Leave Your Phone at the Door: Side Channels that Reveal Factory Floor Secrets

Avesta Hojjati[1,*], Anku Adhikari[1,2,*], Katarina Struckmann[1,*],
Edward J. Chou[1], Thi Ngoc Tho Nguyen[2], Kushagra Madan[1],
Marianne S. Winslett[1,2], Carl A. Gunter[1], William P. King[1]
[1]University of Illinois at Urbana-Champaign, USA [2]Advanced Digital Sciences Center, Singapore
{hojjati2,aadhikr2,struckm2,ejchou2,kushagra,winslett,cgunter,wpk}@illinois.edu
tho.nguyen@adsc.com.sg

ABSTRACT

From pencils to commercial aircraft, every man-made object must be designed and manufactured. When it is cheaper or easier to steal a design or a manufacturing process specification than to invent one's own, the incentive for theft is present. As more and more manufacturing data comes online, incidents of such theft are increasing.

In this paper, we present a side-channel attack on manufacturing equipment that reveals both the form of a product and its manufacturing process, i.e., exactly how it is made. In the attack, a human deliberately or accidentally places an attack-enabled phone close to the equipment or makes or receives a phone call on any phone nearby. The phone executing the attack records audio and, optionally, magnetometer data. We present a method of reconstructing the product's form and manufacturing process from the captured data, based on machine learning, signal processing, and human assistance.

We demonstrate the attack on a 3D printer and a CNC mill, each with its own acoustic signature, and discuss the commonalities in the sensor data captured for these two different machines. We compare the quality of the data captured with a variety of smartphone models. Capturing data from the 3D printer, we reproduce the form and process information of objects previously unknown to the reconstructors. On average, our accuracy is within 1 mm in reconstructing the length of a line segment in a fabricated object's shape and within 1 degree in determining an angle in a fabricated object's shape.

We conclude with recommendations for defending against these attacks.

Categories and Subject Descriptors

K.6.5 [**MANAGEMENT OF COMPUTING AND INFORMATION SYSTEMS**]: Security and Protection

Keywords

Data Security for Manufacturing; Side Channels; Cyber-Physical Systems

1. INTRODUCTION

Hackers have noticed the large amount of valuable information available in the cyber-physical systems on manufacturing factory floors. In addition to straightforward data theft, adversaries can, in theory, take advantage of simple yet effective side-channel attacks based on electromagnetic leaks, acoustic emissions, timing information, light emission, and power consumption [5, 10, 12, 21, 24, 27]. The leaked information can be used to compromise systems and to obtain or infer sensitive data. For example, researchers have successfully partially compromised Diffie-Hellman exponents, factored RSA keys, and broken other cryptosystems by measuring the amount of time required to perform private key operations [11, 17, 18]. Defending against side-channel attacks requires a level of security more advanced and more comprehensive than updating an operating system or installing security patches. Despite the efficacy of firewalls and anti-virus software, manufacturers currently have no effective way to protect against information leakage from their factory floor equipment.

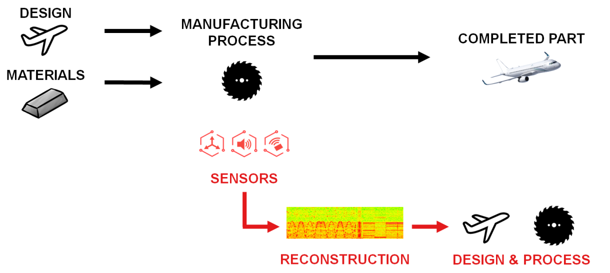

Figure 1: The high level system and attack model. Designs and raw materials are the inputs to a manufacturing process that produces completed parts. By placing phone sensors near the manufacturing process and analyzing the data they collect, the side-channel attack reconstructs the design and the manufacturing process.

In a modern factory, nearly everyone on the manufacturing floor carries a smartphone or similar electronic device. These devices are programmable and come with a growing number of embedded sensors, including a microphone, accelerometer, magnetometer, gyroscope, GPS, and camera. These sensors can capture side-channel information regardless of the level of information technology or security sophistication on the factory floor and inside the manufacturing equipment.

We present a novel attack in which a phone's sensors are deliberately or inadvertently used to capture sensitive information

[*]These authors contributed equally.

CCS'16, October 24 - 28, 2016, Vienna, Austria

ISBN 978-1-4503-4139-4/16/10...$15.00
DOI: http://dx.doi.org/10.1145/2976749.2978323

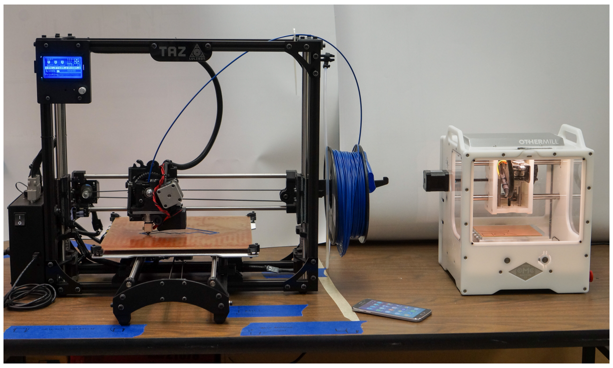

Figure 2: The attack setup. For the phone attack, a phone placed on the same table as a 3D printer (left) or CNC mill (right) records the readings of its sensors. For the phone call attack, an attacker on the other end of a call records the call's audio.

from manufacturing equipment, as shown in Figure 1. We capture the relevant sensor data by deliberately or accidentally placing an attack-enabled phone close to, on top of, or inside a piece of manufacturing equipment while the machinery is fabricating the target object. Figure 2 shows this setup. Alternatively, the relevant audio can be recorded by deliberately or accidentally making or receiving a phone call while standing next to the machinery or by installing malware on any other nearby device that has a microphone.

We provide methods that use the captured data to reconstruct a model of the object being manufactured along with its manufacturing process parameters. We demonstrate the attack on both additive and subtractive manufacturing using a 3D printer and a CNC mill. We demonstrate the reconstruction process with a 3D printer and discuss ways to reduce the attack's effectiveness.

Contributions. We outline the paper's contributions below:

- *New techniques.* We show that the data captured by acoustic and magnetic sensors embedded in a phone can be used to identify specific manufacturing equipment and manufacturing processes, including reconstructing manufactured objects and reproducing the processes used to make them.

- *New understanding.* We demonstrate the feasibility of applying side-channel attacks to manufacturing equipment: in particular, 3D printers and CNC mills. The fundamentally different operating modes of these two types of manufacturing equipment indicate that the attack may be broadly applicable across many types of manufacturing equipment.

- *Implementation and evaluation of reconstruction method.* We provide a method for reconstructing manufactured objects and the processes used to make them; the method is based on machine learning and signal processing. We show that the method accurately reconstructs previously unseen objects.

The paper is organized as follows: Section 2 provides background information and discusses the motivations of potential attackers. Section 3 describes the attack model and reconstruction method. Section 4 provides experimental results, and Section 5 offers recommendations for defending against the attacks and raising the cost of reconstruction.

2. BACKGROUND

Traditional high-value discrete manufacturing relies heavily on subtractive processes: equipment is used to cut, chip, and grind away excess material to form the desired product. Interest is high in the potential for new additive manufacturing processes, which deposit material layer by layer to form objects. Our attack and reconstruction methods target both additive and subtractive manufacturing, represented in our experiments by a 3D printer and CNC mill, respectively.

The manufacturing sector has a rich history of research on obtaining information about a manufacturing process from its acoustic emissions. Recordings have been used to judge parameters including tool wear, tool breakage, chatter, chip formation mechanism, material removal regime, sheet metal material hardness, sheet metal thickness, and the identity of the metal or alloy being machined [6, 9, 13, 16, 19, 20, 22]. Our reconstruction methods use acoustic information for less benign purposes.

Cyberattacks on the manufacturing sector typically fall into one of three categories: theft of intellectual property or processes, disruption of manufacturing operations, or sabotage of products or reputation [28]. These cyberattacks are already widespread: in 2014, 21% of manufacturers reported a loss of intellectual property (IP) [2]. These observed losses may be the tip of the iceberg, as 69% of all 2012 data breaches were carried out within a few hours, but 64% of breaches took months or years to detect [23]. Further, the number of manufacturing cyberattacks is growing fast: the Industrial Control Systems Cyber Emergency Response Team (ICS-CERT), operated by the US Department of Homeland Security, responded to 50% more incidents in the manufacturing sector in 2015 than in 2014 [14, 15].

IP theft is by far the most common motive of attackers, who can target product design information, manufacturing process information, or both. The advantage of stealing design information is clear, but many manufacturers' competitive advantage largely lies in the fact that they know how to manufacture a given design better, faster, or cheaper than their competitors do. Process information may include the details of what materials are used and what machines are used and in what order, plus all the settings of those machines: which tool head was used, its rotation rate, the material feed rate, and so forth.

