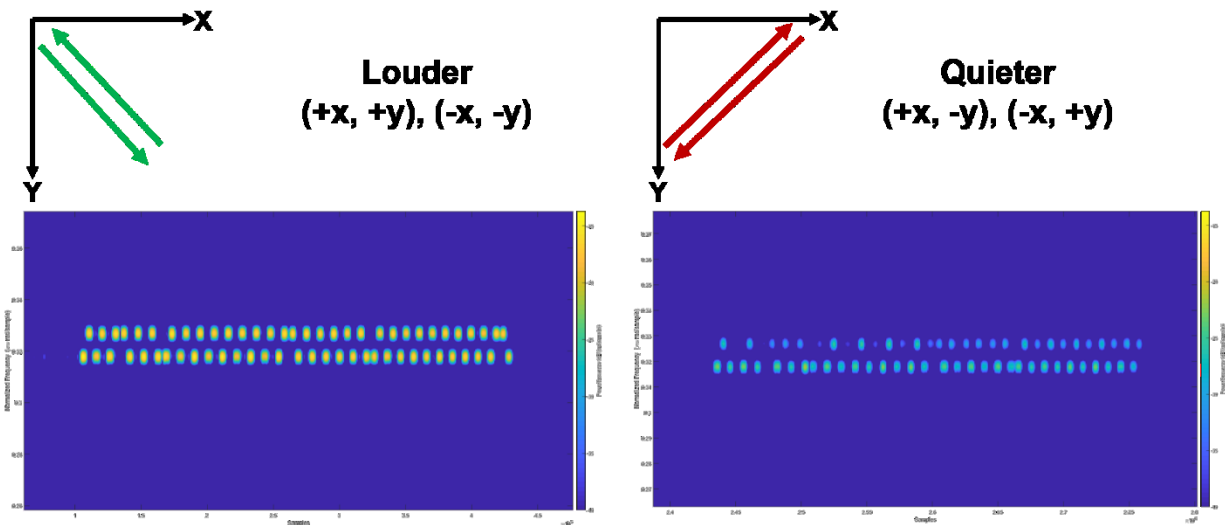


Figure 6.11. Comparison of tone transmission between odd layers (A) and even layers (B). When both stepper motors are moving in the same signed direction the tones are louder. When the stepper motors are moving in opposite signed directions the resulting tones are quieter.

To avoid the effects on reduced tone strength the message was inserted only in odd part layers. While this did reduce the maximum amount of data that could be transmitted it reduced the possibility to missing data due to the tones being too quiet. On systems without this effect the full range of layers could be used. After the tones had been inserted into the g-code the part was fabricated while being monitored. The system then extracted and processed the tones. The first 32 extracted bits are shown in Figure 12. The approach was able to successfully transmit all 160 bits without errors. With a longer message or increased background noise there would be a greater possibility for missed or incorrect bits, resulting in the reed-solomon encoding being used for error correction.




Figure 6.12. Example of message extraction from acoustic emissions. The colored area in the middle is the spectrogram representation of the low modulation. The orange bars at the bottom are the corresponding spectral centroid results (shifted down for visibility) and the ones and zeros are the message bits being interpreted. “01100010100011000100110000001100” = “F1”

6.4.3. Effect on tensile strength

A worst-case scenario (maximum alteration of the toolpath) set of tensile specimens and a set of control specimens were fabricated and destructively evaluated as per section 3.3. The results of the tensile tests are shown in Table 2. Tensile results, of FFF specimen are also included from previous literature as an example of the typical standard deviation of maximum tensile stress found in FFF parts. Figure 13 shows the samples before and after testing. The difference between the control and modulated samples in average maximum tensile stress was 0.86 MPa (2.51% difference) which was within one standard deviation. The set of modulated samples had a slightly larger standard deviation than the control set (0.94 MPa greater); however, both sets of samples had a smaller standard deviation than similar tests reported by Tanikella [77], showing that the variation falls within the standard variation expected for ABS parts on this type of desktop FFF system. The average maximum tensile stress in both groups was higher than that reported by Tanikella[77]; however, this can likely be attributed to differences in print settings (not specifically reported by Tanikella) and material brand. This shows that it is possible to insert

data into the toolpath through modulation without affecting the material properties of the build, resulting in comparable part quality to a non-modulated build.



Figure 6.13. Tensile test specimens before and after testing. (A) modulated specimens, (B) control specimens. There is a minor visual difference between the modulated samples and the control samples. This was caused by the modulation creating slightly more roughness on the top surface.

Table 6.2. Tensile test results for control and modulated ABS parts fabricated using desktop additive manufacturing systems

<i>Sample Group</i>	<i>Average Maximum Load (N)</i>	<i>Average Maximum Tensile Stress (MPa)</i>	<i>Standard Deviation of Maximum Tensile Stress (MPa)</i>
<i>Control</i>	1562.11	34.31	0.94
<i>Modulated</i>	1510.60	33.46	1.88
<i>Tanikella [77]</i>	1196.12	28.75	3.15

6.5. Conclusions

The growing risk of cyber-physical attacks on AM systems calls for the development of secure monitoring systems that are able to validate builds. In situ air-gapped side-channel monitoring systems are a method that has been demonstrated as capable of detecting malicious changes in part/machine properties while at the same time being more robust against being maliciously compromised than inline connected sensors. However, the use of an air-gap to improve security robustness poses a challenge for securely passing data to the monitoring system for comparison against the known good. By inserting information into a part's toolpath using the ATTACH method, it is possible to securely communicate to an air-gapped side-channel monitoring system using physical emissions from the fabrication of the part.

This paper discusses the requirements needed to insert information into a part toolpath using the ATTACH method and presents methods for evaluating parts to determine data storage capacity and suitability. This study successfully demonstrated the ability to insert information into the g-code of an FFF part at a variety of feed rates (12.5mm/s, 15.0mm/s, 17.5mm/s, 20.0mm/s, 22.5mm/s), modulation amplitudes (20Hz, 100Hz), and tone lengths (0.5s, 1.0s, 1.5s) and provides discussion on the criteria for selecting appropriate settings. The storage, transmission, and acoustic processing of a simple quality information message was demonstrated in Section 4.2. In Section 4.3 the ATTACH method allowed up to 5263 bits of data to be stored in an ASTM standard tensile test specimen without having a significant effect on part strength and without changing the fabrication time of the part. This method creates a new way to transmit part specific information to side-channel monitoring systems across machines.

Future work in this area could increase the transmissible data by refining the acoustic processing to detect tones shorter than 0.5s and also identifying the frequency bands that demonstrate the most detectible tones. Additionally, the use of multiple modulation amounts within the same build could be used to greatly increase the amount of data that could be sent per tone by encoding in octal or hex instead of binary. In addition to this, the methods used in the ATTACH method can be generalized to use toolpath modulation as a way to transfer information through the toolpath of parts fabricated on other AM systems besides FFF. One potential application is in powder bed fusion system, where the large number of scan lines has the potential to store large amounts of data. With high resolution laser position tracking, modulation could be inserted into travel movements without affecting build quality.

Acknowledgements

This material is based upon work supported by the National Science Foundation under Grant No. CMMI-1436365 and Grant No. CMMI-1635356.

Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

6.6. References

- [1] 5 Industrial IoT (IIoT) Predictions for 2019 | BehrTech Blog, (n.d.). <https://behrtech.com/blog/5-iiot-predictions-for-2019/> (accessed September 3, 2020).
- [2] S.M. Bridges, S.T.A.T. Hall, S.J. Graves, S.T.A.T. Hall, K. Keiser, S.T.A.T. Hall, N. Sissom, S.T.A.T. Hall, S.J. Graves, Cyber Security for Additive Manufacturing, in: Proc. 10th Annu. Cyber Inf. Secur. Res. Conf., ACM, New York, NY, USA, 2015: pp. 14:1----14:3. <https://doi.org/10.1145/2746266.2746280>.
- [3] S. Goldenberg, J. Brown, J. Haid, J. Ezzard, 3D opportunity and cyber risk management, Deloitte Univ. Press. (2016). <https://doi.org/10.1016/j.jval.2017.05.018>.
- [4] C. Xiao, Security Attack to 3D Printing, (2013). <https://www.claudxiao.net/Attack3DPrinting-Claud-en.pdf>.
- [5] L.D. Sturm, C.B. Williams, J.A. Camelio, J. White, R. Parker, Cyber-Physical Vulnerabilities in Additive Manufacturing Systems, Solid Free. Fabr. Symp. (2014) 951–963.
- [6] L.D. Sturm, C.B. Williams, J.A. Camelio, J. White, R. Parker, Cyber-physical vulnerabilities in additive manufacturing systems: A case study attack on the.STL file with human subjects, J. Manuf. Syst. 44 (2017). <https://doi.org/10.1016/j.jmsy.2017.05.007>.
- [7] M. Yampolskiy, W.E. King, J. Gatlin, S. Belikovetsky, A. Brown, A. Skjellum, Y. Elovici, Security of additive manufacturing: Attack taxonomy and survey, Addit. Manuf. 21 (2018) 431–457. <https://doi.org/10.1016/j.addma.2018.03.015>.
- [8] Y. Pan, J. White, D.C. Schmidt, A. Elhabashy, L. Sturm, J. Camelio, C. Williams, Taxonomies for Reasoning About Cyber-physical Attacks in IoT-based Manufacturing Systems, Int. J. Interact. Multimed. Artif. Intel. 4 (2017) 45–54. <https://doi.org/10.9781/ijimai.2017.437>.
- [9] S.E. Zeltmann, N. Gupta, N.G. Tsoutsos, M. Maniatakos, J. Rajendran, R. Karri, Manufacturing and Security Challenges in 3D Printing, JOM. 68 (2016) 1872–1881. <https://doi.org/10.1007/s11837-016-1937-7>.
- [10] A. Slaughter, M. Yampolskiy, M. Matthews, W.E. King, G. Guss, Y. Elovici, How to Ensure Bad Quality in Metal Additive Manufacturing, Proc. 12th Int. Conf. Availability, Reliab. Secur. - ARES '17. (2017) 1–10. <https://doi.org/10.1145/3098954.3107011>.
- [11] S. Belikovetsky, M. Yampolskiy, J. Toh, Y. Elovici, dr0wned - Cyber-Physical Attack with Additive Manufacturing, CoRR. abs/1609.0 (2016). <http://arxiv.org/abs/1609.00133>.
- [12] S. Trouton, M. Vitale, J. Killmeyer, 3D opportunity for blockchain, Deloitte Univ. Press. (2016). <https://dupress.deloitte.com/dup-us-en/focus/3d-opportunity/3d-printing-blockchain-in-manufacturing.html>.
- [13] A. Angrish, B. Craver, M. Hasan, B. Starly, A Case Study for Blockchain in Manufacturing: “FabRec”: A Prototype for Peer-to-Peer Network of Manufacturing Nodes, (2018). <https://arxiv.org/abs/1804.01083>.
- [14] S. Kurpjuweit, C.G. Schmidt, M. Klöckner, S.M. Wagner, Blockchain in Additive Manufacturing and its Impact on Supply Chains, J. Bus. Logist. (2019) 1–25. <https://doi.org/10.1111/jbl.12231>.

- [15] M. Holland, J. Stjepandic, C. Nigischer, Intellectual Property Protection of 3D Print Supply Chain with Blockchain Technology, 2018 IEEE Int. Conf. Eng. Technol. Innov. ICE/ITMC 2018 - Proc. (2018) 1–8. <https://doi.org/10.1109/ICE.2018.8436315>.
- [16] M. Holland, C. Nigischer, J. Stjepandic, Copyright protection in additive manufacturing with blockchain approach, *Adv. Transdiscipl. Eng.* 5 (2017) 914–921. <https://doi.org/10.3233/978-1-61499-779-5-914>.
- [17] I. Aliende Povedano, P. De Oro Martinez, *Pedagogia, MLearning La Form. En Tu Movil.* 00 (2011) 71–90. <https://doi.org/10.4272/978-84-9745-269-4.ch4>.
- [18] Z.C. Kennedy, D.E. Stephenson, J.F. Christ, T.R. Pope, B.W. Arey, C.A. Barrett, M.G. Warner, Enhanced anti-counterfeiting measures for additive manufacturing: Coupling lanthanide nanomaterial chemical signatures with blockchain technology, *J. Mater. Chem. C* 5 (2017) 9570–9578. <https://doi.org/10.1039/c7tc03348f>.
- [19] N.F. Fadhel, R.M. Crowder, G.B. Wills, Provenance in the additive manufacturing process, *IFAC-PapersOnLine.* 28 (2015) 2345–2350. <https://doi.org/10.1016/j.ifacol.2015.06.438>.
- [20] N.F. Fadhel, R.M. Crowder, G.B. Wills, Approaches to maintaining provenance throughout the additive manufacturing process, 2013 World Congr. Internet Secur. WorldCIS 2013. (2013) 82–87. <https://doi.org/10.1109/WorldCIS.2013.6751022>.
- [21] N.F. Fadhel, R.M. Crowder, F. Akeel, G.B. Wills, Component for 3D printing provenance framework: Security properties components for provenance framework, 2014 World Congr. Internet Secur. WorldCIS 2014. (2014) 91–96. <https://doi.org/10.1109/WorldCIS.2014.7028174>.
- [22] B. Macq, P.R. Alface, M. Montanola, Applicability of watermarking for intellectual property rights protection in a 3D printing scenario, (2015) 89–95. <https://doi.org/10.1145/2775292.2775313>.
- [23] O. Ivanova, A. Elliott, T. Campbell, C.B. Williams, Unclonable security features for additive manufacturing, *Addit. Manuf.* 1 (2014) 24–31. <https://doi.org/10.1016/j.addma.2014.07.001>.
- [24] F. Sharon, 3D Fakes: Chemical Fingerprinting in Additive Manufacturing, from Pharmaceuticals to Engines, *NIP Digit. Fabr. Conf. 2017* (2017).
- [25] S. Flank, A.R. Nassar, T.W. Simpson, N. Valentine, E. Elburn, Fast Authentication of Metal Additive Manufacturing, *3D Print. Addit. Manuf.* 4 (2018) 143–148. <https://doi.org/10.1089/3dp.2017.0018>.
- [26] A. Dachowicz, S.C. Chaduvula, M. Atallah, J.H. Panchal, Microstructure-Based Counterfeit Detection in Metal Part Manufacturing, *Jom.* 69 (2017) 2390–2396. <https://doi.org/10.1007/s11837-017-2502-8>.
- [27] F. Peng, J. Yang, Z.X. Lin, M. Long, Source identification of 3D printed objects based on inherent equipment distortion, *Comput. Secur.* 82 (2019) 173–183. <https://doi.org/10.1016/j.cose.2018.12.015>.
- [28] F. Peng, J. Yang, M. Long, 3-D Printed Object Authentication Based on Printing Noise and Digital Signature, *IEEE Trans. Reliab.* 68 (2019) 342–353. <https://doi.org/10.1109/TR.2018.2869303>.

- [29] S.J. Trenfield, H. Xian Tan, A. Awad, A. Buanz, S. Gaisford, A.W. Basit, A. Goyanes, Track-and-Trace: Novel Anti-Counterfeit Measures for 3D Printed Personalised Drug Products using Smart Material Inks, *Int. J. Pharm.* 567 (2019). <https://doi.org/10.1016/j.ijpharm.2019.06.034>.
- [30] C. Harrison, R. Xiao, S. Hudson, Acoustic Barcodes: Passive, Durable and Inexpensive Notched Identification Tags, *Proc. 25th Annu. ACM Symp. User Interface Softw. Technol. - UIST '12.* (2012) 563. <https://doi.org/10.1145/2380116.2380187>.
- [31] J.-U. Hou, D.-G. Kim, S. Choi, H.-K. Lee, 3D Print-Scan Resilient Watermarking Using a Histogram-Based Circular Shift Coding Structure, (2015) 115–121. <https://doi.org/10.1145/2756601.2756607>.
- [32] D. Li, A.S. Nair, S.K. Nayar, C. Zheng, AirCode: Unobtrusive Physical Tags for Digital Fabrication, (2017). <https://doi.org/10.1145/3126594.3126635>.
- [33] A. Delmotte, K. Tanaka, H. Kubo, T. Funatomi, Y. Mukaigawa, Blind Watermarking for 3D Printed Objects by Locally Modifying Layer Thickness, *IEEE Trans. Multimed. PP* (2019) 1. <https://doi.org/10.1109/TMM.2019.2962306>.
- [34] J.U. Hou, D.G. Kim, H.K. Lee, Blind 3D Mesh Watermarking for 3D Printed Model by Analyzing Layering Artifact, *IEEE Trans. Inf. Forensics Secur.* 12 (2017) 2712–2725. <https://doi.org/10.1109/TIFS.2017.2718482>.
- [35] M. Suzuki, P. Silapasuphakornwong, K. Uehira, H. Unno, Y. Takashima, Copyright Protection for 3D Printing by Embedding Information Inside Real Fabricated Objects, (2015) 180–185. <https://doi.org/10.5220/0005342401800185>.
- [36] K. Uehira, M. Suzuki, Copyright Protection for 3D Printing by Embedding Information Inside 3D-Printed Objects, in: *Digit. Forensic Watermarking, 2016*: pp. 370–378. <https://doi.org/10.5220/0005342401800185>.
- [37] F. Chen, Y. Luo, N.G. Tsoutsos, M. Maniatakos, K. Shahin, N. Gupta, Embedding Tracking Codes in Additive Manufactured Parts for Product Authentication, *Adv. Eng. Mater.* 21 (2019) 1–8. <https://doi.org/10.1002/adem.201800495>.
- [38] K. Willis, A. Wilson, InfraStructs: Fabricating Information Inside Physical Objects for Imaging in the Terahertz Region, *Siggraph 2013.* (2013) 138.
- [39] N. Gupta, F. Chen, N.G. Tsoutsos, M. Maniatakos, ObfusCADE: Obfuscating Additive Manufacturing CAD Models Against Counterfeiting : INVITED, *Proc. 54th Annu. Des. Autom. Conf. 2017 - DAC '17.* (2017) 1–6. <https://doi.org/10.1145/3061639.3079847>.
- [40] F. Chen, G. Mac, N. Gupta, Security features embedded in computer aided design (CAD) solid models for additive manufacturing, *Mater. Des.* 128 (2017) 182–194. <https://doi.org/10.1016/j.matdes.2017.04.078>.
- [41] M. Suzuki, P. Silapasuphakornwong, Y. Takashima, H. Torii, H. Unno, K. Uehira, Technique for protecting copyrights of digital data for 3-D printing, and its application to low infill density objects, *MMEDIA 2016 Eighth Int. Conf. Adv. Multimed. Febr. 21–25, 2016, Lisbon, Port.* (2016) 56–59. https://www.thinkmind.org/index.php?view=article&articleid=mmedia_2016_3_40_50043.

- [42] A. Okada, P. Silapasuphakornwong, M. Suzuki, H. Torii, Y. Takashima, K. Uehira, Non-destructively reading out information embedded inside real objects by using far-infrared light, *Appl. Digit. Image Process.* XXXVIII. 9599 (2015) 95992V. <https://doi.org/10.1117/12.2189486>.
- [43] P. Silapasuphakornwong, M. Suzuki, H. Torii, Y. Takashima, *Nondestructive Readout of Copyright Information Embedded in Objects Fabricated with 3-D Printers*, 2016. <https://doi.org/10.1007/978-3-662-43886-2>.
- [44] S. Yamazaki, S. Kagami, M. Mochimaru, *Extracting Watermark from 3D Prints*, *Proc. - Int. Conf. Pattern Recognit.* (2014) 4576–4581. <https://doi.org/10.1109/ICPR.2014.783>.
- [45] M.S. Tootooni, A. Dsouza, R. Donovan, P.K. Rao, Z.J. Kong, P. Borgesen, *Classifying the Dimensional Variation in Additive Manufactured Parts from Laser-Scanned Three-Dimensional Point Cloud Data Using Machine Learning Approaches*, *J. Manuf. Sci. Eng. Trans. ASME.* 139 (2017). <https://doi.org/10.1115/1.4036641>.
- [46] P.K. Rao, J. (Peter) Liu, D. Roberson, Z. (James) Kong, C. Williams, *Online Real-Time Quality Monitoring in Additive Manufacturing Processes Using Heterogeneous Sensors*, *J. Manuf. Sci. Eng.* 137 (2015) 61007–61012. <https://doi.org/10.1115/1.4029823>.
- [47] C. Gobert, E.W. Reutzel, J. Petrich, A.R. Nassar, S. Phoha, *Application of supervised machine learning for defect detection during metallic powder bed fusion additive manufacturing using high resolution imaging.*, *Addit. Manuf.* 21 (2018) 517–528. <https://doi.org/10.1016/j.addma.2018.04.005>.
- [48] H. Kim, Y. Lin, T.L.B. Tseng, *A review on quality control in additive manufacturing*, *Rapid Prototyp. J.* 24 (2018) 645–669. <https://doi.org/10.1108/RPJ-03-2017-0048>.
- [49] S.K. Everton, M. Hirsch, P. Stravroulakis, R.K. Leach, A.T. Clare, P.I. Stavroulakis, R.K. Leach, A.T. Clare, *Review of in-situ process monitoring and in-situ metrology for metal additive manufacturing*, *Mater. Des.* 95 (2016) 431–445. <https://doi.org/10.1016/j.matdes.2016.01.099>.
- [50] S.A. Shevchik, C. Kenel, C. Leinenbach, K. Wasmer, *Acoustic emission for in situ quality monitoring in additive manufacturing using spectral convolutional neural networks*, *Addit. Manuf.* 21 (2018) 598–604. <https://doi.org/10.1016/j.addma.2017.11.012>.
- [51] L. Sturm, M. Albakri, C.B. Williams, P. Tarazaga, *In-Situ Detection of Build Defects in Additive Manufacturing via Impedance-Based Monitoring*, *27th Annu. Int. Solid Free. Fabr. Symp. - An Addit. Manuf. Conf.* (2016) 1458–1478.
- [52] M.A. Al Faruque, S.R. Chhetri, A. Canedo, J. Wan, *Forensics of thermal side-channel in additive manufacturing systems*, *CECS Tech. Report# 16-01.* (2016).
- [53] M. Guri, M. Monitz, Y. Mirski, Y. Elovici, *BitWhisper: Covert Signaling Channel between Air-Gapped Computers Using Thermal Manipulations*, *Proc. Comput. Secur. Found. Work.* 2015-Septe (2015) 276–289. <https://doi.org/10.1109/CSF.2015.26>.
- [54] S. Belikovetsky, Y. Solewicz, M. Yampolskiy, J. Toh, Y. Elovici, *Digital Audio Signature for 3D Printing Integrity*, *IEEE Trans. Inf. Forensics Secur.* PP (2018) 1. <https://doi.org/10.1109/TIFS.2018.2851584>.

