



A Method for the Detection of Counterfeit, Compromised, or Tampered Electronic Devices

by

Philippe Vibien

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Master of Applied Science
in
Electrical and Computer Engineering

Waterloo, Ontario, Canada, 2022

© Philippe Vibien 2022

Author's Declaration

This thesis consists of material all of which I authored or co-authored: see Statement of Contributions included in the thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Statement of Contributions

Philippe Vibien is the sole author of all material in this thesis, with the following exceptions:

- Portions of Sections 3.1 and 5.5 are adapted from *Counterfeit Detection of FT232RL with Zilky* [1], a white paper co-authored by Dr. Sebastian Fischmeister, Umut Guler, Dr. Carlos Moreno, and Philippe Vibien.
- Portions of Section 3.2 are adapted from *A Method and Apparatus for Detection of Counterfeit Parts, Compromised or Tampered Components or Devices, Tampered Systems Such as Local Communication Networks, and for Secure Identification of Components* [2], a patent application co-authored by Dr. Sebastian Fischmeister, Dr. Carlos Moreno, and Philippe Vibien.
- The Adapter Boards pictured in Figures 4.16c and 4.16d in Chapter 4 are designed and manufactured by Edi Hodzic and Dr. Carlos Moreno.
- Figures 5.4, 5.5, and 5.6 in Chapter 5 are adapted from analysis and figures originally created by Umut Guler.

All work detailed in this thesis was completed under the guidance and supervision of Dr. Sebastian Fischmeister at the University of Waterloo.

Contributing authors are listed in alphabetical order by last name.

Abstract

Modern technological devices, ranging from the pedestrian smartphone to the highly specialized pacemaker, are increasingly dependent on miniaturized electronic components, such as integrated circuits, to accomplish their functions. To source these essential components, manufacturers rely on a complex and fractured global supply chain of distributors, brokers, and gray market intermediaries. The precipitous rise of counterfeit, compromised, or tampered devices being found within electronic supply chains over the past decade poses a threat to the trustworthiness, reliability, and correctness of the technologies that incorporate these components. Current techniques for detecting and mitigating these supply chain threats are unreliable, time-consuming, subjective in nature, and in some instances, cause permanent damage to the component under evaluation.

In response to these concerns, this thesis proposes SILGEO, a system capable of rapidly verifying the authenticity of electronic components and assemblies by analyzing the device’s intrinsic physical properties. Impedance discontinuities and wave propagation characteristics within the device are measured to create a *hardware fingerprint*, which is then used to determine the device’s authenticity. The efficacy of SILGEO at detecting counterfeit microchips was evaluated alongside existing techniques as part of a case study focusing on the commonly counterfeited FT232RL microchip. From an experiment pool containing 260 FT232RL examples obtained from 9 different vendors across the globe, SILGEO was able to determine the authenticity of each microchip with 100% accuracy, correctly identifying a total of 119 counterfeit devices in the experiment pool.

Acknowledgements

I would like to thank my supervisor Professor Sebastian Fischmeister for his support, guidance, and the opportunity to be a member of the Embedded Software Group. Thank you as well to Dr. Carlos Moreno for sharing his expertise, observations, and our many invigorating conversations during my time at the University of Waterloo.

Special thanks to Umut, Jack, and Lukas from the Embedded Software Group for their assistance in collecting, processing, and analyzing the hundreds of gigabytes of measurement data. I also would like to thank Karim, Chris, Ari, and Scott from KA Imaging Inc. for their X-ray expertise and equipment used to image the FT232RL microchip collection.

Dedication

À ma famille et à mes amis — *Merci pour vos encouragements !*

Table of Contents

List of Figures	x
List of Tables	xiii
List of Abbreviations	xiv
1 Introduction	1
1.1 Contributions	2
1.2 Thesis Organization	3
2 Background	4
2.1 Electronic Supply Chains	4
2.1.1 Microelectronic Fabrication	6
2.1.2 Circuit Board Assembly	7
2.1.3 Simplified Electronics Supply Chain	9
2.2 Supply Chain Threats	10
2.2.1 Component-Level Attacks	12
2.2.2 Board-Level Attacks	21
2.2.3 Current Mitigation Efforts	26
2.3 Microwave RF Overview	31
2.3.1 Scattering Parameters	32
2.3.2 Scattering Matrix	32
2.3.3 Two-Port Networks	33

3	SILGEO Overview	37
3.1	Problem Statement	38
3.2	Theory of Operation	38
3.2.1	Single Port Characterization	40
3.2.2	Multiple Port Characterization	41
3.2.3	Determining Authenticity	43
3.3	Data Features	44
3.3.1	Measured Features	44
3.3.2	Derived Features	47
3.3.3	Analysis Methods	48
3.4	Advantages of SILGEO	49
4	SILGEO Acquisition System	50
4.1	Vector Network Analyzer	51
4.2	Multiplexing Board	55
4.2.1	Multiplexing Board Design	57
4.2.2	Multiplexing Board Performance Analysis	68
4.3	Adapter Boards	73
4.4	Acquisition Controller	77
4.4.1	Startup Arguments	78
4.4.2	Functionality	79
5	FT232RL Case Study	81
5.1	FT232RL Overview	81
5.2	Case Study Objective	83
5.3	FT232RL Test Pool	83
5.3.1	Procurement	83
5.3.2	Establishing FT232RL Authenticity	84

5.4	FT232RL Characterization	87
5.5	Case Study Results	89
5.5.1	FT232RL Clustering	90
5.5.2	FT232RL Authenticity	91
5.6	Case Study Conclusions	94
5.6.1	Limitations	95
5.6.2	Lessons Learned	96
6	Future Work	98
6.1	Calibration Improvements	98
6.2	Harnessing the Time-Domain	99
6.3	Additional Applications	101
7	Conclusions	103
	References	104
	APPENDICES	113
A	Complete Set of Multiplexing Board RF Characteristics	114
B	Acquisition Controller Startup Arguments	124
C	Measurement Database Schema	127

List of Figures

2.1	Simplified Supply Chain	9
2.2	Two-Port Network With Incident and Reflected Power Waves	34
3.1	Reflection Response From a Single Pin	40
3.2	Transmission Response Through Two Pins	42
3.3	Simultaneous Reflection & Transmission Response From Two Pins	42
3.4	Characterization of Multiple Pins	43
3.5	Magnitude Transformation	47
3.6	Phase Transformation	48
4.1	Acquisition System	50
4.2	P9370A - 300 kHz to 4.5 GHz VNA	52
4.3	Scattering Parameters for a Two-Port Network	53
4.4	S2P Data File	54
4.5	Multiplexing Board	55
4.6	Multiplexing Board Layout	58
4.7	Multiplexing Topology	60
4.8	Cross Section of a Coplanar Waveguide With Ground Plane	61
4.9	Microstrip Trace Surrounded by a Via Fence	65
4.10	Via Stitching	68
4.11	Multiplexing Board Characterization	69

4.12	Multiplexing Board Insertion Loss	71
4.13	Multiplexing Board Mean Isolation	72
4.14	D38999 Adapter Board Bottom View	74
4.15	Multiplexing Board With Attached D38999 Adapter Board	74
4.16	Adapter Boards for Various DUT's	76
4.17	Acquisition in Progress	77
4.18	Acquisition Controller Startup	78
5.1	FTDI FT232RL	82
5.2	X-ray Images of FT232RL Devices	86
5.3	FT232RL Measurement Setup	88
5.4	FT232RL Clustering, Colored by Vendor	90
5.5	Similarity Score, Colored by Authenticity	92
5.6	Confusion Matrix of the Authenticity Results	93
A.1	Insertion Loss	114
A.2	Mean Isolation	115
A.3	Isolation When SMA ₁ to RF ₁ Selected	115
A.4	Isolation When SMA ₁ to RF ₂ Selected	116
A.5	Isolation When SMA ₁ to RF ₃ Selected	116
A.6	Isolation When SMA ₁ to RF ₄ Selected	117
A.7	Isolation When SMA ₁ to RF ₅ Selected	117
A.8	Isolation When SMA ₁ to RF ₆ Selected	118
A.9	Isolation When SMA ₁ to RF ₁₅ Selected	118
A.10	Isolation When SMA ₁ to RF ₁₆ Selected	119
A.11	Isolation When SMA ₂ to RF ₇ Selected	119
A.12	Isolation When SMA ₂ to RF ₈ Selected	120
A.13	Isolation When SMA ₂ to RF ₉ Selected	120

A.14 Isolation When SMA ₂ to RF ₁₀ Selected	121
A.15 Isolation When SMA ₂ to RF ₁₁ Selected	121
A.16 Isolation When SMA ₂ to RF ₁₂ Selected	122
A.17 Isolation When SMA ₂ to RF ₁₃ Selected	122
A.18 Isolation When SMA ₂ to RF ₁₄ Selected	123

List of Tables

4.1	Multiplexing Board PCB Stackup	63
4.2	Signal Path Lengths	64
5.1	FT232RL Vendors and Quantities Obtained	83
5.2	FT232RL Suspected Category Counts	84
5.3	FT232RL Pin Selection	88
5.4	FT232RL Characterization Criteria	89

List of Abbreviations

AFR automatic fixture removal [99](#)

BMC baseboard management controller [23](#)

BOM bill of materials [8](#), [21–25](#)

CPU central processing unit [6](#), [19](#)

CPWG coplanar waveguide with ground [61](#), [64](#), [65](#), [75](#)

CT computed tomography [29](#), [30](#)

dB decibels [35](#), [36](#), [47](#), [68](#)

DC direct current [57](#)

DTW dynamic time warping [48](#), [91](#), [92](#)

DUE device under evaluation [38](#)

DUT device under test [32](#), [33](#), [38](#), [40](#), [48](#), [50–52](#), [55–57](#), [59](#), [68](#), [72](#), [73](#), [75](#), [88](#), [96](#), [98–100](#), [103](#)

EDA electronic design automation [10](#), [20](#), [65](#)

EEPROM electrically erasable programmable read-only memory [81](#)

EMI electromagnetic interference [25](#)

ENIG electroless nickel immersion gold [63](#)

ESD electrostatic discharge [14](#), [15](#), [27](#), [91](#), [101](#), [102](#)

FFT fast Fourier transform [100](#), [101](#)

FPGA field-programmable gate array [16](#), [51](#)

FTDI Future Technology Devices International [81–86](#), [91](#), [95](#)

GAO Government Accountability Office [1](#)

HDL hardware description language [6](#), [19](#)

Hz hertz [31](#)

IDM integrated device manufacturer [7](#), [20](#)

IF intermediate frequency [79](#), [125](#), [128](#)

IFFT inverse fast Fourier transform [100](#), [101](#)

IP intellectual property [6](#), [17–21](#)

JSON JavaScript Object Notation [125](#)

LDA linear discriminant analysis [49](#)

MB Multiplexing Board [51](#), [55–59](#), [61–70](#), [72](#), [73](#), [75](#), [77](#), [79](#), [80](#), [87](#), [89](#), [96](#), [98](#), [99](#), [124](#), [127](#), [128](#)

NASA National Aeronautics and Space Administration [84](#)

NSA National Security Agency [22–24](#)

OCM original component manufacturer [11](#), [84](#)

OEM original equipment manufacturer [11](#)

OS operating system [77](#), [82](#), [85](#)

PCA principal component analysis [49](#), [90](#), [95](#)

PCB printed circuit board 5, 7–9, 14, 21–25, 39, 61–63, 65, 70, 71, 73, 75, 88, 98

PCBA printed circuit board assembly 7

RF radio frequency 3, 4, 31, 32, 39, 40, 47, 51, 52, 54, 55, 57–62, 65, 69, 70, 73, 79, 99, 100, 125, 127

RTL register-transfer level 6, 17–20

SASC Senate Armed Services Committee 1, 13

SMA SubMiniature version A 52, 55, 56, 58, 61, 63, 69, 70, 73, 75, 98, 99

SMP SubMiniature version P 55, 56, 58, 61, 63, 69–71, 73, 75

SMT surface-mount technology 8, 14, 23, 62

SNR signal-to-noise ratio 73

SSOP shrink small outline package 75, 81, 87–89, 94, 96

t-SNE t-distributed stochastic neighbor embedding 49

TDR time-domain reflectometry 63, 100

TDT time-domain transmissometry 100

TOCTOU time-of-check to time-of-use 37

UART universal asynchronous receiver-transmitter 81, 85

USB Universal Serial Bus 23, 51, 53, 57, 75, 81, 85–87

VHDL VHSIC Hardware Description Language 6

VNA vector network analyzer 32, 51–56, 58, 68–71, 73, 75, 77–80, 98, 100, 127, 128

WLCSP wafer-level chip scale packaging 23

Chapter 1

Introduction

In early 2011, the United States [Senate Armed Services Committee \(SASC\)](#), charged with legislative oversight of the nation’s military, ordered an inquiry into the presence of counterfeit electronic parts within defense supply chains [3]. The findings were stark; over a two-year period, defense contractors had encountered 1,800 incidents of suspected counterfeiting involving more than 1 million individual parts [4]. Counterfeit parts had been found everywhere, from video screens and thermal cameras to safety-critical missile guidance computers and flight avionics for planes and helicopters. As part of the inquiry, the United States [Government Accountability Office \(GAO\)](#), a congressional agency tasked with auditing and investigating, was asked to study how these non-conforming parts had entered defense supply chains [5]. Using a fictitious company, the [GAO](#) joined several electronic component marketplaces and attempted to purchase military-grade electronic parts. Interspersed among quote requests for legitimate part numbers, the GAO had also asked for non-existent and fictitious components, using part numbers that were not associated with any known authentic components. The response was startling; rather than advise the part numbers were not obtainable, nearly 40 vendors, all located in China, stated they could supply the non-existent components.

The message sent by these early reports was clear: supply chains are not friendly territory. They involve untrusted entities with differing sets of priorities, and yet electronic supply chains are essential to the construction of modern technological systems. Miniaturized electronic components, such as integrated circuits, underpin vast areas of our societies, to the point that their reliability, correctness, and trustworthiness is simply assumed. For the trusted hardware used to construct secure and safety-critical systems, this assumption is no longer sufficient [6]. With malicious entities capable of tampering with nearly every stage

of the supply chain, designers, manufacturers, and end-users must have positive verification that hardware is authentic and unmodified.

Current techniques for assuring hardware integrity and authenticity are inadequate. Most methods are haphazard, time-consuming, and overly reliant on anecdotal observations. Human experts are often required to make subjective judgments as to whether an anomaly is the result of malicious actors or just sloppy manufacturing. In practice, many techniques are not 100% reliable in isolation, necessitating the use of multiple different techniques and evaluating the results holistically to obtain a reliable authenticity determination. This has left government and industrially alike in a vulnerable position, having only limited ability to authenticate hardware throughout the stages of manufacturing, testing, and end-use.

To address these challenges, this thesis presents SILGEO, a *hardware fingerprinting* capability that can rapidly determine the authenticity of an electronic device. SILGEO uses physical characteristics intrinsic to all electronic devices, namely the speed of electromagnetic wave propagation within the device and the presence of impedance discontinuities, to create a unique fingerprint. This fingerprint is then statistically compared with the fingerprints of other devices, or a known-good fingerprint, to quantitatively verify the legitimacy of a device. Applicable to individual components and larger assemblies, SILGEO is nondestructive, instant, and can be applied to devices in situ.

1.1 Contributions

The contributions of this thesis to the literature can be summarized as follows:

- Presents an overview of electronic manufacturing and the supply chain entities involved.
- Provides a thorough discussion and taxonomy of the threats currently facing the supply chains that support electronic manufacturing and summarizes the current state of threat mitigation efforts.
- Proposes SILGEO, a novel technology capable of rapidly verifying the integrity of microelectronics and assemblies by analyzing the electronic device’s intrinsic wave propagation and impedance characteristics.
- Confirms the pervasive nature of microelectronic counterfeiting with a case study focusing on the popular FT232RL microchip, encompassing 260 devices sourced from 9 different vendors.

1.2 Thesis Organization

Following the introduction in Chapter 1, Chapter 2 provides background information on electronic manufacturing, supply chain threats, and a brief review of microwave [radio frequency \(RF\)](#) concepts utilized by SILGEO. Chapter 3 then follows with an overview of SILGEO’s method of operation, with SILGEO’s implementation then detailed in Chapter 4. Chapter 5 details the experimental results of testing SILGEO on genuine and counterfeit FT232RL microelectronic devices. Next, Chapter 6 discusses possible avenues of future work and additional applications for SILGEO. Finally, some closing remarks are present in Chapter 7.

Chapter 2

Background

Before detailing the specifics of SILGEO and its method of operation, it is important to first contextualize the challenges currently facing electronic systems and their supply chains. To that end, this chapter opens with an overview of electronic manufacturing and supply chains. It continues with a taxonomy on the types of threats facing electronic supply chains and the current efforts to detect and mitigate them. Additionally, some introductory microwave [RF](#) concepts related to the operation of SILGEO are detailed.

2.1 Electronic Supply Chains

Modern supply chains are increasingly global and fractured affairs, with end-users of products many steps removed from the manufacturer. Advances in communications and technology have allowed businesses to segment and specialize their operations, leading to faster production, increased efficiency, and lower costs [7]. The electronics industry is no exception to this, and in many ways has been at the forefront of this trend. Manufacturing steps that were previously performed by a single vertically integrated entity are now delegated to an armada of assembly houses, subcontractors, distributors, and contract manufacturers. It is not uncommon for an electronic product to be handled by numerous entities, crisscrossing the globe while on its way from designer to final customer [8].

The electronics industry is generally considered to be organized into three manufacturing segments: raw materials, electronic components, and assemblies [9]. *Raw materials* refers to the inputs needed for the subsequent manufacturing of components and circuit boards. Semiconductor fabrication, for example, requires numerous input elements (besides silicon,

of course), such as arsenic, boron, gallium, and phosphorus, in large high-purity quantities. Circuit boards, in turn, require ceramics, plastic, fiberglass, and various metals, including copper, gold, aluminum, and silver.

The *electronic components* segment refers to the manufacturing of electronic devices which are utilized in electronic circuits, ranging from discrete resistors, to highly integrated semiconductor devices. Components are typically categorized as *passive*, such as resistors, capacitors, and inductors; or as *active*, with the distinction that active components can control the voltage or flow of current within a circuit, and passives cannot [9]. Examples of active components include amplifiers, logic circuits, microcontrollers, memory, and processors. A circuit will often contain an interconnected mix of passive and active components, although depending on the application, single component circuits also exist. The manufacturing of components can require a broad spectrum of steps and processes depending on the final electronic device, but in general the segment is considered to be both capital and labor intensive, compounded with demanding quality expectations if a firm wants to be economically profitable [10]. In particular, the design and production of silicon based devices is characterized by having exacting quality requirements, the interaction of numerous separate entities, and long fabrication cycles [11].

The *assemblies* segment varies by product, but most often refers to the packaging of silicon dies into microchips, and separately, the population of components onto a [printed circuit board \(PCB\)](#). For silicon-based devices, assembly entails electrically connecting the bare silicon die to metal leads that will be externally accessible, while simultaneously encapsulating the device in the familiar black epoxy resin. For circuit boards, assembly refers to the placing and soldering of electrical components to the [PCB](#), often performed by an *assembly house*. Usually, the client contracting with the assembly house is responsible for providing the bare circuit boards and components to be soldered. However, some assembly houses also offer *turn-key* services, where they handle the production of the bare circuit boards and sourcing of components. Additionally, some assembly houses offer *box-build* or systems integration services, which goes beyond turn-key, and can include fabrication of enclosures, wiring harnesses, the installation of any sub-assemblies, and final product packaging [9]. Box-build can also include customer-specific product customization, such as the amount of storage in a phone or the choice of graphics card in a laptop.

Within each of the aforementioned market segments, there are entities, ranging from manufacturers, contractors, distributors, specialists, and consultants, who have expertise in their specific area of the supply chain. In order to better understand the interaction between these entities, the following two sections provide an overview of the steps involved in microelectronic fabrication and circuit board assembly. Afterward, Section [2.1.3](#) discusses the supply chain entities in a unified manner.

2.1.1 Microelectronic Fabrication

Microelectronic manufacturing is broadly divided into three major stages: design, fabrication, and packaging. The design of a new microelectronic device, also known as an *integrated circuit*, starts like many other products: with the drafting of a high-level specification. For an integrated circuit, the specification will typically encompass functional requirements, desired performance, power, timing, size, and the environment within which the device will operate [8]. From these specifications, engineers will describe the required circuits in a [hardware description language \(HDL\)](#), such as [VHSIC Hardware Description Language \(VHDL\)](#) [12] or Verilog [13], to produce a [register-transfer level \(RTL\)](#) representation of the circuit. The [RTL](#) acts as a high-level abstraction of the design, comprising of registers and combinational logic; it captures the flows between gates and the logical operations to be performed. The [RTL](#) is then combined with a silicon technology library, a sort of catalog containing standard gate designs for a particular silicon process size, to synthesize a gate-level description of the circuit, commonly called a *netlist*. The netlist is then further combined with process-specific information from the desired silicon foundry to produce a fully placed and routed transistor level *layout*. The layout is an artwork representation of the microelectronic circuit, consisting of geometric shapes for the layers of semiconductor, metal, and oxide that comprise the desired circuit.

Throughout the design process, engineers can choose to incorporate pre-made circuit blocks from external vendors rather than developing everything in-house from scratch. These pre-made blocks are commonly known as [intellectual property \(IP\)](#) cores and are generally sold under a royalty license from their respective designer [8]. The breadth of [IP](#) cores available for sale is quite literally endless, with entire [central processing units \(CPUs\)](#), embedded memory, multimedia accelerators, graphics controllers, networking peripherals, display drivers, and encryption accelerators available, to name a few [14]. Depending on the vendor, the [IP](#) cores are provided as synthesizable [RTL](#) or as completed layout artwork ready to be inserted. The use of [IP](#) cores in designs is increasingly common, particularly with firms looking to avoid costly re-development or reduce their time to market [14]. The most recognizable vendor of [IP](#) cores is ARM Ltd., whose soft microprocessor designs were used in over 25 billion chips in 2020 alone [15].

After the design has been verified and validated, the layout will be sent to a semiconductor fabrication facility, commonly referred to as a *foundry*, for manufacturing [8]. At the foundry, the manufacturing process starts with the production of *lithographic masks*, derived from the supplied artwork, for use in the subsequent photolithography process to image the circuit layout onto the silicon wafer. Starting from a bare silicon wafer, the required semiconductor layers are built up with a repeated process of material deposition, photolithographic

patterning, and chemical etching. Depending on the electronic properties required, stages of dopant introduction and ion implantation are also performed [16]. Once the semiconductor layers are complete, successive *metalization* layers are added to interconnect the various gates, logic sections, and subsystems on the die. Throughout the fabrication process, various automated and manual wafer inspections are performed to check for defects, adjust parameters, and optimize production yields. Completed wafers are then cut into individual silicon dies, with the non-defective dies sent for final assembly.

In order to protect the finished silicon from harsh environments, i.e., anything other than a clean room, and to enable connectivity with the external world, the bare dies are then *packaged*, typically in an epoxy like material. The packaging process involves *wire bonding* the die to a lead frame, using fine wires made of materials such as gold and aluminum. The lead frame is a metal platform which contains the externally accessible electrical pins, provides mechanical rigidity for handling, and thermal heat sinking. The packaged devices are then tested for defects and, depending on the product, binned according to performance, quality, or another metric. The devices are then sent to warehouses or distributors, where they are stored until ordered by customers.

Historically, firms creating silicon components were vertically integrated. In addition to creating the design, firms would operate their own foundries, and the equipment needed final assembly and packaging. However, the increasing complexity of silicon fabrication, which in turn requires ever-increasing capital investments, has given rise to stand-alone fabrication companies, referred to as *pure-play* foundries, who produce silicon under contract for other companies. This has led to an increase in *fabless* semiconductor companies that only focus on designing devices and have no semiconductor fabrication capabilities. Vertically integrated companies, today called [integrated device manufacturers \(IDMs\)](#), still exist, and many have followed in the footsteps of contract manufacturers, opening up their foundries to external orders while still producing their own devices [17]. In a similar vein, the translation of the layout artwork to a set of lithographic masks can be done in-house by the foundry or, as is increasingly common, outsourced to a specialty third party [11].

2.1.2 Circuit Board Assembly

Commonly referred to as circuit boards, [printed circuit boards \(PCBs\)](#), and [printed circuit board assemblies \(PCBAs\)](#) when fully assembled, a PCB serves two main functions; the first is to provide an electrical interconnect, i.e., wiring, for power and signals between the various components. The second is to provide a structural backing to hold the various components in place and to enable the physical mounting of the overall assembly in a larger enclosure

or system. Circuit boards are typically constructed from laminated layers of conducting copper and insulating fiberglass, but more exotic materials are also used, such as ceramics, aluminas, and various blends [18]. When fully assembled, a PCB can contain anywhere from tens to thousands of individual components.

The PCB design process typically starts with the creation of a *schematic*, which details the components used, their inputs and outputs, and how they are interconnected [19]. The schematic will also contain a reference designator for each component, which enables the unique identification of each component which is then used when populating the board with components during assembly. The information from the schematic is then used in the *board layout*, a document which, among other things, depicts the location of the component footprints, dimensions of copper traces, and the location of holes to be drilled on the board [19]. The layout also typically contains a *silkscreen* layer which contains text to be silkscreened, i.e., printed, onto the board, such as product information and reference designators.

After the completion and final verification of the schematic and layout, a set of semi-standardized files are generated and sent to the circuit board manufacturer and board assembly house [19]. Typically the files sent include the *gerbers* which contain the lithographic masks for each board layer; a drill file which details where holes should be drilled into the circuit board and if they should be plated; and a *bill of materials* (BOM) containing all the parts needed for assembly, their position and orientation on the board, and their designator. The board manufacturer then adapts the design to their specific manufacturing process and proceeds with various requested electroplating, etching, silkscreening, and drilling steps. The boards are then typically tested for electrical faults such as shorts and opens.

For component population, circuit boards are typically assembled in two passes, with the *surface-mount technology* (SMT) process first and the through-hole components second. With SMT, the bare circuit boards are first *pasted* with solder using a stencil derived from the gerbers, and then a machine called a *pick-and-place* will populate the board with the requested surface-mount components [20]. The board then passes through a high-temperature oven, known as a *reflow oven*, which flows, i.e., melts, the solder paste into hardened solder, thus forming both an electrical and mechanical bond. After finishing the SMT process, through-hole components are placed, typically by hand, and the board then passes through a wave solder machine which applies solder to the exposed leads [20].

It should be noted that not all of the footprints on a circuit board will always be populated, with these missing components aptly called *unpopulated components*. This typically occurs when the circuit board design is used in multiple product models, where the different versions require boards with fewer or added components. It can also occur with the components related to debugging and testing, which are usually unpopulated on production boards.

As a final step, the assembled boards are often inspected, either manually or automatically, for any abnormalities such as solder bridges and misplaced or missing components [19]. If any are present, the board is pulled and then sent for rework or an equivalent corrective procedure. The completed boards are then sent for *final assembly*, where they are integrated with other parts and sub-assemblies into the desired final product.

2.1.3 Simplified Electronics Supply Chain

While the supply chains that support microelectronic fabrication and circuit board assembly were introduced here separately, the reality is that they are actually one and the same. The fabrication of silicon microelectronic components is upstream from, and directly feeds, the supply chains used in PCB assembly. To better illustrate this, Figure 2.1 presents a flow diagram showing the various entities and relationships, albeit simplified, involved in the production of a microelectronic device which is then incorporated onto a circuit board. What immediately stands out in the diagram is the sheer number of entities that are involved in the electronics industry and the interconnectedness that is present between them.