When a phone illicitly records data on the factory floor, its owner could be intentionally carrying out corporate espionage, or she could be an unwitting dupe with a compromised application or even the innocent maker or receiver of a phone call at an ill-advised moment. In these latter cases, she may have been targeted by a third party such as a rival manufacturer, or swept up in a large net cast by a well-financed backer of economic espionage such as a nation-state. For example, a nation-state hacker might be seeking to increase the competitiveness of its manufacturing sector, gain the ability to manufacture objects viewed as important for the national interest, or learn about its rivals' capabilities and activities. For example, such motivations may have been behind the theft of the design for Lockheed Martin's US F-35 Lightning II fighter jet, stolen by hackers allegedly supported by the Chinese government[1]. The US and Israeli governments allegedly unleashed STUXNET, which targeted the programmable logic controllers of Iran's nuclear centrifuges, causing them to self-destruct. Allegedly, the Chinese government has financed the large-scale theft of industrial IP [1] and the Iranian government has sponsored IT intrusions overseas [3]. The US and Israeli governments have been attributed as potential sources

[1] "Chinese hackers stole F-35 fighter jet blueprints in Pentagon hack, Edward Snowden documents claim": http://goo.gl/Vnvbs2.

of the Flame malware, apparently designed to increase situational awareness of Iran's technical capabilities and activities [26]. In addition to gathering files likely to contain technical information, Flame collected data from the sensors of the devices it infected.

Our approach to factory floor snooping leverages the fact that sensor spyware could spread to anyone's phone and export the data it captures for subsequent analysis and reconstruction of manufacturing activities. While Flame targeted Windows PCs, similar malware can be constructed for phone applications. For example, Cai et al. [7] highlight the capabilities of modern mobile devices for snooping on users by sniffing their smartphone's sensors. For example, the Shedun Android malware provides a framework for automatically downloading and installing undesired new applications and for serving potentially malicious adware; security researchers consider Shedun nearly impossible to remove completely. DroidKungFu, targeted at users in China, offers similar capabilities for the automatic installation of malicious new applications. With over a thousand new infections per day as of this writing, malware like Shedun and DroidKungFu provides a channel to reach factory employees. The malware may even be present when a phone is first purchased; Indian phone manufacturer Gionee has been accused of this[2]. Once a malicious app has been installed, its recording function could be activated by a geofence around the factory, and could run in the background of another app with appropriate permissions, such as a game.

Al Faruque et al. [4] recognized that thermal side channels can be used to infer activities taking place inside a 3D printer. Closer to our work, they also investigated the possibility of attacking manufacturing machinery via audio recordings. They placed a microphone close to additive manufacturing equipment to record fabrication runs, then used machine learning to reconstruct the low-level instructions (G-code) used to manufacture the object, with an accuracy of 89.72% in reproducing the as-designed object's perimeter. While Al Faruque et al. used a high-quality microphone located in a specific location in a controlled environment, our work uses ordinary mobile phones that may be located anywhere near the machine or in the user's hand and can target machinery located in any environment, including public fabrication labs. We employ different reconstruction methods from Al Faruque et al., and our experiments show that we reconstruct perimeters more accurately; however, we also suggest different measures of accuracy that we believe to be more revealing. A final difference is that because G-code is quite low level, reproducing the same object on a different model or type of machine requires nontrivial extra work to rewrite the G-code. For that reason, we provide a higher-level reconstruction suitable for translation into G-code for a variety of machines.

3. RECONSTRUCTION

Different fabrication machines have different process parameters. For example, a grinding wheel can run at different speeds. The wheel could be touching the object being made, or could be away from it, e.g., while repositioning the wheel to a different location on the object. The object could be moving past the wheel at different rates. The same machine could use grinding wheels with different levels of grit. To accurately reproduce a manufacturing process, we need to specify the values for all of its parameters.

As no single paper can reconstruct all process parameters for all major types of equipment, we focus on key parameters related to the location of the tool head with respect to the object being fabricated and its direction of travel. These location and direction parameters are important for both additive and subtractive manufacturing, and must be specified to control machines as disparate as a 3D printer and a CNC mill. Further, while prior research in the manufacturing community has concluded that many aspects of fabrication have inherent acoustic signatures, no signatures that specify these parameters have been established in previous work.

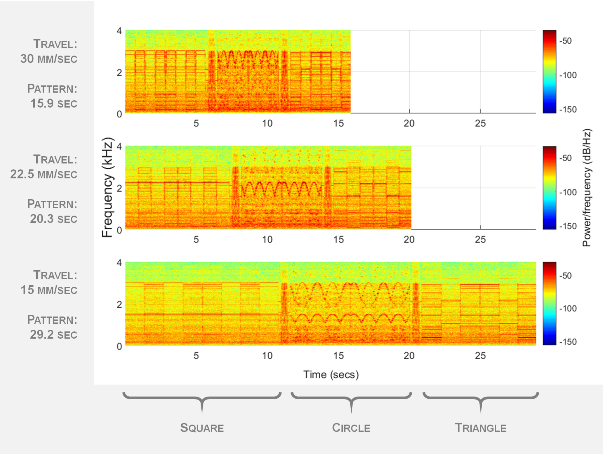

Figure 3: Audio magnitude spectrograms of a 3D printer making the same three geometric primitive objects, a square, circle, and triangle, at three different feed rates (15, 22.5, and 30 mm/sec). The increase in head travel speed changes the spectrogram in a systematic way.

We describe tool head location and direction with respect to the platform of a machine, which defines an implicit XY plane and an associated Z axis. Different machines have different constraints in traversing this space, and reconstruction can take advantage of these constraints to simplify the task. For example, a typical 3D printer builds up an object in horizontal layers. At any given layer, the printer head moves in an XY plane and can trace any angle in that plane with respect to the X axis. The printer slowly works its way up the Z axis, emitting a characteristic sound from this movement. Further, in an object with multiple layers, each layer must either overlap the previous layer or have its own support material, so a layer's shape is constrained by the previous layer. Likewise, a subtractive manufacturing operation generally removes material adjacent to material it has already removed, and subtractive methods typically also work in layers. We take advantage of this layer-focused machine behavior by restricting our attention to the XY plane for a fixed value of Z, i.e., a given layer. Our reference and training data and validation experiments use nearly-planar objects: 3D prints two layers thick and shallow cuts with the mill.

Any planar figure that can be manufactured by machines like CNC mills and 3D printers can be specified as a sequence of tool head movements to be made at particular angles to the X axis for particular straight-line distances (with curves described by short tangential segments). We reconstruct both these angles and the distances. As it can be hard to visualize the tool head trajectories and phone placements we discuss, readers may wish to refer to the videos of our printer and mill in action, and an example reconstruction session, at https://goo.gl/FijZ9T.

The 3D printer head can travel at different speeds (feed rates). As shown in Figure 3, the printer's audio signature for a particular angle changes in a systematic way as the feed rate changes. Zooming in on these high-resolution figures, we see that the pattern in the figure's three spectrograms compresses in time and shifts up in frequency as

[2] "Wenn der Spion in der Hosentasche steckt" (If the Spy Is in Your Pants Pocket), *Die Welt*, 12 October 2014.

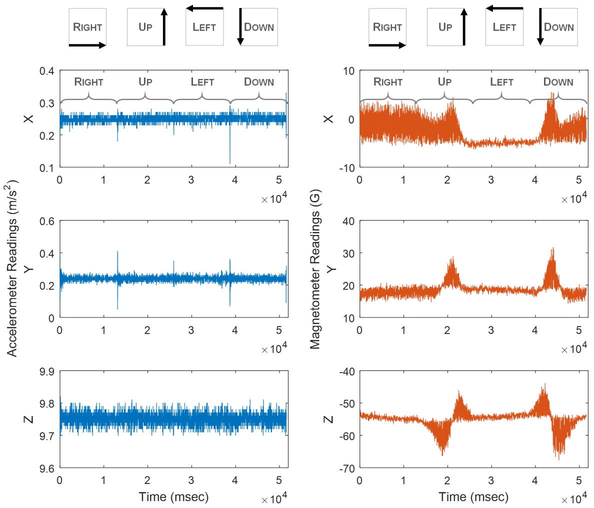

Figure 4: Raw sensor data from the three axes of the accelerometer and magnetometer while 3D printing a square. The readings vary predictably, providing additional information not fully captured in audio recordings. We found readings from the magnetometer to be more accurate than those from the accelerometer in explaining tool head movement.

the feed rate increases, though the human eye quickly recognizes that the high-level pattern is unaffected. For this reason, we focus on the printer's default feed rate of 30 mm/sec.

Figures 4 and 5 show example data from the phone accelerometer, magnetometer, and microphone. The reconstruction method uses audio and magnetometer data when both are available and just audio otherwise. When multiple reconstructed objects are consistent with the results produced by signal processing and machine learning, the reconstruction method uses a search process, domain constraints, and human assistance to rule out unlikely and impossible reconstructions.

