

ABSTRACT

Title of Dissertation: DEGRADATION ASSESSMENT OF
MICROELECTRONIC DEVICES USING
SIDE-CHANNEL POWER MODULATION
ANALYSIS

Devon Robert Richman, Doctor of Philosophy,
2026

Dissertation directed by: Prof. Michael Pecht, Chair
Dr. Michael H. Azarian, Co-Chair

Microelectronic devices degrade over time due to exposure to stress, including temperature, voltage, and ionizing radiation. Rapid and nondestructive methods for detecting and quantifying degradation are increasingly important as our lives become more reliant on electronics, but conventional electrical testing approaches are often time-intensive or lack sensitivity to early-stage degradation.

Side-channels are inadvertent sources of information that originate from a device and contain details about its physical or electrical state. While side-channel analysis has been widely studied in the cryptography and security domains, its application to reliability and quality assessment has remained limited. This dissertation presents a side-channel approach, power modulation analysis (PMA), for the rapid detection and quantification of degradation in microelectronic devices.

A feature extraction methodology along with a suite of engineered features is developed to improve feature traceability, dataset independence, and application relevance relative to existing PCA-based approaches. The proposed method is evaluated on devices that are degraded by being subjected to high-temperature bias stress and total ionizing radiation, demonstrating the ability to detect and quantify multiple stages of degradation.

In addition, PMA is applied to the screening of counterfeit microelectronic devices. Counterfeiting microelectronic devices is an ever-growing threat to the reliability and security of electronic systems. Current standards-based testing methods are time-intensive and require numerous test methods to be performed. The results show that engineered features provide more accurate and repeatable discrimination than PCA-based approaches, enabling faster and more consistent screening. This work expands the use of side-channel analysis beyond security applications and establishes PMA as a practical tool for microelectronic reliability and quality assessment.

DEGRADATION ASSESSMENT OF MICROELECTRONIC DEVICES USING
SIDE-CHANNEL POWER MODULATION ANALYSIS

by

Devon Robert Richman

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park, in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2026

Advisory Committee:

Prof. Michael Pecht, Chair

Dr. Michael Azarian, Co-chair

Prof. Abhijit Dasgupta

Prof. Patrick McCluskey

Prof. Peter Sandborn

Prof. Ankur Srivastava, Dean's Representative

© Copyright by
Devon Robert Richman
2026

Dedication

To my amazingly supportive and loving parents, thank you for always pushing me to be the best I can be. My mother's analytical mind and organizational skills, and my father's craftsmanship and problem-solving ability, helped shape the engineer, scientist, and researcher I am today. Thank you to my younger brother, who has always been by my side, listening to me ramble even when he had no idea what I was talking about. Thank you to my best friend, who has been by my side for more than half my life. To my amazing partner, my biggest cheerleader, who supported me through thick and thin: I love you. And to all those who supported me during the highs and lows of my PhD journey, as well as those who will follow in my footsteps.

Acknowledgements

I would like to thank my advisor and co-advisor, Prof. Michael Pecht and Dr. Michael Azarian, for their guidance and support throughout my doctoral studies. Professor Pecht's direct approach and highly organized management of the research group provided structure and strengthened my professional skills as a researcher. He taught us to question even the literature from which we base our research and to think carefully and critically about both what we say and how we say it.

Dr. Azarian's mentorship style differed from Professor Pecht's yet proved equally impactful. He gave me the freedom to explore various projects until I found one that resonated with me. He pushed me to tackle new and challenging problems, from failure analysis to using specialized test equipment. My involvement in the SAE G-19A committee and work as a technical assessor would not have happened without him. His insights into science, professionalism, and life that were offered at our weekly meetings were invaluable on my journey. I would not be the researcher, writer, and professional I am today without his guidance and support.

I am grateful to the rest of my committee members: Prof. Abhijit Dasgupta, Prof. Patrick McCluskey, Prof. Peter Sandborn, and Prof. Ankur Srivastava. They provided their time, insights, and guidance throughout my graduate studies, and in some cases, dating back to my undergraduate years. I have had the pleasure of working with several committee members beyond my dissertation research, including on projects involving counterfeit microelectronic parts and failure analysis. I would also like to thank Dr. Das, Dr. Osterman, Dr. Utter, and many of my fellow graduate students for their support through my journey.

I would like to thank Joel Christiansen and Alex Robinson from Sandia National Laboratories for their collaboration on research related to the detection of radiation effects in microelectronic parts using PMA. I would also like to thank Christian Eakins and the U.S. Air Force Research Laboratory for their support during Phase I and Phase II of the small business technology transfer (STTR) project, *Quantifiable Assurance for Mission-Critical Microelectronics*. Lastly, I would like to thank Steven Dourmashkin, Connor Dempsey, and Stephen DiBartolomeo from Chiplytics for their collaboration and partnership, including the STTR efforts.

My PhD journey was relatively smooth, aside from a global pandemic and numerous experimental setbacks. In 2019, while pursuing my professional master's degree, I began debating whether to transition to the PhD program. I sought advice from friends, family, instructors, and fellow PhD colleagues. One evening, while visiting my father's business in my old 2002 Buick, I clipped a curb and lost my front right wheel cover. Later that night, annoyed but determined, I returned to the site of the incident and searched for the wheel cover. I eventually found it around 12:30 AM. I decided then that if I was stubborn enough to spend my night searching for a plastic wheel cover, then a PhD would be a perfect fit. I transferred into the doctoral program in January of 2020.

Shortly after entering the doctoral program, the COVID-19 pandemic hit, preventing me from accessing the laboratory or receiving required training for over a year. Fast forward to 2023, and I finally settled on my dissertation topic: the use of side-channels for the detection of degradation in microelectronic devices. I initially planned to graduate in the Fall of 2024, but the degradation of my test parts required

countless attempts and nearly nine months to find a suitable test. I then planned to graduate in the Fall of 2025, but processing and publishing the wealth of data collected over the course of a year, in addition to a delayed radiation study, prevented this. The completion of this PhD is a testament to my determination, persistence, and the support of those around me.

I would also like to thank my friends outside of academia, particularly my best friend, for their unwavering support throughout my graduate studies. Whether through late-night chats, simple encouragement, gaming, or D&D sessions, their presence provided balance and an outlet away from the most challenging periods of my PhD.

Finally, and most importantly, I would like to thank my family and partner. This dissertation is dedicated to you. Without your constant support, this would not have been possible. To borrow a quote, *you stood still so I could look back and see how far I have come.*

Table of Contents

Dedication	ii
Acknowledgements	iii
Table of Contents	vi
List of Tables	viii
List of Figures	xii
List of Abbreviations	xvii
Chapter 1: Introduction	1
1.1 Motivations	3
1.2 Overview of Dissertation	5
Chapter 2: Literature Review	7
2.1 Side-Channels	7
2.2 Counterfeit Microelectronic Devices and Hardware Trojans	10
2.3 Degradation of Microelectronic Devices	16
2.4 Research Gaps	22
2.5 Research Objectives	23
Chapter 3: Side-Channel Power Modulation Analysis Theory and Hardware	25
3.1 Benchtop Equipment Setup	28
3.2 Chiplytics Power Spectrum Analysis (PSA) System	32
Chapter 4: Data Processing, Feature Extraction Methodology, and Engineered Features	38
4.1 Preprocessing of PMA Data	38
4.2 Feature Engineering	39
4.2.1 Spectral Subdivision	41
4.2.2 Engineered Features	43
4.2.3 Interharmonic Removal for PCA	53
4.3 Clustering and Determining the Hyperparameters for Clustering Algorithms	54
4.4 Feature Processing and Selection	56
4.5 Semi-Supervised Feature and Outlier Detection Methods	61
4.6 Supervised Methods	63
4.7 Quantifying the Separation Between Samples Using a Single Feature	63
Chapter 5: Detection and Quantification of Degradation: High-Temperature Bias ...	67
5.1 Experimental Design	68
5.2 Assessment of PMA Repeatability and Variance	74
5.3 Results	86
5.3.1 Clustering	86
5.3.2 Semi-Supervised Methods	93
5.3.3 Conventional Electrical Test Results	97
5.3.4 Regression and Classification	101
5.3.5 Quantification of Separation	107
5.3.6 Qualitative Comparison of Resolution Modes and Shifts in PMA Spectra Over Time	111
5.3.7 Qualitative Comparison of PMA Spectra of HTB and TID Exposed Devices	113

5.4 Conclusions.....	114
Chapter 6: Detection and Quantification of Degradation: Total Ionizing Dosage ...	118
6.1 Experimental Design.....	119
6.2 Results.....	123
6.2.1 Measurement and Socketing Considerations for DMP2104V-7 MOSFETs	125
6.2.2 DMP2104V-7 MOSFET Results	129
6.2.3 TC74HC02APF NOR Logic Results	143
6.2.4 LM741CN/NOPB OpAmp Results.....	150
6.3 Conclusions.....	157
Chapter 7: Screening for Counterfeit Microelectronic Devices	160
7.1 Experimental Design.....	161
7.2 Results.....	163
7.2.1 Logic Gates	164
7.2.2 Flip-Flops.....	171
7.2.3 OpAmps	173
7.3 Hardware Variance	177
7.5 Conclusions.....	184
Chapter 8: Qualitative Study on PMA Sensitivity to Controlled Electrical Perturbations and Physical Modifications	187
8.1 Experimental Design.....	187
8.2 External Electrical Perturbation Results	189
8.3 Internal Physical Modification Results	193
8.4 Conclusions.....	197
Chapter 9: Qualitative Evaluation of PMA Voltage Ranges	199
9.1 Impedance Sweep Results.....	200
9.2 Analysis of PMA Spectra Across Multiple PMA Voltage Ranges.....	204
9.3 Difference Between Power Spectrum and Magnitude Spectrum	212
9.4 Conclusions.....	213
Chapter 10: Future Applications and Work.....	215
Chapter 11: Contributions.....	218
Appendices.....	221
Appendix A. Additional TC74HC02 HTB Study Results	222
Appendix B. Additional TID Study Results	226
B.1 Additional DMP2104V Results	226
B.2 Additional TC74HC02 Results	233
B.3 Additional LM741 Results	234
Appendix C: Additional Counterfeit Screening Study Results.....	235
C.1 Additional NOR Logic Gate Device Results	235
C.2 Additional NAND Logic Gate Device Results	236
C.3 Additional NOR, OR, and NAND Logic Gate Device Results	237
C.4 Additional Flip-Flop Device Results.....	238
C.5 Additional OpAmp Results	239
References.....	243

List of Tables

Table 1: Types of microelectronic devices tested using side-channel counterfeit detection methods.	15
Table 2: Relevant specifications of Keysight EDU33211A arbitrary waveform generator.	29
Table 3: Relevant specifications of Agilent MSO7014B oscilloscope.....	29
Table 4: Relevant specification of Chiplytics PSA system.	34
Table 5: Engineered features and the spectral representation from which they are extracted.....	44
Table 6: Summary of applied HTB stresses and relevant TC74HC02 device specifications.....	70
Table 7: PMA voltage ranges used in the HTB study for measurements performed with the Chiplytics PSA system and the BTE configuration.....	73
Table 8: Mean and standard deviation of MAE and RMSE computed from fifty PMA null spectra across different PMA voltage ranges.	76
Table 9: Mean and standard deviation of MAE and RMSE computed from fifty time-domain null spectra across different PMA voltage ranges.	77
Table 10: Regression model performance for HTB-degraded TC74HC02 devices using different input datasets, including features extracted from the PMA spectra (0–3.5 V).....	103
Table 11: Classification model performance for HTB-degraded TC74HC02 devices using different input datasets, including features extracted from the PMA spectra (0–3.5 V).....	105
Table 12: Ashman’s D and Cohen’s d computed for pairwise comparison between 297-, 520-, and 1014-hour HTB degradation stages.....	110
Table 13: PMA voltage ranges used for TID study.	121
Table 14: Classification model performance for unbiased irradiated DMP2104V devices using different input datasets, including features extracted from the PMA spectra (0–2.0 V).....	142
Table 15: Classification model performance for biased irradiated DMP2104V devices using different input datasets, including features extracted from the PMA spectra (0–2.0 V).....	142
Table 16: Classification model performance of irradiated TC74HC02 using different input datasets, including features extracted from the PMA spectra (0–1.5 V).	150
Table 17: Classification model performance for irradiated LM741 devices using different input datasets, including features extracted from the PMA spectra (0–3.5 V).	157
Table 18: Devices included in counterfeit study.....	162
Table 19: PMA voltage ranges used for counterfeit screening study.	163
Table 20: Relevant specifications of the oscilloscopes and waveform generators used in the CALCE and Chiplytics BTE setups and the Chiplytics PSA system.	179
Table 21: Clustering results for the four hardware setups.	181

Table 22: Resistor and capacitor combinations placed in series with V_{cc} for electrical perturbation study.	188
Table 23: Resistor and capacitor combinations placed in series between inputs 1A and 1B for electrical perturbation study.	188
Table 24: PMA voltage ranges used for electrical and physical modification studies.	189
Table A1: K-means and HDBSCAN clustering accuracy using PCA models fitted to PMA spectra from the HTB study, captured using the Chiplytics PSA system.	223
Table A2: K-means and HDBSCAN clustering accuracy using PCA models fitted to PMA spectra with interharmonic components excluded from the HTB study, captured using the Chiplytics PSA system.	223
Table A3: K-means and HDBSCAN clustering accuracy using engineered features extracted from the PMA spectra from the HTB study, captured using the Chiplytics PSA system.	224
Table A4: K-means and HDBSCAN clustering accuracy using PCA models fitted to PMA spectra from the HTB study, captured using the BTE setup.	224
Table A5: K-means and HDBSCAN clustering accuracy using PCA models fitted to PMA spectra with interharmonic components excluded from the HTB study, captured using the BTE setup.	224
Table A6: K-means and HDBSCAN clustering accuracy using engineered features extracted from PMA spectra from the HTB study, captured using the BTE setup. .	225
Table B1: K-means and HDBSCAN clustering accuracy using PCA models fitted to PMA spectra of unbiased DMP2104V devices from the TID study, captured at CALCE.	226
Table B2: K-means and HDBSCAN clustering accuracy using engineered features extracted from PMA spectra of unbiased DMP2104V devices from the TID study, captured at CALCE.	227
Table B3: K-means and HDBSCAN clustering accuracy using PCA models fitted to PMA spectra of biased DMP2104V devices from the TID study, captured at CALCE.	227
Table B4: K-means and HDBSCAN clustering accuracy using engineered features extracted from PMA spectra of biased DMP2104V devices from the TID study, at CALCE.	228
Table B5: K-means and HDBSCAN clustering accuracy using PCA models fitted to PMA spectra of biased and unbiased DMP2104V devices from the TID study, captured at CALCE.	228
Table B6: K-means and HDBSCAN clustering accuracy using engineered features extracted from PMA spectra of biased and unbiased DMP2104V devices from the TID study, captured at CALCE.	229
Table B7: K-means and HDBSCAN clustering accuracy using PCA models fitted to PMA spectra of unbiased DMP2104V devices from the TID study, captured at Sandia National Laboratories.	229
Table B8: K-means and HDBSCAN clustering accuracy using engineered features extracted from PMA spectra of unbiased DMP2104V devices from the TID study, captured at Sandia National Laboratories.	230

Table B9: K-means and HDBSCAN clustering accuracy using PCA models fitted to PMA spectra of biased DMP2104V devices from the TID study, captured at Sandia National Laboratories.....	230
Table B10: K-means and HDBSCAN clustering accuracy using engineered features extracted from PMA spectra of biased DMP2104V devices from the TID study, captured at Sandia National Laboratories.	231
Table B11: Classification model performance for TID-degraded biased and unbiased DMP2104V devices using different input datasets, including features extracted from the PMA spectra (0–2.0 V) captured at CALCE.	231
Table B12: Classification model performance for TID-degraded unbiased DMP2104V devices using different input datasets, including features extracted from the PMA spectra (0–2.0 V) captured at Sandia National Laboratories.....	232
Table B13: Classification model performance for TID-degraded biased DMP2104V devices using different input datasets, including features extracted from the PMA spectra (0–2.0 V) captured at Sandia National Laboratories.	232
Table B14: K-means and HDBSCAN clustering accuracy using PCA models fitted to PMA spectra of TC74HC02 devices from the TID study.....	233
Table B15: K-means and HDBSCAN clustering accuracy using engineered features extracted from PMA spectra of TC74HC02 devices from the TID study.	233
Table B16: K-means and HDBSCAN clustering accuracy using PCA models fitted to PMA spectra of LM741 devices from the TID study.	234
Table B17: K-means and HDBSCAN clustering accuracy using engineered features extracted from PMA spectra of LM741 devices from the TID study.....	234
Table C1: K-means and HDBSCAN clustering accuracy using PCA models fitted to PMA spectra of NOR logic gate devices from the counterfeit screening study.	235
Table C2: K-means and HDBSCAN clustering accuracy using PCA models fitted to PMA spectra with interharmonic components excluded of NOR logic gate devices from the counterfeit screening study.....	235
Table C3: K-means and HDBSCAN clustering accuracy using engineered features extracted from PMA spectra of NOR logic gate devices from the counterfeit screening study.....	236
Table C4: K-means and HDBSCAN clustering accuracy using PCA models fitted to PMA spectra of NAND logic gate devices from the counterfeit screening study.	236
Table C5: K-means and HDBSCAN clustering accuracy using PCA models fitted to PMA spectra with interharmonic components excluded of NAND logic gate devices from the counterfeit screening study.....	236
Table C6: K-means and HDBSCAN clustering accuracy using engineered features extracted from PMA spectra of NAND logic gate devices from the counterfeit screening study.....	237
Table C7: K-means and HDBSCAN clustering accuracy using PCA models fitted to PMA spectra of NOR, OR, and NAND logic gate devices from the counterfeit screening study.....	237
Table C8: K-means and HDBSCAN clustering accuracy using PCA models fitted to PMA spectra with interharmonic components excluded of NOR, OR, and NAND logic gate devices from the counterfeit screening study.	238

Table C9: K-means and HDBSCAN clustering accuracy using engineered features extracted from PMA spectra of NOR, OR, and NAND logic gate devices from the counterfeit screening study.	238
Table C10: K-means and HDBSCAN clustering accuracy using PCA models fitted to PMA spectra of flip-flop devices from the counterfeit screening study.	238
Table C11: K-means and HDBSCAN clustering accuracy using PCA models fitted to PMA spectra with interharmonic components excluded of flip-flop devices from the counterfeit screening study.	239
Table C12: K-means and HDBSCAN clustering accuracy using engineered features extracted from PMA spectra of flip-flop devices from the counterfeit screening study.	239
Table C13: K-means and HDBSCAN clustering accuracy using PCA models fitted to PMA spectra of 072 and 4558 OpAmp devices from the counterfeit screening study.	239
Table C14: K-means and HDBSCAN clustering accuracy using PCA models fitted to PMA spectra with interharmonic components excluded of 072 and 4558 OpAmp devices from the counterfeit screening study.	240
Table C15: K-means and HDBSCAN clustering accuracy using engineered features extracted from PMA spectra of 072 and 4558 OpAmp devices from the counterfeit screening study.	240
Table C16: K-means and HDBSCAN clustering accuracy using PCA models fitted to PMA spectra of 4558 OpAmp devices from the counterfeit screening study.	240
Table C17: K-means and HDBSCAN clustering accuracy using PCA models fitted to PMA spectra with interharmonic components excluded of 4558 OpAmp devices from the counterfeit screening study.	241
Table C18: K-means and HDBSCAN clustering accuracy using engineered features extracted from PMA spectra of 4558 OpAmp devices from the counterfeit screening study.	241
Table C19: K-means and HDBSCAN clustering accuracy using PCA models fitted to PMA spectra of 072 OpAmp devices from the counterfeit screening study.	241
Table C20: K-means and HDBSCAN clustering accuracy using PCA models fitted to PMA spectra with interharmonic components excluded of 072 OpAmp devices from the counterfeit screening study.	242
Table C21: K-means and HDBSCAN clustering accuracy using engineered features extracted from PMA spectra of 072 OpAmp devices from the counterfeit screening study.	242

List of Figures

Figure 1: Simplified diagram of PMA setup.....	26
Figure 2: PMA data collection and processing steps. The process is repeated n times and the resulting spectra are averaged.	27
Figure 3: Photograph of CALCE BTE setup, including a custom ZIF socket.	29
Figure 4: Magnitude spectrum of 1 V _{DC} signal captured using the MSO7014B, showing sub-Nyquist spurs caused by decimation and filtering.....	32
Figure 5: Image of Chiplytics PSA system.....	33
Figure 6: PMA spectrum captured using the Chiplytics PSA in high-resolution mode, showing roll-off relative to the expected harmonic decay.....	35
Figure 7: Comparison of PMA spectra captured using the high- and low-resolution modes of the Chiplytics PSA system.	37
Figure 8: Example PMA spectrum.	39
Figure 9: Example of PMA subspectra.	41
Figure 10: Example of five spectral bands of a PMA spectrum.	42
Figure 11: Example PMA spectrum (top) and the corresponding ACF and PACF computed from the spectrum (bottom).	48
Figure 12: Example PMA harmonic subspectra (top) and the corresponding normalized cross-correlation between the even and odd harmonic subspectra (bottom).....	50
Figure 13: Comparison of cepstra computed from the positive-frequency portion of a spectrum (top) and the full spectrum (bottom)	52
Figure 14: Example plot of explained variance versus principal component number for a PCA model fitted to a PMA spectrum.....	61
Figure 15: Pinout diagram of the TC74HC02APF.	68
Figure 16: PCB top (left) and bottom (right) with DIP sockets and 1 k Ω resistors used for HTB testing.	70
Figure 17: Schematic of HTB test setup.	71
Figure 18: Oven temperature (left) and approximate current supplied to each device (right) during the HTB study. Each color corresponds to a time interval after which PMA and electrical testing were performed.	72
Figure 19: Comparison of time-domain null spectra using 0–4.5 V square wave (left) and 0–5.0 V square wave (right).....	78
Figure 20: Timeline plots for MAE (dBV) of DtD stability (left) and PtP variance (right) across several PMA voltage ranges.	81
Figure 21: Timeline plot of MAE (mV) for DtD stability.....	81
Figure 22: Box-and-whisker plots of MAE (dBV) for DtD stability (left) and PtP variance (right) across nine PMA voltage ranges.	82
Figure 23: Average PMA spectra (0–3.5 V) of control devices measured on different dates.	83
Figure 24: Box-and-whisker plots of MAE (dBV) for each control device based on ten repeated measurements.	84
Figure 25: Box-and-whisker plots of MAE (dBV) of the DtD stability (left) and PtP variance (right) of the BTE setup.....	85

Figure 26: Average PMA spectra (0–3.5 V) for each degradation stage captured using the Chiplytics PSA system in low-resolution mode.	86
Figure 27: K-means clustering results using two principal components from a PCA model fitted to PMA spectra (0–3.5 V) captured using the Chiplytics PSA system. .	88
Figure 28: K-means clustering results using two principal components from a PCA model fitted to PMA spectra (0–3.5 V) including HTB control devices captured using the Chiplytics PSA system.....	92
Figure 29: MD (left) and MMD (right), trained on PMA spectra (0–3.5 V) of control devices, computed for devices at different stages of HTB degradation. Vertical dashed lines indicate the median of each degradation stage.....	94
Figure 30: LOF (left) and OCSVM (right) metrics, trained on PMA spectra (0–3.5 V) of control devices, computed for devices at different stages of HTB degradation. Vertical dashed lines indicate the median of each degradation stage.....	96
Figure 31: Curve trace measurements of TC74HC02 devices from the HTB study (left) and control devices measured throughout the HTB study (right). Shaded regions in the left plot indicate ± 1 standard deviation.....	98
Figure 32: Impedance sweep measurements using 0–3.5 V input of HTB degraded TC74HC02 devices, showing impedance (left) and phase (right). Shaded regions indicate ± 1 standard deviation.	99
Figure 33: Average difference in propagation delay and transition time between control devices and devices degraded after 1014 hours of HTB exposure.	100
Figure 34: Distribution of a single engineered feature with classes predicted using a GMM model. Vertical dashed lines indicate the median of each class.	108
Figure 35: Distribution of a single engineered feature showing GMM-predicted class distributions and class medians (left), and the same feature with medians based on ground-truth labels (right).....	109
Figure 36: PMA spectra (0–3.5 V) of 1014-hour HTB degraded devices (left) and control samples (right) captured using the Chiplytics PSA system in both low- and high-resolution modes.....	111
Figure 37: PMA spectra (0–3.5 V) of control devices captured 49 days apart.....	113
Figure 38: PMA spectra (0–3.5 V) of TC74HC02 devices after 1014 hours of HTB exposure, 30 krad of TID, 20 krad of TID, and control devices measured during both studies.	114
Figure 39: Photograph of DMP2104V soldered to a test coupon.....	125
Figure 40: Photograph of Chiplytics PSA system with external socket and BNC to ZIF adapter used to perform measurements of DMP2104V devices.....	126
Figure 41: Average PMA spectra (0–2.0 V) of DMP2104V control devices measured throughout the study.	127
Figure 42: PMA spectra (0–1.5 V) of TC74HC02 control devices (left) and PMA spectra (0–3.5 V) of LM741 control devices (right) measured throughout the study.	127
Figure 43: Mean (circles) and ± 1 standard deviation (error bars) of PMA spectra initially acquired for all device types.....	128
Figure 44: Curve trace measurements of all DMP2104V devices prior to irradiation, including control devices measured throughout the study.....	129

Figure 45: Average PMA spectra (0–2.0 V) of unbiased DMP2104V devices at different dose levels (left) and PCA-based K-means clustering results (right).	131
Figure 46: Average PMA spectra (0–2.0 V) of biased DMP2104V devices at different dose levels (left) and PCA-based K-means clustering results (right).	131
Figure 47: MD (left) and MMD (right), trained on PMA spectra (0–2.0 V) of control devices, computed for unbiased DMP2104V devices across different dose levels. Vertical dashed lines indicate the median of each dose level.	135
Figure 48: MD (left) and MMD (right), trained on PMA spectra (0–2.0 V) of control devices, computed for biased DMP2104V devices across different dose levels. Vertical dashed lines indicate the median of each dose level.	135
Figure 49: Curve trace (I_d vs. V_{ds}) measurements of unbiased (left) and biased (right) DMP2104V devices at different dose levels.	136
Figure 50: Curve trace (I_d vs. V_{ds}) measurements performed at both Sandia National Laboratories (SNL) and CALCE of unbiased (left) and biased (right) DMP2104V device. Each dose level is represented by a unique color.	137
Figure 51: Threshold voltage of DMP2104V devices for (A) unbiased, (C) biased, and (E) combined biased and unbiased conditions plotted using logarithmic scale. Corresponding linear fits excluding the 15 Mrad dose level are shown for the (B) unbiased, (D) biased, and (F) combined biased and unbiased devices.	139
Figure 52: Average PMA spectra (0–1.5 V) of TC74HC02 devices at different dose levels (left) and PCA-based K-means clustering results (right).	144
Figure 53: MD (left) and MMD (right), trained on PMA spectra (0–3.5 V) of control devices, computed for TC74HC02 devices across different dose levels. Vertical dashed lines indicate the median of each dose level.	146
Figure 54: Curve trace measurements of TC74HC02 devices at different dose levels.	147
Figure 55: Difference in the transition time (left) and propagation delay (right) between irradiated devices and unirradiated controls at different dose levels. Plotted using logarithmic scale.	149
Figure 56: Average PMA spectra (0–3.5 V) of LM741 devices at different dose levels (left) and PCA-based K-means clustering results (right).	152
Figure 57: MD (left) and MMD (right), trained on PMA spectra (0–3.5 V) of control devices, computed for LM741 devices across different dose levels. Vertical dashed lines indicate the median of each dose level.	153
Figure 58: Curve trace measurements of LM741 devices at different exposure levels (left) and the magnified curve trace from 0.7 to 1.8 V (right).	154
Figure 59: Closed-loop gain measurements of LM741 devices at different dose levels. Plotted using logarithmic scale.	155
Figure 60: K-means clustering results using a PCA model fitted to PMA spectra (0–3.5 V) of NOR logic gate devices (left), and corresponding MD (middle) and MMD (right) trained on PMA spectra (0–3.5 V) of SN74HC02 devices. Vertical dashed lines indicate the median distance for each part number.	165
Figure 61: K-means clustering results using a PCA model fitted to PMA spectra (0–3.5 V) of NAND logic gate devices (left), and corresponding MD (middle) and MMD (right) trained on PMA spectra (0–3.5 V) of SN74HC00 devices. Vertical dashed lines indicate the median distance for each part number.	168

Figure 62: K-means clustering results using a PCA model fitted to PMA spectra (0–3.5 V) of logic gate devices (left), and corresponding MD (middle) and MMD (right) trained on PMA spectra (0–3.5 V) of SN74HC02 devices. Vertical dashed lines indicate the median distance for each part number.	170
Figure 63: K-means clustering results using a PCA model fitted to PMA spectra (0–3.5 V) of flip-flop devices(left), and corresponding MD (middle) and MMD (right) trained on PMA spectra (0–3.5 V) of SN74HC574 devices. Vertical dashed lines indicate the median distance for each part number.	172
Figure 64: K-means clustering results using a PCA model fitted to PMA spectra (0–1.5 V) of OpAmps (left), and corresponding MD (middle) and MMD (right) trained on PMA spectra (0–1.5 V) of TL072IP devices. Vertical dashed lines indicate the median distance for each part number.	174
Figure 65: PMA spectra (0–3.5 V) of a TC74HC02 device measured while inserted into a breadboard and into the ZIF socket of the Chiplytics PSA system.	189
Figure 66: PMA spectra (0–3.5 V) of a TC74HC02 device with different electrical perturbations in series with V_{cc} , including (A) resistive-only changes, (B) capacitive-only changes, and (C) combined RC changes.	191
Figure 67: PMA spectra (0–3.5 V) of a TC74HC02 device with different electrical perturbations in series between inputs 1A and 1B, including resistive-only changes (left) and capacitive-only changes (right).....	192
Figure 68: PMA spectra (0–3.5 V) of SN7414A devices after a series of internal modifications: (A) sample 1, (B) sample 2, (C) sample 3, and (D) sample 4.....	194
Figure 69: Images of SN7414A dies following delidding and surface modifications: (A) sample 1 die and (C) sample 2 die after delidding, and (B) sample 1 die and (D) sample 2 die after surface scratching. Images were captured using dark-field microscopy at 100x magnification.....	194
Figure 70: Images of SN7414A sample 3 die following delidding and surface modifications: after (A) delidding, (B) after shifting the top-left wire bond, (C) after first surface scratching, and (D) after second surface scratching. Images were captured using dark-field microscopy at magnifications ranging from 25x to 100x.	196
Figure 71: PMA spectra (0–3.5 V) of an SN7417 device after a series of internal physical modifications.	197
Figure 72: Images of an SN7417 die following delidding (left) and surface scratching (right). Images were captured using dark-field microscopy at 100x magnification.	197
Figure 73: Average frequency, voltage, impedance, and phase response of TC74HC02 devices. Impedance was measured from (A) 10 kHz to 2 MHz and from (C) 1 kHz to 200 kHz, while phase was measured from (B) 10 kHz to 2 MHz and from (D) 1 kHz to 200 kHz. Voltage corresponds to the maximum amplitude (V_{max}) of the applied V_{min} – V_{max} sinusoidal input waveform.	202
Figure 74: Average frequency, voltage, impedance, and phase response of TL072IP devices. Impedance was measured from (A) 10 kHz to 2 MHz and from (C) 1 kHz to 200 kHz, while phase was measured from (B) 10 kHz to 2 MHz and from (D) 1 kHz to 200 kHz. Voltage corresponds to the maximum amplitude (V_{max}) of the applied V_{min} – V_{max} sinusoidal input waveform.	203
Figure 75: Average frequency, voltage, impedance, and phase response of RC4558IP devices. Impedance was measured from (A) 10 kHz to 2 MHz and from (C) 1 kHz to	

200 kHz, while phase was measured from (B) 10 kHz to 2 MHz and from (D) 1 kHz to 200 kHz. Voltage corresponds to the maximum amplitude ($V_{\max}$) of the applied $V_{\min}$ – $V_{\max}$ sinusoidal input waveform.	204
Figure 76: Mean (circles) and ± 1 standard deviation (error bars) of PMA spectra for DMP2104V across several PMA voltage ranges.....	206
Figure 77: Mean (circles) and ± 1 standard deviation (error bars) of PMA spectra for LM471 across several PMA voltage ranges.	207
Figure 78: Mean (circles) and ± 1 standard deviation (error bars) of PMA spectra for (A) TL072CP, (B) TL072IP, (C) RC4558, and (D) NJM4558 across several PMA voltage ranges.	208
Figure 79: PMA spectra (0–3.5 V) of TL072IP (left) and LM741 (right) obtained using different pins for the modulating input (power) and ground.....	209
Figure 80: Mean (circles) and ± 1 standard deviation (error bars) of PMA spectra for (A) CD74HC02, (B) SN74HC02, and (C) TC74HC02 across several PMA voltage ranges.	210
Figure 81: Mean (circles) and ± 1 standard deviation (error bars) of PMA spectra for SN74HC32 across several PMA voltage ranges.....	211
Figure 82: Mean (circles) and ± 1 standard deviation (error bars) of PMA for SN74HC00 (left) and TC74HC00 (right) across several PMA voltage ranges.....	211
Figure 83: Mean (circles) and ± 1 standard deviation (error bars) of PMA spectra for SN74HC574 (left) and TC74HC574 (right) across several PMA voltage ranges....	212
Figure 84: Comparison of the PMA magnitude spectrum and PMA power spectrum (0–3.5 V) (left) and the difference between the two spectra (right).	213
Figure A1: Timeline plots for RMSE (dBV) of DtD stability (left) and PtP variance (right) across several PMA voltage ranges.	222
Figure A2: Box-and-whisker plots of RMSE (dBV) for DtD stability (left) and PtP variance (right) across nine PMA voltage ranges.	222
Figure A3: Box-and-whisker plots of RMSE (dBV) for each control device based on ten repeated measurements.	222
Figure A4: Box-and-whisker plots of RMSE (dBV) of the DtD stability (left) and PtP variance (right) of the BTE setup.....	223
Figure A5: Autoencoder reconstructor error, trained on PMA spectra (0–3.5 V) of control devices, computed for devices at different stages of HTB degradation. Vertical dashed lines indicate the median of each degradation stage.	225
Figure A6: Impedance sweep measurements using 0–1.0 V input of HTB degraded TC74HC02 devices, showing impedance (left) and phase (right). Shaded region indicates ± 1 standard deviation.....	225
Figure A7: Impedance sweep measurements using 0–2.0 V input of HTB degraded TC74HC02 devices, showing impedance (left) and phase (right). Shaded region indicates ± 1 standard deviation.....	226

List of Abbreviations

ACF	Autocorrelation Function
AD3	Analog Discovery 3
AE	Autoencoder
ASIC	Application Specific Integrated Circuit
BiCMOS	Bipolar Complementary Metal-Oxide-Semiconductor
BTE	Benchtop Equipment
CMOS	Complementary Metal-Oxide-Semiconductor
COTS	Commercial Off-the-Shelf
CWT	Continuous Wavelet Transform
DtD	Day-to-Day
DIP	Dual in-Line Package
DWT	Discrete Wavelet Transform
EPROM	Erasable Programmable Read-Only Memory
ESD	Electrostatic Discharge
FFT	Fast Fourier Transform
FMMEA	Failure Modes, Mechanisms, and Effects Analysis
FPGA	Field Programmable Gate Array
HCI	Hot Carrier Injection
HTB	High-Temperature Bias
IC	Integrated Circuit
LDA	Linear Discriminant Analysis
LOF	Local Outlier Factor
MAE	Mean Absolute Error
MD	Mahalanobis Distance
MMD	Modified Mahalanobis Distance
MOSFET	Metal-Oxide-Semiconductor Field-Effect-Transistor
MSE	Mean Squared Error
NBTI	Negative Bias Temperature Instability
OCSVM	One-Class Support Vector Machine

OpAmp	Operational Amplifier
PtP	Part-to-Part
PACF	Partial Autocorrelation Function
PCA	Principal Component Analysis
PCB	Printed Circuit Board
PHM	Prognostics and Health Management
PLSR	Partial Least Squares Regression
PMA	Power Modulation Analysis
PSA	Power Spectrum Analysis
RBF	Radial Basis Function
RMSE	Root Mean Squared Error
SNL	Sandia National Laboratories
SRAM	Static Random-Access Memory
SVM	Support Vector Machine
THD	Total Harmonic Distortion
TI	Texas Instruments
TID	Total Ionizing Dosage
ZIF	Zero Insert Force

Chapter 1: Introduction

Microelectronic devices, both individual parts and assemblies, are integral to daily life and modern infrastructure. Electronics are used in everything, including mobile phones, lifesaving medical equipment, avionic sensors, and critical internet infrastructure. The degradation of electronics and their eventual failure can have a range of impacts, including possible loss of life [1]. The need to assess degradation in microelectronic devices is driven by safety and financial goals.

Current approaches to assess degradation in microelectronics are contingent on degradation reaching a level that is sufficient to cause detectable changes in the device's electrical characteristics or behavior. Degradation of devices may be driven by operating over long periods, exposure to environmental stresses, including temperature, humidity, and ionizing radiation [2], [3]. Improving the accuracy and sensitivity of conventional monitoring and detection methods is often attempted using machine learning techniques such as neural networks [1]. Different devices and forms of degradation often require tailored detection or monitoring solutions. For example, the approach used to monitor and characterize an analog device may differ from that used for digital or high-power devices [1], [4], [5]. Electrical test methods such as curve trace, parameter measurement, and impedance measurement are relatively rapid ways of obtaining basic information related to the characteristics of a device; however, these methods take longer the greater the number of pin combinations that are tested and often provide limited information. More comprehensive methods, such as functional testing, may require complex circuitry and/or specialized test equipment. New

techniques are needed to complement existing conventional test methods with improved sensitivity to subtle electrical changes, earlier detection of degradation, and differentiation of degradation stages.

Side-channel based methods provide new approaches for assessing the degradation in microelectronic devices. Side-channels are inadvertent sources of information that originate from a device and contain details about its physical state or function. All microelectronic devices exhibit numerous unintended secondary pathways through which information may leak in the form of side-channels. Common examples of side-channels include power consumption, electromagnetic emissions, and thermal signatures. Research on side-channels has been primarily focused on cryptography and hardware security, where the objective is to identify side-channel vulnerabilities and develop techniques to mitigate susceptibility to those attacks [6]. In contrast, the development of methods that use side-channels to detect and quantify degradation has been limited. Potential applications of side-channel-based methods include quantification of degradation, quality screening, defect detection, and counterfeit screening.

Counterfeit microelectronic devices pose a risk to the reliability and security of electronic systems. As counterfeiters use increasingly sophisticated techniques, existing standards-based testing methods may become less effective [7]. These methods primarily rely on visual inspection and other qualitative evaluations, which depend on the inspector's experience and subjective interpretation, increasing the risk of false negatives. Side-channel-based methods have been explored for counterfeit detection with promising results. However, many of these approaches rely on dataset-

dependent dimensionality reduction techniques or black-box models, which limit their consistency, feature traceability, and suitability for practical deployment.