Figure 2.1: Simplified Supply Chain

Figure 2.1 also introduces some entities not discussed in the previous sections, such as *distributors* and the *gray market*. In the electronics ecosystem, manufacturers typically only sell their products to a handful of authorized distributors, which in turn resell them to final customers such as designers or assembly houses. These distributors are “authorized” in the sense that they typically have commercial or legal arrangements with manufacturers to

carry their parts, whereas distributors operating in the gray market, also sometimes referred to as *brokers*, do not [21]. The specifics of these arrangements can vary by manufacturer; however, they generally set out guidelines that the distributor needs to follow for obtaining parts, storing them, and the price they can resell them at. The gray market, put simply, is a collection of entities that sell components outside of the authorized distribution channels. This can include reputable entities who may simply be reselling components bought from authorized distributors, and seedy entities who have obtained components from dubious sources. The concern with the gray market is that is difficult to know exactly what you are buying and where it came from, with brokers having been known to sell counterfeit parts or genuine parts that have been damaged from improper storage and handling practices [22].

Additionally present in the figure are the various equipment and tool vendors that service the silicon and circuit board manufacturing markets. Although only represented as two ellipses in the figure, the scale of entities here too is impressive, ranging from humble sellers of soldering equipment to vendors of multi-million dollar lithograph machines [23]. The same can be seen on the software side, with a wide range of [electronic design automation \(EDA\)](#) tooling available that can assist in areas of schematic capture, layout, simulation, and design verification.

2.2 Supply Chain Threats

As electronic supply chains have become more globalized and interconnected, so too have the threats they face. Malicious entities have taken advantage of the fracturization and limited visibility into supply chains to launch increasingly sophisticated attacks. In particular, techniques created by designers and manufacturers to achieve flexible and low-overhead supply chains, such as fabless foundries, turn-key assembly, and *build-to-order* customization, are starting to be used by malicious entities to implement and carry out attacks [24]. While the malicious entities behind supply chain threats have varying motivations, their objectives generally fall into one of two broad categories: *espionage* or *profit-seeking*.

Espionage is typically carried out by nation-state actors, such as intelligence agencies, with the objective of *sabotaging*, *back-dooring*, or *ex-filtrating data* from a device or system [24]. Supply chains are ideal attack vectors for these entities because alternations can be introduced covertly and in a non-attributable manner. The rise of contract manufacturing and *build-to-order* customization allows for tampered devices to be sent to specific targets with ease. For example, rather than entering a person’s residence to plant a listening device, a target’s phone can be modified at the factory during assembly, or even intercepted mid-shipment, to include a listening device.

Even when nation-state actors are not a concern, supply chains are adversarial places, with customers, assemblers, distributors, manufacturers, and their transient workforce, all having misaligned interests and goals. Low margins and a hyper-competitive environment, leads to all manners of illicit profit-seeking behavior. From the unauthorized use of recycled components to the over-production of devices during the off-hours, designers and end-users alike need to be on the lookout for dishonest behavior from the firms they partner and work with. While threats from profit-seeking actors may appear to be more benign, they are just as insidious to detect and remediate as those coming from nation-state actors.

All areas of the electronic ecosystem are under threat, from individual components to entire systems. The following sections cover these various threats, starting at the component-level and then scaling up to the level of circuit boards. Afterward, an introduction to some of the current detection methods that can be used to mitigate these threats is provided.

Terminology

Before discussing specific supply chain threats, it is necessary first to define some key terms relating to authenticity and counterfeiting, as their meaning and interpretation can vary between government, industry, and end-users. In general, the terminology used in this thesis follows that used in various industry standards and academic literature. Firstly, the terms **original component manufacturer (OCM)** and **original equipment manufacturer (OEM)** refer to the original manufacturer of a device, which can also be the designer of the device or simply a third party that has a license to manufacture the device [11]. The terms **OCM** and **OEM** can be used interchangeably, but **OCM** is generally reserved for electronic components, while **OEM** is used for assemblies and larger devices. Importantly, only the original manufacturer can make an *official* determination as to the authenticity of a component due to the contractual and fraud implications an official authentication implies. Determinations made by entities other than the original manufacturer are generally prefaced with the term *suspected*, indicating that strong evidence as to the authenticity of the component exists; however, the original manufacturer has not made an official determination [11].

At a foundational level, the term *authentic* implies that a device (component, part, circuit board, assembly, etc.) is truly genuine and was produced by the original manufacturer [25]. Essentially, an authentic device means that when a device is *represented* as a particular device, it is the genuine article. Conversely, the term *counterfeit* implies that while a device was represented as a particular device, it is, in fact, not the genuine article. There are various competing definitions as to what truly counts as a “counterfeit” device, but fundamentally the term *counterfeit* indicates that the authenticity of a device has been *misrepresented* in some manner [11]. The range of misrepresentations varies and can cover any aspect of the

device, including type, model, function, date of manufacture, origin, and even supporting documentation. Additionally, the term counterfeit carries legal implications that the buyer of a device who receives a non-genuine article has been misled, deceived, or defrauded [25].

From this starting point, some additional definitions can be introduced. A *known authentic* device, sometimes also referred to as *known-good*, is one which has either been purchased directly from the original manufacturer, their authorized distributors, or has been confirmed as genuine by the original manufacturer [25]. Similarly, a *known counterfeit* device is one which has been confirmed to be non-genuine by the manufacturer. A *suspected authentic* device is one which has been inspected by someone other than the original manufacturer with various methods, and based on the result of the inspection; the device is *suspected* to be genuine [11]. The specifics of the checks that are performed as part of determining device authenticity are discussed further in Section 2.2.3, but essentially they involve looking for inconsistencies in the appearance, construction, or operation of the device. If a device fails one or more authenticity checks, the device is then considered to be a *suspected counterfeit*, indicating the device is possibly not genuine [11].

2.2.1 Component-Level Attacks

Electronic components are the building blocks of electrical systems and thus are tempting targets for attackers. Commonly referred to as *microelectronic counterfeiting* or *component counterfeiting* in the electronics industry, component-level attacks come in a variety of flavors. When thinking of counterfeit products, the traditional idea of a “clone” or lower quality “knock-off” typically comes to mind. While this style of counterfeiting is common with electronic components, and is actually the subject of the case study in Chapter 5, it is not the only type.

For a non-exhaustive taxonomy, component counterfeiting can generally be broken down into the following categories: *recycled*, *remarked*, *defective*, *cloned*, and *maliciously altered* [11]. Subsequent sections will delve into greater detail for each of these categories, discussing the risks posed and malicious entities involved. Component counterfeiting, in all its forms, is by no means limited to just integrated circuits, but it predominantly involves that market segment. For context, in 2017, 72.3% of all reported counterfeit components were integrated circuits [26]. The common consensus is that integrated circuits are heavily targeted by counterfeiters, instead of other electronic segments, because of the higher profit margins that can be earned [27].

The use of counterfeit components within devices and systems can have catastrophic and far-reaching consequences. By unintentionally, or intentionally, incorporating counterfeit

components during manufacturing, a system’s overall safety, security, and reliability can be jeopardized. At the most basic level, counterfeit components pose a threat because while they are imitating a genuine component, they lack the expected quality, correctness, and longevity properties of the genuine equivalent. Counterfeit components can harbor hidden defects due to faulty manufacturing, careless storage or handling, and even prior usage [11]. Because of these defects, counterfeit components are susceptible to random failures, which can potentially cause injury to life and property. Counterfeit components are especially hazardous for secure systems, military equipment, and other systems that rely on trusted hardware, as the provenance of these components is often unknown. Counterfeit components threaten the integrity of these systems as they can contain implants, trojans, and other malicious alterations, which attackers can harness to carry out espionage, sabotage, and other nefarious acts. Beyond these concerns, counterfeit components also cause economic damage [11]. According to the Alliance for Gray Market and Counterfeit Abatement, an electronics trade group, every year, legitimate firms lose out on more than \$100 billion of revenue because of counterfeiting [28].

As electronic supply chains have become global affairs, components are increasingly handled by multiple entities as they flow from manufacturer to end-user. Whereas previously, a component may have been sourced directly from the manufacturer or authorized distributor, today, it is common for components, especially those in short supply, to flow through brokers and other gray market distributors [8]. The concern is that each additional entity involved in a supply chain becomes an opportunity for malicious actors to tamper or introduce counterfeits. The final SASC report on counterfeits within defense supply chains, published in 2012, focused on this concern and documented multiple instances of lengthy and fractured supply chains being responsible, in part, for counterfeit components ending up in military equipment.

A notable example is the United States Navy P-8 Poseidon, a maritime patrol and reconnaissance aircraft, which had five separate subsystems affected by counterfeit parts [4]. One subsystem, the flight critical ice detection module, had an almost comical supply chain. While Boeing was the prime contractor for the P-8 Poseidon, a subcontractor, BAE Systems, oversaw the manufacturing of the ice detection module. As part of the module assembly, BAE Systems hired a subcontractor, Tandex Test Labs in California, to source various components and ensure they were genuine. Tandex Test Labs, in turn, purchased the components from a broker, Abacus Technologies in Florida, which had obtained the components from A Access Electronics in Shenzhen, China [3]. All in, the counterfeit components had flowed through five different companies, with the completed ice detection modules even being delivered to the Navy before the non-conformity was discovered.

Recycled

Considered to be the bulk of counterfeiting incidents, *recycled* components refers to the practice of harvesting previously used genuine components and misrepresenting them as newly manufactured ones. In 2010, it was estimated that recycled components accounted for more than 80% of the counterfeit devices in circulation worldwide [29]. Nearly all recycled components are harvested from previously used electronic waste, also referred to as *e-waste*, that has been sent for recycling. Normally this waste is supposed to be broken down and processed by legitimate recyclers, but in practice, it has been estimated that only around 25% of e-waste is properly recycled [30]. Instead, a majority of the waste ends up being diverted to businesses, often located in underdeveloped countries, that scour the waste stream for components of value, harvesting any they find and then re-selling them as “new” [11]. The fundamental threat posed by recycled components is that while they may be masquerading as new, they are, in fact, not. As such, they lack the reliability and performance qualities associated with new components. Furthermore, recycled components may harbor latent faults or other damage stemming from their prior usage or from the recycling process itself.

The recovery process used by “recyclers” varies, but it is typically very crude and subjects the components to extreme conditions. Some documented methods include heating circuit boards over an open fire and, when the solder starts to melt, banging the PCB aggressively into a bucket to dislodge and collect the components [31]. These components are then often washed with harsh chemical baths in an attempt to remove any remaining solder on the pins. Some recyclers have also been seen using more sophisticated methods to remove components, such as hot air rework stations normally used for fixing SMT board issues [27]. Hot air stations, however, can still cause permanent damage to components, as recyclers will often use air temperature settings that are much higher than ones normally used for rework, as higher heat reduces the time needed to unsolder components. In either case, throughout the removal process, recyclers generally do not follow any electrostatic discharge (ESD), temperature, humidity, and other handling best practices, thus further subjecting the components to harsh conditions [11].

Even if the recovery process is performed with the utmost care to avoid any damage, recycled components still have an elevated risk of age-related degradation due to their prior usage. This is an even greater concern for components that have a limited endurance before they fail, such as with the write cycles of a flash storage device or the total ionization dose for radiation hardened components. These limited endurance components typically set the overfall estimated lifetime of the products they are used in, meaning that if recycled components are accidentally used during manufacturing, an entire product will now have a reduced lifespan [32].

The detection of recycled components can often be done via a simple physical inspection, as the recycling process typically leaves components scarred with a plethora of visual evidence. Solder residue, bent or missing pins, burn marks, packaging cracks, and other damage has commonly been observed on recycled components [33]. Recently, some recyclers have started taking additional steps to clean up the components by resurfacing the package or spraying a coating, commonly called *blacktop*, that approximates or is even identical to the coating used by the original manufacturer [29]. This leaves the components, at least externally, looking visually new; however, since the underlying die is not new, electrical testing can be used in these cases to reveal that the component has been recycled [11].

Remarked

The packaging that encases microelectronic devices typically has markings on it that uniquely identify the part and its capabilities. These markings can detail information such as the manufacturer’s name, device functionality, product code, serial number, date of manufacture, ESD sensitivity, quality grade, and country of origin [32]. In practice, because the silicon die is not visible and since component package sizes are standardized, these markings are effectively the only way that a particular component can be discerned from any other component. The threat of remarked devices occurs when an attacker changes these markings in an attempt to pass it off as a different component [11].

The incentive for remarking is almost always that of profit, in that it allows cheaper components to masquerade as more expensive ones, since the markings essentially determine the value of the component [32]. Depending on the sophistication of the malicious entity, different approaches can be taken as to which components are remarked and in what manner. On the simpler side, an attacker can find a pin equivalent component, one that may not even be electrically or functionally similar to the component they want to counterfeit, sand off all of the existing markings, and then add back markings for the desired component [31]. To avoid detection, since these components do not typically function and, in the worst case, can cause electrical problems, a malicious entity will typically cut the remarked chips into a batch of legitimate chips. By doing so, the inclusion of the remarked components is disguised as the “failure” of a small number of components in the batch, thus drawing less attention, since the failures could plausibly be caused by a soldering issue or damage from ESD [24].

A slightly more subtle remarking attack also exists, where rather than remarking the entire part, a malicious entity only changes the *grade* of the component. Many components, especially microcontrollers and memory, are offered in multiple grades depending on the intended use environment. Grades such as *industrial-grade*, *military-grade*, or *space-grade* are designed to withstand more demanding operating conditions, including a wider range of

temperatures, high humidity, and even radiation. As a result, components at these grades can cost significantly more when compared to lower grades, even if the devices are functionally or electrically similar. To carry out this attack, a malicious entity takes a lower grade version of the component, changes the markings, and then attempts to sell it as the more expensive grade version [11]. What makes this approach much more dangerous is that a component with a remarked grade or quality will generally test perfectly fine at the factory, and only once it is deployed will the problem be noticed, as the lower grade component will often have a shorter lifetime and can even completely fail [22]. Additionally, unlike the previous example where an attacker needs to remove and add back all the markings on the component, changing a component’s grade often only requires minor changes to the markings. Sometimes the only difference between a lower and higher grade component is the presence of an additional character, meaning no existing marks need to be first removed, which makes detection even harder [24].

In general, low-effort attempts at remarking components can easily be detected visually due to surface texture imperfections, inconsistent fonts or lettering, typographical errors, incorrect pin counts, and so forth. However, when perpetrated by a more sophisticated malicious entity, the detection of the remarking attack, in particular the changing of quality grades, becomes significantly more complex and resource-intensive to detect, often requiring X-ray die inspection or equivalent techniques to catch [11]. The occurrence of remarking attacks has been documented publicly numerous times, with one example being the previously mentioned failure of the Navy P-8 Poseidon ice detection module. In this case, after a thorough investigation, it was determined that the military-grade [field-programmable gate array \(FPGA\)](#) devices from Xilinx were, in actuality, an assortment of recycled consumer-grade [FPGA](#) devices that had been remarked [4].

Defective & Out-of-Spec

As part of the silicon production process, microelectronic devices typically undergo several stages of quality control, ranging from on-wafer probing to post-manufacturing tests. Devices that fail these tests outright are considered to be *defective* and are marked for destruction. Devices that are functional but marginally outside of the specified acceptable tolerances, termed *out-of-spec*, are either marked for destruction or, in certain cases, assigned a lower quality grade and sold through official channels [32].

Depending on the security at the fabrication or packaging house, devices marked for destruction are sometimes acquired by malicious entities. These devices are packaged and labeled as needed, and the malicious entity will then attempt to pass them as functional

devices on the open market. End-users then unknowingly integrate the devices into their products, thinking they meet the quality standards of the genuine manufacturer.

Unfortunately, the detection of either defective or out-of-spec devices is not straightforward [11]. In cases where the entirety of the device is defective, functional tests done by the end-user can usually reveal the problem [31]. However, if only a subsection of the device is defective, the anomaly can remain hidden. The detection of out-of-spec devices is particularly challenging, as the testing process to reveal such deficiencies requires intimate knowledge of the component’s functionality and internal design, which end-users nearly never have [11].

Cloned

Sometimes called copies, knock-offs, or advanced counterfeits, *cloned* components are unauthorized recreations of already existing components. The cloning of components is almost always done for reasons of profit, as the design and fabrication of a new component requires considerable time, expertise, and up-front investment, and there is no guarantee that the new component will be commercially profitable [11]. By cloning an already popular and successful component, a malicious actor can ensure that a market will exist for their “product” and simultaneously can take advantage of reduced marketing and customer support costs, as those are handled by the genuine manufacturer. Unlike the previous counterfeiting categories where a malicious entity is working with components that at some point came from an original manufacturer, the creation of a cloned device is entirely done by the counterfeiter. As such, a wide variety of cloning approaches can be taken by the counterfeiter, but most cloned components can be categorized into three flavors: IP theft, emulation, or complete re-implementation [31]. It should be noted that these categories are loose and overlap, with cloned components possibly in a single category or in all three; it mainly depends on the approach taken by the counterfeiter.

IP theft refers to the process of a counterfeiter obtaining component designs created by a genuine manufacturer and using them to fabricate unauthorized copies of the component. The theft of IP from the genuine manufacturer can occur directly with the literal theft of design files or indirectly through the reverse engineering of a component’s silicon layers. Direct theft typically involves the theft (or illegal sale) of layout, netlist, RTL, mask sets, or other design files created by the genuine manufacturer [34]. A counterfeiter will then often modify the stolen files to cover their true origin and adapt the design as needed to different fabrication technologies and process size. Of concern, since the counterfeiter has complete access to the design, it is also possible for the counterfeiter to make functionality changes or even introduce malicious logic. Using these modified files, a counterfeiter will send them off for fabrication as if it was their own design, leaving the foundry completely unaware

that the component they are manufacturing is a clone [35]. Direct theft of IP has become a more common occurrence in recent years, thanks in part to the modern fractured silicon fabrication process, where multiple entities are involved and thus have the opportunity to steal design files [34]. With reverse engineering, a counterfeiter instead reconstitutes the design of a component through the successive imaging of each silicon layer on the die, a process that requires high-magnification microscopes and equipment to precisely sand off each layer of metalization and oxide. The layer images are then stitched together and can be used as the basis for a new set of masks, or alternatively, the images can be used with specialist software to analyze the various features and re-create the original RTL used to create the component [11].

Emulation, on the other hand, is a cloning approach that uses programmable devices, such as microcontrollers, to recreate in software the logic functions originally implemented by the hardware of the genuine device. This phenomenon largely takes advantage of the decrease in cost for mask-programmable and user-programmable microcontrollers, which can be programmed with required functionality and then disguised, via remarking, as a different device [36]. With mask-programmable microcontrollers, the core of the microcontrollers is already completed by the vendor; all that needs to be added by the counterfeiter is a few top mask layers containing the compiled instructions and logic for the component they want to emulate. Compared to other cloning approaches, emulating a device is attractive for counterfeiters as it requires less upfront time and investment since the cloning effort essentially is just writing software and thus can be accomplished by a smaller team.

The last category, re-implementation, is when a counterfeiter creates its own design that attempts to closely match the specifications of an existing component. Depending on the approach taken by the counterfeiter, re-implemented clones can range from true “cleanroom” ground-up re-designs with no re-used IP to designs that heavily re-use IP, which may or may not have been stolen. Of the different cloning methods, re-implementation requires the most effort and resources on the part of the counterfeiter and comes with considerable risk of failure should the newly created implementation not function properly. However, despite these “commercial” risks, over the last decade, more and more high-quality re-implemented clones have been discovered in circulation, with some devices even having better performance characteristics than the original manufacturer’s design, making detection even more complicated [35].

Regardless of the approach taken to create the clone, the main concern with using cloned components is that while they may appear, or even perform, like genuine components, they are not. Cloned components inherently do not carry the same guarantees of quality, reliability, or longevity that are afforded to their genuine counterparts by the original manufacturer. Additionally, since a counterfeiter has a hand in creating the silicon design that goes into

a cloned component, the possibility of adding malicious logic or intentional defects also becomes a concern. It is important to highlight that not all clones are marketed as genuine devices, with some being sold as “compatible” parts carrying their own makings and branding [37]. From a legal perspective, the compatible components are still likely infringing on the IP of the genuine devices. However, they are at least marked to the buyer that the component is not genuine, thus allowing them to make an informed decision as to whether a compatible device would meet their needs.

Maliciously Altered

The possibility of malicious entities making undetected alterations to microelectronic devices is widely considered to be the single most dangerous supply chain threat [11]. Commonly referred to as tampered, back-doored, or trojaned components, malicious alterations to silicon designs are challenging to detect, difficult to attribute, and persistent by default [24]. Due to the numerous entities and hand-offs involved with modern silicon fabrication, the scope of when and where malicious alterations can occur covers nearly the entire manufacturing chain, from HDL design to final die packaging. So far, most evidence for the threat of tampered components has come from academic research, with almost no documented cases of trojaned components being discovered in the field. However, this does not mean tampered components are a rare occurrence; the lack of evidence may simply be due to the difficulty of detecting such components or an unwillingness of victims to publicly discuss being the subject of an attack.

Depending on the attacker’s goal and level of access, varying levels of alteration are possible at the different stages of silicon design and fabrication. At the HDL, RTL, and netlist level, malicious alterations are the easiest and cheapest to perform since it only involves the editing of design files [11]. A malicious entity can easily hide a few thousand additional logic gates within an existing design, which in the case of a modern component, often contains over a billion logic gates. For context on the number of gates required to perform an attack, researchers in 2008 demonstrated a hardware CPU trojan that could dynamically patch a Linux login prompt to create a back-door user account, effectively giving the attacker complete control of the system [38]. To perform this attack, only 1,341 additional gates were added to a CPU design already containing 1,787,958 gates, i.e., a total gate increase of 0.08%. Similarly, in 2009, a separate research team demonstrated an attack that used only 29 additional gates to create a hidden side-channel within a processor, which intentionally leaked information on the secret keys used during encryption and decryption operations [39].

At the layout stage, malicious alterations are still relatively easy to perform, as it only

involves editing artwork files; however, disguising the changes can be more difficult [24]. That being said, during layout, EDA tools will typically add *spare cells*, i.e., unconnected gates, to unused areas of silicon. These spare gates essentially act as an insurance policy, allowing logic bugs discovered after tape-out to be corrected with minimal or no changes to the layout since any additional cells needed for the fix will already be placed. Unfortunately, what is good for the designer is also good for the malicious actor, as the *spare cells* can be re-wired and used to construct hardware trojans [24]. The inclusion of spare cells creates an ideal environment for a malicious actor, as logic can be added to a design without ever touching the source netlist or chaining the total gate count. Furthermore, spare cells greatly complicate the task of verifying that no tampering occurred during netlist to layout translation since layouts that incorporate spare cells will always have more gates than what was originally called for in the netlist.

Alterations to masks, known as *mask editing*, are slightly harder to perform and more expensive, but crucially can still be made [11]. Mask editing has non-malicious applications and is typically performed to improve manufacturing yield or to correct any lithographic issues the masks have in production. For an attacker, mask editing is quite limiting in terms of what can be changed. Adding new gates is challenging and making any modifications to the connective wiring, i.e., the electrical connections between gates, requires making changes across many masks. However, what it lacks in capability, malicious mask editing makes up in how well it can be hidden; locating a minuscule alteration in a mask that contains billions of features is quasi-impossible [24]. To contextualize the danger posed by mask editing, in 2013, researchers demonstrated an attack that compromised the randomness of a secure random number generator by altering the dopant polarity within existing transistors [40]. This extremely stealthy attack only involved a mask alteration on the base silicon layer where dopants are implanted and required no changes to subsequent silicon and metal layers, effectively limiting any possibility of optically checking if an alteration was made.

In general, components fabricated through a fabless design process are more vulnerable to unauthorized semiconductor alterations [8]. This is mainly because the fabless model relies heavily on external entities for large parts of the fabrication process, resulting in a much larger attack surface. Depending on the firm, some fabless companies may do as little as the RTL in-house, leaving everything from netlist generation, layout, and mask creation to a potentially malicious third party [11]. In contrast, for IDM firms, nearly all of the design and fabrication is done in-house, essentially reducing the attack surface down to just insider threats, i.e., malicious alterations done by employees of the company itself.

Furthermore, fabless firms tend to incorporate more externally designed IP, which can potentially include malicious alterations [8]. While IDM firms do also use externally designed IP products, the purchased IP tends to be in the form of RTL or artwork, which, while

tedious, can be audited. The main vulnerability for fabless firms comes from certain types of IP products that are designed and inserted directly by the foundry at the time of fabrication. For foundries, IP products such as electronic fuses, memory, pad rings, and various analog components are part of their competitive edge. As a result, many of these designs are simply not available to audit or inspect by the customer [24]. A fabless firm is essentially required to treat the IP they are buying from the foundry as a black box. As part of the design, empty areas of a predefined size are left in the layout that the foundry later fills with the required IP by means of editing the layout and mask artwork. This creates a perfect opportunity for malicious logic or trojans to be inserted since not only is the foundry altering the layout and masks, but the customer is contractually not allowed to see what was added!

2.2.2 Board-Level Attacks

Since assembled circuit boards contain components, the supply chain threats faced at the board level are essentially the same as those seen at the component level. The accidental inclusion of a counterfeit, compromised, or tampered component will, by extension, threaten the reliability and integrity of the circuit board it is fitted to. For example, should a counterfeit component be fitted to a known-good board, the overall assembly still now becomes untrusted. As these types of threats were discussed in the previous sections, they will not be repeated here; instead, only the threats that uniquely target circuit boards will be introduced.

While it is possible to encounter recycled, defective, and cloned circuit boards within the supply chain, they are rarely seen in practice. Instead, nearly all of the supply chain threats at the board level can be categorized as some form of tampering or malicious alteration. In particular, the threats revolve around *assembled circuit boards* and the components that are mounted to the board, rather than the board layout and construction itself, though this too is possible. Assembled board tampering can be broken down into three categories: the *addition* of components, the *substitution* of components, and the *removal* of components [24]. From the supply chain perspective, this tampering is equally as likely to occur during the stages of assembly or after manufacturing is complete; thus, it is worth first discussing *when* the tampering can occur.

During the stages of assembly, tampering can easily be accomplished by editing the *bill of materials (BOM)*, a critical document used as part of the circuit board assembly process. Prepared by the circuit board designer, a BOM provides an authoritative listing of all components that will be fitted to a PCB during assembly. For each component, a BOM will often contain the part number, quantity, tolerance, reference designator, approved vendor(s), distributor's part number, and even location on the circuit board. Before the assembly process

begins, a customer will send the BOM, unpopulated PCBs, and any specialty components to the assembly house. In turn, the assembly house will then order any needed components on the customer’s behalf and use the BOM as a guide to set up the various assembly steps and machines. As the BOM is essentially the only “source of truth” that an assembly house has for what should be fitted to a board, making changes (assuming they are plausible) can easily be done without detection. If a malicious actor can make changes to a BOM, the consequences can be severe, as any PCBs assembled using that BOM could now have been altered.

Even after a circuit board has been assembled, tested, and inspected, alterations can still be made. Irrespective of the assembled board’s final destination, once it leaves the assembly house, the board is still flowing via a supply chain. Malicious actors can intercept these boards and then make alterations using a technique known as circuit board *rework*, which uses hot air to melt solder leaving minimal trace evidence [24]. Rework can be used to add, replace, or remove components, and is compatible with nearly all component sizes and packaging. For example, it has been extensively documented that the United States National Security Agency (NSA) has rerouted courier packages mid-shipment to NSA facilities so that alterations can be made, a practice known as *interdiction* [41].

Component Addition

Component addition is an alteration where something is added to a PCB without prior approval by the designer. Also known as a *hardware implant*, nearly all component additions are malicious in nature. Hardware implants are highly sought-after by attackers, as they offer a *persistent* foothold on the device, capable of surviving factory resets, firmware changes, and software changes [24].