- [55] S. Rokka Chhetri, M.A. Al Faruque, Side Channels of Cyber-Physical Systems: Case Study in Additive Manufacturing, *IEEE Des. Test.* 34 (2017) 18–25. <https://doi.org/10.1109/MDAT.2017.2682225>.
- [56] Q.Y. Lu, C.H. Wong, Additive manufacturing process monitoring and control by non-destructive testing techniques: challenges and in-process monitoring, *Virtual Phys. Prototyp.* 13 (2018) 39–48. <https://doi.org/10.1080/17452759.2017.1351201>.
- [57] A. Hojjati, A. Adhikari, K. Struckmann, E.J. Chou, T.N.T. Nguyen, K. Madan, M.S. Winslett, C.A. Gunter, W.P. King, Leave your phone at the door: Side channels that reveal factory floor secrets BT - 23rd ACM Conference on Computer and Communications Security, CCS 2016, October 24, 2016 - October 28, 2016, 24-28-Octo (2016) 883–894. <https://doi.org/10.1145/2976749.2978323>.
- [58] S.R. Chhetri, A. Canedo, M.A. Al Faruque, KCAD: Kinetic Cyber-Attack Detection Method for Cyber-Physical Additive Manufacturing Systems, *Proc. 35th Int. Conf. Comput. Des. - ICCAD '16.* (2016) 1–8. <https://doi.org/10.1145/2966986.2967050>.
- [59] C. Song, F. Lin, Z. Ba, K. Ren, C. Zhou, W. Xu, My Smartphone Knows What You Print, in: *Proc. 2016 ACM SIGSAC Conf. Comput. Commun. Secur. - CCS'16*, ACM, New York, 2016: pp. 895–907. <https://doi.org/10.1145/2976749.2978300>.
- [60] S. Belikovetsky, Y.A. Solewicz, M. Yampolskiy, J. Toh, Y. Elovici, Digital audio signature for 3d printing integrity, *IEEE Trans. Inf. Forensics Secur.* 14 (2019) 1127–1141. <https://doi.org/10.1109/TIFS.2018.2851584>.
- [61] S.A. Shevchik, C. Kenel, C. Leinenbach, K. Wasmer, Acoustic emission for in situ quality monitoring in additive manufacturing using spectral convolutional neural networks, *Addit. Manuf.* 21 (2018) 598–604. <https://doi.org/10.1016/j.addma.2017.11.012>.
- [62] S.R. Chhetri, A. Canedo, M.A. Al Faruque, Confidentiality Breach Through Acoustic Side-Channel in Cyber-Physical Additive Manufacturing Systems, *ACM Trans. Cyber-Physical Syst.* 2 (2017) 1–25. <https://doi.org/10.1145/3078622>.
- [63] M. Guri, Y. Solewicz, A. Daidakulov, Y. Elovici, Fansmitter: Acoustic Data Exfiltration from (Speakerless) Air-Gapped Computers, (2016). <http://arxiv.org/abs/1606.05915>.
- [64] M. Guri, D. Bykhovsky, Y. Elovici, Brightness: Leaking Sensitive Data from Air-Gapped Workstations via Screen Brightness, 2019 12th C. Conf. Cybersecurity Privacy, C. 2019. (2019). <https://doi.org/10.1109/CMI48017.2019.8962137>.
- [65] J. Brandman, L. Sturm, J. White, C. Williams, A physical hash for preventing and detecting cyber-physical attacks in additive manufacturing systems, *J. Manuf. Syst.* 56 (2020) 202–212. <https://doi.org/10.1016/j.jmsy.2020.05.014>.
- [66] S. Yu, A.V. Malawade, S.R. Chhetri, M.A. Al Faruque, Sabotage Attack Detection for Additive Manufacturing Systems, *IEEE Access.* 8 (2020) 27218–27231. <https://doi.org/10.1109/ACCESS.2020.2971947>.
- [67] M.A. Al Faruque, S.R. Chhetri, A. Canedo, J. Wan, Acoustic Side-Channel Attacks on Additive Manufacturing Systems, 2016 ACM/IEEE 7th Int. Conf. Cyber-Physical Syst. ICCPS 2016 - Proc.

- (2016) 1–10. <https://doi.org/10.1109/ICCPS.2016.7479068>.
- [68] T. Mativo, C. Fritz, I. Fidan, Cyber acoustic analysis of additively manufactured objects, *Int. J. Adv. Manuf. Technol.* 96 (2018) 581–586. <https://doi.org/10.1007/s00170-018-1603-z>.
- [69] S. Rokka Chhetri, Novel Side-Channel Attack Model for Cyber-Physical Additive Manufacturing Systems, UC Irvine, 2016. <https://doi.org/10.1016/j.mser.2008.09.002>.
- [70] M.A. Al Faruque, S.R. Chhetri, A. Canedo, J. Wan, Acoustic Side-channel Attacks on Additive Manufacturing Systems, in: *Proc. 7th Int. Conf. Cyber-Physical Syst.*, IEEE Press, Piscataway, 2016: pp. 19:1--19:10.
- [71] C. Bayens, T. Le, L. Garcia, C. Bayens, L. Garcia, See No Evil , Hear No Evil , Feel No Evil , Print No Evil ? Malicious Fill Patterns Detection in Additive Manufacturing, *USENIX Secur.* (2017).
- [72] S.R. Chhetri, M. Abdullah, A. Faruque, Side-Channels of Cyber-Physical Systems : Case Study in Additive Manufacturing, *IEEE Des. Test. PP* (2017) 1. <https://doi.org/10.1109/MDAT.2017.2682225>.
- [73] reed-solomon codes, (n.d.). https://www.cs.cmu.edu/~guyb/realworld/reedsolomon/reed_solomon_codes.html (accessed September 17, 2020).
- [74] G. Peeters, A large set of audio features for sound description (similarities and classification) in the CUIDADO project, 2004. http://recherche.ircam.fr/equipes/analyse-synthese/peeters/ARTICLES/Peeters_2003_cuidadoaudiofeatures.pdf (accessed September 17, 2020).
- [75] LulzBot, TAZ-6_spec-sheet.pdf, (n.d.). https://www.lulzbot.com/sites/default/files/TAZ-6_spec-sheet.pdf.
- [76] FiFine, FiFine K669B Microphone, (2020). <https://fifinemicrophone.com/collections/best-seller/products/usb-microphone-with-volume-control-k669-669b>.
- [77] N.G. Tanikella, B. Wittbrodt, J.M. Pearce, Tensile strength of commercial polymer materials for fused filament fabrication 3D printing, *Addit. Manuf.* 15 (2017) 40–47. <https://doi.org/10.1016/j.addma.2017.03.005>.

7. IDEAS (Identify, Define, Establish, Aggregate, Secure): A cyber-physical framework for securing additive manufacturing systems using physical side-channels

Abstract

The distributed nature of additive manufacturing (AM) and its use in producing final parts in aerospace, automotive, and medical device markets means that cyber-physical security is of increasing concern. Side-channel monitoring has been presented as a means for detecting and mitigating attacks by comparing measured responses to a known-good set of responses. IDEAS (Identify, Define, Establish, Aggregate, Secure) is a framework for the design and implementation of a side-channel monitoring system (SCMS) for AM systems that is able to use physical measurements of side-channel emissions to help mitigate the threat of cyber-physical attacks. IDEAS provides a framework to guide its user through the steps needed to implement such a system including, i) identifying attack vectors ii) defining side-channels, iii) establishing baseline values, iv) aggregating SCMS data, and v) securing the transmission of the data to the SCMS. This framework covers a variety of approaches, methods, and techniques for achieving these steps and seeks to be a useful reference for the design and implementation of a SCMS as well as to provide a reference on the challenges and requirements related to cyber-physical security for AM systems and SCMS.

Keywords: Additive manufacturing, side-channel, cyber-physical security, quality monitoring

7.1. The need for side-channel monitoring to secure additive manufacturing systems

Recent cyber-physical security research has shown that additive manufacturing (AM) systems are increasingly at risk of cyber-physical attacks [1,2]. When connecting physical machinery to digital devices, an attack surface is opened for cyber-physical attacks to jump from the digital domain into the physical one. Examples such as Stuxnet [3,4], the German foundry attack [5,6], and the Ukrainian electrical grid [7] have demonstrated the increasing risk of attacks on cyber-physical systems causing real world damage. In the scope of additive manufacturing, this can mean the risk of machine damage, lost production time, increased waste, and the production of defective parts. What makes AM unique compared to traditional manufacturing, is the layer-based process and the selective control of both the shape and the material properties. From part design to fabrication, monitoring, and quality testing, AM systems follow a digital design thread. This thread allows AM systems to be distributed and to quickly fabricate new and complex designs through sharing of digital design files; however, it presents a large attack surface with vectors throughout the process [1].

One of the key strengths of AM systems is the ability to fabricate intricate geometries that traditional manufacturing approaches cannot. These complex parts can prove difficult to validate using traditional nondestructive testing approaches due to internal surfaces being difficult or impossible to access for measurement (calipers, coordinate-measuring machine (CMM), light scanning, etc.).

Completely enclosed defects (voids) – made possible via the layer-by-layer fabrication process - require penetrating techniques (such as x-ray CT) to detect. Localized material property changes – made possible by the process parameters’ (e.g., laser power, scan speed) effects on part microstructure and mechanical properties - may not be detectable at all without destructive testing. Because of the complex interactions between the feed material, machine settings, part geometry, and toolpath two otherwise identical parts may exhibit different final properties.

Prior work has primarily focused on looking at this from a cyber-first perspective. Examples of this include works that have focused on attack vectors such as the model file [9–12], the toolpath [13], and the process parameters [2,14–17]. Other works have taken a more high-level view of the AM process and attempted to identify the full scope of vulnerabilities along the AM process chain [2,18,19]. In particular, Yampolskiv does an excellent job of summarizing work in this space and in demonstrating a taxonomy for metal PBF systems [2,17]. While this works well for identifying cyber based mitigation solutions such as encryption and signing, it is less useful when developing SCMSs to protect against sabotage attacks as these are more focused on the physical characteristics of the system. In addition to following a cyber-based approach to identify vulnerabilities, it is important for manufacturers to perform a physical assessment of their system and the specific needs.

As the adoption of additive manufacturing processes to fabricate end-use products continues to expand, the need for strong qualification and traceability of quality parts becomes increasingly important. However, simply implementing quality control (QC) measures is insufficient to ensure security, as attacks can also be designed to deceive or circumvent existing quality control (QC) techniques [8] causing the manufacturer to unknowingly deliver defective components (e.g. components with internal voids or reduced mechanical properties) to the end user. To ensure quality and security, manufacturers need to have a “root of trust” somewhere in their process chain that can be used to validate other components of the system.

As a result of the digitally controlled process-property interactions and layer-wise fabrication in AM, there exists a need for in situ monitoring approaches that can detect defects as they occur, rather than after they have become enclosed inside of the part. One approach for this is to directly integrate sensors into the machine and use them to implement some level of closed loop control. While this approach can work well for naturally occurring defects, it is not sufficient to protect against malicious attacks designed to introduce intentional defects. Built-in monitoring systems may either not be able to identify a malicious change as a defect, or may even be compromised due to their integration with the systems’ programmable logic controller (PLC) and used as the vector for causing the defect in the first place. To mitigate these risks, it is desirable to have a monitoring system that is difficult for an attacker to compromise, and ideally disconnected (air-gapped) from the AM system to add redundancy. This air-gapped monitoring system can then serve as the “root of trust” that validates parts, machines, and other pieces of the supply chain.

One approach that meets these criteria is the use of side-channel monitoring. An emerging approach for defending against cyber-physical attacks in digitally connected systems, side-channels are indirect, physical emissions from a system that occur as a result of the desired behavior of the system.

[20]. These physical outputs are often difficult to predict, but directly correlate to the operating parameters and quality of the part. In a cybersecurity context side-channels have traditionally been seen as unintended physical leakages of information (such as thermal, electrical, or magnetic emissions) that can be used by an attacker to compromise the security of a device or system. In the context of manufacturing, this paper will use the following extended definition of side-channels: “any physical properties of the part or AM system that can be monitored without requiring a digital connection to the AM system hardware or software”. This extended definition includes the previously mentioned thermal, acoustic, electrical, and magnetic signatures, but also includes more direct channels such as the use of encoders (position) or cameras (visual) to validate the correct operation and properties of builds and parts. While these side-channels are still a potential vector for information leakage, they can also be used to validate the integrity of a part or process. By monitoring these side-channels, it is possible to validate a build indirectly and in a way that is difficult for an attack to compromise.

The basic approach, shown in Figure 1, is that one or more known good parts are fabricated while being monitored by the SCMS. These parts are extensively tested, destructively if necessary, to ensure that they are defect free. The in situ quality information collected by the SCMS is then condensed and stored in a data package which is inserted into the toolpath/model file for the part. When the part is fabricated, this data package causes physical emissions to be generated, which are received by the air-gapped SCMS. The SCMS uses the information contained in the data package to compare against the side-channel data collected during the fabrication process. If the detected signatures match those from the known known good, the part is accepted. If the part or build have been attacked or tampered with, the inherent build signatures will vary from the known-good signatures, and the SCMS will alert the operators that a defect has been detected. A more detailed explanation of this approach, along with a physical embodiment on a material extrusion system, can be found in the authors’ work on cyber-physical hashing [21].

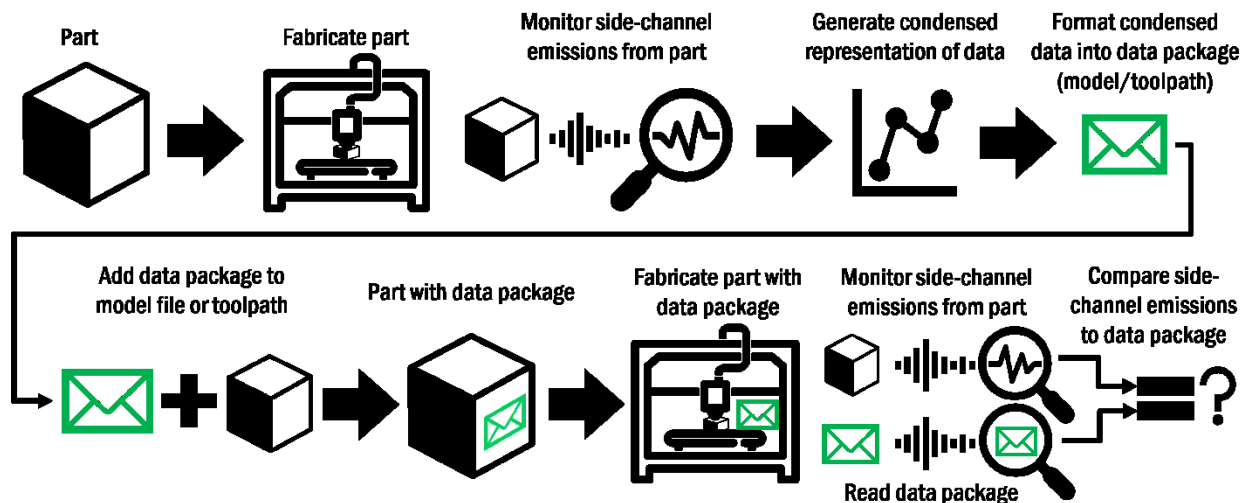


Figure 7.1. Cyber-physical hash overview: Using an air-gapped SCMS to validate AM parts by incorporating a data package in the toolpath or model file to send information using physical emissions during the fabrication process [22].

AM is used in a wide range of domains which have very different requirements for part fabrication and quality, as shown in Figure 7.2. On one extreme is large scale manufacturing, where a high volume of parts with fixed settings are fabricated. In this environment, the SCMS approach presented in Figure 1. Is straightforward to apply. Initial builds are used to establish the baselines that future builds are compared against. In other settings, the task of validation is complicated by a large number of unique parts that cannot be directly compared to each other. In other settings, such as where closed-loop in situ adjustments are made during fabrication, the changes in process parameters between builds (or even within builds) may again prevent direct comparison between parts. Further discussion of these use cases can be found in section 7.4, but it is clear that some type of framework is needed to guide the design and implementation of SCMSs.

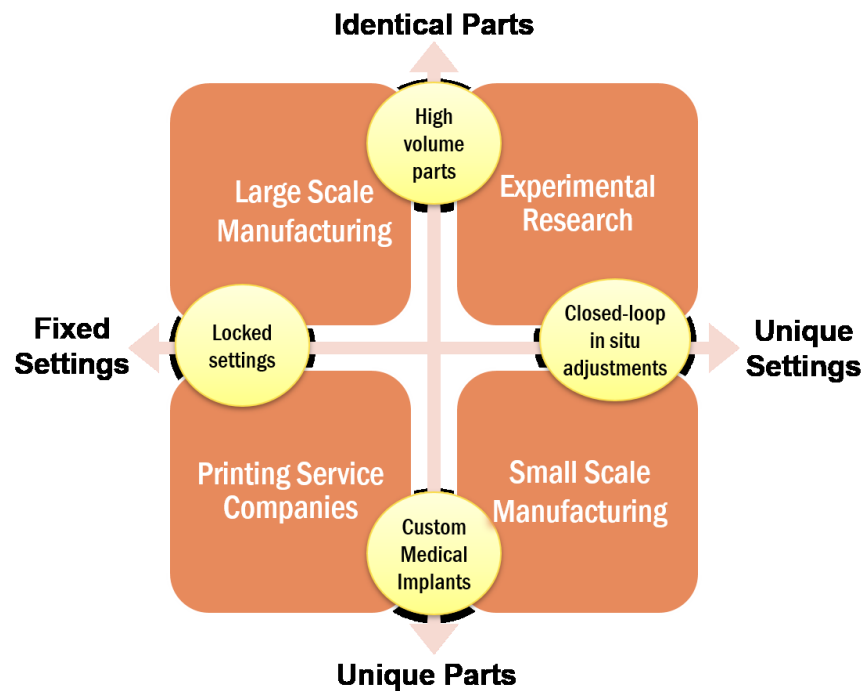


Figure 7.2. Different domains of additive manufacturing based on the part geometries and process parameters used. Fixed geometries and settings are easier to establish baselines for while unique parts with in-situ adjustments are much harder to validate.

While there do exist some cyber-physical security frameworks for manufacturing, such as the NIST framework [23], these are primarily high-level approaches with a strong focus on the cyber components of the system that do not cover the use of side-channels. Yampolskiy’s framework provides a good overview of AM specific considerations and possible attacks, but stops short of making specific recommendations for how to implement systems to mitigate attacks [2]. In this paper, the authors present a framework for identifying and implementing side-channels in AM systems as a way of mitigating cyber-physical attacks and ensure build quality. This new framework seeks to supplement these existing frameworks by providing an AM-specific approach for developing a side-channel monitoring system (SCMS) to detect and mitigate part sabotage attacks, using physical, rather than

cyber monitoring as the core element. The proposed IDEAS (Identify, Define, Establish, Aggregate, Secure) framework (Figure 2) consists of five main steps:

- 1) **Identify** attack vectors
- 2) **Define** side-channels
- 3) **Establish** baselines
- 4) **Aggregate** data
- 5) **Secure** transmission

While the NIST framework provides a broad overview and foundations for security of cyber-physical systems, it lacks specificity on the details of how to implement specific mitigation techniques. This new framework seeks to provide specific guidance in how to design SCMS to mitigate part sabotage attacks on AM systems. The work complements the existing NIST framework by presenting detailed techniques that fall under the broad NIST framework, while not replicating the existing work. In the context of the NIST framework's core functions (Identify, Protect, Detect, Respond, Recover), the first step of the IDEAS framework corresponds to the "identify" function, steps two through four correspond to the "detect" function, and step five correspond to the "protect" function. The "response" and "recover" functions are outside the scope of the design of a SCMS and depend strongly on the manufacture's set up and infrastructure. Figure 3 shows an overview of the IDEAS framework along with some of the core components. The orange column shows the starting point for that step of the process, yellow shows intermediate tasks that need to be completed and the grey column highlights key outputs that should be accomplish in that step.

Due to limited resources in any manufacturing environment, it is unnecessary and impractical to implement side-channels to monitor every possible step and parameter. The goal of this framework is therefore to help manufacturers select a minimal set of side-channels to provide security against cyber-physical attacks and to ensure quality while adding the minimal amount of cost.

In Section 2 a detailed overview of how to identify the most critical attack vectors will be presented. These attack vectors will be used to drive the selection of side-channels using the considerations given in Section 3. With the side-channels established, the Section 4 discusses several approaches for establishing based on the needs of the system. Section 5 covers several approaches for aggregating the data to reduce computation, storage, and transmission requirements. Finally, several approaches for securing the transmission of quality information to the SCMS are provided in Section 6.

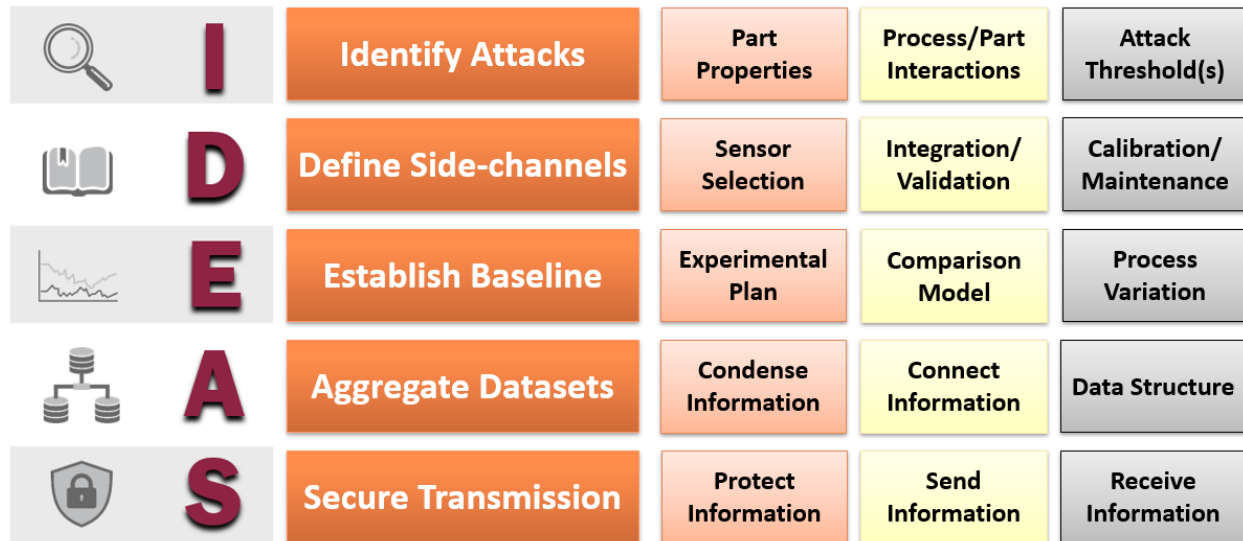


Figure 7.3. Overview of IDEAS framework: The orange column contains the starting objectives/needs for each step, the yellow highlights intermediate tasks, and the grey gives examples of expected outcomes during the step.