This dissertation focuses on furthering the development of a side-channel technique known as power modulation analysis (PMA) for the detection and quantification of degradation in microelectronic devices, as well as for the screening of counterfeit microelectronic devices. PMA is a broad class of side-channel methods that monitor a device’s response to a modulated input. The power spectrum analysis (PSA) method, developed by Sandia National Laboratories (SNL) for counterfeit detection, is encompassed within PMA as a specific frequency-domain implementation [8]. Sandia National Laboratories described PSA as a comparative technique that relies on normalization of the captured spectra [9]. The work presented in this dissertation extends PMA beyond counterfeit screening to include the detection and quantification of degradation in microelectronic devices, while improving its effectiveness as a counterfeit screening technique. These advances include the development of a feature extraction methodology tailored to PMA along with application-relevant and dataset-independent engineered features, enabling improved feature traceability and consistency. The capability of PMA to detect and quantify degradation is demonstrated through studies involving high-temperature bias and ionizing radiation exposure. The application of PMA to counterfeit screening is demonstrated using devices with differing specifications and devices produced by different manufacturers.

1.1 Motivations

The motivation for this research is the increasing dependence of modern systems on microelectronic devices and the challenges associated with assessing their

reliability. As microelectronic devices age, they experience degradation that can lead to reduced performance and eventual failure. Detecting and quantifying degradation is essential for ensuring system reliability, particularly in safety-critical and high-reliability applications. However, existing assessment techniques typically rely on the evaluation of multiple electrical characteristics that often require degradation to progress to a level that produces detectable changes, limiting sensitivity to detect early degradation.

Side-channels provide an alternative and underutilized source of information for assessing the physical state of microelectronic devices. These unintended leakages, such as power consumption or electromagnetic emissions, originate from a device and are inherently coupled to its material properties circuit behavior, and state of health. While side-channel analysis has been extensively studied in cryptography and hardware security, its applications to reliability assessment and the detection and quantification of degradation have been limited. Expanding side-channel analysis for the degradation assessment through power modulation analysis (PMA) offers the potential for a rapid and nondestructive technique capable of capturing subtle changes associated with degradation.

Additionally, side-channel methods are being investigated for counterfeit detection. Counterfeit microelectronic devices pose a risk to system reliability and security. Current standards-based test methods typically rely on subjective evaluation or complex testing procedures that can result in false negatives. Although side-channel-based counterfeit screening methods have shown promise, they frequently depend on dataset-dependent statistical transformations such as principal component analysis

(PCA). As a result, these methods lack feature traceability and exhibit limited consistency across datasets, restricting their suitability for widespread adoption outside of controlled laboratory environments. Furthermore, these approaches rarely explore the repeatability and effect of hardware variance, which must be understood for broader adoption.

This research is motivated by the need for a fast, robust, and repeatable technique capable of quantifying degradation and screening for counterfeit microelectronic devices. The developed approach is intended to complement existing electrical test methodologies by providing rapid information-rich insight with high sensitivity to changes in the device's underlying physical condition, without replacing conventional test practices.

1.2 Overview of Dissertation

The organization of the remainder of this dissertation is as follows. Chapter 2 reviews the current and historical use of side-channels for cryptography, security, failure analysis, counterfeit and Trojan detection, and reliability. The chapter concludes with a summary of the identified research gaps and the research objectives of this dissertation. Chapter 3 describes the methodology and hardware used to perform power modulation analysis (PMA). The data processing techniques, developed feature extraction methods, and feature selection approaches are presented in Chapter 4. Chapter 5 presents the detection and quantification of degradation in complementary metal-oxide-semiconductor (CMOS) logic gate devices subjected to high-temperature bias. Chapter 6 extends this analysis to the degradation of three different devices resulting from exposure to ionizing radiation. The application of PMA to counterfeit

microelectronic device screening is presented in Chapter 7. Chapter 8 examines the effect of controlled electrical and physical perturbations on PMA spectra, while Chapter 9 presents observations regarding the voltage and frequency conditions used for PMA. Future applications and continued development of PMA are discussed in Chapter 10. Finally, the contributions of this dissertation are summarized in Chapter 11.

Chapter 2: Literature Review

Side-channels are inadvertent sources of information that originate from a device and contain details about its physical state or function. These unintended leakages commonly manifest as electromagnetic emissions, variations in power consumption, or thermal signatures. The collection and interpretation of said information is referred to as side-channel analysis.

Historically, side-channel analysis has been primarily studied in the context of cryptography and hardware security, where side-channels are exploited as attack vectors to extract secret keys or plain text information used in the encryption process. These exploitative uses are commonly referred to as side-channel attacks. However, side-channels may also be used to infer other characteristics of a device beyond cryptographic information, including indicators related to degradation, anomalous behavior, or counterfeiting.

This dissertation focuses on side-channel power modulation analysis (PMA), which is described in detail in Chapter 3. This chapter reviews the broader landscape of side-channel research, including applications in security, counterfeit and Hardware Trojan (HT) detection, and reliability and degradation assessment. The chapter concludes with a summary of the identified research gaps and the research objectives of this dissertation.

2.1 Side-Channels

Side-channel attacks can be subdivided into multiple categories based on their characteristics, including whether they are invasive or non-invasive, horizontal or

vertical, profiled or non-profiled, and active or passive [6], [10], [11], [12], [13]. An attack is considered noninvasive if it can be performed without modifying the target device, whereas invasive attacks require physical modification, such as decapsulation. Horizontal attacks exploit information leakage during a single execution of an algorithm, while vertical attacks leverage leakage collected across multiple executions. Profiled attacks require access to a similar or identical device to construct a leakage model, whereas non-profiled attacks rely on more general leakage characteristics. Active methods introduce controlled stimuli that alter a device's normal operation, while passive methods observe leakage without disturbing the intended function of the device.

Passive side-channels arise as a natural byproduct of a device's operation and are inherently linked to its functionality. While such leakages can be mitigated through shielding or design techniques, they cannot be eliminated without altering device functionality, if at all. Common examples of passive side-channels include thermal signatures, electromagnetic emissions, power consumption, mechanical vibrations, and timing variations associated with device operation. In contrast, active side-channels originate in response to externally applied stimuli that deviate from nominal inputs. Examples of active side-channel techniques include fault injection, power modulation analysis, and flash thermography.

Side-channel attacks rely on the analysis of these leakage sources to extract sensitive information, most commonly cryptographic keys. Statistical hypothesis testing methods such as Welch's t -test and the chi-squared (χ^2) test are frequently used to determine the encryption key, although a range of additional techniques, including

machine learning based approaches, have also been employed [10]. Power consumption and electromagnetic emissions appear to be the most studied side-channels in the literature.

In 1999, Kocher *et al.* introduced side-channel differential power analysis (DPA), demonstrating that cryptographic keys could be extracted by statistical analysis of many power consumption traces [14]. Prior to this work, side-channel power analysis was largely limited to simple power analysis (SPA), which involves direct interpretation of individual power consumption traces, often through visual inspection [6], [15]. Subsequent advances in side-channel analysis include correlation power analysis (CPA), template attacks, and mutual information analysis (MIA), which improve attack efficiency and robustness by exploiting statistical relationships [6], [12]. Similar methodologies have been applied using electromagnetic emission rather than power consumption. These techniques typically rely on detailed knowledge of the underlying cryptographic algorithm, such as the Data Encryption Standard (DES) or Advanced Encryption Standard (AES). Knowledge of device inputs and models of expected leakage behavior is often required as well. More recently, machine learning algorithms have been employed to enhance side-channel analysis by learning complex leakage patterns from the measured data.

While side-channel research has historically focused on cryptographic attacks, similar principles have been applied in broader contexts, including counterfeit detection and HT identification. These non-cryptographic applications extend side-channel analysis to binary degradation detection and physical characterization, which are explored in subsequent sections of this chapter.

2.2 Counterfeit Microelectronic Devices and Hardware Trojans

The incorporation of counterfeit microelectronic devices into a system poses risks to reliability, security, and safety. Counterfeit microelectronic devices may be less reliable, defective or non-functional, or fail prematurely due to prior use, incorrect materials, or the lack of rigorous testing applied to authentic devices.

Counterfeit microelectronic devices may originate from a variety of sources and exhibit distinct physical or electrical characteristics. These differences are commonly used to define counterfeit types within established taxonomies. Multiple counterfeit classification schemes have been proposed in the literature. Guin *et al.* [16], [17] developed a taxonomy that closely aligns with the seven counterfeit types defined in SAE AS6171 [18]. Remarked, recycled, and cloned parts (or devices) appear consistently across the taxonomies throughout the literature.

A cloned part is an imitation or replication of an authentic device produced outside the original manufacturing process without authorization. Remarked parts are those whose identifying markings have been altered through processes such as blacktopping, sandblasting, or chemical etching. Recycled parts are previously used devices that have been recovered and reintroduced into the supply chain as new. Because a portion of their useful life has already been consumed, recycled parts are at increased risk of early failure. Parts that have been intentionally modified for malicious purposes, such as sabotage or malfunction, are classified as tampered parts in AS6171. Hardware Trojans, which may be introduced during tampering, have been widely studied in the literature [19], [20]. The detection of tampered devices and HTs presents unique challenges compared to other counterfeit types.

State-sponsored actors allegedly used microelectronic counterfeits to infiltrate dozens of U.S. corporations by installing additional components onto commercial motherboards to enable unauthorized access to sensitive information [21]. Cases such as this underscore the threat posed by counterfeit microelectronic devices. Data on counterfeit microelectronic devices is collected through reporting systems such as the Government-Industry Data Exchange Program (GIDEP) [22] or ERAI [23]. The 2024 ERAI annual analysis documented 1055 counterfeit reports [24]. However, reporting remains incomplete, and the estimates of the economic and security impacts are outdated. A 2011 bulletin published by the National Aeronautics and Space Administration (NASA) and the Jet Propulsion Laboratory (JPL) estimated that counterfeit electronics cost the global semiconductor industry over \$200 billion annually [25]. Given this increasing complexity of modern electronics and the continued shift of manufacturing capacity to regions such as Asia [26], the present-day cost and risk associated with counterfeit microelectronic devices are likely higher. These financial figures do not fully capture the risks posed by counterfeit microelectronic devices to human safety and national security.

Counterfeit testing is required when supply chain traceability and security cannot be assured. Counterfeit testing encompasses a range of techniques used to identify anomalous features associated with the counterfeiting process. Multiple standards have been developed to define recommended practices for counterfeit detection, including SAE AS6171 [18] and IDEA 1010 [27]. Standards-based testing is the use of techniques for counterfeit detection in accordance with the requirements and best practices agreed upon by expert consensus and documented in industry

standards. Common standard-based methods include external visual inspection (EVI) [28], X-ray imaging [29], material identification using X-ray fluorescence spectroscopy [30], internal inspection following decapsulation [31], and electrical testing [32]. The adoption of standards improves consistency across the industry and reduces reliance on inspector experience, which can otherwise introduce subjectivity and variability [33].

Despite their growing acceptance, standards-based methods have limitations. Tehranipour *et al.* identified several challenges in counterfeit detection, including the detection of cloned parts, the dynamic nature of counterfeiting, and the subjective interpretation of test results [33]. The authors hypothesized that certain test methods would eventually become ineffective as countermeasures improve their processes. Richman *et al.* presented results from a 2019 study in which cloned parts passed multiple inspection techniques without exhibiting anomalies. The authors concluded that certain test methods could become less effective if or when counterfeiters improve and make fewer mistakes, such as using the incorrect part termination plating material [7]. Walter performed a series of round-robin studies and identified inconsistencies between observations made by laboratories [34].

Published counterfeit case studies remain limited. Mura *et al.* presented investigations of counterfeit amplifiers and gallium arsenide field-effect transistors [35], [36], while Wang *et al.* identified counterfeit devices as the root cause for observed failures [37]. Counterfeit multi-layer ceramic capacitors and electrolytic capacitors were reported by Shrivastava *et al.* [38], [39], while Kanovsky *et al.* studied counterfeit bipolar transistors and electrolytic capacitors [40]. Additional research on

counterfeit ceramic capacitors was presented by Azarian [41]. Xiao *et al.* reviewed the current landscape of HT detection methods, noting that HT identification often requires different techniques compared to conventional counterfeit parts [42].

While standards-based testing remains the foundation for counterfeit detection, it may be inadequate against increasingly sophisticated conventional counterfeit and tampered parts. Such testing can be time-consuming, costly, and susceptible to variability in interpretation, leading to inconsistent results across laboratories. As counterfeiters adopt more advanced techniques and produce counterfeit microelectronic devices with limited or no observable defects, there is a growing need for complementary test methods that are rapid, repeatable, and capable of detecting subtle anomalies. Side-channel analysis offers a promising approach for identifying anomalies associated with counterfeiting or HT activity.

A variety of side-channels have been used to detect recovered and remarked counterfeit parts, including RF impedance backscattering analysis [43], timing analysis [44], [45], electromagnetic emissions [46], [47], [48], and dynamic current signatures [49]. Thermography has also been explored for counterfeit detection [50], [51], [52], while S-parameter measurements have been used to detect counterfeit microcontrollers [53]. Similar side-channel approaches have been applied to HT detection, including timing analysis [54], [55], [56], electromagnetic emissions [57], [58], [59], power analysis [60], [61], and lock-in thermography [62]. DPA has also been used to detect tampering in microcontrollers [63].

Although many of these side-channel techniques remain confined to laboratory environments due to their complexity and sensitivity, there have been several attempts

to translate these methods into more accessible and field-deployable platforms. Nokomis Advanced Detection of Electronic Counterfeits (ADEC) utilizes electromagnetic emissions for screening [64], [65], while ResCav analyzes the scatter parameters of unpowered devices [66]. Side-channel power analysis approaches include Battelle’s Barricade system [67], [68], Power Fingerprinting (PFP) Cybersecurity’s SigLytics system [69], the Side-Channel Authenticity Discriminant Analysis (SICADA) method developed by the Massachusetts Institute of Technology (MIT) Lincoln Laboratory [70], and Sandia National Laboratories’ Power Spectrum Analysis (PSA) method [8].

The Battelle Barricade system is designed to perform device-level counterfeit screening and is supported by multiple patents, including automated methods for selection of optimal test frequencies and voltages [71]. Sandia National Laboratories also holds patents related to PSA for defect screening in electronic devices [72]. Christiansen *et al.* combined PSA with electromagnetic emissions to evaluate heterogeneous integrated circuits [73], while Bohman *et al.* proposed enhanced excitation strategies using multi-tone inputs to increase PSA sensitivity for counterfeit detection [74]. These systems, which share similarities with the method developed in this dissertation, have been used for discriminating between different parts, date/lot codes, and manufacturers, in addition to counterfeit screening. The Sandia PSA, PFP, and Battelle Barricade systems were evaluated for their counterfeit detection capabilities as part of a Defense Microelectronics Activity (DMEA) study. The overall accuracy was 97%, with PSA achieving 100% accuracy [75]. A summary of the types and applications studied using similar side-channel approaches is provided in Table 1.

Table 1: Types of microelectronic devices tested using side-channel counterfeit detection methods.

Purpose	Device Type
Discriminating between date/lot codes	Microcontroller[8], [68], [76], [77], [78], EPROM [67], NAND gate [79]
Discriminating between manufacturers or technologies	3-bit decoder [67], OpAmp [76], [77], temperature sensor [78], diode [76], processor [76], tantalum capacitor [76], mixed signal chip [47]
Discriminating between grades or speeds	Microcontroller [67], [68], [70], [80]
Counterfeit detection	Hex buffer/inverter [67], [68], 8-bit shift register [68], EPROM [68], FPGA [75], [76], SRAM [76], flash memory [69], differential line drivers [9], voltage regulator [75], serial configuration device [75], pulse width modular [75], NAND gate [75], OpAmp [75], driver/receiver [75], switch regulator [75]
Tampering	Microcontroller [68], FPGA [69]
Aged/Recycled	Microcontroller [68], diode [8], SRAM [76], FPGA [69], mixed signal chip [47], differential line drivers [9], ASIC [9]

Side-channel analysis has also been used to model more complex systems. In his 2017 dissertation, Copos investigated side-channel methods for modeling heterogeneous and Internet of Things (IoT) platforms, demonstrating the ability to identify operation states and analyze encrypted network traffic without prior knowledge [81]. González developed Power Fingerprinting (PFP) as a framework for assessing software integrity and cybersecurity certification, which later became the basis for commercial PFP systems [82].

Machine-learning techniques such as neural networks are often unnecessary for counterfeit detection, as the objective is anomaly identification rather than cryptographic key extraction. Nevertheless, many side-channel counterfeit detection methods rely on principal component analysis (PCA) [83] or similar dimensionality reduction techniques to visually cluster and compare devices. PCA does not explicitly model the physical relationship between side-channel measurements and device

differences. Rather, it generates latent features based on eigenvalue decomposition of the covariance matrix.

The MIT SICADA method is a partial exception, as it incorporates basic statistical features, along with golden samples, within a one-class support vector machine (SVM) framework [80]. These statistical features were extracted from Hilbert analytical signals, correlated to the clock cycle, and used for population-based discrimination [70]. The feature extraction approaches is derived from method used for RF fingerprinting [84]. While the SICADA method does extract features, these were not developed for degradation quantification and are not directly applicable to the PMA methodology, due to the difference in both the applied stimulus and monitored response. The Battelle Barricade method refers to the use of features; however, these correspond to mathematical transformations of the captured signal into other domains, from which no engineered features are explicitly extracted [85]. Feature-based approaches therefore remain uncommon in this domain. Engineered features, unlike PCA-derived components, are consistent across datasets and improve feature traceability.

2.3 Degradation of Microelectronic Devices

Microelectronic devices degrade over time as a result of environmental and operational stresses. The reliability of a microelectronic device is defined as its ability to perform its intended function under specified conditions over a given period of time. Assessing the reliability of a device requires knowledge of operating life conditions, failure mechanisms, and often accelerated testing [3]. A failure modes, mechanisms,

and effects analysis (FMMEA) may be performed alongside reliability assessments to identify high-risk failure modes and mechanisms.

Failures in microelectronic devices arise from either overstress or wear-out mechanisms. Different types of microelectronic devices may be more susceptible to failure caused by certain mechanisms. Common die-level failure mechanisms in semiconductors include hot carrier injection (HCI), negative bias temperature instability (NBTI), time-dependent dielectric breakdown (TDDB), and electromigration [2], [86]. HCI occurs when charge carriers gain sufficient kinetic energy to become trapped in the gate oxide, leading to degraded transistor performance and shifts in threshold voltage. NBTI occurs when a negative gate bias is applied in addition to operating at elevated temperature, generating interface traps that result in threshold voltage shifts. TDDB is driven by the electric field strength across the gate oxide and eventually leads to dielectric failure. Electromigration involves the transport of metal atoms due to high current density, resulting in voids or hillocks. Failure models have been developed for these and other failure mechanisms, but experimental testing is required to determine the prefactors for specific devices [86]. Advanced microelectronic devices utilizing smaller node sizes may be more susceptible to failure caused by mechanisms such as HCI, NBTI, and TDDB [87]. Package-level failure mechanisms, which may not impact electrical performance until catastrophic failure, are also well studied [88].

Exposure to ionizing radiation is also a concern for microelectronic devices, particularly in space applications. Total ionizing dosage (TID), single-event effects (SEE), and displacement damage are the primary failure mechanisms caused by

exposure to ionizing radiation [89], [90]. TID results from the accumulation of trapped charge in gate oxides, leading to threshold voltage shifts and degraded transistor performance. SEE arises from high-energy particle collisions that can cause transient errors such as flipped bits or latch-up. Displacement damage occurs when energetic particles displace atoms in the crystal lattice, causing lattice defects that degrade carrier mobility. Device packaging, shielding, technology node, and architecture influence the radiation susceptibility of a particular device. Radiation-hardened devices are specifically designed to be more resistant to radiation effects, but are often more expensive. The number of low Earth orbit (LEO) missions has increased at a near-exponential rate in the last decade [91]. The use of commercial off-the-shelf (COTS) devices has become increasingly attractive as a means of lowering costs for LEO missions. However, COTS devices must be carefully screened to ensure their reliability for the desired mission [92]. Uncertainty in existing empirical radiation models, including inconsistencies in dose and event-rate predictions, necessitates continued device-level radiation testing [93].

Failure analysis refers to the use of electrical and physical inspection techniques to determine the cause of device failure. Reliability testing often employs standardized test matrices, such as high temperature operating life (HTOL) tests, to accelerate degradation [2], [94]. Physical inspection methods used for failure analysis and reliability testing share similarities with the techniques used in standards-based counterfeit testing, including X-ray imaging, decapsulation, optical microscopy, and die inspection [95]. Some methods are nondestructive, such as optical inspection, while others, such as decapsulation, are destructive processes. Electrical characterization

techniques include DC parameter measurements using curve tracers, picoammeters, and capacitance-voltage or LCR meters [95]. More advanced electrical tests, such as those in MIL-STD-883 [96], can be performed but may require more advanced equipment and time. Reliability qualification tests are defined in MIL-STD and the JEDEC JESD22 standards [97], while radiation qualification tests are similarly outlined in MIL-STD and JEDEC standards [92].

Most conventional electrical tests focus on specification sheet parameters with defined tolerances, such as threshold voltage or propagation delay. While conventional electrical testing can be relatively rapid for simple devices such as diodes or MOSFETs with a few critical parameters, complex devices such as FPGAs and microcontrollers may require verification of hundreds of electrical and functional characteristics, making such methods slow and resource-intensive. Degradation must also be sufficiently advanced to impact these parameters for detection, necessitating multiple tests and time-intensive procedures [94].

Prognostic and health management (PHM) approaches seek to detect degradation by monitoring system variables in the field and identifying deviations from nominal behavior [1]. PHM techniques are useful for predicting maintenance needs but are generally not focused on identifying specific failure mechanisms. On-chip sensors designed to monitor individual failure mechanisms or device health have been proposed as potential solutions, though they are not widely implemented given that they require modifying the design of a device [98].

The application of side-channel methods to failure analysis and degradation assessment has shown promise. Side-channel information is known to change as

devices age. Sutherlin investigated electromagnetic emissions for counterfeit field programmable gate array (FPGA) detection. While Sutherlin achieved limited counterfeit detection accuracy, they concluded that the method could be applied to health diagnosis or failure analysis [99]. Karimi, Moos, and Moradi studied the effect of device aging due to NBTI and HCI on power consumption attacks and observed reductions in power consumption due to aging [100]. Huang, Boyer, and Dhia reported changes in a device's electromagnetic emission resulting from device modification or accelerated aging [47]. Vakil *et al.* used device timing characteristics to detect counterfeit recycled ICs by using HCI and bias temperature instability (BTI) driven degradation, achieving a model that was robust against process variations and drift; however, the device netlist was required [44]. Aging due to NBTI has also been shown to impact power analysis attacks including template attacks [101], which utilize power consumption traces, and cause observable voltage variations during AES [102]. Werner *et al.* demonstrated that RF impedance backscattering side-channels could detect recycled ICs by comparing aged and unaged FPGAs subjected to accelerated aging [43]. Side-channels have also been used to detect anomalous behavior in microelectronic devices [103], [104]. Moos *et al.* examined the influence of measurement conditions and found that temperature, supply voltage, and measurement intervals affected the information content of side-channel power consumption measurements [105]. These studies demonstrate that side-channel information changes as a result of aging, degradation, and environmental conditions, supporting the feasibility of side-channel-based degradation assessment.

Tangyunyong *et al.* analyzed PSA data from Zener diodes subjected to unbiased aging at 140 °C for 500 and 1000 hours. Clear shifts in the clustering were observed when plotting the first three principal components, while only reverse leakage current exhibited measurable changes among all conventional electrical tests performed [78]. Cross-sectioning of the diodes revealed inter-metallic growth at the copper-aluminum interface.

In a subsequent study, Zener diodes were subjected to multiple environmental and bias conditions [9]. Sequential classification on a dataset containing three different stages of degradation achieved accuracies above 90%. The PSA-based classification model outperformed a model trained on conventional curve trace data; however, the conventional model was trained and tested using approximately one-third the number of samples. Attempts at training a regression model using PSA data resulted in high RMSE values and overall poor predictive performance. Without centering the dataset, the control devices were incorrectly predicted as increasing in aging stage. Centering the data resulted in control devices consistently being predicted as the second stage, while still yielding high RMSE for the aged devices.

The same study explored using PSA for screening COTS devices subjected to aging or exposure to radiation [9]. This approach used PCA-based multivariate statistical screening to distinguish healthy, irradiated, annealed, and aged devices. The Q-residuals [106] and Hotelling's T^2 [107] were used for PCA-based screening and anomaly detection. Like other Sandia National Laboratories PSA studies, this approach included normalization to reduce variance, but the required access to control devices

for each data collection event. Classification of differential line drivers exposed to electrostatic discharge events was also performed.

While these studies demonstrate the capability of side-channel techniques to detect degradation, most stop short of quantifying the degradation severity or establishing explicit relationships between side-channel information and the physical state of the device. The primary focus of these studies has been on the impact of device degradation on side-channel attacks or the detection of counterfeit devices, rather than reliability applications. Furthermore, many approaches rely on black-box models or PCA-based analysis. The work performed by Sandia National Laboratories represents the current state of the art, utilizing PCA alongside classification models to distinguish between degradation conditions.

2.4 Research Gaps

There are numerous shortcomings present in the literature regarding side-channel methods and their application to the detection and quantification of degradation. Firstly, most side-channel research has focused on cryptographic applications, where the objective is to extract secret keys rather than to interpret the underlying physical condition of the device. As a result, there is limited emphasis on understanding why side-channel signals differ between devices or how those differences relate to degradation or failure mechanisms. When side-channels are applied to degradation detection, their use is often limited in scope to the identification of recycled devices or concerns related to increased information leakage resulting from aging. The increased interest in utilizing side-channels for counterfeit detection has focused on binary detection approaches rather than on quantifying the degree of

degradation. The work performed by Sandia National Laboratories has demonstrated a correlation between microelectronic device degradation and side-channel information; however, the primary emphasis remains on detection rather than quantification.

Second, many existing side-channel analysis methods rely on dimensionality reduction techniques like PCA. Although these approaches can be effective in certain contexts, they lack the cross-device consistency and transparency required for broader degradation analysis. The development of feature extraction methodology and a suite of generalized, dataset-independent engineered features improve the feature traceability of PMA relative to latent or black-box representations while simultaneously enabling feature-level analysis across datasets.

Lastly, prior work provides limited insight or guidance into the practical deployment of side-channel methods such as PMA, including measurement repeatability, the selection of test parameters such as PMA voltage range, and sensitivity to hardware or device changes. Without evaluating these factors, the ability to confidently assess PMA-based results and limits the ability to translate the method from laboratory demonstrations to broader reliability and screening applications.

2.5 Research Objectives

The research seeks to improve the current understanding of side-channel power modulation analysis and its correlation to the degradation of a device. Furthermore, it aims to overcome many of the shortcomings present in current methods through the following objectives.

First, apply the side-channel power modulation method for the detection and quantification of multiple degradation levels in microelectronic devices.

Second, develop a feature extraction methodology and application-relevant, dataset-independent engineered features tailored to side-channel PMA that provide improved feature traceability and performance. These engineered features are essential for quantifying the degradation of a device regardless of the input data.

Third, apply the feature extraction methodology and engineered features to counterfeit microelectronic device screening. The detection performance should be comparable to or improved relative to current side-channel approaches.

Fourth, evaluate the repeatability of the PMA method, including part-to-part variance, day-to-day stability, short-term repeatability and the effect of hardware variance. Furthermore, investigate the influence of PMA voltage range on performance and the sensitivity of PMA to controlled physical device perturbations.

Chapter 3: Side-Channel Power Modulation Analysis Theory and Hardware

Power modulation analysis (PMA) is a side-channel method performed by modulating the voltage supplied to a device, while its response is monitored. The PSA method, developed by Sandia National Laboratories, is a frequency-domain specific implementation within the broader PMA framework. In this dissertation, PMA is implemented through frequency-domain analysis of the device's voltage response. A simplified diagram of the setup for PMA is shown in Figure 1. The resulting voltage waveform, once converted into the magnitude spectrum and scaled, represents the PMA spectrum. The PMA spectrum differs based on part number, manufacturer, manufacturing date, and location. A device's specific PMA signature also changes in response to degradation.

PMA is commonly performed by modulating the voltage supplied to the device's power termination(s), while the device's ground terminations are connected to common. In devices that have two power terminations and no ground terminations, PMA is performed by supplying the modulating input between the two power terminations. More complex PMA configurations can be performed, including using a power supply to provide a bias voltage to a subset of the device's terminations.




Figure 1: Simplified diagram of PMA setup.

The modulating waveform used for PMA is most commonly a square wave. Mathematically a square wave can be represented as a Fourier series consisting of a fundamental sine wave and an infinite sum of odd harmonics of the fundamental frequency. This harmonic content allows PMA to excite a wide range of frequency components using a single waveform. The rapid transitions of the square wave induce distinct device responses that further enrich the device's PMA signature. Unlike traditional electrical test methods, the device does not need to be active or performing a specific function during the test.

The voltage waveform is monitored using an oscilloscope. A set time span is captured and converted into the frequency domain using the Fast Fourier Transform (FFT). The resulting FFT is converted into the magnitude spectrum by taking the absolute value of the FFT, which is the same as computing $\sqrt{Re(x)^2 + Im(x)^2}$ using the real and imaginary components. A spectrum analyzer can also be used in place of an oscilloscope to bypass the need to convert from the time domain to the frequency domain.

This process is then repeated n times, where n is the desired number of collections that will be averaged. The average time-domain response and magnitude

spectrum are then computed. The data collection and averaging procedure is illustrated in Figure 2. Unless otherwise noted, all results presented in this dissertation used magnitude spectra that were averaged over one hundred acquisitions. The averaged time-domain response and magnitude spectrum are then saved together in a single file. If desired, the average power and phase spectra can also be computed and recorded.




Figure 2: PMA data collection and processing steps. The process is repeated n times and the resulting spectra are averaged.

PMA can be performed at a range of voltages and frequencies. The PSA testing performed by Sandia National Laboratories commonly uses voltages less than or equal to the device's specified operating voltage [8]. PMA may be performed utilizing voltages that exceed the device's specified operating voltage, up to the absolute maximum rating. The applied voltage waveform may be unipolar or bipolar, provided that the voltage does not exceed the device specifications to ensure the method remains nondestructive.

The frequency of the input waveform is often set to 100 kHz, which has been found to be an effective starting point based on prior experimentation; however, a range of input frequencies may be used, and different frequencies can result in improved information content within the PMA spectrum. Sandia National Laboratories reported using PSA frequencies ranging from 1 kHz to 1 MHz [8].

Laboratory benchtop equipment (BTE), such as an arbitrary waveform generator and oscilloscope, can be used to perform PMA. Commercial systems

specifically designed for PMA are also available, including the Chiplytics PSA system. The research presented in this dissertation utilized both BTE and the Chiplytics PSA system, each of which offers distinct benefits and limitations. Using both hardware configurations enabled the identification of practical challenges associated with PMA, verification of experimental observations, and development of normalized PMA procedures.

The Chiplytics PSA system served as the primary hardware platform, as it enabled testing at multiple locations and provided improved measurement repeatability. While both setups were used throughout the research, results obtained using the Chiplytics PSA system are emphasized, as the measurements acquired with the BTE exhibited greater variability due to setup-dependent factors. In addition, the Chiplytics PSA system was utilized for all experiments presented in this dissertation, whereas the BTE was used selectively due to differences in hardware configuration and the focus of specific studies. Results obtained using the BTE are primarily included in the appendices and were used to support methodological development and verification rather than the primary analysis. Both hardware configurations are described in the following subsections.

3.1 Benchtop Equipment Setup

The BTE setup consists of laboratory benchtop equipment, including an arbitrary waveform generator and oscilloscope. An image of the BTE setup used at the Center for Advanced Life Cycle Engineering (CALCE) is shown in Figure 3. The BTE PMA measurements presented in this dissertation were acquired using a Keysight EDU33211A arbitrary waveform generator and an Agilent MSO7014B mixed signal

oscilloscope. The relevant specifications for the equipment are provided in Table 2 [108] and Table 3 [109]. While these instruments were used throughout this work, alternative BTE configurations have been employed in prior studies, including the use of a spectrum analyzer in place of an oscilloscope by Sandia National Laboratories.




Figure 3: Photograph of CALCE BTE setup, including a custom ZIF socket.

Table 2: Relevant specifications of Keysight EDU33211A arbitrary waveform generator.

Specification	Value
Channels	1
Max output frequency (sine wave)	20 MHz
Output range	1 mV _{pp} to 10 V _{pp} into 50 Ω 2 mV _{pp} to 20 V _{pp} into open circuit
Maximum output current (theoretical)	100 mA into 50 Ω load
Max sampling rate	250 MSa/s
Vertical resolution	16-bit (65536)
Square wave rise time	8.4 ns
Overshoot	$\leq 3\%$
Jitter	≤ 5 MHz: 2 ppm of period + 100 ps > 5 MHz: 100 ps
Output impedance	50 Ω

Table 3: Relevant specifications of Agilent MSO7014B oscilloscope.

Specification	Value
Channels	4
Input type	Single ended
Maximum input	300 V _{rms} , 400 V _{pk}
Vertical Range	2 mV/div to 5 V/div 8 divisions
Input impedance	1 M Ω or 50 Ω
Bandwidth	100 MHz
Real-time sample rate	2 GSa/s
Memory depth	8 MSa
Vertical resolution	8-bit (256)

The Keysight EDU33211A arbitrary waveform generator is capable of outputting sine waves with frequencies up to 20 MHz and square waves up to 10 MHz. At frequencies above approximately 1 MHz, the square wave exhibits noticeable edge rounding and becomes increasingly sinusoidal due to bandwidth limitations. At frequency ranges used for PMA, these limitations are not significant.

The generator can supply up to a 20 V_{pp} (peak-to-peak) when driving a high-impedance load. Based on the instrument specifications, an estimated maximum output current can be derived. Assuming a 50 Ω load, a 50 Ω output impedance, and a maximum output of 10 V (10 V_{pp} with a 5 V_{DC} offset), the output current is limited to approximately 100 mA. While lower impedance loads could draw higher current, this would exceed the recommended operating conditions and may result in power compliance limitations, which may be observed as a reduction in output voltage.

The output square wave is specified to have an overshoot of 3% or less and a rise time of 8.4 ns, assuming a 50 Ω load. In practice, both the rise time and overshoot are typically higher due to impedance mismatches within the measurement setup and higher impedance of the device under test.

The Agilent MSO7014B benchtop oscilloscope is designed for high-speed signal acquisition. Its 100 MHz bandwidth and 2 GSa/s real-time sampling rate enable accurate measurement of MHz range signals. The instrument is equipped with four single-ended input channels, each referenced to ground. The sampling rate can be increased to 400 GSa/s using equivalent-time sampling (ETS), in which repetitive signals are reconstructed through repeated acquisitions. This enhanced horizontal

resolution is not required for PMA measurements, as PMA is typically performed using frequencies of 100 kHz or lower.

A bandwidth of 100 MHz enables capture of high-frequency components, including higher-order harmonics of a square wave. Achieving this horizontal resolution, however, comes at the expense of vertical resolution. The voltage resolution is determined by dividing the full-scale voltage range by the vertical resolution. The MSO7014B has an 8-bit vertical resolution, providing 256 discrete amplitude levels. For a 5 V full-scale voltage range, this corresponds to a theoretical voltage resolution of 19.5 mV. Higher effective vertical resolution, up to 12 bits, may be achieved through averaging when the horizontal time base exceeds 20 $\mu\text{s}/\text{div}$. Displaying five periods of a 100 kHz square wave at 100 MSa/s requires a time base of 5 $\mu\text{s}/\text{div}$, while displaying fifteen periods at 33.3 MSa/s requires 15 $\mu\text{s}/\text{div}$. As a result, in practical PMA measurements using the MSO7014B, the effective vertical resolution is typically closer to 8 bits rather than the 12 bits achievable through averaging.

The MSO7014B does not use onboard anti-aliasing when decimating the captured data. The absence of anti-aliasing using techniques, such as finite impulse response (FIR) decimation filtering, can introduce sub-Nyquist spurs in the frequency spectrum. An example of these spurs, occurring at fractions of the Nyquist frequency f_n , is shown in Figure 4. The figure was acquired by capturing 5000 samples of a 1 V_{DC} signal over a 200 μs period. The expected spectrum should consist of a flat line with a peak at 0 Hz from the DC component; however, a spur is observed at half the Nyquist frequency. Increasing the sampling rate by collecting more samples over a shorter

period of time can introduce additional sub-Nyquist spurs. Up to seven spurs have been observed when collecting 20000 samples.




Figure 4: Magnitude spectrum of 1 V_{DC} signal captured using the MSO7014B, showing sub-Nyquist spurs caused by decimation and filtering.

The BTE setup offers several advantages. It can be readily adapted to accommodate a wide range of device packages and pin configurations. It can be expanded to include additional equipment such as power supplies for apply bias voltages or instrumentation for measuring additional signals, including current. These advantages, however, come with tradeoff. The flexibility of the BTE setup increases susceptibility to environmental noise, parasitic effects, and poor impedance control. The exact specifications of a BTE setup may vary due to the use of equipment from different manufacturers and models, which can reduce reproducibility and present a barrier to broader adoption of PMA. The repeatability of the measurements obtained using the BTE configuration study are briefly discussed in Chapter 5.

3.2 Chiplytics Power Spectrum Analysis (PSA) System

The Chiplytics PSA system is a commercial platform designed and manufactured by Chiplytics. The system is based on patents licensed from Sandia National Laboratories to enable commercialization of PSA. An image of the Chiplytics

PSA system is provided in Figure 5. A report published by Sandia National Laboratories assessing the commercialization of PSA compared results obtained using the Chiplytics PSA system with those produced using Sandia National Laboratories' internal PSA implementation and found that the Chiplytics PSA system produced nearly identical clustering results [110].



Figure 5: Image of Chiplytics PSA system.

The Chiplytics PSA system includes a built-in 40-pin zero-insert-force (ZIF) dual in-line socket connected to a set of relays controlled by a multiplexer. An Analog Discovery 3 (AD3) within the system is used to control the multiplexer. The system is powered by an external power supply connected via a coaxial power connector. Communication and control are achieved through a USB connection to the host computer, which can also provide power in the absence of an external power supply. Relevant specifications for the Chiplytics PSA system are provided in Table 4 [111].

The AMD Spartan 7 XC7S25 FPGA onboard the AD3 provides it with a suite of functions including waveform generation and signal acquisition. The AD3 can output sine waves with frequencies up to 20 MHz and a maximum voltage range of

10 V_{pp}. The output current is limited to approximately 30 mA into an effective output impedance of 0 Ω . Similar to the EDU33211A, square waves produced by the AD3 exhibit edge rounding at higher frequencies and become increasingly sinusoidal above approximately 1 MHz. While specifications for the square wave rise time and overshoot are not provided, they have been observed to be slower than the EDU33211A. The limited output voltage range of 10 V_{pp} is more restrictive than that of many benchtop arbitrary waveform generators; however, this range is sufficient for most modern digital microelectronic devices. Exceptions include power devices and certain analog components that require higher operating voltages. PMA may still be effective for power devices when using voltage ranges below their minimum turn-on voltage.

Table 4: Relevant specification of Chiplytics PSA system.