We reconstruct the angle of travel of the machine tool head by comparing its audio to examples in a prerecorded reference library. We found that angles that are just a few degrees apart have very different audio signatures, so the recordings for the library need to include each angle that might be used to fabricate a target object. These recordings could be obtained from a similar machine model that the attacker plans to use to fabricate stolen designs or processes. Alternatively, as discussed in Section 4, the necessary calibration pattern could be hidden in the design of an object fabricated on a machine belonging to the victim or a third party, and recorded in an attack launched specifically to gather that information for the library. The result is a library of audio signatures that the machine produces for each angle of movement.

To build an audio signature library, we first use a cell phone to record the sound that the machine produces when it moves along potential angle of interest. In this study, we used 1 degree of resolution, recording angles from 0 to 359 degrees. The sampling rate of the cell phone recording was 44100 samples/second, the default.

The second step is to transform the recorded audio from the time domain to the frequency domain using a short time Fourier transform (STFT), and then produce a magnitude spectrogram. We used the Matlab function *spectrogram* for this purpose, with a Hann (Hanning) window of length 2048 samples, an overlap of 25% between successive windows, and 2048 frequency points. The use of overlapping Hanning windows is a standard technique in audio processing that helps to reduce the noise in the signal by smoothing it out. The use of 2048 frequency points gave sufficient resolution for reconstruction.

When recordings are made in a manufacturing environment, the audio contains background noise whose energy spreads across all frequency bands, and in general the background noise energy tends to decrease as the frequency increases. We found that for reconstruction to succeed, it is important to reduce this background noise, especially at low frequencies, to emphasize the useful content of the audio signal. Thus the third step in building the audio library is to perform noise normalization in the frequency domain. We estimate the background noise covariance matrix Rnn based on a portion of the recording when the machine is idle. Assuming noise is uncorrelated across frequencies, we use Rnn to normalize the signal spectrogram as follows:

$$X_{white} = diag\left(\frac{1}{\sqrt{diag(Rnn)+\epsilon}}\right) \times X,$$

where X is the magnitude spectrogram for the recording, and $\epsilon = 1 \times e^{-8}$ is a constant used to avoid dividing by zero. To further reduce the effect of background noise and unwanted interference in the library, we average all the frames of the signal of the same angle along the time dimension of the spectrogram. The result is 360 frames, corresponding to the 360 angles of movement illustrated in Figure 6. Each frame is the audio signature of the machine head at a particular angle across frequencies.

Figure 6 shows that most of the information needed to decide what angle the machine is moving at is concentrated at low frequency bands. Further, signal artifacts such as aliasing are visible at high frequency bands. Therefore we select only frequencies below a cutoff frequency f_c for further processing, saving the results in the reference library. We also record the domain constraints specific to that machine, such as its platform size and the value of f_c.

After these signal processing steps, the audio of each 3D printer angle a appears very similar to that of the three other angles created by mirroring the given angle in each quadrant of the plane ($\pm a, 180 \pm a$); this introduces ambiguity into reconstruction. Similarly, each mill angle sounds like 15 other angles, produced by mirroring across the X and Y axes and the ±45 and ±135 degree lines. We suspect that more sophisticated audio signal processing techniques that can pick out the secondary tones visible in our spectrograms (and audible to a keen ear) can be used to tell these angles apart, but that remains for future work, and we rely on two other disambiguation techniques described later: magnetometer data and a search process that exploits domain constraints. For simplicity, the discussion that follows is written as though the library contains just one reference angle audioclip for each set of ambiguous angles, e.g., only first quarter angles for the 3D printer. However, for ease of incorporating new domain constraints and information from other sensors, our actual implementation retains reference data for all angles.

With the reference library in hand and a target object to reconstruct from its fabrication audio, we begin by cleaning up the audio by applying the same first four signal processing steps as for reference audioclips: define overlapping frames, produce a magnitude spectrogram, normalize with respect to background noise, and retain only the relevant frequency band. Then we find the most likely angle for each frame of the cleaned-up audio by comparing it to all of the reference library's angle frames and finding the one it is most correlated with. More precisely, we use a custom matched filter

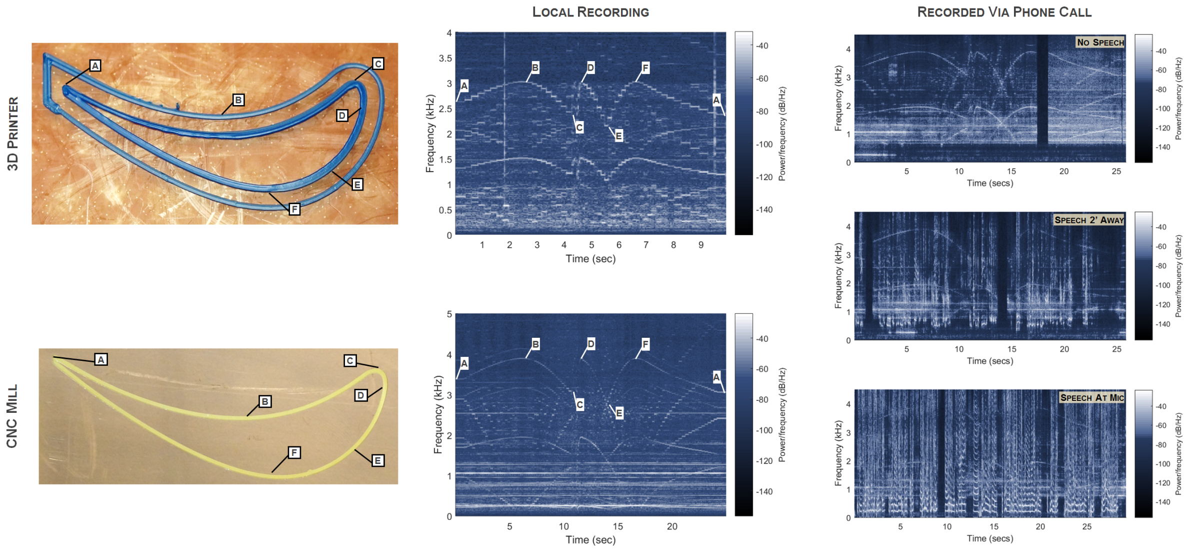

Figure 5: A comparison of the spectrograms from the 3D printer recorded locally (center top), the CNC mill recorded locally (center bottom), and the CNC mill recorded on the other end of a phone call (right column). The right column compares three phone call recordings: no speech (right top), speech 2 ft from the recording microphone (right center), and speech directly at the receiving microphone (right bottom). The annotations illustrate how the spectrograms correspond to the fabrication processes. The spectrograms from recordings of the fabrication of the same turbine blade shape display a trace in a consistent shape, even across different machines; each contains sufficient information to reconstruct the shape. Additionally, while speech overlaps with the frequencies that indicate machining, the traces are not fully obscured.

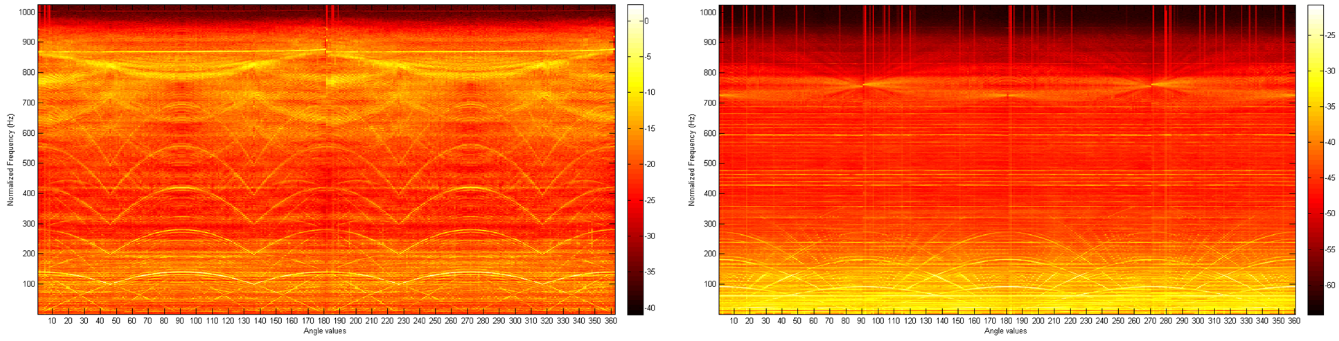

Figure 6: Example magnitude spectrograms for audio data: juxtaposed magnitude spectrograms of 360 different angles of machine head travel with a 3D printer (left) and CNC mill (right). This data is used in a reference library during reconstruction.

function[3] to compute the correlation between the audio frame and each of the reference library frames; this is a standard technique for comparing two audio samples. The result is one value for each combination of a target audio frame and a library angle. For each target audio frame, we select the library angle with the highest cross correlation value for that frame: this is our best guess angle for that moment of the fabrication. Then we present the results to the user, as shown in Figure 7's screenshot of our interactive reconstruction framework. The screen shows the spectrogram of the target audio, aligned along the time axis with a matched filter visualization, where the height of each match head indicates the reference library angle selected for that frame in the audio. When magnetometer data is not available, this completes the automated signal processing phase of the reconstruction.