In theory, hardware implants should be easy to detect, as the alteration leaves behind physical evidence; however, in practice, implants can be challenging to detect. Implants can come in-varying designs and sizes; thus, when inspecting a board it is not clear what exactly to check for, and often implants are hidden within sub-assemblies or other components on a board, requiring destructive analysis to locate. While manual inspection can work for low component count boards, it becomes prohibitively time-consuming for large and complex boards with high and varied component counts. Additionally, it is not always possible to compare a suspect board with a known-good board to rule out the presence of implants. Manufacturing revisions, customer customization, and other genuine changes all limit the ability of locating an equivalent known-good board to compare with [24].

A multitude of implant designs with varying attack capabilities exist. For example, in

2018, a highly sophisticated implant was discovered on the server motherboards used by many large technology companies [42]. The implant was no larger than “a grain of rice” and was capable of modifying the firmware loaded into the [baseboard management controller \(BMC\)](#), an embedded computer fitted to most server motherboards that provides remote management functionality. By changing the firmware, the tightly interconnected [BMC](#) was weaponized to intercept, modify, and sabotage the computations performed on the rest of the machine. It was also documented that the implant was capable of data exfiltration, maintained a connection with a remote command and control server, and allowed the attackers to remote control the server. Similarly, in 2013, an [NSA](#) espionage catalog was leaked, which contained, among a myriad of attack techniques, a set of highly sophisticated implants [43]. The catalog provided an excellent view into the types of implants that could be developed when working with the financial support of a nation-state, with the implants designed to be hidden within [Universal Serial Bus \(USB\)](#) sockets, network sockets, and even video cables.

Implants can also be added to circuit boards by attackers during the design stage. Rather than trying to hide the implant covertly, a *design stage alteration* attack is overt, with the goal of having the implant disguised as just another component on the board. Compared to [BOM](#) swaps, alterations during circuit board design require considerably more effort and planning on the part of the attacker, and often necessitate the cooperation of a rouge employee [24]. However, the reward is that the implant is now part of the design, having its own reference designator and footprint, essentially meaning it can pass any level of board audit or inspection since the design says the implant should be there! *Component piggybacking* is another technique that can be used to add implants to circuit boards [24]. Instead of retrofitting the alternation to the board or editing the board’s design, the implant is designed as a shim that can be inserted in between the circuit board and a component of interest. Acting as a physical man-in-the-middle attack, the shim can now add, remove, or modify the electrical communications seen by both the component and the rest of the board. This style of implant is made possible by recent advancements in package miniaturization, such as [wafer-level chip scale packaging \(WLCSP\)](#), where the silicon die also acts as the external package, making the implant extremely thin and nearly undetectable [24].

Component Substitution

Component substitution is an alteration where a component on a [PCB](#) is changed without prior approval by the designer. Also known as a [BOM swap](#), since this is predominantly how the alteration is made, component substitution is the most common form of board tampering. Detecting swapped components can be especially difficult, particularly for [SMT](#) devices and discrete components, which often have minimal to no markings and come in

minuscule package sizes. Detection is further hampered by the limited accuracy of measuring a component’s value once it has been mounted to a board. Nearby components on the board can mask and change the reading taken with verification equipment, such as a multimeter.

Component substitution can occur for a wide variety of malicious and non-malicious reasons [24]. Swaps sometimes occur by accident and negligence, such as if there is a mistake in the BOM or if the assembly house orders the wrong part number. Accidental swaps can also occur when setting up the pick-and-place machines; reels can easily be put into the wrong slot on the machine, and because many of the components have the same package, the machine will not alarm during assembly [24]. Components can also be swapped legitimately for reasons of availability and cost. For example, when a customer contracts an assembly house to assemble a board, the accompanying BOM will be reviewed to determine which components need to be ordered. Often, larger assembly houses keep stock of common components and their compatible equivalents, making them available to their customers as part of a turn-key option, avoiding the additional effort of component sourcing. The turn-key option can be very advantageous and economical for the customer, especially for simpler components on a board, such as resistors, capacitors, and other *jellybean* components. As these components are available from multiple manufacturers, it is acceptable for part numbers to differ from the ones listed in the BOM, as long as the value, tolerance, etc., is correct.

Unfortunately, the profit motive has caused some assembly houses to abuse the turn-key model. There have been numerous incidents where an assembly house swaps parts without customer approval and mounts compatible, lower cost, or even counterfeits components to boards, while simultaneously billing the customer for the market price for parts listed in the BOM and pocketing the difference [24]. Other times, assembly houses have misrepresented the source of their components, telling customers that only approved distributors or vendors were used, when in fact, the components were procured from gray market brokers [44].

For malicious actors, component substitution is an attractive attack method because it is easy to execute, it simply requires the editing of a BOM, and allows for wide changes in functionality to be made. Because components are often available in many package types, many substitution attacks can be performed without changing the underlying PCB layout or component footprint. For example, various document leaks from the NSA have shown that the agency actively develops “replacements” for certain security-related devices, such as those that store cryptographic keys [43]. These malicious devices have the same packaging and footprints as the components they are replacing, allowing for easy and covert substitution. Similarly, many devices use the presence of resistors to specify boot options, special configuration, maintenance, or debug modes for products [24]. These special modes often have fewer security features enabled and thus are attractive for attackers. Because the presence of a resistor, or lack of one, is simply a line item specified in a BOM, they can easily

be changed with a simple edit. Attackers can swap the resistors for ones that measure 0Ω or have very high resistance, i.e., a short or open, so electrically, the value is now different, but under visual inspection, the appearance is the same.

Component Removal

Component removal is a malicious alteration where a component on a PCB is removed without prior approval by the designer. At first, component removal may seem counterintuitive. Typically, an attacker adds or changes components to create an attack vector, so what can an attacker stand to gain by removing something? Additionally, removing any components surely would cause the circuit board to become non-functional or fail quality control testing. The reality is that component removal is often done in a very subtle manner, typically with the goal of making a circuit board leak more information or making it susceptible to attack [24]. An attacker can remove components from a board in two manners, manually via hot air rework and a soldering iron, or more efficiently by removing them from the BOM so that they are never fitted in the first place. It should be noted that it is not uncommon to see unpopulated components on board, so even with a visual inspection, the missing components may not be flagged as anomalous, allowing the tampered boards to slip through undetected.

As an example of a component removal attack, consider the *decoupling capacitors* are used on circuit boards for noise control and compliance with *electromagnetic interference (EMI)* regulations which are often the smallest components on a board. Typically, they are not critical for the operation of circuits, and if they are removed, most devices will continue to work without issue. By removing the decoupling capacitors, an attacker is able to increase the overall noise generated by a system, enabling certain side-channel based attacks, such as secret key recovery, to be performed faster and easier [24]. Simultaneously, the removal of decoupling capacitors makes the circuit board more sensitive to external interference, opening the door to glitching attacks [45].

Along similar lines, by removing certain components used in radio transmitter filters, a tracking or de-anonymization attack can be performed. Devices that communicate wirelessly, such as smartphones, need to meet various regulatory requirements, such as power output and frequency of operation, in order to minimize harmful emissions. By removing certain passive components, such as inductors and capacitors, that are used to shape the spectral output, i.e., filter out unwanted frequencies, an attack can effectively make a device trackable since it will now have a unique transmission signature [24].

2.2.3 Current Mitigation Efforts

In response to the multifaceted threats introduced in the previous section, manufacturers, distributors, and vendors of electronic devices have developed a set of counterfeit detection techniques. While varied in their approach, fundamentally, these techniques all attempt to determine the authenticity of a device by finding the presence of one or more anomalies.

Although substantial research and effort has gone into developing these supply chain threat mitigations, challenges remain. Currently, many methods are time and labor-intensive, needing many hours to authenticate just one complex component [32]. Additionally, many of the methods are qualitative in nature and are heavily reliant on human experts to interpret if the detected anomalies stem from manufacturing quality issues or are actual indicators of counterfeiting. In practice, many of the techniques are not 100% reliable and thus can not be used in isolation [32]. Obtaining a reliable authenticity result generally requires the use of multiple different techniques and evaluating the results holistically. Moreover, nearly all methods rely at some level on comparing the suspect component to a known-good component, i.e., a component that has been positively identified as being genuine. For modern devices locating known-good components can be done with relative ease, but for components that are no longer manufactured, this can become a real challenge [3].

The following subsections aim to introduce some of the more common risk mitigation techniques and discuss their advantages and shortcomings. Broadly speaking, the methods can be grouped into three main categories: physical inspection, imaging techniques, and electrical tests [32]. Physical inspection focuses on the external appearance of the component, looking at markings, package size, pins, and overall condition. Imaging techniques are a broad category but essentially focus on inspecting the silicon die within the component, either optically or through X-ray radiation. And finally, electrical tests, as the name implies, evaluate how well a component performs electrically against its specification.

Physical Inspection

Physical inspection is a counterfeit detection technique based on the examination of the external structure and markings of a component. Physical inspection is generally the first method used when determining the authenticity of a device due to its low cost and speed [11]. By checking for conspicuous signs of counterfeiting first, entire batches of devices can quickly be examined, leaving only the less obvious counterfeits to be tested by more expensive and time-consuming methods. Physical inspection, when done correctly, is nondestructive and does not affect the certification or the reliability of the component [46].

Depending on the level of thoroughness required, a physical inspection can range from checking for a few types of anomalies to a nearly exhaustive list; but the general idea is to look for anything that might be out of place [11]. An inspection will typically begin with the markings on the package, making sure that any required logos, part numbers, date of manufacturing, lot code, etc., are present, valid, and have no imperfections [47]. Using a low-power microscope, the component's condition is then inspected through a multi-step process. Starting with the texture and finish of the component, signs of blacktop and sanding are checked for, as these almost always signify the component has been remarked [47]. Similarly, the presence of scratches on any of the surfaces could indicate that the component was previously used. Next, the pads or pins of the component are checked for any signs of recycling, such as solder residue or deformed leads [32]. A thorough examination will also include an inspection of any cavities or indents on the package, checking the location, depth, and condition for signs of blacktopping. Cavities, in particular, are by-products of the plastic molding process and counterfeiters have difficulty recreating them or keeping them clean during resurfacing [47].

As part of checking for remarked devices, a solvent test may also be performed to test the permanence of any markings and the presence of blacktop [27]. Using a cotton swab, diluted acetone or a mixture of mineral spirits and alcohol is applied to the surface of the component [5]. If the solvent removes any surface material or causes the markings to streak, the component has been remarked. Physical inspection can also include the packaging material the components are shipped with, such as the shipping box, and more importantly, the tape, reel, or tray that holds the components. For example, the lack of ESD packaging or humidity indicator card for components that are sensitive to moisture can be indicative of counterfeiting, as these are often forgotten by counterfeiters [48].

One of the limitations of physical inspections is their substantial dependence on subject matter experts to interpret whether the anomalies found are valid indicators of counterfeiting [46]. Physical inspections also subject the components to additional handling and moving, increasing the risk of accidental damage over time. Like most counterfeit detection techniques, physical inspections rely on having a known-good component to compare with, which, if not obtainable, limits the tests and authenticity results that can be obtained [46]. Additionally, the increase in high-quality counterfeits has also placed some limits on the efficacy of physical inspection. Recently, some counterfeit manufacturers have started *repackaging* silicon dies to disguise the components recycled or remarked origin [27]. Repackaging involves extracting a die from its original package, wire bonding it to a new lead frame, and then encapsulating it with a new package. If done with sufficient attention to detail, these repackaged components can be made to look essentially new, fooling nearly all physical inspection tests.

Optical Die Inspection

While other counterfeit detection methods rely on external cues in order to guess what lies inside a component, optical die inspection evaluates the authenticity of a component by directly inspecting the silicon die. Also known as decapsulation, decapping, and delidding, optical die inspection involves removing a portion of the component's packaging so that the internal die is visible. Using a microscope, various die features and markings are then examined to make a determination as to the authenticity of a component under evaluation. While decapsulation is a destructive technique, it is one of the only ways to have a definitive confirmation as to what truly is inside a component.

There are various ways that a component can be decapped, and the method used generally depends on the packaging style of the component. For the standard plastic or resin packaging, many methods can be employed, namely: chemical etching, laser etching, and mechanical grinding [49]. *Chemical etching* uses heated corrosive acids to dissolve away the area of interest on the package. *Laser etching* is similar to chemical etching, except that it uses a laser source to essentially evaporate the material. *Mechanical grinding* uses cutting, drilling, and grinding tools to physically create a pocket within the packaging to reveal the silicon die. For metal packages, laser etching is typically used to cut a hole in the top of the device. For ceramic packages with a metal cover, the metal lid can be simply removed to expose the die. For entirely ceramic packages, tedious mechanical grinding, typically with diamond tooling, is required as the ceramic is resistant to chemical etching. Alternatively, laser etching can be used to decapsulate ceramic packages; however, this requires careful cooling of the component, as heat dissipated by the laser energy can lead to temperature gradients on the package and potentially cause damage [49].

Once the silicon die has been exposed, it is typically visually inspected under high power magnification to check for the presence of any anomalies. An authentic or counterfeit determination can then be made by inspecting the die markings, size, orientation, and any discerning features present on the die itself. If a component is decapsulated wide enough, the visual inspection can also include looking for anomalies with the lead frame and bond wires. Depending on the thoroughness of the evaluation, successive sanding of the metalization layers can be performed to inspect each layer within the die [11]. However, in practice, metalization removal is not necessary, as most cloned devices can be quickly be caught by differences in die markings or die size. Similarly, with most remarked devices, an inspection of the markings on the top layer of the die is often sufficient to determine if it matches the exterior markings on the packaging [48].

One of the most significant drawbacks of decapsulation is that it is a destructive verification method; once the component has been verified to be genuine, it is no longer usable.

Because of this, in practice, decapsulation is used on a sampling basis; for example, from a batch of 50 devices, only 2 or 3 would be optically inspected [11]. Additionally, die inspection is not always effective against the threat of remarked components, as differences in part grades are not always due to differences in the silicon design. For example, it is common for commercial and industrial grade versions of a component to contain identical silicon die, with the difference between the two determined by quality-control tests and binning [48].

X-Ray Inspection

X-ray inspection is an approach that uses X-ray radiation to create images of the interior of a device. When a component is placed under an X-ray source, the various materials that compose the component will attenuate X-ray radiation by different amounts. By capturing the X-ray radiation that is not absorbed, i.e., not blocked by the sample, an image can be formed that reveals the structure and materials within the component. This image can then be inspected for a multitude of anomalies in an attempt to determine the component’s authenticity.

Two types of X-ray imaging are used in the detection of counterfeits and tampering: *radiography* which creates a two-dimensional image of a component, and *tomography*, also known as *computed tomography (CT)*, which creates a three-dimensional representation of a component from multiple two-dimensional images [31]. Generally, the two-dimensional images and three-dimensional representations perform equally well when looking for large anomalies related to the bond wire configuration, lead frame, and the die itself [31]. This can include anomalies such as broken or missing bond wires, incorrect or missing lead frame, incorrect die size, location, or rotation, and if the device is missing a silicon die entirely [33]. However, for smaller or more subtle anomalies, these become easier to detect with the CT representations since it includes depth information. In particular, anomalies related bond wires or the presence of fractures within the silicon can only be seen in three dimensions [50].

One of the limitations of X-Ray inspection, whether radiography or tomography, is that when used to determine a component’s authenticity, it requires a known-good component to compare with [51]. For X-Ray inspection, a known-good component typically needs to match the lot code of the device under inspection so that legitimate changes in manufacturing over time can be controlled for. Not having a comparative part is acceptable when only looking for signs of damage, i.e., broken bond wires or chipped die, which may indicate the component was recycled. Additionally, X-ray inspection is generally thought to be nondestructive since it does not require the decapsulation of silicon die or the modification of a package, but this is not always the case. Recent studies have revealed that certain components, namely flash memory, can permanently degrade under larger doses of x-ray radiation [50]. This is of

concern for inspections that use three-dimensional CT scans since it subjects components to irradiation for long periods of time due to the multiple X-ray images required from several different angles.

Electrical Tests

In contrast to the previous methods that focus on the physical aspects of a component, electrical tests aim to evaluate the functionality and performance of a device. Electrical tests are typically nondestructive and can be broken into two main types: functional tests and parametric tests [11].

Functional tests are relatively quick tests that evaluate the component’s functionality and verify that everything is working and operating as designed. The functional tests performed on a component will vary depending on its intended function. For example, memory devices may be subjected to complex read and write cycles, while simpler logic gates could be tested against its expected truth table. The principle of counterfeit detection with functional tests works on the basis that counterfeit or substandard components may fail some of the tests, while genuine devices should have no problem. Functional tests are considered to be one most efficient techniques for evaluating devices, providing a quick “go/no-go” if a component is suitable for use [32]. Typically, a genuine component is not required to establish the expected results for functional tests; however, having one can be helpful when setting up the tests and ensuring they work correctly.

Conversely, *parametric tests* subject the component to a thorough electrical evaluation performed over a wide range of voltages, signals, and operating conditions [32]. Parametric tests aim to establish a dataset of various parameters containing the range of performance values that genuine components can attain. Then when given a suspect component, it can be tested to see how statistically close it is to the genuine devices. The parameters measured as part of the evaluation vary from device to device but can include power consumption, output power, noise, breakdown voltage, rise and fall times, prorogation delay, etc. [11]. Included within parametric tests is a sub-class of techniques that use *side-channel* information, i.e., information that is accidentally leaked or indirectly obtained from a component during its operation, rather than direct measurements, to determine the authenticity of a device. Side-channel techniques operate similarly to traditional parametric methods and determine authenticity by comparing how a suspect component’s side-channel response differs from the side-channel response of a known-good component. Examples of side channels that have been successfully used to detect counterfeit devices include power analysis [52], near-field electromagnetic emissions [53], and radiated clock timing emissions [54].

Parametric tests are one of the only techniques that can catch recycled components; since as silicon devices age, various parameters, such as its curve trace, will no longer match those from genuine devices [32]. Unlike functional tests, parametric tests do require a set of known-good components to compare with, as data-sheet performance values are typically not detailed enough. This leads to the problem of locating many known-good devices to build the initial dataset, which can be a real challenge for components that are no longer produced.

One of the limitations of electrical tests is that they can not detect all types of counterfeit components. Cloned, maliciously altered, and even recycled devices can all pass electrical testing as long as their parameters are within specification [11]. Electrical tests also have high setup costs, essentially requiring custom test fixtures and measurement setups for each device model. Parametric tests, in particular, have a limited ability to discern naturally occurring manufacturing variation from actual changes introduced by a malicious actor. Differences between foundry, technology process, and silicon revisions can all lead to differences in the performance of a component [32]. Depending on the number and thoroughness of the parameters evaluated, the time required to perform the parametric tests can quickly balloon, sometimes leading to test times in the tens of hours for a single device [11].

2.3 Microwave RF Overview

Within the field of electrical engineering, the term **radio frequency (RF)** refers to the study and application of oscillating non-ionizing electromagnetic radiation. This radiation is principally defined by its *frequency*, measured in **hertz (Hz)**, which quantifies the number of oscillations that occur in one second. The frequency range, or *spectrum*, of **RF** radiation is immense, with signals ranging anywhere from 3 kHz to over 300 GHz [55]. The term *microwave* refers to a subset of this spectrum, namely the frequencies between 300 MHz and 40 GHz [56]. Because of the high frequencies involved, and thus small wavelengths, traditional circuit analysis techniques do not hold at microwave frequencies. In particular, it becomes challenging to measure voltages and currents directly, and reference “grounds” are no longer uniform. As a result, microwave networks, i.e., circuits, are instead defined in terms of waves and power flows. The following sections aim to introduce some of the microwave **RF** concepts utilized by SILGEO.

2.3.1 Scattering Parameters

In microwave theory, *scattering parameters*, often called “s-parameters”, are used to quantify how RF power propagates through multi-port networks. As an aside, the term *port* in this context means a current-carrying electrical terminal that connects a network to the external world [56]. The principle behind scattering parameters is the following: when an RF wave is incident on a port, a portion of that wave’s power will be reflected back. The remaining non-reflected power of the wave enters the incident port and *scatters*, i.e., propagates, through the network, sometimes being attenuated or amplified, and finally exits at all or some of the other ports of the network [56]. Scattering parameters capture this input-output relationship of the network and are calculated as the ratios of the incident and reflected power waves [56]. For simple circuits, scattering parameters can be obtained manually by measuring the standing waves and computing the ratios. However, this process is tedious and slow; today, measurement devices, such as [vector network analyzers \(VNAs\)](#), can be used to obtain the scattering parameters for a [device under test \(DUT\)](#) directly [57].

Scattering parameters allow arbitrarily complex networks and circuits to be represented as a “black box” and thus are equally applicable for representing the effects of a single component or a network of multiple interconnected devices. In essence, scattering parameters allow the external behavior of a network to be modeled accurately without any regard as to the underlying composition of the network. This also allows the interaction of multiple networks to be computed analytically, such as how an amplifier performs with a particular antenna, without ever having to *physically* connect or measure any devices.

2.3.2 Scattering Matrix

The scattering parameters for a network are commonly displayed in a matrix format, with an n port network having n^2 scattering parameters. Equations 2.1 and 2.2 show the scattering matrices for a two-port and four-port network, which respectively contain 4 and 16 scattering parameters.

$$[S] = \begin{bmatrix} S_{11} & S_{12} \\ S_{21} & S_{22} \end{bmatrix} \quad (2.1)$$

$$[S] = \begin{bmatrix} S_{11} & S_{12} & S_{13} & S_{14} \\ S_{21} & S_{22} & S_{23} & S_{24} \\ S_{31} & S_{32} & S_{33} & S_{34} \\ S_{41} & S_{42} & S_{43} & S_{44} \end{bmatrix} \quad (2.2)$$

Within the scattering matrix S , each element (S_{11} , S_{12} , S_{21} , etc.) represents a scattering parameter for the network under test. The element subscripts are read in the following manner: for a parameter S_{ij} , the j subscript indicates the network input port that is excited by a wave, while the i subscript indicates the output port where the wave amplitude is measured [56]. For example, S_{11} refers to the ratio of the amplitude of the wave that is incident onto Port 1, to the amplitude of the wave that is reflected by Port 1. Similarly, S_{21} refers to the ratio of the amplitude of the wave that is incident onto Port 1, to the amplitude of the wave that is transmitted and exits from Port 2. The scattering parameters that have $i = j$, such as S_{11} or S_{22} , are called *reflection coefficients*, as they describe how much of a wave is reflected when it is incident on a port [56]. Denoted by the symbol Γ , reflection coefficients essentially indicate how much power is rejected by a port due to impedance discontinuities. All other elements, i.e., $i \neq j$, are called *transmission coefficients* as they describe how the wave, and thus power, travels through the network to the various ports [56]. Transmission coefficients are denoted by the symbol T .

As microwave circuits are complex, i.e., contain components that affect both the magnitude and phase of a wave, scattering parameters are correspondingly complex numbers. It is also important to note that the scattering matrix and its parameters capture a network's behavior at a single frequency and system impedance, typically $Z_0 = 50 \Omega$ [57]. For non-ideal networks, when the frequency or system impedance is varied, the network's response will as well. In practice, it is useful to know the network's response over a range of frequencies; thus, it is common to characterize a DUT at multiple frequencies and create a scattering matrix for each one. Essentially, each scattering parameter (S_{11} , S_{12} , etc.) is treated as a one-dimensional array containing a value for each frequency of interest.

2.3.3 Two-Port Networks

The relationship between power waves and scattering parameters is best illustrated when considering the case of the two-port network. As the name implies, a two-port network is an arbitrarily complex microwave network, with only two external ports. Two-port networks are an important building block of microwave analysis because of their wide applicability. Many commonly used components either have two terminals, such as attenuators, or have two

relevant terminals, such as with amplifiers where only the input and output are of interest. Additionally, two-port networks serve as a starting point for the analysis of more complex networks with higher port counts.

Figure 2.2 provides an example of a two-port network, having a scattering matrix S . Waves a_1 and a_2 represent the incident waves, and b_1 and b_2 represent the reflected waves. By convention, power waves incident onto a network are designated as a_n , and the reflected power waves as b_n [56].




Figure 2.2: Two-Port Network With Incident and Reflected Power Waves

Using Figure 2.2, the relationship between the scattering parameter, incident waves, and reflected waves can be obtained. Equation 2.3 shows the matrix setup for the derivation, where a pair of incident waves interacts with a two-port network S , resulting in a pair of reflected waves.

$$\begin{pmatrix} b_1 \\ b_2 \end{pmatrix} = \begin{bmatrix} S_{11} & S_{12} \\ S_{21} & S_{22} \end{bmatrix} \begin{pmatrix} a_1 \\ a_2 \end{pmatrix} \quad (2.3)$$

Using matrix expansion, generalized linear equations for the reflected waves can be obtained.

$$b_1 = S_{11} \cdot a_1 + S_{12} \cdot a_2 \quad (2.4)$$

$$b_2 = S_{21} \cdot a_1 + S_{22} \cdot a_2 \quad (2.5)$$

Note that the reflected waves in Equations 2.4 and 2.5 contain contributions from both ports in the network, not just the initial port where the wave was incident [56]. E.g., the reflected wave b_1 is created by both the reflection of a_1 at Port 1, but also the transmission of a_2 from Port 2 through the network to Port 1. In simpler terms, the reflection seen at a port is dependent on itself *and* the state of any other ports in the network. By rearranging, the definitions for each scattering parameter can be obtained.

$$S_{11} = \left. \frac{b_1}{a_1} \right|_{a_2=0} = \Gamma_1 \quad (2.6)$$

$$S_{21} = \left. \frac{b_2}{a_1} \right|_{a_2=0} = T_{21} \quad (2.7)$$

$$S_{12} = \left. \frac{b_1}{a_2} \right|_{a_1=0} = T_{12} \quad (2.8)$$

$$S_{22} = \left. \frac{b_2}{a_2} \right|_{a_1=0} = \Gamma_2 \quad (2.9)$$

Where Γ_1 and Γ_2 are the reflection coefficients for Port 1 and Port 2, respectively. T_{21} is the forward transmission coefficient for Port 1 to Port 2, and T_{12} is the reverse transmission coefficient for Port 2 to Port 1. Observe that the parameters are evaluated at either $a_1 = 0$ or $a_2 = 0$; this is fundamental to the definition of scattering parameters. When measuring a network, any ports not involved in the measurement must be terminated by the system impedance Z_0 , or *matched*, to minimize the possibility of stray reflections [56]. This is because if a wave incident on Port 1 travels through the network and exits via Port 2 into an impedance discontinuity, a portion of that wave will be reflected back through the network to Port 1, essentially “increasing” the reflection coefficient for Port 1 (S_{11}). If unconnected ports are not terminated, or are terminated with the wrong impedance, the scattering parameters will be incorrect.

Notice that the scattering matrix for a two-port network has both a *forward* and *reverse* transmission coefficient. While generally $T_{21} = T_{12}$, it is perfectly valid for $T_{21} \neq T_{12}$, indicating that the network’s transmission effects are *directional*. At lower frequencies, directionality is rare and mainly only seen with active components, i.e., an operational amplifier, but at microwave frequencies, even purely passive components can be directional, a property that is heavily used by couplers and other microwave components [56].

From the scattering parameters, additional microwave parameters can be obtained. *Insertion loss* quantifies the additional attenuation, or loss, between any two points in a network. Measured in **decibels (dB)**, insertion loss is typically used to indicate the additional loss that will be incurred by adding a component to the signal path. The formula for insertion loss can be seen in Equation 2.10.

$$IL = -20 \cdot \log_{10}|S_{21}| \text{ dB} \quad (2.10)$$

Return loss quantifies, in dB, the amount of power reflected by a network. This is typically caused by a discontinuity or impedance mismatch between the network and the circuit it is connected to. In essence, it is the amount of power the network “rejected” and sent back to the source. The formula for insertion loss can be seen in Equation 2.11.

$$RL = -20 \cdot \log_{10}|S_{11}| \text{ dB} \quad (2.11)$$

Chapter 3

SILGEO Overview

SILGEO, a portmanteau of *silicon* and *geometry*, aims to provide a solution to the [time-of-check to time-of-use \(TOCTOU\)](#) dilemma facing hardware devices. In short, the [TOCTOU](#) problem arises from the inability of verifying the integrity of a hardware device at the time of use. Instead, end-users are reliant on the attestation of the manufacturer or designer, made at some previous point in time, that the hardware is conformant. This results in a window of time where malicious entities can tamper with hardware knowing that end-users have limited recourse to detect.