Specification	Value
Waveform Generator	
Channels	2
Max output frequency (sine wave)	10 MHz
Output range	Up to 10 V _{pp}
Maximum output current	30 mA DC
Max sampling rate	125 MSa/s
Vertical resolution	14-bit (16384)
Output impedance	0 Ω
Oscilloscope	
Channels	2
Input type	Differential Single ended when using BNC adapter
Maximum input	± 2.5 V ± 25 V
Vertical Range	200 μ V/div to 5 V/div 10 divisions
Input impedance	1 M Ω
Bandwidth	9 MHz 30 MHz with BNC
Real-time sample rate	125 MSa/s
Memory depth	Up to 32,768
Vertical resolution	14-bit (16384)

The AD3 oscilloscope supports a maximum sampling rate of 125 MSa/s or 100 MSa/s, depending on the boot configuration. The 100 MSa/s configuration was used for all testing performed in this dissertation. According to the manufacturer specifications, the AD3 oscilloscope has a bandwidth of 9 MHz when not using a BNC adapter and greater than 30 MHz when a BNC adapter is used. The Chiplytics PSA system does not contain a BNC adapter. A PMA spectrum acquired using the Chiplytics PSA system while measuring a 0–3.5 V square wave is shown in Figure 6, along with the predicted harmonic decay of the measured square wave. The measured PMA spectrum exhibits a gradual roll-off beginning near 20 MHz. This roll-off occurs at a frequency higher than the 9 MHz bandwidth specified without a BNC adapter, but below the bandwidth expected when using an adapter. Traditionally, oscilloscope bandwidth is characterized using sine waves; however, the approach used here provides a practical estimate of the Chiplytics PSA system’s effective bandwidth. The PMA spectrum shown in Figure 6 was obtained following the preprocessing steps, including normalization and conversion to dBV scale, as described in Chapter 4.




Figure 6: PMA spectrum captured using the Chiplytics PSA in high-resolution mode, showing roll-off relative to the expected harmonic decay.

The horizontal resolution of the AD3 is lower than that of the MSO7014B; however, this tradeoff enables higher vertical resolution. The AD3 oscilloscope has a vertical resolution of 14 bits, corresponding to 16384 discrete amplitude levels. This resolution can increase to 16 bits when averaging multiple acquisitions, and the sampling rate is less than or equal to 25 MSa/s. For example, the AD3 oscilloscope can achieve a voltage resolution of approximately 3.05 mV when operating with a 50 V full-scale range. The Chiplytics PSA system can be configured into a “high-resolution” mode in which the full-scale voltage range is limited to 5 V, increasing the expected voltage resolution to approximately 0.305 mV. These estimates align with the AD3’s specifications, which list an absolute resolution of 3.36 mV when the voltage scale is greater than 0.5 V/div and 0.336 mV when the voltage scale is less than 0.5 V/div [111].

A comparison of the PMA spectra captured using both low- and high-resolution configurations is shown in Figure 7. The voltage resolution of the Chiplytics PSA system is approximately 64 times greater than the resolution of the MSO7014B at a 5 V full-scale range. This higher voltage resolution enables detection of more subtle changes in a device’s voltage response. In addition, the AD3 employs a FIR filter for anti-aliasing during decimation, which may explain why no sub-Nyquist spurs were observed when using the Chiplytics PSA system.




Figure 7: Comparison of PMA spectra captured using the high- and low-resolution modes of the Chiplytics PSA system.

The Chiplytics PSA system offers several advantages. Its compact form factor allows the system to be easily transported and set up quickly. This enabled the PMA measurements performed during the study presented in Chapter 6 that would have otherwise not been feasible. The hardware within the Chiplytics PSA system is relatively consistent across units, contributing to improved system-to-system variance. In addition, short internal signal paths reduce parasitic effects and cable losses that may be encountered in BTE configurations. The Chiplytics PSA system is actively being improved based on research performed as part of this dissertation, including ongoing hardware and software revisions. These benefits come at the cost of flexibility, as the capabilities of the Chiplytics PSA system cannot be expanded or modified like BTE configurations.

Chapter 4: Data Processing, Feature Extraction Methodology, and Engineered Features

Once the data collection is completed, the captured magnitude spectrum needs to be processed. This chapter describes the data preprocessing steps applied to the captured spectrum to obtain the PMA spectrum, followed by the feature extraction methodology and engineered features. This methodology is designed to produce features that capture the characteristics of the PMA spectrum to be used for clustering, degradation quantification, and counterfeit screening.

4.1 Preprocessing of PMA Data

The voltage waveform is real-valued, meaning the resulting magnitude spectrum is Hermitian symmetric. Only the positive frequency portion of the spectrum is retained for further analysis. The magnitude spectrum is scaled and converted to dBV using (1). The magnitude is multiplied by a factor of two to account for the omission of the negative frequency portion of the spectrum, while the division by the number of samples N_s ensures proper amplitude normalization and maintains energy conservation in accordance with Parseval's theorem [112], [113]. The factor of two is applied to the positive frequency components except for the DC component at 0 Hz and the Nyquist frequency, which do not have corresponding negative frequency terms. The work in this dissertation does not retain the Nyquist frequency. There are multiple approaches to scaling frequency spectra, including the use of the unscaled magnitude spectrum. The choice of scaling does not affect the relative behavior of the later extracted

engineered features, but it does affect their absolute values. Thus, a consistent approach to scaling is required to ensure meaningful comparisons across measurements.

$$20\log_{10}\left(|FFT| \times \frac{2}{N_s}\right) \quad (1)$$

The scaled magnitude spectrum represents the device's PMA spectrum, from which the engineered features are extracted. An example PMA spectrum for a TC74HC02 NOR logic gate device is presented in Figure 8. A common processing approach is to impose a noise floor, in which all points below the desired noise floor are set equal to the noise floor. This approach should be avoided as it can be arbitrary and can result in the removal of meaningful information. Instead, the feature extraction method developed divides the PMA spectrum into multiple components from which features are extracted. This eliminates the need to impose a noise floor and retains potentially meaningful information.




Figure 8: Example PMA spectrum.

4.2 Feature Engineering

Principal component analysis (PCA) is a commonly used dimensionality reduction technique that is sometimes employed in place of engineered features. PCA

transforms the input data into a new set of variables called principal components, which are linear combinations of the original inputs. These components are ordered by the amount of variance they explain in the dataset. Because this transformation is data-dependent, the resulting components vary with the training data and often require refitting when applied to new datasets. PCA maximizes variance and assumes linear relationships, which may not align with the dataset characteristics most sensitive to the changes of interest. As a result, the extracted components lack application relevance and feature traceability, and vary depending on the composition and variance structure of the dataset. For example, changes in preprocessing steps, such as assigning a noise floor or removing portions of the PMA spectrum, can alter the variance distribution and lead to different clustering outcomes when PCA is used. Similar limitations apply to other linear and nonlinear dimensionality reduction techniques [114].

In contrast, the presented feature extraction methodology and engineered features are derived from domain knowledge and retain consistent meaning across datasets. These features provide a traceable representation of the PMA spectrum, enabling improved detection and quantification of degradation in microelectronic devices.

PCA was used as an exploratory tool to assess the relative importance of different components within the PMA spectrum. Although PCA does not provide a consistent physical interpretation across datasets, examining the component loadings provides insights into which frequency bins influence the principal components. This assessment identified even and odd harmonic components in the PMA spectrum as dominant contributors, while interharmonic frequency components had minimal

influence. This led to the use of spectral subdivisions for improved feature extraction, as explained in the following subsections.

4.2.1 Spectral Subdivision

The PMA spectrum, as shown in Figure 8, can be divided into subspectra and spectral bands. The subspectra consist of harmonic and interharmonic components. For feature extraction, the DC component at 0 Hz is excluded from the interharmonics and treated as an individual feature. The harmonic subspectrum can be further subdivided into the even and odd harmonic subspectra. An example PMA spectrum divided into subspectra is shown in Figure 9, where the interharmonic components are shown in purple, the odd harmonic components in orange, and the even harmonic components in green. This spectral subdivision enables improved feature extraction by isolating harmonic content from interharmonic components.




Figure 9: Example of PMA subspectra.

Spectral bands are adjacent portions of a spectrum, each defined by an upper and lower frequency bound. These bands are evenly spaced across the spectrum's frequency range. Typically, three to five spectral bands are used, corresponding to

divisions representing 33%, 25%, or 20% of the spectrum. Unless otherwise noted, this dissertation used five spectral bands. An example of a PMA spectrum divided into five spectral bands is shown in Figure 10. In addition to the PMA spectrum, the subspectra can also be divided into spectral bands. However, spectral bands are only extracted from the even and odd harmonic subspectra. Dividing the harmonic subspectrum or the interharmonic subspectrum into bands did not provide additional information, as the features were redundant.




Figure 10: Example of five spectral bands of a PMA spectrum.

Extraction and analysis of even and odd harmonic components from a spectrum are common in multiple domains, including vibration, acoustic, and power analysis. Even harmonics originate from waveform asymmetry. For signals exhibiting half-wave symmetry, such as an ideal square wave, there will be no even harmonics. Odd harmonics, by contrast, are associated with the repetitive structure of a waveform and the extent to which it deviates from a pure sinusoid. As a result, analysis of the harmonic components is often limited to features such as the total harmonic distortion (THD), harmonic centroid, harmonic energy ratio, and harmonic spread [115], [116].

In power systems, harmonics are a concern as their presence can negatively impact power quality and cause damage to machinery [115]. In vibration analysis, harmonic components are used to identify faults, such as those in bearings [117], [118]. In acoustic analysis, harmonic structure is closely related to timbre, which encompasses characteristics such as loudness, pitch, and duration [119]. The analysis of the PMA spectrum shares many similarities with these established approaches in speech and audio analysis domains. While many features in this dissertation are inspired by these established fields, as well as RF fingerprinting [84], the systematic extraction of engineered features from the odd harmonic, even harmonic, and interharmonic subspectra appears to be novel, particularly in the context of side-channel analysis, where most prior work relies on PCA.

4.2.2 Engineered Features

Existing feature extraction tools such as the time series feature extraction (tsfresh) [120] and time series feature extraction (tsfel) [121] Python libraries can be used to extract a large number of features from time-domain datasets. Although these tools are designed for time series analysis, tsfresh and tsfel were applied to PMA spectra to identify candidate feature categories. The identified feature categories were further explored for inclusion in the set of engineered features. In addition to analysis of the PMA spectrum, the voltage waveform, frequency spectrum, and power spectrum were examined using multiple approaches; however, these were found to be redundant with the features extracted from the PMA spectrum.

The engineered features can be categorized based on shared characteristics or the process used to extract them. A complete list of the engineered features and the

spectra they are extracted from is provided in Table 5. Most of the engineered features are extracted from the dBV scaled spectra unless the feature contains the term “linear”, which indicates the feature is extracted from the linear scaled magnitude spectrum.

Zytek *et al.* discussed the qualities of interpretable features and other important feature characteristics including trackability, meaningfulness, simulatability, and understandability, while noting that interpretability is relative to both the domain and end application [122]. In this dissertation, the feature extraction methodology and engineered features are application-relevant, dataset-independent, and traceable. Traceability encompasses the ability to consistently link feature values to the underlying PMA spectra and to reproduce feature values using fixed, dataset-agnostic mathematical formulations. Traceability combines the notions of trackability and simulatability described by Zytek *et al.* The developed features are tailored to the PMA spectrum and therefore encode domain-specific information, enabling consistent feature computation across datasets and measurement events.

Table 5: Engineered features and the spectral representation from which they are extracted.

Feature Category	Feature	PMA	Harm. ¹	Odd Harm.	Even Harm.	Inter. Harm.	Spectral Bands
Statistical	Mean	✓	✓	✓	✓	✓	✓
	Standard deviation	✓	✓	✓	✓	✓	✓
	Skewness	✓	✓	✓	✓	✓	✓
	Kurtosis	✓	✓	✓	✓	✓	✓
	Dynamic range	✓	✓	✓	✓	✓	✓
	Peak-to-trough index	✓	✓	✓	✓	✓	✓
Spectral descriptors	Spectral centroid	✓	✓	✓	✓	✓	✓
	Spectral spread	✓	✓	✓	✓	✓	✓
	Spectral skewness	✓	✓	✓	✓	✓	✓

¹ “Harm.” denotes harmonic

Feature Category	Feature	PMA	Harm. ¹	Odd Harm.	Even Harm.	Inter. Harm.	Spectral Bands
	Spectral kurtosis	✓	✓	✓	✓	✓	✓
	Spectral crest factor	✓	✓	✓	✓	✓	✓
	Spectral decrease	✓	✓	✓	✓	✓	✓
	Linear spectral flatness	✓	✓	✓	✓	✓	✓
	Spectral ratio	-	✓	✓	✓	✓	✓
	Odd-to-Even index	-	-	✓	✓	-	-
	Linear energy	✓	✓	✓	✓	✓	✓
	Sum of squares	✓	✓	✓	✓	✓	✓
	Fundamental frequency magnitude	✓	-	-	-	-	-
	DC component magnitude	✓	-	-	-	-	-
	Cumulative sum of squares frequency: 25%, 50%, 75%	✓	-	-	-	-	-
	Total harmonic distortion (THD)	✓	-	-	-	-	-
	1 st order polynomial slope	✓	✓	✓	✓	✓	-
	Log-frequency slope	✓	✓	✓	✓	-	-
	2 nd order polynomial fit	✓	✓	✓	✓	✓	-
	3 rd order polynomial fit	✓	✓	✓	✓	✓	✓
	3 rd order polynomial residual sum of squares	✓	✓	✓	✓	✓	-
	3 rd order polynomial R ²	✓	✓	✓	✓	✓	-
	3 rd order polynomial residual statistics	✓	✓	✓	✓	✓	-
	Harmonic decay model residual RMS	-	-	✓	-	-	-
	Harmonic decay model residual error	-	-	✓	-	-	-
ACF	Mean	✓	✓	✓	✓	✓	-
	Standard deviation	✓	✓	✓	✓	✓	-

Feature Category	Feature	PMA	Harm. ¹	Odd Harm.	Even Harm.	Inter. Harm.	Spectral Bands
	Skewness	✓	✓	✓	✓	✓	-
	Kurtosis	✓	✓	✓	✓	✓	-
	Sum of squares	✓	✓	✓	✓	✓	✓
PACF	Mean	✓	✓	✓	✓	✓	-
	Standard deviation	✓	✓	✓	✓	✓	-
	Skewness	✓	✓	✓	✓	✓	-
	Kurtosis	✓	✓	✓	✓	✓	-
	Sum of squares	✓	✓	✓	✓	✓	-
Harmonic cross-correlation	Mean	-	-	✓	✓	-	-
	Max	-	-	✓	✓	-	-
	Sum of squares	-	-	✓	✓	-	-
CWT	Sum of squares	✓	-	-	-	-	-
Cepstral Coefficients	Sum of squares	✓	-	-	-	-	-
Derivative	Mean of 1 st and 2 nd derivative	✓	✓	✓	✓	✓	✓
	Standard deviation of 1 st and 2 nd derivative	✓	✓	✓	✓	✓	✓
	Mean spectral roughness	✓	✓	✓	✓	✓	✓

Statistical features include the mean, standard deviation, skewness, kurtosis, and dynamic range. The skewness (non-spectral weighted) captures the asymmetry of the magnitude distribution, while the kurtosis quantifies the peakedness of the distribution. In most cases, the spectral skewness and kurtosis provided more meaningful information. The peak-to-trough index is defined as the maximum magnitude divided by the minimum magnitude of the PMA spectrum.

Spectral descriptors are features that characterize the shape, distribution, and structure of the PMA spectrum. These include weighted statistical measures such as the centroid, spread, skewness, and kurtosis. The spectral centroid represents the weighted mean frequency and corresponds to the center of mass of the spectrum, while the spectral spread describes the dispersion of energy about this center. Spectral skewness

captures the asymmetry of the spectral energy distribution, and spectral kurtosis quantifies whether the energy is concentrated within a narrow region or broadly distributed.

Additional descriptors include spectral decrease, which measures how quickly the spectrum decays relative to the fundamental component, and spectral flatness, also referred to as Wiener entropy or tonality coefficient, is computed as the ratio of the geometric mean to the arithmetic mean of the linear magnitudes. The crest factor is defined as the squared peak magnitude divided by the root mean squared (RMS) value of the spectrum.

Linear energy is computed as the sum of squared linear magnitudes, while an analogous sum of squares is computed from the dBV-scaled PMA spectrum. Spectral indices are the ratio of the sum of squares of one spectrum to another, including the odd-to-even index, which is similar to features used in audio and speech analysis [119], [123]. THD is computed using every harmonic component of the PMA spectrum, whereas THD is often limited to a fixed number of harmonics in other applications [124]. The cumulative sum of squares is defined as the frequency at which a specified percentage of the total sum of squares is reached. Many of these spectral descriptors are inspired by established features used in audio [119], [123], [125] and speech analysis [126].

Shape-based features are extracted by fitting the first, second, and third-order polynomials to the PMA spectrum, subspectra, and spectral bands. Each polynomial coefficient, except for the offset term, is treated as an independent feature. In addition to using polynomial coefficients directly, the third-order polynomial fit is compared to

the original spectrum or subspectra to compute residuals. Statistical measures of these residuals, including the sum of squares and coefficient of determination (R^2) are extracted as features. Because the PMA input waveform is a square wave, the linear odd harmonic magnitudes are expected to decay approximately as $\frac{1}{n_{harmonic}}$. This expectation can be used to construct an ideal decay curve using the fundamental component from the linear PMA spectrum, which is then converted to dBV and compared with the odd harmonic subspectrum. The RMS of the residuals and the relative error are computed as features that quantify the degree of deviation from the expected harmonic decay. Similar harmonic deviation metrics have been used in other domains [116].

The autocorrelation function (ACF) measures the correlation of a signal with a delayed version of itself [112]. The partial autocorrelation function (PACF) is similar to the ACF but isolates the contribution of each individual lag by removing the effects of all preceding lags [127], as illustrated in Figure 11. The ACF and PACF are used to identify repetitive structures in the spectrum, such as harmonic spacing.




Figure 11: Example PMA spectrum (top) and the corresponding ACF and PACF computed from the spectrum (bottom).

The number of lags analyzed for the ACF and PACF was selected based on $10 \log_{10} \left(\frac{N_s}{2} \right)$, where N_s is the number of spectral samples. In practice, the first 20 to 50 lags are commonly analyzed, as including additional lags often yields diminishing returns for feature extraction while increasing computational overhead, particularly for the PACF. The optimal number of lags can vary depending on the application [127].

Features extracted from the ACF and PACF include statistical measures such as the mean, standard deviation, skewness, and kurtosis. In addition, the sum of squares of the ACF and PACF is computed. While features derived from the ACF and PACF may be strongly correlated, including both provides a more comprehensive representation of repetitive and harmonic structures within the spectrum.

The cross-correlation [126] between the even and odd harmonic subspectra provides insight into their similarity and shared repetitive structure. Subtle changes in a device's PMA signature may manifest differently within the odd and even harmonic components, and cross-correlation captures how these components relate as a function of frequency shift. Summary features, including the mean, maximum correlation, and sum of square are extracted from the harmonic cross-correlation. An example of the harmonic cross-correlation is shown in Figure 12.




Figure 12: Example PMA harmonic subspectra (top) and the corresponding normalized cross-correlation between the even and odd harmonic subspectra (bottom).

Unlike the Fourier transform, which maps a signal into the frequency domain, wavelet analysis represents a signal in the time-frequency domain, providing information about both spectral content and its localization [128]. The discrete wavelet transform (DWT) decomposes a signal into a finite number of levels using discrete scales and translations of a selected mother wavelet, with each corresponding to a specific frequency band. Guidance for selecting the number of decomposition levels and the ideal mother wavelet is well established in the literature. In contrast, the continuous wavelet transform (CWT) uses a continuous range of scales and translations, resulting in a highly resolved time-frequency representation.

Applying wavelet analysis directly to frequency-domain signals is not typical. This approach is conceptually related to cepstral analysis, although it is not directly equivalent. Empirical evaluation of both the DWT and CWT for PMA spectra indicated that the CWT provided more informative feature extraction, whereas the multi-level decomposition of the DWT produced an excessive number of highly correlated features across the spectrum, subspectra, and spectral bands. The Ricker (Mexican hat) wavelet

was selected based on the evaluation of several wavelets and its widespread use in CWT applications. Among the CWT-based features examined, the sum of squares of the wavelet coefficients was found to be one of the most impactful.

A cepstrum is a signal representation obtained by applying the FFT to the logarithm of a spectrum, often referred to as the spectrum of a spectrum [129], [130]. The cepstrum allows for analysis of periodic structures within the spectral domain. Two primary forms of the cepstrum are commonly used: the power cepstrum and the complex cepstrum [130]. The real cepstrum is derived from the complex cepstrum by discarding the phase information [117]. Cepstral analysis, the analysis of a cepstrum, is typically performed using the complex spectrum, including both positive and negative frequency components.

For analysis of the PMA spectra in this dissertation, the real cepstrum is computed using the positive frequency portion of the magnitude spectrum, excluding the Nyquist frequency. Figure 13 illustrates a comparison between the real cepstrum obtained from the complete magnitude spectrum and the real cepstrum obtained using only the positive frequency portion. The two cepstra appear identical, although the x-axis is different. The x-axis of the cepstrum, referred to as quefrency, has units of time and corresponds to periodic spacing in the frequency domain, analogous to lags in ACF. The sum of squares of the cepstral coefficients is extracted as a feature to capture repetitive spectral structure, similar to the features extracted from the ACF and PACF.

The mel-frequency cepstrum is widely used in audio and speech analysis [125], [131]. It was not considered for PMA feature extraction, as the Mel scale is designed to model the frequency perception of the human ear.




Figure 13: Comparison of cepstra computed from the positive-frequency portion of a spectrum (top) and the full spectrum (bottom)

Derivative-based features are extracted from the first and second derivatives of the PMA spectrum, subspectra, and spectral bands. The mean and standard deviation of the derivatives are computed for both the first and second order derivatives. Spectral roughness, which characterizes variability within a spectrum, is estimated using the mean absolute spectral slope, defined as the average of the absolute value of the first derivative.

Collectively, the presented feature extraction methodology and engineered features provide a comprehensive characterization of a device's PMA spectrum by capturing its structure, harmonic behavior, energy distribution, and stability. The features are designed to capture both global trends and local deviations within the PMA spectrum, balancing generalization with sensitivity to subtle changes. Statistical features and spectral descriptors characterize the overall shape and distribution of the

spectrum, while the harmonic decay and derivative-based features capture distortion and variability. Sensitivity to periodic behavior and localized structure is achieved by including correlation-, wavelet-, and cepstrum-based features. Together, these features provide application-relevant information that corresponds directly to known spectral characteristics and are computed using consistent mathematical formulations.

Additional engineered features and extraction methods could be explored, including time-domain filtering, wavelet analysis of the time-domain response, Hilbert-space transformations, and analysis of the phase spectrum. While such approaches may provide complementary information and additional sensitivity, they also increase methodological complexity beyond an approach that is shown to be effective in this dissertation.

4.2.3 Interharmonic Removal for PCA

The interharmonic components within the PMA spectrum are often a major source of low-information variance and noise. Removing these interharmonic components can increase the effective signal-to-noise ratio and reduce the chances that PCA is dominated by variance that does not meaningfully separate classes within the dataset. This also reduces the dimensionality of the PMA spectrum, which can improve computation efficiency for PCA and other machine-learning methods.

The feature extraction framework presented in this dissertation includes interharmonic-based features to preserve potentially relevant information, while mitigating the influence of interharmonic noise by analyzing the subspectra and spectral bands separately. However, in cases where engineered features are not

employed, removing the interharmonic components can improve the accuracy and confidence of PCA-based approaches.

Care must be taken when removing interharmonic components from PMA spectra acquired at lower sampling rates. At lower sampling rates, high-order harmonic components at frequencies above the Nyquist frequency may alias into lower-frequency interharmonic regions. Such aliasing is clearly visible within the PMA spectrum as evenly spaced periodic peaks at interharmonic frequencies. This behavior was observed when using a sampling rate of 33.3 MSa/s, but not at 100 MSa/s. In these cases, indiscriminate removal of interharmonic components can result in the loss of useful information.

4.3 Clustering and Determining the Hyperparameters for Clustering Algorithms

Clustering, which is predominantly an unsupervised learning technique, is the process of grouping data points based on similar characteristics [132]. Some clustering approaches require the number of clusters to be specified, such as K-means clustering [133] and spectral clustering [134]. Other methods determine clusters based on a distance threshold or neighborhood size, including Density-Based Spatial Clustering of Applications of Noise (DBSCAN) [135]. Hierarchical clustering methods, such as agglomerative clustering [136], iteratively merge data points or clusters, while hierarchical DBSCAN (HDBSCAN) [137], [138] constructs a hierarchy of clusters based on density variations. Affinity propagation is another clustering approach that identifies representative examples in the dataset by passing messages between data points [139]. Gaussian mixture models (GMM) are probabilistic, model-based

clustering methods that represent the data as a mixture of normal distributions, with each cluster characterized by its own mean and covariance.

Each clustering algorithm has inherent advantages and limitations. For example, K-means clustering is computationally efficient and scalable to large datasets, but requires the number of clusters to be specified and performs best when clusters are approximately spherical, well separated, and similar in size. In contrast, affinity propagation does not require the number of clusters to be predefined and instead relies on a damping factor to control convergence; however, it does not scale well to large datasets. The suitability of a clustering algorithm for a given application depends on several factors, including the shape, distribution, size, and noise characteristics of the dataset. As a result, no clustering algorithm is universally optimal, and algorithm selection often involves tradeoffs.

Numerous techniques have been developed to estimate the number of clusters, particularly for K-means clustering, including the elbow criterion and gap statistic [140], [141]. The elbow method is widely used but is inherently subjective, as it relies on visual inspection and can yield ambiguous results, making it less rigorous than mathematically defined approaches. Attempts have been made to automate the elbow method; however, results are often mixed and similarly inconsistent. The gap statistic selects the number of clusters by comparing the clustering performance of the observed data to that of a reference dataset generated from a uniform distribution.

Other approaches determine the number of clusters by evaluating clustering quality using validation metrics [140]. These clustering metrics are either extrinsic or intrinsic. Extrinsic metrics require knowledge of the ground-truth labels, which are

often unavailable in practice. When the ground truth is known, the number of clusters is usually also known. Intrinsic metrics assess the structure of the clustering without labels. Common intrinsic clustering metrics include the silhouette score [142], Davies–Bouldin index [143], Calinski–Harabasz index [144], and Dunn index [145]. The silhouette score is bounded between negative and positive one, where higher values indicate better separation and cohesion. Higher values of the Dunn and Calinski–Harabasz indices correspond to improved clustering quality, whereas lower Davies–Bouldin index values indicate better cluster separation. Similar metric-based approaches are also used to guide hyperparameter selection for density-based clustering methods such as HDBSCAN. Moulavi *et al.* developed a density-based clustering validation framework for determining HDBSCAN hyperparameters [146].

For GMM, the Akaike Information Criterion (AIC) and the Bayesian Information Criterion (BIC) are commonly used to select the number of components. AIC estimates the relative information loss of a model, with lower values indicating a better fit, while BIC penalizes model complexity and favors simpler models [147], [148]. However, when comparing different preprocessing methods or PMA voltage ranges, using the known or expected number of clusters provides a more consistent basis for comparison.

4.4 Feature Processing and Selection

Extraction of engineered features, such as those defined in this dissertation, can result in redundant or uninformative features. This is particularly true when similar features are extracted across multiple subspectra and spectral bands. Feature selection refers to the process of identifying the most relevant features for a given analysis.

Numerous feature selection approaches have been developed, including supervised and unsupervised techniques [149]. Supervised approaches require knowledge of the ground-truth labels, which are often unavailable for counterfeit detection and reliability applications. Consequently, this work focuses on unsupervised feature processing and selection methods.

Normalization of features is often required prior to applying statistical or machine learning methods, as extracted features can vary in scale and variance. Z-score normalization standardizes each feature to have a mean of zero and a standard deviation of one, improving comparability across features. While normalization removes absolute scale information and, by extension, physical meaning, it is useful for feature selection. In cases where the unscaled magnitude of features is important for interpretation or modeling, the raw feature values can be retained after feature selection is performed using normalized features.

Redundant features are removed using a correlation-based pruning approach, which constitutes an initial, unsupervised form of feature selection. The Spearman rank correlation is used, as the engineered features include both linear and nonlinear relationships and may originate from different scales. The absolute pairwise Spearman correlation coefficient is computed for all features, and for each pair with a correlation exceeding 0.95, one feature is removed. The remaining set of selected features is referred to as the FS-SCC features. This approach is consistent with multivariate statistical filtering methods described by Dwivedi *et al.* [150].

An unsupervised feature selection method was developed based on PCA. PCA is used exclusively to rank and evaluate the importance of the extracted engineered

features. Although this approach is referred to as PCA-based feature selection, it relies on the eigenvalue decomposition of the feature covariance matrix solely for feature scoring, rather than for feature extraction or dimensionality reduction that produces latent variables.

After FS-SCC, each remaining feature is scored using the sum of squared PCA loadings, weighted by the explained variance of the retained principal components. This approach is conceptually similar to spectral filter methods, such as the Laplacian score-based approach outlined by Dwivedi *et al.* [150]. The PCA-based feature selection accounts for the covariance and global relationship between features while retaining the benefits of the engineered features.

In this dissertation, the first two principal components were retained, as the inclusion of additional components was found to reduce the performance of downstream machine learning models. Based on this feature ranking, the top 50 features were retained and referred to as the FS-PCA feature set. The number of retained features was selected based on empirical observations across multiple studies to provide a consistent basis for comparison. The number of retained features may be treated as a hyperparameter that can be optimized for specific applications.

A second feature selection approach was investigated based on evaluating clustering performance using engineered features. One-dimensional clustering is first performed using individual features and evaluated using intrinsic clustering metrics such as the silhouette score, Calinski–Harabasz index, Davies–Bouldin score, or Dunn index. Similar feature evaluation approaches using the silhouette score and Calinski–Harabasz index have been outlined by Dwivedi *et al* [150].

Following feature ranking, the number of features to retain can be estimated using a clustering-based evaluation. Clustering, using a selected algorithm, is performed repeatedly while incrementally increasing the number of top-ranked features included. Each clustering result is evaluated using an intrinsic clustering metric. Among the metrics evaluated, the Dunn index was empirically observed to provide a consistent approximation of the true clustering separation. This approach is conceptually similar to clustering-driven wrapper methods described by Dwivedi *et al.* [150] and can also be applied to PCA-ranked features.

Although both the clustering-based feature ranking and the clustering-based method for determining the number of features to retain were explored, neither approach was used in the final analysis presented in this dissertation. These methods are included for completeness and may be useful for future refinement of feature selection strategies for PMA.

Unsupervised feature selection can be challenging, as common methods may not perform well for a given feature set, often necessitating the use of multiple feature selection approaches or the development of methods tailored to the application. While this dissertation primarily relies on unsupervised feature selection, a supervised feature-ranking approach based on linear discriminant analysis (LDA) was also explored to represent a more tailored approach to feature selection [151]. LDA is a supervised dimensionality reduction method that determines linear combinations of features that maximize separation between classes while minimizing within-class scatter.

Similar to the PCA-based feature ranking approach discussed, LDA can be used to rank features based on their contribution weighted by the discriminative strength of the components. As with the FS-PCA method, this method is performed using FS-SCC features. The number of retained features is selected by repeatedly performing clustering while incrementally increasing the number of top-ranked features to identify the maximum clustering accuracy using the fewest features, referred to as the FS-LDA features. However, because LDA relies on class labels, this approach is susceptible to overfitting and may retain features that do not generalize beyond the labeled dataset.

While this dissertation focuses on the use of engineered features, PCA-based analysis is included as a baseline for comparison. To ensure consistency across studies, either a fixed number of principal components must be used, or a systematic method for selecting the number of principal components retained must be defined. Unlike engineered features, principal components are ordered by explained variance. A common approach to selecting the number of components to retain is to plot the explained variance versus the component number and select the point at which the explained variance begins to change minimally, analogous to the elbow method used in K-means clustering.

To provide a consistent and automated approach, a method for selecting the number of retained principal components was devised, referred to as the slope rule. This method evaluates the local slope of the explained variance curve between consecutive principal components. Specifically, the slope between components p_i and p_{i+1} is computed as the difference in explained variance. A threshold slope of 0.01 was used throughout this dissertation. Principal component p_{i+1} is retained if the slope

exceeds this threshold. This criterion is not intended to maximize clustering accuracy or model performance for a specific dataset, but rather provide a consistent, automated, and reproducible baseline for PCA-based analyses. A visual example of the slope rule is shown in Figure 14, where the slope rule recommends five principal components be retained.




Figure 14: Example plot of explained variance versus principal component number for a PCA model fitted to a PMA spectrum.

4.5 Semi-Supervised Feature and Outlier Detection Methods

In some scenarios, reference devices such as known authentic devices or healthy samples may be available. When present, these reference devices enable the use of semi-supervised methods that model the expected behavior of a given device. Deviations from this reference model can be quantified using distance or reconstruction error metrics, which can be extracted as features for further analysis. These distance-based features provide a means of quantifying anomaly or degradation relative to a known baseline.

Several semi-supervised methods were explored in this dissertation, including the Mahalanobis distance (MD) [152], a modified version of the Mahalanobis distance

(MMD) [153], autoencoders (AE) [154], one-class SVM (OCSVM) [155], [156], and local outlier factor (LOF) [157]. Each method was trained using reference PMA spectra and used to calculate a scalar distance or error score relative to reference training data. Engineered features or principal components could alternatively be used in place of the PMA spectra; however, doing so may limit the model's ability to capture nontrivial patterns within the spectra. Conversely, training methods like MD using high-dimensional data can result in increased sensitivity to small variations, potentially leading to incorrect differentiation of samples that are assumed to be identical. While this dissertation trains these semi-supervised methods using PMA spectra, the reported results may not be directly representative of the model performance when trained using principal components or engineered features.

The MMD method proposed by Ott *et al.* [153] diagonalizes the covariance matrix and normalizes each feature by its variance, thereby avoiding matrix singularity issues when the number of reference samples is small relative to the number of features. This approach also improves computational efficiency compared to the traditional MD. However, diagonalization removes information contained in the off-diagonal covariance terms, which may reflect important feature correlations. For comparison, the traditional Mahalanobis distance (MD), which measures the distance between a sample and reference devices, was also used.

AEs are neural networks that use an encoder to compress the input data and a decoder that reconstructs the input at the output. In this work, AEs were implemented as a feedforward neural network trained on reference data using the mean squared error

(MSE) loss function. The reconstruction error between the input and output was used as a distance-based feature to quantify deviation from the reference population.

One-class classification methods attempt to determine whether a sample belongs to the reference training class [156]. In this work, an OCSVM was trained using reference PMA spectra, and the signed distance from the decision function was used as a distance-based feature representing anomaly magnitude.

The LOF algorithm is a density-based method that compares the local density of a data point to its nearest neighbors. A sample is considered an outlier if it has sufficiently lower density than its neighbors. The inverse of the local density was used as an anomaly score for each sample.

4.6 Supervised Methods

In rare scenarios where fully labeled data are available, such as datasets containing multiple known stages of degradation or both authentic and counterfeit devices, supervised learning methods can be employed. Common supervised approaches include SVM, random forest (RF), multilayer perceptron (MLP), and partial least squares regression (PLSR). These methods can be used for classification or regression tasks, and in some cases, both. In this dissertation, supervised methods, such as RF, are used as benchmarks for comparing different electrical test methods and extracted feature sets.

4.7 Quantifying the Separation Between Samples Using a Single Feature

Fully supervised methods typically require labeled training data spanning multiple degradation stages or a monotonic progression in the measured response

across stages of degradation. While clustering can provide a coarse indication of separation between groups, the resulting distance between clusters is influenced by the chosen algorithm and hyperparameters, which can impact the outcome. Clustering can indicate whether groups are separable, but it does not provide a stable, feature-level scalar measure of separation that can be consistently compared across datasets. Intrinsic clustering metrics evaluate the quality of clustering but do not directly quantify the separation or progression between classes. Semi-supervised methods, such as MD, offer a more robust approach to quantifying deviation from the reference dataset but require access to reference samples for training.

Both clustering and semi-supervised approaches assess the dataset holistically by considering all features simultaneously. This provides a global perspective of the dataset but can be influenced by a handful of features or samples, making it challenging to identify which features are driving separation in higher-dimensional space. To address these limitations, an approach was developed that utilizes the distribution of the engineered features and domain knowledge to determine whether multiple populations exist within the dataset and to quantify the separation between those populations on a feature-by-feature basis.

Non-normalized engineered features retained after feature selection can be analyzed individually by examining their empirical distributions. Plotting these features as histograms enables direct computation of distances between populations on a per-feature basis. Unlike PCA, where the distances between samples are defined in latent space that depends on the dataset composition, distances computed from engineered features are grounded in physically and mathematically meaningful

quantities that are consistent across datasets. This enables the definition of acceptable feature ranges and supports outlier detection based on domain-relevant thresholds.

This feature-wise analysis can be extended through modality assessment to determine whether multiple classes exist within the dataset that is otherwise assumed to be homogeneous. GMMs can be used to estimate the number of classes present in a given feature distribution. Model selection criteria such as the AIC and BIC are used to determine the appropriate number of components, consistent with standard GMM-based clustering practices. When multiple classes are identified, the separation between the corresponding feature distributions can be quantified using the distance between their centers, with or without access to ground-truth labels. While other clustering algorithms could be used, GMMs are well-suited for this task because they directly model the distribution of the desired feature.

Several statistical metrics have been developed to quantify separation between bimodal or multimodal distributions, including Ashman's D and Cohen's d , which are metrics that can be generalized to multimodal datasets using pairwise comparisons [158], [159]. Ashman's D measures the parametric distance between two Gaussian distributions, with values greater than 2.0 indicating well-separated classes. Cohen's d is an effect size metric that quantifies the standardized difference between means using the pooled standard deviation of the populations, where values greater than 1.2 indicate good separation and values greater than 2.0 suggest minimal overlap. Fisher's discriminant ratio (FDR), which is equal to the square of Ashman's D , provides a related metric applicable to any two groups or classes regardless of assumed distribution. Hypothesis test methods such as ANOVA can also be applied to assess

whether two distributions differ statistically; however, the resulting p-values do not provide a scalar measure of separation comparable to the other metrics discussed.

This approach shares conceptual similarities with one-dimensional clustering and semi-supervised methods but differs in its feature-wise formulation and incorporation of domain knowledge. By operating directly on application-relevant and dataset-independent engineered features, the approach enables quantification of separation between degradation stages, detection of multiple populations, and identification of features driving the separation. When paired with domain knowledge, the approach can be further improved to incorporate confidence bonds for specific features or to use ground-truth labels for direct class comparison rather than relying solely on modality assessment.

Chapter 5: Detection and Quantification of Degradation: High-Temperature Bias

Degradation of microelectronic devices is an unavoidable byproduct of their operation. When devices are used within specified limits, degradation is typically slow and may have a negligible impact on performance for decades. This makes the detection of early-stage degradation challenging and is of particular concern for devices that are intended for long service lifetimes or operated under conditions that accelerate degradation mechanisms. As technology nodes continue to shrink transistor dimensions, devices may become more susceptible to failure by mechanisms including HCI, NBTI, and TDDB [87].