Next, the search phase begins, with optional human assistance to steer the framework's search process. To prepare for this role, a user requires only brief training in how to recognize changes of angles and the start/stop of tool work in audio magnitude spectrograms.

In our current framework implementation, the user has two tasks. First is to click on the points in the audio at which the tool head changes its angle; these points can be seen quite easily as the edges of the vertical bars in the spectrogram. The identified points divide the manufactured object into a series of straight-line tool head runs, which we call *segments*. From watching a video of the 3D printer, framework users learned several constraints that were useful to them during segmentation. The first two are generic to 3D printing: the printer begins and ends each run with a particular movement sequence; and the printer fabricates a 1-layer bounding box around all the closed figures it will subsequently construct during the run. Users also observed a third constraint specific to the 2-layer objects we were building: after fabricating the first layer of an object, the

[3] For an explanation of matched filters, see goo.gl/Nrjojv.

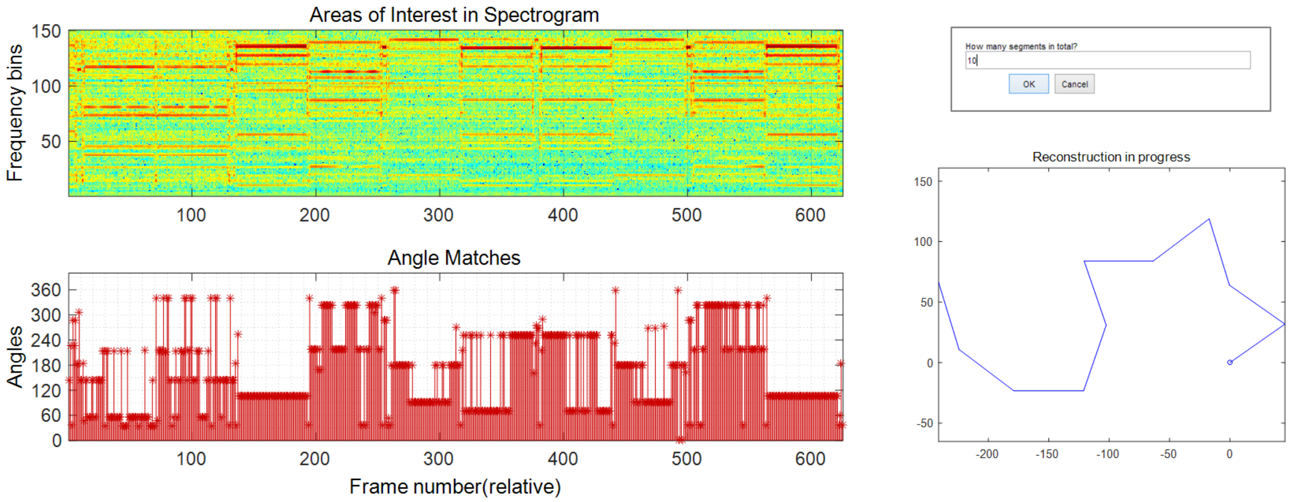

Figure 7: In-progress reconstruction of a 3D-printed star, with spectrogram above and matched filter beneath. Changes of tool head direction are visible as yellowish vertical bars in the spectrogram. The audio has been segmented and the search for a reconstruction that satisfies domain constraints is underway. The reconstruction at the lower right will be automatically rejected because it is not a closed figure; the first erroneous angle is off by 180 degrees, and its mirror image will be considered later in the search.

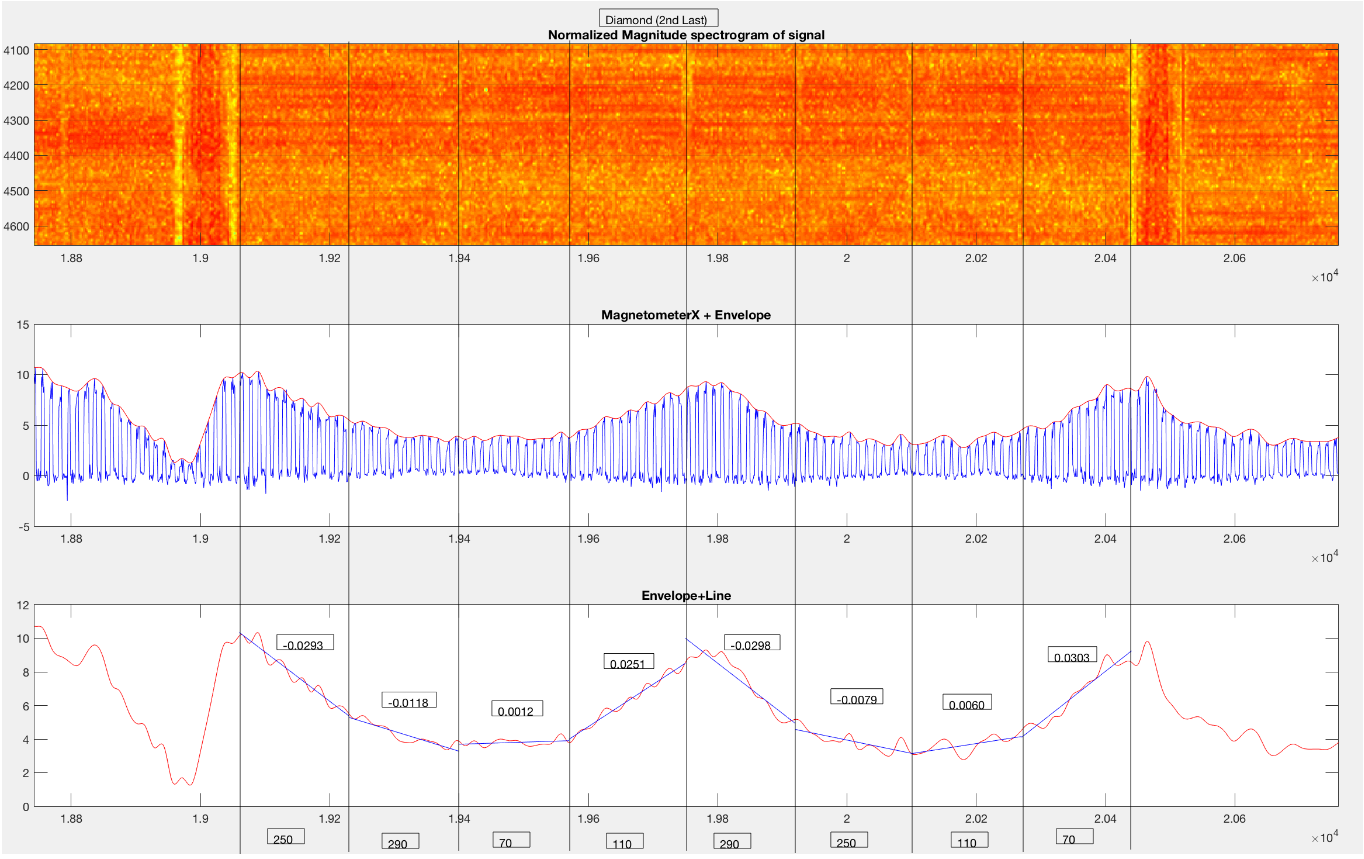

Figure 8: Segmented audio magnitude spectrogram, magnetometer signal in the X dimension, and regression lines for each segment. The 3D printer is fabricating a 2-layer diamond shape, and the sign of the slope of each regression line indicates whether the angle for that segment lies above or below the printer platform's X axis. The magnitude of a segment's peaks indicates how far forward on the platform that segment's fabricated line lies, which can be helpful for establishing a canonical orientation for the object being manufactured.

tool head retraces its path in reverse to construct the second layer. Together, these three constraints helped the user identify the first and last segment of each fabricated object.

The duration of each segment multiplied by the machine's feed rate gives the physical length of each linear segment in the reconstructed shape, so the accuracy of segmentation affects the accuracy of the final reconstruction. In our experiments, we did not focus on trying to get segment lengths exactly right. We expect that signal processing techniques can be used to automate segmentation in the future, and may be more accurate than a human.

The interactive framework automatically shows the user-selected segment boundaries superimposed on the spectrogram and its accompanying matched filter timeline. The most common matched filter head height in a segment is the best guess angle for that segment. For example, the third and tenth segments of the audio in Figure 7 will have only one suggested angle, while the sixth and seventh will have two.