7.2. Identifying Attack Vectors

In order to effectively secure an AM system using a SCMS, one must first identify the most critical attack vectors. These can then be used to inform the design of the monitoring system. There are two ways of approaching this task:

- *Cyber-physical approach:* The first, is to define the attack surface by looking at the cyber components of the system, how they are exposed to attacks, and how those attacks could affect the process and parts [9,24–29]. This is the traditional cybersecurity approach to identifying attack vectors and has been used in numerous works addressing advanced manufacturing security.
- *Physical-cyber approach:* The second approach is to identify the key desired physical specifications of the final part, the allowable variation in these properties, how changes to the process affect these properties, and the allowable deviation in the process. This second approach frames the problem as physical-cyber issue rather than a cyber-physical one, by putting the emphasis on the use of physical controls and sensors to mitigate sabotage attacks on parts. The focus on the physical properties in this approach allows the system to be designed to perform both security and quality monitoring functions.

The physical-cyber approach to identifying attack vectors has the following six steps:

1. Identify critical part properties (e.g., tensile strength, surface finish, color, geometry, etc.)
2. Establish allowable tolerance for properties
3. Identify high level process overview (mass and energy inputs)
4. Identify process/machine specific implementation and correlation to geometry (extrinsic) and parameter (intrinsic) effects on part properties

5. Quantify relationships between process settings and part properties to determine maximum allowable variation
6. Define attack detection threshold(s)

7.2.1. Identifying critical part properties

The first step is to identify the critical parameters for the final part and to establish suitable criteria for validating them. It is important to understand the product needs and the implications of the AM process on those needs. For example, additive manufacture of prototype models may not require precise validation of material microstructure, but do require precise geometric tolerances. Medical models may require accurate color matching. Functional parts may require tight dimensional tolerances in some areas, but not in others, and will require accurate material property performance. As another example, in an AM part that is post-process machined, the surface finish of the as-built part may not be important. However, poor surface finish may have a negative effect on the machinability of the part and have a deleterious effect on the surface integrity of the resulting part [30].

During this process, properties that are essential to the part function need to be distinguished from those that are secondary. For example, an aerospace bracket might be required to support a certain load, which is affected by both the geometry and the tensile strength of the material. In this example, the critical property is not the tensile strength of the material, but the load bearing ability of the bracket. The load bearing functionality is a result of the interaction between the tensile strength and the part geometry. These critical properties should be measurable and have testing methods in place to validate parts. Clearly defining and understanding these requirements is essential in identifying attack vectors and designing a SCMS.

7.2.2. Establishing tolerances for part properties

Once critical part functionality/properties have been identified, an allowable tolerance range for each property should be established to aid in determining detection thresholds. When establishing tolerances for AM parts, designers should be aware that there may be challenges in transferring tolerances from existing manufacturing processes [31,32]. Work has been done to translate existing geometric dimensioning and tolerancing (GD&T) practices into a more AM specific context [33] and a variety of research have looked into how to apply geometric tolerances and verification to AM parts [34–38]. In addition to geometric tolerances, additively manufactured parts also need to set requirements for the properties of the material, since the material properties are defined in the processing of the part. Two parts may have identical geometries, but significantly different material properties due to differences in processing parameters. Material properties may further vary not just between different parts, but within the same part due to localized effects of geometry; e.g. a thin fin may experience a different thermal profile than a thick solid body and layers may be affected by prior and subsequent layer geometries and properties. Work has also been done in this area to support qualification and certification based on models and testing of material properties [39,40]

7.2.3. Identify high-level process overview (mass and energy inputs)

While often referred to as a group, there are seven modalities of AM processes, each with their own methods of creating a part in a layer-wise fashion [41]. Once the required part properties and

tolerances are determined, the manufacturer needs to identify how the AM system can be manipulated to affect these properties. Before investigating detailed machine specifics, the manufacturer should start with a high-level overview of how material and energy are fed into the system to create the part. As an example, Table 1 shows how raw material and energy are supplied and patterned in the seven different AM categories system to fabricate the part. “Voxel” indicates that material/energy is patterned on a point-by-point basis, “Voxel/vector” represents a continuous road of material being placed, “Pass” indicates that material/energy is applied to some finite width subsection of the build along an axis within a layer, and “Layer” indicates that the material/energy is broadly applied to the entire layer without the ability for selective modification of individual sections. As the granularity of the patterning decreases (i.e., voxel -> vector -> pass -> layer), the greater the difficulty for an attacker to hide a change, since it is more difficult to hide large changes in a process/part. For example, the removal of an entire layer would cause a much larger and more obviously detectable change to a part than the removal of a single voxel of material, particularly if the voxel was located inside of the part. Details of each AM process are given in the following subsections (2.3.1-2.3.7).

Table 7.1. Overview of AM process types and the physical material and energy inputs into the system along with their dimensionality (i.e. voxel, pass, layer, or volume).

AM Process	Material Patterning	Energy Patterning	Secondary Material	Secondary Energy
Vat Polymerization	(Resin) Layer	(UV Irradiation) Voxel	-	Heating
Material Jetting	(Resin) Voxel	(UV Irradiation) Pass	-	-
Binder Jetting	(Powder) Layer	(IR Irradiation) Pass/Layer	Voxel (Binder)	Heating
Material Extrusion	(Thermoplastic/Paste) Voxel	(Heat) Linked	-	Bed/Volume Heating
Powder Bed Fusion	(Powder) Layer	(Laser/Electrons) Voxel	Inert Gas	Surface/Volume Heating
Sheet Lamination	(Sheet) Pass	(Vibration) Pass	-	-
Directed Energy Deposition	(Powder/Wire) Voxel	(Laser/Arc) Linked	Inert Gas	Surface/Volume Heating

7.2.3.1. Vat photopolymerization (VP)

In VP, the material is recoated over each layer (or continuously in the case of continuous liquid interface printing) in a uniform distribution (i.e., without selective application of resin) for the next layer. The energy is patterned either through the use of one or more lasers or through a projected mask (LCD, DMD, etc.). This allows the material properties of the part to be altered on a voxel-by-voxel basis through the over- or underexposure of the material by the energy source.

7.2.3.2. Material Jetting (MJ)

In MJ, the material is patterned at a voxel level through the control of small jetting nozzles in conjunction with the movements of the print head. In many MJ systems, a key capability is the ability to control not just the location of the material, but also the type of material at each location by using two or more sets of jetting heads with different materials. By selectively patterning these materials at the voxel level, the macroscale properties of the part can be adjusted. The amount energy (typically a broad-spectrum UV light source) provided to the system can be adjusted by increasing/decreasing the intensity of the light source and by increasing the amount of time the light is over a particular location, either through changing pass speed or through changing the number of passes.

7.2.3.3. Binder Jetting (BJ)

BJ is somewhat unique in that it is by nature a two-material process. The part's powder material is recoated over each layer, by using a roller or other mechanism to spread the powder. The secondary, liquid binding material is patterned at a voxel level in the same manner as in material jetting. The interaction between the binder and the powder bed determines the material composition of the printed part. The energy of the system is added in a layer basis through the use of surface heaters. These heaters are used to speed up the evaporation of the solvent in the binder and facilitate quicker drying.

7.2.3.4. Material Extrusion (ME)

In ME, the material is patterned at a pseudo-voxel level, through the use of material patterned in streamlines. In ME the location, direction, and length of extruded material can affect final part properties [42,43]. changing the input energy applied to the extruded material (through an embedded heater in the extrusion nozzle) can affect the material flow rate and the interlayer entanglement and adhesion between the printed layers, selectively throughout the part. Secondary energy sources in the system can provided surface heating (heated bed) or volumetric heating (heated build volume).

7.2.3.5. Powder Bed Fusion (PBF)

In PBF, the material is recoated each layer by using a roller or other mechanism to spread to powder. The energy is patterned through the use of laser(s) or electron beam to selectively melt the material at the voxel level. The energy input can be adjusted by changing the input power, scanning speed, scan count, and scan spacing. Secondary sources of energy include surface and volumetric heating that can affect the part more broadly, such as layer warping in polymer PBF.

7.2.3.6. Sheet Lamination (SL)

In sheet lamination (SL) the material is patterned out in a pass-by-pass or layer-by-layer approach, where a thin strip of material is rolled out (potentially cut to shape) in selected areas. Energy can then be applied to this material to fuse to previous layers (e.g., via an ultrasonic horn) in a pass-by-pass fashion.

7.2.3.7. Directed Energy Deposition (DED)

In DED the material is patterned at a voxel level through the deposition of molten wire or powder. The input energy is linked to the material being placed on the part. If the material feed rate is increased, the energy input must increase proportionally. By altering the material feed rate or energy being input, it is possible to change the material properties of the part.

7.2.3.8. Comparing Process Similarities

When designing the SCMS it is useful to keep in mind the high-level similarity between some process types. As an example, while vat polymerization and powder bed fusion are two very different AM technologies, they share many similarities at a high level. Both VP and PBF have material that is repatterned in bulk on a layer-by-layer basis and can be attacked by affecting the recoating mechanism (e.g. changing the layer height or changing the roller/wiper settings to result in poor/uneven recoating). Both of these systems also rely on the precise patterning of energy both to generate the geometry and to control the final material properties[44–52]. Attacking the energy source intensity or movement is the primary vector through which part geometry and properties can be affected and therefore the most critical to monitor with side-channels. By considering these high-level similarities it is possible to develop and architecture for the SCMS that is flexible enough to be shared between different process types. While the specific sensors and settings may vary by approach, the techniques used for processing the side-channel data can remain the same

7.2.4. Identify process/machine specific implementation and correlation to geometry (extrinsic) and parameter (intrinsic) effects on part properties

After identifying the high-level inputs of mass and energy into the system, the next step is to identify how these specific inputs, generally manipulated through process parameters, can be modified/attacked by a malicious actor. For example, a laser source could have its power, speed, pattern, hatch spacing, focal distance, or number of scanning passes altered in ways that would affect the part properties. In this step, the manufacturer should list the process parameters of their AM system and identify their effects on final part properties. One common approach for accomplishing this is to use an Ishikawa (fishbone) diagram to identify cause and effect. Figure 4 and 5 show examples of Ishikawa diagrams for ME and metal PBF systems, respectively. The first figure shows the impact of process parameters on part characteristics for ME systems and illustrates that some properties are dependent on more factors than others causing these properties to have a larger attack surface. Figure 5 is a more detailed example that includes not just process properties, but inputs along the AM process chain such as part preparation, system properties, and material inputs. These inputs may be important from a quality monitoring standpoint, a security standpoint, or both.

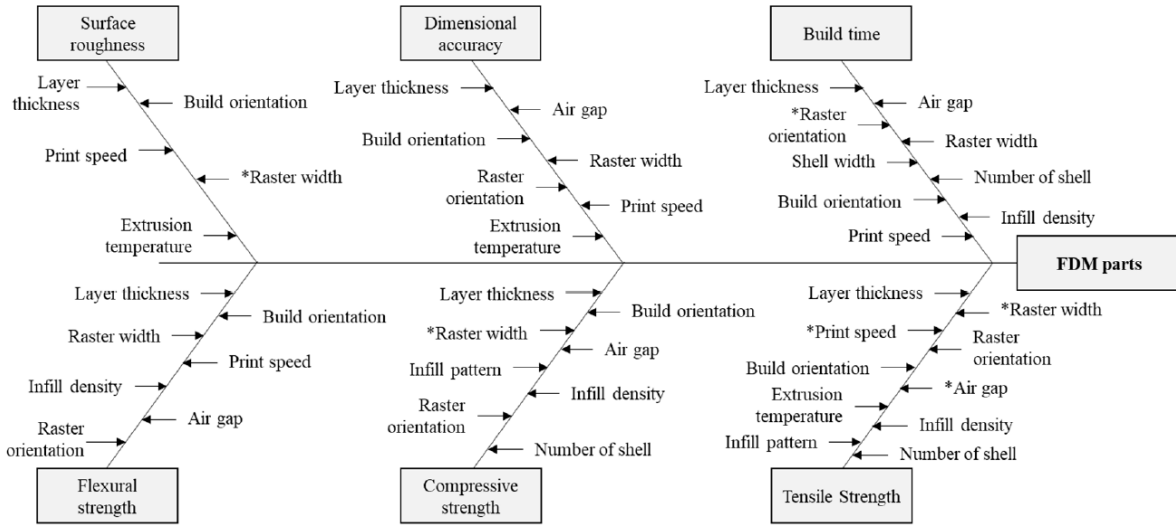


Figure 7.4. Fishbone diagram of the FDM process and how different process parameters affect various part properties [53].

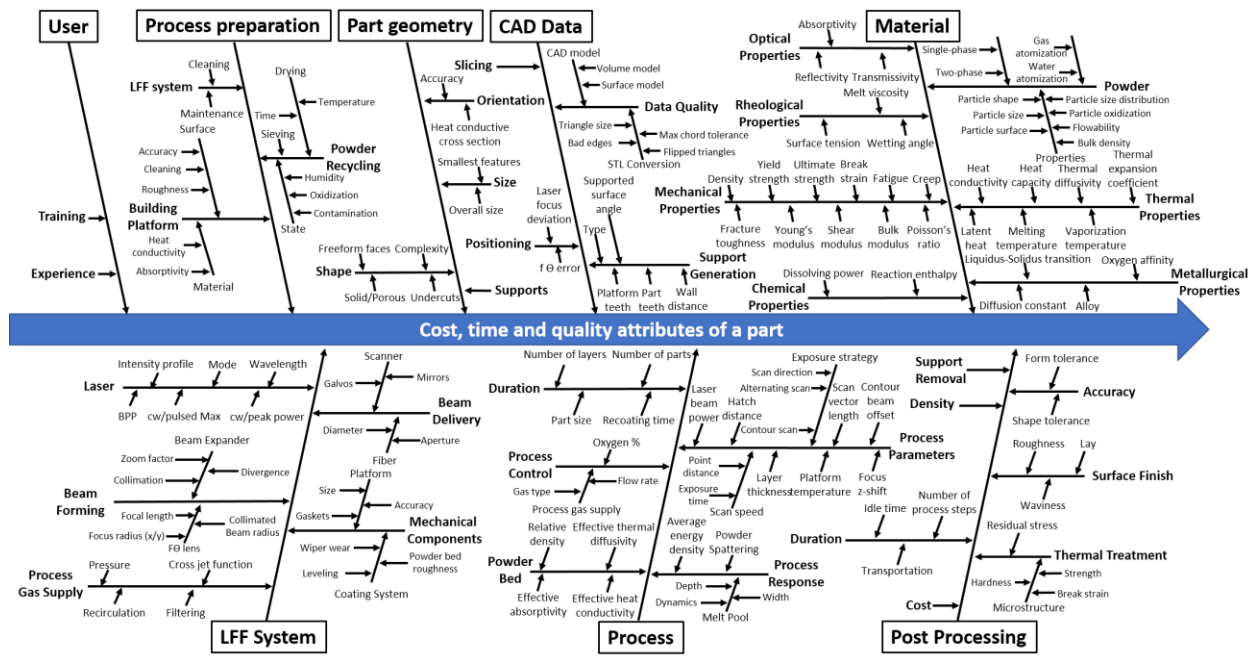


Figure 7.5. Fishbone diagram of a laser powder bed fusion system.[54]

7.2.5. Quantify relationships between process settings and part properties to determine maximum allowable variation

Once the relevant interactions between the process and the part properties have been identified, the next step is to quantify how these inputs physically affect the part properties (e.g., recoating a thinner layer in VP will reduce the part cross sectional geometry in the z-direction). The degree by which a change in process parameters has been altered the part properties/functionality is the magnitude of the change. A good example of how the magnitude of change to process parameters

affects part properties in a PBF system can be seen in work by Khorasani et al. [55]. The process parameters that were studied were laser power (90-110 W), scan speed (600-800mm/s), hatch spacing (65-85 μ m), scan pattern angle (36-72 $^{\circ}$), and post process heat treatment (20-1050C). The material properties that were investigated were hardness (HB), relative density (%), tensile strength (MPa), elongation at break (%), and surface roughness (μ m). Khorasani et al. found that hardness was most strongly affected by post process heat treatment, with the rest of the processing parameters having a relatively minor impact, relative density mostly correlates with the trends seen for hardness, with harder samples usually having a greater density. Elongation at break and average surface roughness were both strongly affected by heat treatment, while other process parameters had a small affect within the tested ranges. Tensile strength was primarily driven by laser power with other process parameters having a much smaller (scan speed, scan pattern) or nearly negligible (hatch space, heat treatment) effect within the tested ranges [55].

This demonstrates the need to define the critical part properties and their interaction with the process parameters when establishing attack vectors. If tensile strength was the critical part property, a SCMS that monitors the laser power would be preferable to one that monitored part density. Based on this data, a manufacturer whose critical part property was tensile strength might require their laser power to be maintained at 95W (+/- 1W), scan speed at 725mm (+/- 75mm/s), hatch speed at 75 μ m (+/- 10 μ m), and scan angle at 40 $^{\circ}$ (+/- 5 $^{\circ}$). These would provide a threshold at which variation would cause a significant effect on the critical part properties and would need to be detected by the SCMS.

Another consideration for manufacturers when determining maximum allowable variation is the interaction between multiple factors. While any one of these factors might be able to vary within the prescribed range, an attack could also alter multiple process parameters within their acceptable ranges to cause a combined effect that would exceed the allowable part property tolerance. Surface plots such as those given by Khorasani et al. can be used investigate these interactions and to further refine the attack detection thresholds. An entire separate domain of research, there have been significant studies that investigate the process-property relationships of AM processes [55–61] .

7.2.6. Define attack detection threshold(s)

Once the impact of process parameters on part properties has been established, an attack detection threshold can be defined. The detection threshold is the smallest change in the process parameters or part properties that should be detected by the SCMS. In order to do this, it is necessary to understand the effect that the process parameters have on the part properties. The primary aspects that affect the detection threshold are:

- How large is the change in the process? (significance)
- What is the area of the part that is affected? (volume)
- Where does the change occur? (location)
- How often is the change repeated? (number)
- How long is the change in effect? (duration)
- How long does it take for the change to occur? (timescale)

There are two sets of severity classifications that need to be considered 1) how will the change to the process parameter affect the part properties and 2) how will the change affect the detectability of the attack.

7.2.6.1. Significance of change

As discussed in Section 2.5, the significance of the change is the amount at which the change in the process will affect the final part properties. A setting that requires a large change to have a significant effect on the part properties is considered low risk, while a setting that requires little change to have a significant affect the part properties is considered high risk.

7.2.6.2. Volume of change

Research studies investigating the effects of process parameters on part properties typically apply a fixed magnitude change to the parameter across the entire volume of the part; however, in the case of a malicious attack, it is possible that an attacker might only alter a select portion of the part. An example of this is lack of extrusion in ME process. A naturally occurring defect might be a clogged nozzle that would result in no material extrusion after the clog. In contrast to this, an attack might turn off the extruder for a short amount of time in order to cause an internal void. This is a large change in extrusion over a small area (the void). While small voids occur naturally in the ME process between roads, a large void can have a significant effect on the part strength. Similarly, in PBF an affect to material properties of a single layer may not be significant enough to require detection. A maximum allowable size for defects should be defined and used in setting the detection thresholds. A small area affected is considered low risk, while a large area affected is considered a high risk.

7.2.6.3. Location of change

Unlike subtractive manufacturing, which primarily affects the surface of a part, AM is able to affect the geometry and the material properties volumetrically throughout a part. Some of these areas may be critical to the performance of the part (e.g., the gauge section of a tensile test specimen), while other locations may have little to no effect on the final properties (e.g., the gripped section of a tensile test specimen). This means that the detection threshold might be smaller in some areas and larger in others, which may be able to reduce the requirements of the SCMS. A defect occurring in a non-critical location is considered low risk, while a defect that occurs in a critical location is considered high risk.

7.2.6.4. Number of changes

While a single defect, such as a small void, may have an insignificant effect on the part properties and thus fall within the allowable area of change, numerous small voids in aggregate could have a significant effect on the part properties. This occurs most commonly as porosity in AM parts where it affects the part strength [62,63] , but it is also possible for an attacker to insert numerous small defects in an attempt to avoid detection. Small, but repeatedly delayed/modified commands or small power fluctuations could cause many small defects that might not normally trigger an alarm. When identifying attack vectors the manufacture should consider the allowable quantity of small defects. A small number of changes is considered low risk, while a large number is considered high risk.

7.2.6.5. Duration of change

The duration for which a change occurs can have a significant effect on the final properties of a part. A brief fluctuation in laser power might result in a small localized defect, while a sustained change in the power could affect a much larger area of the part. Other factors such as the environmental temperature might not significantly affect the part properties over a short duration, but long-term changes could affect properties such as crystallinity (polymers) or grain growth (metals). When thermal or material factors are involved the duration of which an attack occurs can have a significant effect on the final material properties. This has been demonstrated in studies showing that the dwell time between layers affects the material properties and the thermal behavior of the system [64–66]. How long the process can remain outside of normal operating ranges is an important consideration when designing the SCMS. A short duration is considered low risk, while a large duration is considered high risk.

7.2.6.6. Timescale of change

Timescale represents how fast a process setting can be changed by an attacker. Some process settings react slowly due to physical limitations, such as thermal mass. If the change occurs over a long timescale (such as the volumetric heating of a build volume), it will not require a particularly fast monitoring system to detect the change. However, if the change can occur very quickly (such as changes to laser power or speed) the monitoring system needs to be able to sample and detect changes at a much faster rate. The timescale and duration are closely related considerations when designing the SCMS. A long timescale is considered low risk, while a short timescale is considered high risk.

7.2.6.7. Severity threshold

The goal of the severity threshold is to allow the manufacturer to evaluate the risks of various potential attacks in order to be able to select a SCMS that is most effectively able to mitigate them. Two illustrative example attacks are presented here. Table 2 shows a summary of the evaluated severity of the example attacks.

- **ME Void Attack:** Consider an attack on an ASTM tensile test specimen as demonstrated in the authors' previous work and shown in Figure 6 [9]. A void can be considered an attack that reduces the tensile strength of the affected area to zero; this has a high significance risk. The size of the void (relative to the part) determines the volume risk; a 3.3 mm³ void might be considered a medium risk. A void located in the gauge section of the test specimen would be a high risk' a void located in the grip section would be a low risk. In this case, the defect is located in the neck section so it is considered a medium risk. There is only a single void, so the quantity of the defect is low. Due to the size of the void, there is a moderate amount of time when the system is not depositing material which puts it at a medium risk, if the void was caused by altering the toolpath this could cause a larger effect on the part that could indicate a larger risk. Similarly, the timescale of the change is moderate and so the corresponding risk value is medium. As seen in Table 2, the effect on the final part properties is moderate to high. The difficulty in detecting the attack is low to moderate due to the significant change to the process settings and the moderate size of the defect. This means that the monitoring system would not require extremely fine resolution to detect this attack.