Conventional electrical test methods evaluate the functionality or parametric characteristics of a device to determine whether it is operating within its specified limits or has drifted from a predefined state. These test methods are often performed as part of qualification and reliability assessments outlined in standards such as MIL-STD-883 or JEDEC JESD22. While such testing can be relatively rapid for simple devices, such as diodes and MOSFETs with a limited number of critical electrical parameters, complex devices, including FPGAs and microcontrollers, may require verification of hundreds of electrical and functional parameters, making conventional testing slow and resource intensive. Side-channel techniques such as PMA offer a complementary approach that enables rapid assessment of device degradation, with sensitivity to subtle changes that may not yet manifest in the device's conventional electrical parameters, while requiring reduced measurement complexity.

5.1 Experimental Design

Evaluating the reliability of microelectronic devices often requires accelerated testing to induce degradation that would otherwise take years to manifest under normal operating conditions. Accelerated degradation is typically achieved by applying stress at or near the device's absolute maximum ratings. Stresses include environmental conditions such as elevated humidity and temperature, as well as electrical stresses.

The Toshiba TC74HC02APF quad two-input NOR logic gate is a 14-pin dual in-line (DIP) CMOS device selected for its low transistor count and simple, well-understood architecture that is representative of many planar CMOS devices used today. The relatively low complexity of the TC74HC02 was essential to ensuring that changes observed in the PMA spectra due to degradation were understandable and could be plausibly linked to underlying mechanisms. A pinout diagram of the TC74HC02 is provided in Figure 15 [160].




Figure 15: Pinout diagram of the TC74HC02APF.

A preliminary study was conducted using three TC74HC02 devices to identify an appropriate bias voltage that would accelerate degradation without inducing catastrophic failure. The primary experiment presented in this chapter was conducted

using twenty individual TC74HC02 devices. Of these, four devices were kept at room temperature and electrically unbiased to serve as control samples. The remaining sixteen devices were subjected to elevated temperature, a DC bias applied to V_{cc} , and a switching signal applied to the inputs. This combination of stresses is referred to as a high-temperature bias (HTB) test.

To facilitate electrical stimulation during aging and allow each device to be periodically removed for measurement, DIP sockets were mounted to a printed circuit board (PCB). The sockets are rated for operation up to 105 °C, while the PCB is assumed to be rated for approximately 110 °C based on comparable products. A test temperature of 125°C was selected based on the limitations imposed by the test hardware and to align with standard accelerated testing conditions defined in MIL-STD-883 [96] and JEDEC JESD22 [161], [162]. Prior to HTB testing, the unpopulated PCB and sockets were placed in the oven at 125 °C for 24 hours to verify mechanical stability and ensure no premature electrical failure occurred.

All devices on the PCB were connected in parallel to ensure identical electrical stress conditions. An Agilent E3634A DC power supply provided a 12.7 V bias to the V_{cc} pin of each device, while a Keysight EDU33211A arbitrary waveform generator supplied a 1 MHz 0–9 V square wave to the inputs. A switching frequency of 1 MHz was chosen, as higher frequencies resulted in increased distortion and a more sinusoidal waveform.

A 1 Ω shunt resistor was placed in series between the power supply and the devices to monitor the supply current during testing. The voltage drop across the shunt was compensated by setting the power supply output to 13.5 V, ensuring that the

intended 12.7 V bias was applied. To limit current, 1 k Ω resistors were placed in series between each device output and ground. These resistors were selected to allow for voltages above the absolute maximum rating while maintaining current near, or slightly above, the device's specified 50 mA limit. 2 W resistors were used to ensure adequate power dissipation at elevated temperatures.

A summary of the applied stress conditions and corresponding device specifications is provided in Table 6. Photographs of the populated PCB assembly, showing both the top and bottom views, are presented in Figure 16.

Table 6: Summary of applied HTB stresses and relevant TC74HC02 device specifications.

Parameter	Experiment	Absolute Maximum Rating of Device
Operating Temperature	125 °C	−40 °C to 85 °C
Absolute Supply Voltage (V_{cc})	12.7 V	−0.5 V to 7 V
Input Voltage	0–9 V	0– V_{cc}
Input Frequency	1 MHz	83.3 MHz ²




Figure 16: PCB top (left) and bottom (right) with DIP sockets and 1 k Ω resistors used for HTB testing.

Conventional electrical testing and PMA measurements were performed on all twenty devices prior to HTB testing. Sixteen devices were then inserted into the DIP sockets on the PCB and placed in the oven. A data acquisition unit (DAQ) recorded the

² Theoretical maximum switching speed based on propagation delay

voltage across the shunt and the temperature measured by two T-type thermocouples in the oven at five-minute intervals. A schematic of the experimental setup is shown in Figure 17.




Figure 17: Schematic of HTB test setup.

The current supplied to each device, computed from the voltage drop across the shunt resistor, and the temperature recorded during the HTB testing are shown in Figure 18. Periods during which the HTB test was not running have been removed for clarity. The shaded regions indicate distinct stages of degradation, after which PMA and electrical testing were performed. During the gaps between stages, the devices were unbiased and kept at room temperature. The average current supplied to each device decreased during the initial degradation stages before stabilizing at later stages, with the exception of two spikes.




Figure 18: Oven temperature (left) and approximate current supplied to each device (right) during the HTB study. Each color corresponds to a time interval after which PMA and electrical testing were performed.

PMA testing was conducted using both the Chiplytics PSA system and BTE, as described in Chapter 3; however, the majority of the results presented in this chapter were acquired using the Chiplytics PSA system. The modulating waveform was applied to the device's V_{cc} pin, while the device's ground pin was connected to ground. A 100 kHz square wave was used for all PMA measurements. Two sampling rates were used, including 100 MSa/s with five time-domain waveform periods and 33.3 MSa/s with fifteen time-domain waveform periods. Unless otherwise noted, the results presented in this chapter focus on the 100 MSa/s configuration.

The PMA voltage ranges differed between the BTE and Chiplytics PSA system, with a larger number of PMA voltage ranges evaluated using the Chiplytics PSA system. This increased voltage granularity enables improved observation of degradation-driven changes in the PMA spectrum and assessment of PMA sensitivity as a function of voltage range. A summary of the PMA voltage ranges used for each hardware configuration is provided in Table 7, where ranges are listed as the minimum voltage (V_{min}) to maximum voltage (V_{max}) for clarity. For all PMA testing, the DC offset was set to half of the peak-to-peak voltage, resulting in a minimum voltage 0 V.

The Chiplytics PSA system is limited to a maximum output waveform amplitude 5 V in low-resolution mode, which is why the device's specified maximum operating voltage of 6 V was not included. Three of the included PMA voltage ranges have maximum voltages below the device's specified 2 V minimum operating voltage, enabling assessment of PMA sensitivity at or below the minimum operating voltage.

Table 7: PMA voltage ranges used in the HTB study for measurements performed with the Chiplytics PSA system and the BTE configuration.

$V_{\min} - V_{\max}$ (V)	Chiplytics PSA system	Benchtop Equipment
0.0–0.5	✓	✓
0.0–1.0	✓	✓
0.0–1.5	✓	✓
0.0–2.0	✓	✓
0.0–2.5	✓	-
0.0–3.0	✓	-
0.0–3.5	✓	✓
0.0–4.0	✓	-
0.0–4.5	✓	-
0.0–5.0	✓	✓

Curve trace measurements were performed at each PMA measurement interval by sweeping the voltage applied between the device's V_{cc} and ground pin from -0.5 V to 7 V with a current compliance of 50 mA. All curve trace measurements were performed using an Agilent B1505A Curve Tracer with a custom 40-pin ZIF socket. The voltage sweep range was selected to align with the absolute minimum and maximum supply voltage specification for the TC74HC02.

Basic functional testing was conducted at each PMA measurement interval using Boolean logic table verification. Automated functional testing was performed using an Arduino Uno Rev3 and a benchtop power supply. Devices were inserted into a breadboard, where the benchtop power supply provided 4.5 V to V_{cc} , while the

Arduino swept the input logic states using digital inputs. The device outputs were measured using the Arduino's analog inputs over multiple acquisitions, then averaged.

Transition time and propagation delay measurements were performed after devices reached 1014 hours of cumulative HTB exposure. These measurements were conducted using a benchtop power supply set to 4.5 V, an Agilent 33250A waveform generator, and a benchtop oscilloscope. The Agilent 33250A was selected over the Keysight EDU33211A used for BTE PMA measurements due to its faster rise time, which is required to meet the testing specifications outlined in the device datasheet.

Although not a conventional electrical test, LCR sweep measurements were performed to characterize the impedance and phase response between the device's V_{cc} and ground pins. An Agilent E4980A was used to perform LCR sweeps from 100 kHz to 2 MHz with a step size of 10 kHz, resulting in 200 impedance and phase measurements per sweep. Measurements were performed using 0–1 V, 0–2 V, and 0–3.5 V inputs to sample device behavior below, near, and above the specified minimum operating voltage.

5.2 Assessment of PMA Repeatability and Variance

Assessing the repeatability of PMA measurements is essential to ensuring that observed changes in the PMA spectra are attributable to the applied stress conditions rather than measurement variability or environmental factors. Both systematic sources of variation and random noise must be quantified to support confident interpretation of PMA results, particularly for subtle changes associated with degradation.

Repeatability, also referred to as gauge variation, describes the precision of a measurement system over short time intervals under consistent test conditions. It is

evaluated by repeatedly measuring the same characteristic on the same set of devices using an identical measurement setup. In practice, repeatability is often assessed alongside longer-term stability through gauge repeatability and reproducibility (GR&R) studies conducted over multiple days, sometimes using multiple operators [163], [164], [165]. The day-to-day (DtD) stability, part-to-part (PtP) variance, and short-term measurement repeatability were evaluated using the control devices included in the HTB study, while measurement system variance was evaluated using null spectra.

Quantification of DtD stability, PtP variance, and measurement repeatability was performed using the mean absolute error (MAE) and root mean squared error (RMSE) computed from both the PMA spectra and time-domain voltage response. The RMSE between a response of interest y_a and a reference response y_r was computed using (2), while the MAE was computed using (3). The summation index i corresponds to either the frequency-bin or time sample, and n denotes the total number of samples in the measured response. The Chiplytics PSA system was used as the primary platform for repeatability evaluation due to issues encountered with the BTE setup, as will be discussed later in this section.

$$\text{RMSE} = \sqrt{\sum_{i=1}^n \frac{(y_{a_i} - y_{r_i})^2}{n}} \quad (2)$$

$$\text{MAE} = \frac{\sum_{i=1}^n |y_{a_i} - y_{r_i}|}{n} \quad (3)$$

A PMA spectrum measured with no device present captures the waveform generator output as well as any effects associated with socketing and internal circuitry.

This PMA spectrum is referred to as a null spectrum. Fifty null spectra were captured using the low-resolution mode of the Chiplytics PSA system across several PMA voltage ranges. The mean and standard deviation of the MAE and RMSE were computed from the 1,225 pairwise comparisons between the null spectra. Pairwise comparisons were used to avoid any bias that may emerge from selecting a single reference spectrum and to capture the full distribution of variability. The error metrics computed from the PMA spectra are summarized in Table 8, while the corresponding time-domain metrics are provided in Table 9.

Table 8: Mean and standard deviation of MAE and RMSE computed from fifty PMA null spectra across different PMA voltage ranges.

$V_{\min}-V_{\max}$ (V)	MAE (dBV)		RMSE (dBV)	
	Mean	Std Dev	Mean	Std Dev
0.0–1.0	0.49	0.01	0.63	0.01
0.0–1.5	0.48	0.01	0.62	0.01
0.0–2.0	0.48	0.01	0.62	0.01
0.0–2.5	0.48	0.01	0.62	0.01
0.0–3.0	0.48	0.01	0.62	0.01
0.0–3.5	0.48	0.01	0.62	0.01
0.0–4.0	0.48	0.01	0.62	0.01
0.0–4.5	0.48	0.01	0.62	0.01
0.0–5.0	0.47	0.01	0.61	0.01

Table 9: Mean and standard deviation of MAE and RMSE computed from fifty time-domain null spectra across different PMA voltage ranges.

$V_{\min}-V_{\max}$ (V)	MAE (mV)		RMSE (mV)	
	Mean	Std Dev	Mean	Std Dev
0.0–1.0	1.57	0.29	1.95	0.31
0.0–1.5	1.55	0.25	1.93	0.27
0.0–2.0	1.51	0.20	1.88	0.22
0.0–2.5	1.55	0.22	1.92	0.24
0.0–3.0	1.59	0.26	1.98	0.29
0.0–3.5	1.54	0.19	1.92	0.21
0.0–4.0	1.61	0.25	2.00	0.28
0.0–4.5	1.68	0.28	2.09	0.32
0.0–5.0	2.88	1.23	12.72	11.79

The AD3 within the Chiplytics PSA system has an absolute resolution of 3.36 mV and an accuracy of $\pm 100 \text{ mV} \pm 0.5\%$ when operating in low-resolution mode, or 0.336 mV and $\pm 10 \text{ mV} \pm 0.5\%$ when operating in high-resolution mode [111]. The time-domain results indicate that the null spectrum variance is below the nominal resolution of the AD3, which is possible due to averaging across repeated measurements. Interpretation of the PMA spectrum error metrics is less intuitive due to their representation in dBV; however, quantifying the MAE and RMSE of the PMA spectrum is critical, as these spectra are used directly for feature extraction and PCA-based analysis.

A difference of +6 dBV corresponds to a doubling of the signal magnitude, while a difference of +20 dBV corresponds to a tenfold increase. Because the RMSE and MAE are computed using differences between logarithmic values, these metrics represent the amplitude ratios between spectra through the logarithmic quotient rule. The linear amplitude ratios can be obtained using (4), and the percentage difference in the amplitude can be computed using (5). Using these conversions, the average

amplitude difference between null spectra was 5.7% using the MAE and 7.4% based on the RMSE, indicating minimal variation between the measured PMA spectra.

$$\text{Linear ratio} = 10^{\frac{\text{amplitude ratio}}{20}} \quad (4)$$

$$\text{Percentage difference} = |\text{Linear Ratio} - 1| \times 100\% \quad (5)$$

The Chiplytics PSA system exhibited similar MAE and RMSE values across all evaluated PMA voltage ranges, except for the time-domain results obtained using a 0–5 V waveform. The increased error observed at 0–5 V is attributed to variations in the recorded voltage during voltage transitions, as well as the overshoot behavior, possibly caused by the AD3 operating near its supply limits. As shown in Figure 19, the fifty individual 0–4.5 V time-domain responses overlap closely, while substantial variance is observed among the fifty 0–5 V responses. This difference is most apparent when comparing the average response, shown in red, to the individual responses, shown in green. Notably, this increase in variance is not observed in the corresponding PMA spectra, likely because the transient behavior during voltage transitions has a minimal impact on the resulting FFT.




Figure 19: Comparison of time-domain null spectra using 0–4.5 V square wave (left) and 0–5.0 V square wave (right).

The null spectra were further assessed by removing the interharmonic components and focusing on the harmonic ones. When the interharmonic components were excluded, the average MAE across all PMA voltage ranges was 0.34 dBV compared with 0.48 dBV from the complete PMA spectrum, while the standard deviation remained unchanged. Similarly, the average RMSE was 0.50 dBV compared with 0.62 dBV; however, the standard deviation increased from 0.01 dBV to 0.02 dBV. These changes correspond to an average linear amplitude difference of 4.0% based on the MAE, representing a 29.8% improvement relative to the complete PMA spectrum analysis. An average linear amplitude difference of 5.9% is observed based on the RMSE, corresponding to a 20.3% improvement. These results indicate that the interharmonic components act as a significant source of variability, while providing minimal information outside of the noise floor. Because this dissertation focuses on engineered features, all results include the interharmonic components unless stated otherwise. The null spectrum analysis of the Chiplytics PSA system establishes a baseline against which the DtD stability, PtP variance, and short-term device repeatability can be assessed.

The DtD stability of the control devices was evaluated by comparing the PMA spectra of an individual device on a given day to the spectra of the same device captured on the first measurement day (03-17). Using the first measurement day as a reference provides insight into system or device drift that may occur over the course of the HTB study. Because four control devices were included, the DtD error metrics reported for each day represent the average of four pairwise comparisons. The PtP variance was

evaluated by comparing the four control devices to one another on the same day, resulting in six pairwise comparisons per day.

An MAE timeline plot illustrating the average DtD stability and PtP variance for each PMA voltage range is shown in Figure 20. Error bars represent the minimum and maximum MAE values obtained across the corresponding pairwise comparisons. The minimum and maximum were plotted instead of the standard deviation due to the small sample size. While only MAE results are presented in this chapter, the RMSE results were observed to follow the same overall trends and therefore are omitted here for brevity. Corresponding RMSE plots are provided in Appendix A. The final measurement day (07-28) is included in the analysis to provide additional insight into longer-term repeatability, despite not corresponding to a degradation stage.

Over the course of the HTB study, the PtP variance remained relatively constant, indicating that the control devices did not diverge from one another over time and remained consistent on each measurement day. The PtP MAE values for most PMA voltage ranges are less than or equal to the MAE computed from the null spectrum analysis, demonstrating the stability of the TC74HC02 devices. In contrast, the DtD stability gradually increased over the study, with spikes in MAE observed for certain PMA voltage ranges on 05-19 and 07-28. The DtD MAE values for most PMA voltage ranges during the first three measurement days are comparable to the MAE values observed in the null spectrum analysis; however, after the fourth measurement day (04-07), the DtD for multiple PMA voltage ranges exceeded the null analysis MAE. This gradual increase is more apparent in the MAE computed from the time-domain voltage response, as shown in Figure 21. The increasing MAE trend is not observed in

the PtP results, indicating the DtD trend may be related to external factors, such as environmental conditions or measurement system drift, rather than changes in the devices.




Figure 20: Timeline plots for MAE (dBV) of DtD stability (left) and PtP variance (right) across several PMA voltage ranges.




Figure 21: Timeline plot of MAE (mV) for DtD stability.

Although timeline plots are useful for identifying temporal trends, they are less effective for comparing repeatability across different PMA voltage ranges. The study-wide DtD stability was computed for each PMA voltage range using all available DtD pairwise combinations, rather than exclusively using the first measurement day as a reference. A total of 21 pairwise comparisons exist for each control device across the

seven measurement days. Similarly, the study-wide PtP variance was computed by comparing each device measured on a specific day to every other device measured on every other day, resulting in 42 pairwise comparisons per device.

The study-wide DtD and PtP MAE distributions are presented using box-and-whisker plots in Figure 22. The DtD results show a decrease in the median MAE value with increasing PMA voltage range with relatively consistent tails from 0–1 V to 0–2 V, followed by an increase that remains relatively constant at higher PMA voltage ranges. Notably, the 0–3.5 V DtD results exhibit lower variance relative to neighboring PMA voltage ranges. The same behavior is not observed in the PtP results, which show a consistent decrease in median MAE with increasing peak voltage. This trend suggests that PMA measurements performed using lower peak voltages are more susceptible to systematic noise originating from variations between devices or the measurement equipment.




Figure 22: Box-and-whisker plots of MAE (dBV) for DtD stability (left) and PtP variance (right) across nine PMA voltage ranges.

The consistency of control devices over the HTB study can be visualized by plotting the average 0–3.5 V PMA spectra on each measurement day, as shown in Figure 23. The average PMA spectra on most days overlap, making them difficult to

distinguish, with the exception of the last measurement day (07-28). A small deviation between the average PMA spectrum measured on 07-28 and those measured on earlier days is observed around 20 MHz. This deviation is consistent with the increase in MAE observed on 07-28 for PMA voltage ranges above 0–2.0 V.




Figure 23: Average PMA spectra (0–3.5 V) of control devices measured on different dates.

Short-term measurement repeatability was evaluated by acquiring ten PMA spectra for each of the four control devices. Measurements were performed in a randomized order, with each device being removed from and reinserted into the socket between measurements, thereby capturing both electrical and mechanical sources of measurement variability. The MAE was computed for the 45 pairwise combinations for each control device, the results of which are shown using a box-and-whisker plot in Figure 24. There is a gradual decrease in the median MAE as the maximum voltage amplitude increases, but the median remains relatively consistent across the voltages, particularly when compared to the trends observed in the PtP results. For all PMA voltage ranges, the obtained MAE values are less than or equal to those obtained from the null spectrum analysis.




Figure 24: Box-and-whisker plots of MAE (dBV) for each control device based on ten repeated measurements.

A similar short-term analysis was performed using PMA spectra acquired using the high-resolution mode on the Chiplytics PSA on 06-04. The average PMA-based MAE in high-resolution is 0.43 dBV compared to 0.46 dBV in low-resolution mode, while the average time-domain voltage-derived MAE is 0.5 mV compared to 1.98 mV. This reduction in MAE demonstrates the improved measurement precision of the high-resolution mode.

Collectively, the null spectrum analysis, short-term measurement repeatability, DtD stability, and PtP variance demonstrate the low repeatability error of the Chiplytics PSA system, supporting its suitability for scientific analysis of PMA results. The TC74HC02APF control devices also exhibited strong repeatability and stability throughout the HTB study. However, the gradual increase in DtD MAE observed over time warrants further investigation to identify the root cause. Understanding the source of this trend will provide vital information for the development of mitigation strategies to ensure that PMA measurements acquired on different dates remain comparable,

which is essential for long-term usability of PMA in applications including distribution centers.

In contrast, the MAE values obtained using the BTE setup are an order of magnitude higher for PMA voltage ranges exceeding 0–2.0 V. Because these values are reported on a dBV scale, this corresponds to change of approximately a factor of two. Below this voltage range, the DtD and PtP MAE values are comparable to those from PMA measurements from the Chiplytics PSA system. Box-and-whisker plots of the DtD and PtP MAE for the BTE setup, shown in Figure 25, highlight the large increase in MAE at higher PMA voltage ranges. Interharmonic components were removed from the PMA spectra based on the null spectrum analysis results; however, this did not affect the resulting MAE values, with the exception of the 0–2.0 V voltage range, for which the MAE increased. The increase in MAE at or above 0–2.0 V may be related to device behavior near its minimum operating voltage interacting with the measurement circuit.




Figure 25: Box-and-whisker plots of MAE (dBV) of the DtD stability (left) and PtP variance (right) of the BTE setup.

5.3 Results

Initial PMA measurements were performed on all twenty devices prior to HTB exposure and were measured again after 46, 136, 297, 520, and 1014 hours of cumulative exposure. The control devices were measured at each point the sixteen HTB-exposed devices were. The average PMA spectra of the sixteen devices at each degradation stage, acquired using a 0–3.5 V square wave, are shown in Figure 26.

The PMA spectra exhibit exposure-dependent changes. The trough located near 15 MHz decreases in amplitude with increasing cumulative HTB exposure, while the trough located near 35 MHz shifts toward higher frequencies as exposure time increases. Although visual inspection of the PMA spectra provides qualitative evidence of degradation-driven changes, further analysis is required to quantitatively characterize the degradation stages.




Figure 26: Average PMA spectra (0–3.5 V) for each degradation stage captured using the Chiplytics PSA system in low-resolution mode.

5.3.1 Clustering

Visual clustering of PMA results provides a direct comparison to how PMA methods, such as PSA, are typically used. In this work, the analysis was extended by

applying clustering algorithms rather than relying solely on visual interpretation. K-means clustering was selected due to its widespread use and simplicity. The number of clusters was defined to match the number of known degradation stages, providing a consistent basis for comparison across PMA voltage ranges and post-processing approaches.

HDBSCAN was also selected as a representative hierarchical clustering method that may be more applicable in cases where the number of classes may be unknown. The minimum cluster size for HDBSCAN was set to two, representing the most permissive choice and ensuring comparability between PMA voltage ranges. However, this permissive approach can lead to overfitting to noise in the dataset. Clustering results obtained using both K-means and HDBSCAN across several PMA voltage ranges are summarized in Appendix A. Analysis of the clustering results in this chapter focuses on K-means, which typically performed better than HDBSCAN. Clustering was performed using both PCA fitted to the PMA spectra and engineered features. The number of retained principal components was determined using the slope rule discussed in Section 4.4.

A clustering accuracy of 93.8%, defined as the agreement between cluster assignment and known degradation stages, was achieved across the six degradation stages by using K-means clustering with two principal components from PCA fitted to the 0–3.5 V PMA spectra. Five of the sixteen samples from the 520-hour degradation stage were misclustered, with two samples assigned to the 297-hour stage devices and three samples assigned along with the 1014-hour stage devices. The first two principal

components are plotted in Figure 27, with each degradation stage indicated by a unique color.

Clustering quality was evaluated using intrinsic clustering metrics, yielding a silhouette score of 0.69, a Calinski–Harabasz index of 3808.3, a Davies–Bouldin score of 0.36, and a Dunn index of 0.11. Of these metrics, the silhouette score is well-suited for comparison as the score is bounded. The Dunn index and Davies–Bouldin score, while unbounded, are also good candidates for relative comparisons. While every clustering metric is influenced by the dimensionality of the dataset, the Calinski–Harabasz is particularly sensitive, with scores that can differ by multiple orders of magnitude.




Figure 27: K-means clustering results using two principal components from a PCA model fitted to PMA spectra (0–3.5 V) captured using the Chiplytics PSA system.

Perfect clustering accuracy was achieved for the 0–3.5 V PMA measurements when twelve principal components were retained; however, retaining more than two or three components was not justified in this context, as additional components individually accounted for a negligible fraction of the explained variance. When twelve principal components were used, the silhouette score decreased to 0.3, the Calinski–

Harabasz index decreased to 122.5, the Davies–Bouldin score increased to 1.44, and the Dunn index increased to 0.37. These scores demonstrate that improved clustering accuracy using higher-dimensional space does not necessarily correspond to improved clustering quality.

K-means clustering was also performed using the extracted engineered features. Perfect clustering of the 0–3.5 V PMA data was achieved using the FS-PCA features, yielding a silhouette score of 0.42, a Calinski-Harabasz index of 318.9, a Davies-Bouldin score of 1.1, and a Dunn index of 0.42. Overall clustering quality using the FS-PCA features was lower than PCA-based clustering using two principal components, with the exception of the Dunn index, which increased. Clustering performance using the FS-PCA features closely matched the PCA-based results obtained using twelve principal components. Achieving similar clustering quality with four times the number of features demonstrates that the engineered features encode degradation-relevant information without relying on high-dimensional latent representations. K-means clustering using the FS-SCC features achieved an accuracy of 92.7% for 0–3.5 V data, which is the only PMA voltage range for which clustering accuracy was lower than the PCA-based approach. On average, the FS-SCC features increased clustering accuracy by $10.5\% \pm 11.9\%$, while the FS-PCA features increased clustering accuracy by $11.6\% \pm 12.9\%$. Clustering accuracy for the 0–5.0 V FS-PCA feature set was lower than that achieved using principal-component-based clustering.

Improved feature selection was simulated using an LDA-based approach, as discussed in Section 4.4. On average, K-means clustering accuracy increased by

19.2% $\pm$ 12.6% across the PMA voltage ranges when using the LDA-based method compared to the PCA-based approach.

Clustering using HDBSCAN resulted in lower clustering accuracy compared to K-means. The average PCA-based clustering using HDBSCAN was 41.8% $\pm$ 18.2% lower than the corresponding K-means results. The FS-PCA features exhibited a similar decrease of 41.9% $\pm$ 11.6% when comparing HDBSCAN and K-means. FS-SCC features were not used with HDBSCAN due to their relatively high dimensionality. However, an 11.5% $\pm$ 13.3% improvement in HDBSCAN clustering accuracy was observed when using the FS-PCA features instead of principal components. The HDBSCAN clustering accuracy could likely have been improved through hyperparameter optimization, including adjusting the minimum cluster size.

The same clustering analysis was repeated using the PMA data acquired with the BTE setup. The 520-hour degradation stage measurements were excluded due to data acquisition issues encountered during that testing interval. This exclusion reduced the number of classes from six to five. A clustering accuracy of 57.5% was achieved using five principal components, which increased to 72.5% after removing the interharmonic components and retaining four principal components. The interharmonic components were removed to assess whether they were a primary contributor to variance that degraded PCA-based clustering performance. In contrast, removing the interharmonic components from the Chiplytics PSA system data resulted in a decrease in average clustering accuracy of 2.6% $\pm$ 11.6%, with the only notable improvement occurring at 0–2.5 V. These results indicate that interharmonic components can

introduce variance that negatively impacts PCA-based clustering, particularly for PMA measurements acquired using BTE configurations.

Clustering using the FS-SCC features extracted from the BTE-acquired PMA spectra improved clustering accuracy by $19.2\% \pm 20.5\%$ on average relative to PCA-based clustering. The FS-SCC features outperformed the FS-PCA features by an average margin of $9.2\% \pm 14.7\%$. Perfect clustering of the 0–3.5 V results was achieved using the FS-SCC features, compared to 57.5% accuracy achieved using PCA-based clustering.

Finally, clustering using the Chiplytics PSA system acquired measurements was performed again with the 0–3.5 V PMA measurements of the control devices, to confirm the conclusions drawn from the repeatability assessment. The number of clusters remained unchanged, as the control devices should be identical to the 0-hour measurements. Clustering accuracy using PCA improved to 96.7%, with only four of the 520-hour devices being misclustered, all assigned to the 1014-hour cluster. The control and 0-hour stage cluster exhibited similar spread to the other degradation-stage clusters. The overall trend remained the same as showed in Figure 28. While the pattern remained the same, the clustering accuracy decreased at other PMA voltage ranges, as the controls and 0-hour devices were split into two clusters and the 520-hour and 1014-hour devices were clustered together.




Figure 28: K-means clustering results using two principal components from a PCA model fitted to PMA spectra (0–3.5 V) including HTB control devices captured using the Chiplytics PSA system.

Removing the interharmonic components resulted in a more cohesive control and 0-hour cluster; however, the overall clustering accuracy decreased to 95.0%. While this accuracy is higher than that obtained using PCA without the inclusion of control devices, the number of misclustered samples increased from five to six, with one 1014-hour device being misclustered with the 520-hour devices. The increased cohesion of the control and 0-hour cluster reaffirms the null spectrum analysis, which demonstrated that interharmonic components contributed to increased error.

Clustering using the engineered features resulted in reduced accuracy, with the FS-SCC features and FS-PCA achieving an accuracy of approximately 75.0% due to clustering of the 520-hour and 1014-hour devices into a single group while separating control devices into two groups. In contrast, LDA-based feature selection achieved 100% clustering accuracy, while using the first eight of the FS-PCA features achieved 95.0% accuracy. These results confirm that any changes observed in the control devices are minor relative to the separation between degradation stages, although feature-based

clustering exhibited greater sensitivity to class overlap among the later degradation stages when the control devices were included. Despite the decrease in clustering accuracy, the trend between the degradation stages remained the same with and without the control devices, and the relative separation between the controls and 0-hour devices was minimal.

5.3.2 Semi-Supervised Methods

In cases where data from healthy devices are available as well as devices of unknown health status, semi-supervised methods can provide a framework for quantifying degradation. Several semi-supervised approaches were evaluated, including the traditional Mahalanobis distance (MD), a modified version of the Mahalanobis distance (MMD) developed by Ott *et al.*, LOF, OCSVM, and autoencoders. These semi-supervised techniques were trained using data from healthy devices and generate distance- or error-based metrics that indicate deviation from the learned healthy reference. While engineered features or principal components extracted from the PMA spectra could alternatively be used as model inputs, doing so constrains the learning process to predefined representations and may limit the model's ability to capture nontrivial patterns within the spectra. However, using methods like the MD with high dimensionality data can result in the model being too sensitive. The resulting metrics were visualized as histograms, with each degradation stage indicated by a unique color. The median of each degradation stage is indicated by a vertical dashed line.

MD is commonly used in PHM and anomaly detection applications. Distances obtained using both the MD and MMD are presented as histograms in Figure 28, with

the MD shown on the left and the MMD shown on the right. Both approaches demonstrate clear separation between the 0-hour, 46-hour, and later-stage devices. The relative separation between the 136-, 297-, 520-, and 1014-hour devices is smaller in the MD results compared to the MMD; however, the difference between these later stages remains minimal for both methods. Neither MD nor MMD yields non-overlapping separation among the 297- to 1014-hour degradation stages. The use of MD- or MMD-based metrics in conjunction with engineered features may offer improved discrimination or downstream model performance.




Figure 29: MD (left) and MMD (right), trained on PMA spectra (0–3.5 V) of control devices, computed for devices at different stages of HTB degradation. Vertical dashed lines indicate the median of each degradation stage.

Unlike clustering-based approaches, MD and MMD provide scalar distance metrics that enable quantification of deviation from a healthy reference. The MD results indicate that the separation between healthy and 46-hour devices is comparable to the separation between the 46- and 297-hour devices, indicating that MD is sensitive to early degradation-driven changes. In contrast, the MMD results show a smaller separation between the healthy and 46-hour devices, while a more pronounced

separation emerges between the healthy and 136-hour devices. The MD results suggest that PMA can capture additional degradation progression between the 46- and 136-hour stages. The relative separation among the 297-, 520-, and 1014-hour stages for both MD and MMD is small compared to earlier degradation stages. These observations are consistent with visual inspection of the K-means clustering plot but provide a more quantitative assessment of degradation progression.

In addition to serving as a quantitative degradation metric, MD can be applied for outlier detection. Under a multivariate normal reference assumption, the squared MD, MD^2 , can be compared to the upper $(1-\alpha)$ quantile of a chi-squared distribution with p degrees of freedom, where p is the number of features used to compute the MD. Here, MD was computed using the PMA spectra containing 2500 frequency bins, for which the 95th-percentile threshold is $\chi^2_{2500,0.95} \approx 2617$. Thus, an MD of approximately 51.2 or greater corresponds to a $p \leq 0.05$. The MD for all degradation stages exceed this threshold. The distance between the healthy and 46-hour devices is large enough that early degradation may be detectable. The comparison between MD^2 and chi-square is not directly applicable when using the MMD.

LOF and OCSVM were also investigated as methods to quantify degradation stages. OCSVM is traditionally used to determine whether a sample belongs to a reference class or represents a novel observation, while LOF is commonly used for density-based anomaly detection, analogous to MD. The scores produced from LOF and OCSVM, using a radial basis function (RBF) kernel for OCSVM, are used as metrics. Metrics obtained from the LOF and OCSVM models are shown using histograms in Figure 29, with the LOF results shown on the left and the OCSVM results

shown on the right. While the magnitude of the LOF and OCSVM-derived metrics is different, the relative separation and ordering of each degradation stage are nearly identical. These metrics can be used similarly to MD or MMD, although the separation between degradation stages more closely resembles the MMD results than the MD results.




Figure 30: LOF (left) and OCSVM (right) metrics, trained on PMA spectra (0–3.5 V) of control devices, computed for devices at different stages of HTB degradation. Vertical dashed lines indicate the median of each degradation stage.

Investigation of quantifying degradation using autoencoders was also performed. The results, provided in Appendix A, are qualitatively similar to those obtained using MMD, LOF, and OCSVM, with no observed improvement in separation between degradation stages.

While the investigated semi-supervised methods provided similar results, each has advantages and limitations. OCSVM and AE are capable of learning nonlinear behaviors within the dataset; however, the resulting metrics are difficult to relate directly to underlying data characteristics. LOF is sensitive to local effects within large datasets and requires minimal hyperparameter tuning. MD is simple, requires no

hyperparameters, and produces results that can be interpreted as distances from a reference population; however, the results should be interpreted within a statistical framework that accounts for the dimensionality of the feature space. For anomaly detection, MD or LOF are well-suited when simplicity and transparency are primary objectives. The Q-residual and Hotelling's T^2 statistics used by Sandia National Laboratories represent an alternative approach for outlier detection; however, these methods assess samples relative to a PCA-based model and therefore inherit the dataset dependence and limited feature traceability associated with PCA.

5.3.3 Conventional Electrical Test Results

Curve trace measurements were performed in conjunction with PMA measurements. The average curve trace response at each degradation stage is the left plot presented in Figure 30, where the shaded area represents ± 1 standard deviation. The "controls" curve, shown in cyan, includes measurements of the control devices performed at each stage of degradation. While the same socket, cable, and adapter were used for each measurement, the control curve does not align with the 0-hour curve, indicating that the electrical response of the control devices changed over the course of the HTB study. The variance of the curve trace measurements is visible in the large shaded area around the average control measurement. The control devices appear to shift after the 136-hour results, as shown in the right plot in Figure 30. This shift is likely attributable to environmental factors, such as temperature.

The curve trace results show nonmonotonic progression, with the 46-hour devices exhibiting lower current draw than the 0-hour devices, while the 136-hour devices have less current draw than the 297-hour devices. Curve trace measurement of

the different degradation stages cannot be visually separated, even when color-coded. In general, degradation resulted in lower current draw and an apparent shift in the threshold region, which may indicate carrier trapping and interface state generation.

The impedance and phase curves for 0–3.5 V are provided in Figure 31, while the 0–1 V and 0–1.5 V curves are provided in Appendix A. The impedance curves shown on the left, with the shaded area representing ± 1 standard deviation, overlap across different degradation stages, making them challenging to visually separate. However, the 0-hour and control curves are visually distinct from the degradation curve. Unlike the curve trace results, the control impedance measurements do not exhibit the same variance observed in the curve trace results, despite being performed using the same socket and devices. This suggests that the variance was isolated to the curve trace measurement system. Similar overlap is observed in the phase curves, which are shown on the right. The control and 0-hour devices were found to have a more resistive phase compared to the degraded devices. Collectively, the impedance and phase measurements provide validation of the PMA results by confirming exposure-dependent changes in the electrical characteristics of the devices.




Figure 31: Curve trace measurements of TC74HC02 devices from the HTB study (left) and control devices measured throughout the HTB study (right). Shaded regions in the left plot indicate ± 1 standard deviation.




Figure 32: Impedance sweep measurements using 0–3.5 V input of HTB degraded TC74HC02 devices, showing impedance (left) and phase (right). Shaded regions indicate ± 1 standard deviation.

Propagation delay and transition time were measured for the first gate of each device after 1014 hours of HTB exposure. Both parameters were measured when input A was switching with input B held at 0 V, and when input B was switching with input A held 0 V. The rising- and falling-edge behavior was measured, resulting in four timing measurements per device. Due to hardware configuration-related concerns, the degraded devices were compared to the control devices rather than using the raw measured timing characteristics. The average differences in the rising- and falling-edge behavior are plotted in Figure 32, along with error bars indicating the standard deviation. The rising-edge transition time measured with input B switching exhibited longer settling times and overshoot, making accurate measurements challenging. As a result, several of these measurements were excluded from the analysis.




Figure 33: Average difference in propagation delay and transition time between control devices and devices degraded after 1014 hours of HTB exposure.

If an increase in propagation delay or transition time were observed exclusively on either the rising- or falling-edges, it would indicate degradation localized to the NMOS or PMOS portions of the CMOS gates. An increase in the timing characteristic of the degraded devices relative to the controls was observed for both the falling- and rising-edges. This behavior indicates that degradation occurred in both the PMOS and NMOS portions of the devices.