Whereas with software, many opportunities exist for verifying the authenticity of files, libraries, or entire systems. For example, source code files can easily be verified for data integrity or tampering with *checksum algorithms*, a procedure which generates a cryptographically secure and repeatable mathematical sum from a given block of data. By comparing the checksum of a given file with a known-good checksum, authenticity can quickly be established. The same can be done for external libraries or even entire executables; a checksum can be computed and then compared against publicly available known-good values.

For hardware, the challenge is that no such checksum capability exists. Instead, hardware authenticity is established via a haphazard set of anecdotal observations and time-consuming techniques that attempt to detect the presence of anomalies. SILGEO attempts to rectify this situation by providing a *hardware fingerprinting* capability that can provide a rapid determination of an electronic device's authenticity.

3.1 Problem Statement

The problem statement that drove the development of SILGEO is best described as follows: Given a set of arbitrary, blackbox electronic targets, of which some are authentic and some are counterfeit, separate all targets into the appropriate category. The solution should not damage or affect the integrity of the target, or require the removal of the target from its circuit board or larger assembly. The solution should be efficient and operate reliably without any user input or guidance.

To be precise, *target* means the electrical device whose authenticity is under evaluation, also known as a [device under evaluation \(DUE\)](#) or, more commonly, a [device under test \(DUT\)](#). The solution should be applicable to discrete components, microchips, assemblies, such as circuit boards, and complete devices. *Arbitrary* means the solution should be compatible with common fabrication processes, packages, layouts, and gate technologies utilized to manufacture targets. It should support both analog and digital circuits, and targets constructed from non-silicon semiconductors, such as germanium or gallium arsenide. *Blackbox* means the solution should require no disclosure of proprietary information on the function or use of the electronic target. *Authentic* means a genuine target produced by the original manufacturer. *Counterfeit* means a target not produced by the original manufacturer and may have been cloned, remarked, or altered.

3.2 Theory of Operation

Conceptually, SILGEO operates similarly to, but not exactly like, the technique of *system identification*, where a mathematical model is constructed for blackbox system by means of statical methods and collected measurement data. To determine the authenticity of a target, a set of frequency-domain measurements are taken of the target to collect the key identifying electrical characteristics. The characterization process is blackbox and only requires access to the electrical interfaces of the target; no knowledge of the target’s function or internal construction is needed. Notably, the characterization process is performed with the target *unpowered*, further avoiding the need for any target-specific knowledge on supply voltage or logic voltage levels. The collected measurements are then processed and cleaned, to create what is essentially a hardware fingerprint of the target. Depending on the application of SILGEO, this fingerprint can be used and analyzed in different manners. For the detection of counterfeits, the fingerprint can be compared for similarity with known-good examples or a collection of other devices. In the case of detecting alterations or tampering, the fingerprint of the target can be compared with previous fingerprints of the same target.

The theory underpinning SILGEO, is that traveling electromagnetic waves are affected by the *medium* within which they travel. The physical characteristics of a medium, such as dimensional geometry, relative permittivity, and impedance, contribute to the medium’s propagation characteristics [56]. In simple terms, propagation characteristics describe how a wave will *change* as a consequence of moving through a medium. The propagation characteristics of a medium can be quantified in different ways, but most commonly used are the medium’s propagation velocity v_p , which describes how fast a wave travels in the medium, and reflection and transmission coefficients, describing how power waves act when incident onto a medium. Importantly, should any part of the medium change, the propagation characteristics will change as well. Within the context of RF, the term medium typically refers to a transmission line of some sort, i.e., cables, planar PCB traces, or free space, which are specifically designed to propagate waves with minimal power losses and distortion. The key to SILGEO’s fingerprinting approach is to treat the electrical target under evaluation, which is likely not an actual transmission line, as if it was, in fact, a transmission line. The assortment of semiconductors, electronic components, wires, and circuit boards that make up an electrical device all act as mediums, albeit poor ones, for electrical signals to travel. The transmission properties of these “mediums” can then be measured and characterized in the same manner that any other transmission line, such as coaxial cable, would be.

SILGEO’s ability to determine the authenticity of a target relies on two principles. Firstly, targets that are physically identical will have identical wave transmission characteristics. This means that two semiconductors from the same lot or two circuit boards from the same batch should have the same transmission behavior, after controlling for natural manufacturing variations. By extension, this also means that targets that are *not* physically identical will *not* have identical wave transmission characteristics. A genuine component and its counterfeit counterpart are physically constructed differently and thus will have different transmission behavior. Secondly, the alternation of a target causes the target, or more precisely: the target’s medium, to have slightly different wave transmission characteristics. Alterations come in different forms, but all involve changing some physical aspect of the target, which in turn causes the transmission behavior to change.

To measure these changes in transmission medium behavior and thus create the hardware fingerprint, SILGEO applies the frequency-domain approach used to obtain scattering parameters. Waves of varying frequency and power are generated and then injected into the target. As the incident wave interacts with the internals of the target, a set of reflection and transmission waves are created. Using sophisticated equipment, these waves are captured, and the final scattering parameter ratios representing the target are computed. SILGEO essentially treats the electrical target as if it was any other RF device that needs to be characterized.

3.2.1 Single Port Characterization

Before detailing the characterization process for targets with multiple electrical interfaces, the characterization process for a single electrical interface, or *port*, on a target will be considered first, as it provides a foundation for the more complex measurements. Figure 3.1 shows the measurement setup, with the component labeled “IC” representing the target. Prior to commencing the measurement, the selected electrical interface on the target, in this case, a pin, is connected to a *signal generator* and a *signal capture* device via high-frequency RF grade cables. The signal generator is configured with the desired power and frequency parameters needed to characterize the DUT. The selection of these parameters and their effects is further detailed in Section 3.3. At the time of measurement, the signal generator creates an RF *frequency sweep*, shown in red, consisting of a single continuous tone changing in frequency or multiple discrete frequency tones in rapid succession.




Figure 3.1: Reflection Response From a Single Pin

This frequency sweep is then applied to the port of the target, termed the *incident port*, as this is where the sweep is incident upon. As the frequency sweep moves into the target’s port, a portion of the power will be reflected, shown in blue. The *amount* of power reflected is frequency-dependent and is intrinsically linked to the impedance characteristics and wave propagation properties of the medium within the target [56]. In essence, the portions of the frequency sweep that are reflected now contain information about the target’s internal physical composition. In the single port case, the reflected power returns via the same electrical pin that the sweep was incident upon, where it is then sampled by the signal capture device. Not shown in the figure is a connection from the signal generator to the capture device that allows the original incident sweep to also be sampled for use in computing scattering parameter ratios.

The result of this measurement is the reflection coefficient S_{11} , at each frequency of interest in the sweep, for that particular pin on the target. Because each electrical interface on a target is a different “perspective” on the transmission medium with the target, this single

interface characterization can also be repeated on the other pins of the target, yielding a unique set of reflection coefficients for each pin. Additionally, because the response of a target is often non-linear with respect to incident power, i.e., the reflection obtained from a target changes with the amount of incident power used, the characterization process can be performed multiple times at different power levels, again yielding additional unique sets of reflection coefficients for the target.

Throughout the characterization process, the incident power levels are carefully controlled. The power levels need to be high enough that a response is elicited from the target but simultaneously need to be kept below the logic threshold levels of the target to avoid activating any of the target’s functionality or damaging the circuits within and potentially altering the measurement. Additionally, the power levels need to be kept constant across the different characterization frequencies used within the sweep to ensure that effects related to power level non-linearity do not accidentally distort the measurements.

3.2.2 Multiple Port Characterization

Since targets often have multiple electrical ports, using a characterization process restricted to only measuring a single port at a time would be inefficient and have limited practicality. The simultaneous characterization of multiple ports is not only faster than the single port approach, but also allows a more robust verification of the target to be performed, as it captures additional uniquely identifying information from the target that the single port approach can not. The process for characterizing multiple ports on a target is quite similar to the process for a single port, reusing the same equipment and setup, with the only difference being where the signal capture device is connected to the target.

Consider the case of characterizing two ports, which for the target shown in Figure 3.2, labeled “IC”, are Pins 1 and 3. Similar to the single port example, the incident frequency sweep is shown in red. However, rather than connecting the signal capture device to the same port where the frequency sweep is incident upon, it is now connected to a different electrical pin on the target. In doing so, the captured frequency sweep response, shown in blue, will have traveled *through* the target and interacted with much more of the target. This response represents the transmitted power, and similar to the reflected power, the *amount* of power transmitted through the target is frequency-dependent and is intrinsically linked to the impedance characteristics and wave propagation properties of the medium within the target. The transmitted power can be thought of as the complement of the reflected power; any incident power that is not reflected will be transmitted to the other ports on the target. For the fingerprinting of a target, the transmitted power can be very important, as

it will often contain a response from areas within the target that the reflected power does not interact with. Though it varies from target to target, a target's reflection behavior is typically determined by the impedance effects at the edges of the target, meaning that the reflection response will not have entered very far into the target, whereas the transmission response will have.




Figure 3.2: Transmission Response Through Two Pins

After the transmitted sweep has been captured, it can then be compared with the original incident sweep to compute the transmission coefficient S_{21} , at each frequency of interest in the sweep and for the particular pair of ports chosen on the target; in this example Pin 1 and Pin 3. Though not shown in the figure, the reflection sweep seen in the single port example is still, in fact, present when measuring two (or more) ports, due to the properties of microwave propagation within networks, as discussed in Section 2.3.3. By adding a second connection from the signal capture device to the port where the sweep is incident, shown in Figure 3.3, the reflection of Pin 1 and transmission from Pin 1 to Pin 3 responses can be obtained simultaneously from the target.




Figure 3.3: Simultaneous Reflection & Transmission Response From Two Pins

Note that the measurement setup shown in Figure 3.3 is *unidirectional*, in that the reflection of Pin 3 or the transmission from Pin 3 to Pin 1 is not obtained. However, with

the addition of a directional switch to the measurement setup, allowing the signal generator and capture device to swap the target pins they are connected to, the characterization process would be able to capture the effects in the reverse direction. In doing so, the two remaining scattering parameters, S_{22} and S_{12} , would be obtained, thus fully quantifying the propagation characteristics of that two port subset of the target.

For targets with multiple electrical ports, many port combination subsets can be created, such as Pin 1 to Pin 3 as shown in Figure 3.3, Pin 1 to Pin 5, Pin 8 to Pin 15, and so on. For a target with 16 pins, there are a total of 120 two-port combinations that can be made, with each combination effectively acting as its own distinct medium due to the pins providing different propagation perspectives on the target. This means that each port combination will have its own individual response characteristics, yielding considerable identifying information on the physical construction and transmission medium within the target.

Port combinations are by no means limited to just two ports, with one incident port and one response port. Port combinations can be created with multiple incident ports, multiple response ports, or both. Figure 3.4 shows an example where the frequency sweep is simultaneously incident upon two pins of the target, and the signal capture device is connected to three pins. In doing so, supplementary unique information from the target can be obtained, with the multiple incident sweeps interacting together and with the target's medium, creating a sort of hardware-dependent convolution response.




Figure 3.4: Characterization of Multiple Pins

3.2.3 Determining Authenticity

Once a target has been characterized, its fingerprint can then be analyzed and compared to determine the target's authenticity. Depending on the application of SILGEO, there are two approaches as to how authenticity can be determined: *verification* and *classification*. Verification, also known as one-to-one matching, is when a target's fingerprint is compared

to a previous fingerprint of the same target to see if any changes occurred. Verification is typically used in tamper detection applications, such as checking if any unauthorized additions or modifications were made to a circuit board during shipment. For verification tasks, SILGEO does not require fingerprints or characterizations from other targets; all that is required is a previous fingerprint of the target in a known-good state. To verify the authenticity of the target, SILGEO compares the current fingerprint to the fingerprint from the known-good state and calculates a similarity score indicating the likelihood of tampering.

On the other hand, classification, also known as one-to-many matching, is when a fingerprint from a target of an unknown class, in this case, authenticity, is compared with the fingerprints from other similar and different targets, in order to see which class the unknown target most closely belongs to. Classification is typically used in counterfeit detection applications, such as determining if a collection of microchips contains any non-genuine devices. For classification tasks, SILGEO requires a set of fingerprints from a single, or preferably multiple, known-good examples of the target. To verify authenticity, SILGEO compares the fingerprint from the target to the known-good fingerprints and calculates a similarity score with each candidate fingerprint, indicating the likelihood of belonging to that class. Having fingerprints from multiple known-good examples is preferable, as it provides SILGEO with additional information as to the naturally occurring manufacturing variations, increasing the accuracy and confidence of the classification results. When multiple known-good fingerprints from the same class are available, SILGEO calculates a similarity between the target and each of the known-good fingerprints and then averages the values together, thus allowing the manufacturing variations to be accounted for.

3.3 Data Features

As part of the SILGEO fingerprinting process, various *features* are collected and used to determine the authenticity of the target. In simple terms, a feature is an observation or measurable property that can be recorded and used for analysis. SILGEO has both *measured features*, which are obtained at the time of target characterization, and *derived features*, which are obtained via mathematical transformation.

3.3.1 Measured Features

For SILGEO, all *measured features* are obtained from the scattering parameters captured during the characterization process. Each feature consists of a complex number, typically

expressed as *magnitude* and *angle*, that describes how the propagating wave changed as a result of moving through the medium. The magnitude portion describes the amplitude change of the wave, and the angle, also known as the *phase*, describes the propagation delay experienced by the wave. Each feature is defined for a single *characterization criterion*, consisting of the frequency, power level, port combination, scattering parameter direction, and other characterization settings used to obtain the feature. Parts of the characterization criteria process can be adjusted to increase or decrease the desired number of features obtained from a target.

Characterization Criteria

The characterization criteria helps to contextualize how a measured feature was obtained and is comprised of all the adjustable parameters of the characterization process. By varying the characterization criteria, the thoroughness of a target's characterization can be adjusted. Increasing the feature count for a target can help ensure that the fingerprint is comprehensive and that no features that may aid in determining authenticity were missed. Large feature counts can also help guard against naturally occurring manufacturing and measurement variations from causing errant authenticity results. However, as the feature count increases, so too does the storage space needed to store all the measurement data and the computation power required to analyze it. Additionally, increasing the feature count can substantially increase the time required to acquire all the measurements.

The parameters that comprise the characterization criteria and their roles are detailed below.

- **Frequency Range** describes the frequency bands that are used in the frequency sweep. By characterizing a target across multiple frequencies, additional features can be obtained since a target's response often changes with frequency. This is due to the different frequency bands, i.e., 10 MHz, 100 MHz, 1 GHz, or 3 GHz, having different wave propagation properties. For the same target, certain bands may elicit responses with less noise and variation, leading to a more consistent and accurate characterization.
- **Frequency Point Count** is related to the frequency range and describes the number of points, i.e., distinct frequencies, used in the frequency sweep. By increasing the number of frequencies used in the characterization, greater resolution can be obtained into how the target's response changes with respect to frequency. For example, the range of 1 GHz to 2 GHz, could be swept with 10 points to obtain a coarse overview or with 500 points to obtain a more granular view.

- **Frequency Point Distribution** is related to the frequency points and describes how they are distributed within the frequency range. The steps between frequencies can be linear, logarithmic, or of arbitrary size, depending on the complexity of the target. Changes in wave propagation behavior tend to be logarithmic, meaning that going from 1 MHz to 10 MHz will exhibit a more noticeable change in propagation than going from 100 MHz to 110 MHz. With a logarithmic distribution, more frequency points are sampled in the lower portion of the frequency range. For example, the frequency range 100 MHz to 1 GHz would be sampled as: 100 MHz, 177 MHz, 316 MHz, 562 MHz, and 1 GHz. For comparison, the same range could also be split linearly into five equal steps: 100 MHz, 325 MHz, 550 MHz, 775 MHz, and 1 GHz.
- **Power Level** describes the amplitude of the stimulus wave that is incident upon the target's ports as part of the measurement. Because a target's response is often non-linear with respect to incident power, characterizing at multiple power levels can reveal additional unique features. This is especially true for the transmission portion of the scattering parameters, as the power level can influence the propagation path the wave takes within the target. Thus even without changing the port, measuring at multiple power levels can result in different responses and thus additional features.
- **Port Combination** describes the subset of ports on the target that the frequency sweep is incident upon during the measurement. Characterizing multiple port combinations increases the overall coverage of the target, ensuring that no uniquely identifying features are missed. Additionally, using multiple port combinations helps improve SILGEO's accuracy by reducing the risk of false positives, as each additional port combination provides a different perspective of a target's medium. For each of these perspectives, a unique set of features can be created, further reducing the likelihood of an errant authenticity result, even for high-quality counterfeits.
- **Scattering Parameter Direction** describes whether a response captures the reflection or transmission portion of the wave propagation effects for a network. As the reflected and transmitted responses result from different physical interactions with the target's medium, they yield different sets of unique features, both of which are valuable for determining authenticity. Reflected waves tend to be generated electrically close to where ports are located, resulting in a response that is consistent and noise-free but has not interacted with much of the target's internal medium. Transmitted waves, on the other hand, contain the cumulative effects of traveling port to port through the target and interact with more of the medium within. Transmitted responses thus generally yield more unique features, but at the cost of additional noise and naturally occurring measurement variability.

3.3.2 Derived Features

By post-processing the measured features, additional *derived features* can be obtained: *logarithmic* & *linear* magnitude and *unwrapped* phase. These features provide alternative mathematical representations of the measured features, enabling them to be processed more efficiently and better formatted for additional analysis. For RF measurements, power values are typically expressed in **decibels (dB)** as they better represent the non-linear effects of changes in power. **dB** values are calculated as the logarithmic ratio of two power values and are the ideal representation when directly comparing two features or plotting responses. However, when applying certain statistical analysis methods, a linear representation of the magnitude may lend itself to easier interpretation, particularly for machine-learning algorithms, as the non-linear transformation added by the logarithm could be considered a distortion of the feature. To illustrate the differences between the two representations, Figure 3.5 shows the same set of power features, plotted as linear magnitude and as logarithmic magnitude. Notice that while the overall response is visually similar in both representations, the heights of many features, especially near peaks and valleys, have been alerted.




Figure 3.5: Magnitude Transformation

Unwrapped phase is a transformation that can be applied to the phase portion of each feature in order to make the change in phase between multiple features continuous. Typically, the phase is expressed in radians and is constrained to the range of $\pm\pi$, which causes the phase to wrap around once $-\pi$ or $+\pi$ has been reached, introducing a discontinuity. These discontinuities make a side-by-side comparison of the phase tedious and, worse, can mask differences leading to errors. Phase unwrapping removes these discontinuities by adding back

the multiples of 2π that were removed. Figure 3.6 shows an example of a set of phase features before and after unwrapping. The unwrapping of a set of phase features often starts at the lowest frequency in the sweep, and is used as the reference point for all subsequent phase features to be unwrapped against.




Figure 3.6: Phase Transformation

3.3.3 Analysis Methods

Throughout the research and development stages of SILGEO, various statistical methods have been used to analyze and determine the authenticity of targets. Two main types of analysis were performed with the collected features: *classification* and *data exploration*. For classification tasks, such as determining if a target is genuine or counterfeit or if a target is part of a specific class, [dynamic time warping \(DTW\)](#) was primarily utilized. Though [DTW](#) is typically used with time-series data, it was successfully adapted to work with SILGEO’s frequency-based measurement features [\[58\]](#). For determining the authenticity of a [DUT](#), [DTW](#) was used to compute similarity scores between the features in a [DUT’s](#) fingerprint and those from a known-authentic device. These similarity scores were then subjected to a threshold test, and if a sufficient number were above the threshold, the target was classified as genuine. A similar approach was used for determining the class of a [DUT](#), where [DTW](#) computed similarity scores between the [DUT’s](#) features and features from an example of each class. The set of similarity scores for each candidate class was then condensed down into a single aggregate value, summarizing the overall similarity between the [DUT](#) and that

specific class. The aggregated values were then ranked from highest to lowest, with the class having the highest value, being the most probable class the target belongs to.

For data exploration, the analysis goals were generally more open-ended and focused on understanding the roles of various features. As part of the early research, [principal component analysis \(PCA\)](#) was employed to aid in feature selection and dimensionality reduction. This was especially important for evaluating the applicability of SILGEO to various larger targets, such as vehicles, where millions of features could be collected due to the size and complexity of the targets. [PCA](#) and [t-distributed stochastic neighbor embedding \(t-SNE\)](#) were also used to visualize the relations of certain features and the clustering effects of multiple target classes. Additionally, clustering was used to examine the variability of features within classes and their distance to features of other classes. [Linear discriminant analysis \(LDA\)](#) and [PCA](#) were used to develop *variance masks*, a dimensionality reduction technique that can boost or minimize the effect of certain features. Sets of low and high variance masks were created to test the exclusion of certain features, such as the response from certain bands that exhibited significant variability due to measurement noise.

3.4 Advantages of SILGEO

SILGEO offers several benefits compared to the threat mitigation approaches discussed in [2.2.3](#). Unlike more traditional electrical tests and side-channel-based approaches, SILGEO is unique in that the authenticity assessment is performed with the target *unpowered*, thus avoiding the need for information on a target’s functionality, operating voltages, timing, or electrical connections. This means that SILGEO can be used to authenticate devices even when functional and parametric tests can not be used, such as when the original manufacturer no longer exists or is unwilling to share the required proprietary information. Compared to decapsulation and X-ray imaging, SILGEO does not require any interpretation of images, leaves no temporary or permanent damage, and is not hazardous to the operator performing the evaluation. In contrast to physical inspection, SILGEO requires no subject matter experts and is significantly faster, with authenticity results available in seconds, not hours or days. Furthermore, SILGEO is not reliant on qualitative or subjective interpretation of detected anatomies to determine authenticity; instead, a clear numerical score indicating authenticity is provided. While SILGEO does require a threshold to be defined in order to classify the raw similarity scores, in practice, the numerical difference between the classes is significant enough that selecting a threshold is not a problem. Additionally, unlike physical inspection or electrical tests, SILGEO can evaluate the authenticity of targets in situ, avoiding the need to desolder or remove components from the larger systems they are a part of.

Chapter 4

SILGEO Acquisition System

The following chapter details the current implementation of SILGEO, commonly referred to as the SILGEO *acquisition system*. As a preface, it should be noted that the approach outlined here is just one way to implement the concepts introduced by the SILGEO technology. Other configurations and approaches to implementing SILGEO are possible and are indeed briefly touched upon in Chapter 6. In short, the acquisition system is responsible for generating stimulus signals, delivering them to the [DUT](#), i.e., the target under evaluation, and capturing the corresponding response signals.




Figure 4.1: Acquisition System

Figure 4.1 shows the components that comprise the acquisition system, which primarily consists of a **VNA**, **Multiplexing Board (MB)**, adapter boards, and control software. The dashed lines in the figure indicate **RF** connections, which carry the stimulus and response signals, and the solid lines indicate the data transfer connections between the various components. Frequency sweeps are generated by the **VNA** and sent through the **MB**, which allows the signals to be delivered to a specific subset of the electrical connections on the **DUT**. The stimulus signals then flow through an adapter board, which is mated on top of the **MB** via a standardized interconnect and contains the specific sockets needed to electrically connect to the **DUT**. The response signals from the **DUT** then propagate back to the **VNA** for digitization via the same components but in the reverse direction.

The **VNA** and **MB** are commanded by a custom piece of software named the *acquisition controller*, which runs on a dedicated control computer. The acquisition controller is responsible for the setup and operation of the characterization process, including **VNA** configuration, selection of **MB** output ports, and storage of captured traces. The control computer communicates with the **VNA** via a high-speed **USB 3.1** connection, which is essential for the fast transfer of captured data, providing SILGEO with a fast measurement speed. The **MB** is similarly connected to the control computer via a **USB 2.0** connection. The following sections cover each of these aforementioned components in greater detail.

4.1 Vector Network Analyzer

In the SILGEO acquisition system, the **vector network analyzer (VNA)** is responsible for generating, capturing, and digitizing all stimulus and response signals from a target under evaluation. A **VNA** is an **RF** test instrument used to measure the scattering parameters of electrical networks in the frequency-domain [57]. **VNAs** are complex devices composed of multiple high-precision **RF** devices, including directional couplers, frequency synthesizers, mixers, power amplifiers, and splitters. As for non-**RF** components, **VNAs** also have various onboard high-speed processors, typically **FPGAs**, memory, and analog-to-digital converters. Nearly all **VNAs** have at least two ports, as a single port can only capture reflection, whereas two ports can capture both the reflection and transmission of **DUT**. Multi-port **VNAs** also exist but are less common and cost considerably more due to the additional complexity. The term *vector* in the name **vector network analyzer**, indicates that a **VNA** measures both the amplitude and phase properties of the **DUT**, whereas traditional *network analyzers* only measure magnitude.

For SILGEO, a **VNA** was used instead of a custom measurement setup because the integrated form factor of the **VNA** offered greater accuracy and reliability than a custom

setup could provide. While discreet signal generators and capture equipment could have been used to save on cost, the setup would still have required many specialized RF components, requiring time and effort to evaluate and source them. Several components, such as frequency synthesizers and power amplifiers, require tuning and calibration to produce accurate results, further increasing the complexity of the undertaking. Additionally, the amount of background noise, also known as the *noise floor*, for a custom measurement setup would have likely been far inferior to that of a VNA. For these reasons, it was decided to purchase an off-the-shelf VNA, allowing development efforts to be focused on other areas of the acquisition system.

In order to choose the make and model of the VNA, a survey of the units available on the market was first performed. From this, a set of requirements was developed to guide the selection of the VNA. The VNA needed to support remote control or programmatic control so that it could be integrated into the larger measurement setup. Because the VNA would not be used to measure “traditional” RF devices, it needed to support a wide range of power levels and have a sizable supported frequency range. Based on the noise floor of the other equipment used in the acquisition system, the VNA needed to have a noise floor of at least -75 dBm. The VNA also needed to have at least two ports, with higher port counts being ideal but not a requirement.



Figure 4.2: P9370A - 300 kHz to 4.5 GHz VNA

From this, the Keysight P9370A two-port VNA was selected and is shown in Figure 4.2 [59]. It offers a frequency range of 300 kHz to 4.5 GHz, with an average noise floor of -103 dBm and a sweep time of 14 ms for a 200-point measurement [59]. For interfacing with the DUT, the P9370A has two 3.5 mm connectors, which are mechanically identical and compatible with the more traditional SubMiniature version A (SMA) connectors, except that 3.5 mm connectors use an air dielectric instead of polytetrafluoroethylene. The

P9370A interfaces with a host control computer via a high-speed [USB 3.1](#) connection, and most importantly, the P9370A supports the programmatic control of nearly all operations and settings. The P9370A also comes with a set of software development libraries and examples, significantly reducing the development time and effort required to integrate the [VNA](#) with the rest of the acquisition system.




Figure 4.3: Scattering Parameters for a Two-Port Network

Figure 4.3 provides an example of the scattering parameters that can be obtained with the P9370A [VNA](#). The response of a two-port network was measured over the frequency range of 300 kHz to 4.5 GHz. Within each plot, the magnitude is shown in blue, and the unwrapped phase in orange.