Figure 7.6. Cross sectional slice of a dogbone infected with a void.

- **LPBF Laser Power Attack:** Consider an attack in which a bad actor attack would be to alters the laser power by a moderate amount to affect the material properties. This would be a medium significance risk. If it was applied over the entire part it would be a high size risk. The location would include the gauge section so this would also be considered a high location risk. Since the defect affects the entire part the quantity criteria is not applicable as it only applies to defects with distinct instances. The duration risk is high since the change applies to the entire part. The timescale risk is low since the change is occurring over a long period of time (leading to greater chance of detection). As seen in Table 7.2, the effect on the final part properties is high, while the difficulty of detect is low. If the change in laser power was instead small repeated pulses instead of a constant effect, the significance and location risks would remain the same, the size risk would become low, the quantity risk would be high, the duration risk would be low and the timescale risk would be high. The effect on part properties would be reduced, but the attack would become significantly more difficult to detect due to the short timescale during which the change was occurring and the small size of the defects.

This approach can be used to identify the attack vectors that pose the greatest security risks and the process settings that are most critical to monitor. This informs the selection of the side-channels In addition, using quantitative data on the relationship between process parameters and part properties can be used to estimate the required minimum resolution for the SCMS. Once the significance of the attack vectors has been established (either through literature review or experimentation), the next step is to move into defining side-channels for the system that can monitor these vectors.

Table 7.2 Example attacks and their corresponding risk evaluations.

<i>Example Attacks</i>	<i>Significance</i>	<i>Size</i>	<i>Location</i>	<i>Quantity</i>	<i>Duration</i>	<i>Timescale</i>
<i>ME Void</i>	High	Medium	Medium	Low	Medium	Medium
<i>LPBF Laser Power (Entire build)</i>	High	High	High	NA	High	Low
<i>LPBF Laser Power (Pulses)</i>	High	Low	High	High	Medium	High

7.3. Defining Side-Channels

Once the attack vectors have been identified and the severity has been considered, the next step is to define and select side-channels that will be able to monitor the process parameters most critical to part performance and detect attacks causing an unacceptable reduction in part functionality. To define appropriate side-channels, it is important to identify the primary purpose of the monitoring system. The three use cases of a monitoring system are:

- 1) To detect changes to process parameters due to inherent process noise and variation
- 2) To detect intentional changes to parameters caused by malicious cyber-based attacks
- 3) To simultaneously be able to detect the changes in both previous cases (inherent variation and malicious attack)

While there is some overlap between these roles, there are important distinctions in how they need to operate in order to accomplish each roll. A system may perform one or both of these roles depending on how it is designed.

In the first case, it can be assumed that the input part file matches the designer's original intent. The monitoring system needs to be able to check for common printing defects such as a clogged nozzle, warping, peeling, keyholing, or any number of standard AM process failure modes. The side-channel(s) needs to have sufficient resolution to detect defects that occur in the process and can be sampled at set intervals to check that the process is behaving as expected. This system may need to monitor physical/mechanical parameters that do not have a cyber input (e.g., foreign matter in the powder bed or a clogged nozzle) [67–72]. While a system designed solely for quality monitoring purposes may be able to detect a cyber-physical attack, it should not be relied upon to do so as a sophisticated attacker might be able to exploit gaps in the operation of the monitoring system. The original monitoring system was only designed to check for expected features and is not robust enough to detect defects of an unexpected type.

In the second case, the SCMS only needs to validate that the AM machine is operating as expected and that it has not been maliciously tampered with in some way. This approach needs to validate that the input files have not been modified (i.e. STL, toolpath, calibrated/set parameter values) and that digital signals on the machine are not tampered with during printing. An example of this is the stepper motor driving filament extrusion on a ME printer. An attacker might change the extrusion rate by sending a modified signal to the stepper motor or by delaying the signal that is being sent. Monitoring this signal input would detect any malicious changes that were made; however, it would not detect if the extrusion stopped or slowed for another reason such as a clogged nozzle or a reduced filament diameter. A side-channel system designed to detect cyber-attacks only needs to monitor that the digital signals being sent and received by the system are operating as expected, it does not necessarily need to be able to monitor that the fabrication of the part occurs successfully, since natural process error may cause failures that are not detected.

Side-channel monitoring systems designed for security need to consider failure modes that may not be expected to occur during normal operation or small defects intentionally inserted at regular intervals between normal quality control sampling. The system may need to implement a continuous or

random sampling interval to ensure that an attacker is unable to insert defects between normal sampling intervals. This type of system works well in a setting where there already exists a quality monitoring system (such as a closed loop control system) and the side-channel monitoring only needs to ensure that the printer and monitoring system are operating as expected and have not been maliciously modified. By implementing this type of system on top of an existing system it may be able to improve security while incurring less cost than a complete overhaul of the quality system.

The third case is where the SCMS provides both in the quality monitoring and the security. This is the case where a new system is being designed or where an existing system (without any quality control) is being upgraded. In this case the SCMS needs to meet all the quality needs from the second case while at the same time also meeting security requirements from the first (e.g., random sampling). This design needs to be able to validate all relevant part properties, either through inspection of the part itself or through the monitoring of all of the relevant machine parameters during fabrication. The system needs to be able to connect design intent with the monitoring and quality data and requires a strong understanding of both attack vectors and natural process failures. In addition to this the system should have redundancy built-in to avoid being a single point of failure.

7.3.1. Monitoring approaches

Once the use case for the SCMS has been established, the sensor selection process can begin. There are three different approaches to monitoring that can be used:

- 1) Monitoring of the part
- 2) Monitoring of the AM system
- 3) Monitoring of the part/process interactions

These approaches can be used separately or in conjunction with each other depending on the role and requirements of the SCMS.

7.3.1.1. Part monitoring

Monitoring the part of the part is done through the use of in-situ nondestructive evaluation techniques. Examples included in situ X-ray CT [73], ultrasonic [74,75], capacitive [76], impedance-based [77], and vision techniques [78–81]. Part monitoring has the advantage of being able to directly inspect the properties of the final part and do well in the quality monitoring case. These approaches can be capable of detecting defects such as voids, altered geometry, and modified material properties; however, these approaches may be difficult to implement in the AM system, and the part geometry may limit the effectiveness of these approaches. For example, machine vision techniques cannot detect changes to material properties, capacitive methods may have limited depth, ultrasonic methods may miss features hidden behind other features, X-ray CT may have beam hardening and other effects reduce detection resolution. While these methods can effectively detect certain types of defects, they may be costly or difficult to implement with multiple parts or with the restrictions/environment of the AM system.

7.3.1.2. System Monitoring

Directly monitoring the AM system is a good approach when the primary goal of the SCMS is to detect cyber-physical attacks. This includes approaches such as using encoders to monitoring stepper motor positions, thermocouples to monitor extruder temperatures [82,83] , tracking laser position [84–86], monitoring power consumption [87] , electromagnetic emissions [88], acoustic emissions [13,88,89], visual tracking tool position [78,90], and more. Monitoring these system movements can ensure that the system is receiving commands as expected and that the toolpath/parameters have not be altered; however, it is often unable to detect natural physically occurring defects in the part such as lack of adhesion/stringing (ME), defects in the powder bed (PBF, BJ), clogging nozzles (ME, MJ, BJ), and many more process defects. These types of systems can be easier to implement than direct monitoring of parts and work well in supplementing existing in-situ monitoring systems to add security against cyber-physical attacks.

7.3.1.3. Interaction Monitoring

The final type of monitoring is looking at the interactions between the process and the physical part. This includes approaches such as meltpool monitoring [85,86,91–95], tracking acoustic emissions [89,93,96–98], thermal imaging [91,99], and other techniques that capture energy/material being added to the part. What distinguishes these approaches from part monitoring is that they are only monitoring the surface layer (not the whole part volume) and are typically not directly measuring part properties. What distinguishes them from the system monitoring is that they are only inspecting the localized effect on the part and not tracking the overall movements or settings of the system. There can be some overlaps between this approach and the system monitoring approach, such as in the case of co-axial monitoring of the laser (which gives both the laser position information as well as information about the meltpool). However, these approaches can also detect emissions, such as spatter, that part or system monitoring may be unable to observe.

7.3.2. Sensor requirements

The most critical requirement in selecting sensors for the SCMS is that those sensors are able to detect the attack vectors identified in Section 2. If the primary concern is with material property changes due to thermal effects, the SCMS might monitor the power input of system (current/voltage), the temperature of the part/interaction/system (IR camera, pyrometer, thermocouple, etc.), or the part properties directly (impedance-based [77]). To further refine the selection after identifying appropriate sensing modalities, the manufacture needs to consider the following:

- Resolution
- Sample Rate
- Bandwidth
- Robustness
- Integration
- Cost and maintenance

7.3.2.1. Resolution

Resolution can apply to the degree of change (such as smallest resolvable temperature change), the spatial change (smallest detectable area), and the time of change (shortest time of change that can be detected). For parameter that change slowly or evenly, such as volumetric temperature, the resolution can be quite low both spatially and temporally, and still be able to detect changes of a few degrees. For parameters that are both small and change quickly, such as the meltpool temperature, the spatial and temporal resolution must be quite high, while the detectable change in temperature might only need to be accurate to the nearest 10 degrees (due to the proportionally smaller change having less potential impact on part properties).

7.3.2.2. Sample Rate

The sample rate is closely related to the temporal resolution. The temporal resolution is how fast the physical sensor reacts, while the sample rate is how often the data is reported. A high-speed sensor, sampled at a low rate might still detect rapid changes if they occurred at the sampling time, but these changes could also easily be undetected if they occurred between the sampling times. Similarly, if slow speed sensor (e.g. a large thermocouple) could be sampled at a high rate, however this would simply generate unnecessary data that failed to provide any additional useful information. The sample rate should be chosen to ensure that enough resolution from the sensor is maintained, while minimizing the amount of unnecessary data.

7.3.2.3. Bandwidth

The bandwidth is another constraint that is closely linked to the sample rate. Bandwidth is the amount of data that the system is able to sample and process. A thermocouple is a single data point and reacts slowly to changes in temperature, resulting in a small amount of data that needs to be processed, but also a small amount of information that can be transmitted (e.g., 1 bit per second). A thermal camera can stream hundreds to thousands of data points at a much faster rate, resulting in a large amount of data that needs to be processed. This can mean that a much larger amount of data can be sent from the toolpath/AM system to the SCMS, but only if the AM system is able to generate changes at a similar rate. A PBF system could send significantly more information to an IR camera compared to a single thermocouple using a laser to rapidly heat material, but a ME system communicating by changing the nozzle temperature would still be limited to the same rate as the thermocouple since the limiting factor is the AM system and not the SCMS. While higher sample rates can be desirable for detection purposes, it is important to keep in mind the volume of data that can be processed and stored by the SCMS.

7.3.2.4. Robustness

The robustness of the sensor is how sensitive it is to external noise in the environment that it will be used. A microphone in a noise shop floor would likely have low signal-to-noise ratio as external noise would reduce the accuracy of the measurements. A linear encoder in the same environment would not be affected by external noise and would be quite robust, regardless of the surrounding activity. Camera-based systems are often affected by lighting conditions that may reduce their effectiveness, while sensors mounted inside of the AM system are more likely to be isolated from external interference.

When designing the SCMS it is important to consider the possible use cases and select sensors/implementations that will result in as much robustness to noise as possible.

7.3.2.5. Integration

Integration concerns how difficult it is to physically integrate the SCMS with the AM system. In the case of a new AM system design it may be possible to design around the SCMS making integration less of a concern, however in the case of existing systems, difficulty of access and modification are important. A microphone or camera mounted outside of an AM system may not even physically touch the system and requires little to no integration. An encoder mounted to a stepper motor requires physical access inside the machine, potential modification of the system, and the running of new wires out from the system. This is significantly more intrusive and a more difficult integration problem. In addition to physical access, another concern is the impact of the environment on the sensor performance. In high temperature PBF systems it may not be possible to mount many sensors directly inside of the heated build volume and alternative methods (i.e., using infrared thermography through the system's window) will need to be used.

7.3.2.5. Cost and maintenance

There are several costs associated with the implementation of a SCMS. The first is the direct cost of the system and sensors. The second is the cost of the maintenance and upkeep of the system, ensuring that it is calibrated and functioning correctly. When designing and implementing a SCMS it is important to establish a clear maintenance and validation schedule. As the root of trust in the system it is essential to ensure that the SCMS remains functioning within its designed parameters and has not deteriorated or been compromised by a cyber-attack.

7.4. Establish Baselines

Once appropriate side-channels have been selected for the AM system it is important to have a testing method in place to establish what the baseline "good" response is from the SCMS. Determining how to establish this baseline signature will vary based on the use case, as discussed in Section 1. To guide this decision, the author presents a framework for identifying use case, which is guided by two primary axes (Figure 7): the geometry of the part and/or the process settings for the part/build.

- The upper left quadrant (Figure 2) is the domain of traditional large-scale manufacturing. A single geometry is fabricated in large quantities, with fixed process settings. This domain is the simplest in which to identify attacks because parts can be fabricated and destructively tested to make a known good set of quality data which can be directly compared to the monitoring data from future builds to detect any deviations.
- The upper right quadrant (Figure 2) is the research domain where various process parameters are tested on a controlled geometry (such as an ASTM standard tensile test specimen) to evaluate and quantify their effects.
- The lower left quadrant (Figure 2) is where many print shops and service companies operate. They fabricate a wide variety of geometries, by using well-known set of process settings.

- Finally, in the bottom right quadrant (Figure 2) is the small scale, high value manufacturing environment in which the products being fabricated are each unique (e.g., custom medical implants).

When developing an experimental testing plan for the baseline and creating a comparison model it is important to consider the expected types of variation and how a set of test parts can be used to validate future builds. A simple example of this would be an experimental plan that fabricated a singular test specimen and then used this data to validate a tray containing multiple of that same specimen. The part to part comparison would be the same (no additional test parts need to be fabricated) and the comparison model would just need to be able to account for translations movements and multiple instances. As geometric and parameter changes increase, creating an adequate baseline will require a larger set of experimental data and/or a more sophisticated comparison model. The extreme of this is the use of in situ, closed loop control the adjust process parameters on the fly which means that even if the same build was fabricated twice, the resulting side-channel signatures might vary significantly at one or more points.

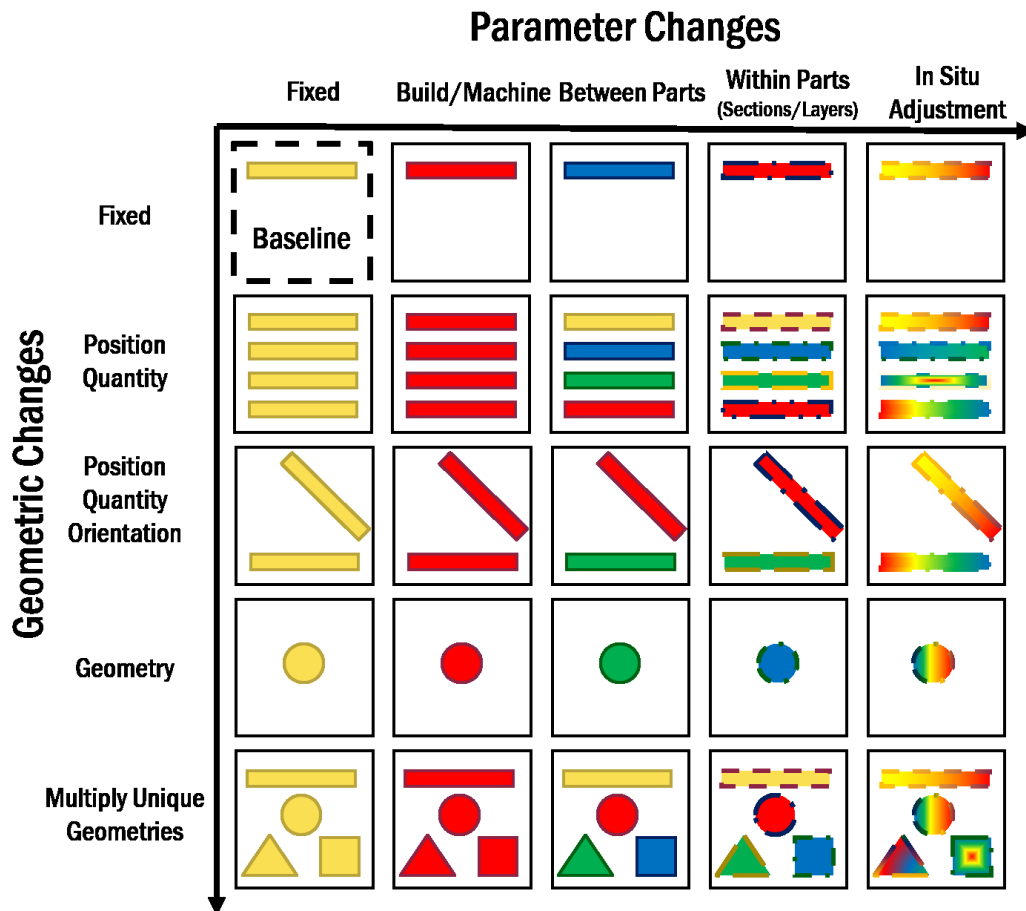


Figure 7.7. Examples of different types of geometric and parameter changes for AM. Establishing a baseline needs to start by comparing fixed settings and geometries and then build up to more complex changes and interactions.

7.4.1. Direct comparison

The simplest and most straightforward approach to establishing a baseline is to use direct comparison. In direct comparison a number of control builds are monitored during fabrication and tested to ensure quality. These builds are then used to establish a baseline value for the side-channels and an expected level of variation and noise. Future builds are compared against this baseline and, if the deviation is outside of the expected range for the side-channel, an alert is given. While this works well for large-scale production with fixed settings it is not adaptable to different parts or large changes in process parameters. Changes away from the fixed baseline, such as adding additional parts or fabricating on a different machine might be able to be accommodated through recalibration, but any large changes will require a new set of control parts to be fabricated. This is the least flexible approach to monitoring.

7.4.2. Machine learning

A variety of machine learning based methods have been demonstrated on AM systems for the purpose of process monitoring and quality control [100]. While machine learning approaches take a large initial dataset to train, once they have been trained on a sufficiently broad set of data they can be used to predict responses from entirely new part/process setting combinations without requiring a previously identical build to have been fabricated as a baseline. The advantage of this approach is that a manufacturer can establish a large training set and then use that training set to validate future parts or process parameter changes, trading off a larger initial investment for less long-term cost being involved. Manufacturers such as printing services companies who do a high volume of varied parts can benefit significantly from this approach. A further benefit is that the manufacturer does not need to collect the training data if there is a trusted third party (such as the machine OEM) who has already gathered data and trained the system. Trained models can simply be passed along to customers who can quickly use the SCMS to begin validating parts.

7.4.3. Physics-based modeling/Digital twin

The most difficult case is where the part geometry is always unique and when the process parameters are varying in-situ. In this case a machine learning -based model could still be beneficial, but may have difficulty correctly predicting small real-time parameter changes that could differ between two identical parts. In this situation, there is increasingly the need for an accurate physics-based model that can take the side-channel measurements and accurately predict the resulting part properties. This is both more difficult to create and requires much more sophisticated monitoring in order to accurately capture the relevant information to predict the part properties. While some advanced research labs are establishing this type of capability, it is still in its infancy [101,102].

7.5. Aggregate Datasets

The high speed and/or resolution needed to detect defects in some AM processes can lead to extremely large datasets. For this reason, it is important to aggregate collected data into a manageable format. This makes data processing, storage, and transmission easier and can save both time and money. When multiple side-channels are used in a monitoring system, data aggregation is also important to ensure that the various data streams are synchronized and coupled in appropriate ways. Well aggregated data sets from multiple sensors can improve detection performance, but poorly grouped or stored data may reduce effectiveness and sensitivity of the SCMS.

7.5.1. Downsampling

One of the simplest ways to reduce data volume is through downsampling [103,104], where only a fraction of the data collected is actually transmitted/compared. The drawback of this approach is that it can significantly reduce the resolution of the monitoring system. A system sampling at 10kHz that is down sampled by a factor of 10, will require only 1/10 of the data, but will effectively only be sampling at a rate of 1kHz. If defects occur below this timescale they will likely avoid detection by the SCMS. If an attacker is familiar with the sampling procedure, they may be able to design defects that will escape detection.

7.5.2. Compression

Another method for reducing data volume is compression. There are a variety of compression algorithms that exist, with the most basic being simple downsampling approaches. Lossless compression can reduce the amount of data without losing any information; however, these approaches are limited in how much they can compress the data. Some data may not be compressible at all using these methods, while highly repetitive data may be very compressible [105]. Lossy compression algorithms (such as downsampling) can reduce the data volumes much more, but suffers from the same problem as downsampling where the loss of information will result in reduced resolution [106]. While both of these types of compression can be useful in combinations with other approaches for reducing data volumes, they are unlikely to be sufficient enough to be used by themselves, unless the original data volumes were already quite low, highly repetitive, or high-resolution data is not needed.

7.5.3. Hashing

One approach that can condense an indeterminately large amount of data into a finite length string is the use of a hash function. Hash functions are used to generate a signature (hash) for a file and can be used to quickly check to see if a file matches the original source of the hash. Cryptographic hashing is a subcategory of hash functions that offers several advantages in a security environment, with a key one being the irreversibility of the information stored inside. The security aspects of cryptographic hashing will be discussed more in detail in section 6.3. A drawback of cryptographic hash functions when applied to physical systems is that small changes in the input values (unavoidable in a manufacturing environment) will cause large changes in the output hash, resulting in two similar parts having very different hash strings. There are a few approaches for dealing with the limitations that this imposes. One technique is through the use of fuzzy hash functions. These functions are designed to allow some variation in the input parameters while still generating an output that can be matched between similar inputs. [107]

Another way to avoid these problems is through the use of predetermined ranges by the SCMS. In this setup, a list of ranges and thresholds are predetermined and a selection of these ranges is converted into a hash that is transmitted to the SCMS using physical emissions. The SCMS then records the emissions from the fabrication of a part and maps these emissions into one of the predefined ranges. The mapped ranges are then converted in a hash string that is compared to the hash string sent to the SCMS. If the two strings match, then the AM system is operating within acceptable parameters. If the two strings do not match, then the system is operating outside of expected parameters. This approach, dubbed a “cyber-physical hash” is expanded more in detail in the authors’ prior work [21].