HCI is an NMOS-dominated mechanism and primarily impacts the falling-edge behavior [166], while NBTI is a PMOS dominant mechanism that primarily impacts the rising-edge behavior [167]. The results presented here provide strong evidence that both failure mechanisms contributed to the observed degradation, which is also consistent with the applied stresses. Although the applied electrical and thermal stress could have driven time-dependent dielectric breakdown (TDDB), the impact is likely minimal due to the relatively low current densities employed in the HTB study.

Functionality of each device was verified by testing its Boolean logic table. All devices produced the correct logical outputs at every stage of degradation, indicating that functional failure never occurred.

5.3.4 Regression and Classification

Regression and classification were performed to evaluate and benchmark the predictive capability of different datasets. Regression aims to identify relationships between features in order to predict continuous-valued outputs, while classification assigns unlabeled samples to one of the classes observed during training. Regression was performed using four machine learning models: a multi-layer perceptron (MLP) regressor, a random forest (RF) regressor, partial least squares regression (PLSR), and support vector regression (SVR). Similarly, classification was performed using MLP, RF, and support vector classification (SVC).

The MLP network consisted of two fully connected hidden layers, a sixteen node dense layer connected to an eight node dense layer, both utilizing the rectified linear unit (ReLU) activation function. Both regression and classification were performed using five-fold cross-validation, with each fold repeated three times, resulting in a total of fifteen unique training and test splits. Hyperparameter tuning for the SVR and SVC was accomplished using nested cross-validation with an inner grid search. A linear kernel was used for both SVM-based methods. The reported $\pm$ values represent the standard deviation across the five cross-validation folds.

Different input datasets were used to train and evaluate each model, including PMA FS-SCC features, PMA spectra, principal components from a PCA model fitted to the PMA spectra, the FS-PCA feature set, and the PMA spectra with the

interharmonic components removed (referred to as the harmonic spectra). In addition to PMA-based inputs, curve trace measurements and impedance sweep measurements were also evaluated as model inputs.

Regression performance was evaluated using the R^2 and RMSE, while classification performance was evaluated using global classification accuracy and the F1 score. The regression results are provided in Table 10 while the classification results are provided in Table 11. All PMA-based results presented in this section were obtained using the 0–3.5 V PMA spectra, which were selected to align with the clustering and semi-supervised analysis performed.

The FS-SCC features and FS-PCA features provided the strongest regression performance across the evaluated models. For the MLP regressor, the FS-SCC feature achieved both the highest R^2 (0.93 ± 0.03) and lowest RMSE (86.9 ± 24.6 hours), while PCA-based inputs yielded comparable R^2 values but higher prediction error. Although PCA achieved the second-highest R^2 , the FS-PCA features produced the second-lowest RMSE, indicating that variance compression alone does not guarantee improved predictive accuracy.

FS-SCC features also resulted in the best overall regression performance using RF, which achieved the lowest RMSE of 56.9 ± 11.5 hours among all evaluated models and inputs. Relative to the total degradation duration of 1014 hours, this corresponds to an average prediction error of approximately 5.6%, demonstrating strong quantitative performance.

Table 10: Regression model performance for HTB-degraded TC74HC02 devices using different input datasets, including features extracted from the PMA spectra (0–3.5 V).

Method	Input	R ²	RMSE (hours)
MLP	FS-SCC Features	0.93 ± 0.03	86.92 ± 24.59
	PMA Spectra	0.79 ± 0.08	154.78 ± 42.98
	PCA	0.93 ± 0.04	99.51 ± 12.54
	FS-PCA Features	0.91 ± 0.05	95.96 ± 26.31
	Harmonic Spectra	0.87 ± 0.11	113.60 ± 48.79
	Impedance	0.91 ± 0.05	97.34 ± 25.70
	Phase	0.83 ± 0.10	133.09 ± 42.21
	Curve Trace	0.38 ± 0.32	251.53 ± 73.77
RF	FS-SCC Features	0.97 ± 0.02	56.90 ± 11.46
	PMA Spectra	0.95 ± 0.02	71.86 ± 22.39
	PCA	0.95 ± 0.03	68.84 ± 27.15
	FS-PCA Features	0.95 ± 0.03	66.61 ± 19.25
	Harmonic Spectra	0.95 ± 0.03	71.61 ± 21.59
	Impedance	0.89 ± 0.05	112.95 ± 35.26
	Phase	0.80 ± 0.13	143.50 ± 53.14
	Curve Trace	0.72 ± 0.15	175.13 ± 58.48
PLSR	FS-SCC Features	0.90 ± 0.04	102.18 ± 14.16
	PMA Spectra	0.76 ± 0.07	161.12 ± 29.46
	PCA	0.73 ± 0.08	169.55 ± 13.32
	FS-PCA Features	0.90 ± 0.04	101.37 ± 15.96
	Harmonic Spectra	0.91 ± 0.04	96.75 ± 15.43
	Impedance	0.79 ± 0.20	142.49 ± 52.99
	Phase	0.83 ± 0.07	132.84 ± 25.42
	Curve Trace	0.67 ± 0.15	187.60 ± 49.94
SVR	FS-SCC Features	0.85 ± 0.07	126.15 ± 17.56
	PMA Spectra	0.76 ± 0.07	161.10 ± 29.51
	PCA	0.67 ± 0.09	193.12 ± 40.17
	FS-PCA Features	0.88 ± 0.07	112.91 ± 22.78
	Harmonic Spectra	0.91 ± 0.03	99.60 ± 15.45
	Impedance	0.72 ± 0.29	162.24 ± 65.52
	Phase	0.79 ± 0.09	146.58 ± 26.52
	Curve Trace	0.67 ± 0.13	188.46 ± 46.74

In contrast, the PMA harmonic spectra provided the strongest regression performance when using PLSR and SVR, with FS-PCA features performing comparably. This trend suggests that interharmonic components may provide limited information related to degradation for linear or kernel-based regression models. SVR using an RBF kernel resulted in lower overall regression performance for every input dataset, except for PCA.

The phase and impedance consistently provided outperformed the curve trace measurements in regression performance, which exhibited the poorest performance across all models. The MLP and RF models achieved lower RMSE when using the impedance compared to the phase, whereas the PLSR and SVR models performed better with the phase. This behavior is consistent with the differing sensitivities of tree-based, neural network, and linear regression models to input domains.

Overall, the results indicate that the engineered features provided a comprehensive overview of the degradation behavior. The FS-PCA features outperformed the FS-SCC features in three of the four regression models, suggesting that additional feature selection strategies may further improve predictive accuracy. For example, excluding non-spectral skewness and kurtosis features reduced RMSE for both the PLSR and RF models when using the FS-PCA feature set.

Similar to the regression results, the FS-SCC features and FS-PCA features provided strong classification performance with RF and SVC models but performed poorly with the MLP model. The strongest MLP performance was achieved using the harmonic spectra and phase measurements, while the FS-SCC features, FS-PCA features, PCA, and PMA spectra all performed markedly worse. This suggests that the MLP architecture used may have struggled with the features related to interharmonic components, rather than indicating a fundamental limitation of the engineered features.

Table 11: Classification model performance for HTB-degraded TC74HC02 devices using different input datasets, including features extracted from the PMA spectra (0–3.5 V).

Method	Input	Accuracy	F1
MLP	FS-SCC Features	0.76 ± 0.07	0.70 ± 0.07
	PMA Spectra	0.57 ± 0.08	0.52 ± 0.10
	PCA	0.41 ± 0.18	0.33 ± 0.18
	FS-PCA Features	0.69 ± 0.12	0.64 ± 0.14
	Harmonic Spectra	0.91 ± 0.08	0.91 ± 0.09
	Impedance	0.73 ± 0.09	0.67 ± 0.08
	Phase	0.83 ± 0.15	0.81 ± 0.17
	Curve Trace	0.67 ± 0.10	0.60 ± 0.11
RF	FS-SCC Features	0.99 ± 0.02	0.99 ± 0.02
	PMA Spectra	0.98 ± 0.03	0.98 ± 0.03
	PCA	0.98 ± 0.04	0.97 ± 0.04
	FS-PCA Features	0.99 ± 0.02	0.99 ± 0.02
	Harmonic Spectra	0.99 ± 0.03	0.98 ± 0.03
	Impedance	0.98 ± 0.02	0.98 ± 0.03
	Phase	1.00 ± 0.0	1.00 ± 0.0
	Curve Trace	0.89 ± 0.07	0.88 ± 0.08
SVC	FS-SCC Features	0.98 ± 0.03	0.98 ± 0.03
	PMA Spectra	0.93 ± 0.06	0.92 ± 0.07
	PCA	0.95 ± 0.06	0.95 ± 0.07
	FS-PCA Features	0.99 ± 0.03	0.99 ± 0.03
	Harmonic Spectra	0.97 ± 0.03	0.96 ± 0.04
	Impedance	0.98 ± 0.03	0.98 ± 0.03
	Phase	1.00 ± 0.01	1.00 ± 0.02
	Curve Trace	0.88 ± 0.06	0.87 ± 0.06

Both the RF and SVC models performed well using the FS-PCA and FS-SCC features, yielding consistently high classification accuracy and F1 scores. Surprisingly, the phase measurement achieved perfect classification for both RF and SVC models. The RF model demonstrated strong robustness across nearly every input, achieving accuracies of approximately 0.98 or higher for every input except curve trace measurements. In contrast, the SVC model exhibited greater sensitivity to inputs, although classification accuracy exceeded 0.95 in most cases. SVC using an RBF kernel was also evaluated and produced similar results to the linear kernel, indicating limited benefit from nonlinear approaches for classification of this dataset.

Overall, these classification results demonstrate that engineered PMA features provide strong and consistent classification performance, particularly when compared to using the PMA spectra. As observed from the regression results, additional or refined feature selection may further improve classification accuracy. For example, excluding non-spectral skewness and kurtosis features increased the classification accuracy to 1.0 when using FS-PCA features with the RF model. The strong classification performance achieved by using phase measurements for both the RF and SCV models suggests that a multi-modal approach combining PMA-derived features with impedance or phase measurements could provide a robust and accurate framework for degradation classification.

Regression and classification were repeated with the control devices included to evaluate whether their inclusion would negatively impact performance, as was observed in the K-means results when using engineered features. RF regression using the FS-SCC feature achieved an R^2 of 0.98 ± 0.02 and RMSE of 46.42 ± 14.16 hours, while the FS-PCA features achieved an R^2 of 0.97 ± 0.02 and RMSE of 44.35 ± 18.86 hours. Similarly, RF-based classification achieved an accuracy of 0.99 ± 0.01 using the FS-PCA features and 1.0 using the FS-SCC features.

The inclusion of the control devices improved both regression and classification performance, in contrast to the decrease observed for K-means clustering. This behavior suggests that class imbalance may have disproportionately affected K-means clustering, whereas supervised methods are more robust. Class imbalance also influences feature ranking using PCA, potentially resulting in top-ranked features that were more sensitive to variations in the control devices rather than to differences

between later degradation stages. This effect may explain why using a subset of the FS-PCA features or LDA-ranked features achieved clustering accuracy comparable to that obtained without the inclusion of control devices.

5.3.5 Quantification of Separation

The clustering, anomaly detection, regression, and classification results presented thus far have demonstrated that PMA is effective for identifying and characterizing degradation in microelectronic devices. However, each method has inherent limitations when the objective is to quantify the relative separation between degradation stages. Fully supervised methods, such as classification and regression, require labeled training data spanning multiple degradation stages and sufficiently large sample sizes. Semi-supervised approaches can quantify deviation relative to a reference population but do not directly provide a framework for comparing degradation stages to one another. Clustering algorithms can be used to estimate separation between groups using distances between cluster centroids; however, these distances depend strongly on the chosen algorithm, feature scaling, and hyperparameters. Furthermore, clustering results in high-dimensional feature spaces are difficult to relate to domain-relevant or physically meaningful characteristics.

To address these limitations, the feature-wise quantification method described in Section 4.7 was applied to the HTB dataset. This approach operates on individual engineered features and uses the feature distributions to identify classes and quantify their separation. Unlike clustering or PCA-based distance measures, this approach generates metrics derived from the feature domain, providing a consistent and comparable method across datasets.

An example implementation of this approach was performed using the spectral crest factor of the third spectral band of the PMA spectra, which was the third highest-ranked FS-PCA feature. The two highest-ranked features required logarithmic scaling for visualization and were therefore excluded from this example. A GMM was fitted to the normalized feature, which predicted four distinct distributions. The distribution of the features, along with the predicted classes, is provided in Figure 33. Across the ten highest-ranked features, the GMM predicted either three or four classes, indicating consistent multimodality in the feature space.




Figure 34: Distribution of a single engineered feature with classes predicted using a GMM model. Vertical dashed lines indicate the median of each class.

Figure 34 includes two complementary plots. The left plot shows the GMM predicted classes with the ground truth of each sample identified by color, while the right plot shows the distribution of the feature along with the median of each stage. The GMM accurately separated the 0-, 46-, and 136-hour degradation stages, but was unable to correctly distinguish between the later degradation stages. This is consistent with the overlap and minimal difference in the medians of the 297-, 520-, and 1014-hour degradation stages observed in the right plot. While this approach does not

correctly identify all six classes, the same was true for clustering performed using a single feature. These later degradation stages are easily distinguished from the 0-hour data but not from each other using a single feature. Importantly, this method demonstrates why later stages of degradation are challenging to distinguish by showing the overlap in their feature distributions.




Figure 35: Distribution of a single engineered feature showing GMM-predicted class distributions and class medians (left), and the same feature with medians based on ground-truth labels (right).

Although one-dimensional in formulation, this approach enables quantification of the separation between degradation stages using engineered features. Separation may be quantified using the distance between the median of each class or relative to a reference distribution. For example, the orange class in Figure 33 is separated from the green class by a difference of approximately 0.01 in spectral crest factor. This basic distance measure can be extended by evaluating the feature distributions using statistical separation metrics such as Ashman's D or Cohen's d.

To quantify separation between the later degradation stages, the ground-truth labels were used instead of the GMM-predicted labels to compute Ashman's D and Cohen's d for pairwise comparisons between the 297-, 520-, and 1014-hour stages. The

results are summarized in Table 12. The separation between the 297- and 1014-hour yielded an Ashman's D of 1.30, below the threshold of 2.0 for well-separated classes. A Cohen's d of 1.84 was yielded for the same pairwise comparison, indicating strong but overlapping separation. The remaining pairwise comparisons resulted in even smaller values, signifying poor separation using a single feature.

These results quantitatively demonstrate that later degradation stages are distinguishable relative to early stages but are not well separated from one another using any single feature. This explains why accurate separation of these degradation stages was achieved only when multiple features were used, as previously demonstrated in the regression, classification, and K-means clustering results.

Table 12: Ashman's D and Cohen's d computed for pairwise comparison between 297-, 520-, and 1014-hour HTB degradation stages.

Class A	Class B	Ashman's D	Cohen's d
297 hours	520 hours	0.62	0.87
297 hours	1014 hours	1.30	1.84
520 hours	1014 hours	0.50	0.71

While hypothesis testing approaches such as ANOVA could be applied to assess whether two classes are statistically different, the resulting p-values would not provide a scalar measure of separation like Ashman's D or Cohen's d. Uncertainty bounds and acceptance thresholds could be incorporated into this approach by constructing expected distributions of features from a sufficient number of samples, analogous to the approach discussed for MD.

5.3.6 Qualitative Comparison of Resolution Modes and Shifts in PMA Spectra Over Time

The PMA spectra of the devices exposed to 1014 hours of HTB were acquired using both the low- and high-resolution modes of the Chiplytics PSA system. Healthy control devices were measured during the same PMA measurement interval. The low- and high-resolution PMA spectra obtained using a 0–3.5 V square wave are overlaid in Figure 35, with the 1014-hour devices shown on the left and the control devices shown on the right.




Figure 36: PMA spectra (0–3.5 V) of 1014-hour HTB degraded devices (left) and control samples (right) captured using the Chiplytics PSA system in both low- and high-resolution modes.

There is additional content presented in the high-resolution spectrum that is absent in the low-resolution spectrum due to its higher noise floor. These differences are present in the spectra of both the degraded and healthy devices. While the trough located at 15 MHz appears similar in both resolution modes, the trough near 20 MHz exhibits a lower amplitude in the high-resolution spectra, and the trough located near 35 MHz exhibits a substantially lower amplitude. These additional portions of the spectra, which are partially obstructed by the noise floor in the low-resolution spectra, may have provided increased sensitivity to degradation, potentially resulting in greater

separation of the 297-, 520-, and 1014-hour stages in addition to improved clustering quality.

In addition to amplitude differences, a frequency shift is present in both the degraded and control device spectra when comparing low- and high-resolution results. While the amplitude change and frequency shift may be partially attributed to the increased voltage resolution of the high-resolution mode, the impact of changes in the AD3 configuration required to achieve this mode cannot be ruled out. Further investigation is required to isolate the contributing factors and determine the extent of their influence.

PMA spectra of the control devices were acquired on 06-09 and 07-28 using the high-resolution mode of the Chiplytics PSA system. As discussed in Section 5.2, the DtD MAE gradually increased over the course of the HTB study, with a spike in MAE on 07-28. This behavior is observable in both the low- and high-resolution spectra but is more apparent in the high-resolution spectra due to the lower noise floor. The PMA spectra from 06-09 and 07-28 are overlaid in Figure 36, revealing a small frequency shift at approximately 20 MHz and 35 MHz. Although the cause of this shift remains unknown, it explains the spike in MAE and RMSE 07-28 relative to the first measurement interval.




Figure 37: PMA spectra (0–3.5 V) of control devices captured 49 days apart.

5.3.7 Qualitative Comparison of PMA Spectra of HTB and TID Exposed Devices

The PMA spectra of TC74HC02 devices exposed to HTB presented in this chapter, and TID, presented in Chapter 6, were qualitatively compared to assess whether different failure mechanisms produced distinct signatures within the PMA spectra. Visual inspection of the PMA spectra revealed that the 30 krad devices were the most similar to the 1014-hour degraded devices, while the 20 krad devices exhibited a difference in propagation delay between the control and irradiated devices that was most similar to the 1014-hour HTB devices. The PMA spectra of these three degradation conditions, along with the spectra of the control device measured during both studies, are shown in Figure 37.

The 1014-hour HTB and 30 krad spectra exhibit clear differences, including the location of peaks and troughs, as well as differences in the amplitude of certain peaks and troughs. Similarly, the 20 krad spectrum has a deeper trough near 19 MHz and lacks the trough near 15 MHz observed in the other two degradation spectra. These

results indicate that PMA has the potential not only to quantify degradation but also differentiate between degradation mechanisms.

The PMA spectra of the control devices, which were acquired months apart and at different test locations, appear nearly identical. Only minor amplitude differences are observed at a small number of frequencies. The consistency of the control spectra across studies further supports that the observation differences between the HTB- and TID-exposed device spectra arise from the applied stresses.




Figure 38: PMA spectra (0–3.5 V) of TC74HC02 devices after 1014 hours of HTB exposure, 30 krad of TID, 20 krad of TID, and control devices measured during both studies.

5.4 Conclusions

The presented HTB study was conducted to assess the capability of PMA to detect and quantify degradation in microelectronic devices. Because degradation is an unavoidable consequence of device operation and is accelerated by environmental and electrical stresses, this study evaluated PMA across multiple HTB degradation stages using quantitative analysis methods rather than qualitative inspection alone.

Prior to interpreting degradation-driven changes, the repeatability and stability of the Chiplytics PSA system were evaluated to ensure that observed trends were attributable to applied stresses rather than measurement uncertainty or device variability. Analysis of the null spectra, control devices, and repeated measurements demonstrated low measurement variability across all evaluated PMA voltage ranges. Although a gradual increase in DtD variation was observed over the course of the study, its magnitude was negligible relative to the separation between degradation stages, as demonstrated through clustering, regression, and classification results.

Multiple input datasets and PMA voltage ranges yielded perfect clustering accuracy, with engineered features consistently achieving higher clustering accuracy than PCA-based inputs. Nevertheless, PCA-based clustering achieved accuracies exceeding 90% for several PMA voltage ranges, indicating that PCA provides a strong baseline for degradation separation, but with reduced application relevance and feature traceability, limited comparability across datasets, and poorer compatibility with downstream quantification approaches. Dimensionality reduction techniques can provide improved visualization relative to engineered features, which may be beneficial for qualitative analysis of results.

Quantification of the separation between degradation stages and healthy reference data was performed using semi-supervised methods. These approaches indicated sensitivity to intermediate degradation stages and demonstrated PMA's ability to resolve progressive degradation rather than binary health states. These approaches could be further enhanced by defining acceptable deviation thresholds for anomaly detection in applications such as screening.

Supervised regression and classification models were used to benchmark PMA input datasets and conventional electrical measurements. Engineered PMA features consistently provided the strongest regression performance, whereas curve trace measurements consistently yielded the worst results. Classification performance was strong using engineered features; however, phase measurements achieved perfect classification accuracy when used with the RF and SVC models. These results demonstrate the sensitivity of PMA and the engineered features to device degradation.

Conventional electrical testing in the form of curve trace and device timing characterization was performed in addition to impedance sweep measurements to complement the PMA analysis. Curve trace measurements and impedance sweeps did not provide clear separation between degradation stages; however, they confirmed exposure-dependent changes in the devices' electrical behavior. Propagation delay and transition time measurements revealed increases on both the rising and falling-edges of the first gate of degraded devices relative to controls. The results are consistent with combined NBTI- and HCI-driven degradation and further validate that observed changes in the PMA spectra were the result of device degradation.

Collectively, the feature extraction methodology and engineered features provided improved performance relative to PCA-based approaches, while remaining dataset-independent, traceable, and application-relevant. Feature-wise quantification using Ashman's D and Cohen's d provided insight into why separation of later degradation stages was more challenging, complementing clustering and regression analyses.

Comparison of results acquired using the low- and high-resolution modes of the Chiplytics PSA system suggests that separation of the later degradation stages may have been more significant if the high-resolution mode had been employed throughout the study. This comparison also revealed minor differences between the low- and high-resolution PMA spectra that may be attributable to AD3 configuration rather than device degradation.

Finally, the PMA spectra of devices exposed to 1014 hours of HTB and 30 krad of TID were qualitatively compared. Observed differences between the spectra suggest that PMA may be capable of distinguishing between dominant failure mechanisms in addition to quantifying stages of degradation.

Overall, the HTB study results demonstrate PMA's ability to repeatably and sensitively differentiate between multiple stages of device degradation. The inclusion of engineered features, derived using the developed feature extraction methodology, improves feature traceability and application relevance, and provides dataset-agnostic representations, while simultaneously enhancing clustering, regression, and classification performance. This supports the broader applicability of PMA for degradation analysis of microelectronic devices.

Chapter 6: Detection and Quantification of Degradation: Total Ionizing Dosage

Exposure to ionizing radiation has a wide range of effects that can negatively impact the reliability and performance of microelectronic devices. This issue is particularly critical in space applications, where exposure to radiation is unavoidable and must be accounted for during design and qualification. Shielding and radiation hardening provide devices with a layer of protection, but they cannot fully negate the effects of ionizing radiation. Testing devices exposed to ionizing radiation is therefore vital to ensure reliability and to verify that radiation hardening or shielding provides the intended level of protection [168]. In recent years, there has been increased demand for COTS devices for space applications to reduce costs and meet the demand for LEO missions, which are increasing at a near-exponential rate.

Traditional radiation assurance screening exposes devices to ionizing radiation dosages that are two to three times those anticipated to be encountered during the course of the mission [169]. Electrical testing of devices and recommended practices are outlined in standards such as MIL-STDs and JEDEC standards. Typically, a bias voltage is applied to devices during ionizing radiation exposure to increase the effective electric field across the dielectric regions, which maximizes the effects of radiation by causing increased oxide and interface trap buildup [170]. Unbiased radiation exposure is less severe in most cases. This is why NASA's radiation testing and much of the radiation-related testing reported in the literature includes a bias voltage applied to devices during ionizing radiation exposure [171].

Finazzi *et al.* analyzed multiple LEO missions and found that the TID accumulated over 1100 days ranged from 500 rad to 1.9 krad [172]. Gutiérrez *et al.* outline the growing interest in utilizing COTS devices for LEO missions, noting that shielded satellites, such as the Jupiter icy moons explorer (JUICE), are expected to be exposed to approximately 1 Mrad over a 12-year mission. In contrast, LEO missions are comparatively shorter in duration and may be exposed to approximately 300 rad per year. Gutiérrez *et al.* further predicted that COTS devices may be used in future interplanetary or Moon-based missions, as TID on the order of 100 krad is achievable with COTS devices [173]. Consequently, following established radiation assurance practices, COTS devices can be screened using ionizing radiation dosages on the order of 200 to 300 krad for certain applications.

6.1 Experimental Design

Three different microelectronic devices were included in this TID radiation study. The first device was the DMP2104V-7, a P-channel metal-oxide-semiconductor field-effect-transistor (MOSFET) packaged in a SOT-563 (small outliner transistor) package. The second device was the Toshiba TC74HC02APF quad two-input NOR logic gate device in a 14-pin DIP package constructed using planar CMOS technology. This is the same device used in the HTB study presented in Chapter 5. Lastly, the LM741CN/NOPM 8-pin DIP bipolar operational amplifier was included.

PMA and curve trace measurements were acquired for each individual device prior to exposure to ionizing radiation. Five control devices of each device type were also measured. For each dosage level, four TC74HC02 devices, four LM741 devices, and four DMP2104V-7 devices were placed into an electrostatic discharge (ESD) bag

and irradiated within a Co-60 gamma cell. An additional four DMP2104V-7 devices were inserted into sockets on a custom PCB used to apply a voltage bias of -1.4 V to the gate of each device during irradiation. This bias voltage was selected because higher amplitude voltages could have drawn more current than the test equipment could reliably supply. This process was repeated using previously unexposed devices for each desired dose: 10 krad, 20 krad, 30 krad, 50 krad, 75 krad, 100 krad, and 15 Mrad. Exposures were performed in groups, with 10 krad, 20 krad, and 30 krad constituting the first group; 50 krad, 75 krad, and 100 krad constituting the second group; and 15 Mrad representing the final group. Four TC74HC02 devices and four LM741 devices were exposed to 250 krad, in addition to the other dosages.

Four devices per exposure level may be less than ideal depending on the size of the ionizing radiation effect on the PMA spectra. However, time and cost limitations coupled with the available number of samples prevented including additional devices per dose level. The presented results demonstrate that four devices per dose level were sufficient for the NOR logic gate devices across several PMA voltage ranges, and in select cases for the MOSFETs and OpAmps.

PMA and curve trace measurements were performed after each dose level. PMA was first performed using the Sandia National Laboratories BTE configuration with a spectrum analyzer, followed by curve trace measurements and subsequent PMA measurements using the Chiplytics PSA system. Control devices were measured after each radiation event, along with the exposed devices. Measurements of the first dose group were performed in the morning, while measurements of the second dose group were performed in the afternoon. The 15 Mrad dose group was measured on the

following day, as devices were required to remain in the gamma cell for approximately 17 hours to reach the desired TID. The 250 krad dose group was irradiated after the 15 Mrad exposure, and measurements for both dose levels were performed during the same measurement sessions.

The results presented here are exclusively those acquired using the Chiplytics PSA system. A 100 kHz square wave was used for the modulating input, and a sampling rate of 100 MSa/s was used for data acquisition. PMA was performed using the PMA voltage ranges listed in Table 13. Because the high-resolution mode of the Chiplytics PSA system was utilized, the maximum PMA voltage range used was limited to 0–4 V.

Table 13: PMA voltage ranges used for TID study.

$V_{\min} - V_{\max}$ (V)
–0.05–0.05
0.0–0.10
0.0–0.25
0.0–0.5
0.0–1.0
0.0–1.5
0.0–2.0
0.0–2.5
0.0–3.0
0.0–3.5
0.0–4.0

PMA measurements of the TC74HC02 devices were acquired by applying the modulating input to the device’s V_{cc} pin, while the device’s ground pin was connected to ground. PMA measurements of the DMP2104V devices were performed by applying the modulating input to the device’s ground, while the device’s source pin was connected to ground. Finally, PMA measurements of the LM741 devices were conducted by applying the modulating input to the device’s V^+ pin, while the device’s V^- pin was connected to ground.

The ZIF socket on the Chiplytics PSA system enables direct testing of DIP devices, such as the TC74HC02 and LM741 devices, without the need for an adapter or additional socketing. In contrast, the DMP2104V 563-SOT package required the use of an adapter and an additional socket. Furthermore, the relatively small size of the DMP2104V devices necessitated mounting them to test coupons for irradiation and testing. These test coupons were then inserted into a custom socket for measurement. An additional PCB populated with four of these sockets was used to apply the bias voltage during irradiation.

Curve trace measurements were performed during the same measurement intervals as PMA for all three device types. The TC74HC02 devices were measured by ramping the voltage applied to V_{cc} from -0.5 V to 7 V, while the GND pin was connected to ground. A 50 mA compliance was used. The socket utilized in Chapter 5 for curve trace and impedance sweep measurements was also used for curve trace measurements performed on the TC74HC02 devices in this study. The DMP2104V devices were measured by applying a bias voltage of -1.4 V to the device's gate, while the device's drain was swept from 0 V to -1.5 V with a compliance of 1 A. The device's source pin was connected to ground. The same socket used for PMA measurements was also used for curve trace measurements. The LM741 devices were measured by sweeping the voltage applied to V^+ from -1 V to 0.5 V, while V^- was connected to ground. A compliance of 2.8 mA was used. Each curve trace measurement consisted of 101 voltage steps, with a current measurement acquired at each step.

Additional functional testing was performed on all three device types at CALCE. The transition time and propagation delay of the first gate were measured for

a single TC74HC02 device from each dose level, five days after irradiation. Timing characterization was performed by applying 5 V to the device's V_{cc} pin, while one of the gate's inputs was switched at 1 kHz and the other was held at 0 V. This procedure is nearly identical to the timing characterization performed in Chapter 5.

The threshold voltage of the DMP2104V devices was measured using an Agilent B1505A seven days after irradiation, with additional PMA measurements performed concurrently to capture any annealing effects. Every device, both irradiated and unirradiated controls, was characterized. The DMP2104V's specification sheet defines threshold voltage as the point where the drain current reaches $-250\text{ }\mu\text{A}$, with a minimum of -0.45 V and a maximum of -1.0 V .

The closed-loop gain of the LM741 devices was measured seven days after irradiation. Each device was powered using a $\pm 12\text{ V}$ power supply, which applied $+12\text{ V}$ to V^+ and -12 V to V^- . A $1\text{ k}\Omega$ resistor was placed in series between the device's ground pin and the device's inverting input, while a $100\text{ k}\Omega$ resistor was placed in series between the device's inverting pin and the device's output pin. This configuration was used to achieve a gain of approximately 101. The gain of each device was then measured by applying a 100 mV DC signal to the non-inverting input, measuring the voltage of the output pin, and computing the ratio of the output voltage to the input voltage.

6.2 Results

The following subsections present measurement issues encountered during PMA testing of the DMP2104V MOSFETs, after which the results from the three device types are presented. The results for the DMP2104V-7 devices include both

biased and unbiased irradiated devices. Each device-type subsection presents the acquired PMA spectra, followed by clustering using multiple input datasets, after which several machine learning models are used for classification.

The PMA voltage range that yielded the highest clustering accuracy is presented in each subsection; however, results for all evaluated PMA voltage ranges are provided in Appendix B. The number of clusters for K-means was defined to match the number of classes, while the minimum cluster size for HDBSCAN was set to two. The number of principal components retained for clustering or classification was selected using the slope rule defined in Section 4.4.

The MD and MMD were trained using the PMA spectra of non-irradiated control devices. MLP, RF, and SVC models were used for classification. Each classification model was trained using various datasets, including the PMA spectra, engineered features, curve trace measurements, and PCA models fitted to the PMA spectra. Classification performance was evaluated using global classification accuracy and the F1 score.

Classification was performed using three-fold cross-validation, with each fold repeated four times for the DMP2104V devices, while four-fold cross-validation, with each fold repeated three times, was used for the TC74HC02 and LM741 devices. The cross-validation resulted in a total of twelve unique training and test splits, where the reported $\pm$ values represent the standard deviation across the cross-validation folds. The MLP network consisted of two fully connected hidden layers, a sixteen node dense layer connected to an eight node dense layer, both utilizing the ReLU activation function. Hyperparameter tuning for the SVC, with a linear kernel, was accomplished

using nested cross-validation with an inner grid search. Regression was not performed due to the limited number of samples per dose level.

6.2.1 Measurement and Socketing Considerations for DMP2104V-7 MOSFETs

A test coupon to which a DMP2104V MOSFET is soldered is shown in Figure 38. All tested DMP2104V devices were mounted to test coupons, which were inserted into one of four sockets mounted on a PCB board that provided electrical connections for applying a gate bias during irradiation. A separate PCB with a single socket was used to perform conventional electrical testing and was also used, along with a ZIF-to-BNC adapter, for PMA measurements on the Chiplytics PSA system. This socket and ZIF-to-BNC adapter are shown connected to the Chiplytics PSA system in Figure 39. Multiple coaxial cables with BNC or SMA connectors were required in addition to banana cables to electrically connect the test hardware. Similar configurations would be required to perform PMA on most surface-mount devices using the Chiplytics PSA system.




Figure 39: Photograph of DMP2104V soldered to a test coupon




Figure 40: Photograph of Chiplytics PSA system with external socket and BNC to ZIF adapter used to perform measurements of DMP2104V devices.

The PMA spectra of the five control devices, measured whenever one of the exposure groups was measured, were qualitatively compared to identify any observable variance over the course of the study. The average PMA spectra of the control devices for the DMP2104V devices are shown in Figure 40, while the TC74HC02 control devices and LM741 control devices are provided in Figure 41.

The initial PMA spectra of the DMP2104V devices, including those irradiated at later stages of the study, are shown in green and exhibit relatively high PtP variance while maintaining a similar overall spectral shape. The control devices measured in conjunction with the 10 krad to 30 krad exposure group are similar to the initial PMA spectra, with a localized discrepancy observed near 45 MHz. In contrast, the PMA spectra of the control devices acquired in conjunction with the 50 krad to 100 krad exposure group and the 15 Mrad exposure group exhibit a pattern that is clearly distinct from both the initial and the 10 krad to 30 krad exposure group. By comparison, the PMA spectra of the TC74HC02 and LM741 control devices remain consistent over the course of the study, as shown in Figure 41.




Figure 41: Average PMA spectra (0–2.0 V) of DMP2104V control devices measured throughout the study.




Figure 42: PMA spectra (0–1.5 V) of TC74HC02 control devices (left) and PMA spectra (0–3.5 V) of LM741 control devices (right) measured throughout the study.

These observed differences in the PMA spectra of the DMP2104V control devices can only be attributed to changes in the devices themselves, the measurement system, or the electrical configuration, including adapters and socketing. The control devices should exhibit minimal variations as they were not irradiated. The DMP2104V devices exhibited higher PtP variance, as demonstrated by visible error bars around the average PMA spectrum shown in Figure 42. However, this variance does not account for the observed changes in the number and location of peaks and troughs in the control device PMA spectra. Furthermore, the stability of the TC74HC02 and LM741 control device PMA spectra demonstrates that the Chiplytics PSA system itself did not

introduce the observed changes. Thus, the observed differences are attributed to the external socketing and adapter required to perform PMA on the DMP2104V devices.




Figure 43: Mean (circles) and ± 1 standard deviation (error bars) of PMA spectra initially acquired for all device types.

PMA measurements performed using the Chiplytics PSA system required the devices, external socket, and the ZIF-to-BNC adapter to be physically relocated to another laboratory. After PMA measurements using the Chiplytics PSA system were completed, the hardware was transferred back to the original laboratory. This relocation was repeated for each measurement interval. It is therefore feasible that one or more electrical connections on the external socket loosened during transport, resulting in the observed spectral differences. To facilitate a robust assessment of the results, the primary PMA results for the DMP2104V devices were acquired at CALCE approximately one week after irradiation, without relocation of the test hardware during measurements. The PMA results captured at Sandia National Laboratories are provided in Appendix B.

The same external socket used to perform PMA measurements of the DMP2104V devices was also used to acquire curve trace measurements of the same

devices. Notably, the issues observed in the PMA spectra of control devices are not observed in the curve trace measurements. The initial curve trace measurements of all tested devices, shown in Figure 43, exhibit PtP variability similar in appearance to that observed in the PMA spectra. However, the curve trace measurements performed on the DMP2104V control devices remain consistent with the initial measurements. Curve tracing relies on DC voltages, whereas PMA utilizes a 100 kHz square wave. As a result, minor changes in the external socket or electrical connections are more likely to influence the frequency response of the circuit rather than the DC behavior, provided electrical continuity is maintained.




Figure 44: Curve trace measurements of all DMP2104V devices prior to irradiation, including control devices measured throughout the study.

6.2.2 DMP2104V-7 MOSFET Results

The PMA spectra of the unbiased DMP2104V devices at each dose level are presented in Figure 44, while the biased devices are shown in Figure 45. Both presented spectra were captured using a 0–2.0 V square wave. The unbiased and biased PMA spectra appear similar across the dose levels. The 15 Mrad PMA spectra are distinctly different from one another, with differences in the number and location of peaks and

troughs. Notably, the bias and unbiased PMA spectra for dose levels other than 15 Mrad exhibit observable differences near 17 MHz, 30 MHz, and 45 MHz, which are most apparent in the 100 krad spectra. Each dose level can be visually separated when they are color-coded. Furthermore, there is a consistent change from unirradiated to 75 krad in both the bias and unbiased PMA spectra. The 100 krad PMA spectra follow the trend observed in the lower dose levels, with the exception of the amplitude of the trough near 17 MHz, where the amplitude increases relative to the 75 krad spectra.

There are only three unbiased and three biased devices for the 100 krad dose level, rather than the four devices for every other dose level. This is because two devices failed prior to the study. This condition was only identified after all PMA measurements were complete and the curve trace results were analyzed.

K-means clustering was performed using principal components from a PCA model fitted to the PMA spectra. A clustering accuracy of 78.1% was achieved for both the biased and unbiased devices. The clustering results for the 0–2.0 V PMA spectra are provided in Figure 44 for the unbiased devices, and Figure 45 for the biased devices, along with the corresponding average PMA spectra for each dose level.

The 15 Mrad cluster is easily separable from the other dose levels for both biased and unbiased results. For the unbiased devices, 50 krad, 75 krad, and 100 krad dose levels are largely separated, with a 75 krad device misclustered with the 100 krad devices and a 50 krad device misclustered with the 30 krad devices. In contrast, for the biased devices, 50 krad, 75 krad, and 100 krad device dose level clusters visually overlap, with the 50 krad devices correctly grouped and three of four of the 75 krad devices correctly grouped. Similarly, the 20 krad and 30 krad dose levels are grouped

together, while the unirradiated and 10 krad are correctly grouped, with the exception of a single 10 krad device.




Figure 45: Average PMA spectra (0–2.0 V) of unbiased DMP2104V devices at different dose levels (left) and PCA-based K-means clustering results (right).




Figure 46: Average PMA spectra (0–2.0 V) of biased DMP2104V devices at different dose levels (left) and PCA-based K-means clustering results (right).