Since the SILGEO characterization process requires measuring the response of a target at many different power levels, frequency ranges, and port combinations, a fast and efficient mechanism was needed to save the scattering parameters captured by the VNA. The P9370A offered two main methods of doing so: direct memory copy or the creation of *SnP* data files. While direct memory access was more straightforward, through testing, it was determined that the creation of *SnP* files was faster, at the cost of having to manage and rejoin the files later when performing the subsequent analysis.

SnP Data File

A *SnP* file, also known as a TouchStone file, is an industry-standard format used by VNAs and other RF test equipment to store the scattering parameters of a network. The *n* in *SnP* indicates the number of network ports contained in the file. For example, a two-port network will have a file extension of **S2P**, a three-port network will have **S3P**, and so forth. The *SnP* format also supports alternate network representations, such as admittance parameters and impedance parameters. In the case of the P9370A, which has two ports, **S2P** files are utilized.

Figure 4.4 shows a portion of a **S2P** file. The header section, denoted by the exclamation marks, contains meta-data about the measurement, including calibration status, data and time the measurement was taken, and information about the instrument used. The remainder of the file contains the scattering parameters for the network, with one line for each frequency the network was characterized at. The first column indicates the frequency at which each measurement was taken, with the remaining eight columns alternating between the magnitude and phase values for each of the four scattering parameters.

```

1 !Keysight Technologies,P9370A,MY58100360,A.15.20.07
2 !Date: Tuesday, September 14, 2021 18:46:49
3 !Correction: S11(C 2-Port)
4 !S21(C 2-Port)
5 !S12(C 2-Port)
6 !S22(C 2-Port)
7 !S2P File: Measurements: S11, S21, S12, S22:
8 # Hz S dB R 50
9 9000 1.79407 0.04668 -89.3978 152.606 -87.1135 109.282 1.95783 1.37171
10 2610036.8 -0.25737 -112.994 -47.1078 32.4466 -47.0741 32.5319 0.04824 -1.34687
11 5211073.6 -0.24864 -145.001 -45.9501 15.8098 -45.9223 15.8884 0.02737 -2.52014
12 7812110.4 -0.23984 -158.506 -45.6807 8.52011 -45.6426 8.54942 0.02967 -3.69643
13 10413147.2 -0.23902 -166.198 -45.5789 4.03124 -45.5404 4.06123 0.02593 -4.86037
14 13014184.0 -0.23922 -171.504 -45.5638 0.64566 -45.5198 0.60218 0.02279 -5.99574
15 15615220.8 -0.24171 -175.554 -45.5284 -1.66047 -45.4837 -1.76406 0.01230 -7.13825
16 18216257.7 -0.24209 -178.839 -45.5035 -4.02739 -45.4723 -4.02990 0.00080 -8.28181

```

Figure 4.4: S2P Data File

4.2 Multiplexing Board

Within the SILGEO acquisition system, the **Multiplexing Board (MB)** plays a critical role in terms of accuracy, measurement throughput, and repeatability. Sitting in between the **VNA** and an adapter board, the **MB** enables a **DUT** which may have multiple electrical interfaces, to be sampled in an efficient manner by the **VNA**. The **MB** functionally consists of two isolated bidirectional 8-to-1 **RF** multiplexing pathways with an operating bandwidth of 300 kHz to 4 GHz, power supplies, and an embedded microcontroller for control and communication with the host computer. The **MB** features two sets of **RF** grade connectors: 2 **SubMiniature version A (SMA)** ports for interfacing with the **VNA**, and 16 **SubMiniature version P (SMP)** ports for interfacing with the **DUT** via an adapter board.

The multiplexing capability of the **MB** works as follows: each **SMA** input can be electrically connected to one of 8 **SMP** output connectors. In doing so, the stimulus and response signals of the **VNA** are transparently sent to a subset of the **DUT's** electrical ports. The **MB's** two multiplexing paths are implemented using solid-state **RF** switches, which provide fast switching speeds and high levels of isolation between adjacent channels. The selection of the multiplexing path is sent to the **MB** from the host computer, and the onboard microcontroller sequences the switches in the appropriate manner to set up the request path.




Figure 4.5: Multiplexing Board

Figure 4.5 shows an assembled MB. The blue cables on the right of the figure are mated with the two SMA connectors. The square array of gold devices in the center of the board are the 16 SMP connectors, with the top half forming the first 8-to-1 multiplexer and the bottom half the second 8-to-1 multiplexer. The silver-colored shields on the upper and lower left side of the board contain the microcontroller and power regulation components, respectively.

Importance of the Multiplexing Board

The design and incorporation of the MB in the acquisition system was done primarily for reasons of measurement throughput, eliminating the time-consuming process of repeated SMA connection and disconnection required when measuring a DUT directly with the VNA. The problem was most acute when measuring DUTs that had more than two electrical ports since the VNA has only two measurement ports. The complete characterization of a DUT for assessment with SILGEO requires $\binom{n}{2}$ connections, where n is the number of DUT ports. As n increases, the time required for complete characterization quickly becomes prohibitive. For example, the complete characterization of a DUT with 8 ports would require 28 separate SMA connections and disconnections to be performed, each taking minutes to complete. By incorporating the MB, the measurement time for the 8-port DUT can be reduced by an order of magnitude, going from tens of minutes to tens of seconds. For DUTs with higher port counts, the reduction in measurement time can be even more significant.

Additionally, while SMA connectors are designed to provide repeatable measurements, that repeatability is only achieved when the connectors are mated using the appropriate torque wrench [57]. As detailed above, measuring all combinations of a 8 port DUT without the MB would require 28 connections and disconnections, each having the possibility of a slight misconnection leading to an errant measurement. Whereas with the MB, the DUT is connected only once at the start of a measurement cycle, any number and combination of measurements can then be taken or even repeatedly taken, and only after the DUT is fully characterized is the DUT then disconnected once.

Finally, the MB was developed as part of a solution to the challenge of characterizing DUTs that are not fitted with SMA connectors and, therefore, can not be connected to a VNA. Rather than designing a custom multiplexing setup for each DUT, a cost-prohibitive and time-consuming endeavor, the solution was to compartmentalize the expensive portions of the measurement setup onto a single board, i.e., the MB, which then exposes a standardized interconnect. An additional adapter board then mates to this interconnect and contains the traces and specific sockets or ports needed to finalize the electrical connection with the DUT. As the adapter board has a simple layout and no active electronics, it is significantly less expensive to manufacture than the MB. This approach allows SILGEO to characterize

a wide variety of DUTs at a low cost since the incremental effort required to support a new type of DUT is limited to the creation of a new adapter board. Section 4.3 covers the adapter boards in greater detail.

4.2.1 Multiplexing Board Design

As the MB is a *mixed-signal* circuit board, meaning it contains both digital and analog circuits, careful consideration was given to its design to ensure that noise and signals from the digital components did not cause disturbances to the noise-sensitive analog components. The MB consists of three main areas: control, power, and RF multiplexing.

The control section is composed primarily of an Atmel ATmega32U4 microcontroller [60] and its supporting components. To reduce the possibility of spurious emissions from the microcontroller affecting the analog signals on the multiplexing channels, the control section is galvanically isolated from the rest of the board. The galvanic isolation is achieved by having separate power and ground references for the control section, an isolation gap between the copper planes, and the use of *optoisolators*, an electronic component that allows signals to pass between two electrically isolated circuits. This results in the MB having separate digital and analog sections, with only the optoisolators bridging the two sections, in order to move the low-frequency control signals from the microcontroller to the RF switches. This isolation becomes even more important when considering the microcontroller is electrically connected to a host computer via USB during its operation, which could allow coupled emissions from the computer to be introduced to the board [61]. As an extra isolation measure, the majority of components in the control section are encased in an RF shield to minimize the possibility of any radiated emissions coupling back to the board and causing distortion [62].

The power section of the MB is responsible for supplying the RF switches with clean and low noise power at the required voltages, and is comprised of three voltage regulators, various smoothing capacitors, and support components. The main input voltage regulator accepts a range of 7 V to 35 V from an external *direct current* (DC) power supply and converts it to a base 5 V supply. The 5 V is then further converted by two additional regulators to generate the final ± 3 V and -2.5 V power rails used by the RF switches. Like the control section, nearly all of the power section is encased in an RF shield to minimize any radiated emissions.

The remainder of the space on the MB is taken up by the multiplexing section, which principally consists of microstrip RF traces, RF switches, and the RF connectors for the input and output signals. Two models of RF switches are used to construct the multiplexers: the ADRF5021 2-to-1 switch [63] and the ADRF5047 4-to-1 switch [64]. Both ADRF switches are surface-mount devices manufactured by Analog Devices. Each multiplexing pathway

has one “input” SMA connector and eight “output” SMP connectors. The input and output here are put in quotes since the multiplexing pathways are bi-directional; the SMA connector is simultaneously the input for the stimulus signal from the VNA and the output for the response signal returning from the target. To interconnect the RF switches, SMA, and SMP connectors, microstrip traces are used on the top layer of the board, with its corresponding ground plane reference placed on the layer immediately below. A variety of measures to increase isolation in the multiplexing section are also used, most notably *via stitching* and *via fencing*.




Figure 4.6: Multiplexing Board Layout

Figure 4.6 shows the copper layers of the MB, with the RF multiplexers visible on the right half of the board. The red in the figure represents areas of copper on the top layer, with the yellow representing the ground plane immediately below on the first inner layer. The microcontroller and supporting components are in the top left, and the power regulation is located below in the lower left corner. The reverse “L” shaped band of white space is the aforementioned galvanic isolation gap, with the optoisolators on the vertical portion labeled

as devices U11 through U20.

As part of designing the MB, various topologies, components, and trade-offs were considered. Front of mind was ensuring the design had high isolation between channels and low attenuation for the carried stimulus and response signals. The following subsections detail the engineering choices and provide greater specificity into the design.

Solid-State Multiplexing

Among the first of design choices considered was whether to use electromechanical or solid-state, i.e., semiconductor-based, RF switches to build the multiplexers. When compared to solid-state switches, electromechanical switches have higher isolation and lower insertion loss, both desirable qualities for the MB. However, this comes at a cost, as electromechanical switches are physically bulky, often many times larger than solid-state ones, and suffer from repeatability issues over time because they have moving parts [65]. Ultimately for reasons of space, cost, and repeatability, it was decided to use solid-state switches.

Multiplexing Topology


Before deciding on a specific model of solid-state switch, it was first necessary to determine the multiplexer topology and, specifically, how many inputs the multiplexers would have. Selecting the number of inputs is a trade-off as it drives both the price of the MB, the more inputs, the higher the cost, and sets a maximum to the number of pins on a DUT that can be simultaneously characterized. Designs ranging from 10 inputs all the way up to 20 inputs were considered, but it was decided to go with 16 inputs, split into two equal 8 input multiplexers. A large driver of not increasing the input count past 16 was the switch pricing per input being non-linear, with part costs from manufacturers nearly doubling when going from 16 to 18 inputs. Additionally, there were concerns that high input count switches tended to have worse isolation and insertion loss characteristics when compared to 8 or 4 input switches.

With the number of inputs decided, a parts search was conducted to determine what switches were available on the market and the topology to best configured them in. As part of this search, it was decided that the 8-to-1 multiplexers should be constructed from multiple switches, rather than just one. While there were single device 8-to-1 multiplexers available from some vendors, they often had poor isolation between channels, a small operating frequency range, and added unacceptable amounts of attenuation to the carried signals. Based on this finding, calculations were performed to estimate the performance of connecting

multiple switches together, namely connecting two 4-to-1 switches to a single 2-to-1 switch. This alternative topology to achieve an 8-to-1 multiplexer was vastly superior in all aspects, including price, compared to a single-device solution.

Of the available switches on the market, the ADRF5021 and ADRF5047 were selected primarily due to their superior isolation and low insertion loss. The ADRF5021 2-to-1 switch has an operating range of 9 kHz to 30 GHz [63], and the ADRF5047 4-to-1 switch has an operating range of 9 kHz to 44 GHz [64]. One may wonder why these two selected switches have maximum frequencies many times higher than the board’s maximum operating frequency of 4 GHz. This choice was intentional; as for the operating band of the board, 300 kHz to 4 GHz, these devices had the best performance on the market. At 4 GHz, the ADRF5021 and ADRF5047 have a mean insertion loss across all paths of 0.81 dB and 0.96 dB, respectively, with the ideal insertion loss being 0 dB. In terms of isolation, the ADRF5021 and ADRF5047 have a mean input to input isolation of -70 dB and -65 dB, respectively, at 4 GHz, with lower values indicating better isolation.

Also of note is that the ADRF5021 is a non-reflective switch, meaning that when a path is not selected, the RF input is terminated with a $50\ \Omega$ matched load. This helps to increase the isolation performance of the switch and the overall measurement accuracy of SILGEO, as it ensures signals from non-selected paths do not leak into the currently selected path. Unfortunately, the ADRF5047 is a reflective switch, meaning that non-selected paths are not terminated by a $50\ \Omega$ matched load, and instead, the inputs are left free-floating. This lack of the built-in termination is reflected in the isolation numbers above, with the ADRF5047 having an isolation figure 5 dB higher than the ADRF5021.



multiplexers. Switches MUX_1 , MUX_2 , MUX_3 , and MUX_4 use the ADRF5047. Switches MUX_5 and MUX_6 use the ADRF5021. For reasons related to trace layout, RF ports 1, 2, 3, 4, 5, 6, 15, and 16 map to the first multiplexer, and RF ports 7, 8, 9, 10, 11, 12, 13, and 14 map to the second multiplexer, rather than just increasing linearly. Not shown in the figure are the control lines used to set the state of the switches, which connect back to the microcontroller.

Microstrip Traces

To electrically connect the RF switches, SMA, and SMP connectors, coplanar waveguide with ground (CPWG) microwave transmission lines were utilized. CPWGs are planar microstrip structures, constructed from a center conductor surrounded on both sides with ground planes on top of a dielectric medium, the bottom of which is on top of an additional ground plane [56]. For greater clarity, a cross-section of a CPWG trace is shown in Figure 4.8, along with the various design parameters. In the figure, W indicates the width of the center conductor; G is the gap between the conductor and the coplanar ground planes, and t is the thickness of the center conductor. The light gray slab represents the dielectric medium, with its relative permittivity denoted as ϵ_r and height as h . The dark gray plane on the bottom is the additional ground plane. Of the various CPWG parameters, W and G are the most flexible for the designer, and are generally selected in response to the more rigid parameters of h and ϵ_r , which are a function of PCB material. Amongst the family of planar transmission lines, CPWG traces were chosen because of their ease of construction and the high levels of isolation they offer, which is largely due to the ground planes surrounding the center conductor on three out of four sides [56].




Figure 4.8: Cross Section of a Coplanar Waveguide With Ground Plane

To ensure the functional goals of the MB were met, the microstrip traces were designed to have a target impedance of $Z_0 = 50\ \Omega$ at 4 GHz. To achieve this, rough calculations were first performed for values of W and G . These estimates were then combined with the h and ϵ_r of the selected PCB material (discussed further in the next section), and iteratively

refined with the aid of simulation software. The complete set of simulation steps performed is beyond the scope of this document, but after several rounds of simulation, the final values of $W = 0.35$ mm and $G = 0.2$ mm were obtained.

PCB Layers and Stackup

The MB is constructed as a four-layer PCB, meaning it uses four copper layers, separated by insulating layers, to route the required electrical connections. Each layer is generally reserved for a specific function, with the top layer used for RF traces, the first inner layer for an RF ground plane, the second inner layer for signals, and the bottom layer for power.

The RF layer contains all of the microstrip traces and SMT footprints for the RF switches, power supplies, and microcontroller. The ground plane layer immediately below is a contiguous area of copper and is used as an analog ground reference by the microstrip traces and RF switches. The contiguous nature of the ground plane is important as it provides a low impedance return path for all of the high-frequency signals in the layer above [56]. The ground plane also helps to reduce the possibility of any noise or spurious emissions from the signaling and power layers from coupling to the microstrip traces and causing distortion [62]. The signal layer carries the control lines used by the microcontroller, connected via the optoisolators, to set the state of the ADRF switches. And finally, the power layer is reserved for traces moving power from the voltage regulators to the microcontroller and RF switches. It is important to note that the galvanic isolation discussed previously is on *all* layers; for example, the bottom power layer is two electrically isolated areas, both with their own set of power traces.

As part of designing the MB, careful attention was paid to the materials and layer thicknesses that were used to construct the board, commonly referred to as the PCB *stackup*. To ensure minimal attenuation of the carried high-frequency stimulus and response signals, the MB was designed against and constructed from the Isola 370HR material [66]. The 370HR material was selected as it has superior permittivity and loss tangent characteristic stability for signals over 2 GHz when compared to traditional FR4. For example, at 4 GHz, the 370HR material has a known ϵ_r value of 3.98, whereas ϵ_r for FR4 can range anywhere from 4.0 up to 5.0 depending on the manufacturer [67]. This unpredictability in FR4 characteristics creates a challenge for the selection of microstrip parameters, as higher than expected ϵ_r values can cause impedance mismatches and unintended signal attenuation [18]. The control of layer thickness, particularly when using microstrip traces, is also important, as changes in dielectric thickness, h , directly affect the trace impedance [56].

With the expected ϵ_r now defined, the remaining microstrip parameter, h the dielectric

thickness, could now be calculated. Holding W , G , and ε_r as constant, the thickness h was varied to obtain the exact trace impedance of $50\ \Omega$. Working with the PCB manufacturer, this information was then codified into a custom stackup, which detailed the ordering of layers, materials to be used, and expected layer thicknesses after manufacturing. Table 4.1 shows the final PCB stackup used. For clean and finished look, the MB has a white-colored silkscreen on a matte black-colored solder mask. Exposed copper areas on the top and bottom layers are finished with an [electroless nickel immersion gold \(ENIG\)](#) plating for superior solderability and smoothness [68].

Name	Material	Thickness (Mils)
Solder Mask	PSR-4000BN	0.80
Layer 1	1 oz copper	1.40
Pre-preg	370HR	9.12
Layer 2	1 oz copper	1.40
Core	370HR	39.00
Layer 3	1 oz copper	1.40
Pre-preg	370HR	9.12
Layer 4	1 oz copper	1.40
Solder Mask	PSR-4000BN	0.80

Table 4.1: Multiplexing Board PCB Stackup

Phase Matching

Another design objective for the MB was to have relatively similar trace lengths for each of the 16 multiplexing pathways, thus ensuring that the phase delay, also called the propagation delay, imparted onto the carried signals was uniform across all channels. Having a uniform phase delay offers various benefits, but principally it simplifies the task of *de-embedding*, i.e., nulling out, the effects of the MB on the saved measurements [57]. With phase machining, the number of de-embedding calibrations required is simplified from one for every port combination to one for each channel, which on the MB means only 16 are needed instead of 64. Uniform phase is also important for any experiments using [time-domain reflectometry \(TDR\)](#) or the phase component in general, as a uniform phase delay enables cross-channel comparisons to be made [69]. The lengths for each multiplexing pathway are listed in Table 4.2, with the trace lengths measured from the SMA connector board edge to the center pad of the SMP connector.

Path	Length (mm)	Path	Length (mm)
SMA ₁ to RF ₁	134.054	SMA ₂ to RF ₇	134.053
SMA ₁ to RF ₂	134.054	SMA ₂ to RF ₈	134.053
SMA ₁ to RF ₃	134.054	SMA ₂ to RF ₉	134.053
SMA ₁ to RF ₄	134.054	SMA ₂ to RF ₁₀	134.053
SMA ₁ to RF ₅	134.045	SMA ₂ to RF ₁₁	134.053
SMA ₁ to RF ₆	134.059	SMA ₂ to RF ₁₂	134.053
SMA ₁ to RF ₁₅	134.059	SMA ₂ to RF ₁₃	134.053
SMA ₁ to RF ₁₆	134.045	SMA ₂ to RF ₁₄	134.053

Table 4.2: Signal Path Lengths

The difference between the shortest trace and longest trace is 0.014 mm or 14 μm , which at 4 GHz results in a phase difference of only 0.1065°. The calculations to obtain this value are shown in Equation 4.1.

$$\begin{aligned}
\phi &= \beta \cdot l \\
&= \sqrt{\varepsilon_e} \cdot k_0 \cdot l \\
&= \sqrt{\varepsilon_e} \cdot \frac{2 \cdot \pi \cdot f}{c} \cdot l \\
&= \sqrt{2.513} \cdot \frac{2 \cdot \pi \cdot 4 \text{ GHz}}{2.99 \cdot 10^8 \text{ m s}^{-1}} \cdot 14 \mu\text{m} \\
&= 0.00186 \text{ rad} \cdot \frac{180}{\pi} \\
&= 0.1065^\circ
\end{aligned} \tag{4.1}$$

Where l is the largest difference in microstrip trace length, β is the phase constant which is equal to $\sqrt{\varepsilon_e} \cdot k_0$ for CPWG microstrip lines, k_0 is the *free-space* wave cut-off number for the maximum operating frequency f of 4 GHz, c is the speed of light, and ε_e is the *effective* dielectric constant of the MB, which was obtained via computer simulation.

Via Fence

To increase the isolation between channels and reduce the susceptibility of coupling environmental noise, each microstrip trace on the MB was surrounded by a *via fence*. Constructed by placing rows of vias on both sides of a microstrip, a via fence electrically connects the

coplanar ground planes that surround the microstrip with the ground plane below it. When placed tightly enough, the rows of vias form a barrier to any stray electromagnetic wave propagation in the [PCB](#) substrate itself [70]. Figure 4.9 shows a via fence applied to a trace on the top layer of the [MB](#). The microstrip trace is colored in pink, with the top layer coplanar ground planes in red and the additional ground plane below the microstrip colored in yellow. The vias on both sides of the trace, shown as gray circles, are connected to analog ground (GNDA), the reference used by the [CPWG](#). The smaller concentric white circle inside the vias represents the via hole that will be drilled. Note that the aliasing seen on the edges and lines in the figure is a visualization artifact from the [EDA](#) tool and is not present on the actual board.




Figure 4.9: Microstrip Trace Surrounded by a Via Fence

A via fence has two main design parameters: the *offset*, which is the distance from the [RF](#) trace to via, and the *pitch*, which is the distance between subsequent vias in each row. The minimum offset, i.e., the closest the vias can be to the trace, is calculated as the sum of the width and the gap parameters of the microstrip [71]. For the microstrips on the [MB](#), this value was then rounded up to the nearest integer. The calculations for the offset can be seen in Equations 4.2 and 4.3.

$$\begin{aligned}
\text{offset}_{min} &= W + G \\
&= 0.35 \text{ mm} + 0.2 \text{ mm} \\
&= 0.55 \text{ mm}
\end{aligned} \tag{4.2}$$

$$\begin{aligned}
\text{offset} &= \lceil \text{offset}_{min} \rceil \\
&= \lceil 0.55 \text{ mm} \rceil \\
&= 1 \text{ mm}
\end{aligned} \tag{4.3}$$

Where W is the width of the microstrip and G is the gap between the microstrip and the coplanar ground planes.

The pitch parameter is selected by first calculating the maximum allowed distance between subsequent vias. For the MB, a rule of thumb was used in which the maximum pitch is defined as the wavelength at the maximum operating frequency divided by 20 [62]. This value was then rounded down, thus increasing the via density, to have an easier number to work with. The calculations for the pitch can be seen in Equation 4.4.

$$\begin{aligned}
\text{pitch}_{max} &= \frac{\lambda}{20} \\
&= \frac{37.468 \text{ mm}}{20} \\
&= 1.8734 \text{ mm} \\
&\approx 1.75 \text{ mm}
\end{aligned} \tag{4.4}$$

Where λ is the wavelength of the maximum operating frequency of 4 GHz. For completeness, the calculations for obtaining λ are shown in Equation 4.5.

$$\begin{aligned}
\lambda &= \frac{c}{f \cdot \sqrt{\varepsilon_r}} \\
&= \frac{2.99 \cdot 10^8 \text{ m s}^{-1}}{4 \text{ GHz} \cdot \sqrt{3.98}} \\
&= 37.468 \text{ mm}
\end{aligned} \tag{4.5}$$

Where c is the speed of light, f is the maximum operating frequency of 4 GHz, and ε_r is the relative permittivity of the circuit board at 4 GHz. Like all other vias on the MB, the via

fences are comprised of vias with a copper diameter size of 0.6 mm and a drill diameter size of 0.3 mm.

Via Stitching

As a final design step, *via stitching* was applied to the ground planes and board edges of the MB to ensure that a uniform and low impedance return path was present across the board. Via stitching is a technique in which large quantities of vias are added to a layout in order to electrically bond together areas of overlapping ground planes located on different layers of a circuit board [72]. The added vias provide vertical connections throughout the board, causing the ground planes to be tightly coupled at the same potential, thus creating a low impedance path for the return current of the high-frequency signals to flow through [62]. Furthermore, the tight coupling ensures that none of the ground planes on the board can act as potential stubs or transmission line segments, leading to trace impedance discontinuities [72].

When applied to the perimeter of the board, sometimes also referred to as a *picket fence*, the via stitching has a secondary benefit of reducing the possibility that any internally propagating signals escape via the edge of the board [73]. The occurrence of these edge emissions is of particular concern for board designs that have multiple parallel ground planes, such as the MB, as the planes can act as pseudo-waveguides, funneling spurious energy from the center of the board to the edges, where it can then radiate.

Via stitching has one main design parameter: the maximum *spacing* in-between vias. For the via stitching on the ground planes and on the board edge, the same rule of thumb was used, in which the maximum distance between vias is defined as the wavelength at the maximum operating frequency divided by 8 [72]. The calculations for the maximum distance can be seen in Equation 4.6.

$$\begin{aligned} \text{spacing}_{max} &= \frac{\lambda}{8} \\ &= \frac{37.468 \text{ mm}}{8} \\ &= 4.6835 \text{ mm} \end{aligned} \tag{4.6}$$

Where λ is the same wavelength as calculated in Equation 4.5. To have simpler numbers to work with during board layout, the value from Equation 4.6 was rounded down to 3 mm spacing for the ground planes and 2 mm for the board edges. Rounding down leads to greater

via density per unit area, which has a positive benefit as it further increases the isolation effects [72].

Figure 4.10 shows two examples of via stitching from the MB, with the red representing the top ground layer, the vias shown as gray circles, and the drill holes as white circles. Note that in Figure 4.10a, the edge stitching was also applied to the mounting holes of the MB, shown as a larger white circle, as this technically is another board edge that could be radiating.




Figure 4.10: Via Stitching

4.2.2 Multiplexing Board Performance Analysis


Since the MB sits in the signal path between the VNA and DUT, via an adapter board, it has a direct impact on the acquired measurements and the overall detection capability of SILGEO. Therefore, it was important to first verify the performance of the MB before it could be used to measure and assess targets. The performance, or quality, of the MB is best characterized by two main figures of merit: insertion loss and isolation, both of which are measured in decibels (dB). Insertion loss quantifies the amount of signal power attenuated, i.e., lost, as a consequence of adding the MB to the signal path between the VNA and the DUT. Isolation quantifies the amount of signal power that unintentionally leaks into the selected signal path from the other non-selected multiplexing paths and from the background

environment. It should be noted that since the MB contains 16 distinct signal paths, the figures of merit are not expressed as a single value but instead as 16 distinct insertion loss values and 16 distinct isolation values.


Characterization of the MB had additional benefits as well, namely that the saved measurements could be post-processed to create a de-embedding profile, which allows the effects of the MB to be removed from the collected measurements. And finally, a key trade-off faced during the design of the MB was how to minimize insertion loss while simultaneously maximizing isolation. Since these two goals can sometimes be at odds, measuring the MB helps to validate if the assumptions made during design were justified and if changes should be made for subsequent board design revisions.