7.5.4. Control Charting

Control charting is a tool used in manufacturing to determine whether or not a process is operating in a controlled manner and within the desired range. The traditional approach to control charting takes historical data and generates a set of statistical bounds that are used to compare future data against [108]. If a certain amount of data points in a row exceed these thresholds then the process is said to be out of control. This can work quite well for some emissions, such as temperature or laser power, that may fluctuate within a set range. Small or short-term deviations from the set value may not have a significant effect on part quality and the amount of data that needs to be sent is much smaller than sending raw sampled data. The natural fluctuations of raw data can also mean that control charting can be a better way of comparing two builds, rather than trying to directly compare each individual data point, since the natural variations make it unlikely that these will line up exactly, even with two identical builds. Various types of control charting can be used depending on the nature of the system being monitored as some control charting methods are more robust to outliers in the process [109]. Control charts can be applied to a wide range of sensor types including image data [110]. By sending control chart parameters and combining it with other methods for reducing data volumes, the total amount of data being sent can be reduced significantly [109,111–113].

7.5.5. Data Fusion

In systems monitoring side-channels another consideration is data fusion. By using multiple sensors and combining the data, the overall system can perform better than the sum of its components. Data fusion systems can increase the resolution, detection range, accuracy, and reliability of the system. Data fusion approaches may use using multiples of a sensor instead of a single one or can include multiple different types of sensors that are used to better predict build defects and the operating state of an AM system. Using multiple sensors will increase the data volume, however using data fusion approaches can combine these separate data streams and reduce the overall volume while at the same time increasing the performance of the system [82,114–118].

7.6. Secure Transmission

The final step in the IDEAS framework is securing the transmission of information to the side-channel monitoring system. While securing other steps of the AM process such as the input file and attaching the quality information to a part are important, these are outside the scope of this paper. There are three ways that information can be passed to a SCMS, as shown in Figure 7.8.

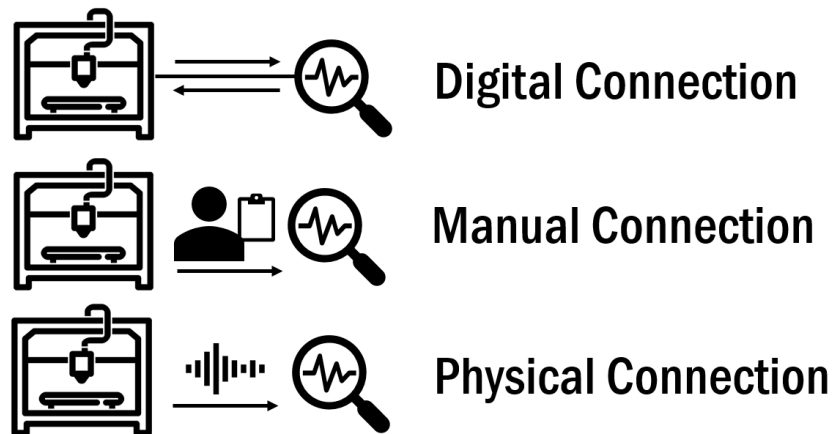


Figure 7.8. Ways of transmitting build information to a SCMS [22].

The first is by having a direct digital connection between the SCMS and the printer. This is the most common type of setup and has the advantage of allowing the monitoring system direct access to information about the part/build/toolpath. Most existing quality control and in-situ monitoring systems operate this way. Communication may be one or two directional, in the latter the monitoring system receives information from the AM system about the build and sends feedback to the AM system. In the former the monitoring system receives information about the build from the AM system, but does not send feedback to the AM system. While this type of connection works well from a quality monitoring standpoint, it has limitations from a security point of view. If the AM system becomes compromised, the digital connection means that malicious software has the potential to jump from the AM system to the monitoring system, compromising it as well. If the monitoring system provides feedback an attack could even use the monitoring system to cause the defect in the build by generating false values that force the system out of normal ranges. This is similar to how Stuxnet caused centrifuges to operate outside of regular parameters, while still showing normal parameter ranges to operators on their instruments.

To reduce the risk of a single attack compromising both the AM system and the SCMS it is desirable to digitally disconnect the two systems. However, build information still needs to be transferred to the monitoring system. A straightforward way to accomplish this is to do so manually, either through physically transferring some form of digital media such as a USB drive or optical disc, or by having an operator manually enter data into the SCMS. In both cases the process adds additional work and has the potential for operator error. Manual entry is time consuming and error prone, while digital media has the potential to expose the SCMS to attack, potentially negating many of the benefits of air-gapping the system. While these types of methods might be used occasionally (such as to update

the firmware of the SCMS), it is preferable to avoid having to manually transfer information on a regular basis.

A second way to transmit information to an air-gapped SCMS, without requiring manual interaction, is to encode that information into the toolpath in such a way that it will generate physical emissions that the SCMS can detect. These emissions provide a physical way of transmitting data, without relying on a digital connection. However, a drawback of this approach is that the amount and rate at which data can be transferred is significantly reduced. Side channels can vary from extremely low bitrate (thermal [119], fan/smaller [120]) to more moderate bitrates (acoustic, visual); however, these fall far short of the Mb/sec - Gb/sec rates that can be achieved in modern digital systems. For this reason, it is necessary to sufficiently reduce the amount of data required to be transmitted while still maintaining sufficient information for in-situ defect detection.

When using this approach, it is important to develop both the data package and the monitoring system at the same time. The data package needs to be able to communicate sufficient information to the monitoring system and the monitoring system needs to be able to receive that data while at the same time monitoring the relevant part properties and emissions from the process in order to detect attacks. Figure 3. shows a flowchart of the selection process for sensors when using an air-gapped SCMS. Section 5 discuss how to reduce the amount of data that needs to be stored and sent, reducing the requirements of the SCMS.

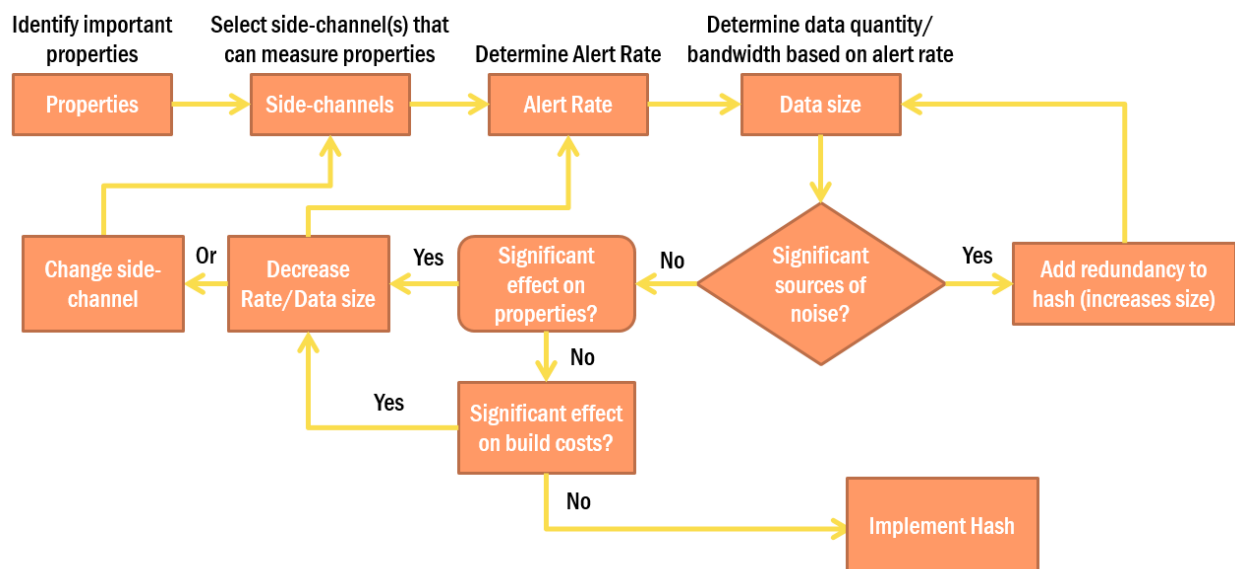


Figure 7.9. Flowchart for selecting side-channels for use with a data package, stored in the toolpath, that contains part quality information for validating the build against sabotage attacks.

7.6.1. Security Approaches

There are several areas in which a manufacture may want to ensure the security of information in the system. The first and most critical is to ensure that no changes can be made to the monitoring

system information without resulting in a rejection of the part by the system. The second area is in protecting the potentially valuable IP of the monitoring/parameter settings being used to validate the build. If a part file is stolen it is preferable not to have the process/quality settings also stolen, which could allow the attacker to potentially replicate the parts. The third area is protecting the part data itself. This is the primary consideration for many manufacturers worried about their IP being stolen and counterfeit products being produced.

There are two sets of data that need to be considered (Figure 9). The first set is the part data which determines whether or not the machine is given the correct instructions to fabricate a good part. If the part data is attacked, the system will produce a defective part and this should be detected by the SCMS. If the part data is not attacked, the system will produce a good part and the SCMS should mark it as acceptable. The second set of data is the information sent to the SCMS that is used to tell the system what values to check the part data against. If the SCMS data has not been attacked, the output from the system should be the same as the values of the input data (i.e., pass good parts, reject modified parts). However, if the data sent to the SCMS is modified, the system will reject good parts that consist of unmodified data. While these false positives may generate scrap in the form of rejected good parts, they should also alert the system that a modification (potential attack) has occurred. The most dangerous case is the one in which the monitoring system data has been modified and the part data has been modified to match, as it would enable production (and qualification) of bad parts. . These false negatives are the most dangerous case from a security sabotage standpoint because they are the case where defective parts could be put into use. To avoid this scenario, the data sent to the SCMS needs to be secured in some way to prevent an attack from modifying it in such a way as to allow a defect part to be considered good. There are several approaches that can be used to protect the transmission of this information.




Figure 7.10. Examples of different cases for the SCMS where the part/data package are modified/unmodified. The manufacturer needs to ensure that an attacker cannot alter the data package to match a modified part or the system could reject or pass the part.

7.6.1.1. Encryption

One of the most straightforward ways of protecting data is through the use of encryption. By encrypting the part/toolpath file and decrypting it at the AM system, manufacturers can significantly reduce the points at which an attacker could compromise the data. However, if an attack has compromised the AM system it could also compromise the data on the AM system once it has been decrypted, allowing it to be modified both the part data and the data being sent to the SCMS. A simple way to mitigate this is by separately encrypting the data sent to the SCMS before adding it to the toolpath. In this case, the toolpath is decrypted when it is fabricated, but the information that is generated during fabrication is encrypted and transmitted through the physical transmissions. The SCMS then receives and decrypts the monitoring information and validates the build. If an attacker compromises the build and not the monitoring data then the SCMS will detect the defect, if the attacker attacks the part data and the monitoring data the encryption will mean that they will be unable to generate a matching set of monitoring data (the changes will be effectively random and extremely unlikely to generate a collision) this will result in the SCMS rejecting the part as defective. Any attempt to modify the encrypted monitoring data will result in the system rejecting the part due to the monitoring data no longer matching the part data unless the attacker is able to break the encryption and modify the plain text monitoring information.

7.6.1.2. Hashing

Another way of securing the transmission of the monitoring data is through the use of a hash-string. Unlike encryption, cryptographic hashing is an irreversible process, meaning that even given effectively infinite computational power an attacker would be unable to recover the original information from the hash string. Key cryptographic hash features are:

- 1) Pre-image resistance: Given a specific hash it is infeasible to find a message that maps to that hash
- 2) Collision resistance: It should be infeasible to find two messages that map to the same hash
- 3) Avalanche effect: Small changes to the input should generate a large change to the output value

These features prevent an attacker from extracting information from a hash and also from creating a forged hash. This approach is both compact and secure, making it a useful in many situations where there are limitations on data transmission and concerns about the theft of process/quality settings IP. In the case of using predetermined ranges for generating the hash strings (as discussed in Section 5.3), it is important to include some type of key when generating the hash string. The reason for this is that in the event of the theft of the file an attacker could potentially reverse the hash ranges through brute force. While this would likely be time consuming and expensive (due to having to physically fabricate parts using different settings), the limited number of physical possibilities means that given sufficient time and resources a determined attacker might be able to reverse the system. By including a key, the complexity of brute forcing the system can be substantially increased (both the key and the parameters must now be brute forced). With a reasonably secure key (such as those commonly used in cryptographic systems) the manufacturer can make the system infeasible to brute force.

7.6.1.3. Steganography

Steganography is a technique for concealing information inside of another object or message. In the case of AM, it can be used to conceal information inside of the model file in a way that can later be extracted. This is most commonly used for storing/extracting information digitally, such as in the .STL mech and significant work has been done with 3D meshes [121–125]. Physical steganography can be more difficult (due to the small size of the modifications and greater difficulty of extracting physical information), but can also be implemented to allow for the physical part to contain the information [121,126]. The drawback of steganography is that it only provides security while the approach is unknown. If an attacker (or other party) is aware that steganography is being used they may be able to remove of the modifications containing the information or extract stenographic information using steganalysis. While stenography can provide some security, particularly if combined with other approaches, it should not be relied on as a primary security method.

7.6.1.4. Watermarking

Watermarking can be used on both digital files and physical parts and can either be visible (easily detected/seen) or invisible (hidden to make detection difficult. Invisible watermarks can be used as a form of steganography. The primary goal of watermarking is to validate the origin or authenticity of the watermarked file (e.g. a watermarked image or photograph). In the context of manufacturing, watermarks may be used as a method of preventing counterfeiting. While this approach may have some efficacy, such as making it difficult to replicate parts through 3D scanning, the drawback is that if the model files containing the watermark are stolen (or otherwise acquired by a 3rd party), then counterfeit parts can be fabricated containing the original watermark. This may be useful in tracking or pursuing legal recourse against the 3rd party, but does not prevent the consumer from accidentally acquiring counterfeit parts. Various examples of watermarking for AM applications have been proposed [127–131]. While watermarking may be helpful for mitigating counterfeiting, it does not help prevent attacks that seek to sabotage parts directly.

7.6.1.5. Physical unclonable functions

A physical unclonable function (PUF) uses inherent randomness to establish a unique signature that a would-be attacker is unable to duplicate. They have found considerable use in electronics as a security measure as a method[132–136] . In a manufacturing environment, PUFs have application as an anticounterfeiting method to uniquely identify individual parts and validate their authenticity. Unlike predesigned watermarks, which have a fixed representation and can be replicated with access to similar equipment, PUFs are unique to each part. In the context of AM, a variety of approaches have been proposed, based on both chemical and physical techniques. Approaches include the use of UV florescent quantum dots [137–140] as well as other types of chemical fingerprinting [141,142]. Other proposed methods use the microstructure [143] or unique printing noise [144,145] to identify the part or system used for fabrication. Overall, PUFs are an area of continued interest in AM security. While encryption and other digital techniques can secure the digital components of the AM process chain and in situ monitoring and sensing can validate the physical quality of a part, PUFs allow quality data to be linked directly to parts in a way that cannot be replicated by an attacker. This is a key connection between the

digital data and the physical part that can be used to ensure that a received part is the same as the one fabricated and validated, and not a counterfeit.

7.7. Conclusions

The growth of networked manufacturing and an increasing number of targeted attacks has led to the need for improved security and monitoring of AM systems. While a variety of techniques and methods have been proposed for improving the security of AM systems, there remains a need for a comprehensive approach for detecting cyber-physical attacks using physical methods. In this paper the authors have proposed a novel framework for securing AM systems through the use of side-channel monitoring. The IDEAS (Identify, Define, Establish, Aggregate, Secure) framework presents an approach for *identifying* attack surfaces in the system, *defining* side-channels to monitor, *establishing* a nominal baseline, *aggregating* the output data, and *securing* the transmission of that data. Steps for identification of attack vectors were presented in Section 2 by defining critical part properties and tolerances, identifying high level process inputs and system specific process-part interactions, quantifying the interactions between process parameters and part properties, defining an attack detection threshold, and evaluating the severity of attacks. An overview of how to appropriately select side-channels based on the previously identified attack vectors was provided in Section 3 by determining the monitoring approach and sensor requirements. flowchart for sensor selection and highlighted the need for a calibration and maintenance schedule for maintaining the monitoring system. A discussion on how to establish baselines for the SCMS was presented in Section 4. This discussion included an overview of different use cases of additive manufacturing and techniques than can be used for these cases. The importance of data aggregation and several approaches for doing so such as compression, hashing, control charting, and data fusion were discussed in Section 5 along with the advantages and drawbacks of each. Section 6 addressed the need for securing the transmitted data to avoid tampering or IP theft. Both digital techniques (encryption, hashing) and physical techniques (watermarking, PUFs) were introduced along with a discussion of their role and usefulness in securing AM systems.

Overall, the IDEAS framework provides a basis for manufacturers to design and implement an air-gapped SCMS as a way to physically detect part sabotage attacks. Unlike previous work in this area, which often focuses on a subset of this framework, IDEAS provides a comprehensive framework that addresses the relevant cyber-physical security concerns in designing a SCMS from start to finish. The objective of this framework is to provide a guide for evaluating and securing AM systems using physical monitoring and to provide a centralized source of information in this area and to highlight areas where additional work may need to be done.

Acknowledgements

This material is based upon work supported by the National Science Foundation under Grant No. CMMI-1436365 and Grant No. CMMI-1635356.

Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation

7.8. References

- [1] L.D. Sturm, C.B. Williams, J.A. Camelio, J. White, R. Parker, Cyber-physical vulnerabilities in additive manufacturing systems: A case study attack on the.STL file with human subjects, *J. Manuf. Syst.* 44 (2017). <https://doi.org/10.1016/j.jmsy.2017.05.007>.
- [2] M. Yampolskiy, W.E. King, J. Gatlin, S. Belikovetsky, A. Brown, A. Skjellum, Y. Elovici, Security of additive manufacturing: Attack taxonomy and survey, *Addit. Manuf.* 21 (2018) 431–457. <https://doi.org/10.1016/j.addma.2018.03.015>.
- [3] N. Falliere, L.O. Murchu, E. Chien, W32.Stuxnet Dossier, 4 (2011) 1–69.
- [4] M. Holloway, Stuxnet Worm Attack on Iranian Nuclear Facilities, (2015).
- [5] K. (Wired) Zetter, A Cyberattack Has Caused Confirmed Physical Damage for the Second Time Ever | WIRED, (2015). <http://www.wired.com/2015/01/german-steel-mill-hack-destruction/> (accessed January 7, 2016).
- [6] B. für S. in der I. BSI, Die Lage der in Deutschland 2011, *Informationstechnik.* (2011).
- [7] R. Lee, M. Assante, T. Conway, Analysis of the Cyber Attack on the Ukrainian Power Grid, *SANS Ind. Control Syst. Secur. Blog.* (2016) 1–26.
- [8] A.E. Elhabashy, J.A. Camelio, L.J. Wells, W.H. Woodall, V. Tech, Misuse of Quality Control Tools in Manufacturing Abstract ID : 2226, (2018) 2–4.
- [9] L.D.L.D. Sturm, C.B.C.B. Williams, J.A.J.A. Camelio, J. White, R. Parker, Cyber-physical vulnerabilities in additive manufacturing systems: A case study attack on the.STL file with human subjects, *J. Manuf. Syst.* 44 (2017) 154–164. <https://doi.org/10.1016/j.jmsy.2017.05.007>.
- [10] S. Belikovetsky, M. Yampolskiy, J. Toh, Y. Elovici, drOwned - Cyber-Physical Attack with Additive Manufacturing, *CoRR.* abs/1609.0 (2016). <http://arxiv.org/abs/1609.00133>.
- [11] S.E. Zeltmann, N. Gupta, N.G. Tsoutsos, M. Maniatakos, J. Rajendran, R. Karri, Manufacturing and Security Challenges in 3D Printing, *JOM.* 68 (2016) 1872–1881. <https://doi.org/10.1007/s11837-016-1937-7>.
- [12] J. Straub, An approach to detecting deliberately introduced defects and micro-defects in 3D printed objects, *Pattern Recognit. Track. XXVIII.* 10203 (2017) 102030L. <https://doi.org/10.1117/12.2264588>.
- [13] S. Belikovetsky, Y. Solewicz, M. Yampolskiy, J. Toh, Y. Elovici, Digital Audio Signature for 3D Printing Integrity, *IEEE Trans. Inf. Forensics Secur.* PP (2018) 1. <https://doi.org/10.1109/TIFS.2018.2851584>.
- [14] A. Cui, M. Costello, S.J. Stolfo, When Firmware Modifications Attack : A Case Study of Embedded Exploitation, *Ndss.* (2013).
- [15] S. Moore, P. Armstrong, T. McDonald, M. Yampolskiy, Vulnerability analysis of desktop 3D printer software, *Proc. - 2016 Resil. Week, RWS 2016.* (2016) 46–51. <https://doi.org/10.1109/RWEEK.2016.7573305>.
- [16] A. Slaughter, M. Yampolskiy, M. Matthews, W.E. King, G. Guss, Y. Elovici, How to Ensure Bad Quality in Metal Additive Manufacturing, *Proc. 12th Int. Conf. Availability, Reliab. Secur. - ARES '17.* (2017) 1–10. <https://doi.org/10.1145/3098954.3107011>.