PMA results acquired at Sandia National Laboratories were also evaluated. A K-means clustering accuracy of 87.5% was achieved using PCA for the unbiased PMA spectra, while an accuracy of 75.0% was achieved for the biased devices. It is not known whether annealing impacted these results. However, the shift in the PMA spectra that resulted in better separation of certain dose levels is consistent with the observed differences in clustering accuracy. Notably, this improvement was observed only for the unbiased results, while the biased devices exhibited a lower clustering accuracy compared to the CALCE-acquired data.

For the CALCE-acquired 0–2.0 V PMA spectra, the FS-PCA features and FS-SCC features were also evaluated using K-means clustering. The FS-PCA features achieved a clustering accuracy of 81.2%, which is greater than the accuracy achieved

using PCA. In contrast, the FS-PCA feature extracted from the biased PMA spectra achieved an accuracy of 75.0% compared to 78.1% achieved using PCA.

LDA was used to simulate ideal feature extraction and achieved a peak clustering accuracy of 90.6% for the unbiased devices and 96.9% for the biased devices for a PMA voltage range of 0–2.0 V. The maximum PCA-based clustering accuracy was determined by evaluating the clustering accuracy using differing numbers of retained principal components. The maximum clustering accuracy for the unbiased devices was 81.2% using seventeen retained principal components, and 78.1% using five components for the biased devices.

For the unbiased devices, the engineered features achieved lower clustering accuracy overall compared to PCA. On average, the FS-SCC features achieved $0.5\% \pm 4.8\%$ decreased clustering accuracy compared to PCA, while FS-PCA features achieved a $1.4\% \pm 6.0\%$ decrease. In contrast, ideal feature selection using LDA increased the clustering accuracy across all PMA voltage ranges by $18.5\% \pm 7.0\%$ on average.

For the biased devices, the FS-SCC features decreased clustering accuracy by $2.3\% \pm 7.1\%$ on average compared to PCA, while the FS-PCA features had a minimal difference in clustering accuracy of $0.0\% \pm 5.2\%$ on average across the evaluated PMA voltage ranges. The LDA-based features increased clustering accuracy across all PMA voltage ranges by $15.6\% \pm 7.1\%$ on average. These results show that the engineered features performed similarly or better than PCA, particularly when ideal feature selection was used.

Clustering was also performed using both biased and unbiased devices simultaneously. Using the 0–2.0 V PMA spectra, PCA achieved a clustering accuracy of 67.8%, FS-SCC features achieved an accuracy of 62.7%, and FS-PCA features achieved an accuracy of 66.1%. Across all PMA voltage ranges, the FS-SCC features decreased clustering accuracy by $2.5\% \pm 4.7\%$ compared to PCA, while the FS-PCA features decreased accuracy by $7.1\% \pm 6.3\%$ on average. The LDA-based feature selection methods increased clustering accuracy for all evaluated PMA voltage ranges by an average of $10.6\% \pm 5.0\%$.

The observed overlap between several dose levels suggests that these dose levels are too similar to be differentiated. Accordingly, multiple dose levels were binned together and treated as a single group. The 10 krad, 20 krad, and 30 krad devices were binned together, as were the 50 krad, 75 krad, and 100 krad devices. This reduced the number of classes from eight to four. K-means clustering of the binned dose levels achieved an accuracy to 65.6% for the unbiased devices and 71.9% for the biased devices using PCA of the 0–2.0 V PMA spectra. The PMA spectra acquired at Sandia National Laboratories achieved 100% accuracy for both biased and unbiased devices.

For the CALCE-acquired results, the FS-SCC features and FS-PCA features achieved an accuracy of 65.6%, matching PCA. Notably, the FS-SCC and FS-PCA feature achieved perfect clustering for the biased devices when the dose levels were binned. This increase in clustering accuracy for the biased devices and decrease for the unbiased devices follows the predicted outcome based on visual inspection of the PMA spectra and K-means clustering plots. Clustering was also performed using

HDBSCAN, the results of which are provided in Appendix B, and achieved lower clustering accuracies compared to K-means.

The MD and MMD provide a relative distance between irradiated devices and the controls. The distribution of these distances is plotted using histograms, with the unbiased results provided in Figure 46, while the biased devices are provided in Figure 47. The vertical dashed lines represent the median of each dose level. The distances of the control devices, which were used for training, were also extracted and plotted.

The MD is shown on the left in both figures, while the MMD is shown on the right. The MD results show that all dose levels have an MD greater than 51.2, corresponding to the 95% chi-squared threshold for the reference population in this multivariate case. There is greater separation between the 15 Mrad devices and the other dose levels in the MMD results compared to the MD results. The 100 krad and 75 krad devices are also further separated from each other in the MMD results compared to the MD results. In contrast, the 10 krad, 20 krad, and 30 krad devices are more closely grouped in the MMD results than in the MD results. These observations suggest that neglecting the off-diagonal terms of the covariance matrix amplifies marginal feature differences, improving visual separation at certain dose levels while reducing grouping consistency among closely-related low-dose devices. While not explicitly evaluated in this study, interharmonic components may also contribute to the observed differences between the MD and MMD. Overall, the MD and MMD results reaffirm the clustering observations, while providing a more quantitative assessment of said observations.




Figure 47: MD (left) and MMD (right), trained on PMA spectra (0–2.0 V) of control devices, computed for unbiased DMP2104V devices across different dose levels. Vertical dashed lines indicate the median of each dose level.




Figure 48: MD (left) and MMD (right), trained on PMA spectra (0–2.0 V) of control devices, computed for biased DMP2104V devices across different dose levels. Vertical dashed lines indicate the median of each dose level.

In addition to PMA measurements, conventional electrical testing was performed, including curve trace and threshold voltage measurements. The hardware and electrical configuration for these conventional electrical tests are outlined in Section 6.1.

The curve trace measurements are plotted in Figure 48, with the unbiased devices shown on the left and the biased devices shown on the right. The curve trace measurement configuration was selected to mimic the drain current versus drain-source voltage curve within the DMP2104V specifications. A gate-source bias of -1.4 V was selected, as higher-amplitude voltages would have exceeded the current supply capabilities of the available test equipment, consistent with the bias voltage applied during irradiation.

The 50 krad, 75 krad, and 100 krad devices are easily distinguishable from one another for both biased and unbiased results. The unirradiated, 10 krad, 20 krad, and 30 krad curves are not fully separable, but they followed a consistent trend of decreased current with increased radiation exposure. Overall, these results align with the observed trends in the other PMA analysis results. However, the higher dose levels appear more separable in the curve trace measurements compared to other analyses.




Figure 49: Curve trace (I_d vs. V_{ds}) measurements of unbiased (left) and biased (right) DMP2104V devices at different dose levels.

As noted in Section 6.1, the threshold voltage was measured approximately one week after irradiation, while the curve trace measurements were performed within hours of exposure. Roughly half of the irradiated devices were provided to Sandia National Laboratories (SNL) for further failure analysis. Curve trace measurements were reperformed on the remaining half of the devices approximately one month and

one week after the original curve trace measurements were performed. The repeated curve trace measurements were performed at CALCE using the same external socket used at Sandia National Laboratories; however, the test equipment was different. Both sets of curve trace measurements are presented in Figure 49, with the unbiased results shown on the left and biased results shown on the right.

The measurements acquired at Sandia National Laboratories, indicated with solid lines, are distinct from those acquired at CALCE, indicated by dashed lines. The point where a MOSFET begins to enter saturation is where the slope changes from relatively flat to steep. This saturation onset has shifted to lower voltages for the curve trace measurements performed at CALCE compared to the ones performed at Sandia National Laboratories. Notably, the unirradiated control devices measured at Sandia National Laboratories exhibit distinct differences compared to when they were measured at CALCE. Because these devices are unirradiated, any observed differences can be attributed to the external socketing or differences in measurement hardware. It is not known why this difference is more pronounced for the unirradiated and lower dose levels. Nevertheless, this difference suggests that the observed discrepancies were not caused by annealing.




Figure 50: Curve trace (I_d vs. V_{ds}) measurements performed at both Sandia National Laboratories (SNL) and CALCE of unbiased (left) and biased (right) DMP2104V device. Each dose level is represented by a unique color.

The threshold voltage of each device was also measured. The DMP2104V has a specified threshold voltage range of -0.45 V to -1.0 V and defines the threshold voltage as the voltage at which the drain current reaches $-250\text{ }\mu\text{A}$. The measured threshold voltages were then plotted, with the biased, unbiased, and combined biased and unbiased devices being plotted separately. The three left plots (A, C, E) in Figure 50 include the 15 Mrad irradiated devices, while the three right plots (B, D, F) exclude the 15 Mrad devices and include a linear line of best fit. Horizontal dashed lines indicate the specified threshold limits for the DMP2104V devices.

The threshold voltage of both the biased and unbiased devices increased as the dose level increased, with the exception of the biased 15 Mrad devices for which the threshold voltage decreased to approximately 0 V . The unbiased 15 Mrad devices were found to have a threshold voltage of approximately -2.0 V . Both the biased and unbiased 15 Mrad devices exceeded the specification limits. These results indicate that the unbiased devices required approximately twice the specified voltage to switch the gate, while the gate of the biased devices remained open. Notably, two of the three unbiased 100 krad devices were within specification limits, while all three of the biased 100 krad devices exceeded the specification limits. The single unbiased 100 krad device exceeded the specification by 0.0005 V . Overall, the biased devices exhibited higher threshold voltages for a given dose level compared to the unbiased devices, which can be visually observed by comparing the circular markers to the star markers in the bottom-most plots of Figure 50. These results are consistent with ionizing-radiation-induced degradation of metal-oxide-semiconductor devices [170], [174], [175].




Figure 51: Threshold voltage of DMP2104V devices for (A) unbiased, (C) biased, and (E) combined biased and unbiased conditions plotted using logarithmic scale. Corresponding linear fits excluding the 15 Mrad dose level are shown for the (B) unbiased, (D) biased, and (F) combined biased and unbiased devices.

The line of best fit for both the biased and unbiased devices appears similar, with the unbiased device fit achieving an R^2 of 0.96, the biased device fit achieving an R^2 of 0.97, and the combined bias and unbiased device fit achieving an R^2 of 0.94. These results indicate a limited difference between the threshold voltage of the biased and unbiased devices. This observation is consistent with previous trends identified using clustering and MD/MMD analyses.

The unbiased classification results are presented in Table 14, while the biased results are presented in

Table 15. Both tables include all evaluated datasets. Although not presented in its entirety here, the classification results using the combined unbiased and biased devices are provided in Appendix B.

The MLP model was the worst-performing of the three classification models used across the unbiased, biased, and combined results. The harmonic spectra yielded the highest performance for the unbiased devices with an accuracy of 0.38 ± 0.13 , while the FS-SCC features performed similarly, yielding an accuracy of 0.34 ± 0.07 . The harmonic spectra also yielded the best results for the biased devices, with an accuracy of 0.42 ± 0.14 . Notably, the curve trace measurements yielded the second-best performance for the biased devices with an accuracy of 0.34 ± 0.04 , while the FS-SCC features yielded 0.14 ± 0.06 . Subsequent empirical evaluation indicated that increasing the number of nodes in each hidden layer improved classification performance; however, the MLP model remained the worst-performing overall.

The RF and SVC models achieved similar overall performance. Using the FS-SCC features extracted from the unbiased devices, the RF model achieved a classification accuracy of 0.72 ± 0.12 , while the SVC achieved an accuracy of 0.66 ± 0.14 . Both the RF and SVC models performed similarly using the FS-PCA features. Notably, the curve trace measurements yielded the highest classification accuracy for the RF model, achieving 0.82 ± 0.08 . Similarly, the SVC model performed best with the curve trace measurements, achieving an accuracy of 0.79 ± 0.09 . In both cases, the standard deviation was lower when using the curve trace data compared to PMA-based data. This trend was also observed for the biased devices; however, the SVC model outperformed the RF model for the biased devices.

The classification accuracy achieved when using the combined biased and unbiased devices was lower than that achieved using either bias condition individually, with a peak classification accuracy of 0.66 ± 0.10 . Unlike the individual bias conditions, the curve trace measurements did not yield the highest classification accuracy for the combined dataset; instead, the FS-PCA features and PCA achieved higher accuracy.

The limited number of samples per class, in combination with the aggressive k-fold cross-validation employed, may have limited the models' ability to learn within-class variance. A larger number of training and testing samples would be expected to improve classification performance.

As performed in the clustering analysis, multiple dose levels were binned together, and the RF and SVC models were reevaluated using the binned classes. The FS-PCA features yielded a classification accuracy of 1.0 for the unbiased devices for both the RF and SVC models. For the biased devices, the FS-PCA features yielded a classification accuracy of 0.97 ± 0.07 with the RF model and 0.96 ± 0.09 with the SVC model. The classification accuracy using the curve trace measurements increased to 0.95 ± 0.07 for unbiased devices and 0.88 ± 0.13 for biased devices. Binning the dose levels removed poorly separated classes and increased the number of samples per class.

Overall, the classification results demonstrate that PMA can accurately distinguish between different dose levels; however, additional testing is required to determine its sensitivity at lower dosages, particularly at 10 krad and 20 krad, in MOSFET devices.

Table 14: Classification model performance for unbiased irradiated DMP2104V devices using different input datasets, including features extracted from the PMA spectra (0–2.0 V).

Method	Input	Accuracy	F1
MLP	FS-SCC Features	0.34 ± 0.07	0.19 ± 0.08
	PMA Spectra	0.15 ± 0.06	0.08 ± 0.04
	PCA	0.16 ± 0.04	0.03 ± 0.01
	FS-PCA Features	0.29 ± 0.09	0.18 ± 0.07
	Harmonics Spectra	0.38 ± 0.13	0.28 ± 0.13
	Curve Trace	0.28 ± 0.08	0.13 ± 0.04
RF	FS-SCC Features	0.72 ± 0.12	0.70 ± 0.14
	PMA Spectra	0.55 ± 0.08	0.48 ± 0.09
	PCA	0.69 ± 0.10	0.66 ± 0.10
	FS-PCA Features	0.72 ± 0.13	0.70 ± 0.14
	Harmonics Spectra	0.64 ± 0.14	0.63 ± 0.13
	Curve Trace	0.82 ± 0.08	0.80 ± 0.07
SVC	FS-SCC Features	0.66 ± 0.14	0.63 ± 0.14
	PMA Spectra	0.77 ± 0.10	0.70 ± 0.12
	PCA	0.65 ± 0.14	0.62 ± 0.13
	FS-PCA Features	0.75 ± 0.13	0.72 ± 0.15
	Harmonics Spectra	0.72 ± 0.12	0.69 ± 0.12
	Curve Trace	0.79 ± 0.09	0.77 ± 0.09

Table 15: Classification model performance for biased irradiated DMP2104V devices using different input datasets, including features extracted from the PMA spectra (0–2.0 V).

Method	Input	Accuracy	F1
MLP	FS-SCC Features	0.14 ± 0.06	0.04 ± 0.03
	PMA Spectra	0.20 ± 0.10	0.12 ± 0.08
	PCA	0.15 ± 0.04	0.03 ± 0.01
	FS-PCA Features	0.18 ± 0.14	0.13 ± 0.11
	Harmonics Spectra	0.42 ± 0.14	0.34 ± 0.13
	Curve Trace	0.34 ± 0.04	0.20 ± 0.02
RF	FS-SCC Features	0.64 ± 0.13	0.56 ± 0.14
	PMA Spectra	0.57 ± 0.09	0.52 ± 0.11
	PCA	0.64 ± 0.13	0.58 ± 0.12
	FS-PCA Features	0.64 ± 0.07	0.58 ± 0.08
	Harmonics Spectra	0.67 ± 0.10	0.61 ± 0.11
	Curve Trace	0.80 ± 0.08	0.77 ± 0.12
SVC	FS-SCC Features	0.75 ± 0.10	0.71 ± 0.12
	PMA Spectra	0.73 ± 0.11	0.69 ± 0.11
	PCA	0.75 ± 0.06	0.68 ± 0.05
	FS-PCA Features	0.77 ± 0.11	0.72 ± 0.12
	Harmonics Spectra	0.71 ± 0.10	0.63 ± 0.09
	Curve Trace	0.77 ± 0.07	0.74 ± 0.07

6.2.3 TC74HC02APF NOR Logic Results

The PMA spectra for the TC74HC02 devices at each dose level, presented in the left plot of Figure 51, were acquired using a 0–1.5 V square wave. This voltage yielded the best clustering results and was therefore used for all presented analyses. The peak near 18 MHz decreases in amplitude with increasing radiation dosage until 250 krad and 15 Mrad, where the peak shifts to lower frequencies and increases in amplitude relative to the other dose levels. The behavior exhibited by all dose levels near 45 MHz is nonmonotonic, decreasing in amplitude from the unirradiated devices to the 50 krad, before the amplitude increases for the 75 krad and higher dose levels. The relative differences between the 50 krad, 75 krad, and 100 krad PMA spectra are smaller than those observed between the lower dose levels.

K-means clustering was performed using several input datasets across multiple PMA voltage ranges. A summary of the clustering results, including K-means and HDBSCAN, is provided in Appendix B. A plot of the K-means clustering results using the principal components from a PCA model fitted to the PMA spectra is presented in the right plot of Figure 51. Each dose level is separated from the others, with the unirradiated control devices plotted at the far right. There is an approximate linear trend from the unirradiated devices to the 30 krad devices along the first principal component axis. This trend transitions with the 50 krad to 15 Mrad dose levels, which follow a linear pattern along the second principal component axis. These observations may indicate that a transition in ionizing-radiation-induced degradation occurred between 30 krad and 50 krad. In contrast to the DMP2104V MOSFETs, the 15 Mrad dose level does not appear to be significantly different from the other dose levels.




Figure 52: Average PMA spectra (0–1.5 V) of TC74HC02 devices at different dose levels (left) and PCA-based K-means clustering results (right).

Perfect clustering accuracy was achieved using both K-means and HDBSCAN for PCA fitted to the 0–1.5 V PMA spectra. Mixed results were obtained with the engineered features, with the FS-SCC features yielding an accuracy of 75.7%, while the FS-PCA features yielded an accuracy of 100%. On average, the FS-SCC features decreased K-means clustering accuracy by $8.3\% \pm 14.3\%$ across the evaluated PMA voltage ranges compared to PCA. In contrast, the FS-PCA features increased the average clustering accuracy by $1.7\% \pm 9.9\%$ compared to PCA. Notably, the FS-PCA features increased clustering accuracy compared to PCA by 24.3% for the 0–0.25 V PMA spectra and 16.2% for the 0–3.5 V PMA spectra. Meanwhile, the clustering accuracy decreased when using the FS-PCA features compared to PCA by 8.1% for 0–0.5 V PMA spectra and 10.8% for 0–4.0 V PMA spectra.

When ideal feature selection using LDA was applied, the clustering accuracy increased by $17.6\% \pm 12.2\%$ on average across the evaluated PMA voltage ranges. HDBSCAN achieved higher clustering accuracy on average compared to K-means for the PCA-based results, with an average increase of $7.9\% \pm 9.2\%$. The opposite behavior was observed when using the FS-PCA features, with the average clustering accuracy decreasing by $14.5\% \pm 18.4\%$. The difference in feature-space dimensionality

may explain why the PCA-based input dataset performed better with HDBSCAN compared to the FS-PCA features.

The MD and MMD computed for the irradiated devices as well as the unirradiated controls, are shown in Figure 52, with the median of each dose level indicated by a vertical dashed line. The MD plot, shown on the left, and the MMD plot, shown on the right, are relatively similar. The MD results show that all dose levels have an MD greater than 51.2, corresponding to the 95% chi-squared threshold for the reference population in this multivariate case. Clear separation is observed between the unirradiated, 10 krad, 20 krad, and 30 krad devices in both plots, while the 50 krad to 15 Mrad devices are more closely grouped. The relative separation between the median of the higher dose levels is greater in the MMD results compared to the MD results. Notably, the median of the 15 Mrad devices is closer to the median of the unirradiated devices than that of the 50 krad, 75 krad, 100 krad, or 250 krad devices. The separation of the higher dose levels appears smaller using both MD and MMD compared to the clustering results.

While the MD and MMD results are similar overall, there are a few notable differences. The relative separation between the unirradiated and 10 krad devices is larger in the MD results than in the MMD results. The medians of the 50 krad and higher dose levels are more closely grouped in the MD results compared to the MMD results. These observations suggest that neglecting the off-diagonal terms of the covariance matrix may provide additional sensitivity to lower dose levels, which is consistent with the DMP2104V observations. However, the off-diagonal terms also

appear to decrease sensitivity to differences between higher dose levels, a behavior that was not observed in the DMP2104V results.




Figure 53: MD (left) and MMD (right), trained on PMA spectra (0–3.5 V) of control devices, computed for TC74HC02 devices across different dose levels. Vertical dashed lines indicate the median of each dose level.

The curve trace measurements, performed concurrently with PMA measurements, are shown in Figure 53. High PtP variance is observed for most dose levels, except for unirradiated and 10 krad devices. The PtP variance between devices from the same dose level becomes apparent above approximately 2 V, as does the separation between dose levels. The general pattern of each dose level appears to be consistent across devices, although the current draw and location of shoulders shift with increasing exposure. The unirradiated and 10 krad devices appear to overlap, each exhibiting smooth curves. The 20 krad dose level appears to include an outlier that is more similar to the higher dose levels. Excluding this outlier, there is a consistent increase in current draw from the unirradiated devices to the 30 krad devices.

For dose levels greater than 30 krad, shoulders begin to appear and become more pronounced above 50 krad. The location of these shoulders differs between dose

levels, with the shoulder position shifting to higher voltages as the dose level increases. Notably, the current associated with the peak of the shoulders increases until 15 Mrad, where the current draw decreases relative to the preceding dose levels.

Comparison to the PMA spectra, particularly the 0–4.0 V spectra, revealed similarly increased variance between devices exposed to the same dosage. Compared to the PMA spectra, the curve trace results appear more sensitive to higher dose levels.




Figure 54: Curve trace measurements of TC74HC02 devices at different dose levels.

Timing characterization, in the form of propagation delay and transition time measurements, was performed approximately one week after irradiation. The propagation delay and transition time of four unirradiated control devices and a single device from each dose level were measured. Because each of the four NOR gates within the TC74HC02 device is identical, it was assumed that each gate was equally affected by the ionizing radiation. As such, the timing characteristics of the first gate of each device were measured. The difference in the transition time of the irradiated devices relative to the unirradiated controls versus dosage is plotted on the left of Figure 54,

while the difference in the propagation delay of the irradiated devices relative to the unirradiated controls versus dosage is plotted on the right.

As noted in Chapter 5, the rising-edge transition time of the TC74HC02 often exhibited ringing, which reduced the accuracy and consistency of measurements. The transition time generally appears to increase for the rising-edge for irradiated devices. However, the unirradiated control devices exhibited variability in the measured transition times. For the rising-edge of input A switching, there was an outlier around 25 ns, while the remaining control devices were below 10 ns. For the rising-edge of input B switching, the control devices were closer together, but the minimum transition time was approximately 14 ns, and the maximum was 22.5 ns. Combined with the use of a single device characterized for each dose level, the observed increase in the transition time of the rising-edge with increasing dosage cannot be conclusively attributed solely to the effects of ionizing radiation.

The overshoot and ringing issues were not encountered while measuring the propagation delay. The propagation delay of the rising-edge increased with increasing dosage, while the propagation delay of the falling-edge remained constant or decreased. This timing asymmetry is consistent with CMOS devices exposed to ionizing radiation in the literature [174], [176], and provides independent validation of the degradation trends observed in the PMA results. While all devices remained functional, the propagation delay of the 250 krad and 15 Mrad devices exceeded the device-rated specification.




Figure 55: Difference in the transition time (left) and propagation delay (right) between irradiated devices and unirradiated controls at different dose levels. Plotted using logarithmic scale.

Classification was performed using several PMA-based datasets in addition to the curve trace measurements. A summary of the performance of the three evaluated classification models using the different input datasets is presented in Table 16.

The MLP model was the worst performing of the three, as was also observed for the DMP2104V MOSFETs. The FS-SCC features yielded a classification accuracy of 0.23 ± 0.10 ; however, the corresponding F1 score was lower at 0.13 ± 0.06 . Similar to the DMP2104V results, hyperparameter optimization yielded improved performance, although the MLP model remained the worst-performing model overall. The RF and SVC models performed significantly better, with both achieving perfect classification accuracy in select cases. The RF model performed better than the SVC when using the FS-SCC features, achieving an accuracy of 0.99 ± 0.03 , while the SVC achieved an accuracy of 0.68 ± 0.08 . The RF model achieved perfect classification using the FS-PCA features, while the SVC model achieved an accuracy of 0.99 ± 0.03 . Using an RBF kernel rather than a linear kernel improved the classification accuracy for several input datasets.

The harmonic spectra and curve trace measurements yielded perfect classification for both the RF model and the SVC models. The performance of the curve

trace measurements is consistent with the DMP2104V results. However, the relative performance gap between the harmonic spectra and the PMA spectra when using the SVC model is more pronounced than that observed for the DMP2104V devices. Overall, the RF model achieved the best classification performance, with perfect or near-perfect classification accuracy across all evaluated input datasets.

Table 16: Classification model performance of irradiated TC74HC02 using different input datasets, including features extracted from the PMA spectra (0–1.5 V).

Method	Input	Accuracy	F1
MLP	FS-SCC Features	0.23 ± 0.10	0.13 ± 0.06
	PMA Spectra	0.17 ± 0.12	0.10 ± 0.09
	PCA	0.11 ± 0.01	0.02 ± 0.0
	FS-PCA Features	0.14 ± 0.05	0.07 ± 0.02
	Harmonics Spectra	0.19 ± 0.05	0.08 ± 0.03
	Curve Trace	0.11 ± 0.01	0.05 ± 0.01
RF	FS-SCC Features	0.99 ± 0.03	0.99 ± 0.04
	PMA Spectra	0.98 ± 0.04	0.98 ± 0.06
	PCA	0.97 ± 0.05	0.96 ± 0.06
	FS-PCA Features	1.00 ± 0.0	1.00 ± 0.0
	Harmonics Spectra	1.00 ± 0.0	1.00 ± 0.0
	Curve Trace	1.00 ± 0.0	1.00 ± 0.0
SVC	FS-SCC Features	0.68 ± 0.08	0.60 ± 0.10
	PMA Spectra	0.78 ± 0.12	0.72 ± 0.16
	PCA	0.98 ± 0.04	0.98 ± 0.06
	FS-PCA Features	0.99 ± 0.03	0.99 ± 0.04
	Harmonics Spectra	1.00 ± 0.0	1.00 ± 0.0
	Curve Trace	1.00 ± 0.0	1.00 ± 0.0

6.2.4 LM741CN/NOPB OpAmp Results

The PMA spectra for the LM741 devices at each dose level, acquired using a 0–3.5 V square wave, are shown on the left of Figure 55. The 15 Mrad devices are visually distinct from the other exposures, which largely overlap. No additional peaks or troughs are observed that differentiate the dose levels, leaving only changes in the amplitude of the even harmonic components observed in the lower curve in the spectra.

K-means clustering was performed using several input datasets, with clustering results using the principal components from a PCA model fitted to the PMA spectra

shown on the right of Figure 55. The unirradiated through 100 krad devices generally cluster together, with three of the 250 krad devices visually separating. The 15 Mrad devices are clearly separated, as shown on the right side of the plot. Four of the nine clusters, represented by their centroids, include a single 15 Mrad device. This indicates that the relative intra-class variance of the 15 Mrad devices is greater than the variance between the other dose levels. While most of the dose levels group, there are multiple centroids presented. Closer inspection reveals the 50 krad, 75 krad, and 100 krad devices cluster together, while the 10 krad devices, three of the four 20 krad devices, the 30 krad devices, and one 250 krad device form a cluster. The unirradiated devices are grouped together with a single 20 krad device.

K-means clustering of the 0–3.5 V PMA spectra results using PCA yielded a clustering accuracy of 43.2%. In comparison, the FS-SCC features yielded an accuracy of 46.0%, while the FS-PCA features yielded an accuracy of 43.2% for the same PMA voltage range. LDA-based feature selection yielded an accuracy of 75.7%, representing an increase of over 30%, with a peak clustering accuracy of 78.4% achieved using LDA-based feature selection with the 0–2.5 V PMA spectra.

Across the evaluated PMA voltage ranges, the FS-SCC features increased the average clustering accuracy by $3.5\% \pm 3.3\%$ compared to PCA, while the FS-PCA features decreased the accuracy by $2.2\% \pm 4.7\%$. Notably, HDBSCAN achieved higher clustering accuracy on average compared to K-means, with an average increase of $4.3\% \pm 8.3\%$ using PCA. Similarly, HDBSCAN increased the average clustering accuracy by $14.5\% \pm 6.4\%$ using the FS-PCA features compared to K-means.

The K-means clustering was re-evaluated after binning the 10 krad, 20 krad, and 30 krad devices into a single group, and the 50 krad, 75 krad, and 100 krad devices into a second group. This reduced the total number of classes from nine to five. Using the FS-PCA features, a clustering accuracy of 78.4% was achieved for the 0–3.5 V PMA spectra, representing an increase of 34.8%. The FS-SCC features yielded the same accuracy. PCA yielded a marginal increase of 45.9% from 43.2%, with a maximum accuracy of 78.4% achieved when eight principal components were retained. Notably, the LDA-based feature selection achieved a maximum accuracy of 97.3%, approaching perfect clustering.




Figure 56: Average PMA spectra (0–3.5 V) of LM741 devices at different dose levels (left) and PCA-based K-means clustering results (right).

The MD and MMD were trained using unirradiated control devices, after which the relative distance between irradiated devices and the reference controls was computed. The MD results are shown on the left of Figure 56, while the MMD results are shown on the right. The MMD results are plotted on a logarithmic scale to improve visualization due to the large dynamic range observed between some dose levels.

The MD results show that all dose levels have an MD greater than 51.2, corresponding to the 95% chi-squared threshold for the reference population in this multivariate case. The medians, represented by vertical dashed lines, of most dose levels for both the MD and MMD are similar, with the exception of the 15 Mrad

devices. Modest separation is observed between the medians of the 10 krad, 20 krad, and 30 krad devices in the MD results. Both the MD and MMD results show separation between the medians of the 10 krad, 20 krad, and 30 krad devices and those of the 50 krad, 75 krad, and 100 krad devices. However, the distributions of these dose levels exhibit substantial overlap. The limited separation between dose levels is consistent with the K-means clustering results and analysis of the PMA spectra.

The 250 krad devices exhibit large spread, with several devices showing increased separation relative to prior dose levels. The 15 Mrad devices exhibit a similar spread in the MD result but appear compressed in the MMD results, which is attributable to the use of logarithmic scaling. Overall, both the MD and MMD demonstrate greater separation between the unirradiated and irradiated devices than is observed using K-means clustering.




Figure 57: MD (left) and MMD (right), trained on PMA spectra (0–3.5 V) of control devices, computed for LM741 devices across different dose levels. Vertical dashed lines indicate the median of each dose level.

Curve trace measurements were captured concurrently with PMA measurements. The results are presented in Figure 57, with the complete curve trace results shown on the left and a magnified view from 0.7 V to 1.8 V shown on the right.

No clear, monotonic trend is observed between the curve trace results and radiation dosage. While some devices exhibit slight shifts toward higher voltages, these shifts do not correlate consistently with dose level. Examination of the magnified region reveals two groups near 0.8 V. One group begins to decrease in current after 0.8 V, while the second group exhibits lower current and begins decreasing slightly earlier than the first. Both groups converge near 1.4 V, with the exception of the 15 Mrad devices, which remain offset until 2.0 V or more.

The 30 krad and 250 krad devices appear exclusively in the second group, whereas the other dose levels appear in both groups. Overall, the curve trace measurements do not exhibit a clear or systematic dependence on dose. Compared to PMA, the curve trace measurements appear to have reduced sensitivity to radiation-induced degradation.




Figure 58: Curve trace measurements of LM741 devices at different exposure levels (left) and the magnified curve trace from 0.7 to 1.8 V (right).

The gain of each device was measured approximately one week after irradiation. The measured gain values are shown in Figure 58. Additional unirradiated devices were also characterized to establish baseline variability. There is no observable change in gain between the unirradiated devices and the irradiated devices up to 250 krad, with an average gain of 100.3. In contrast, the gain increased to an average of 106.3 for the 15 Mrad devices. This increase is isolated to the highest dose level.

These results are consistent with other analyses performed, including PMA and curve trace measurements.




Figure 59: Closed-loop gain measurements of LM741 devices at different dose levels. Plotted using logarithmic scale.

These results are inconsistent with ionizing-radiation-induced degradation trends in bipolar devices reported in the literature, where the open-loop gain typically decreases with increasing radiation dosage [177], [178]. This discrepancy indicates that the measured increase in apparent closed-loop gain is likely reflective of changes in the LM741 devices' parametric shifts in the LM741 device's electrical characteristics rather than a true change in gain, such as shifts in the DC offset voltage or input bias current. For example, the input bias current corresponding to devices with an average gain of 100.3 is estimated to be approximately 700 nA, while the 15 Mrad devices, which exhibited an average gain of 106.3, correspond to an estimated input bias current of approximately 5.7 μ A. The increased current through the 100 k Ω feedback resistor would generate an output offset of approximately 0.6 V, producing the observed increase in measured gain. Additionally, open-loop gain is commonly measured instead of closed-loop gain, typically using an AC input signal rather than the DC signal

employed in this study. Direct measurement of the input bias current and offset voltage may have provided greater sensitivity to radiation-induced changes in the devices due.

Classification was performed using three models and several input datasets, including the 0–3.5 V PMA spectra, engineered features, and curve trace measurements. The classification performance of each evaluated model is summarized in Table 17.

Consistent with the DMP2104V and TC74HC02 results, the MLP model was the worst-performing classifier, achieving a maximum classification accuracy of 0.29 ± 0.07 using the PMA spectra. The linear SVC model performed better, achieving a peak classification accuracy of 0.61 ± 0.11 using the curve trace measurements, with the FS-PCA features yielding a comparable accuracy of 0.59 ± 0.08 . The RF model performed best overall, with the curve trace measurements yielding a classification accuracy of 0.80 ± 0.10 . The harmonic spectra achieved a lower but comparable accuracy of 0.70 ± 0.08 , while the FS-PCA features yielded an accuracy of 0.64 ± 0.06 .

Unlike the TC74HC02 and DMP2104V devices, the RF model exhibited reduced performance when using the FS-PCA features, performing better when trained directly on PMA spectra. In contrast, the SVC model performed better using the FS-PCA features compared to utilizing the PMA spectra. These results suggest that the degradation signatures present in the LM741 PMA spectra are weaker and less separable than those observed in the other devices studied, or that the LM741 devices are less sensitive to ionizing radiation. Additional features or alternative feature selection strategies may be required to improve classification performance for bipolar devices such as the LM741.

To further evaluate class separability, dose levels were binned, and the RF model using FS-PCA features was reevaluated. Following binning, the classification accuracy increased to 0.90 ± 0.13 , indicating that grouping poorly separable classes and increasing the number of samples per class substantially improves performance. Binning the curve trace measurements resulted in a classification accuracy of 1.0. These results are consistent with earlier clustering and MD/MMD analyses, which showed that many intermediate dose levels exhibited similar degradation behavior in the LM741 devices.

Table 17: Classification model performance for irradiated LM741 devices using different input datasets, including features extracted from the PMA spectra (0–3.5 V).

Method	Input	Accuracy	F1
MLP	FS-SCC Features	0.04 ± 0.05	0.01 ± 0.02
	PMA Spectra	0.29 ± 0.07	0.17 ± 0.07
	PCA	0.11 ± 0.01	0.03 ± 0.0
	FS-PCA Features	0.19 ± 0.05	0.11 ± 0.05
	Harmonics Spectra	0.18 ± 0.05	0.08 ± 0.04
	Curve Trace	0.19 ± 0.05	0.07 ± 0.02
RF	FS-SCC Features	0.58 ± 0.13	0.52 ± 0.13
	PMA Spectra	0.69 ± 0.15	0.63 ± 0.17
	PCA	0.44 ± 0.13	0.38 ± 0.14
	FS-PCA Features	0.65 ± 0.06	0.57 ± 0.06
	Harmonics Spectra	0.70 ± 0.08	0.63 ± 0.07
	Curve Trace	0.80 ± 0.10	0.75 ± 0.12
SVC Linear	FS-SCC Features	0.43 ± 0.06	0.37 ± 0.07
	PMA Spectra	0.50 ± 0.09	0.41 ± 0.09
	PCA	0.34 ± 0.11	0.25 ± 0.08
	FS-PCA Features	0.59 ± 0.08	0.51 ± 0.09
	Harmonics Spectra	0.53 ± 0.09	0.43 ± 0.10
	Curve Trace	0.61 ± 0.11	0.54 ± 0.12

6.3 Conclusions

The presented TID study was conducted to assess PMA's ability to detect and quantify ionizing-radiation-driven degradation in microelectronic devices, particularly for space applications. As the number of LEO missions continues to increase, so does the demand for low-cost, high-availability COTS devices. Accordingly, this study

evaluated PMA across multiple total ionizing radiation dose levels for three different device types: DMP2104V MOSFETs, TC74HC02 NOR logic gate devices, and LM741 OpAmps.

Radiation exposure using a Co-60 gamma cell, PMA measurements, and curve trace measurements were performed at Sandia National Laboratories, with additional PMA measurements performed at CALCE to address repeatability concerns observed for the DMP2104V devices. These inconsistencies were attributed to external socketing and adapter configuration, highlighting the challenge of achieving repeatable measurements with non-fixed hardware configurations and the importance of fixture control and stable electrical environments.

For the DMP2104V devices, both biased and unbiased irradiation conditions produced dose-dependent changes in the PMA spectra, with the higher dose levels consistently being separable from lower dose levels. Clustering performance was strongest using principal components, while FS-SCC and FS-PCA features yielded comparable results. Binning of dose levels increased clustering accuracy for the biased devices but degraded performance for unbiased devices. The MD results were consistent with the clustering results, while providing improved separation between unirradiated and irradiated devices. These results were consistent with the curve trace and threshold measurements.

The PMA spectra of the TC74HC02 devices exhibited a clear, monotonic trend with increasing radiation dose, particularly at the lower and intermediate dose levels. Clustering results demonstrated strong linear degradation trends between multiple dose levels, while engineered features provided a small but consistent improvement in

clustering accuracy compared to PCA. The MD exhibited enhanced sensitivity to lower dose levels compared to the MMD. While the curve trace measurement did not exhibit a consistent dose-dependent trend, propagation delay measurements independently validated the PMA observations, showing increasing delay with increasing dose.

In contrast, the LM741 devices exhibited limited visual separability across all dose levels except 15 Mrad. Clustering performance using K-means was generally low but improved when dose levels were binned and when LDA-based feature selection was applied. Conventional electrical measurements showed minimal dose-dependent trends, with the exception of 15 Mrad.

Classification results were consistent with clustering trends. RF models consistently outperformed MLP models across all device types. Engineered features generally yielded strong classification accuracy, although curve trace measurements yielded the highest classification accuracy for the DMP2104V and LM741 devices. Binning of dose levels resulted in near-perfect classification accuracy in several cases, reinforcing that several dose levels exhibit similar degradation signatures and that limited per-dose sample size constrained model performance.

Overall, this study demonstrates that PMA can effectively complement conventional electrical test methods for radiation screening of microelectronic devices. PMA was particularly effective for CMOS devices, with strong performance observed for the MOSFET and logic gate devices. In contrast, the LM741 devices exhibited reduced sensitivity to ionizing radiation. These results support the use of PMA as a rapid screening technique for radiation effects in COTS devices intended for space applications.