Measurement Setup

In order to characterize the insertion loss and isolation of each multiplexing path on the MB, a custom measurement setup was created. Since the MB consists of two identical 8-to-1 multiplexers, they were measured separately using the same setup. Figure 4.11a shows an overview of the setup consisting of the VNA on the top right, the multiplexing board, and a pair of 24-inch experiment-grade SMA-to-SMA RF cables. Port 1 of the VNA is connected to the input of the 8-to-1 multiplexer, in this case, SMA₁, via a SMA cable. Port 2 of the VNA is then connected to one of the SMP outputs on the same 8-to-1 multiplexer, such as RF₉ as shown in Figure 4.11b, via a SMA cable and a special adapter.



(a) Measurement Setup



(b) SMA-to-SMP Adapter

Figure 4.11: Multiplexing Board Characterization

Once the connections were made, the control computer cycled through the 16 multiplexing pathways and saved a **VNA** measurement sweep for each one, for a total of 16 sweeps. Port 2 of the **VNA** was then moved to the next output on the 8-to-1 multiplexer, in this case, RF_2 following the example from above, and another 16 sweeps were saved. This was repeated until all 8 **SMP** outputs of the same 8-to-1 multiplexer had been measured. The above procedure was then repeated for the second 8-to-1 multiplexer, with Port 1 of the **VNA** now connected to the second **SMA** input on the **MB**.

To guard against any measurement variability, connector misalignment, or environmental noise, the outlined experiment procedure was repeated five times. To ensure accurate and true results, prior to making any measurements with **VNA**, a full two-port calibration was completed utilizing a Keysight N7551A 2-Port ECal module, which has a calibration range of DC to 6.5 GHz [74]. The calibration was conducted at the end of the 24-inch **SMA** cables and stored as a profile in the **VNA**, thus nulling out any effects from the cables and calibrating the power level and dynamic range of the **VNA** [57].

As noted above, in order to connect the **SMA** cables to the **SMP** ports on the **MB**, an adapter was used. This adapter, shown in Figure 4.11b, consists of a **SMA** female to **SMP** male adapter and a short length of semi-rigid low loss coax cable. Since this adapter lies beyond the calibration plane, that is, after the end of the **SMA** cables, the adapter was specially ordered from a vendor who advertised they could provide both the adapter itself and its de-embedding profile. This profile would have enabled the removal of the effects imparted by the **SMA-to-SMP** adapter during the post-processing of the captured measurements. However, due to unknown reasons, the vendor was unable to provide a characterization of the adapter. As a result, the collected measurements have the effects of the 24-inch **SMA-to-SMA** cables removed but include the effects of the **SMA-to-SMP** adapter.

Figures of Merit

To calculate the figures of merit, a series of custom MATLAB scripts were created to process and plot the collected measurement data. Figure 4.12 shows the insertion loss of each signal path on the multiplexing board. To provide additional context, the insertion loss for the ADRF5021 and ADRF5047 multiplexing chips are also plotted, as well as the cumulative insertion loss when ADRF5021 and ADRF5047 are placed back to back in the signal path, as done in the 8-to-1 multiplexer design. This provides a sort of ideal baseline, indicating the lowest insertion loss that could be achieved with the chosen ADRF devices if the **PCB** and **RF** connectors were ideal. Looking at the signal paths, there is an overall trend of the insertion loss increasing as the frequency increases. This is to be expected as the ϵ_r of the

PCB material typically decreases as the frequency increases, thus leading to higher losses caused by absorption [56].




Figure 4.12: Multiplexing Board Insertion Loss

Figure 4.13 shows the mean isolation of each signal path on the multiplexing board, calculated as the average power received at Port 2 of the VNA when each of the other outputs on the same 8-to-1 multiplexer carries stimulus and response signals. In essence, this represents the power unintentionally leaking into the selected SMP output from the other outputs on the same multiplexer. For example, the mean value plotted for SMA₁ to RF₁ is the average power measured at RF₁ when SMA₁ is used to send power to outputs 2, 3, 4, 5, 6, 15, and 16.

The mean isolation for most of the signal paths was at or below -60 dB, with the exception of around 2.5 GHz for SMA₁ to RF₅ and at 4.5 GHz for SMA₁ to RF₄ and SMA₂ to RF₉. Overall there is a trend that as the frequency increases, the isolation between channels decreases. This is largely to be expected, as higher frequencies tend to radiate more and




Figure 4.13: Multiplexing Board Mean Isolation

require additional design efforts to contain [56]. The complete set of performance figures and characteristics for the MB can be found in Appendix A.

Analysis & Discussion

Looking at Figures 4.12 and 4.13 together, it should be noted that while the MB does have an insertion loss that is slightly higher than desired for frequencies greater than 4 GHz, the excellent overall isolation of -60 dB indicates the MB is suitable for use in the SILGEO acquisition system and that the MB does not significantly distort the measurements of a connected DUT. It is also important to point out that the MB was designed with a maximum operating frequency of 4 GHz in mind, so the performance concerns beyond that frequency are somewhat to be expected. If the maximum of 4 GHz is respected during operation, these aforementioned issues should not cause any measurement degradation.

The two figures of merit discussed are also both contributors to the overall measurement **signal-to-noise ratio (SNR)**, which is defined as the ratio of the signal power, i.e., the signal of interest, to background noise power [56]. Measurements with a high **SNR** allow for increased accuracy, sensitivity, and better detection of minute changes in the **DUT**, critical to SILGEO’s operation. Low insertion loss helps to maximize the amount of **VNA** stimulus power that is delivered to the **DUT**, increasing the *signal* portion of the **SNR** ratio. High levels of isolation ensures that a **DUT** measurement on a particular electrical pin only contains responses from that pin and no others, in effect reducing the *noise* component of the **SNR** ratio. If lower insertion loss were required, the design of the **MB** could be altered by considering alternative **PCB** materials such as ceramics or alumina [62]. Alternatively, if increased isolation were required, electromechanical **RF** switches could be utilized instead of the solid-state silicon-based multiplexing chips. As previously mentioned, electromechanical switches can provide increased isolation and lower insertion loss when compared to their silicon-based counterparts, at the expense of being physically much larger and costing more.

And finally, the accuracy of the calculated figures of merit could be improved by removing the effects of the **SMA-to-SMP** adapter. This could be achieved by obtaining a characterization for the adapter from the original vendor or by switching to an adapter from another vendor who can provide the characterization. Alternatively, a custom test jig could be created as a means to measure the adapter provided by the first vendor, although this method may not yield measurements that have the same accuracy as vendor-provided characterizations.

4.3 Adapter Boards

Adapter boards are interchangeable circuit boards that mount on the **MB** and enable a wide variety of targets to be electrically connected to the acquisition system. Adapter boards are responsible for transferring the high-frequency stimulus and response signals between the **MB** and the target with minimal distortion and attenuation. On the top of the adapter board is the specific connector needed to electrically interface with the **DUT**. On the bottom of the adapter board is the standardized interconnect designed to mate with the companion interconnect on the **MB**, consisting of 16 **SMP** connectors, one for each **RF** channel. The **SMP** connectors not only provide a low-impedance pathway for the **RF** signals to flow through but also supply mechanical rigidity to ensure the adapter board does not move with respect to the **MB**. Figure 4.14 shows the underside of an adapter board with the 16 **SMP** connectors in gold and the pins for connecting to the **DUT**, in this case, the D38999, in the center. Figure 4.15 shows the same D38999 adapter board mounted onto the **MB**.




Figure 4.14: D38999 Adapter Board Bottom View





Figure 4.15: Multiplexing Board With Attached D38999 Adapter Board


Along with the MB, adapters boards are part of the solution used by the acquisition system to connect targets that are not fitted with SMA connectors to the VNA. By limiting the active electronics of the acquisition system to the MB, adapter boards require minimal components, just the SMP and DUT connectors, and thus are easy to design and cheap to manufacture. The adapter boards benefit from the design efforts that went to the MB and are constructed similarly, with PCB material, stackup, and CPWG trace parameters all identical. Additionally, the PCB layout of the adapter board is designed to be easily adapted to new targets, requiring only the DUT connector itself and some connecting traces to be updated for each new target. This approach allows the acquisition system to interface with a wide variety of DUTs, as the incremental effort and cost to evaluate a new target with SILGEO, is limited to the creation of a new adapter board.

The modularity and flexibility of the split MB and adapter board approach is best seen in Figure 4.16, which displays a collection of adapter boards that were created as part of evaluating various targets with SILGEO. Each adapter board shown in the collection has a different electrical connector fitted in the center of the board and interfaces with DUTs that have the matching electrical connector. Figure 4.16a has a D38999 circular connector, a standardized interface used on various pieces of military equipment. Figure 4.16b has a 28 lead shrink small outline package (SSOP) socket, allowing a non-permanent connection to be made to components that otherwise would require a soldered connection. Figures 4.16c and 4.16d are fitted with Ethernet and USB connectors, respectively, and are slightly different in how they connect to the DUT. Rather than connect directly like the previous two examples, these adapter boards are instead designed to connect to the DUT through an Ethernet or USB cable. As these are standardized connectors, they are often present on many targets, making them an ideal starting point for a target's characterization.


Of note, Figure 4.16d shows a half-sized adapter board, using 8 channels instead of the full 16, for interfacing with USB devices. Half-sized adapter boards were created after encountering multiple targets that did not require the full 16 channels. By using 8 channels, not only is the board more efficient space and cost wise, it additionally allows two half-sized adapter boards to be simultaneously mounted to the MB. This enables DUTs with multiple types of connectors to be characterized in a more thorough manner. For example, by measuring the USB and Ethernet interfaces of a target simultaneously, rather than inter-dependently, the interaction effects of the two ports can be captured, yielding additional features.



(a) D38999



(b) SSOP-28



(c) Ethernet




(d) USB

Figure 4.16: Adapter Boards for Various DUT's

4.4 Acquisition Controller

The *acquisition controller* is a software program running on the control computer and is responsible for controlling and sequencing measurement acquisition as part of the SILGEO characterization process. The acquisition controller's responsibilities range from accepting user input, configuring the [VNA](#) with requested characterization settings, changing the selected ports on the [MB](#), storing the collected measurements, and more. The acquisition controller is limited to only running on the Windows [operating system \(OS\)](#) because of the [VNA's](#) software libraries, which provide programmatic control of the [VNA](#), only operating on that [OS](#). The acquisition controller is a command-line application and was developed using the C# language.



```
D:\x64\Acquisition_Controller.exe
Using Interactive Mode
Press CTRL+C to exit
Type in Target to Scan: DeviceA

Sample: 1, BW: 30 KHz, Power: 15 dBm, Port 1 to Port 7 Saved : ( 1/4800)
Sample: 1, BW: 30 KHz, Power: 15 dBm, Port 1 to Port 8 Saved : ( 2/4800)
Sample: 1, BW: 30 KHz, Power: 15 dBm, Port 1 to Port 9 Saved : ( 3/4800)
Sample: 1, BW: 30 KHz, Power: 15 dBm, Port 1 to Port 10 Saved : ( 4/4800)
Sample: 1, BW: 30 KHz, Power: 15 dBm, Port 1 to Port 11 Saved : ( 5/4800)
Sample: 1, BW: 30 KHz, Power: 15 dBm, Port 1 to Port 12 Saved : ( 6/4800)
Sample: 1, BW: 30 KHz, Power: 15 dBm, Port 1 to Port 13 Saved : ( 7/4800)
Sample: 1, BW: 30 KHz, Power: 15 dBm, Port 1 to Port 14 Saved : ( 8/4800)
Sample: 1, BW: 30 KHz, Power: 15 dBm, Port 2 to Port 7 Saved : ( 9/4800)
Sample: 1, BW: 30 KHz, Power: 15 dBm, Port 2 to Port 8 Saved : (10/4800)
Sample: 1, BW: 30 KHz, Power: 15 dBm, Port 2 to Port 9 Saved : (11/4800)
Sample: 1, BW: 30 KHz, Power: 15 dBm, Port 2 to Port 10 Saved : (12/4800)
Sample: 1, BW: 30 KHz, Power: 15 dBm, Port 2 to Port 11 Saved : (13/4800)
Sample: 1, BW: 30 KHz, Power: 15 dBm, Port 2 to Port 12 Saved : (14/4800)
Sample: 1, BW: 30 KHz, Power: 15 dBm, Port 2 to Port 13 Saved : (15/4800)
Sample: 1, BW: 30 KHz, Power: 15 dBm, Port 2 to Port 14 Saved : (16/4800)
Sample: 1, BW: 30 KHz, Power: 15 dBm, Port 3 to Port 7 Saved : (17/4800)
Sample: 1, BW: 30 KHz, Power: 15 dBm, Port 3 to Port 8 Saved : (18/4800)
Sample: 1, BW: 30 KHz, Power: 15 dBm, Port 3 to Port 9 Saved : (19/4800)
Sample: 1, BW: 30 KHz, Power: 15 dBm, Port 3 to Port 10 Saved : (20/4800)
Sample: 1, BW: 30 KHz, Power: 15 dBm, Port 3 to Port 11 Saved : (21/4800)
Sample: 1, BW: 30 KHz, Power: 15 dBm, Port 3 to Port 12 Saved : (22/4800)
Sample: 1, BW: 30 KHz, Power: 15 dBm, Port 3 to Port 13 Saved : (23/4800)
Sample: 1, BW: 30 KHz, Power: 15 dBm, Port 3 to Port 14 Saved : (24/4800)
Sample: 1, BW: 30 KHz, Power: 15 dBm, Port 4 to Port 7 Saved : (25/4800)
Sample: 1, BW: 30 KHz, Power: 15 dBm, Port 4 to Port 8 Saved : (26/4800)
Sample: 1, BW: 30 KHz, Power: 15 dBm, Port 4 to Port 9 Saved : (27/4800)
Sample: 1, BW: 30 KHz, Power: 15 dBm, Port 4 to Port 10 Saved : (28/4800)
Sample: 1, BW: 30 KHz, Power: 15 dBm, Port 4 to Port 11 Saved : (29/4800)
Sample: 1, BW: 30 KHz, Power: 15 dBm, Port 4 to Port 12 Saved : (30/4800)
Sample: 1, BW: 30 KHz, Power: 15 dBm, Port 4 to Port 13 Saved : (31/4800)
Sample: 1, BW: 30 KHz, Power: 15 dBm, Port 4 to Port 14 Saved : (32/4800)
Sample: 1, BW: 30 KHz, Power: 15 dBm, Port 5 to Port 7 Saved : (33/4800)
Sample: 1, BW: 30 KHz, Power: 15 dBm, Port 5 to Port 8 Saved : (34/4800)
```


Figure 4.17: Acquisition in Progress

A screenshot of the acquisition controller can be seen in [Figure 4.17](#). The screenshot shows the output created by the acquisition controller during the characterization of a target, with each line of text indicating a different measurement that was performed. On the

rightmost column, a progress counter shows the number of measurements taken and the amount remaining to complete the characterization.

4.4.1 Startup Arguments

The acquisition controller accepts a set of startup arguments which allows fine-grain control over various aspects of the characterization process. The arguments include parameters related to the operation of the [VNA](#), the characterization criteria to use during the measurement, the location where to store the collected data, and various pieces of metadata to be saved with the measurement. The full listing of startup arguments is available in [Appendix B](#). As part of the initialization of the acquisition controller, the startup arguments are displayed back to the operator to ensure no mistakes were made, an example of which can be seen in [Figure 4.18](#).



```
D:\x64\Acquisition_Controller.exe
Acquisition Controller - Version 1.3.0.0
-----
Using the following settings for acquisition:
DB File: results.db
Samples: 5
Points: 2500
RF Power Levels(dBm): 15,5,0,-5,-20
IF Bandwidths(kHz): 30,50,100
VNA Input 1 Selected RF Port(s): 1,2,3,4,5,6,15,16
VNA Input 2 Selected RF Port(s): 7,8,9,10,11,12,13,14
-----
Dataset: USB_Boards
Target:
Adapter: USB3
Adapter Revision: 1
Sample Notes: {scan:pass1}
-----
Successfully opened SQLite DB File: results.db
-----
Found Acquisition Selector on port: COM59
Connected to: COM59
Multiplex Board ID:B3, Software Version:1.2
-----
Found P9370A VNA
Connection: PXI10::CHASSIS2::SLOT1::FUNC0::INSTR
VNA Serial #: MY57211881
Firmware Rev: A.15.20.07
Driver Vendor: Agilent Technologies
Driver Ident: AgNA
Driver Rev: 3.4.0.0
-----
Startup Complete
-----
Using Interactive Mode
Press CTRL+C to exit
Type in Target to Scan:
```

Figure 4.18: Acquisition Controller Startup

The startup arguments can also dictate the measurement collection mode, which can either be *interactive* or *non-interactive*. In the interactive mode, the operator is prompted for a target name before each characterization is started. After the requested measurements are captured, the acquisition controller stays open and prompts the operator if another target is to be characterized or if the program should be closed. In the non-interactive mode, the target name and identification are provided via the startup arguments, the requested measurements are collected, and the program then automatically exits. The non-interactive mode allows SILGEO to be used within automated test environments.

4.4.2 Functionality

The functionality of the acquisition controller is twofold; it is responsible for initializing the various components in the acquisition system for operation and then controls these components during the capture of measurements.

Initialization

During the initialization stage, the acquisition controller initializes the various components that comprise the acquisition system and configures them for operation. First, the acquisition controller locates and opens the SQLite database within which all the measurements will be saved. If the database does not exist or is formatted incorrectly, a new database with the appropriate schema is created. Next, a connection is opened to the [VNA](#), the required measurement support software is started, and the [VNA](#) is set to the full two-port measurement mode. Finally, a connection is opened with the [MB](#), and the selected [RF](#) ports are set to the defaults. Once initialization is complete, the acquisition controller will either prompt the operator for the name of the target to be characterized if in the interactive mode or immediately begin collecting measurements if in the non-interactive mode.

Measurement Capture

During the characterization of a target, the acquisition controller sequences the [VNA](#), [MB](#), and database to collect and save the requested measurements. With respect to the [VNA](#), the acquisition controller is responsible for the setup of *each* measurement, specifying the [intermediate frequency \(IF\)](#), sweep frequency range, number of points within the frequency sweep, and power level to be used. On the [MB](#) side, the acquisition controller is responsible

for issuing commands to the [MB](#) to select the correct signal path through the multiplexers as part of measuring the different port combinations on a target.

The acquisition controller additionally handles the capture and export of the scattering parameters from the [VNA](#) each time a sweep is completed. Each time the [VNA](#) completes a frequency sweep, the acquisition controller instructs the [VNA](#) library to export the captured scattering parameters as a **S2P** file. This file is then immediately read by the acquisition controller and inserted into the database, along with metadata about the measurement, including the characterization settings, port combination, and power level. The complete list of captured metadata and the overall database schema can be seen in [Appendix C](#). Once this is complete, the acquisition controller loads the next set of measurement settings into the [VNA](#), updates the selected ports on the [MB](#), and instructs the [VNA](#) to take the next measurement. This cycle is then repeated until all the measurements specified by the characterization criteria are captured.

Chapter 5

FT232RL Case Study

To assess the capabilities of SILGEO at the microelectronic level, a comprehensive case study was conducted with the FT232RL microchip manufactured by [Future Technology Devices International \(FTDI\)](#). The case study aimed to quantitatively evaluate SILGEO's ability to discern between authentic and non-authentic microchips. The FT232RL was chosen as the subject of the case study due to its common usage in various industrial and consumer electronics, and because it was widely known to be a target of counterfeiting [75]. Consequently, this allowed SILGEO's performance to be benchmarked against a real threat affecting the industry today.

5.1 FT232RL Overview

The FT232RL is a [USB 2.0 to universal asynchronous receiver-transmitter \(UART\)](#) converter microchip, allowing microcontrollers and other low-speed serial devices to communicate with a [USB](#) host device [76]. In effect, the FT232RL acts as a translator, accepting simple line-level serial communications, adding the necessary format and padding, and then handling all the timing and sequencing to send the information over the high-speed [USB](#) connection. Duplex communication is also supported, so messages can be sent from the [USB](#) host side and converted into messages on the serial [UART](#) side.

The FT232RL is packaged as a 28 lead [shrink small outline package \(SSOP\)](#) and sold in either a tube or reel format. Included on the FT232RL die are the various support components that would typically be needed when used in-circuit, such as a clock oscillator, [electrically erasable programmable read-only memory \(EEPROM\)](#) storage, and [USB](#) termination resistors. Figure 5.1 shows the [FTDI](#) 28 lead [SSOP](#) FT232RL microchip before

soldering [77]. The various die markings on the microchip mean the following: “1803” is the date code, indicating it was manufactured in the third week of 2018; “C” is the revision code; “D9263701” is the lot code, indicating the foundry and manufacturing batch the device is from; and “FT232RL” is the part number.




Figure 5.1: FTDI FT232RL

FT232RL Counterfeiting History

It is unknown when the counterfeiting of the FT232RL first began, but die markings present on some of the counterfeits would indicate around 2011. The existence of counterfeit FT232RL devices was first brought to light by two events in 2014. In early 2014, an independent microelectronics blogger released the first publicly available die shots of genuine and counterfeit FT232RL devices, confirming the presence of counterfeit FT232RL devices in supply chains [78]. Later that year, in October, FTDI released an updated software driver for the Windows OS that would check for the connection of any counterfeit FT232RL devices and, if found, attempt to *brick*, i.e., permanently disable, the counterfeit device [79]. This generated significant pushback from the electronic industry and community because the update was silently distributed via the Windows Update mechanism without notice and because the update effectively punishes the end-users instead of the manufacturers of the products for the inclusion of counterfeit parts. The incident became known as “FTDI-gate” within the news reporting, and FTDI eventually relented and removed the bricking mechanism from future updates.

5.2 Case Study Objective

The goal of this case study was to quantitatively answer the following question regarding the efficacy of SILGEO: Given a set of authentic and non-authentic FT232RL microchips, can SILGEO correctly separate the devices into the appropriate category?

5.3 FT232RL Test Pool

As part of conducting the case study, it was necessary to create a test pool of FT232RL microchips that was representative of the devices available for sale on the market. To ensure the test pool had a diverse amount of manufacturing dates, lot codes, and possibly counterfeit devices, the FT232RL devices were sourced in small quantity orders from multiple different vendors. As part of the sourcing effort, each device was thoroughly inspected to determine its authenticity and was assigned a unique number for identification.

5.3.1 Procurement

Over a period of five months in mid-2020, a total of 260 FT232RL microchips were sourced from nine different vendors. Using [FTDI's](#) list of authorized distributors [80], 80 devices were obtained from Mouser Electronics and Digikey Electronics, spread over three different orders. The remaining 180 devices were obtained from seven distinct sellers on the auction-based websites AliExpress and Ebay. Table 5.1 shows quantities obtained from each vendor.

Vendor	Quantity
Digikey	40
Mouser	40
AliExpress Vendor 1	40
AliExpress Vendor 2	40
Ebay Vendor 1	20
Ebay Vendor 2	20
Ebay Vendor 3	20
Ebay Vendor 4	20
Ebay Vendor 5	20

Table 5.1: FT232RL Vendors and Quantities Obtained

5.3.2 Establishing FT232RL Authenticity

In order to properly assess the efficacy of SILGEO, the authenticity of each FT232RL device in the test pool needed to be independently determined. This involved an inspection of each device for any visible anomalies, functional tests, and X-ray imaging of the internal die and lead frame. Using the results and indications from these tests, each device was assigned a label: *suspected authentic*, *suspected counterfeit*, or *suspected recycled*. The final count for each category is shown in Table 5.2.

Suspected Category	Quantity
Authentic	100
Counterfeit	119
Recycled	41

Table 5.2: FT232RL Suspected Category Counts

Note that for this case study, the category of non-authentic devices was split across two labels: counterfeit and recycled. While recycled devices are generally considered to be counterfeit, the labels were split, in part, to see how the SILGEO characterization process handled recycled devices. Because recycled devices are simply previously used authentic devices, their internals will closely match those of an authentic device and thus may have responses similar to that of an authentic device. By using a separate label, the response behavior and detectability of recycled devices can be better understood and potentially excluded from experiments if needed. Additionally, the term *suspected* is used because the OCM, in this case, FTDI, has not made an official determination as to the authenticity of the devices. The following subsections cover the aforementioned authenticity inspections in greater detail.

Physical Inspection

Using lower-power visual magnification, each FT232RL device was inspected for any signs of previous use, defects, damage, or any other anomalies. The physical inspection followed the National Aeronautics and Space Administration’s (NASA) visual inspection best practices guide [27] and focused on the package, markings, and pins of each device. As part of the inspection, the height, width, and length of the package were measured, and then the surfaces were checked for any signs of sanding or blacktopping. Similarly, for any mold marks present on the package, the depth and diameter were noted, as well as the texture and location of

the mold marks. Next, the identifying markings on the package were recorded and then inspected. Using a powerful light source, the laser engraved markings were checked for any fading or other inconsistencies, including placement and size. Finally, the pins were inspected for any signs of damage or previous use. This includes checking for signs of damage such as missing or bent pins, as well as solder residue or metal wear on the pads of pins.

From these indications, each device was holistically classified. Devices that had package anomalies were classified as possibly non-authentic. Devices that had scratches or damage to the package, markings, or pins were classified as recycled. Devices that had faded, missing, or non-laser engraved markings were classified as non-authentic. Devices that had rough surface finishes on pins or solder residue were classified as recycled.

Functional Test

For each device in the test pool, a functional test was performed to determine if the device initialized correctly and was capable of transmitting information. To expatiate the testing, a custom test jig with a socket to easily swap the FT232RL device was created. The test jig supplied power, a [USB](#) connection to the host computer running the Windows [OS](#), and the necessary support components needed by the FT232RL to operate.

The functional testing of each device involved two checks. Upon initial connection of the FT232RL to the host computer, the Windows “Device Manager” tool was used to verify that the device was correctly identified and the [FTDI](#) driver software was able to properly initialize the device. The second test involved a data loop-back test, where the [UART](#) output (TXD) of the FT232RL is connected back to its [UART](#) input (RXD). This causes any characters sent from the host computer to the FT232RL to be echoed back to the computer, providing a bi-directional verification that the device functioned correctly. A device that failed either of these two tests was deemed to be *non-functional*, which indicated the device is damaged and might be recycled or non-authentic,

FT232RL Unique ID

Hardcoded into each FT232RL during manufacturing is a permanent and unique identification number. This unique identification number is commercially branded as “FTDIDchip-ID” and is intended for use in security and dongle applications of the FT232RL. Because the number is supposed to be globally unique and not user-changeable, it offers a good indicator of the authenticity of each device. Using the same test jig as the functional tests, a custom application was developed to access and read back the unique number from each FT232RL.

The unique value can only be accessed over the FT232RL’s [USB](#) connection and requires using a custom software library developed by [FTDI](#).

Devices with no identification number or had an identification number that was seen multiple times, i.e., a *non-unique* unique identification number, were classified as non-authentic. Devices with a non-repeated unique identification number were classified as *conditionally* authentic, since having a unique identification number is not definitive proof of authenticity. This is because of the possibility that the device’s identification number is, in fact, not unique, but that the repetition was just not seen in this test pool.

X-ray Imaging

Using X-ray radiation, each FT232RL device was imaged and then examined for lead frame, die size, and other package abnormalities. The X-ray images were taken with an X-ray forward voltage of 60 kV and tube current of 4 mA.




Figure 5.2: X-ray Images of FT232RL Devices

Figure [5.2](#) shows two examples of FT232RL devices imaged via X-ray radiation, one genuine and one counterfeit. The images have been magnified but are dimensionally to scale and have the same contrast and brightness levels. Immediately noticeable are the design and size differences between the lead frames on the genuine and counterfeit devices. More subtle is that the genuine die, in Figure [5.2a](#), is rectangular shaped and larger than the Counterfeit circular shaped die of Figure [5.2b](#).

Devices with die size anomalies were classified as non-authentic. Devices with lead frame abnormalities were classified as indeterminate and possibly non-authentic because of two complications. Firstly, multiple styles of lead frames were encountered on the known-genuine devices, complicating the authenticity process. Secondly, visually identical lead frames were seen on some known genuine and suspected counterfeit devices, meaning the lead frames are likely commodity parts, and thus the appearance of a lead frame on its own can not determine authenticity.