- [17] M. Yampolskiy, L. Schutzle, U. Vaidya, A. Yasinsac, Security Challenges of Additive Manufacturing with Metals and Alloys, in: M. Rice, S. Shenoj (Eds.), Crit. Infrastruct. Prot. IX 9th IFIP 11.10 Int. Conf. ICCIP 2015, Arlington, VA, USA, March 16-18, 2015, Revis. Sel. Pap., Springer International Publishing, Cham, 2015: pp. 169–183.
- [18] N. Gupta, A. Tiwari, S.T.S. Bukkapatnam, R. Karri, Additive Manufacturing Cyber-Physical System: Supply Chain Cybersecurity and Risks, *IEEE Access*. 8 (2020) 47322–47333. <https://doi.org/10.1109/ACCESS.2020.2978815>.
- [19] G. Pope, M. Yampolskiy, A Hazard Analysis Technique for Additive Manufacturing, (2017). <http://arxiv.org/abs/1706.00497>.
- [20] D. Agrawal, B. Archambeault, J.R. Rao, P. Rohatgi, The em Side-Channel(s), *Lect. Notes Comput. Sci. (Including Subser. Lect. Notes Artif. Intell. Lect. Notes Bioinformatics)*. 2523 (2003) 29–45. https://doi.org/10.1007/3-540-36400-5_4.
- [21] J. Brandman, L. Sturm, J. White, C. Williams, A physical hash for preventing and detecting cyber-physical attacks in additive manufacturing systems, *J. Manuf. Syst.* 56 (2020) 202–212. <https://doi.org/10.1016/j.jmsy.2020.05.014>.
- [22] fdm, Cartesian, printer, enclosed icon, (n.d.). <https://www.shareicon.net/fdm-cartesian-printer-enclosed-112026> (accessed September 3, 2020).
- [23] M. Barrett, Framework for improving critical infrastructure cybersecurity, *Proc. Annu. ISA Anal. Div. Symp.* 535 (2018) 9–25.
- [24] A. Wegner, J. Graham, E. Ribble, A New Approach to Cyber-physical Security in Industry 4.0, in *Cybersecurity for Industry 4.0*, 2017. <https://doi.org/10.1007/978-3-319-50660-9>.
- [25] H. Turner, B. Amos, J. White, J. Camelio, C. Williams, Bad parts: Are our manufacturing systems at risk of silent cyber-attacks, *IEEE Secur. Priv.* (2015) 40–47. <https://doi.org/10.1109/MSP.2015.60>.
- [26] T. Huelsman, E. Powers, S. Peasley, R. Robinson, Cyber risk in advanced manufacturing, *Deloitte MAPI*. (2016) 53. <https://www2.deloitte.com/us/en/pages/manufacturing/articles/cyber-risk-in-advanced-manufacturing.html>.
- [27] M. Mcgrath, Cyber Security for Advanced Manufacturing Protecting the Digital Thread Software and Supply Chain Assurance (SSCA) Continuation of Discussion Started at SSCA Winter 2016 Session, (2017).
- [28] S.M. Bridges, S.T.A.T. Hall, S.J. Graves, S.T.A.T. Hall, K. Keiser, S.T.A.T. Hall, N. Sissom, S.T.A.T. Hall, S.J. Graves, Cyber Security for Additive Manufacturing, in: *Proc. 10th Annu. Cyber Inf. Secur. Res. Conf.*, ACM, New York, NY, USA, 2015: pp. 14:1----14:3. <https://doi.org/10.1145/2746266.2746280>.
- [29] C. Xiao, Security Attack to 3D Printing, (2013). <https://www.claudxiao.net/Attack3DPrinting-Claud-en.pdf>.
- [30] O. Oyelola, P. Crawforth, R. M'Saoubi, A.T. Clare, Machining of Additively Manufactured Parts: Implications for Surface Integrity, *Procedia CIRP*. 45 (2016) 119–122. <https://doi.org/10.1016/j.procir.2016.02.066>.
- [31] G. Ameta, P. Witherell, S. Moylan, R. Lipman, Tolerance specification and related issues for

- additively manufactured products, Proc. ASME Des. Eng. Tech. Conf. 1A-2015 (2015).
<https://doi.org/10.1115/DETC2015-47531>.
- [32] G. Ameta, S. Moylan, P. Witherell, R. Lipman, Challenges in tolerance transfer for additive manufacturing, Proc. - ASPE 2015 Spring Top. Meet. Achiev. Precis. Toler. Addit. Manuf. (2015) 129–135.
- [33] P. Witherell, J. Herron, G. Ameta, Towards Annotations and Product Definitions for Additive Manufacturing, Procedia CIRP. 43 (2016) 339–344. <https://doi.org/10.1016/j.procir.2016.01.198>.
- [34] I.F. Ituarte, E. Coatanea, M. Salmi, J. Tuomi, J. Partanen, Additive Manufacturing in Production: A Study Case Applying Technical Requirements, Phys. Procedia. 78 (2015) 357–366.
<https://doi.org/10.1016/j.phpro.2015.11.050>.
- [35] J. Xiao, N. Anwer, A. Durupt, J. Le Duigou, B. Eynard, Information exchange standards for design, tolerancing and Additive Manufacturing: a research review, Int. J. Interact. Des. Manuf. 12 (2018) 495–504. <https://doi.org/10.1007/s12008-017-0401-4>.
- [36] R. Lipman, S. Moylan, P. Witherell, Investigating the Role of Geometric Dimensioning and Tolerancing in Additive Manufacturing, (2015). <https://doi.org/10.1115/1.4031296>.
- [37] W. Polini, G. Td, I.F. Moroni, S. Petro, CIRP Annals - Manufacturing Technology Geometrical product specification and verification in additive manufacturing, 66 (2017) 157–160.
<https://doi.org/10.1016/j.cirp.2017.04.043>.
- [38] T. Lieneke, V. Denzer, G.A.O. Adam, D. Zimmer, Dimensional Tolerances for Additive Manufacturing: Experimental Investigation for Fused Deposition Modeling, Procedia CIRP. 43 (2016) 286–291. <https://doi.org/10.1016/j.procir.2016.02.361>.
- [39] M. Seifi, M. Gorelik, J. Waller, N. Hrabe, N. Shamsaei, S. Daniewicz, J.J. Lewandowski, Progress Towards Metal Additive Manufacturing Standardization to Support Qualification and Certification, Jom. 69 (2017) 439–455. <https://doi.org/10.1007/s11837-017-2265-2>.
- [40] A.D. Peralta, M. Enright, M. Megahed, J. Gong, M. Roybal, J. Craig, Towards rapid qualification of powder-bed laser additively manufactured parts, Integr. Mater. Manuf. Innov. 5 (2016) 154–176.
<https://doi.org/10.1186/s40192-016-0052-5>.
- [41] ISO/ASTM, INTERNATIONAL STANDARD ISO / ASTM 52900 Additive manufacturing — General principles — Terminology, Int. Organ. Stand. 5 (2015) 1–26.
<https://doi.org/10.1520/ISOASTM52900-15>.
- [42] J.E. Seppala, K.D. Migler, Infrared thermography of welding zones produced by polymer extrusion additive manufacturing, Addit. Manuf. 12 (2016) 71–76.
<https://doi.org/10.1016/j.addma.2016.06.007>.
- [43] S.H. Ahn, M. Montero, D. Odell, S. Roundy, P.K. Wright, Anisotropic material properties of fused deposition modeling ABS, Rapid Prototyp. J. 8 (2002) 248–257.
<https://doi.org/10.1108/13552540210441166>.
- [44] V. Meenakshisundaram, L.D. Sturm, C.B. Williams, Modeling A Scanning-Mask Projection Vat Photopolymerization System For Multiscale Additive Manufacturing, J. Mater. Process. Technol. 279 (2020) 116546. <https://doi.org/10.1016/j.jmatprotec.2019.116546>.

- [45] P.F. Jacobs, Rapid prototyping & manufacturing— Fundamentals of stereolithography, *J. Manuf. Syst.* 12 (1993) 430–433. [https://doi.org/10.1016/0278-6125\(93\)90311-g](https://doi.org/10.1016/0278-6125(93)90311-g).
- [46] S. Jayanthi, M. Keefe, E. Gargiulo, Studies in stereolithography: influence of process parameters on curl distortion in photopolymer models, *Solid Free. Fabr.* (1994) 250–258. <http://www.dtic.mil/cgi-bin/GetTRDoc?AD=ADA290949#page=259>.
- [47] R.B.S. Gowda, C.S. Udayagiri, D.D. Narendra, Studies on the Process Parameters of Rapid Prototyping Technique (Stereolithography) for the Betterment of Part Quality, *Int. J. Manuf. Eng.* 2014 (2014) 1–11. <https://doi.org/10.1155/2014/804705>.
- [48] I.H. Lee, D.W. Cho, An investigation on photopolymer solidification considering laser irradiation energy in micro-stereolithography, *Microsyst. Technol.* 10 (2004) 592–598. <https://doi.org/10.1007/s00542-003-0337-4>.
- [49] C. Kamath, B. El-Dasher, G.F. Gallegos, W.E. King, A. Sisto, Density of additively-manufactured, 316L SS parts using laser powder-bed fusion at powers up to 400 W, *Int. J. Adv. Manuf. Technol.* 74 (2014) 65–78. <https://doi.org/10.1007/s00170-014-5954-9>.
- [50] Y.M. Arisoy, L.E. Criales, T. Özel, B. Lane, S. Moylan, A. Donmez, Influence of scan strategy and process parameters on microstructure and its optimization in additively manufactured nickel alloy 625 via laser powder bed fusion, *Int. J. Adv. Manuf. Technol.* 90 (2017) 1393–1417. <https://doi.org/10.1007/s00170-016-9429-z>.
- [51] C.A. Chatham, T.E. Long, C.B. Williams, A review of the process physics and material screening methods for polymer powder bed fusion additive manufacturing, *Prog. Polym. Sci.* 93 (2019) 68–95. <https://doi.org/10.1016/j.progpolymsci.2019.03.003>.
- [52] T.L. Starr, T.J. Gornet, J.S. Usher, The effect of process conditions on mechanical properties of laser-sintered nylon, *Rapid Prototyp. J.* 17 (2011) 418–423. <https://doi.org/10.1108/13552541111184143>.
- [53] A. Dey, N. Yodo, A systematic survey of FDM process parameter optimization and their influence on part characteristics, *J. Manuf. Mater. Process.* 3 (2019). <https://doi.org/10.3390/jmmp3030064>.
- [54] Design - Sintavia, (n.d.). <https://sintavia.com/dfam/> (accessed September 3, 2020).
- [55] A.M. Khorasani, I. Gibson, U.S. Awan, A. Ghaderi, The effect of SLM process parameters on density, hardness, tensile strength and surface quality of Ti-6Al-4V, *Addit. Manuf.* 25 (2019) 176–186. <https://doi.org/10.1016/j.addma.2018.09.002>.
- [56] C.U. Brown, G. Jacob, A. Possolo, C. Beauchamp, M. Peltz, M. Stoudt, A. Donmez, The Effects of Laser Powder Bed Fusion Process Parameters on Material Hardness and Density for Nickel Alloy 625, *NIST Adv. Manuf. Ser.* (2018) 100–119. <https://doi.org/10.6028/NIST.AMS.100-19>.
- [57] C. Qiu, N.J.E. Adkins, M.M. Attallah, Microstructure and tensile properties of selectively laser-melted and of HIPed laser-melted Ti-6Al-4V, *Mater. Sci. Eng. A.* 578 (2013) 230–239. <https://doi.org/10.1016/j.msea.2013.04.099>.
- [58] S.M.H. Hojjatzadeh, N.D. Parab, Q. Guo, M. Qu, L. Xiong, C. Zhao, L.I. Escano, K. Fezzaa, W. Everhart, T. Sun, L. Chen, Direct observation of pore formation mechanisms during LPBF additive manufacturing process and high energy density laser welding, *Int. J. Mach. Tools Manuf.* 153

- (2020) 103555. <https://doi.org/10.1016/j.ijmachtools.2020.103555>.
- [59] E. Liverani, S. Toschi, L. Ceschini, A. Fortunato, Effect of selective laser melting (SLM) process parameters on microstructure and mechanical properties of 316L austenitic stainless steel, *J. Mater. Process. Technol.* 249 (2017) 255–263. <https://doi.org/10.1016/j.jmatprotec.2017.05.042>.
- [60] Z. Wang, T.A. Palmer, A.M. Beese, Effect of processing parameters on microstructure and tensile properties of austenitic stainless steel 304L made by directed energy deposition additive manufacturing, *Acta Mater.* 110 (2016) 226–235. <https://doi.org/10.1016/j.actamat.2016.03.019>.
- [61] N. Kang, Y. Li, X. Lin, E. Feng, W. Huang, Microstructure and tensile properties of Ti-Mo alloys manufactured via using laser powder bed fusion, *J. Alloys Compd.* 771 (2019) 877–884. <https://doi.org/10.1016/j.jallcom.2018.09.008>.
- [62] X. Wang, L. Zhao, J.Y.H. Fuh, H.P. Lee, Effect of porosity on mechanical properties of 3D printed polymers: Experiments and micromechanical modeling based on X-ray computed tomography analysis, *Polymers (Basel)*. 11 (2019). <https://doi.org/10.3390/polym11071154>.
- [63] J.A. Slotwinski, E.J. Garboczi, K.M. Hebenstreit, Porosity measurements and analysis for metal additive manufacturing process control, *J. Res. Natl. Inst. Stand. Technol.* 119 (2014) 494–528. <https://doi.org/10.6028/jres.119.019>.
- [64] E.R. Denlinger, J.C. Heigel, P. Michaleris, T.A. Palmer, Effect of inter-layer dwell time on distortion and residual stress in additive manufacturing of titanium and nickel alloys, *J. Mater. Process. Technol.* 215 (2015) 123–131. <https://doi.org/10.1016/j.jmatprotec.2014.07.030>.
- [65] Y. Lei, J. Xiong, R. Li, Effect of inter layer idle time on thermal behavior for multi-layer single-pass thin-walled parts in GMAW-based additive manufacturing, *Int. J. Adv. Manuf. Technol.* 96 (2018) 1355–1365. <https://doi.org/10.1007/s00170-018-1699-1>.
- [66] M. Faes, E. Ferraris, D. Moens, Influence of Inter-layer Cooling time on the Quasi-static Properties of ABS Components Produced via Fused Deposition Modelling, *Procedia CIRP*. 42 (2016) 748–753. <https://doi.org/10.1016/j.procir.2016.02.313>.
- [67] S.K. Everton, M. Hirsch, P. Stravroulakis, R.K. Leach, A.T. Clare, P.I. Stavroulakis, R.K. Leach, A.T. Clare, Review of in-situ process monitoring and in-situ metrology for metal additive manufacturing, *Mater. Des.* 95 (2016) 431–445. <https://doi.org/10.1016/j.matdes.2016.01.099>.
- [68] G. Tapia, A. Elwany, A Review on Process Monitoring and Control in Metal-Based Additive Manufacturing, *J. Manuf. Sci. Eng.* 136 (2014) 60801–60810. <https://doi.org/10.1115/1.4028540>.
- [69] Q.Y. Lu, C.H. Wong, Additive manufacturing process monitoring and control by non-destructive testing techniques: challenges and in-process monitoring, *Virtual Phys. Prototyp.* 13 (2018) 39–48. <https://doi.org/10.1080/17452759.2017.1351201>.
- [70] I.T. Cummings, M.E. Bax, I.J. Fuller, A.J. Wachtor, J.D. Bernardin, A framework for additive manufacturing process monitoring & control, in: *Conf. Proc. Soc. Exp. Mech. Ser.*, Springer New York LLC, 2017: pp. 137–146. https://doi.org/10.1007/978-3-319-54810-4_14.
- [71] L. Saerens, C. Vervaet, J.P. Remon, T. De Beer, Process monitoring and visualization solutions for hot-melt extrusion: a review, *J. Pharm. Pharmacol.* 66 (2014) 180–203. <https://doi.org/10.1111/jphp.12123>.

- [72] H. Kim, Y. Lin, T.L.B. Tseng, A review on quality control in additive manufacturing, *Rapid Prototyp. J.* 24 (2018) 645–669. <https://doi.org/10.1108/RPJ-03-2017-0048>.
- [73] A. Thompson, I. Maskery, R.K. Leach, X-ray computed tomography for additive manufacturing: A review, *Meas. Sci. Technol.* 27 (2016). <https://doi.org/10.1088/0957-0233/27/7/072001>.
- [74] H. Rieder, M. Spies, J. Bamberg, B. Henkel, On- and offline ultrasonic characterization of components built by SLM additive manufacturing, *AIP Conf. Proc.* 1706 (2016). <https://doi.org/10.1063/1.4940605>.
- [75] H. Rieder, A. Dillhöfer, M.S. Bamberg, H. Rieder, A. Dillhöfer, M. Spies, J. Bamberg, T. Hess, laser melting Ultrasonic Online Monitoring of Additive Manufacturing Processes Based on Selective Laser Melting, 184 (2016). <https://doi.org/10.1063/1.4914609>.
- [76] G. Diamond, D.A. Hutchins, K.K. Leong, T.H. Gan, Electrostatic capacitive imaging: A new NDE technique, in: *AIP Conf. Proc.*, AIP, 2007: pp. 689–694. <https://doi.org/10.1063/1.2718037>.
- [77] L. Sturm, M. Albakri, C.B. Williams, P. Tarazaga, In-Situ Detection of Build Defects in Additive Manufacturing via Impedance-Based Monitoring, 27th Annu. Int. Solid Free. Fabr. Symp. - An Addit. Manuf. Conf. (2016) 1458–1478.
- [78] J. Straub, Identifying positioning-based attacks against 3D printed objects and the 3D printing process, *Pattern Recognit. Track. XXVIII.* 10203 (2017) 1020304. <https://doi.org/10.1117/12.2264671>.
- [79] M. Aminzadeh, A machine vision system for in-situ quality inspection in metal powder-bed additive manufacturing, (2016). <https://smartech.gatech.edu/handle/1853/56291>.
- [80] J. Straub, 3D printing cybersecurity: detecting and preventing attacks that seek to weaken a printed object by changing fill level, *Dimens. Opt. Metrol. Insp. Pract. Appl.* VI. 10220 (2017) 1022000. <https://doi.org/10.1117/12.2264575>.
- [81] M. Wu, V. V Phoha, Y.B. Moon, A.K. Belman, Detecting Malicious Defects in 3D Printing Process Using Machine Learning and Image Classification, in: Vol. 14 *Emerg. Technol. Mater. Genet. to Struct. Saf. Eng. Risk Anal.*, ASME, 2016: p. V014T07A004. <https://doi.org/10.1115/IMECE2016-67641>.
- [82] P.K. Rao, J. (Peter) Liu, D. Roberson, Z. (James) Kong, C. Williams, Online Real-Time Quality Monitoring in Additive Manufacturing Processes Using Heterogeneous Sensors, *J. Manuf. Sci. Eng.* 137 (2015) 61007–61012. <https://doi.org/10.1115/1.4029823>.
- [83] Z. Li, Z. Zhang, J. Shi, D. Wu, Prediction of surface roughness in extrusion-based additive manufacturing with machine learning, *Robot. Comput. Integr. Manuf.* 57 (2019) 488–495. <https://doi.org/10.1016/j.rcim.2019.01.004>.
- [84] A.G. Demir, C. De Giorgi, B. Previtali, Design and implementation of a multi-sensor coaxial monitoring system with correction strategies for selective laser melting of a maraging steel, *J. Manuf. Sci. Eng.* 140 (2017) 1–14. <https://doi.org/10.1115/1.4038568>.
- [85] J.C. Fox, B.M. Lane, H. Yeung, Measurement of process dynamics through coaxially aligned high speed near-infrared imaging in laser powder bed fusion additive manufacturing, in: P. Bison, D. Burleigh (Eds.), *Thermosense Therm. Infrared Appl.* XXXIX, SPIE, 2017: p. 1021407. <https://doi.org/10.1117/12.2263863>.

- [86] B.A. Fisher, B. Lane, H. Yeung, J. Beuth, Toward determining melt pool quality metrics via coaxial monitoring in laser powder bed fusion, *Manuf. Lett.* 15 (2018) 119–121. <https://doi.org/10.1016/j.mfglet.2018.02.009>.
- [87] J. Gatlin, S. Belikovetsky, S.B. Moore, Y. Solewicz, Y. Elovici, M. Yampolskiy, Detecting sabotage attacks in additive manufacturing using actuator power signatures, *IEEE Access.* 7 (2019) 133421–133432. <https://doi.org/10.1109/ACCESS.2019.2928005>.
- [88] S.R. Chhetri, A. Canedo, M.A. Al Faruque, KCAD: Kinetic Cyber-Attack Detection Method for Cyber-Physical Additive Manufacturing Systems, *Proc. 35th Int. Conf. Comput. Des. - ICCAD '16.* (2016) 1–8. <https://doi.org/10.1145/2966986.2967050>.
- [89] S.A. Shevchik, C. Kenel, C. Leinenbach, K. Wasmer, Acoustic emission for in situ quality monitoring in additive manufacturing using spectral convolutional neural networks, *Addit. Manuf.* 21 (2018) 598–604. <https://doi.org/10.1016/j.addma.2017.11.012>.
- [90] M.A. Al Faruque, S.R. Chhetri, A. Canedo, J. Wan, Forensics of thermal side-channel in additive manufacturing systems, *CECS Tech. Report# 16-01.* (2016).
- [91] B. Lane, S. Moylan, E.P. Whinton, L. Ma, Thermographic measurements of the commercial laser powder bed fusion process at NIST, in: *Rapid Prototyp. J.*, Emerald Group Publishing Ltd., 2016: pp. 778–787. <https://doi.org/10.1108/RPJ-11-2015-0161>.
- [92] P.A. Hooper, Melt pool temperature and cooling rates in laser powder bed fusion, *Addit. Manuf.* 22 (2018) 548–559. <https://doi.org/10.1016/j.addma.2018.05.032>.
- [93] Y. Zhang, G.S. Hong, D. Ye, K. Zhu, J.Y.H. Fuh, Extraction and evaluation of melt pool, plume and spatter information for powder-bed fusion AM process monitoring, *Mater. Des.* 156 (2018) 458–469. <https://doi.org/10.1016/j.matdes.2018.07.002>.
- [94] L. Scime, J. Beuth, Using machine learning to identify in-situ melt pool signatures indicative of flaw formation in a laser powder bed fusion additive manufacturing process, *Addit. Manuf.* 25 (2019) 151–165. <https://doi.org/10.1016/j.addma.2018.11.010>.
- [95] J.C. Fox, S.P. Moylan, B.M. Lane, Effect of Process Parameters on the Surface Roughness of Overhanging Structures in Laser Powder Bed Fusion Additive Manufacturing, *Procedia CIRP.* 45 (2016) 131–134. <https://doi.org/10.1016/j.procir.2016.02.347>.
- [96] D. Koupryanoff, N. Luwes, E. Newby, I. Yadroitsava, I. Yadroitsev, On-line monitoring of laser powder bed fusion by acoustic emission: Acoustic emission for inspection of single tracks under different powder layer thickness, in: *2017 Pattern Recognit. Assoc. South Africa Robot. Mechatronics Int. Conf. PRASA-RobMech 2017*, Institute of Electrical and Electronics Engineers Inc., 2017: pp. 203–207. <https://doi.org/10.1109/RoboMech.2017.8261148>.
- [97] S.A. Shevchik, G. Masinelli, C. Kenel, C. Leinenbach, K. Wasmer, Deep learning for in situ and real-time quality monitoring in additive manufacturing using acoustic emission, *IEEE Trans. Ind. Informatics.* 15 (2019) 5194–5203. <https://doi.org/10.1109/TII.2019.2910524>.
- [98] J. Williams, P. Dryburgh, A. Clare, P. Rao, A. Samal, Defect Detection and Monitoring in Metal Additive Manufactured Parts through Deep Learning of Spatially Resolved Acoustic Spectroscopy Signals, *Smart Sustain. Manuf. Syst.* 2 (2018) 20180035. <https://doi.org/10.1520/SSMS20180035>.
- [99] E. Rodriguez, J. Mireles, C.A. Terrazas, D. Espalin, M.A. Perez, R.B. Wicker, Approximation of