The user's second task is to provide optional guidance to speed up the search process. For example, suppose that the first choice angles for all segments do not produce a reconstruction. While the search process can automatically identify the segments with the most uncertainty in the matches, and automatically identify the second most likely angle for each, the user can also guide this process by clicking on the matched filter head heights (angles) that she would like the search to consider next. This flexibility is particularly useful when the framework does not have a full set of domain constraints. For example, we implemented a domain constraint that the tool head must move at different angles in adjacent segments. Without this constraint, the framework might assign the same angles to segments 6 and 7 in Figure 7, but the user could click to force the use of different angles.

Each matched filter head height may correspond to several different mirrored angles in the reference library. Thus in our experiments, a k-sided 3D-printed object has an audio-only reconstruction search space of roughly 4^k potential objects. (When magnetometer readings are available, the search space shrinks to roughly 2^k, as explained below.) Fortunately, manufacturing domain constraints allow us to prune away most of the search space. We implemented two constraints generic to 3D printing. The first constraint is that a layer should not cross over itself. More precisely, each segment in a layer should intersect exactly two other segments, one at each of its endpoints, except that the first and last can intersect either one or two other segments. The second constraint is that at the end of a segment, the printer head should change its angle of travel rather than continuing in a straight line or (unless it is the end of a layer) doubling back on itself. We implemented a third constraint specific to the kinds of objects we were building: each object layer should form a closed figure in the plane. The framework automatically explores the search space not eliminated by these constraints and displays the resulting reconstructions to the user, who can accept or reject them. If unhappy with all the reconstructions shown, the user can click on additional matched filter head heights for a segment, so that additional angles will be considered.

In theory, a phone's magnetometers could tell us whether each nearby motor in a machine is accelerating, decelerating, or holding steady, and for how long; from that information we could determine the exact path the tool head traces. In practice, however, we only found magnetometer data useful for reliably distinguishing between angles a and $-a$ for the 3D printer. This means that when we have both magnetometer and audio data for a fabrication run on the 3D printer, each angle appears very similar to only *one* other angle, its mirror image across the X axis. In other words, de facto, with the phone located near the corner of the printer, its magnetometer registers the machine platform's forward and back movement during fabrication, but does not pick up a signal from the machine's other motors and movements.

To illustrate the disambiguation, consider the example three-dimensional magnetometer signal in Figure 8, which was recorded with the phone lying flat near the corner of the printer. If the peaks of the magnsetometer signal in the dimension most closely aligned with the machine platform's Y axis are decreasing in a segment where the machine head traverses angle a, then the peaks of the magnetometer signal in that dimension will be increasing as the machine head traverses angle $-a$. More precisely, our algorithm for processing magnetometer data uses the magnetometer's dimension with the strongest signal overall (X for the 3D printer, Z for the mill). Then for that dimension in each segment, the algorithm identifies the peaks in the magnetometer's magnetic field strength measurements. Our implementation uses the Matlab function *envelope*$(x, np, ' peak')$, which uses spline interpolation over local maxima separated by at least np samples; $np = 8$ worked well for our phone's magnetometer data. Then we find the regression line that minimizes the peak points' average squared distance to the line. Our implementation uses the Matlab function *polyfit*$(x, y, 1)$, which returns the coefficients for a line $p(x)$ that is a best fit (in a least-squares sense) for the data in y. If the slope of the resulting line is negative, then the angle is between 0 and 180 degrees. If the slope is positive, then the angle is between 180 and 360 degrees. Intuitively, a positive slope means that the platform of the printer is moving toward the phone. A negative slope means that the platform is moving away from the phone. A slope very close to zero (with respect to the amplitude of the signal) means that the platform is not moving closer to or further away from the phone. We found that background noise normalization was not helpful in analyzing the magnetometer data for either the printer or the mill.

We found that the sound associated with travel at a particular angle to the X axis did not depend on where the tool head was located on the Y axis, so audio for the reference library and target objects could be recorded with the tool head anywhere on the machine platform and the phone anywhere nearby. In contrast, magnetometer readings fall off with the cube of the distance to the source, and we found that the phone needs to be within a foot of the platform to pick up useful data. Magnetometer readings are also sensitive to the phone's orientation; flipping the phone around essentially reverses its reading. Still, as long as the phone's orientation remains approximately the same while recording, its magnetometer readings will reliably distinguish between a and $-a$. More generally, we expect that the best way to use magnetometer data will vary greatly for different types of machines, depending on the configuration of their motors and where it is natural to set down a phone. For example, if our phone had picked up on the head's side-to-side movement only, then we would have been able to distinguish between a and $180 - a$, rather than registering only the movement of the platform forward and back. In fact, when we printed an entire platform full of diamond shapes, the phone's magnetometer did seem to register additional information (perhaps generated at the tool head) while a diamond was printed in the extreme corner of the platform, very close to the phone. When magnetometer data is not available (e.g., an audio-only recording, the phone changing orientation during recording as the user moves around, or the magnetometer being too far away to pick up readings), each angle still appears similar to three others, its mirror images across the X and Y axes.

We found that analysis of the phone's accelerometer data did not improve the accuracy of reconstruction, so we used only audio and magnetometer data in the experiments. The accelerometer data does indicate the times at which the tool head changes direction at a segment intersection, which can be incorporated in the future to segment the data from the other sensors. Further analysis may reveal additional information the reconstruction method could utilize.

For a particular reconstruction task, additional domain or product constraints may be useful. For example, a reconstructed object should not extend beyond the machine's platform. If we know the general shape of the item being manufactured, such as a turbine blade, this context can inform the reconstruction process. If all constraints are met, we show the reconstructed object to the user; otherwise we move on to the next candidate reconstruction.

4. EXPERIMENTAL RESULTS

Setup. We conducted experiments with the Lulzbot Taz 5 3D printer and Other Machine Co. Othermill CNC mill shown in Figure 2, hereafter referred to as the "printer" and the "mill." The X axis of each machine is controlled by a stationary stepper motor that drives a carriage on which the tooling (the printer's extruder and the mill's spindle) rides. The Y axis of each machine is controlled by a second stationary stepper motor that moves the platform. The printer's Z axis is controlled by two stepper motors, one on each end, that raise and lower the full X axis. The mill's Z axis is controlled by a single stepper motor that controls the height of the spindle relative to the X carriage, which remains fixed in height. All experiments in this section used the printer's default feed rate, 30mm/second.

We built an Android app that monitors and records the sensor data on a phone and used it to record the audio and magnetometer data used in reconstruction. The audio is collected at a 44100 Hz sampling rate.

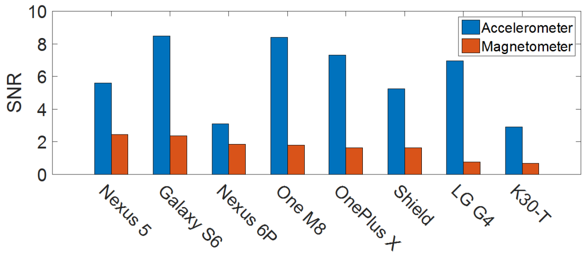

Figure 9: The signal-to-noise ratio (SNR) of the accelerometer Y axis and magnetometer Z axis readings on different devices. The Samsung Galaxy S6, with the highest SNR for its accelerometer and second-highest SNR for its magnetometer, is the best overall.

4.1 Data Quality with Different Devices

To evaluate the quality of the sensor data produced by different devices, we installed the recording app on seven smartphones and one tablet, listed in Table 1. To ensure a fair comparison across devices, we placed each device with its lower right corner 2.5 centimeters from the rear left corner of the printer. Then we enabled the app's recording function while the printer fabricated a simple geometric shape resembling a trapezoid. The printed object and its process parameters were identical in each trial.

We compared the signal to noise ratio (SNR) of the accelerometer and magnetometer in each device. The Samsung Galaxy S6 performed the best overall, with the highest accelerometer SNR and second-highest magnetometer SNR. Surprisingly, the Nexus 6P, the newest model, had the second-lowest accelerometer SNR. The full results are shown in Figure 9. The placement of the sensors inside each device varies. While our other experiments suggest that for the accelerometer this variation will have a negligible impact on the reconstruction quality compared to the impact of the variation inherent in the sensor, the magnetometer works at a much shorter range, and it may be affected. In evaluating the other sensors' data quality, we focused on the Galaxy S6.

Manufacturer	Model	OS	Form
HTC	One M8	Android 6.0	Phone
Huawei	Nexus 6P	Android 6.0	Phone
LG	G4	Android 5.1	Phone
LG	Nexus 5	Android 5.1	Phone
OnePlus	X	Android 5.1	Phone
Samsung	Galaxy S6	Android 5.1	Phone
Lenovo	K30-T	Android 4.4	Phone
Nvidia	Shield	Android 5.1	Tablet

Table 1: Compared Devices

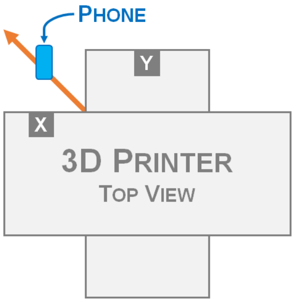

Figure 10: The setup for testing the quality of sensor recordings at different distances from the printer. Starting from the rear left corner, approximately midway between the X and Y motors, the phone was moved away at a 135° angle.