Chapter 7: Screening for Counterfeit Microelectronic Devices

Counterfeit microelectronic devices pose a threat to the security, reliability, and safety of electronic systems. Under ideal conditions, microelectronic devices should be procured from the original manufacturer or an authorized distributor. However, this is not always possible, particularly due to device obsolescence. In such cases, counterfeit testing is used to detect and intercept suspect counterfeit devices from the supply chain. While standards-based test methods are effective, they can be slow to perform and often rely on inspector experience and subjective interpretation. Rapid screening methods are needed to provide a comprehensive analysis of all devices that complement existing standards-based testing approaches.

The capability of PMA for counterfeit screening was assessed during an Air Force Small Business Technology Transfer (STTR) project. The STTR focused on evaluating the Chiplytics PSA system's capability to meet the needs of the Air Force for counterfeit screening, as well as identifying what improvements were required to reach those goals. PMA performed using BTE was also evaluated during one portion of the larger counterfeit screening study. In this study, counterfeit screening was evaluated using devices from different manufacturers and with differing specifications, rather than known counterfeit samples. The work presented in this chapter extends beyond the original STTR objectives by incorporating engineered features extraction and evaluating additional PMA voltage ranges.

7.1 Experimental Design

Several microelectronic device types were selected to encompass a range of technologies and functionalities in order to demonstrate and validate the performance of the Chiplytics PSA system. Device types were also selected based on considerations of availability and cost, resulting in the selection of devices of lower complexity. Attempts were made to ensure that the date and lot codes for each part number were identical. Devices packaged as DIPs were favored, as they can be tested using the Chiplytics PSA system without the need for additional socketing.

Flip-flop devices, OpAmp devices, and NOR, OR, and NAND logic gate devices were selected. Two different part numbers with equivalent function, pin configuration, and electrical specification were obtained for each device type, with the exception of the OR logic gate devices. Devices with the same functionality differed only in operational temperature rating or manufacturer. A summary of the selected devices, including their function, package, manufacturer, and operational temperature rating, is provided in Table 18. Twenty-five samples of each part number were included in the study.

PMA measurements were performed using the Chiplytics PSA system and a BTE setup, as described in Chapter 3. The results presented in this chapter were acquired using a 100 kHz square wave and a sampling rate of 100 MSa/s. The focus of this chapter is on the measurements acquired using the Chiplytics PSA system, which was primarily operated in high-resolution mode. Although not presented here, PMA measurements were also performed using the Chiplytics PSA system operating in low-resolution mode for a subset of the evaluated part numbers and PMA voltage ranges.

The clustering accuracy was comparable for the TL072 OpAmp and flip-flops devices; however, the NOR logic gate devices exhibited lower clustering accuracy in low-resolution mode compared to the high-resolution results presented in this chapter.

Table 18: Devices included in counterfeit study.

Part Number	Function (Technology)	Manufacturer	Operating Temperature Range	Package
TC74HC02APF	NOR (CMOS)	Toshiba	−40–85 °C	DIP-14
SN74HC02N	NOR (CMOS)	Texas Instruments	−40–85 °C	DIP-14
CD74HC02E	NOR (CMOS)	Texas Instruments	−55–125 °C	DIP-14
SN74HC00N	NAND (CMOS)	Texas Instruments	−40–85 °C	DIP-14
TC74HC00APF	NAND (CMOS)	Toshiba	−40–85 °C	DIP-14
SN74HC32N	OR (CMOS)	Texas Instruments	−40–85 °C	DIP-14
SN74HC574AN	Flip-Flop (CMOS)	Texas Instruments	−40–85 °C	DIP-20
TC74HC574APF	Flip-Flop (CMOS)	Toshiba	−40–85 °C	DIP-20
TL072CP	OpAmp (Bipolar)	Texas Instruments	0–70 °C	DIP-8
TL072IP	OpAmp (Bipolar)	Texas Instruments	−40–85 °C	DIP-8
RC4558IP	OpAmp (BiCMOS)	Texas Instruments	−40–85 °C	DIP-8
NJM4558D	OpAmp (BiCMOS)	Nisshinbo Micro Devices	−40–85 °C	DIP-8

For the logic gate devices and flip-flops, the modulating input was connected to the device's power pin, while the device's ground pin was connected to ground. For the OpAmp devices, the modulating input was applied to V^+ while V^- was connected to ground. A summary of the PMA voltage ranges used for the Chiplytics PSA system is provided in Table 19.

Table 19: PMA voltage ranges used for counterfeit screening study.

$V_{\min} - V_{\max}$ (V)
0.0–0.10
0.0–0.25
0.0–0.5
0.0–1.0
0.0–1.5
0.0–2.0
0.0–2.5
0.0–3.0
0.0–3.5
0.0–4.0

Conventional electrical testing, including curve trace and functional testing, was not performed, as the focus of the counterfeit screening study was to evaluate PMA. The objective of standards-based counterfeit electrical testing is to verify whether a device is operating within its specification limits. The effectiveness of electrical testing in counterfeit detection, alongside other standards-based methods, was discussed by Richman *et al.* [7], who found that across five different counterfeit part numbers, only two exhibited electrical characteristics outside of their specification limits. Furthermore, all devices tested in this study were procured from DigiKey and were known to be functional, rendering standards-based electrical testing unnecessary for the objectives of this study.

7.2 Results

The following subsections present the clustering, MD, and MMD results for the different device types included in the study. Clustering was performed using both K-means and HDBSCAN, with K-means serving as the primary focus. PCA models fitted to the PMA spectra with and without interharmonic components, as well as engineered features, were used as inputs to the clustering algorithms. The number of clusters for K-means was defined to match the number of classes, while the minimum

cluster size for HDBSCAN was set to two. MD and MMD were trained using the PMA spectra of a single part number, and the relative distance between the training devices and other devices was computed. Additional results, including analyses performed across all evaluated PMA voltage ranges, are provided in Appendix C.

7.2.1 Logic Gates

Six different logic gate part numbers, including three NOR logic gates, one OR logic gate, and two NAND logic gates, were included. Two of the NOR logic gate part numbers, SN74HC02 and CD74HC02, were manufactured by TI, each with a different operational temperature rating, along with TC74HC02 NOR logic gate devices manufactured by Toshiba. SN74HC32 OR logic gate devices manufactured by TI were included as a comparison to the NOR devices, as the primary functional difference is the addition of an inverter stage within each logic gate. Lastly, two NAND logic gate part numbers, SN74HC00 manufactured by TI and TC74HC00 manufactured by Toshiba, were included.

The K-means clustering results obtained using principal components from a PCA model fitted to the 0–3.5 V PMA spectra of the NOR logic gate part numbers are shown on the left of Figure 59. The TI-manufactured SN74HC02 and CD74HC02 devices and the Toshiba-manufactured TC74HC02 devices are separated along the first principal component axis, while the two TI part numbers, SN74HC02 and CD74HC02, are separated along the second principal component axis.

PCA-based K-means clustering achieved perfect clustering accuracy for five of the ten PMA voltage ranges evaluated, including 0–1.0 V and PMA voltage ranges greater than 0–2.0 V. In contrast, PCA-based HDBSCAN achieved lower clustering

performance, with a peak clustering accuracy of 76.0%. Reperforming PCA-based clustering using the PMA spectra with the interharmonic components excluded resulted in minimal changes to the clustering accuracy. For K-means, the clustering accuracy decreased by 6.7% for the 0–0.1 V PMA spectra but increased by 1.3% for the 0–1.5 V PMA spectra. The impact on HDBSCAN performance was larger, with an average increase in accuracy of $2.3\% \pm 12.8\%$.




Figure 60: K-means clustering results using a PCA model fitted to PMA spectra (0–3.5 V) of NOR logic gate devices (left), and corresponding MD (middle) and MMD (right) trained on PMA spectra (0–3.5 V) of SN74HC02 devices. Vertical dashed lines indicate the median distance for each part number.

The FS-SCC features yielded comparable results to PCA-based clustering, with the average clustering accuracy decreasing by $1.1\% \pm 7.4\%$. Notably, the engineered features extracted from the 0–2.0 V PMA spectra yielded clustering accuracy that was 20.0% lower than PCA. Excluding this PMA voltage range, the FS-SCC features yielded higher clustering accuracy, with an average increase of 1.0%. The FS-PCA features yielded greater improvements, with an average increase in clustering accuracy of $4.0\% \pm 12.8\%$ compared to PCA. As with the FS-SCC features, the engineered features extracted from the 0–2.0 V PMA spectra yielded clustering accuracy that was 18.7% lower than PCA, but provided a 20.4% increase in clustering accuracy on average for the three lowest PMA voltage ranges. HDBSCAN also performed well

using the FS-PCA features, with an average increase in clustering accuracy of $47.2\% \pm 21.0\%$ compared to PCA.

LDA-based feature selection yielded perfect clustering accuracy for nine of the ten PMA voltage ranges, achieving an accuracy of 92.0% for the 0–2.0 V PMA spectra. Only two features were required to achieve perfect clustering accuracy using K-means clustering.

The MD and MMD were trained using the 0–3.5 V PMA spectra of SN74HC02 devices. The relative distance between the training devices and the other devices was computed. The distribution of these distances is presented using histograms in the middle (MD) and right (MMD) of Figure 59. The MD results show that all devices have an MD greater than 51.2, corresponding to the 95% chi-squared threshold for the reference population in this multivariate case. The MD results exhibit approximately equal distance between all three part numbers. In contrast, the MMD results exhibit a smaller relative distance between the SN74HC02 and CD74HC02 devices compared to the TC74HC02 devices. This is notable as both the SN74HC02 and CD74HC02 devices are manufactured by TI, whereas the TC74HC02 devices are manufactured by Toshiba. These results suggest that the off-diagonal covariance terms may contain information that enables improved separation between devices manufactured by the same vendor.

The MD and MMD were retrained using the TC74HC02 devices, and the relative distance between the training devices and the SN74HC02 and CD74HC02 devices was computed. Both the MD and MMD results separated the training and test

devices; however, the distances of the SN74HC02 and CD74HC02 devices overlapped. Despite this overlap, the medians of the two TI part numbers were visually separable.

Two NAND logic gate part numbers, SN74HC00 manufactured by TI and TC74HC00 manufactured by Toshiba, were assessed using clustering. The PCA-based K-means clustering results for the 0–3.5 V PMA spectra are shown on the left of Figure 60. The two part numbers are separated along the first principal component axis, while exhibiting within-device spread along the second principal component axis. K-means clustering achieved perfect accuracy across all ten evaluated PMA voltage ranges, using PCA models fitted to the PMA spectra with and without the interharmonic components, as well as the FS-SCC features and FS-PCA features. In contrast, HDBSCAN achieved an average clustering accuracy of 21.0% when using PCA, while the FS-PCA features yielded perfect clustering accuracy.

The MD and MMD were trained using the SN74HC00 devices; however, the results were insensitive to the choice of training part number. The distribution of these distances based on the 0–3.5 V PMA spectra is presented using histograms in the middle (MD) and right (MMD) of Figure 60. The MD results show that all devices have an MD greater than 51.2, corresponding to the 95% chi-squared threshold for the reference population in this multivariate case. Consistent with the clustering results, the two part numbers are well separated. Along with the NOR logic gate results, these observations indicate that differentiating between manufacturers is straightforward using PMA.




Figure 61: K-means clustering results using a PCA model fitted to PMA spectra (0–3.5 V) of NAND logic gate devices (left), and corresponding MD (middle) and MMD (right) trained on PMA spectra (0–3.5 V) of SN74HC00 devices. Vertical dashed lines indicate the median distance for each part number.

The three NOR logic gate part numbers, the two NAND logic gate part numbers, and one OR logic gate part number were combined into a single dataset. The K-means clustering results obtained using a PCA model fitted to 0–3.5 V PMA spectra are shown on the left of Figure 61. The two TI NOR logic gate part numbers, SN74HC02 and CD74HC02, are closely grouped near the lower left corner, with the two part numbers forming distinct clusters. The Toshiba NOR logic gate part number, TC74HC02, and the Toshiba NAND logic gate part number, TC74HC00, are also closely grouped, but are correctly clustered separately. This is notable, as the two part numbers implement different logic functions. In contrast, the TI-manufactured NAND logic gate devices are located near the top middle of the plot, while the OR logic gate devices are located near the bottom right. The separation between the OR logic gate devices and the remaining logic gate devices is consistent with differences in logic implementation and the associated number of transistors per gate.

K-means clustering achieved perfect accuracy using PCA models fitted to the PMA spectra across the same five PMA voltage ranges that yielded perfect clustering accuracy for the NOR logic gate results. Similarly, perfect clustering accuracy was achieved using PCA models fitted to the PMA spectra with the interharmonic

components excluded for the same six PMA voltage ranges that yielded perfect clustering accuracy for the NOR logic results. These results indicate that the NOR logic gate part numbers dominate the clustering results, while the remaining part numbers are readily separable.

The FS-PCA features yielded an average increase in clustering accuracy of $4.0\% \pm 6.5\%$ compared to using PCA, consistent with the NOR logic gate results. In contrast, the FS-SCC features decreased the average clustering accuracy by $5.1\% \pm 14.7\%$ compared to PCA. Unlike the NOR logic gate results, this decrease in clustering accuracy was not dominated by a single PMA voltage range; instead, clustering accuracy decreased by 12.6% on average across eight of the ten PMA voltage ranges. Ideal feature selection using LDA yielded an average increase in clustering accuracy of $9.7\% \pm 15.5\%$ compared to PCA. Relative to the maximum achievable accuracy using PCA by retaining additional principal components, LDA-based feature selection increased the average clustering accuracy by $4.9\% \pm 9.7\%$.

The MD and MMD distributions trained using the 0–3.5 V PMA spectra of the SN74HC02 devices are shown in the middle (MD) and right (MMD) of Figure 61. The MMD is plotted on a logarithmic scale due to the large dynamic range of the distances. All six part numbers are visually separable in both the MD and MMD results. In both cases, the OR logic gate devices exhibit the largest distance from the training devices, while the CD74HC02 devices are the closest. This behavior is consistent with the clustering results. Similar to the clustering results, the TC74HC02 and TC74HC00 devices are relatively close, while the SN74HC00 devices are a greater distance from the training devices. The MD results show that all devices have an MD greater than

51.2, corresponding to the 95% chi-squared threshold for the reference population in this multivariate case.




Figure 62: K-means clustering results using a PCA model fitted to PMA spectra (0–3.5 V) of logic gate devices (left), and corresponding MD (middle) and MMD (right) trained on PMA spectra (0–3.5 V) of SN74HC02 devices. Vertical dashed lines indicate the median distance for each part number.

The MD and MMD were retrained using different part numbers, and the relative distance between the training devices and remaining devices was computed. Using the TC74HC02 devices for training resulted in overlap between the CD74HC02 and SN74HC02 device distributions in both the MD and MMD results. The TC74HC00 devices were less separable from the SN74HC02 and CD74HC02 devices in both distance metrics. Using the SN74HC32 devices for training resulted in separation between the TC74HC02 and TC74HC00. However, the CD74HC02 devices remained close to, but separable from, the TC74HC02 and SN74HC02 devices, whose distributions overlapped. This behavior is notable, as the logic gate devices exhibited strong separability between manufacturers across most other clustering and MD/MMD results. Outside of the mixed logic gate device MD and MMD results, the logic gate devices consistently exhibited strong separability by manufacturer and operating temperature rating across multiple PMA voltage ranges.

The influence of harmonic and interharmonic components on the PMA spectra was also evaluated. The RMSE between the PMA spectra of the 25 samples for each

device type was computed. Across the six device types, the average RMSE was 0.405 dBV for the PMA spectra, 0.113 dBV for the harmonic subspectra, and 0.449 dBV for the interharmonic subspectra. On average, exclusion of the interharmonic components prior to performing PCA decreased clustering accuracy by $0.05\% \pm 9.58\%$, while increasing explained variance by 0.18 ± 0.14 , relative to using the PMA spectra. In contrast, removing the harmonic components yielded an average decrease in clustering accuracy of $11.36\% \pm 15.92\%$ and reduced the explained variance by 0.45 ± 0.27 , relative to the PMA spectra. These results demonstrate that the interharmonic components provide minimal additional discriminatory information, while simultaneously increasing variance within the dataset.

7.2.2 *Flip-Flops*

Two D-type flip-flop part numbers, the TI SN74HC574 and the Toshiba TC74HC574, were tested. The K-means clustering results obtained using a PCA model fitted to the 0–3.5 V PMA spectra are shown on the left of Figure 62. Similar to the NOR and NAND logic gate results, the two part numbers from different manufacturers are separated along the first principal component axis. The SN74HC574 devices form a compact cluster, while the TC74HC574 devices exhibit higher within-class variance along both plotted principal component axes. The source of this increased variance is not known; however, it may be attributable to state changes occurring during PMA measurements.

K-means clustering using PCA models fitted to the PMA spectra with and without the interharmonic components achieved identical clustering accuracy, with perfect accuracy obtained for eight of the ten evaluated PMA voltage ranges. These

results indicate that the interharmonic components did not influence PCA-based K-means clustering performance for the flip-flop devices. However, exclusion of the interharmonic components increased the explained variance by 0.14 ± 0.15 on average relative to using the PMA spectra. In contrast, excluding the harmonic components decreased the average explained variance by 0.41 ± 0.17 relative to using the PMA spectra. HDBSCAN achieved lower clustering accuracy, with an average accuracy of 45.2% when using PCA with interharmonic components and 27.0% when using PCA without interharmonic components.

The FS-SCC features and FS-PCA features yielded identical clustering accuracy, with both producing 4.0% lower clustering accuracy for the 0–1.5 V and 0–4.0 V PMA spectra compared to PCA. Perfect clustering accuracy was yielded for all other PMA voltage ranges using the engineered features. Unlike the PCA-based HDBSCAN results, the FS-SCC features and FS-PCA features increased the average clustering accuracy by $48.0\% \pm 20.1\%$.




Figure 63: K-means clustering results using a PCA model fitted to PMA spectra (0–3.5 V) of flip-flop devices(left), and corresponding MD (middle) and MMD (right) trained on PMA spectra (0–3.5 V) of SN74HC574 devices. Vertical dashed lines indicate the median distance for each part number.

The MD and MMD were trained using the 0–3.5 V PMA spectra of the SN74HC574 devices. The resulting distance distributions are shown using histograms in the middle (MD) and right (MMD) of Figure 62. The two part numbers are well

separated; however, several TC74HC574 devices exhibit smaller distances from the training devices than the remaining TC74HC574 devices, consistent with the separation observed in the clustering results. This behavior is more pronounced in the MMD results than in the MD results. Retraining the MD and MMD using the TC74HC574 devices reduced both the observed variability and the magnitude of the distances. The MD between the two part numbers decreased by approximately a factor of two, while the MMD decreased by more than two orders of magnitude. In either case, MD of reach devices was greater than 51.2, corresponding to the 95% chi-squared threshold for the reference population in this multivariate case.

7.2.3 OpAmps

Four different OpAmp part numbers were tested, including two 4558 BiCMOS OpAmps and two 072 bipolar OpAmps. One of the 4558 OpAmp part numbers, RC4558IP, is manufactured by TI, while the other, NJM4558D, is manufactured by Nisshinbo Micro Devices. The two TI-manufactured 072 OpAmp part numbers, TL072IP and TL072CP, differ in operating temperature rating.

K-means clustering was performed using a PCA model fitted to the 0–1.5 V PMA spectra, the results of which are shown on the left of Figure 63. Unlike the logic gate and flip-flop results, the two 4558 OpAmp part numbers are separated along the second principal component axis. The bipolar and BiCMOS OpAmp part numbers are separated along the first principal component axis, with two 072 OpAmp part numbers overlapping. Despite this visual overlap, a clustering accuracy of 82.0% was achieved. Clustering using PCA fitted to the PMA spectra with the interharmonic component excluded had minimal effect, resulting in changes to clustering accuracy across four

PMA voltage ranges with an average decrease of $0.1\% \pm 2.5\%$ compared to including the interharmonic components. Consistent with previous observations, excluding the interharmonic increased the average explained variance by 0.31 ± 0.21 relative to using the PMA spectra, while excluding the harmonic components decreased the explained variance by 0.31 ± 0.29 relative to the PMA spectra. HDBSCAN achieved substantially lower clustering accuracy compared to K-means, with an average decrease of $48.2\% \pm 23.4\%$ relative to PCA with interharmonic components and $49.9\% \pm 18.2\%$ relative to PCA without interharmonic components.

The FS-SCC features yielded an average decrease in clustering accuracy of $1.8\% \pm 14.8\%$ compared to PCA, while the FS-PCA features increased average accuracy by $0.7\% \pm 10.5\%$. The improvement obtained using the FS-PCA features was primarily isolated to the 0–2.5 V PMA spectra, with the accuracy decreasing or remaining unchanged for eight of the ten PMA voltage ranges. LDA-based feature selection yielded an average increase in accuracy of $9.2\% \pm 8.5\%$ compared to PCA, with a peak clustering accuracy of 93.0% achieved for the 0–1.5 V PMA spectra. Consistent with the results for the other device types, the FS-PCA features increased the average clustering accuracy obtained using HDBSCAN by $42.6\% \pm 17.7\%$.




Figure 64: K-means clustering results using a PCA model fitted to PMA spectra (0–1.5 V) of OpAmps (left), and corresponding MD (middle) and MMD (right) trained on PMA spectra (0–1.5 V) of TL072IP devices. Vertical dashed lines indicate the median distance for each part number.

The MD and MMD were trained using the 0–1.5 V PMA spectra of the TL072IP devices. The resulting distance distributions are shown using a histogram in the middle (MD) and right (MMD) of Figure 63. The MMD results exhibit no separation between the two 072 OpAmp part numbers, while the two 4558 OpAmp part numbers are separable. In contrast, the MD results exhibit clear separation between the two 072 OpAmp part numbers, with no visible overlap in their distributions. The difference between the MD and MMD results is attributable to the use of the full covariance matrix in the MD, including the off-diagonal terms, which are excluded in the MMD. The MD results show that all devices have an MD greater than 51.2, corresponding to the 95% chi-squared threshold for the reference population in this multivariate case. However, the TL072IP devices had a median of 54.3, which is relatively close to the threshold. Retraining the MD and MMD using either the NJM4558 or RC4558 devices resulted in no separation between the two 072 OpAmp part numbers.

The same clustering and MD/MMD analyses were performed with the two 4558 OpAmp part numbers and two 072 OpAmp part numbers evaluated separately. K-means achieved perfect clustering accuracy for the 4558 OpAmp part numbers across all evaluated PMA voltage ranges using PCA with the interharmonic components, PCA without the interharmonic components, the FS-SCC features, and the FS-PCA features. Consistent with the combined OpAmp results, HDBSCAN achieved an average clustering accuracy of 27.0% using PCA, while perfect clustering accuracy was achieved using the FS-PCA features. The MD and MMD results

exhibited separation between the two 4558 OpAmp part numbers, regardless of which part number was used for training.

The 072 OpAmp results exhibited more mixed behavior. PCA-based K-means achieved an average accuracy of $58.0\% \pm 5.8\%$, and PCA without the interharmonic components yielded a comparable $55.8\% \pm 3.6\%$. The FS-SCC features further decreased the average clustering accuracy by $3.0\% \pm 7.3\%$ compared to PCA, while the FS-PCA features provided a similar $4.0\% \pm 6.1\%$ decrease compared to PCA. On average, the LDA-based feature selection method yielded an increase in clustering accuracy of $20.0\% \pm 10.5\%$. Notably, the LDA-based approach yielded a peak clustering accuracy of 98.0% for the 0–1.5 V PMA spectra. This near-perfect clustering accuracy was achieved using only three engineered features, was 8.0% higher than the next-best LDA-based result, and 28.0% higher than the best PCA-based result.

The MD and MMD were trained using the TL072IP devices; however, the results were insensitive to the choice of training part number. Consistent with the combined OpAmp results, the MD results separated the two 072 OpAmp part numbers while the MMD results exhibited overlapping distributions. The distance distributions were identical to those presented in Figure 63, excluding the 4588 OpAmp devices, as the training samples were unchanged, and the distance of each device is computed individually. Removal of the interharmonic components reduced the relative distance between the 072 devices; however, all devices remained above the threshold of 23.5.

The separation between the TL072IP and TL072CP devices was further quantified using Ashman's D and Cohen's d, as described in Section 4.7. Pairwise comparison of a single feature, the odd harmonic third-order polynomial fit residual

sum of squares, corresponding to the first LDA-ranked feature, yielded an Ashman's D of 1.82 and a Cohen's d of 2.57. The computed Ashman's D is below the threshold of 2.0 commonly used to indicate strong separability, although it is close. This corresponds to a Fisher discriminant ratio of 3.3, for which no formal threshold is defined. The Cohen's d value of 2.57 indicates very strong separation, with no overlap between the two distributions.

Texas Instruments was contacted to investigate potential manufacturing or design differences between the TL072IP and TL072CP part numbers. TI confirmed that the two part numbers are identical, including the die. The only difference between the tested devices was their date code, which differed by four weeks. This confirmation explains why the majority of the results exhibited limited or no separation between the two part numbers. Two methods successfully separated the two part numbers: the MD and LDA-based feature selection applied at a single PMA voltage range.

While these results demonstrate the high sensitivity of PMA, they also highlight a potential challenge for its adoption as a counterfeit screening method. Two sets of authentic devices manufactured on different dates should pass screening. The results presented here suggest that PMA may incorrectly identify such devices as suspect, increasing the risk of false positives. However, both methods that successfully separated the two part numbers require at least one labeled class, which is uncommon in practice for counterfeit screening and testing scenarios.

7.3 Hardware Variance

A study was performed to assess the influence of hardware variance and to compare the clustering performance of the Chiplytics PSA system and BTE. Testing

was performed by CALCE and Chiplytics, with each organization using its own Chiplytics PSA system and BTE setup. A summary of the hardware configurations and relevant specifications is provided in Table 20.

The most notable differences between Chiplytics' and CALCE's oscilloscopes is the maximum sampling rate. Sampling rates in the gigasamples-per-second range are not required for PMA because the frequencies of the modulating waveform are typically less than 100 kHz. Modern digital oscilloscopes utilize an analog-to-digital converter (ADC) to convert the analog input into digital values that can be recorded. The choice of an ADC often requires a tradeoff between sampling speed and bit depth. Faster ADCs, which enable higher sampling rates, are often limited in bit depth, resulting in lower resolution, whereas slower ADCs can often achieve high bit depths. Both benchtop oscilloscopes have a vertical resolution of 8 bits, corresponding to a voltage resolution equal to the full-scale voltage divided by 256 (2^8). In contrast, the Chiplytics PSA system has a vertical resolution of 14 bits, corresponding to a voltage resolution equal to the full-scale voltage divided by 16,384 (2^{14}). For the same full-scale voltage, the Chiplytics PSA system offers 64 times finer voltage resolution than the benchtop oscilloscopes.

Table 20: Relevant specifications of the oscilloscopes and waveform generators used in the CALCE and Chiplytics BTE setups and the Chiplytics PSA system.

Specification	Oscilloscope Specifications		
	Chiplytics PSA system Oscilloscope	Chiplytics Rigol SD1102Z-E	CALCE Agilent MSO7014B
Input type	Differential	Single ended	Single ended
Vertical resolution	14-bit	8-bit	8-bit
Bandwidth	9 - 20 MHz	100 MHz	100 MHz
Maximum sampling rate	100–125 MSa/s	1 GSa/s	2 GSa/s
Specification	Arbitrary Waveform Generator (AWG) Specifications		
	Chiplytics PSA system AWG	Chiplytics Siglent SDG 1032X	CALCE Keysight EDU33211A
Max output frequency	10 MHz	30 MHz	20 MHz
Vertical resolution	14-bit	14-bit	16-bit
Square wave rise time	~20 ns	16.8 ns	8.4 ns

Differences in the frequency bandwidth of the oscilloscopes are also noteworthy.

The 100 MHz bandwidth of the benchtop oscilloscopes results in reduced attenuation of higher harmonic components from the square wave. In contrast, lower bandwidth results in earlier roll-off of higher-frequency components, which can reduce the information content of the resulting PMA spectrum. The Chiplytics PSA system has a maximum sampling rate that is one-tenth that of the Chiplytics oscilloscope and one-twentieth of the CALCE oscilloscope.

There are multiple differences between the arbitrary waveform generators (AWGs); however, the most notable difference is the square wave rise time. The rise time of the Chiplytics' AWG and the Chiplytics PSA system is similar and is approximately twice as long as the rise time of the CALCE AWG. Faster rise times can result in increased ringing and noise.

CALCE and Chiplytics used different wiring configurations for their respective BTE setups. The placement of the BNC tee and the lengths of the coaxial cables differed.

While these differences affect circuit impedance and reflections, they were expected to have minimal influence on the final results.

Twenty-five samples were tested by both CALCE and Chiplytics, with each sample of a given part number sharing the same date and lot code. The three quad 2-input NOR logic gate part numbers used in the counterfeit screening study were selected.

PMA was performed using a 100 kHz square wave and three PMA voltage ranges: 0–1.0 V, 0–1.5 V, and 0–3.5 V. Three principal components were retained from PCA models fitted to the PMA spectra for each PMA voltage range. Three principal components were retained for consistency across the hardware configurations, although the slope rule could have been applied. Clustering was performed using K-means with the three retained principal components, and the number of clusters was defined using the ground truth. This approach enabled a more direct comparison of the hardware configurations; however, the true number of clusters is rarely known in practical counterfeit screening applications.

Clustering performance was evaluated using clustering accuracy, computed from the predicted labels and the ground truth. Clustering quality was evaluated using the silhouette score, and the explained variance (EV) of the three retained principal components was recorded. The accuracy, silhouette score, and explained variance for the four hardware configurations across the three evaluated PMA voltage ranges are provided in Table 21.

Table 21: Clustering results for the four hardware setups.

Metric	Chiplytics PSA system		Benchtop equipment	
	Chiplytics	CALCE	Chiplytics	CALCE
	0–1.0 V			
Accuracy	100%	100%	61.3%	94.7%
Explained variance	0.34	0.40	0.28	0.23
Silhouette score	0.73	0.73	0.43	0.55
	0–1.5 V			
Accuracy	100%	98.7%	57.3%	53.3%
Explained variance	0.65	0.81	0.31	0.08
Silhouette score	0.69	0.75	0.49	0.46
	0–3.5 V			
Accuracy	100%	100%	82.7%	65.3%
Explained variance	0.70	0.55	0.62	0.52
Silhouette score	0.71	0.79	0.62	0.45

Perfect clustering accuracy was achieved by Chiplytics across all three PMA voltage ranges for PMA acquired using the Chiplytics PSA system. CALCE achieved similar results, obtaining perfect clustering accuracy with the exception of the 0–1.5 V PMA spectra, which yielded a clustering accuracy of 98.7%. Both CALCE and Chiplytics were able to separate the TI-manufactured devices from the Toshiba-manufactured devices using the Chiplytics PSA system, with CALCE misclustering a single CD74HC02 device. Similar clustering quality was achieved by both organizations, although CALCE’s results yielded higher silhouette scores. CALCE’s results also yielded higher EV for the 0–1.0 V and 0–1.5 V PMA spectra, while Chiplytics EV was higher for the 0–3.5 V results.

Both BTE setups performed worse than their Chiplytics PSA system counterparts. CALCE achieved a maximum clustering accuracy of 96.0% using the 0–1.0 V PMA spectra. In contrast, Chiplytics achieved a maximum accuracy of 82.7%

using the 0–3.5 V PMA spectra. Both CALCE and Chiplytics observed the lowest clustering accuracy using the 0–1.5 V PMA spectra. Chiplytics observed reduced clustering accuracy for PMA voltage ranges below the device minimum operating voltage, while the PMA voltage range above the minimum yielded the highest accuracy. In contrast, CALCE achieved the highest clustering accuracy using the lowest PMA voltage range and observed a fourfold increase in PtP variance between the lowest and highest PMA voltage range when using its BTE setup. Silhouette scores were lower in all cases compared to the corresponding Chiplytics PSA system results.

Several factors may have contributed to the observed differences between the CALCE and Chiplytics BTE results. The BTE setups used multiple feet of coaxial cabling, which increased circuit inductance and provided additional pathways for to enter the signal path. Differences in the BTE configuration between CALCE and Chiplytics likely contributed to variations in the magnitude and frequency of signal reflections. These effects were not observed in the Chiplytics PSA system results due to the short metallized traces and consistent hardware configuration. CALCE's AWG has a rise time that is approximately twice as fast as that of the Chiplytics AWG, which can influence overshoot and ringing and may have contributed to the higher variance observed by CALCE. Additionally, differences in grounding configuration between the Chiplytics PSA system and the BTE may contributed to the observed discrepancies. A more detailed investigation would be required identify the primary contributors.

Analysis of the CALCE results was repeated after the interharmonic components were removed. The Chiplytics PSA system-based results improved, achieving 100% clustering accuracy across all PMA voltage ranges. Notably, the BTE-based results also

improved, with the 0–1.5 V accuracy increasing to 94.7% and the 0–3.5 V accuracy increasing to 86.7%. This analysis retained three principal components; however, retaining two principal components was recommended by the slope rule for the 0–1.0 V and 0–1.5 V results, while six components were recommended for the 0–3.5 V results. Furthermore, using the slope rule recommended number of retained principal components for the PMA spectra that included interharmonics yielded an accuracy of 50.7% for the 0–1.5 V PMA spectra and 73.3% for the 0–3.5 V PMA spectra. These results are more consistent with those obtained by Chiplytics using their BTE setup.

Overall, the comparison of clustering performance using PMA acquired with different hardware configurations demonstrates the improved and more consistent clustering accuracy achieved using measurements acquired with the Chiplytics PSA system relative to BTE-based results. CALCE and Chiplytics achieved nearly identical clustering performance using their respective Chiplytics PSA systems; however, this consistency was not observed for the BTE results. The higher bandwidth of the BTE setups did not result in improved clustering performance, nor did the lower bandwidth of the Chiplytics PSA system lead to reduced clustering accuracy. The higher vertical resolution and shorter conduction path lengths of the Chiplytics PSA system likely contributed to the observed performance improvements. In contrast, the additional cabling and circuit impedance of the BTE appear to have resulted in increased interharmonic noise relative to the Chiplytics PSA system results.

These results highlight the challenges associated with adopting side-channel techniques that rely on laboratory-specific BTE setups, which may vary between

organizations. Standalone test platforms and/or calibration approaches for side-channel measurement systems provide a clearer path toward broader adoption.

7.5 Conclusions

The presented counterfeit screening study was conducted as an expansion of an STTR project focused on evaluating the capability of the Chiplytics PSA system to be used for counterfeit screening. As counterfeits adopt increasingly sophisticated techniques and the demand for microelectronic devices continues to increase, new methods are needed to rapidly screen for counterfeits. Accordingly, this study evaluated PMA as a screening technique using several device types and multiple part numbers with equivalent function, pin configuration, and electrical specifications to emulate counterfeit screening scenarios.

Overall, PMA demonstrated strong capability to differentiate between equivalent part numbers with differing manufacturers, technologies, or operating temperature ratings. Perfect or near-perfect clustering accuracy was achieved across multiple types of devices, while MD and MMD provided complementary quantitative separation between part numbers.

Logic gate and flip-flop devices exhibited consistent separability, particularly between manufacturers and/or logic functions. The NOR and NAND logic part numbers were successfully separated by manufacturers, and in the case of the NOR logic gate devices, by operating temperature rating. When all logic gate part numbers were combined, the NOR logic gate devices dominated the clustering structure; however, strong clustering performance and MD/MMD separation were still achieved

across several PMA voltage ranges. Flip-flop devices exhibited similarly strong separability.

The OpAmp devices exhibited more nuanced behavior. The two BiCMOS 4558 OpAmp part numbers exhibited perfect accuracy, whereas the bipolar 072 OpAmp part numbers reduced separability when combined with the 4558 devices. When evaluated independently, the 072 OpAmp part numbers were generally not separable using unsupervised clustering, with separation observed only when LDA-based feature selection or MD was applied. These results are consistent with TI's confirmation that the two part numbers are physically identical.

Across all device types, PCA-based K-means clustering provided a strong baseline for screening performance. Engineered features, particularly the FS-PCA features, improved clustering performance in several cases, while LDA-based feature selection consistently increased clustering accuracy using a small number of retained features. HDBSCAN performed substantially better using the FS-PCA features compared to PCA. Removal of the interharmonic components prior to using PCA resulted in minimal changes to the overall PCA-based clustering accuracy, while increasing the explained variance and reducing the input dimensionality. In contrast, removing the harmonic components and retaining the interharmonic components decreased clustering accuracy and reduced the explained variance, demonstrating that the interharmonic components can be removed with minimal negative effect for counterfeit screening.

MD and MMD results generally aligned with the clustering results, with MD demonstrating increased sensitivity in several cases, most notably for the 072 OpAmp

devices. The choice of training part number influenced the resulting distance distributions.

PMA demonstrated sufficient sensitivity to separate two authentic part numbers manufactured only weeks apart. This highlights both the sensitivity and the potential risk of false positives in practical screening applications using PMA. The methods that successfully achieved this separation relied on labeled data, which may not be available in real-world counterfeit screening applications.

Evaluation of hardware variance demonstrated that measurements acquired using the Chiplytics PSA system yielded more consistent and higher clustering performance than the BTE-based setups. Removal of the interharmonic components increased the BTE-based clustering accuracy, making results more comparable to those obtained using the Chiplytics PSA system. The results demonstrate that side-channel methods such as PMA are more reproducible and consistent when performed using dedicated test platforms. Collectively, these results demonstrate that PMA is a sensitive, non-destructive, and scalable technique for counterfeit screening and support its use as a rapid counterfeit screening tool.

Chapter 8: Qualitative Study on PMA Sensitivity to Controlled Electrical Perturbations and Physical Modifications

The influence of subtle changes in electrical characteristics or internal circuitry on a device's PMA spectrum is not well understood. The ability of PMA to detect degradation, ionizing radiation included effects, and differences between similar devices has been demonstrated in the preceding chapters. However, while the physical and electrical mechanisms driving degradation are generally understood, the impact of small, localized changes to a device on the PMA spectrum has not been studied. Accordingly, two studies were performed to investigate the effects of controlled electrical perturbations and physical modifications on microelectronic devices and to assess how these manifest in the PMA spectra.

8.1 Experimental Design

External electrical perturbations were performed using a combination of resistors and capacitors placed in series with the PMA modulating input and V_{cc} or between the inputs 1A and 1B of a TC74HC02 device. In both cases, the modulating input was connected to V_{cc} while the ground pin of the device was connected to ground. These changes simulate the impact of external configuration differences such as fixturing and wiring. The resistor-capacitor combinations placed in series with V_{cc} are listed in Table 22, while the combinations placed between inputs 1A and 1B are listed in Table 23. Testing was performed by placing the device, resistor, and capacitor on a breadboard along with jumper wires that connected to a BNC cable, which was connected to the Chiplytics PSA system using the ZIF-to-BNC adapter.