- absolute surface temperature measurements of powder bed fusion additive manufacturing technology using in situ infrared thermography, *Addit. Manuf.* 5 (2015) 31–39. <https://doi.org/10.1016/j.addma.2014.12.001>.
- [100] C. Gobert, E.W. Reutzel, J. Petrich, A.R. Nassar, S. Phoha, Application of supervised machine learning for defect detection during metallic powder bed fusion additive manufacturing using high resolution imaging., *Addit. Manuf.* 21 (2018) 517–528. <https://doi.org/10.1016/j.addma.2018.04.005>.
- [101] W. King, A.T. Anderson, R.M. Ferencz, N.E. Hodge, C. Kamath, S.A. Khairallah, Overview of modelling and simulation of metal powder bed fusion process at Lawrence Livermore National Laboratory, *Mater. Sci. Technol. (United Kingdom)*. 31 (2015) 957–968. <https://doi.org/10.1179/1743284714Y.0000000728>.
- [102] T. DebRoy, W. Zhang, J. Turner, S.S. Babu, Building digital twins of 3D printing machines, *Scr. Mater.* 135 (2017) 119–124. <https://doi.org/10.1016/j.scriptamat.2016.12.005>.
- [103] A. V. Oppenheim, J.R. Buck, R.W. Schafer, *Discrete-time Signal Processing. Vol.2*, (2001). http://repository.vnu.edu.vn/handle/VNU_123/34218 (accessed September 3, 2020).
- [104] L. Tan, Multirate DSP, part 1: Upsampling and downsampling | *EE Times*, (n.d.). <https://www.eetimes.com/multirate-dsp-part-1-upsampling-and-downsampling/> (accessed September 3, 2020).
- [105] K. Sayood, *Lossless Compression Handbook*, Elsevier, 2003. <https://doi.org/10.1016/b978-0-12-620861-0.x5000-1>.
- [106] K. Sayood, *Introduction to Data Compression*, Elsevier Inc., 2006. <https://doi.org/10.1016/B978-0-12-620862-7.X5000-7>.
- [107] M. Kassner, Fuzzy hashing helps researchers spot morphing malware - *TechRepublic*, (n.d.). <https://www.techrepublic.com/blog/it-security/fuzzy-hashing-helps-researchers-spot-morphing-malware/> (accessed September 25, 2020).
- [108] 6.3.1. What are Control Charts?, (n.d.). <https://www.itl.nist.gov/div898/handbook/pmc/section3/pmc31.htm> (accessed September 3, 2020).
- [109] M. Riaz, R. Mehmood, R.J.M.M. Does, On the performance of different control charting rules, *Qual. Reliab. Eng. Int.* 27 (2011) 1059–1067. <https://doi.org/10.1002/qre.1195>.
- [110] F.M. Megahed, W.H. Woodall, J.A. Camelio, A review and perspective on control charting with image data, *J. Qual. Technol.* 43 (2011) 83–98. <https://doi.org/10.1080/00224065.2011.11917848>.
- [111] L.S. Zimmer, D.C. Montgomery, G.C. Runger, Guidelines for the application of adaptive control charting schemes, *Int. J. Prod. Res.* 38 (2000) 1977–1992. <https://doi.org/10.1080/002075400188447>.
- [112] S. Ahmad, M. Riaz, S.A. Abbasi, Z. Lin, On efficient median control charting, *J. Chinese Inst. Eng. Trans. Chinese Inst. Eng. A.* 37 (2014) 358–375. <https://doi.org/10.1080/02533839.2013.781794>.
- [113] C.A. ACOSTA-MEJIA, J.J. PIGNATIELLO, B.V. RAO, A comparison of control charting procedures for

- monitoring process dispersion, *IIE Trans.* 31 (1999) 569–579.
<https://doi.org/10.1023/A:1007606524244>.
- [114] D.L. Hall, J. Llinas, An introduction to multisensor data fusion, *Proc. IEEE.* 85 (1997) 6–23.
<https://doi.org/10.1109/5.554205>.
- [115] B. Khaleghi, A. Khamis, F.O. Karray, S.N. Razavi, Multisensor data fusion: A review of the state-of-the-art, *Inf. Fusion.* 14 (2013) 28–44. <https://doi.org/10.1016/j.inffus.2011.08.001>.
- [116] A. Vandone, S. Baraldo, A. Valente, Multisensor data fusion for additive manufacturing process control, *IEEE Robot. Autom. Lett.* 3 (2018) 3279–3284.
<https://doi.org/10.1109/LRA.2018.2851792>.
- [117] M. Grasso, F. Gallina, B.M. Colosimo, Data fusion methods for statistical process monitoring and quality characterization in metal additive manufacturing, in: *Procedia CIRP*, Elsevier B.V., 2018: pp. 103–107. <https://doi.org/10.1016/j.procir.2018.04.045>.
- [118] L. Kong, X. Peng, Y. Chen, P. Wang, M. Xu, Multi-sensor measurement and data fusion technology for manufacturing process monitoring: a literature review, *Int. J. Extrem. Manuf.* 2 (2020) 022001. <https://doi.org/10.1088/2631-7990/ab7ae6>.
- [119] M. Guri, M. Monitz, Y. Mirski, Y. Elovici, BitWhisper: Covert Signaling Channel between Air-Gapped Computers Using Thermal Manipulations, *Proc. Comput. Secur. Found. Work.* 2015-Septe (2015) 276–289. <https://doi.org/10.1109/CSF.2015.26>.
- [120] M. Guri, Y. Solewicz, A. Daidakulov, Y. Elovici, Fansmitter: Acoustic Data Exfiltration from (Speakerless) Air-Gapped Computers, (2016). <http://arxiv.org/abs/1606.05915>.
- [121] A. Kuznetsov, O. Stefanovych, Y. Gorbenko, O. Smirnov, V. Krasnobaev, K. Kuznetsova, Information Hiding Using 3D-Printing Technology, in: *Proc. 2019 10th IEEE Int. Conf. Intell. Data Acquis. Adv. Comput. Syst. Technol. Appl. IDAACS 2019*, Institute of Electrical and Electronics Engineers Inc., 2019: pp. 701–706. <https://doi.org/10.1109/IDAACS.2019.8924352>.
- [122] Y.M. Cheng, C.M. Wang, An adaptive steganographic algorithm for 3D polygonal meshes, in: *Vis. Comput.*, Springer, 2007: pp. 721–732. <https://doi.org/10.1007/s00371-007-0147-2>.
- [123] Y.M. Cheng, C.M. Wang, A high-capacity steganographic approach for 3D polygonal meshes, *Vis. Comput.* 22 (2006) 845–855. <https://doi.org/10.1007/s00371-006-0069-4>.
- [124] A. Bogomjakov, C. Gotsman, M. Isenburg, Distortion-free steganography for polygonal meshes, *Comput. Graph. Forum.* 27 (2008) 637–642. <https://doi.org/10.1111/j.1467-8659.2008.01161.x>.
- [125] Z. Li, S. Beugnon, W. Puech, A.G. Bors, Rethinking the high capacity 3D steganography: Increasing its resistance to steganalysis, in: *Proc. - Int. Conf. Image Process. ICIP*, IEEE Computer Society, 2018: pp. 510–514. <https://doi.org/10.1109/ICIP.2017.8296333>.
- [126] M. Papas, T. Houit, D. Nowrouzezahrai, M. Gross, W. Jarosz, The magic lens: Refractive steganography, *ACM Trans. Graph.* 31 (2012). <https://doi.org/10.1145/2366145.2366205>.
- [127] M. Suzuki, P. Silapasuphakornwong, Y. Takashima, H. Torii, H. Unno, K. Uehira, Technique for protecting copyrights of digital data for 3-D printing, and its application to low infill density objects, *MMEDIA 2016 Eighth Int. Conf. Adv. Multimed. Febr.* 21–25, 2016, Lisbon, Port. (2016) 56–59.

- https://www.thinkmind.org/index.php?view=article&articleid=mmedia_2016_3_40_50043.
- [128] K. Uehira, M. Suzuki, Copyright Protection for 3D Printing by Embedding Information Inside 3D-Printed Objects, in: *Digit. Forensic Watermarking*, 2016: pp. 370–378. <https://doi.org/10.5220/0005342401800185>.
- [129] A. Delmotte, K. Tanaka, H. Kubo, T. Funatomi, Y. Mukaigawa, Blind Watermarking for 3D Printed Objects by Locally Modifying Layer Thickness, *IEEE Trans. Multimed. PP* (2019) 1. <https://doi.org/10.1109/TMM.2019.2962306>.
- [130] J.-U. Hou, D.-G. Kim, S. Choi, H.-K. Lee, 3D Print-Scan Resilient Watermarking Using a Histogram-Based Circular Shift Coding Structure, (2015) 115–121. <https://doi.org/10.1145/2756601.2756607>.
- [131] J.U. Hou, D.G. Kim, H.K. Lee, Blind 3D Mesh Watermarking for 3D Printed Model by Analyzing Layering Artifact, *IEEE Trans. Inf. Forensics Secur.* 12 (2017) 2712–2725. <https://doi.org/10.1109/TIFS.2017.2718482>.
- [132] N. Beckmann, M. Potkonjak, Hardware-based public-key cryptography with public physically unclonable functions, in: *Lect. Notes Comput. Sci. (Including Subser. Lect. Notes Artif. Intell. Lect. Notes Bioinformatics)*, Springer, Berlin, Heidelberg, 2009: pp. 206–220. https://doi.org/10.1007/978-3-642-04431-1_15.
- [133] R. Maes, I. Verbauwhede, Physically unclonable functions: A study on the state of the art and future research directions, in: *Inf. Secur. Cryptogr.*, Springer International Publishing, 2010: pp. 3–37. https://doi.org/10.1007/978-3-642-14452-3_1.
- [134] F. Armknecht, R. Maes, A.R. Sadeghi, B. Sunar, P. Tuyls, Memory leakage-resilient encryption based on physically unclonable functions, in: *Inf. Secur. Cryptogr.*, Springer International Publishing, 2010: pp. 135–164. https://doi.org/10.1007/978-3-642-14452-3_6.
- [135] K.B. Frikken, M. Blanton, M.J. Atallah, Robust authentication using physically unclonable functions, in: *Lect. Notes Comput. Sci. (Including Subser. Lect. Notes Artif. Intell. Lect. Notes Bioinformatics)*, Springer, Berlin, Heidelberg, 2009: pp. 262–277. https://doi.org/10.1007/978-3-642-04474-8_22.
- [136] S. Katzenbeisser, Ü. Kocabaş, V. Rožić, A.R. Sadeghi, I. Verbauwhede, C. Wachsmann, PUFs: Myth, fact or busted? A security evaluation of Physically Unclonable Functions (PUFs) cast in silicon, in: *Lect. Notes Comput. Sci. (Including Subser. Lect. Notes Artif. Intell. Lect. Notes Bioinformatics)*, Springer, Berlin, Heidelberg, 2012: pp. 283–301. https://doi.org/10.1007/978-3-642-33027-8_17.
- [137] A.M. Elliott, O.S. Ivanova, C.B. Williams, T.A. Campbell, Inkjet Printing of Quantum Dots in Photopolymer for Use in Additive Manufacturing of Nanocomposites, *Adv. Eng. Mater.* 15 (2013) n/a-n/a. <https://doi.org/10.1002/adem.201300020>.
- [138] O. Ivanova, A. Elliott, T. Campbell, C.B. Williams, Unclonable security features for additive manufacturing, *Addit. Manuf.* 1 (2014) 24–31. <https://doi.org/10.1016/j.addma.2014.07.001>.
- [139] O. Ivanova, C. Williams, T. Campbell, Additive Manufacturing (AM) and Nanotechnology: Promises and Challenges, ...*An Addit. Manuf.* 19 (2011) 353–364. <http://www.emeraldinsight.com/10.1108/RPJ-12-2011-0127%5Cnhttp://www.academia.edu/download/30862632/2011-56-Ivanova.pdf>.

- [140] C.D. Brubaker, T.M. Frecker, J.R. McBride, K.R. Reid, G.K. Jennings, S.J. Rosenthal, D.E. Adams, Incorporation of fluorescent quantum dots for 3D printing and additive manufacturing applications, *J. Mater. Chem. C* 6 (2018) 7584–7593. <https://doi.org/10.1039/c8tc02024h>.
- [141] S. Flank, A.R. Nassar, T.W. Simpson, N. Valentine, E. Elburn, Fast Authentication of Metal Additive Manufacturing, *3D Print. Addit. Manuf.* 4 (2018) 143–148. <https://doi.org/10.1089/3dp.2017.0018>.
- [142] S.J. Trenfield, H. Xian Tan, A. Awad, A. Buanz, S. Gaisford, A.W. Basit, A. Goyanes, Track-and-Trace: Novel Anti-Counterfeit Measures for 3D Printed Personalised Drug Products using Smart Material Inks, *Int. J. Pharm.* 567 (2019). <https://doi.org/10.1016/j.ijpharm.2019.06.034>.
- [143] A. Dachowicz, S.C. Chaduvula, M. Atallah, J.H. Panchal, Microstructure-Based Counterfeit Detection in Metal Part Manufacturing, *Jom.* 69 (2017) 2390–2396. <https://doi.org/10.1007/s11837-017-2502-8>.
- [144] F. Peng, J. Yang, M. Long, 3-D Printed Object Authentication Based on Printing Noise and Digital Signature, *IEEE Trans. Reliab.* 68 (2019) 342–353. <https://doi.org/10.1109/TR.2018.2869303>.
- [145] F. Peng, J. Yang, Z.X. Lin, M. Long, Source identification of 3D printed objects based on inherent equipment distortion, *Comput. Secur.* 82 (2019) 173–183. <https://doi.org/10.1016/j.cose.2018.12.015>.
- [146] M. Mani, B. Lane, M.A. Donmez, S.C. Feng, S.P. Moylan, R. Fesperman, Measurement Science Needs for Real-time Control of Additive Manufacturing Powder Bed Fusion Processes; NIST Interagency/Internal Report (NISTIR) - 8036, NIST Interagency/Internal Rep. 8036 (2015) 50. <https://doi.org/10.6028/NIST.IR.8036>.
- [147] M. Bisht, N. Ray, F. Verbist, S. Coeck, Correlation of selective laser melting-melt pool events with the tensile properties of Ti-6Al-4V ELI processed by laser powder bed fusion, *Addit. Manuf.* 22 (2018) 302–306. <https://doi.org/10.1016/j.addma.2018.05.004>.

8. Conclusions and broader impacts

8.1. Summary of research

Overall Goal
To improve the security of AM systems by understanding the potential nature of attacks and by developing new techniques for in-situ monitoring to detect build defects as they occur.

The motivation behind this research was to improve the security of additive manufacturing systems by identifying and addressing the unique vulnerabilities that layer-wise fabrication enables. While existing work had explored cyber vulnerabilities and quality control problems in traditional manufacturing, little work had been done to investigate cyber-physical vulnerabilities in AM systems. This gap motivated the initial research into vulnerability in the AM process chain (Chapter 3), with the goal being to raise awareness, to better understand the security (or lack thereof) in AM systems, and to help better understand the problems that needed to be addressed for future work in mitigating vulnerabilities. A second research gap was the need for new in-situ process monitoring techniques to detect internal defects and material property changes in AM parts due to both intrinsic process variation and to the unique CPS vulnerabilities identified earlier. that could detect internal defects as well as material property changes. The use of impedance-based monitoring techniques on AM parts was presented as a new approach for solving these challenges (Chapter 4).

After finding ways of using physical monitoring of AM systems to detect defects caused by cyber-physical vulnerabilities, a second research gap existed in protecting and securing the monitoring systems against attack. While some side-channel monitoring techniques have been proposed for detecting build defects, there remained the need to integrate these approaches into AM systems without exposing them to vulnerabilities. To address this gap, the concept of cyber-physical hashing as a technique for storing quality information in the part toolpath and transmitting it through physical emissions to an air-gapped SCMS was presented and demonstrated for two types of AM systems, metal PBF (Chapter 5) and ME (Chapter 6). After completing each of the previous steps, the final goal was to use the knowledge gain to create a generalizable framework to guide people in the select/design of a SCMS and to use this framework to disseminate the results of this research to the broader AM community. Finally, to help organize and disseminate the information to the broader manufacturing and research communities, a framework for the design of SCMS to protect against part sabotage was created.

A breakdown of the specifics of the contributions and their motivating objectives is given in the following subsections.

8.1.1 Objective 1.

Objective 1

To understand the vulnerabilities of AM systems

A detailed analysis of the AM process chain showed that there were numerous areas where cyber-physical attacks could occur. A discussion of the difficulty of attacking each step was presented as well as some methods for mitigating these attacks. Due to the ubiquitous nature of the .STL file format, it was chosen as a case study to investigate, in detail, the process an attacker might use to design and implement an attack. This case study discussed the types of attacks that could be performed on an .STL file and selected a void attack due to the difficulty of detection and uniqueness to AM. A variety of considerations in implementing a void attack were covered and an algorithm was developed to automatically perform attacks on .STL files.

The work showed that it was possible to use information about the aspect ratio and relative size of the triangles in the .STL file to automatically select void locations in the model that were more likely to cause part failure. By inserting 200 bytes of additional data into the .STL file a 3.3mm³ tetrahedral void was created in an ASTM Standard D638-10 tensile test specimen that resulted in an average reduction of 14% in yield load, from 1085N to 930N, and the strain at failure was reduced from 10.4% to 5.8%. This demonstrated the feasibility of an automated attack algorithm to semi-intelligently attack model files through the insertion of scaled voids.

The human subjects study showed that operators lacked sufficient awareness of the possibility of malicious sabotage attacks. Many operators failed to detect the presence of defects during fabrication, with most only detecting the defect at the point of failure. Even upon discovering the defect, operators did not identify malicious interference as a possible source of the error.

The work achieved both its goal to improve understanding of vulnerabilities in the AM process chain and its goal of raising awareness and interest in the broader AM research community. Since the initial publication of this work, its premises were further validated in work by Beliovetsky, et al. who used the type of attack detailed in the work to sabotage a 3D printed propeller, causing the crash of a quadcopter [32].

8.1.2. Objective 2

Objective 2

To research a non-destructive side-channel monitoring technique that can detect changes to material properties and geometry

Impedance-based monitoring was chosen as a nondestructive monitoring technique due to its low sensor cost, detection of both geometry and material property changes, and ability to travel with a fabricated part. The effectiveness of this approach was studied by investigating Research Question 2. It was hypothesized that by attaching piezo sensors (either directly to the part or indirectly through the use of a fixture) it would be possible to detect internal changes to material properties in AM parts during fabrication.

Impedance-based monitoring was chosen as a nondestructive monitoring technique due to its low sensor cost, detection of both geometry and material property changes, and ability to travel with a fabricated part. This was a new technique for defect detection in AM parts. Two approaches for sensor attachment were presented, direct embedding of the sensor and attachment to a fixture which the parts were fabricated on. Both approaches were shown to be implementable on a material jetting system. The study provided insight into how defects of increasing size affected the impedance signature of parts during fabrication.

The study was able to identify internal defects (caused by change material) with embedded sensors when they affected 2.28% of the part volume (95.6 mm^3) and with steel fixture-based piezos when they affected 1.38% of the part volume (53.8 mm^3). This work showed that in situ detection of small material change defects could be detected; however, there were limitations to the smallest defect that could be detected due to the variation between control samples being used as a baseline.

A secondary benefit of this work was the identification of impedance-based signatures as a potential physically unclonable function for AM parts, due to the natural variation observed between test specimens.

8.1.3. Objective 3

Objective 3

To utilize an air-gapped side-channel measurement system on powder bed fusion that can detect changes to process parameters and toolpath

To further develop the techniques for SCMSs and data transmission using a physical hash, a study was done on a MPBF system. The goal of this study was to investigate how changes to the geometry and material properties of a part affected the side-channel emissions and to use the cyber-physical hash concept to store and transmit information for validating parts. This study used galvo position and thermal emissions from the build as the monitored side-channel. A new technique, a frequency representation of the scan lines, was developed to reduce the quantity of data generated by the SCMS while still allowing the detecting of geometry and process parameter changes. This study resulting in a better understanding on how changes to the process parameters and toolpath affected the scan pattern of the laser and the corresponding frequency representation of the scan lines. This approach converts the scan data into a format that is more easily inspected by an operator and makes it easier to identify both the type and location of defects in the toolpath.

To securely transmit information to the air-gapped SCMS a new “ghost” QR code stack approach was developed. This approach allowed quality information to be inserted in each layer of a build (rather than once for an entire build) and without requiring the physical fabrication of the QR code, eliminating material consumption and reducing the impact on build time. An algorithm for extracting the “ghost” QR code data from the scan lines recorded by the SCMS was created. Using these techniques combined with the physical-hash a demonstrated was presented that was able to detect both internal and external geometry changes as well as process parameter changes of 50% to both laser power and speed. A smaller change of 2% scan speed was visible using the frequency representation approach; however, this was too small to be automatically identified using the current implementation.

8.1.4. Objective 4

Objective 4

To develop techniques for securely communicating a physical-hash to a side-channel measurement system

To expand the use of side-channels as a method for transmitting quality information to an air-gapped monitoring system, acoustic tones were embedded in the toolpath of ME parts. To evaluate the efficacy of using acoustic tones to transmit quality information to an air-gapped SCMS it was necessary to develop a set of methods for evaluating storage capacity, encoding the information in the toolpath, and interpreting the generated acoustic emissions.

An approach for calculating the maximum amount of information that can be stored by toolpath modulation in an arbitrary toolpath file was demonstrated (Section 6.2). This approach allows for the rapid screening of parts for the viability of using toolpath modulation.

Next an algorithm for inserting information into the fill lines through modulation was created. This algorithm used paired tones to both increase the robustness of the signal as well as to avoid changing the time to fabricate each fill line (6.2.3). A second algorithm was created for extracting the tones during fabrication. This approach filtered the acoustic information into signal bins and used a frequency analysis to determine the acoustic intensity in a series of windowed timesteps. A spectral centroid was then calculated and compared against a noise threshold to determine whether a tone had occurred. Analysis of the time and tone information was then used to extract the binary information stored in the toolpath (Section 6.2.4).

Using these techniques, an experimental study was done with a variety of tonal parameters and an approach was presented for selecting settings for improved signal quality (Section 6.3). In the study presented, it was found that 1s tones, at a 20mm/s feed rate, with 20Hz modulation provided the best balance between tone intensity, balance, and ease of detection (Section 6.4). It was also found that the orientation of the scan lines in the printer could affect the quality of the acoustic emissions, with signal quality alternating between layers. Using these settings, the transmission of a message of 160 bits was transmitted without errors.

To determine if this approach compromised the part quality, a study was done on the effect of modulation on tensile strength. The experiment showed that the inclusion of modulation throughout the thickness of a tensile test specimen did not cause a statistically significant effect on part strength. This demonstration showed that toolpath modulation could feasibly be used as a data transmission vector without affecting part quality in a meaningful way.