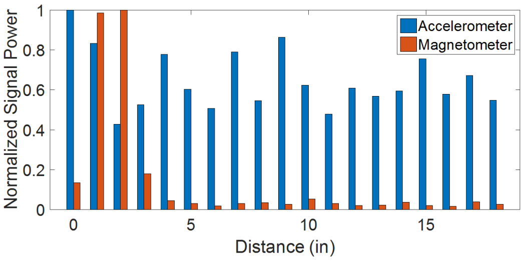

Figure 11: The signal power of a phone's sensor recordings at different locations from the machine. While the magnetometer readings drop off sharply with distance, the accelerometer readings are strong at all locations on the table.

4.2 Data Quality at Different Locations

To determine the impact of distance on data quality, we compared readings from the Samsung Galaxy S6 at different locations relative to the 3D printer. Beginning at the rear left corner, the phone was used to record the same fabrication activity (a simple 45-degree line) as its distance from the printer was incremented by 2.5 cm. We moved the phone away from the printer in a line approximately 135 degrees from horizontal so that it remained approximately equidistant from the X and Y motors, as illustrated in Figure 10.

We calculated the signal power of the accelerometer and magnetometer readings at each location as a measurement of the effect of distance, as shown in Figure 11. The accelerometer, which was measuring the movement of the table the printer was placed on, had strong readings at all distances from 0 to 18 inches. The readings decreased slightly with distance, but the output was clear at all distances. In contrast, the magnetometer was measuring a magnetic field, and the strength of a magnetic field drops with the distance cubed. The magnetometer readings were unusable at distances greater than 4 inches. This limitation affects our ability to use magnetometer data to distinguish between mirrored angles.

4.3 Data from Different Machines

While the data generated by the 3D printer and mill is similar enough that our reconstruction methods can be applied to both, the data is also distinct. More generally, each machine has a unique signature, and types of machines will have distinct sounds corresponding to their manufacturing processes. For example, the sound of a mill's spindle spinning and tool cutting is absent in audio from a 3D printer. The spindle noise alone is sufficient to distinguish between the printer and the mill used in our experiments. Additionally, each motor of a machine has a signature. Though they are nominally identical, depending on the configuration of the machine, each motor moves a different amount of weight. This distinction is already apparent within our printer and mill: the machine's X and Y motors are nominally identical but display distinct signatures. The frequencies at which the motors emit noise, as a function of the work they are doing, allows us to differentiate between different machine models. While this technique would also work to distinguish between different models of the same machine type—say, two printers instead of a printer and a mill—a more complex technique would be needed to distinguish between two same-model machines.

To substantiate these claims, we compared recordings of the 3D printer with recordings from the mill. The mill's and printer's X and Y movements are driven by similar motors in similar configurations, and we found that the mill's movements exhibit a clear, consistent, and uniquely identifiable audio signature, analogous to that of the 3D printer. We demonstrate this signature in Figure 5, comparing the spectrograms of the same turbine blade shape made on the printer and on the mill. The trace is shifted in frequency on the two machines, since each motor on each machine has its own signature, but the two traces exhibit the same pattern.

These results suggest that a recording of a simple calibration pattern is all that is needed to train either of our reconstruction methods on most 3D printers and desktop mills, as well as other types of subtractive manufacturing methods operated by stepper motors. This calibration pattern could be hidden in the interior of an object and designed to look like typical 3D printer infill, or hidden in the toolpath of a subtractive manufacturing operation. If the attacker asks the operator to manufacture this object and makes a recording, she now has all the information necessary to reconstruct objects and machining conditions from that machine.

4.4 Data Quality in a Phone Call Attack

The previous sections focused on the case where phone sensor data was captured by a malicious phone application. If the data was instead captured during a phone call, the audio signal will have been altered by the phone's noise reduction. Conveniently, the key audio frequencies of factory floor machinery tend to lie in the same range as the human voice, so the phone's noise reduction does not simply remove the signal.

We tested the phone call attack on both the printer and mill. The results from each recording, while noisier, are clear and consistent with the audio recorded directly on a phone located near the machine. For example, the frequency magnitude spectrogram in Figure 5 shows that the same pattern is visible whether a recording is made next to the mill or recorded through a phone call. We also tested the phone call attack while people were speaking. Figure 5 shows the difference in the audio when a person is speaking two feet from the device next to the machine and speaking directly into the microphone of the device far from the machine; even though the speech overlaps the information-rich regions of the spectrogram, the trace is not obscured completely and the shape is still clearly visible.

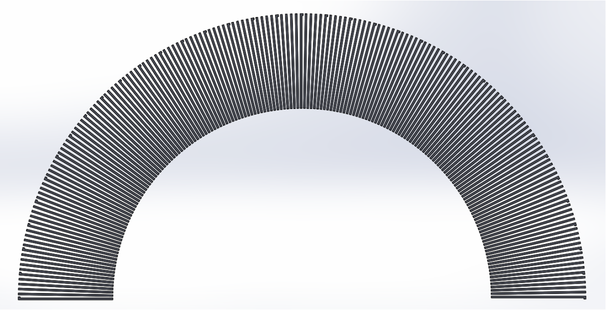

Figure 12: The fan shape fabricated to provide example audio for the reference library.

While reconstruction following a phone call attack must rely on audio only, this method greatly broadens the scope of the attack. Capturing information from multiple sensors at once requires an appropriate app to be present on the phone; in contrast, the phone call attack allows any phone to capture factory audio with no prior preparation beyond the attacker being prepared to record the call on the remote end. More generally, an audio-only attack can be executed using any device with a microphone, which expands the attack to not only phones but also tablets, laptops, and other computers, either through malware or by recording a voice-over-IP call.

4.5 Accuracy of Reconstruction

All training and test data for the reconstruction methods was recorded using a Samsung Galaxy S6 placed within 4 inches of the printer, i.e., close enough to collect usable magnetometer data. We did not try to place the phone in the exact same position for each run.

We built the interactive framework using Matlab, Adobe Audition, and Python. For both the printer and the mill, we constructed the reference library from the audio of one pass of the machine head over the left-hand half of the 2-layer planar fan shape shown in Figure 12; this shape has 360 different angles of machine head travel in each mirrored half. A spectrogram of the resulting library is shown in Figure 6. As mentioned earlier, to prune the reconstruction search space for the signal processing method, we implemented three domain constraints in the signal processing interactive framework: the reconstructed object layer should be within .5 feed units of being a closed planar object with no mid-segment self-crossings, and there should be a change of angle at the end of each segment.

The framework's user was an EE Master's student with no prior experience with 3D printers or mills and no prior experience in audio analysis. She prepared for her reconstruction tasks by watching and listening to videos of a 3D printer and mill traversing a square, circle, triangle, and turbine blade outline (see https://goo.gl/FijZ9T). From examining the resulting spectrograms, she learned to recognize the visual signatures in the spectrogram corresponding to the start and stop of the machine's work on an object, the vertical bars

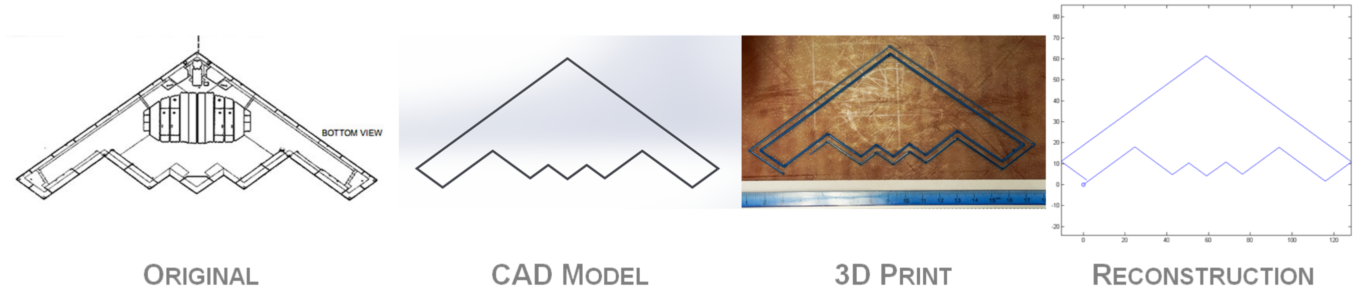

Figure 13: Results from reconstructing an airplane. The original design model (far left); the CAD design (center left); the fabricated object (center right); and the reconstructed model (far right).

where the machine head changed its direction of travel, and tool head up/down. She practiced by using the framework to reconstruct the triangle and square made by the 3D printer and the mill; due to angle ambiguities, her results included the actual objects as well as their mirror reflections, which we find acceptable for all reconstructions. We anticipate that crowdsourced workers with the same training as our user could segment the data equally well. To test this claim, a second-year computer science undergraduate with no prior experience with signal processing or audio also experimented with segmenting, and found it easy.