5.4 FT232RL Characterization

In order to characterize the devices in the test pool with SILGEO, an adapter board was created to electrical interface with the [SSOP-28](#) pin layout of the FT232RL. Since the FT232RL has 28 pins, and the [MB](#) supports a maximum of 16, a subset of the pins needed to be selected before the adapter board could be designed. It should be noted that throughout the evaluation of various targets with SILGEO, the process of selecting a subset of pins from a target has generally remained more art than science. There is no exact method to determine the “importance” of some pins over others; therefore, the selection process partially relies on assumptions of how a target is constructed internally.

To start the pin selection process, the FT232RL’s datasheet was used to group the pins by function so that duplicative pins could be removed from the selections. For example, “CBUS0”, “CBUS1”, and “CBUS2” all provide the same functionality; thus, it is likely that by measuring only one of the pins, the behaviors of all three can be obtained. From each of these groups, pins were then selected by the importance of their function and assumptions as to how the pins and logic interconnect within the silicon die. Power pins, in particular, were of interest, as the supplied voltages are likely needed in all parts of the silicon die, and thus a stimulus signal on a power pin is likely to propagate through all parts of the die.

Additionally, assumptions were made about the likelihood of a pin’s functionally being implemented similarly on both the genuine and counterfeited device. For example, the internal circuitry used to implement [USB](#) communications is more likely to be shared between genuine and counterfeit devices than, say, the internal logic connected to input-output pins or clock pins. This is because the internal circuitry needed for [USB](#) is dictated, in part, by the [USB](#) specification. In contrast, the circuitry implemented for input-output ports or clock distribution is more likely to be custom designed for each application and thus is more likely to produce a unique response on a given target. The final list of selected pins can be seen in [Table 5.3](#), along with each pin’s function and signal direction.

Pin	Function	Direction	Pin	Function	Direction
1	UART TX Data	Ouput	16	USB -	Input / Output
4	VCC	Power	19	Reset	Input
5	UART RX Data	Input	20	VCC	Power
6	Ring Indicator	Input	21	GND	Power
7	GND	Power	25	GND	Power
11	UART CTS	Input	26	Test	Input
14	Configurable	Input / Output	27	Clock In	Input
15	USB +	Input / Output	28	Clock Out	Output

Table 5.3: FT232RL Pin Selection

To electrically connect with the FT232RL devices, an [SSOP-28](#) compatible *burn-in* socket was used on the adapter board. Burn-in sockets provide a non-permanent electrical connection between a [PCB](#) and a component by mechanically holding the [DUT](#) in place on a set of metal spring contacts. The use of sockets avoids the need for soldering devices to dedicated test boards and allows devices to be measured in rapid succession, reducing the effort and




Figure 5.3: FT232RL Measurement Setup

time required to measure the entire test pool. Figure 5.3 shows the measurement setup used to characterize the FT232RL devices. The black circuit board at the bottom of the figure is the MB, and the green circuit board connected to it on top is the adapter board. In the center of the adapter board is the SSOP-28 socket, shown in the with socket lid in the open position and with an FT232RL inserted. When characterizing a device, the socket lid is latched in the closed position to place mechanical pressure on the FT232RL and the spring contact below it, ensuring the electrical connections are consistent and uninterrupted during a measurement.

Characterization Procedure

To characterize each FT232RL device in the test pool, the following procedure was used. Each device was inserted into the socket on the adapter board, and the socket lid was closed securely. The identification of the device was recorded, and a series of measurements was taken with the characterization criteria shown in Table 5.4.

Characterization Parameter	Value
Frequency Range	300 kHz to 4.5 GHz
Frequency Point Count	10,000 points
Frequency Point Distribution	Linear
Power Level	−5 dBm
Port Combination	All 16 ports
Scattering Parameter Direction	Reflection & Transmission
VNA Intermediate Frequency Bandwidth	100 kHz

Table 5.4: FT232RL Characterization Criteria

5.5 Case Study Results

The analysis of the collected FT232RL characterizations was done in two parts. The measurements were first visualized with a clustering algorithm to understand the overall association between the different devices. Next, the authenticity of each device was predicted and compared against the independently determined authenticity label.

5.5.1 FT232RL Clustering

Starting with the raw characterizations for all 260 devices, [principal component analysis \(PCA\)](#) was used to obtain a set of *principal components*, that described the variability between the FT232RL devices. With [PCA](#), each principal component describes a portion of the variance found within the underlying features, with the first component describing the greatest amount and each additional component afterward describing successively smaller portions of the variance. Using the first and second [PCA](#) components values as x and y , a scatter plot was created of all 260 devices, and is shown in Figure 5.4. The color of each point indicates the vendor the device was sourced from, and the symbol indicates the independently determined authenticity, with a *circle* for authentic devices, a *cross* for counterfeit devices, and a *triangle* for recycled devices.




Figure 5.4: FT232RL Clustering, Colored by Vendor

Discussion

Several noteworthy clustering patterns can be observed in Figure 5.4. All the authentic devices were clustered together around -50 on the first component and 0 on the second component, and no authentic devices were found within clusters composed of primarily counterfeit devices. The large distance between the various clusters appears to indicate that the internals of the authentic and various counterfeit designs are significantly different, potentially meaning the counterfeits are ground-up reimplementations. Based on the several distinct clusters of counterfeit devices, the test pool may contain three or four different variants of counterfeits.

The main cluster of recycled devices can be seen just left of the main authentic cluster, with several recycled devices appearing very close to, but not within, the main cluster of authentic devices. The distance between the main recycled cluster and the authentic cluster may be attributable to the aging of the recycled devices due to prior usage or may indicate an older version of the silicon design. Note that the separation between recycled and authentic devices occurred mainly on the first component and not the second component. As for the recycled devices that appear very close to the main authentic cluster, is possible those devices were newer or had different usage patterns. Three outliers, two counterfeit and one recycled, can be seen at the top of the figure. While both of the counterfeit devices were found to be non-functional, the recycled device was functional. It is possible these devices are outliers due to internal damage caused by [electrostatic discharge \(ESD\)](#).

5.5.2 FT232RL Authenticity

Due to the close distance between the clusters of recycled and authentic devices, it was decided to withhold the 41 suspected recycled FT232RL devices from the authenticity analysis. This decision reduced the test pool to a total of 219 FT232RL devices and now only included the devices suspected of being authentic and counterfeit. The authenticity of each FT232RL device was predicted by computing a similarity score between the characterizations of each device with the characterizations of four known-authentic devices. The four known-authentic devices were picked randomly from a shortlist of devices that were sourced from authorized [FTDI](#) vendors and exhibited no physical, functional, or X-ray anomalies.

A similarity score was then computed between the characterization of each device and the set of known-authentic characterizations. The similarity score was based on the [dynamic time warping \(DTW\)](#) distance metric [58], which computed the difference between the collected frequency sweep responses. More precisely, at each port combination and scattering

parameter, DTW computed the distance between the response of the device under evaluation and the response from each of the four known-authentic devices. The four distance metrics were then averaged together to create a single distance metric between the device and the set of known-authentic devices, for each port combination and scattering parameter. The final similarity score for the device was then obtained by computing the average of the single distance metrics at each port combination and scattering parameter.

Figure 5.5 shows a visualization of the computed similarity scores. The color of each point indicates the independently determined authenticity, with green for suspected authentic devices and red for suspected counterfeit devices. A similarity score of 1 on this scale indicates a device that is indistinguishable from the known-authentic examples, and a score of 0 indicates a device that is entirely inconsistent with the known-authentic examples.




Figure 5.5: Similarity Score, Colored by Authenticity

Using the similarity scores, a threshold was applied to predict the authenticity of the devices. Devices with similarity scores higher than or equal to 0.8 were predicted to be authentic. Devices with scores less than 0.8 were predicted to be counterfeit. To ensure the randomly selected four known-authentic devices did not impart bias onto the predictions, the entire procedure outlined above, including the computation of similarity scores for each FT232RL device, was repeated 100 times, with each round using a different random selection of four known-authentic devices. The accuracy of the predictions can be seen in the confusion matrix, shown in Figure 5.6. The true labels of the confusion matrix represent the independently determined authenticity labels, and the predicted labels represent the predictions obtained with SILGEO.




Figure 5.6: Confusion Matrix of the Authenticity Results

Discussion

From the confusion matrix shown in Figure 5.5, SILGEO had 100% accuracy when predicting the authenticity of the FT232RL devices. This 100% accuracy was maintained across all 100 rounds where different selections of known-authentic devices were used. This accuracy was likely achieved, in part, because of the construction and implementation differences between the authentic and counterfeit FT232RL devices. The SILGEO characterization technique detected these internal differences and reflected them in the similarity scores, where there is a significant difference between the scores for authentic and counterfeit devices.

From the similarity score plot shown in Figure 5.5, a linear separation can be seen between the devices suspected of being authentic and devices suspected of being counterfeit. Clustering can also be seen among the similarity scores, with a large counterfeit cluster around 0.17, one at 0.07, and two more around 0.5. The similarity scores for authentic devices ranged from 0.85 to 0.95.

5.6 Case Study Conclusions

Overall, SILGEO performed well during the quantitative evaluation and correctly classified 100% of the FT232RL microchips within the test pool. From the reduced pool of 219 devices, SILGEO was able to correctly discern between the 100 suspected authentic FT232RL devices and 119 suspected counterfeit FT232RL devices. When including the 41 suspected recycled devices, which were held out during the classification experiment, SILGEO was able to correctly cluster all three classes of FT232RL devices, totaling 260 devices.

Examining the characteristics of the FT232RL devices in the test pool, the suspected counterfeit devices tended, but not always, to have poor or misaligned laser markings, with some even appearing to have markings that were applied via silk-screen instead of laser etched. The suspected counterfeit devices also tended to come in package sizes that differed from the suspected authentic devices. Some devices were shorter in length and others slightly too wide, causing them to get stuck in the [SSOP-28](#) socket. Nearly all the suspected counterfeit devices had mold markings that did not match the size and location of those present on the suspected authentic devices or had no mold markings at all. All suspected counterfeit devices failed the unique ID test, and 58% of the suspected counterfeit devices were found to be entirely non-functional. Under X-ray imaging, the suspected counterfeit devices had varying lead-frames size and four discernible unique die sizes. All 41 suspected recycled were electrically functional, had a unique ID, and had evidence of prior use, including bent pins and solder residue. The suspected authentic devices all appeared physically similar, including under X-ray imaging, with the only noticeable difference being the lot code value.

This case study broadly confirmed the dangers of sourcing components from online auction-based websites. Of the 260 FT232RL devices sourced, more than 60% of the devices were suspected of being non-authentic, with 119 suspected of being counterfeit and 41 suspected of being recycled. Notably, two of the five auction-based vendors did, in fact, sell what appeared to be authentic devices; therefore, not all auction vendors can be assumed to be malicious. Using the SILGEO clustering results and die size obtained via X-ray imaging, it is estimated that the test pool contains three or four different counterfeit variants. This case study also revealed evidence of auction vendors operating multiple profiles on the same site. After comparing the purchasing records and shipping information, on two occasions, orders from two entirely different vendors shipped from the same address and had identical customs information and post markings. It is suspected that the sellers are operating multiple pages for business continuity reasons; if one of their pages gets closed due to selling counterfeits, they still have other pages they can use to sell from.

As an aside, one may argue that ordering microchips from auction-based websites is asking for trouble or unrealistic as no reputable company would purchase in this manner; and

that the whole problem of counterfeit microchips in the supply chain can simply be avoided by only buying from authorized distributors. With the electronic supply chain disruptions experienced in 2021 and 2022, the idea of ordering from non-authorized or auction-style distributors no longer is as far-fetched and, for some, has become the norm [81]. To drive this point home, at the time of writing this thesis (December 2021), the official [FTDI](#) website shows that of the eight authorized distributors for the FT232RL, *none* have any stock, and there is a 44-week lead time before any become available again [80].

5.6.1 Limitations

This case study revealed some limitations with respect to the applicability of SILGEO. Firstly, SILGEO has a limited ability to identify recycled devices. While in this case study, the suspected recycled devices did form their own distinct clusters, which were separate from the genuine devices, shown in Figure 5.4, there is no guarantee that this will occur for all targets evaluated with SILGEO. Even though the clusters of genuine and recycled were separate, the [PCA](#) distance in-between the clusters was minimal, with some individual recycled devices being closer and nearly overlapping with the genuine cluster. The problem posed to SILGEO is that recycled devices are essentially just genuine devices that have been used, and relying on the effects of device aging to separate recycled devices from new devices is unlikely to be reliable across all classes of devices. In particular, the limitation will be most acute for devices that exhibit little or no wear over their life cycles, such as components designed for harsh environments.

Secondly, the naturally occurring manufacturing variability between authentic devices poses a concern to SILGEO’s accuracy. Since SILGEO determines the authenticity of a device by comparing it to a known-good device or collection of known-good devices, classes of targets that have high manufacturing variability, such as devices that are hand assembled, may lead to authentic parts being erroneously marked as counterfeit and vice versa. Having high variability among the known-good devices can cause “differences” that are attributable to the counterfeit nature of a device to be hidden among the valid range of responses from the known-good devices, thus increasing the difficulty of spotting counterfeit devices. This potential limitation was partially confirmed in this case study with the similarity scores variability for the known authentic FT232RL devices, shown in Figure 5.5. Rather than clustering as a point near the value of 1.0, indicating the devices are perfectly similar, the scores for the known authentic devices were spread over a range of 0.82 to 0.91, covering 9.4% of the total similarity scale.

5.6.2 Lessons Learned

The lessons discussed in this section are drawn from this case study and the broader experience of evaluating targets with SILGEO. The first, and most important lesson, is that the availability of known-good parts to compare against potential suspected parts is critical to SILGEO's operation, and their availability can not be assumed. As part of estimating SILGEO's suitability on potential targets, difficulty was encountered when attempting to source and identify multiple known-good devices, particularly for old or military-related classes of targets. The lack of multiple known-good devices reduces SILGEO's tolerance to measurement and manufacturing variability, leading to reduced confidence when classifying the authenticity of devices. However, if no known-good devices are available, then SILGEO is restricted to simply detecting that tampering occurred, comparing the target with previous characterizations of the same target.

Another lesson, which was gleaned through hours of target characterization, is that sockets are a definitive solution to the problem of interfacing with targets quickly and are superior to the previous approach of soldering targets to custom circuit boards. However, sockets are not without their problem, and issues begin to appear when measuring many targets. In particular, when inserting devices into the socket in a rapid manner, it is possible to accidentally insert the device in the wrong direction or have target pins slightly misaligned with the socket contacts. In either case, the improper placement of the device within the socket can lead to increased measurement variability and even improper authenticity classifications. Further complicating matters is that while a socket can be correctly sized for the genuine devices, there is no guarantee that the non-genuine devices will also fit. This was encountered with certain non-genuine FT232RL devices, which had slightly larger package sizes, causing them to occasionally get stuck in the socket.

One drawback that was noticed when using sockets, compared to soldering targets, is that sockets can negatively affect the channel-to-channel isolation of the measurements. This is due to the closeness of the metal spring contacts to each other, needed as part of contacting the DUT, which allows signals to couple across contacts. Additionally, because sockets are mechanical devices, there is a concern that they can wear out after a certain number of device insertions which can attenuate and distort the carried signals. However, in practice, this number is typically large enough that it takes multiple years of use before a performance impact is seen. For example, the SSOP-28 socket used to interface with the FT232RL devices is rated for a minimum of 10,000 insertions [82].

The final takeaway is that while creating an adapter board for each new target is certainly less work than having to remake an entire MBs each time, designing and manufacturing adapter boards remains a non-trivial task. Specifically, the sourcing of electrical connectors

needed to interface with targets and determining which subset of pins to connect with electrically presents an ongoing challenge. On more than one occasion, significant effort and time were required to determine the correct counterpart connector needed to interface with a device, with military connectors in particular, such as the D38999, being especially convoluted. Furthermore, the process of selecting a subset of 16 pins from targets with high pin counts is fraught and overly reliant on assumptions and guesses as to how a target is constructed internally. Selecting a subset of pins means there will be pins on the target that are not characterized, thus potentially missing valuable information that could be used to determine authenticity.

Chapter 6

Future Work

While significant efforts have gone into the development of SILGEO, many directions of future work and research remain. At a high level, these areas can be broken down into three main categories: improvements to the current implementation of SILGEO, alternative characterization techniques, and additional applications where SILGEO could be of use. This chapter aims to briefly introduce each of these categories and provide some suggestions for next steps.

6.1 Calibration Improvements

Focusing on the SILGEO implementation presented in Chapter 4, one of the most pressing concerns is the lack of a complete signal path calibration capability when measuring DUTs via the MB and companion adapter boards. Currently, at the time of measurement, a calibration plane can only be established at the end of the SMA cables, meaning that while the VNA can reliably remove the effects of the cables, anything that lies beyond will not be error-corrected in the saved measurement. As a post-processing step, a de-embedding profile can be used to remove the effects of the MB; however, this profile is static and represents the board and its effects at a fixed point in time, and therefore cannot correct for ambient temperature, internal changes within the PCB structure, or connector wear. It should be noted that the de-embedding profile for the MB does not remove the effects or account for the adapter board mated with MB, as no de-embedding profiles exist for these boards.

The main challenge, and thus the opportunity for future development, with moving the calibration plane from the end of the SMA cables to the DUT connector on the adapter board

is largely twofold. Firstly, since the connectors that interface directly with the [DUT](#) are often not designed to carry microwave [RF](#) signals, they are not reliable places to establish a calibration plane. Issues of signal leakage and mechanical play in the connector construction results in calibrations that are effectively not repeatable [57]. This challenge is the same reason why no de-embedding profiles have been created for the various adapter board designs. Secondly, the number of channels on the [MB](#), and consequently on the adapter board, creates a logistical challenge when manually performing the full signal path calibration. For 16 channels, a calibration at the time of measurement correspondingly requires at minimum 8 individual 2-port calibrations for the reflection paths, and potentially up to 64 when accounting for all transmission paths [57]. For context, calibrating at the end of the [SMA](#) cables requires only one 2-port calibration.

To solve these concerns, the use of multi-port calibration standards, which can calibrate 4 or 8 ports at once, should be evaluated. As a short-term compromise, dedicated calibration ports could also be added to the adapter boards, moving the calibration plane to include the adapter board, but leaving the [DUT](#) connector itself not error-corrected. Long term, the use of in-situ calibration techniques, such as [automatic fixture removal \(AFR\)](#), which use time-domain measurements and gaiting to synthesize an error-correction profile, should be investigated. If successful, these methods would completely sidestep the aforementioned issues, as they do not involve the connection and measuring of a calibration standard as part of establishing a calibration plane [83].

6.2 Harnessing the Time-Domain

To date, SILGEO has relied exclusively on characterizing and analyzing the [DUTs](#) in the frequency-domain. While this has yielded excellent results, research should be conducted to see if operating in the time-domain offers SILGEO any benefits or additional capabilities. As a brief review, the time-domain studies the change of signals and functions with respect to time, whereas the frequency-domain does so with respect to frequency. Given the characterization of an ideal [DUT](#), the time-domain response and the frequency-domain response will each present the captured information differently, but crucially *both* will contain the same information, and in fact, conversion between the two is possible via mathematical transformation. For most types of problems, there is a preferred domain to work in, as it can greatly simplify the mathematics involved, and in certain cases, the analyses can only be performed in the time-domain, or vice versa, in the frequency-domain. By considering both representations, it can be ensured that SILGEO is utilizing all of the information present in the collected measurements.

For SILGEO, two possible applications of the time-domain should be explored: the use of time-domain equipment in the characterization of DUTs, and the application of time-domain analysis to improve the frequency-domain measurements captured by the VNA. To characterize a DUT directly in the time-domain, a measurement technique known as *time-domain reflectometry (TDR)* & *time-domain transmissometry (TDT)* (or TDR for short, with transmission variant being implied) can be used, where a step pulse with a very fast rising edge is injected into a DUT, and the resulting reflected and transmitted waveforms are recorded [57]. One of the main opportunities of measuring and analyzing with TDR is that since the time component is involved, the physical location of where an impedance change occurs along the signal path can be known. It remains to be tested, but this could enable SILGEO to pinpoint within the DUT itself where a modification or change was made. Another potential benefit of characterizing in the time-domain is that it could enable the increased miniaturization of SILGEO since the equipment required for TDR is generally simpler and more compact than the frequency-domain equivalents.

Separately, the utilization of time-domain analysis with the existing frequency-domain measurements may yield improvements in characterization accuracy and stability. Specifically, a concept called *gating* could enable SILGEO to better recognize and remove changes present in the measurement signal path that are not related to the DUT [57]. The process of gating works in the following manner: a DUT measurement is taken in the frequency-domain, and the *inverse fast Fourier transform (IFFT)* algorithm is then applied to obtain a time-domain representation of the response. Using this representation and speed of propagation information, regions of response in the time-domain can then be mapped to where they physically occur along the signal path. Unwanted regions can be identified, removed, and the modified measurement is then converted back to the frequency-domain via the *fast Fourier transform (FFT)* algorithm. The result is a frequency-domain representation of the DUT with the effects contributed by the undesired sections removed as if they were never there [57].

This powerful technique could allow a DUT to be decomposed into multiple smaller characterizations, or *slices*, taken at different depths of propagation. This would potentially help with the fingerprinting of complex devices and circuit boards, where small malicious changes can be hidden by larger benign manufacturing variations. By slicing, a target can be evaluated by a set of smaller responses, rather than just one cumulative response. Gating may also help with the problem of connector unreliability, where at microwave frequencies, minor movements of the non-RF grade DUT connectors introduce discontinuities that SILGEO then recognizes as potential changes. By gating the VNA measurement at a position right after the unreliable DUT connector, the characterization should be free of any connector variations. Of course, it should also be noted that gating is not a perfect process, and neither

are the [FFT](#) and [IFFT](#) algorithms, so some imperfections and distortion will be imparted onto the measurements. As part of evaluating the use of gating, it will be important to verify that these distortions are not significant enough to cause issues of their own during the characterization process.

6.3 Additional Applications

Beyond the detection of counterfeit electronic devices, the applicability of SILGEO to other problems should be investigated. One area where SILGEO may be of use is the challenge of detecting [electrostatic discharge \(ESD\)](#) damage in microelectronic devices. When [ESD](#) current flows through a component, severe damage can be done to the internal structures resulting in either catastrophic failure or a latent defect. Catastrophic failures, such as the melting of junctions, cratering of oxide layers, and fusing of bond wires, often results in the partial or complete loss of device functionality, and as such, can easily be detected with methods such as functional testing. On the other hand, latent defects are significantly harder to detect and pose a more significant threat, as the component has been weakened but continues to function [\[84\]](#). Depending on the discharge, latent defects can cause components to have shorter lifespans, reduced tolerance to mechanical and environmental disturbances, and altered operating tolerances. The opportunity for applying SILGEO is that [ESD](#) damage, both catastrophic and latent, creates permanent changes in the physical geometry of the component. Thus the same technique that currently translates device geometry into fingerprints for counterfeit detection could be repurposed for the detection of [ESD](#) damage.

Since catastrophic failures are readily detectable with existing techniques, research should focus on how well SILGEO can detect the presence of latent defects. To reinforce the need for a latent defect detection capability, to date, most examples of latent damage have only been discovered through component failure investigations that included die decapsulation and imaging [\[84\]](#). Moreover, since die imaging and the accompanying analysis are time-consuming and expensive, they may not always be performed, leading to many incidents of latent defect caused component failure to never be caught. To draw additional parallels to the counterfeit detection application of SILGEO, the only way to know with absolute certainty that a component has no latent [ESD](#) damage is through decapsulation, metalization removal, and inspection with a scanning electron microscope. Of course, the problem now is that this component is no longer functional. Assuming latent defects have enough magnitude to be detected, SILGEO could be employed in various capacities. It could provide a spot-checking capability as part of randomized quality control checks, or entire batch checking, such as when a tray of components is suspected of being mishandled. The idea here is that unintentional

changes caused by [ESD](#) are in principle no different from intentional changes caused by an attacker, and the same fingerprinting technique should be applicable to both.

Similarly, SILGEO may be of use in verifying that *intentional* permanent changes to the silicon geometry have actually occurred, such as with electronic fuses. Also referred to as e-fuses, electronic fuses are memory cells that can be irreparably “blown” in order to permanently write a value. Devices often use e-fuses to store serial numbers, version numbers to prevent software down-grading, and other security-related values. Here, the functionality that SILGEO offers can be one of independent verification, or value extraction, in the sense of making sure that a value returned by the software on a device matches the actual value blown into the silicon of the device. This line of inquiry is, of course, conditionate on the setting of e-fuses being a change in silicon geometry that SILGEO can reliably detect with a high level of accuracy. To further complicate matters, some e-fuse writes are single-bit changes, which may be undetectable even at the highest power levels.

Chapter 7

Conclusions

This thesis introduced SILGEO, a hardware fingerprinting technique that counters the growing problem of counterfeit, compromised, or tampered devices being found within electronic supply chains. Compared to other approaches used for determining the authenticity of electronic devices, SILGEO offers several benefits in that its assessment approach is blackbox, nondestructive, and can be performed on devices in situ. Unlike the functional or parametric electrical testing approaches, SILGEO requires no information on the device’s functionality, operating voltages, timing, or external electrical connections to determine its authenticity. This allows SILGEO to authenticate devices even when the original manufacturer no longer exists or is unwilling to share the special or proprietary information typically needed to perform functional or parametric tests. In contrast to decapsulation and X-ray imaging, SILGEO is not hazardous to the [device under test \(DUT\)](#) or the operator performing the evaluation and leaves no temporary or permanent damage. Among other detection approaches, SILGEO is unique in its ability to evaluate the authenticity of targets in situ, avoiding the need to desolder or remove components from the larger systems they are a part of.

As part of a case study evaluating SILGEO, this thesis confirmed the increasingly common threat of finding counterfeit devices within supply chains. Of the 260 FT232RL devices sourced, more than 60% of the devices were found to be non-authentic, with 119 suspected of being counterfeit and 41 suspected of being recycled. From this pool of FT232RL devices, SILGEO was able to correctly discern between the authentic and counterfeit devices with 100% accuracy. This thesis additionally presented an overview of the electronic supply chains and manufacturing steps involved in creating modern technological devices. The threat landscape facing these supply chains, as well as the current state of mitigation methods, was also detailed.

References

- [1] Carlos Moreno, Sebastian Fischmeister, Umut Guler, and Philippe Vibien. *Counterfeit Detection of FT232RL with Zilky*, August 2020.
- [2] Carlos Moreno, Sebastian Fischmeister, and Philippe Vibien. *A Method and Apparatus for Detection of Counterfeit Parts, Compromised or Tampered Components or Devices, Tampered Systems Such as Local Communication Networks, and for Secure Identification of Components.*, 2020. Patent Application: PCT/CA2020/051266.
- [3] *The Committee’s Investigation Into Counterfeit Electronic Parts in the Department of Defense Supply Chain: Hearing before the Committee on Armed Services*. Senate Hearing: 112-340. U.S. G.P.O., Washington D.C., November 2012. <https://www.govinfo.gov/content/pkg/CHRG-112shrg72702/pdf/CHRG-112shrg72702.pdf>.
- [4] *Inquiry into Counterfeit Electronic Parts in the Department of Defense Supply Chain: Report for the Committee on Armed Services*. Senate Hearing: 112-167. U.S. G.P.O., Washington D.C., May 2012. <https://www.armed-services.senate.gov/imo/media/doc/Counterfeit-Electronic-Parts.pdf>.
- [5] *DOD Supply Chain: Suspect Counterfeit Electronic Parts Can Be Found on Internet Purchasing Platforms*. GAO-12-375. United States Government Accountability Office, Washington D.C., February 2012. <https://www.gao.gov/assets/gao-12-375.pdf>.
- [6] *Defense Supplier Base: DOD Should Leverage Ongoing Initiatives in Developing Its Program to Mitigate Risk of Counterfeit Parts*. GAO-10-389. United States Government Accountability Office, Washington D.C., March 2010. <https://www.gao.gov/assets/gao-10-389.pdf>.
- [7] Deborah Elms. *Global Value Chains In A Changing World*. World Trade Organization and Temasek Foundation Centre for Trade & Negotiations, Geneva, January 2013. doi: [10.30875/0b68ab34-en](https://doi.org/10.30875/0b68ab34-en).