8.1.5. Objective 5

Objective 5

To develop a framework for designing a side-channel monitoring system to mitigate cyber-physical attacks on AM systems

In keeping with the overall goal of the work, the final objective was to create a framework that could be used to design a SCMS for securing AM systems. IDEAS is a comprehensive framework for this purpose that consists of five main steps 1) **Identify** attack vectors, 2) **Define** side-channels, 3) **Establish** baselines, 4) **Aggregate** data, 5) **Secure** transmission. IDEAS builds off of the work presented in Chapters 3-6 and provides a generalizable approach for implementing these techniques on any AM system.

First, IDEAS expands the discussion on identifying attack vectors, presenting an approach for using this assessment to inform the selection of side-channels. Next, IDEAS provides more detailed guidelines on the selection of side-channels for the monitoring system along with the criteria for their selection. In the third section, IDEAS discusses how to define a baseline dataset for comparison by the SCMS. This discussion expands part the direct comparison approach used in Chapters 4-6 and presents techniques that can be used in a more varied manufacturing environment, such as one producing one off custom parts. The aggregation section covers techniques for reducing data volumes such as by physical hashing, to improve processability and allow quality data to be stored in the model file/toolpath. The final section addresses techniques that can be used to secure the transmission of quality information to the SCMS as well as to help mitigate counterfeiting and IP theft. In total this IDEAS serves as a guide for the design of SCMS as well as a reference on cyber-physical concerns in AM systems.

8.2. Limitations and future work

While the work executed in this research study has expanded the capabilities of AM security, there are several limitations on the methods used in the work presented. These limitations and proposed areas of future work are presented in the following subsections.

8.2.1. Techniques for data aggregation and transmission

A significant challenge of high-resolution side-channel monitoring of AM systems is dealing with the quantity of data generated. To secure an AM system using a SCMS, the data needs to be reduced to a manageable amount while still being able to detect the occurrence of defects and tampering. If quality data is stored in the part toolpath to be transmitted by physical emission to an air-gapped SCMS, the physics of the process significantly limit the amount of information that can be feasibly sent. The frequency representation of scan data and hashing into predefined ranges (Chapter 5, 7) is one approach that can be used to accomplish this. There are however drawbacks and limitations to these approaches. Hashing into predefined ranges when dealing with continuous and variable data sets (such as position) is likely to have values that fall on the edges between ranges. These locations will cause false detection events when natural process variation causes the value to be binned into an adjacent range. While the frequency representation is able to reduce the data volume significantly, it is

susceptible to over detection when there are small changes in the toolpath that affect the signature, but not the part quality. Additional techniques should be investigated to reduce the potential of false detection events while maintaining sufficient protection against attacks. Further, the large variety of AM systems and potential side-channels means that additional techniques are needed for this purpose. Future work to investigate methods such as multivariate control charting and data fusion are needed to improve the usability of the cyber-physical hash approach. Additionally, more work is needed to develop techniques for storing and transmitting the quality information to expand the range of systems and sensors that the air-gapped SCMS approach can be implemented with.

8.2.2. Techniques for establishing SCMS baselines on systems with greater variability

In the work presented in Chapters 4 and 5, defect detection was achieved by the comparison of defective part signatures against the known good signature of the same part from a previous control build. As discussed in Chapter 7 (Section 4), while this approach works well for larger scale, more traditional manufacturing of parts, it is not feasible for a manufacturing environment where each part is unique or where the process parameters are changing significantly between builds or between systems. In these cases, alternative techniques such as machine learning or physics-based models (digital twin) are needed to effectively implement a SCMS. While a brief discussion of these techniques has been given, additional work is needed to develop them in a way that is robust against cyber-physical attacks as typical machine learning approaches may be vulnerable to manipulation by an attacker.

8.2.3. Development and integration of physically unclonable functions

While the SCMS approach developed in this work serves as an effective means of validating physical parts against sabotage attacks during manufacturing, it does not address the security of AM parts in the supply chain after manufacturing. To address counterfeiting and traceability concerns in the AM supply chain, it is necessary to securely link the SCMS validation data to each part instance. While naïve approaches such as serial numbers or barcodes can be affixed to AM parts as unique identifiers, these approaches are not secure and may be replicated or altered by a third party. To ensure that the end-user receives the part that was validated by the SCMS, it is necessary to develop a unique, irreplicable signature that can travel with a part and be linked back to the quality information. To this end, physically unclonable functions have been presented as a solution that can uniquely link the physical part to the digital information in a way that cannot be replicated. Further work is needed to develop these methods for AM systems and to integrate them into the SCMS approach to bridge the physical-digital gap.

8.2.4. Obfuscation of side-channels

While the work presented in this dissertation focused on the use of side-channels as an approach for mitigating cyber-physical sabotage attacks on AM parts, other research has use side-channels as a technique for stealing valuable IP. One approach for mitigating this risk is through the use of side-channel obfuscation (e.g. by adding noise into the system to hide the signal). While these techniques can increase the difficulty for an attacker attempting to steal IP, they would also serve to

make it more difficult to use these same side-channels as a means of defect detection. Further work is need, both to investigate techniques of obfuscating side-channels to prevent information leakage, and in selecting side-channels for the monitoring system that are difficult for an attack to access. For example, acoustic emissions on a material extrusion system can be used to monitor movement; however, these emissions can be easily detected by bringing a smartphone or other microphone into nearby proximity, increasing the risk of IP theft. Linear encoders placed inside the printer can also monitor movement, but are not easily accessed outside of the system. To improve the overall security a speaker might be used to generate noise to obfuscate the easily detected acoustic emissions, while a set of linear encoders were used by the SCMS to validate against sabotage attacks. These encoders would be unaffected by the additional acoustic noise and would allow the system to improve security against both sabotage attacks and IP theft attacks.

Publications

Five core papers have been generated in accomplishing this work. A list of related journal publications and conference proceedings is also provided.

- Cyber-physical vulnerabilities in additive manufacturing systems: A case study attack on the. STL file with human subjects, LD Sturm, CB Williams, JA Camelio, J White, R Parker, Journal of Manufacturing Systems 44, 154-164
- In situ monitoring of material jetting additive manufacturing process via impedance based measurements, LD Sturm, MI Albakri, PA Tarazaga, CB Williams, Additive Manufacturing 28, 456-463
- In-situ Detection of Build Tampering in Metal Additive Manufacturing Using a Cyber-physical Hash, LD Sturm, CB Williams, Additive Manufacturing (In preparation)
- ATTACH: Additive Toolpath Transmission of an Acoustic Cyber-Physical Hash, LD Sturm, N Raeker-Jordan, Additive Manufacturing (In preparation)
- IDEAS (Identify, Define, Establish, Aggregate, Secure): A cyber-physical framework for securing additive manufacturing systems using physical side-channels, LD Sturm, CB Williams, Additive Manufacturing (In preparation)

Other publications related to this work are:

- A physical hash for preventing and detecting cyber-physical attacks in additive manufacturing systems J Brandman, L Sturm, J White, C Williams Journal of Manufacturing Systems 56, 202-212
- Internal porosity detection in additively manufactured parts via electromechanical impedance measurements C Tenney, MI Albakri, J Kubalak, LD Sturm, CB Williams, PA Tarazaga ASME 2017 Conference on Smart Materials, Adaptive Structures and Intelligent ...
- Impedance-based non-destructive evaluation of additively manufactured parts MI Albakri, LD Sturm, CB Williams, PA Tarazaga Rapid Prototyping Journal

- Taxonomies for Reasoning About Cyber-physical Attacks in IoT-based Manufacturing Systems. Y Pan, J White, DC Schmidt, A Elhabashy, L Sturm, J Camelio, C Williams International Journal of Interactive Multimedia & Artificial Intelligence 4 (3)
- In-situ detection of build defects in additive manufacturing via impedance-based monitoring L Sturm, M Albakri, CB Williams, P Tarazaga 27th Annual International Solid Freeform Fabrication Symposium—An Additive ...
- Non-destructive evaluation of additively manufactured parts via impedance-based monitoring M Albakri, L Sturm, CB Williams, PA Tarazaga Solid Freeform Fabrication Symposium, 1475-1490 SFF 2014
- Cyber-physical vulnerabilities in additive manufacturing systems L Sturm, C Williams, J Camelio, J White, R Parker

Research Contributions

In accomplishing this work, several contributions have been made towards improving the security of AM systems. Techniques for identifying attacks, performing in situ monitoring, analyzing data, and transmitting information to air-gapped monitoring systems have been demonstrated. A fundamental understanding of the cyber-physical security issues facing AM systems was developed through the study and development of these techniques.

This work has brought major contributions to the field of cyber-physical security for AM systems by i) highlighting vulnerabilities in existing systems ii) developing new systems which can mitigate these vulnerabilities and iii) creating techniques which can be used to secure these new systems against attacks while maintaining system usability.

The follow major contributions have been achieved by this work:

- A detailed analysis of vulnerabilities in the AM process chain
- An increased understanding on how attacks against AM systems may be designed and on how operator awareness of cyber-physical issues needs to be improved
- The development of a new impedance-based technique for in situ monitoring of AM systems
- An understanding of the effects of defects and attachment modality on the in situ impedance response of AM parts
- An understanding of how geometry and process parameters changes affect side-channel emissions on a MPBF system
- A technique for reducing the volume of meltpool/scan data generated by a SCMS on a MPBF system while still maintaining the ability to detect geometry and process parameter changes on fabricated parts both through automatic measures and visual inspection by an operator
- A technique, the “ghost” QR code for storing information in the model file/toolpath of a part and for transmitting this information to an air-gapped SCMS through the use of physical emissions without requiring the physical fabrication of additional parts
- A technique for embedding quality information into the g-code of a ME part without altering the build time or finished properties of the part and a method for using a microphone to extract the information from the acoustic emission generated by the system

- A greater understanding of the effects of toolpath changes on the acoustic emissions of a ME AM system and how these changes affect transmissibility of information from the toolpath to a SCMS
- A framework for the design of SCMS to secure AM systems against malicious cyber-physical attacks designed to sabotage parts.

Broader impacts

- This work has helped to define and launch a new area of research into cyber-physical security for additive manufacturing systems and more broadly for advance manufacturing.
- This work has drawn attention to, promoted dialogue, and encouraged further investigation of cyber-physical vulnerabilities in additive manufacturing systems. Both in the research community and the manufacturing sector a broader awareness of the potential for cyber-physical attacks has been achieved.
- Dissemination of this work has help to increase the awareness of the AM workforce to the potential of cyber-physical attacks on part quality.
- Helped to improve the security of AM parts against malicious cyber-attacks improving trust and reliability in the supply chain by improving methods for preventing sabotage attacks with dangerous physical consequences.
- Created an approach that can be used to secure dated AM systems that cannot be secured using modern cyber-focused solutions due to limitations imposed by machine hardware. The addresses a need that cyber-focused solutions may neglect.
- Increased understanding of the necessity of integrating security into quality control systems to ensure quality, security, and traceability throughout the AM process chain.

References

- [1] Griffor, Edward R, Greer, Chris, Wollman, David A, Burns, Martin J, Framework for Cyber-Physical Systems: Volume 1, Overview, NIST Spec. Publ. 1 (2017) 1201–1500. <https://doi.org/10.6028/NIST.SP.1500-201>.
- [2] International Data Corporation, New IDC Smart Home Device Tracker Forecasts Solid Growth for Connected Devices in Key Smart Home Categories, (2018). <https://www.idc.com/getdoc.jsp?containerId=prUS43701518> (accessed October 11, 2018).
- [3] S. Sridhar, A. Hahn, M. Govindarasu, Cyber-physical system security for the electric power grid, *Proc. IEEE*. 100 (2012) 210–224. <https://doi.org/10.1109/JPROC.2011.2165269>.
- [4] R. Charette, This Car Runs on Code, *IEEE Spectr.* (2009). <https://spectrum.ieee.org/transportation/systems/this-car-runs-on-code> (accessed October 11, 2018).
- [5] F. Pasqualetti, F. Dörfler, F. Bullo, Attack Detection and Identification in Cyber-Physical Systems, *IEEE Trans. Automat. Contr.* 58 (2013) 2715–2729. <https://doi.org/10.1109/TAC.2013.2266831>.
- [6] L. Sha, S. Gopalakrishnan, X. Liu, Q. Wang, Cyber-physical systems: A new frontier, *Mach. Learn. Cyber Trust Secur. Privacy, Reliab.* (2009) 3–13. https://doi.org/10.1007/978-0-387-88735-7_1.
- [7] R. (Raj) Rajkumar, I. Lee, L. Sha, J. Stankovic, Cyber-physical systems: The Next Computing Revolution, in: *Proc. 47th Des. Autom. Conf. - DAC '10*, ACM Press, New York, New York, USA, 2010: p. 731. <https://doi.org/10.1145/1837274.1837461>.
- [8] R. Baheti, H. Gill, Cyber-physical Systems, *Impact Control Technol.* (2011) 161–166.
- [9] D. Evans, The Internet of Things: How the Next Evolution of the Internet is Changing Everything, *CISCO White Pap.* 1 (2011) 1–11. http://www.cisco.com/c/dam/en_us/about/ac79/docs/innov/IoT_IBSG_0411FINAL.pdf.
- [10] J. Bayuk, D. Cavit, E. Guerrino, J. Mahony, B. McDowell, W. Nelson, R. Snelvel, P. Staarfanger, *Malware Risks and Mitigation Report*, BITS Financ. Serv. Roundtable, Washington, DC. (2011).
- [11] M. Watin-Augouard, Prospective Analysis on Trends in Cybercrime from 2011 to 2020, *Natl. Gendarm.* (2011). <http://www.mcafee.com/us/resources/white-papers/wp-trends-in-cybercrime-2011-2020.pdf>.
- [12] N. Falliere, L.O. Murchu, E. Chien, *W32.Stuxnet Dossier*, 4 (2011) 1–69.
- [13] C. Li, A. Raghunathan, N.K. Jha, Hijacking an insulin pump: Security attacks and defenses for a diabetes therapy system, in: *E-Health Netw. Appl. Serv. (Healthcom)*, 2011 13th IEEE Int. Conf., 2011: pp. 150–156. <https://doi.org/10.1109/HEALTH.2011.6026732>.

- [14] K. (Wired) Zetter, A Cyberattack Has Caused Confirmed Physical Damage for the Second Time Ever | WIRED, (2015). <http://www.wired.com/2015/01/german-steel-mill-hack-destruction/> (accessed January 7, 2016).
- [15] B. für S. in der I. BSI, Die Lage der in Deutschland 2011, Informationstechnik. (2011).
- [16] National Defence Industrial Association's, Cybersecurity for Advanced Manufacturing, 2014.
- [17] A. Cherepanov, WIN32/INDUSTROYER: A new threat for industrial control systems, ESET. (2017). https://www.welivesecurity.com/wp-content/uploads/2017/06/Win32_Industroyer.pdf.
- [18] R. Lee, M. Assante, T. Conway, Analysis of the Cyber Attack on the Ukrainian Power Grid, SANS Ind. Control Syst. Secur. Blog. (2016) 1–26.
- [19] L. Wang, M. Törngren, M. Onori, Current status and advancement of cyber-physical systems in manufacturing, *J. Manuf. Syst.* 37 (2015) 517–527. <https://doi.org/10.1016/j.jmsy.2015.04.008>.
- [20] B. Dworschak, H. Zaiser, Competences for cyber-physical systems in manufacturing – first findings and scenarios, *Procedia CIRP.* 25 (2014) 345–350. <https://doi.org/10.1016/j.procir.2014.10.048>.
- [21] L.J. Wells, J.A. Camelio, C.B. Williams, J. White, Cyber-physical security challenges in manufacturing systems, *Manuf. Lett.* 2 (2014) 74–77. <https://doi.org/10.1016/j.mfglet.2014.01.005>.
- [22] H. Turner, B. Amos, J. White, J. Camelio, C. Williams, Bad parts: Are our manufacturing systems at risk of silent cyber-attacks, *IEEE Secur. Priv.* (2015) 40–47. <https://doi.org/10.1109/MSP.2015.60>.
- [23] Cybersecurity for Manufacturing Networks a White Paper, 2017. <http://www.ndia.org/-/media/sites/ndia/divisions/working-groups/cfam/ndia-cfam-2017-white-paper-20171023.ashx?la=en>.
- [24] K. V Wong, A. Hernandez, A Review of Additive Manufacturing, *ISRN Mech. Eng.* 2012 (2012) 1–10. <https://doi.org/10.5402/2012/208760>.
- [25] A. Bellini, S. Güçeri, Mechanical characterization of parts fabricated using fused deposition modeling, *Rapid Prototyp. J.* 9 (2003) 252–264. <https://doi.org/10.1108/13552540310489631>.
- [26] J. Choi, Y. Chang, Characteristics of laser aided direct metal/material deposition process for tool steel, *Int. J. Mach. Tools Manuf.* 45 (2005) 597–607. <https://doi.org/10.1016/j.ijmachtools.2004.08.014>.
- [27] NIST, Measurement Science Roadmap for Metal-based Additive Manufacturing, 2013. http://www.nist.gov/el/isd/upload/NISTAdd_Mfg_Report_FINAL-2.pdf.

- [28] D. Howie, High powered Trent XWB-97, (2015) 12–15. <http://www.rolls-royce.com/media/insights/simon-burr.aspx>.
- [29] Stratasys, FAA-Approved Air Duct for ‘Flying Eye Hospital’ Produced in Just Days,’ (2015). <http://blog.stratasys.com/2015/03/05/3d-printed-air-duct-flying-eye-hospital/> (accessed January 31, 2017).
- [30] L.D. Sturm, C.B. Williams, J.A. Camelio, J. White, R. Parker, Cyber-Physical Vulnerabilities in Additive Manufacturing Systems, *Solid Free. Fabr. Symp.* (2014) 951–963.
- [31] B. Berman, 3-D printing: The new industrial revolution, *Bus. Horiz.* 55 (2012) 155–162. <https://doi.org/10.1016/j.bushor.2011.11.003>.
- [32] S. Belikovetsky, M. Yampolskiy, J. Toh, Y. Elovici, dr0wned - Cyber-Physical Attack with Additive Manufacturing, *CoRR*. abs/1609.0 (2016). <http://arxiv.org/abs/1609.00133>.
- [33] S.E. Zeltmann, N. Gupta, N.G. Tsoutsos, M. Maniatakos, J. Rajendran, R. Karri, Manufacturing and Security Challenges in 3D Printing, *JOM*. 68 (2016) 1872–1881. <https://doi.org/10.1007/s11837-016-1937-7>.
- [34] C. Song, F. Lin, Z. Ba, K. Ren, C. Zhou, W. Xu, My Smartphone Knows What You Print: Exploring Smartphone-based Side-channel Attacks Against 3D Printers, in: *Proc. 2016 ACM SIGSAC Conf. Comput. Commun. Secur.*, ACM, New York, NY, USA, 2016: pp. 895–907. <https://doi.org/10.1145/2976749.2978300>.
- [35] M. Yampolskiy, L. Schutzle, U. Vaidya, A. Yasinsac, Security Challenges of Additive Manufacturing with Metals and Alloys, in: M. Rice, S. Sheno (Eds.), *Crit. Infrastruct. Prot. IX 9th IFIP 11.10 Int. Conf. ICCIP 2015*, Arlington, VA, USA, March 16-18, 2015, Revis. Sel. Pap., Springer International Publishing, Cham, 2015: pp. 169–183.
- [36] S. Goldenberg, J. Brown, J. Haid, J. Ezzard, 3D opportunity and cyber risk management, *Deloitte Univ. Press.* (2016). <https://doi.org/10.1016/j.jval.2017.05.018>.
- [37] F. Chen, G. Mac, N. Gupta, Security features embedded in computer aided design (CAD) solid models for additive manufacturing, *Mater. Des.* 128 (2017) 182–194. <https://doi.org/10.1016/j.matdes.2017.04.078>.
- [38] O. Ivanova, A. Elliott, T. Campbell, C.B. Williams, Unclonable security features for additive manufacturing, *Addit. Manuf.* 1 (2014) 24–31. <https://doi.org/10.1016/j.addma.2014.07.001>.
- [39] D. Li, A.S. Nair, S.K. Nayar, C. Zheng, AirCode: Unobtrusive Physical Tags for Digital Fabrication, (2017). <https://doi.org/10.1145/3126594.3126635>.
- [40] C. Wei, Z. Sun, Y. Huang, L. Li, Embedding anti-counterfeiting features in metallic components via multiple material additive manufacturing, *Addit. Manuf.* 24 (2018) 1–12. <https://doi.org/10.1016/j.addma.2018.09.003>.
- [41] S.R. Chhetri, M. Abdullah, A. Faruque, Side-Channels of Cyber-Physical Systems : Case

- Study in Additive Manufacturing, IEEE Des. Test. PP (2017) 1.
<https://doi.org/10.1109/MDAT.2017.2682225>.
- [42] S.B. Moore, J. Gatlin, S. Belikovetsky, M. Yampolskiy, W.E. King, Y. Elovici, Power Consumption-based Detection of Sabotage Attacks in Additive Manufacturing, (2017) 1–19. <http://arxiv.org/abs/1709.01822>.
- [43] S.R. Chhetri, A. Canedo, M.A. Al Faruque, KCAD: Kinetic Cyber-Attack Detection Method for Cyber-Physical Additive Manufacturing Systems, Proc. 35th Int. Conf. Comput. Des. - ICCAD '16. (2016) 1–8. <https://doi.org/10.1145/2966986.2967050>.
- [44] S. Belikovetsky, Y. Solewicz, M. Yampolskiy, J. Toh, Y. Elovici, Digital Audio Signature for 3D Printing Integrity, IEEE Trans. Inf. Forensics Secur. PP (2018) 1.
<https://doi.org/10.1109/TIFS.2018.2851584>.
- [45] S.A. Shevchik, C. Kenel, C. Leinenbach, K. Wasmer, Acoustic emission for in situ quality monitoring in additive manufacturing using spectral convolutional neural networks, Addit. Manuf. 21 (2018) 598–604. <https://doi.org/10.1016/j.addma.2017.11.012>.
- [46] M.A. Al Faruque, S.R. Chhetri, A. Canedo, J. Wan, Forensics of thermal side-channel in additive manufacturing systems, CECS Tech. Report# 16-01. (2016).
- [47] C. Bayens, T. Le, L. Garcia, C. Bayens, L. Garcia, See No Evil , Hear No Evil , Feel No Evil , Print No Evil ? Malicious Fill Patterns Detection in Additive Manufacturing, USENIX Secur. (2017).
- [48] P.K. Rao, J. (Peter) Liu, D. Roberson, Z. (James) Kong, C. Williams, Online Real-Time Quality Monitoring in Additive Manufacturing Processes Using Heterogeneous Sensors, J. Manuf. Sci. Eng. 137 (2015) 61007–61012. <https://doi.org/10.1115/1.4029823>.