We tested the reconstruction effectiveness on the 3D-printed outlines of a star, an airplane, and a gun. Our user had never seen any of these designs before.

The resulting reconstructions, along with the original design, are shown in Figures 13 and 14, and we discuss their accuracy below.

The user reconstructed all three objects, plus their mirror reflections. However, she described the airplane, which was a B2 stealth bomber, as a "fish mouth" and was quite dissatisfied with the result, even revisiting the matched filter diagram to consider second-choice values for angles and look for other potential reconstructions. In other words, our user successfully reconstructed a mystery object, even though the mystery object was not something she could recognize in real life so she was not assisted by context. For the gun, our user reconstructed the original object, but did not a priori prefer it to variants produced by mirroring the angles in the very short segments of the gun.

Figure 15 shows the length of each side and degree of each angle in the original and reconstructed objects. On average, angles are within a degree of the actual.

For the mill, we correctly reconstructed all angles in a square (not shown in the figure); for a triangle, we reconstructed one angle exactly and the other two with 1 degree of error. The other objects in the figure are from the 3D printer, and have similar accuracy in angle reconstruction.

As we currently perform segmentation manually, error in computing segment lengths is independent of the segment lengths and reflects human judgment and focus. The framework's user was off by approximately 1mm on average in indicating segment lengths for the plane and the star. This error could probably be reduced, as our main concern was getting the angles right and slightly-too-short segments do allow fully accurate angle reconstruction. We were highly accurate in perimeter reconstruction, the measure used by [4]. However, this measure is quite sensitive to the segment lengths of the target object, as the total error in manual segmentation is directly proportional to the number of segments, rather than to their lengths.

Even with fully-automated reconstruction, non-expert users could provide useful guidance during reconstruction. For example, we could hear someone clicking a pen in one of the machine recordings. The click was quite distinct from ordinary machine sounds, and we knew that it did not indicate a machine event such as an angle change; but signal processing techniques or a regression model might have been fooled by it. More generally, when machine learning and signal processing fail due to irrelevant background or foreground noise, a human may be able to salvage the reconstruction.

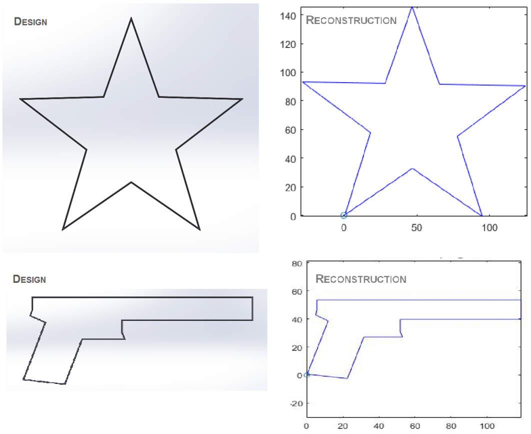

Figure 14: Original design and reconstructed shape, for the star (top) and gun (bottom).

5. RECOMMENDATIONS

We designed and tested a defense that obfuscates the acoustic emissions from manufacturing equipment by playing recordings during production. Since noise reduction has been studied extensively[4], instead of playing a random signal, we chose to play recordings of variations of the part being produced that have small dimensional deviations from it. The attacker would still be able to determine the general shape of item being manufactured, which may provide situational awareness about a manufacturer's capabilities and the current activities in their factory. On the other hand, obfuscation can make it harder for the attacker to separate the target audio stream from the others and reconstruct the object's exact dimensions or process parameters. Because often the small details of the process or design are exactly the information that an attacker would like to obtain, it is worthwhile to make them harder to identify. For example, in high-value manufacturing, there may be a hundred wrong ways to make an object and one way to make it correctly. Obfuscation could

[4] For an introduction to the topic, see goo.gl/IFnJ0r.

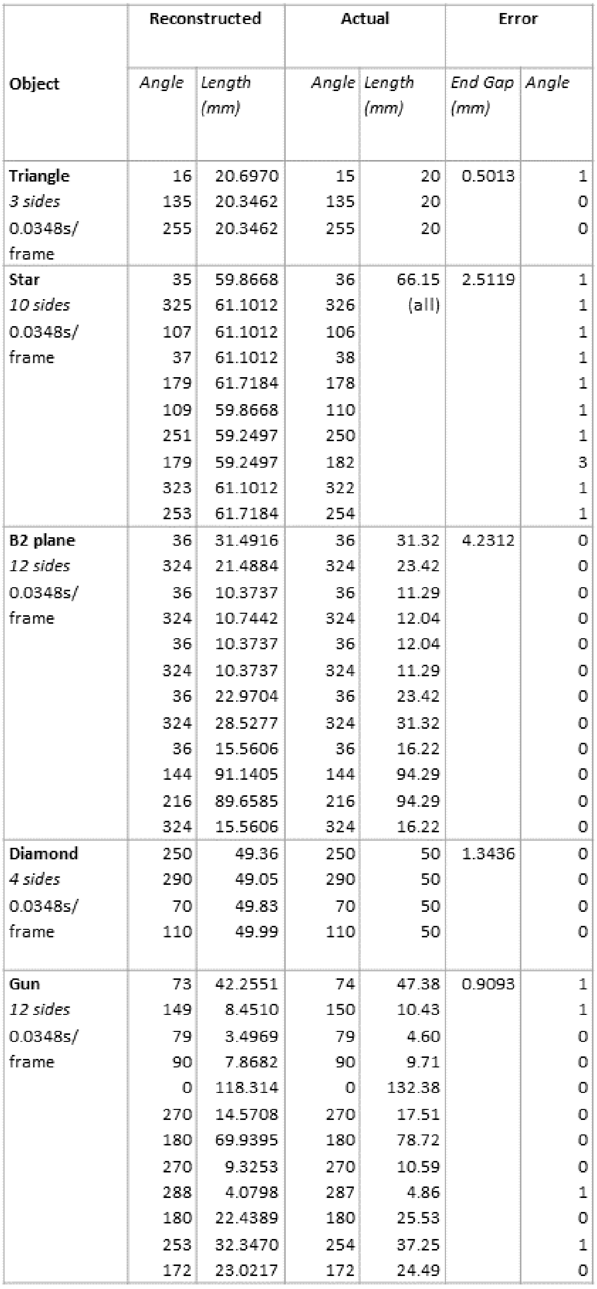

Object	Reconstructed		Actual		Error	
	Angle	*Length (mm)*	*Angle*	*Length (mm)*	*End Gap (mm)*	*Angle*
Triangle	16	20.6970	15	20	0.5013	1
3 sides	135	20.3462	135	20		0
0.0348s/	255	20.3462	255	20		0
frame						
Star	35	59.8668	36	66.15	2.5119	1
10 sides	325	61.1012	326	(all)		1
0.0348s/	107	61.1012	106			1
frame	37	61.1012	38			1
	179	61.7184	178			1
	109	59.8668	110			1
	251	59.2497	250			1
	179	59.2497	182			3
	323	61.1012	322			1
	253	61.7184	254			1
B2 plane	36	31.4916	36	31.32	4.2312	0
12 sides	324	21.4884	324	23.42		0
0.0348s/	36	10.3737	36	11.29		0
frame	324	10.7442	324	12.04		0
	36	10.3737	36	12.04		0
	324	10.3737	324	11.29		0
	36	22.9704	36	23.42		0
	324	28.5277	324	31.32		0
	36	15.5606	36	16.22		0
	144	91.1405	144	94.29		0
	216	89.6585	216	94.29		0
	324	15.5606	324	16.22		0
Diamond	250	49.36	250	50	1.3436	0
4 sides	290	49.05	290	50		0
0.0348s/	70	49.83	70	50		0
frame	110	49.99	110	50		0
Gun	73	42.2551	74	47.38	0.9093	1
12 sides	149	8.4510	150	10.43		1
0.0348s/	79	3.4969	79	4.60		0
frame	90	7.8682	90	9.71		0
	0	118.314	0	132.38		0
	270	14.5708	270	17.51		0
	180	69.9395	180	78.72		0
	270	9.3253	270	10.59		0
	288	4.0798	287	4.86		1
	180	22.4389	180	25.53		0
	253	32.3470	254	37.25		1
	172	23.0217	172	24.49		0

Figure 15: Angles and segment lengths for the original and reconstructed figures. The triangle is from the mill and the other shapes are from the 3D printer.

greatly raise the cost of finding the right method, though this kind of obfuscation has inherent limits: since every speaker has a unique acoustic signature, in principle a set of played-back recordings could be identified as such and peeled away, revealing the desired audio within. However, if the obfuscation greatly increases the cost of the attack, it will make many kinds of manufacturing espionage not worth the price.