- [8] Edlyn V. Levine. *The Die is Cast: Hardware Security is Not Assured*. 18-4. Association for Computing Machinery, New York, USA, August 2020. doi: [10.1145/3424302.3431245](https://doi.org/10.1145/3424302.3431245).
- [9] Stacey Frederick, Penny Bamber, Lukas Brun, Jaehan Cho, Gary Gereffi, and Joonkoo Lee. *Korea in Global Value Chains: Pathways for Industrial Transformation*. Duke University Global Value Chains and Korea Institute For Industrial Economics & Trade, September 2017. <https://gvcc.duke.edu/wp-content/uploads/Duke-GVCC-KIET-Korea-in-GVCs-Pathways-for-Industrial-Transformation.pdf>.
- [10] Gary May and Costas Spanos. *Fundamentals of Semiconductor Manufacturing and Process Control*. IEEE, Piscataway, NJ, 2006. ISBN: 9780471784067.
- [11] Mark Tehranipoor, Ujjwal Guin, and Domenic Forte. *Counterfeit Integrated Circuits, Detection and Avoidance*. Springer International Publishing, 2015. doi: [10.1007/978-3-319-11824-6](https://doi.org/10.1007/978-3-319-11824-6).
- [12] IEEE 1076-2019 Standard for VHDL Language - Reference Manual. IEEE, 2019. doi: [10.1109/IEEESTD.2019.8938196](https://doi.org/10.1109/IEEESTD.2019.8938196).
- [13] IEEE 1800-2017 Standard for SystemVerilog - Unified Hardware Design, Specification, and Verification Language. IEEE, 2017. doi: [10.1109/IEEESTD.2018.8299595](https://doi.org/10.1109/IEEESTD.2018.8299595).
- [14] Wei-Ting Chen, Mohammad Saud Khan, and Rehan Iftikhar. Value Creation Analysis: Semiconductor Intellectual Property Business Models and Dynamic Capabilities. *IEEE Transactions on Engineering Management*, pages 1–12, 2021. doi: [10.1109/TEM.2020.3044283](https://doi.org/10.1109/TEM.2020.3044283).
- [15] Phil Hughes. *Arm Partners Are Shipping More Than 900 Arm-Based Chips Per Second Based On Latest Results*. Arm Ltd., May 2021. <https://www.arm.com/company/news/2021/05/arm-partners-are-shipping-more-than-900-arm-based-chips-per-second>.
- [16] Saif M. Khan, Alexander Mann, and Dahlia Peterson. *The Semiconductor Supply Chain: Assessing National Competitiveness*. Center for Security and Emerging Technology, January 2021. doi: [10.51593/20190016](https://doi.org/10.51593/20190016).
- [17] Hui-Chih Hung, Yu-Chih Chiu, and Muh-Cherng Wu. Analysis of Competition Between IDM and Fabless-Foundry Business Models in the Semiconductor Industry. *IEEE Transactions on Semiconductor Manufacturing*, 30(3):254–260, 2017. doi: [10.1109/TSM.2017.2699739](https://doi.org/10.1109/TSM.2017.2699739).

- [18] Tiffany Zinn. *PCB Material Selection for RF, Microwave and Millimeter-wave Design*. Isola, 2016. <http://www.isola-group.com/wp-content/uploads/PCB-Material-Selection-for-RF-Microwave-and-Millimeter-wave-Designs.pdf>.
- [19] Swarup Bhunia and Mark Tehranipoor. Printed Circuit Board (PCB): Design and Test (Chapter 4). In *Hardware Security*, pages 81–105. Morgan Kaufmann, 2019. doi: [10.1016/B978-0-12-812477-2.00009-5](https://doi.org/10.1016/B978-0-12-812477-2.00009-5).
- [20] Robert L. Myers. *The Electronics Assembly Handbook*. Springer Berlin Heidelberg, Berlin, 1998. doi: [10.1007/978-3-662-13161-9_92](https://doi.org/10.1007/978-3-662-13161-9_92).
- [21] Yongmin Chen and Keith Maskus. Vertical Pricing and Parallel Imports. *The Journal of International Trade & Economic Development*, 14(1):1–18, 2005. doi: [10.1080/0963819042000333225](https://doi.org/10.1080/0963819042000333225).
- [22] Giovanna Mura. Reliability Concerns From The Gray Market. In *Microelectronics Reliability - 29th European Symposium on Reliability of Electron Devices, Failure Physics and Analysis (ESREF 2018)*, volume 88-90, pages 26–30, September 2018. doi: [10.1016/J.MICROREL.2018.06.098](https://doi.org/10.1016/J.MICROREL.2018.06.098).
- [23] Don Clark. *The Tech Cold War’s ‘Most Complicated Machine’ That’s Out of China’s Reach*. New York Times, July 2021. <https://www.nytimes.com/2021/07/04/technology/tech-cold-war-chips.html>.
- [24] Andrew Huang. *Supply Chain Security: “If I were a Nation State...”*. BlueHat IL, February 2019. <https://www.youtube.com/watch?v=RqQhWitJ1As>.
- [25] AS5553A - Fraudulent/Counterfeit Electronic Parts; Avoidance, Detection, Mitigation, and Disposition. SAE International, January 2013. <https://saemobilus.sae.org/content/AS5553A>.
- [26] Damir Akhoundov. *2017 ERAI Reported Parts Analysis*. ERAI, Inc., June 2018. https://www.eraai.com/eraai_blog/3139/_2017_eraai_reported_parts_analysis.
- [27] Leblanc Abesamis. *NASA Counterfeit Parts Awareness and Inspection*. MTTC Training Series, Manufacturing Technology Transfer Center. JPL & NASA. https://mttc.jpl.nasa.gov/pdf/NASA_Counterfeit_Training_unlimited_distribution_handout.pdf.
- [28] M. Pecht and S. Tiku. Bogus: Electronic Manufacturing and Consumers Confront a Rising Tide of Counterfeit Electronics. *IEEE Spectrum*, 43(5), May 2006. doi: [10.1109/MSPEC.2006.1628506](https://doi.org/10.1109/MSPEC.2006.1628506).

- [29] Lawrence W. Kessler and Thomas Sharpe. Faked Parts Detection - New Techniques to ID Counterfeit ICs. *The Journal for Surface Mount and Electronics Assembly*, June 2010. <https://www.era1.com/customuploads/ca/sc/newtechnquestoidcounterfeitics.pdf>.
- [30] *Electronics Waste Management in the United States Through 2009*. EPA 530-R-11-002. United States Environmental Protection Agency, Washington D.C., May 2011. <https://nepis.epa.gov/Exe/ZyPDF.cgi/P100BKKL.PDF?Dockkey=P100BKKL.PDF>.
- [31] Domenic Forte and Rajat Chakraborty. *Counterfeit Integrated Circuits: Threats, Detection, and Avoidance*. 2018. <https://ches.iacr.org/2018/slides/ches2018-tutorial1-slides.pdf>.
- [32] Ujjwal Guin, Ke Huang, Daniel DiMase, John M. Carulli, Mohammad Tehranipoor, and Yiorgos Makris. Counterfeit Integrated Circuits: A Rising Threat in the Global Semiconductor Supply Chain. *Proceedings of the IEEE*, 102(8):1207–1228, 2014. doi: [10.1109/JPROC.2014.2332291](https://doi.org/10.1109/JPROC.2014.2332291).
- [33] Ujjwal Guin, Daniel Dimase, and Mohammad Tehranipoor. A Comprehensive Framework for Counterfeit Defect Coverage Analysis and Detection Assessment. *Journal of Electronic Testing*, 30(1):25–40, 2014. doi: [10.1007/S10836-013-5428-2](https://doi.org/10.1007/S10836-013-5428-2).
- [34] United States District Court, For the Northern District of California - San Jose. *United States of America v. United Microelectronics Corporation et al.*, December 2018. <https://www.justice.gov/usao-ndca/press-release/file/1107341/download>.
- [35] Martin Rowe. *Cloned ICs: Better Specs Than Authentics*. Electronics Purchasing and The Supply Chain - EPSNews, December 2018. <https://epsnews.com/2018/12/27/cloned-ics-better-specs-than-authentics/>.
- [36] Maya Posch. *STM32 Clones: The Good, The Bad And The Ugly*. Hackaday, October 2020. <https://hackaday.com/2020/10/22/stm32-clones-the-good-the-bad-and-the-ugly/>.
- [37] Johannes Obermaier, Marc Schink, and Kosma Moczek. One Exploit to Rule them All? On the Security of Drop-in Replacement and Counterfeit Microcontrollers. In *14th USENIX Workshop on Offensive Technologies (WOOT 20)*. USENIX Association, August 2020. <https://www.usenix.org/system/files/woot20-paper-obermaier.pdf>.

- [38] Samuel T. King, Joseph Tucek, Anthony Cozzie, Chris Grier, Weihang Jiang, and Yuanyuan Zhou. Designing and Implementing Malicious Hardware. In *First USENIX Workshop on Large-Scale Exploits and Emergent Threats (LEET 08)*, San Francisco, CA, 2008. USENIX Association. https://www.usenix.org/legacy/event/leet08/tech/full_papers/king/king.pdf.
- [39] Lang Lin, Markus Kasper, Tim Güneysu, Christof Paar, and Wayne Burleson. Trojan Side-Channels: Lightweight Hardware Trojans through Side-Channel Engineering. In *Lecture Notes in Computer Science*, pages 382–395. Springer Berlin Heidelberg, 2009. doi: [10.1007/978-3-642-04138-9_27](https://doi.org/10.1007/978-3-642-04138-9_27).
- [40] Georg T. Becker, Francesco Regazzoni, Christof Paar, and Wayne P. Burleson. Stealthy Dopant-Level Hardware Trojans. In *Cryptographic Hardware and Embedded Systems - CHES 2013*, pages 197–214. Springer Berlin Heidelberg, 2013. doi: [10.1007/978-3-642-40349-1_12](https://doi.org/10.1007/978-3-642-40349-1_12).
- [41] Sean Gallagher. *Photos of an NSA "upgrade" Factory Show Cisco Router Getting Implant*. Ars Technica, May 2014. <https://arstechnica.com/tech-policy/2014/05/photos-of-an-nsa-upgrade-factory-show-cisco-router-getting-implant/>.
- [42] Jordan Robertson and Michael Riley. *The Big Hack: How China Used a Tiny Chip to Infiltrate U.S. Companies*. Bloomberg, Oct 2018. <https://www.bloomberg.com/news/features/2018-10-04/the-big-hack-how-china-used-a-tiny-chip-to-infiltrate-america-s-top-companies>.
- [43] National Security Agency. *Advanced Network Technology Catalog*, 2013. https://www.eff.org/files/2014/01/06/20131230-appelbaum-nsa_ant_catalog.pdf.
- [44] Superior Court of the State of California, For the County of Santa Clara. *Xilinx, Inc. v. Flextronics International, Ltd*, December 2013. <https://www.era1.com/customuploads/20131218135319582.pdf>.
- [45] Brett Giller. *Implementing Practical Electrical Glitching Attacks*. NCC Group, November 2015. <https://www.blackhat.com/docs/eu-15/materials/eu-15-Giller-Implementing-Electrical-Glitching-Attacks.pdf>.
- [46] *Advanced Physical Inspection Methods for Counterfeit IC Detection*, volume ISTFA 2014: Conference Proceedings from the 40th International Symposium for Testing and Failure Analysis of *International Symposium for Testing and Failure Analysis*, 11 2014. doi: [10.31399/ASM.CP.ISTFA2014P0055](https://doi.org/10.31399/ASM.CP.ISTFA2014P0055).

- [47] Robb Hammond. *Counterfeit Electronic Component Detection*. AERI, September 2012. <https://www.aeri.com/counterfeit-electronic-component-detection/>.
- [48] Harold Hewett. *Methods Used In The Detection Of Counterfeit Electronic Components*. Electro-Comp Services, October 2010. https://www.circuitinsight.com/pdf/methods_detection_counterfeit_components_smta.pdf.
- [49] I. B. Leukhin, A. O. Shirin, L. N. Kessarinskiy, F. F. Taiibov, A. S. Kameneva, and A. P. Durakovskiy. Electronic Components Authentication Methods Based on Decapsulation and Chip Photo Analysis. In *2021 IEEE 32nd International Conference on Microelectronics (MIEL)*, pages 127–130, 2021. doi: [10.1109/MIEL52794.2021.9569204](https://doi.org/10.1109/MIEL52794.2021.9569204).
- [50] Mahbub Alam, Haoting Shen, Navid Asadizanjani, Mark Tehranipoor, and Domenic Forte. Impact of X-Ray Tomography on the Reliability of Integrated Circuits. *IEEE Transactions on Device and Materials Reliability*, 17(1):59–68, 2017. doi: [10.1109/TDMR.2017.2656839](https://doi.org/10.1109/TDMR.2017.2656839).
- [51] Art Ogg. *Using X-ray Systems To Detect Counterfeit And Reworked Electronic Components*. World Micro, March 2021. https://www.circuitinsight.com/pdf/xray_detect_counterfeit_smta.pdf.
- [52] Thomas D. Bergman, Cyber Program Manager, and Katie T. Liszewski. Battelle Barricade: A Nondestructive Electronic Component Authentication and Counterfeit Detection Technology. In *2016 IEEE Symposium on Technologies for Homeland Security (HST)*, pages 1–6, 2016. doi: [10.1109/THS.2016.7568901](https://doi.org/10.1109/THS.2016.7568901).
- [53] H. Huang, A. Boyer, and S. Ben Dhia. The Detection of Counterfeit Integrated Circuit by the Use of Electromagnetic Fingerprint. In *2014 International Symposium on Electromagnetic Compatibility*, pages 1118–1122, 2014. doi: [10.1109/EMCEurope.2014.6931070](https://doi.org/10.1109/EMCEurope.2014.6931070).
- [54] Andrew Stern, Ulbert Botero, Fahim Rahman, Domenic Forte, and Mark Tehranipoor. EMFORCED: EM-Based Fingerprinting Framework for Remarked and Cloned Counterfeit IC Detection Using Machine Learning Classification. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 28(2):363–375, 2020. doi: [10.1109/TVLSI.2019.2949733](https://doi.org/10.1109/TVLSI.2019.2949733).
- [55] R Hitchcock. *Radio-Frequency and Microwave Radiation*. American Industrial Hygiene Association, Fairfax, Va, 2004. ISBN: 9781931504553.

- [56] David Pozar. *Microwave Engineering, 4th Edition*. Wiley, Hoboken, NJ, 2012. ISBN: 9788126541904.
- [57] Joel Dunsmore. *Handbook of Microwave Component Measurements: With Advanced VNA Techniques*. John Wiley & Sons, Hoboken, NJ, 2020. ISBN: 9781119477129.
- [58] H. Sakoe and S. Chiba. Dynamic Programming Algorithm Optimization for Spoken Word Recognition. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 26(1):43–49, 1978. doi: [10.1109/TASSP.1978.1163055](https://doi.org/10.1109/TASSP.1978.1163055).
- [59] Keysight. *Keysight Streamline Series USB Vector Network Analyzer P937XA 2-port, up to 26.5 GHz Datasheet*, 2018. <https://www.keysight.com/us/en/assets/7018-06033/data-sheets/5992-2765.pdf>.
- [60] Atmel. *Atmel ATmega16U4 & ATmega16U4 - 8-bit Microcontroller with 16/32K bytes of ISP Flash and USB Controller - Datasheet*, 2016. https://ww1.microchip.com/downloads/en/DeviceDoc/Atmel-7766-8-bit-AVR-ATmega16U4-32U4_Datasheet.pdf.
- [61] M.K. Armstrong. PCB Design Techniques For Lowest-Cost EMC Compliance. Part 1. *Electronics & Communication Engineering Journal*, 11(4):185–194, August 1999. doi: [10.1049/ECEJ:19990402](https://doi.org/10.1049/ECEJ:19990402).
- [62] John Ardizzoni and Dennis Falls. A Practical Guide to High-Speed Printed Circuit Board Layout. *Analog Dialogue*, 39(3), September 2005. <https://www.analog.com/media/en/analog-dialogue/volume-39/number-3/articles/high-speed-printed-circuit-board-layout.pdf>.
- [63] Analog Devices. *ADRF5021 9 kHz to 30 GHz, Silicon SPDT Switch Datasheet*, 2016. <https://www.analog.com/media/en/technical-documentation/data-sheets/ADRF5021.pdf>.
- [64] Analog Devices. *ADRF5047 Silicon SP4T Switch, Reflective, 9kHz to 44GHz Datasheet*, 2019. <https://www.analog.com/media/en/technical-documentation/data-sheets/ADRF5047.pdf>.
- [65] Pat Hindle. The State of RF and Microwave Switches. *Microwave Journal*, 53(11):20–, August 2010. <https://www.microwavejournal.com/articles/10269-the-state-of-rf-and-microwave-switches>.
- [66] Isola. *370HR Laminate and Prepreg Datasheet*, 2016. <https://www.isola-group.com/wp-content/uploads/data-sheets/370hr-laminate-prepreg.pdf>.

- [67] J. Paleček, M. Vestenický, P. Vestenický, and J. Spalek. Frequency Dependence Examination of PCB Material FR4 Relative Permittivity. *IFAC Proceedings Volumes*, 46(28):90–94, 2013. doi: [10.3182/20130925-3-CZ-3023.00020](https://doi.org/10.3182/20130925-3-CZ-3023.00020).
- [68] George Milad. *New Developments in DC Acid Copper for Vertical Plating Tanks*. Uyemura International Corporation, 2007. <https://www.uyemura.com/articles/lead-free-surface-finishes.html>.
- [69] Agilent Technologies. *Limitations and Accuracies of Time and Frequency Domain Analysis of Physical Layer Devices*, 2016. <http://literature.cdn.keysight.com/litweb/pdf/5989-2421EN.pdf>.
- [70] I. J. Bahl. *Lumped Elements for RF and Microwave Circuits*. Artech House, Boston, 2003. ISBN: 9781580536615.
- [71] G.E. Ponchak, Donghoon Chun, Jong-Gwan Yook, and L.P.B. Katehi. The Use of Metal Filled Via Holes for Improving Isolation in LTCC RF and Wireless Multichip Packages. *IEEE Transactions on Advanced Packaging*, 23(1):88–99, 2000. doi: [10.1109/6040.826766](https://doi.org/10.1109/6040.826766).
- [72] Steve Hageman. *Via Spacing On High-Performance PCBs*. AspenCore, 2013. <https://www.edn.com/via-spacing-on-high-performance-pcbs/>.
- [73] Waylon Lindseth. Effectiveness of PCB Perimeter Via Fencing: Radially Propagating EMC Emissions Reduction Technique. In *2016 IEEE International Symposium on Electromagnetic Compatibility (EMC)*. IEEE, July 2016. doi: [10.1109/ISEMC.2016.7571721](https://doi.org/10.1109/ISEMC.2016.7571721).
- [74] Keysight. *Keysight N755xA Series Electronic Calibration (ECal) Module*, 2021. <https://www.keysight.com/us/en/assets/7018-05909/product-fact-sheets/5992-2547.pdf>.
- [75] Sean Gallagher. *FTDI On Counterfeit Chip Bricking: "Our intentions were honorable"*. Ars Technica, October 2014. <https://arstechnica.com/information-technology/2014/10/ftdis-anti-counterfeiting-efforts-sit-between-a-rock-and-a-hard-place/>.
- [76] Future Technology Devices International Limited. *FT232R USB UART IC Datasheet*, May 2020. https://ftdichip.com/wp-content/uploads/2020/08/DS_FT232R.pdf.

- [77] ElectronicsComp.com. *FT232RL (SMD SSOP-28 Package) USB to Serial UART Interface IC*. <https://www.electroniccomp.com/ft232rl-smd-ssop-28-package-usb-to-serial-uart-interface-ic>.
- [78] Mikhail Svarichevsky. *FTDI FT232RL: Real vs. Fake*. ZeptoBars.com, February 2014. <https://zeptobars.com/en/read/FTDI-FT232RL-real-vs-fake-supereal>.
- [79] Violet Blue. *FTDI Admits To Bricking Innocent Users' Chips in Silent Update*. ZDNet, October 2014. <https://www.zdnet.com/article/ftdi-admits-to-bricking-innocent-users-chips-in-silent-update/>.
- [80] Future Technology Devices International Limited. *FT232R - USB UART IC - Inventory from our Authorised Distributors*. <https://www.ftdichip.com/old2020/Products/ICs/FT232R.htm>.
- [81] Stephanie Yang. *Chip Shortage Brings Frustration but More Business to Industry's Middlemen*. The Wall Street Journal, June 2021. <https://www.wsj.com/articles/chip-shortage-brings-frustration-but-more-business-to-industrys-middlemen-11623582003>.
- [82] Yamaichi Electronics. *Test and Burn-in Sockets - Series Overview - SOP (Small Outline Packages, Gullwing Leads)*, September 2019. https://www.mouser.com/datasheet/2/448/YAMAICHI_TS-Katalog-2017_01_1-1523424.pdf.
- [83] Joel Dunsmore, Ning Cheng, and Yong-xun Zhang. Characterizations of Asymmetric Fixtures With a Two-Gate Approach. In *77th ARFTG Microwave Measurement Conference*, October 2011. doi: [10.1109/ARFTG77.2011.6034555](https://doi.org/10.1109/ARFTG77.2011.6034555).
- [84] Shri Agarwal, Jennifer Hill, Peter Majewicz, Gene Monroe, and John Pandolf. *Compendium Special Edition on Electrostatic Discharge (ESD)*, volume 10-2 of *EEE Parts Bulletin*. JPL & NASA, April 2019. https://nepp.nasa.gov/files/30202/EEE-Parts-Blt_n_2018-v10-Issue2-2018Jun-to-2018Sept--ESD-Compendium-URS282874-CL19-2548.pdf.

APPENDICES

Appendix A

Complete Set of Multiplexing Board RF Characteristics




Figure A.1: Insertion Loss




Figure A.2: Mean Isolation




Figure A.3: Isolation When SMA₁ to RF₁ Selected




Figure A.4: Isolation When SMA₁ to RF₂ Selected




Figure A.5: Isolation When SMA₁ to RF₃ Selected




Figure A.6: Isolation When SMA₁ to RF₄ Selected




Figure A.7: Isolation When SMA₁ to RF₅ Selected




Figure A.8: Isolation When SMA₁ to RF₆ Selected




Figure A.9: Isolation When SMA₁ to RF₁₅ Selected




Figure A.10: Isolation When SMA₁ to RF₁₆ Selected




Figure A.11: Isolation When SMA₂ to RF₇ Selected




Figure A.12: Isolation When SMA₂ to RF₈ Selected




Figure A.13: Isolation When SMA₂ to RF₉ Selected




Figure A.14: Isolation When SMA₂ to RF₁₀ Selected




Figure A.15: Isolation When SMA₂ to RF₁₁ Selected




Figure A.16: Isolation When SMA₂ to RF₁₂ Selected




Figure A.17: Isolation When SMA₂ to RF₁₃ Selected




Figure A.18: Isolation When SMA₂ to RF₁₄ Selected

Appendix B

Acquisition Controller Startup Arguments

The startup arguments supported by the Acquisition Controller are provided below and are listed in the following manner: the name of the command is first in bold, followed by a short description of the command's purpose, whether the command is required or optional, and the default value if one exists. The listing for each command also includes an example of how the command is used.

Interactive Commands

Supplying any of the following commands will override the default value used by the Acquisition Controller. In the interactive mode, the operator is prompted for a target name before each capture cycle is started. After the requested measurements are captured, the Acquisition Controller stays open and prompts the operator if they want to measure another target or exit.

-adapter Name of the adapter connected to [MB](#).

Optional, Default is NULL

Example: -adapter obd2

-adapterv Version of the adapter connected to [MB](#).

Optional, Default is NULL

Example: -adapterv 2

- dataset** Name of the dataset the measurements are a part of.
Optional, Default is NULL
Example: -dataset car_test
- db** Name of the database file to save measurements in.
Required, If the file does not already exist a new empty database will be created
Example: -db file.db
- ifbw** Comma-separated list of [IF](#) bandwidth values to sample with.
Optional, Default is 100 kHz, Minimum is 1 kHz, Maximum is 1200 kHz
Example: -ifbw 5,500,200
- notes** [JavaScript Object Notation \(JSON\)](#) formatted notes to save with the measurement.
Optional, Default is NULL
Example: -notes {'capture':'car_test'}
- points** Number of points to use in a sweep.
Optional, Default is 2500
Example: -points 1000
- rfports** Comma-separated list of [RF](#) ports to include in the sample.
Optional, Default is all [RF](#) ports
Example: -rfports 1,2,3,4,5
- rfportsf** Similar to -rfports, loads a file containing a comma-separated list of [RF](#) ports to include in the sample. Overrides value from -rfports if the switch is also present.
Optional, Default is all [RF](#) ports
Example: -rfportsf file.txt
- rfpower** Comma-separated list of power levels, in dBm, to sample at. Use 'off' as a value to sample with [RF](#) input power off. To avoid parsing errors, negative power levels should be surrounded by single quotation marks, e.g. '-5' for -5 dB
Optional, Default is -5 dB, Minimum is -43 dB, Maximum is 20 dB
Example: -rfports 5,'-10','-30',off
- samples** Number of samples to acquire.
Optional, Default is 10
Example: -samples 5

Non-Interactive Command

Supplying the following command will start the Acquisition Controller in the non-interactive mode, which captures the requested measurements under the provided target name and then automatically exits.

-target Name of the target being measured.

Required for Non-Interactive Mode

Example: -target ftdi301

Appendix C

Measurement Database Schema

The database schema used to store the SILGEO measurements is listed below. The column name is in bold, with the data type in square brackets, followed by a short description of the column contents.

DATETIME [Integer]

A UNIX formatted timestamp, in milliseconds, of the date and time the measurement was taken.

DATASET [String]

Name of the dataset the collected measurement belongs to.

TARGET [String]

Name of the target being measured.

ADAPTER [String]

Name of the adapter board connected to the [MB](#).

ADAPTER_CONFIG [String]

Version or orientation of the adapter board connected to the [MB](#).

PORT1 [Integer]

Selected [RF](#) port of the [MB](#) connected to Port 1 of the [VNA](#).

PORT2 [Integer]

Selected [RF](#) port of the [MB](#) connected to Port 2 of the [VNA](#).

POWER [Integer]

Power setting used by the [VNA](#).

IF_BANDWIDTH [Integer]

[IF](#) bandwidth setting used by the [VNA](#).

S2P [String]

Touchstone formatted measurement sweep data acquired by the [VNA](#).

SAMPLE_IN_SERIES [Integer]

Sequence number indicating the order of repeated samples.

SAMPLE_NOTES [String]

Notes or comments relevant to measurement collection.

ACQUISITION_CONTROLLER_REV [Integer]

Software version of the Acquisition Controller.

BASEBOARD_SERIAL [String]

Serial number of the [MB](#).

BASEBOARD_SOFTWARE_REV [Integer]

Software version of the firmware on the [MB](#).

VNA_SERIAL [String]

Serial number of the [VNA](#).

VNA_FIRMWARE_REV [Integer]

Software version of the firmware on the [VNA](#).

VNA_LIB_REV [Integer]

Software version of the [VNA](#) libraries running on the host computer.