To test this hypothesis, we selected eleven similar turbine blade profiles and scaled them so that the print time was approximately the same. The first ten were recorded as they printed individually. The audio recordings from these prints were combined and aligned with a slight stagger at the beginning, and the resultant composite audio was played while the eleventh profile printed. Analysis of the composite audio shows that while the fundamental frequencies were reproduced, the harmonics were lost during the combination step. In the eleventh recording, the fundamentals from the first ten turbine blades obscure the data necessary to reconstruct the eleventh, but the harmonics from the eleventh appear clearly; this harmonic data is sufficient for an audio reconstruction. In future work, we will experiment with combining the audio tracks in a way that preserves the harmonics and matches other features such as amplitude, to obfuscate the recording to a state that will significantly raise the cost of reconstruction.

Limiting the electromagnetic field generated by manufacturing machinery can raise the cost of reconstruction by making it expensive or impossible for reconstruction methods to determine which quadrant an angle of travel lies in. Since magnetometer readings drop off with the cube of the distance from the source, one option is to increase the size of a machine's enclosure. We tested this hypothesis with a large high-end new-model mill at the Digital Manufacturing Design and Innovation Institute, and found that when the phone was placed on the machine's enclosure, the magnetometer was too far away from the motors to pick up useful readings.

When it is not practical to enlarge an enclosure, improving motor shielding can help. For example, recent research on interference shielding has shown that polymer-matrix composites are effective for electromagnetic interference shielding due to their light weight, resistance to corrosion, flexibility, and modest cost. These composites have been used for many purposes (see, e.g., [8]). Additionally, researchers have shown that composites such as carbon nanofiber-polymer can provide effective shielding for the frequency range of $8.2 - 12.4\,\text{GHz}$ [25]. We suggest the use of composites to cover the stepper motors in manufacturing equipment with a shield thin enough that the motor is not damaged by excessive heat retention, but thick enough to protect it from broadcasting sensitive information to an adversary.

6. CONCLUSIONS

Factory floors produce vast quantities of proprietary data that can be stolen by adversaries intent on learning about manufacturing process specifications, product designs, and factory activities. Nation-states' activities to obtain this kind of information have grown over the years, and attacks on the manufacturing sector have become common. Using a CNC mill and a 3D printer to represent subtractive and additive manufacturing, respectively, we demonstrated that ordinary mobile phones can effectively capture these machines' acoustic and electromagnetic information on a factory floor, and the recordings can be used to reconstruct the objects being manufactured and the processes used to make them.

As both additive and subtractive manufacturing tend to be layer-oriented, we adopted a layer-oriented approach to reconstruction, using 3D printer data to reconstruct the shape and process information for three previously unseen manufactured objects. Our method uses signal processing and machine learning techniques, coupled with an interactive framework that uses manufacturing domain constraints and a non-expert human to help guide the reconstruction process. Experiments showed that the method was highly accurate in reconstructing a star, a gun, and an airplane shape from recordings of a 3D printer, even though the human did not recognize the airplane as such after reconstruction.

As mobile phones are ubiquitous, so is the potential for carrying out phone-based attacks, regardless of the state of IT security in a factory floor's systems. Phone recordings may be made deliberately by an attacker or inadvertently by an individual with a compromised application on their phone or other microphone-enabled device. Good audio data can even be captured in the background of a phone call placed or received on the factory floor. For these reasons, we recommend that manufacturers consider obfuscating the side-channel signals emanating from their equipment by playing audio recordings of similar but flawed processes and shielding tool head motors or enlarging machines' enclosures. Most importantly, manufacturers should consider asking their employees and visitors to leave their phones at the factory door.

7. ACKNOWLEDGMENTS

This work was supported by the U.S. Department of Homeland Security under grant 2015-ST-061-CIRC01, NSF under grant CNS 12-23967, and by Singapore's Agency for Science, Technology, And Research under the Human Centered Cyberphysical Systems program. The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the U.S. Department of Homeland Security and NSF.

8. REFERENCES

[1] APT1: Exposing one of China's cyber espionage units. Mandiant Intelligence Center, 2013.

[2] Kaspersky lab survey: One in every five manufacturing businesses has lost intellectual property to security breaches within the past year. Kaspersky Lab press release, August 2014.

[3] Iranian cyber attack on new york dam shows future of war. Time, 2016. [Online; accessed 13-May-2016].

[4] M. A. Al Faruque, S. R. Chhetri, A. Canedo, and J. Wan. Acoustic side-channel attacks on additive manufacturing systems. In *International Conference on Cyber-Physical Systems*, 2016.

[5] D. Asonov and R. Agrawal. Keyboard acoustic emanations. In *IEEE Symposium on Security and Privacy*, 2004.

[6] T. Bifano and Y. Yi. Acoustic emission as an indicator of material-removal regime in glass micro-machining. *Precision Engineering*, 14(4):219–228, 1992.

[7] L. Cai, S. Machiraju, and H. Chen. Defending against sensor-sniffing attacks on mobile phones. In *ACM Workshop on Networking, Systems, and Applications for Mobile Handhelds*, pages 31–36. ACM, 2009.

[8] H.-C. Chen, K.-C. Lee, and J.-H. Lin. Electromagnetic and electrostatic shielding properties of co-weaving-knitting fabrics reinforced composites. *Composites Part A: Applied Science and Manufacturing*, 35(11):1249–1256, 2004.

[9] R. Y. Chiou and S. Y. Liang. Analysis of acoustic emission in chatter vibration with tool wear effect in turning. *International Journal of Machine Tools and Manufacture*, 40(7):927–941, 2000.

[10] D. Foo Kune and Y. Kim. Timing attacks on pin input devices. In *ACM Conference on Computer and Communications Security*, pages 678–680. ACM, 2010.

[11] D. Genkin, I. Pipman, and E. Tromer. Get your hands off my laptop: Physical side-channel key-extraction attacks on PCs. *Journal of Cryptographic Engineering*, 5(2):95–112, 2015.

[12] G. Goller and G. Sigl. Side channel attacks on smartphones and embedded devices using standard radio equipment. In *Constructive Side-Channel Analysis and Secure Design*, pages 255–270. Springer, 2015.

[13] S. Hayashi, C. Thomas, D. Wildes, and G. Tlusty. Tool break detection by monitoring ultrasonic vibrations. *CIRP Annals–Manufacturing Technology*, 37(1):61–64, 1988.

[14] ICS-CERT. ICS-CERT Monitor September 2014 – February 2015. Technical report, March 2015.

[15] ICS-CERT. ICS-CERT Monitor November/December 2015. Technical report, May 2016.

[16] B. Kim. Punch press monitoring with acoustic emission (AE) Part I: signal characterization and stock hardness effects. *Journal of Engineering Materials and Technology*, 105(4):295–300, 1983.

[17] P. Kocher, J. Jaffe, B. Jun, and P. Rohatgi. Introduction to differential power analysis. *Journal of Cryptographic Engineering*, 1(1):5–27, 2011.

[18] P. C. Kocher. Timing attacks on implementations of Diffie-Hellman, RSA, DSS, and other systems. In *Advances in Cryptology*, pages 104–113. Springer, 1996.

[19] S. Liang and D. Dornfeld. Tool wear detection using time series analysis of acoustic emission. *Journal of Engineering for Industry*, 111(3):199–205, 1989.

[20] J. K. Nelson. Acoustic emission detection of metals and alloys during machining operations. Master's thesis, Purdue University, 2012.

[21] D. X. Song, D. Wagner, and X. Tian. Timing analysis of keystrokes and timing attacks on SSH. In *USENIX Security Symposium*, 2001.

[22] K. Uehara and Y. Kanda. Identification of chip formation mechanism through acoustic emission measurements. *CIRP Annals-Manufacturing Technology*, 33(1):71–74, 1984.

[23] Verizon. 2013 data breach investigations report. Technical report, 2013.

[24] M. Vuagnoux and S. Pasini. Compromising electromagnetic emanations of wired and wireless keyboards. In *USENIX Security Symposium*, pages 1–16, 2009.

[25] Y. Yang, M. C. Gupta, K. L. Dudley, and R. W. Lawrence. Novel carbon nanotube-polystyrene foam composites for electromagnetic interference shielding. *Nano Letters*, 5(11):2131–2134, 2005.

[26] K. Zetter. Meet 'Flame,' the massive spy malware infiltrating Iranian computers. *Wired*, 2012.

[27] L. Zhuang, F. Zhou, and J. D. Tygar. Keyboard acoustic emanations revisited. *ACM Transactions on Information and System Security*, 13(1), 2009.

[28] S. Zimmerman and D. Glavach. Applying and assessing cybersecurity controls for direct digital manufacturing systems. In *Cybersecurity for Direct Digital Manufacturing Symposium*, pages 51–64. NIST, 2015.