Physical modificationss were performed in multiple stages. First, ceramic devices were de-lidded using a hammer and chisel, after which PMA was performed. Second, a tungsten probe was used to displace a single wire bond, after which PMA was performed again. Lastly, the tungsten probe was used to intentionally mar a portion of the die surface. Four TI SN7414A bipolar inverters with Schmitt trigger inputs and a single TI SN7417 bipolar buffer were subjected to these internal modificationss.

Table 22: Resistor and capacitor combinations placed in series with V_{cc} for electrical perturbation study.

Resistance (Ω)	Capacitance (nF)										
	0	100	150	220	330	470	680	1000	2200	4700	10e3
0	-	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
1.5	✓	✓	-	✓	-	✓	-	✓	-	-	-
4.7	✓	✓	-	✓	-	✓	-	✓	-	-	-
10	✓	-	-	-	-	-	-	-	-	-	-
47	✓	-	-	-	-	-	-	-	-	-	-
100	✓	-	-	-	-	-	-	-	-	-	-
220	✓	-	-	-	-	-	-	-	-	-	-
330	✓	-	-	-	-	-	-	-	-	-	-
1000	✓	-	-	-	-	-	-	-	-	-	-

Table 23: Resistor and capacitor combinations placed in series between inputs 1A and 1B for electrical perturbation study.

Resistance (Ω)	Capacitance (nF)					
	0	100	150	220	330	470
0	-	✓	-	✓	-	✓
1.5	✓	-	-	-	-	-
4.7	✓	-	-	-	-	-

PMA was performed using the Chiplytics PSA system operating in high-resolution mode with a 100 kHz modulating input. The PMA voltage ranges used are listed in Table 24, although all presented results correspond to the 0–3.5 PMA spectra results. PMA was performed by applying the modulating input to V_{cc} while the ground pin was connected to ground.

Table 24: PMA voltage ranges used for electrical and physical modification studies.

$V_{\min} - V_{\max}$
0 – 0.5
0 – 1.0
0 – 1.5
0 – 2.0
0 – 2.5
0 – 3.0
0 – 3.5
0 – 4.0

8.2 External Electrical Perturbation Results

The use of the breadboard and additional wiring drastically altered the PMA spectra of the TC74HC02 device, as shown in Figure 64. The breadboard added additional capacitance, resistance, and inductance to the circuit, resulting in the observed changes. The ZIF-to-BNC adapter and additional wiring further compounded these effects. As such, all electrical perturbations must be compared to the PMA spectra acquired with the TC74HC02 installed on the breadboard. All PMA measurements were performed twice for each perturbation, and the average of the two measurements is presented.




Figure 65: PMA spectra (0–3.5 V) of a TC74HC02 device measured while inserted into a breadboard and into the ZIF socket of the Chiplytics PSA system.

The external electrical perturbations performed can be divided into three categories: resistor-, capacitor-only, and combined resistor-capacitor perturbations. Plots showing the acquired PMA spectra for all three categories are provided in Figure 65, with the resistor-only perturbations shown in plot (A), the capacitor-only perturbations shown in plot (B), and the resistor-capacitor perturbations shown in plot (C). The resistor-only results exhibit the highest variance of the three categories. The 1.5 Ω , 4.7 Ω , and 10 Ω resistors had minimal impact on the resulting PMA spectra and were comparable to the no-perturbation PMA spectra. Above these resistance values, the even harmonic components exhibited lower amplitude, while the odd harmonic response exhibited similar overall amplitude but fewer troughs. Even harmonic and interharmonic components above 20 MHz exhibited increased variance, with no clear trend related to the resistance value.

The capacitor-only perturbations exhibited lower variance compared to the resistor-only perturbations, with the PMA spectra of different capacitors behaving similarly. The addition of a capacitor in series with V_{cc} resulted in a decrease in the number of peaks and troughs below 25 MHz, along with additional ringing in the interharmonic components beginning around 17 MHz and increasing in amplitude with frequency. Above 25 MHz, the harmonic response becomes sporadic, with no clear trends. This behavior is partially attributed to the charging and discharging of the capacitor during PMA measurements, particularly between the first and second repletion, resulting in an average of two nonidentical spectra.

The resistor-capacitor perturbation PMA spectra are dominated by the capacitive contribution, as evidenced by the similarity to the capacitive-only PMA

spectra. These results highlight the sensitivity of PMA to the surrounding circuit configuration.




Figure 66: PMA spectra (0–3.5 V) of a TC74HC02 device with different electrical perturbations in series with V_{cc} , including (A) resistive-only changes, (B) capacitive-only changes, and (C) combined RC changes.

A similar set of perturbations was performed by placing a resistor or capacitor in series between inputs 1A and 1B, electrically connecting the two inputs. The resulting PMA spectra are provided in Figure 66, with the resistor-only perturbations shown on the left and capacitor-only perturbations shown on the right.

The smallest-valued resistors from the V_{cc} -based perturbations were used, both of which resulted in a more significant deviation from the no-perturbation PMA spectra compared to their effect when placed in series with V_{cc} . The 1.5 Ω resistor PMA spectrum exhibits an additional peak near 11 MHz and reduced amplitude of the trough near 15 MHz. In contrast, the 4.7 Ω resistor PMA spectrum does not exhibit an additional peak near 11 MHz and reduces the amplitude of the trough near 15 MHz by

half of that observed for the $1.5\ \Omega$ case. At frequencies above 20 MHz, the $4.7\ \Omega$ PMA spectrum exhibits reduced amplitude compared to the $1.5\ \Omega$ PMA spectrum. Several through also shift in frequency also occur as a result of the added resistance.

Three capacitor values were evaluated, with the resulting PMA spectra exhibiting nearly identical behavior. This suggests that the selected capacitor values were too large to introduce distinguishable capacitance-dependent effects. Notably, the capacitor-only perturbations between the inputs exhibit less variance and noise than the V_{cc} capacitor-only perturbations, likely due to the absence of charging and discharging during PMA measurements. The capacitor-only PMA spectra are nearly identical to the $4.7\ \Omega$ resistor PMA spectra.




Figure 67: PMA spectra (0–3.5 V) of a TC74HC02 device with different electrical perturbations in series between inputs 1A and 1B, including resistive-only changes (left) and capacitive-only changes (right)

The TC74HC02 device would be expected to produce the correct logic output when a resistor and/or capacitor are placed in series with V_{cc} , assuming sufficient voltage is supplied for proper operation. By adding an electrical pathway between the two inputs, the inputs are constrained to similar potentials. For a two-input NOR logic gate, an input combination of low-low produces a high output, while all other input combinations produce a low output. Thus, the additional electrical pathway between the inputs has no impact on the NOR logic output. While placing a resistor or capacitor

in series between the two inputs would result in the correct logic operation for a NOR logic gate, this would not necessarily be the case for other logic functions, such as AND logic gates.

PMA is generally performed with all pins floating except for ground and power. While connecting the inputs of a gate still leaves them floating with respect to ground, the variability between their floating potentials is reduced. The relative influence of altering input floating behavior compared to the changes in circuit resistance or capacitance could not be isolated in this study.

8.3 Internal Physical Modification Results

PMA measurements were acquired for four SN7414A bipolar inverters after a series of internal physical modifications. The resulting PMA spectra are provided in Figure 67. Sample 3 was subjected to two surface scratches, whereas the remaining samples were scratched once. Images showing the die after delidding and die surface scratching for samples 1 and 2 are provided in Figure 68, while images showing the displaced wire bond and die surface scratches for sample 3 are provided in Figure 69.

The PMA spectra after wire bond displacement exhibit no observable difference from the delidded PMA spectra in all four samples. No electrical changes would be expected from wire bond displacement unless the wire was mechanically deformed or electrically compromised. Therefore, the absence of observable changes in the PMA spectra is unsurprising. High-frequency inputs, such as radio frequencies, may provide increased sensitivity to subtle geometric or surface-morphology changes caused by the tungsten probe.




Figure 68: PMA spectra (0–3.5 V) of SN7414A devices after a series of internal modifications: (A) sample 1, (B) sample 2, (C) sample 3, and (D) sample 4.




Figure 69: Images of SN7414A dies following delidding and surface modifications: (A) sample 1 die and (C) sample 2 die after delidding, and (B) sample 1 die and (D) sample 2 die after surface scratching. Images were captured using dark-field microscopy at 100x magnification.

Scratching the die surface resulted in observable deviations within the PMA spectra. The surface scratches near the middle left of sample 1 appear to cross multiple metallized paths that were previously unconnected. However, the scratched-die PMA spectrum exhibits only minor deviations relative to the delidded PMA spectrum, including a reduced amplitude of the even harmonic components above 20 MHz. Similarly, smaller deviations were observed in the PMA spectrum of sample 4 following die surface scratching.

More pronounced deviations are observed in the PMA spectrum of sample 2. Relative to the delidded PMA spectrum, the scratched-die PMA spectrum exhibits an increased amplitude of the even harmonic components beginning near 7.5 MHz. The odd harmonic components also exhibit a change in shape between 45 MHz and 50 MHz, with the scratched-die spectrum lacking the elbow observed in the delidded spectrum. Microscopy images indicate that the scratching of sample 2 was more severe than that of sample 1, with visible damage across much of the right side of the die. This increased damage severity likely explains the larger relative deviations observed.

The first surface scratch performed on the die of sample 3 is located near the left-middle region. Above the location scratched in sample 1. Consistent with the PMA spectrum of sample 1, the overall amplitude decreases slightly relative to the delidded PMA spectrum. The scratch also crosses multiple previously unconnected metallized traces. Similar to sample 2, the odd harmonic components from 45 MHz to 50 MHz approach a near-zero slope. This slope change is also observed in the PMA spectrum of the second scratch, which affected approximately half of the die surface. In contrast

to sample 2, the amplitude of the even harmonic response decreases further, rather than increasing.




Figure 70: Images of SN7414A sample 3 die following delidding and surface modifications: after (A) delidding, (B) after shifting the top-left wire bond, (C) after first surface scratching, and (D) after second surface scratching. Images were captured using dark-field microscopy at magnifications ranging from 25x to 100x.

The same modifications were performed on a single SN7417 bipolar buffer. The resulting PMA spectra are shown in Figure 70, while die microscopy images after delidding and scratching are shown in Figure 71. The primary scratch is located near the center of the die, with minor surface damage in the surrounding area. As observed in the SN7414A devices, wire bond displacement produced no observable change in the PMA spectrum. After die surface scratching, the PMA spectrum exhibits qualitatively similar behavior to the SN7414A devices, including a reduction in amplitude and a decrease in slope at higher frequencies.




Figure 71: PMA spectra (0–3.5 V) of an SN7417 device after a series of internal physical modifications.




Figure 72: Images of an SN7417 die following delidding (left) and surface scratching (right). Images were captured using dark-field microscopy at 100x magnification.

8.4 Conclusions

The presented perturbation and modification study was conducted to provide insight into the sensitivity of PMA to subtle physical and electrical changes in a device. This study was intended as a proof of concept, and the analysis is primarily qualitative. Nevertheless, the results improve the current understanding of PMA behavior and provide guidance for subsequent investigations.

Electrical perturbation results demonstrate the effect of circuit configuration on PMA, highlighting the importance of consistent and repeatable test hardware. Placing resistors and/or capacitors in series with V_{cc} produced highly variable and noisy results, particularly for the capacitor-only perturbations. In contrast, perturbations applied between the inputs produced more consistent PMA spectra, with one resistor-only perturbation exhibiting behavior nearly identical to all capacitor-only perturbations.

Physical modifications performed on inverters and buffers resulted in measurable but subtle deviations within the PMA spectra. Wire bond displacement exhibited no observable effect, while die surface scratching produced changes in amplitude and shape of the PMA spectra relative to the delidded devices.

Overall, the presented study represents an initial investigation into the effect of electrical perturbations and physical modifications on PMA spectra. Finer control of resistance, inductance, and capacitance values, along with more controlled die-surface modification techniques, would have enabled improved isolation of perturbation effects and more repeatable results.

Chapter 9: Qualitative Evaluation of PMA Voltage Ranges

As demonstrated in this dissertation, PMA is effective for a range of applications, including the detection and quantification of degradation and counterfeit screening. The performance of PMA and corresponding analysis techniques is influenced by the selected PMA voltage range. In some cases, performance differences between PMA voltage ranges are significant, such that devices either separate correctly or fail to separate entirely.

The PMA results presented throughout this dissertation utilized a trial-and-evaluation approach, often coupled with performing PMA across multiple voltage ranges. While performing PMA across multiple voltage ranges remains relatively rapid, each additional voltage range increases total testing time and the volume of data requiring analysis. Identification of the best-performing PMA voltage range(s) in this dissertation required prior knowledge of the ground truth, including the number of classes and labels. A robust, unsupervised methodology for identifying effective PMA voltage ranges could improve overall performance while reducing unnecessary data collection from lower-performing voltage ranges.

The presented investigation of effective PMA voltage ranges included two complementary components. The first component used impedance sweep measurements performed over a range of voltages and frequency spans to determine whether such measurements could provide predictive insight into effective PMA voltage ranges. The second component utilized the PMA measurements acquired across the studies presented in Chapters 5, 6, and 7 to identify qualitative visual indicators associated with higher clustering performance or improved downstream

model performance. Finally, a brief inspection of both the magnitude and power spectra derived from PMA measurements was performed to determine whether the power spectrum provides additional or distinct information.

9.1 Impedance Sweep Results

Impedance sweep measurements were performed using an Agilent E4980A, which is capable of sweeping across 201 frequencies ranging from 20 Hz to 2 MHz. To capture evenly spaced frequencies, only 200 frequencies were swept. Two frequency spans were used. A span of 10 kHz to 200 kHz with a step size of 1 kHz was used to isolate lower-frequency responses, while a span of 100 kHz to 2 MHz with a step size of 10 kHz was used to investigate higher-frequency responses. Three part numbers were tested, including four TC74HC02 NOR logic gate devices, one TL072IP OpAmp, and one RC4558IP OpAmp.

The positive terminal of the LCR meter was connected to V_{cc} of the TC74HC02 devices, while the negative terminal was connected to ground. For the OpAmp devices, the positive terminal was connected to V^+ , while the negative terminal was connected to V^- . Both frequency spans were performed using input voltage ranges of 0–0.5 V to 0–7.0 V with a voltage step size of 0.5 V for the TC74HC02 devices. The 1 kHz to 200 kHz frequency span also included voltage ranges of 0–0.05 V and 0–0.10 V, while the 100 kHz to 2 MHz frequency span included a range of 0–0.02 V. The voltage ranges used for the OpAmp devices consisted of 0–0.5 V to 0–4.0 V. The impedance and phase measurements are presented using 3D surface plots.

The TC74HC02 results are shown in Figure 72. The 100 kHz to 2 MHz phase results, shown in (B), exhibit three distinct regions. In the first region, between 0 V and

2 V, the phase is negative, indicating that the impedance is dominated by capacitance. The impedance is highest in this region as well. This is followed by a transition region, between 2 V and 4 V, where the phase rapidly increases, indicating that the impedance becomes more resistive. At the same time, the impedance continues to decrease. The observed transition is frequency- and voltage-dependent, with a positive correlation between the voltage and frequency at the transition point. A plateau in the phase is observed in the third region, between 4 V and 7 V. The impedance maintains a consistent decrease in the third region. The same three regions are observed in the phase measurements for the 1 kHz to 200 kHz span. The additional frequency resolution reveals significant decreases in the measured impedance and phase near 60 kHz, 80 kHz, and 100 kHz. It is not known why these decreases occurred.

The impedance sweep results for the TL072IP OpAmp are provided in Figure 73, while the RC4558IP OpAmp results are provided in Figure 74. While the TL072IP measured impedance from 100 kHz to 2 MHz does not exhibit any voltage dependence, it does exhibit frequency dependence, with decreasing impedance as frequency increases. The phase measurements appear to be nearly constant, except for frequencies below 100 kHz. The 1 kHz to 200 kHz phase measurements are consistent with this observation, exhibiting an increase in phase with increasing voltage below 100 kHz.

The RC4558 device exhibits results more similar to the TC74HC02 than the TL072IP, exhibiting three distinct regions. The impedance decreases with voltage from 1 kHz to 100 kHz but remains relatively constant above 100 kHz. The maximum impedance for both the RC4558 and TL072IP occurs at the lowest frequency and voltage. The similarity between the RC4558 and TC74HC02 may be a result of their

similar technologies, with the RC4558 being a BiCMOS device and the TC74HC02 being a CMOS device. In contrast, the TL072IP is a bipolar device.

While these impedance sweep results have provided insight into how specific devices may respond to certain frequencies and voltages, they alone cannot be used to determine which PMA voltage ranges may be most effective. None of the presented results were observed to correlate with clustering performance or other downstream model performance; however, a more extensive quantitative assessment could provide more meaningful information.




Figure 73: Average frequency, voltage, impedance, and phase response of TC74HC02 devices. Impedance was measured from (A) 10 kHz to 2 MHz and from (C) 1 kHz to 200 kHz, while phase was measured from (B) 10 kHz to 2 MHz and from (D) 1 kHz to 200 kHz. Voltage corresponds to the maximum amplitude ($V_{\max}$) of the applied $V_{\min}$ – $V_{\max}$ sinusoidal input waveform.




Figure 74: Average frequency, voltage, impedance, and phase response of TL072IP devices. Impedance was measured from (A) 10 kHz to 2 MHz and from (C) 1 kHz to 200 kHz, while phase was measured from (B) 10 kHz to 2 MHz and from (D) 1 kHz to 200 kHz. Voltage corresponds to the maximum amplitude ($V_{\max}$) of the applied $V_{\min}$ – $V_{\max}$ sinusoidal input waveform.




Figure 75: Average frequency, voltage, impedance, and phase response of RC4558IP devices. Impedance was measured from (A) 10 kHz to 2 MHz and from (C) 1 kHz to 200 kHz, while phase was measured from (B) 10 kHz to 2 MHz and from (D) 1 kHz to 200 kHz. Voltage corresponds to the maximum amplitude ($V_{\max}$) of the applied $V_{\min}$ – $V_{\max}$ sinusoidal input waveform.

9.2 Analysis of PMA Spectra Across Multiple PMA Voltage Ranges

The results presented in the prior chapters focused on results from individual PMA voltage ranges. However, a device's PMA spectra evolve based on the PMA voltage range used. Qualitatively assessing this evolution may provide insight into why certain PMA voltage ranges were more effective and how to identify these voltage ranges. This analysis can be further extended by investigating the PtP variance across PMA voltage ranges, which may reveal voltage-range-specific or frequency-dependent

variance within the PMA spectrum. Only the PMA voltage ranges that were used in all three studies presented in Chapters 5, 6, and 7 were assessed here. The average PMA spectra from multiple identical devices are presented in each plot, with ± 1 standard deviation shown using error bars.

The DMP2104V PMA spectra, shown in Figure 75, exhibit voltage-dependent evolution and frequency-dependent variance. The 0–1.0 V to 0–3.0 V spectra exhibit similar odd harmonic responses, with the amplitude differing between voltage ranges. The number of peaks and troughs within the even harmonic response increases with voltage range. The 0–3.5 V and 0–4.0 V PMA spectra are significantly different from the other PMA voltage ranges, with differences in amplitude and shape.

There is an increase in standard deviation between 45 MHz and 50 MHz for multiple PMA voltage ranges across the final 5 MHz. Notably, the 0–3.5 V PMA spectrum exhibits higher standard deviation within the troughs located at 20 MHz and 30 MHz compared to the rest of the spectrum and the spectra of other PMA voltage ranges. The 0–0.5 V and 0–3.5 V PMA spectra yielded the lowest K-means clustering performance among all evaluated PMA voltage ranges. The lack of peaks and troughs in the 0–0.5 V spectrum likely explains the poor clustering accuracy. Likewise, the increased standard deviation at specific frequencies in the 0–3.5 V spectrum likely contributed to the worse clustering performance by increasing within-class variance, resulting in misclustering.




Figure 76: Mean (circles) and ± 1 standard deviation (error bars) of PMA spectra for DMP2104V across several PMA voltage ranges.

In total, five different OpAmp part numbers were included across various studies presented in this dissertation. The average PMA spectra across several PMA voltage ranges for five LM741 OpAmps are plotted in Figure 76. Across the evaluated PMA voltage ranges, no new peaks or troughs emerge. The primary difference between the spectra is their amplitude, which is observed to increase with voltage amplitude. The frequency at which the even harmonic response merges with the noise floor also changes with PMA voltage range, as does the odd harmonic response near 50 MHz.

The PMA spectra of the remaining four OpAmps part numbers are shown in Figure 77. A trough emerges in the PMA spectra for both the TL072CP (A) and TL072IP (B) around 45 MHz when the voltage range exceeds 0–3.0 V. The 0–2.5 V spectra exhibit higher standard deviation than the other voltage ranges. There is also an increase in standard deviation for multiple voltage ranges above 40 MHz. Visually, the two TL072 part numbers exhibit identical behavior. The 0–1.5 V PMA spectra yielded the highest clustering accuracy, while the 0–2.5 V PMA spectra yielded the worst.

Visual inspection of the 0–1.5 V PMA spectra does not reveal why they yielded the best clustering performance.

Unlike the TL072 PMA spectra, increased standard deviation is not observed in either the RC4558 or NJM4558 PMA spectra. Both 4588 OpAmp part numbers exhibit a voltage-dependent frequency shift of the trough located between 35 MHz and 45 MHz, with the location of the trough shifting to lower frequencies with increasing voltage amplitude. Additionally, no trough is present for voltage ranges of 0–2.5 V or less. PMA voltage range did not influence clustering or downstream model performance, as perfect performance was yielded across all voltage ranges.




Figure 77: Mean (circles) and ± 1 standard deviation (error bars) of PMA spectra for LM471 across several PMA voltage ranges.




Figure 78: Mean (circles) and ± 1 standard deviation (error bars) of PMA spectra for (A) TL072CP, (B) TL072IP, (C) RC4558, and (D) NJM4558 across several PMA voltage ranges.

The PMA spectra of the LM741 and TL072 devices exhibited similar behavior to null spectra. As such, a simple experiment was performed to determine whether using different pin configurations for PMA measurements resulted in additional peaks, troughs, or other spectral features. The results of this limited-scope experiment are presented in Figure 78.

The red spectrum, in both plots, represents the typical PMA pin configuration used for OpAmp devices. Three additional pin configurations were used for each part number. For the TL072IP, supplying the modulating input (power) to V^+ while grounding Input1— exhibits the most distinct spectrum, with a trough appearing in the odd harmonic response near 45 MHz. Amplitude differences are observed between the different pin configurations. The interharmonic components have notably higher amplitude in the modulating input supplied to V^+ and Input1— connected to the ground pin configuration. The LM741 spectra exhibit identical behavior across the different pin configurations, with amplitude differences between the spectra. It is possible that

powering the OpAmps with a differential power supply and performing PMA measurements between an input and output could produce additional peaks and troughs.




Figure 79: PMA spectra (0–3.5 V) of TL072IP (left) and LM741 (right) obtained using different pins for the modulating input (power) and ground.

Multiple different logic gate devices were included across the studies presented in this dissertation. The NOR logic gate devices, shown in Figure 79, exhibit several notable features. Both the SN74HC02 and CD74HC02 manufactured by TI show significantly higher standard deviation in the 0–2.0 V spectra than any other voltage range. The standard deviation is large enough that the error bars form a shaded area. The Toshiba-manufactured TC74HC02 devices did not exhibit this increased standard deviation at 0–2.0 V.

The clustering accuracy yielded for the 0–2.0 V NOR logic gate PMA spectra was one of the worst, with only the PMA spectra below 0–1.0 V yielding lower performance. The increased standard deviation explains why the clustering performance was worse using the 0–2.0 V PMA spectra. The low clustering accuracy below 0–1.0 V is attributed to a lack of peaks or troughs that help to distinguish between the two operating temperature ratings, as separating the two manufacturers did not present a challenge. Notably, the 0–1.5 V PMA spectra yielded the best clustering performance for the irradiated devices, while the 0–3.0 V PMA spectra yielded the best

performance for the HTB study. The PMA spectra across multiple voltage ranges provide no clear indication of why these voltages were ideal. This demonstrates that the optimal PMA voltage range is both device- and application-dependent.

The remaining logic gate devices, including the OR logic gate devices shown in Figure 80, and two NAND logic gate devices are shown in Figure 81, do not exhibit voltage-dependent standard deviation. Across all six logic gate devices, peaks and troughs exhibit frequency shifting based on voltage range, in addition to new peaks and troughs emerging as the voltage range increases. Consistent with the OpAmp results, the amplitude of the PMA spectra changes with voltage range. When performing clustering of all six logic gate devices, the 0–2.0 V PMA spectra provided the worst clustering accuracy. This is consistent with the increased standard deviation observed in the NOR logic gate devices, which dominated the clustering structure.




Figure 80: Mean (circles) and ± 1 standard deviation (error bars) of PMA spectra for (A) CD74HC02, (B) SN74HC02, and (C) TC74HC02 across several PMA voltage ranges.




Figure 81: Mean (circles) and ± 1 standard deviation (error bars) of PMA spectra for SN74HC32 across several PMA voltage ranges.




Figure 82: Mean (circles) and ± 1 standard deviation (error bars) of PMA for SN74HC00 (left) and TC74HC00 (right) across several PMA voltage ranges.

The flip-flop device PMA spectra, shown in Figure 82, exhibit similar voltage-dependent patterns and frequency shifting to the logic gate devices. The 0–1.5 V PMA spectra of the SN74HC574 devices exhibit increased standard deviation, similar to the 0–2.0 V PMA spectra of the TI-manufactured NOR logic gate devices. The standard deviation is observed to increase for several PMA voltage ranges above 30 MHz. Higher standard deviation is observed across the spectra for the TC74HC574 devices. This is consistent with the observed within-class clustering variance of the TC74HC574, while the SN74HC574 devices exhibited tighter clustering. It is not clear

from visual inspection why the 0–1.5 V and 0–4.0 V PMA spectra yielded near-perfect clustering accuracy, while all other PMA voltage ranges yielded perfect accuracy.




Figure 83: Mean (circles) and ± 1 standard deviation (error bars) of PMA spectra for SN74HC574 (left) and TC74HC574 (right) across several PMA voltage ranges.

9.3 Difference Between Power Spectrum and Magnitude Spectrum

In most domains, the power spectrum is used rather than the magnitude spectrum, as it provides an accurate description of energy. As such, the magnitude and power spectra derived from PMA measurements were compared to assess whether there was any difference between them.

The magnitude and power spectra of a TC74HC02 NOR logic gate device captured using the Chiplytics PSA system operating in high-resolution mode are presented on the left of Figure 83. The magnitude spectrum is identical to the PMA spectra presented in this dissertation. The power spectrum was acquired by computing the FFT of the voltage response, computing the magnitude, and then squaring the result. As was performed for the PMA magnitude spectrum, the power spectrum was averaged over one hundred acquisitions. This is critical, as the conversion from magnitude to power is nonlinear, owing to the square-law relationship. Scaling to dBm is therefore not equivalent to scaling to dBV, with the $\frac{2}{N_s}$ in (1) squared for dBm.

Visual comparison of the magnitude and power spectra reveals nearly identical behavior, with slightly higher amplitude observed in the interharmonic components of the power spectrum compared to the magnitude spectrum. The difference between the magnitude and power spectra is shown on the right of Figure 83. While this represents a comparison between dBV and dBm, the results indicate minimal distinction between the two spectra. A maximum difference of approximately 1.0 dB is observed, primarily in the interharmonic components. Other observable differences occur in the peaks and troughs of the spectra; however, these differences are minimal. Overall, the magnitude and power spectra exhibit identical behavior and convey the same information.




Figure 84: Comparison of the PMA magnitude spectrum and PMA power spectrum (0–3.5 V) (left) and the difference between the two spectra (right).

9.4 Conclusions

The presented qualitative analysis of the relationship between PMA voltage range and the resulting PMA spectra was performed to provide insights into how effective PMA voltage ranges may be identified and to explain several observations reported in earlier chapters. While the impedance sweep measurements did not provide a direct indication of which frequencies and voltage ranges would be most effective for PMA, the results may prove useful in future studies by identifying voltage regions in which devices transition between more resistive and more capacitive behavior. Identifying these voltage-dependent transitions may provide a path toward reducing the

number of PMA voltage ranges required. Voltage ranges on both sides of these transitions represent distinct spectral regimes, thereby capturing the full evolution of a device's PMA spectra.

Analysis of the PMA spectra across multiple PMA voltage ranges did not yield a clear, generalizable criterion for identifying an effective PMA voltage range. However, the result helped explain why certain voltage ranges performed worse, enabling the elimination of low-value PMA voltage ranges in future testing applications. Lastly, comparison of the PMA the magnitude and power spectra revealed no meaningful difference in spectral content or conveyed information, indicating that either representation is suitable for PMA-based analysis. Collectively, these results indicate that PMA voltage range selection remains device- and application-dependent. However, qualitative inspection of the spectra voltage evolution and variance provides a practical basis for excluding certain PMA voltage ranges.

Chapter 10: Future Applications and Work

The work presented in this dissertation has demonstrated the capability of PMA for the detection and quantification of degradation, radiation assurance testing, and screening for counterfeit microelectronic devices. Collectively, this work has progressed the state of the art in both the application and fundamental understanding of PMA, while also laying the foundation for future research and expanded applications.

Several limitations of the presented work should be acknowledged to contextualize the results and inform future studies. Multiple studies relied on a limited number of samples, particularly the TID study, which constrained statistical confidence and downstream model performance. Furthermore, the inclusion of additional intermediate stages of degradation, for both TID and HTB exposure, could provide deeper insight into PMA's ability to distinguish between closely spaced degradation states. Likewise, investigation of more advanced degradation states approaching functional failure would improve understanding of near-failure responses, while lower radiation doses and shorter HTB exposure durations could provide insight into PMA's sensitivity to subtle or early-stage degradation. Lastly, future work should include devices with a wider range of functionality and complexity. Investigation of PMA applied to assemblies or multi-device systems is also an important area for future research.

While the engineered features derived from the developed feature extraction methodology yielded strong performance and provided improved feature traceability, dataset independence, and application relevance relative to PCA-based approaches,

future work could focus on the development of additional features. Features that explicitly capture the evolution of the PMA spectrum across multiple PMA voltage ranges may provide increased sensitivity for distinguishing between similar degradation stages or closely related devices. Similarly, features extracted from additional signal domains, such as the time domain or phase response, may yield improved sensitivity in specific applications.

The perturbation and modification study results indicated that PMA may be sensitive to subtle electrical and physical changes; however, the degree and limits of this sensitivity were not quantified. Future work should expand on this study by introducing controlled modifications and more complex devices using precise modification techniques, such as focused ion beam (FIB) processing for die-level alterations. Alternatively, modifications could involve controlled alteration of semiconductor doping profiles. Validation of these modifications using complementary techniques, such as quantum diamond microscopy (QDM) to map current distributions, could enable direct correlation between localized physical changes and the PMA spectra.

Additional research to assess the repeatability and reproducibility of PMA measurements, particularly when using the Chiplytics PSA system, would further strengthen confidence in the technique. Furthermore, the development of calibration methodologies for PMA would facilitate broader adoption and improved cross-platform comparability. Such studies could also explore the influence of environmental factors, including temperature, on PMA measurement stability and accuracy. The effect of temperature on a device's PMA spectra, and whether testing over a range of

temperatures provided improved performance, could also be evaluated. Similarly, exploration of alternative modulating waveforms, such as sawtooth or triangle waveforms, may provide new mechanisms for extracting device-specific information. Future work should also consider how PMA may be integrated alongside existing standards-based counterfeit detection and reliability test methods to serve as a rapid pre-screen or complementary diagnostic tool.

Finally, exploration of PMA as a prognostics and health management (PHM) tool is of interest. This would likely require adaptation of PMA toward low-amplitude modulating signals combined with larger DC offsets, enabling in-situ monitoring without disrupting normal device operation. Such an approach would require transitioning PMA from a discrete test methodology to a continuous or embedded monitoring framework.

Chapter 11: Contributions

1. Development of a PMA-based side-channel method for detection and quantification of microelectronic device degradation.
 - a. Evaluated degradation detection and quantification using clustering, semi-supervised metrics, classification, and regression. Regression error results demonstrated the accuracy of PMA-based predictions.
 - b. Demonstrated differentiation between five distinct stages of high-temperature bias (HTB) degradation in CMOS logic gate devices.
 - c. Demonstrated differentiation between six or more ionizing radiation dose levels across multiple device types, including MOSFETs, logic gates, and operational amplifiers.
2. Development of a feature extraction methodology tailored to PMA from which application-relevant and dataset-independent features are acquired that provide improved feature traceability and performance compared to existing PCA-based approaches.
 - a. The developed feature extraction methodology uses spectral subdivision to break the PMA spectrum into spectral bands, and even harmonic, odd harmonic, and interharmonic subspectra.
 - b. Feature categories include statistical features, spectral descriptors, correlation-based features, wavelet-based features, cepstral-based features, and derivative-based features.

- c. Established that the feature extraction methodology and resulting engineered features are dataset-independent, enabling reuse across devices, stress mechanisms, and hardware configurations.
 - d. Demonstrated that the feature extraction methodology and engineered features provided comparable or improved separation and model performance relative to existing PCA-based approaches.
- 3. Development of an interharmonic removal approach for PMA-based counterfeit screening using PCA.
 - a. Identified interharmonic components as a counterfeit screening specific source of non-informative variance in PMA spectra, demonstrating that their removal produces minimal negative impact on clustering accuracy across a range of devices and voltage ranges, while improving performance in select cases.
 - b. Demonstrated that interharmonic removal increases the proportion of explained variance captured by the leading principal components.
 - c. Successfully discriminated between manufacturers, technologies, and device variants for both analog (operational amplifiers) and digital devices (logic gates and flip-flops) using both engineered features and PCA-based analysis.
- 4. Quantification of PMA equipment and device response repeatability, including short-term repeatability, part-to-part, and day-to-day variance.
 - a. Repeatability analysis validated that the observed HTB results were driven by degradation.

- b. First quantitative assessment of PMA repeatability.
- 5. Establishment that optimal PMA voltage range is device- and stress-specific.
 - a. Demonstration of the role of PMA voltage range in clustering performance and part-to-part measurement variance.
- 6. Demonstration that variations in measurement hardware and circuit electrical characteristics can produce PMA spectral differences comparable to device-level differences and significantly affect clustering performance.
 - a. Comparative evaluation of PMA results obtained using benchtop equipment and Chiplytics PSA system highlighted hardware-specific clustering performance.
 - b. Demonstration of the impact of electrical perturbations on PMA spectral shape.

Appendices

The following appendices provide additional results from the studies presented in Chapters 5–7, including the HTB study in Chapter 5, the TID study presented in Chapter 6, and the counterfeit screening study presented in Chapter 7. Additional PMA results from the HTD study indicate whether the data were captured using the Chiplytics PSA system or the BTE equipment. Unless otherwise stated, all the remaining PMA results were captured using the Chiplytics PSA system.

Unless otherwise noted, values in parentheses within tables indicate the number of retained principal components or engineered features used in the corresponding model. For the PCA-based results labeled “Expected”, the number of retained principal components was selected using the slope rule defined in Section 4.4. Columns labeled “Maximum” report the minimum number of retained principal components required to achieve the highest clustering accuracy.

The FS-SCC feature set refers to the engineered features remaining after Spearman correlation pruning, as defined in Section 4.4. Results using FS-PCA and FS-LDA feature selection are also included and were performed as defined in Section 4.4. The columns labeled FS-PCA retained the top 50 features, while the FS-LDA retained the minimum number required to achieve the highest clustering accuracy.

Appendix A.

Additional TC74HC02 HTB Study Results




Figure A1: Timeline plots for RMSE (dBV) of DtD stability (left) and PtP variance (right) across several PMA voltage ranges.




Figure A2: Box-and-whisker plots of RMSE (dBV) for DtD stability (left) and PtP variance (right) across nine PMA voltage ranges.




Figure A3: Box-and-whisker plots of RMSE (dBV) for each control device based on ten repeated measurements.




Figure A4: Box-and-whisker plots of RMSE (dBV) of the DtD stability (left) and PtP variance (right) of the BTE setup

Table A1: K-means and HDBSCAN clustering accuracy using PCA models fitted to PMA spectra from the HTB study, captured using the Chiplytics PSA system.

$V_{\min} - V_{\max}$	K-means Accuracy		HDBSCAN Accuracy	
	Expected	Maximum	Expected	Maximum
0 – 0.5	56.2 (2)	63.8 (25)	17.5 (2)	28.8 (20)
0 – 1.0	69.8 (2)	75.0 (4)	30.2 (2)	33.3 (4)
0 – 1.5	79.2 (3)	97.9 (16)	37.5 (3)	6.7 (4)
0 – 2.0	61.5 (2)	96.9 (4)	34.4 (2)	66.7 (8)
0 – 2.5	69.9 (3)	99.0 (7)	59.4 (3)	66.7 (5)
0 – 3.0	96.9 (3)	100 (11)	43.8 (3)	67.7 (4)
0 – 3.5	93.8 (2)	100 (12)	25.0 (2)	52.1 (4)
0 – 4.0	82.3 (3)	94.8 (13)	63.5 (3)	63.5 (3)
0 – 4.5	94.8 (3)	100 (4)	40.6 (3)	62.5 (4)
0 – 5.0	92.7 (2)	100 (10)	27.1 (2)	69.8 (4)

Table A2: K-means and HDBSCAN clustering accuracy using PCA models fitted to PMA spectra with interharmonic components excluded from the HTB study, captured using the Chiplytics PSA system.

$V_{\min} - V_{\max}$	K-means Accuracy		HDBSCAN Accuracy	
	Expected	Maximum	Expected	Maximum
0 – 0.5	40.0 (2)	45.0 (4)	15.0 (2)	26.2 (8)
0 – 1.0	53.1 (2)	62.5 (5)	20.8 (2)	33.3 (4)
0 – 1.5	70.8 (3)	76.0 (13)	47.9 (3)	68.8 (9)
0 – 2.0	69.8 (3)	88.5 (11)	42.7 (3)	66.7 (5)
0 – 2.5	91.7 (3)	99.0 (4)	58.3 (3)	66.7 (5)
0 – 3.0	95.8 (3)	99.0 (7)	66.7 (3)	74.0 (7)
0 – 3.5	95.8 (3)	96.9 (13)	34.4 (3)	66.7 (5)
0 – 4.0	78.1 (4)	81.2 (22)	50.0 (4)	65.6 (5)
0 – 4.5	97.9 (3)	100 (5)	46.9 (3)	81.2 (7)
0 – 5.0	78.1 (2)	86.5 (3)	24.0 (2)	91.7 (11)

Table A3: K-means and HDBSCAN clustering accuracy using engineered features extracted from the PMA spectra from the HTB study, captured using the Chiptytics PSA system.

$V_{\min} - V_{\max}$	K-means Accuracy			HDBSCAN Accuracy	
	FS-SCC	FS-PCA	FS-LDA	FS-PCA	FS-LDA
0 – 0.5	57.5 (324)	60.0 (50)	91.3 (17)	25.0 (50)	66.2 (21)
0 – 1.0	80.2 (323)	88.5 (50)	97.9 (72)	33.3 (50)	33.3 (2)
0 – 1.5	91.7 (309)	81.3 (50)	100 (51)	52.1 (50)	81.2 (13)
0 – 2.0	99.0 (275)	100 (50)	100 (60)	50.0 (50)	66.7 (9)
0 – 2.5	94.8 (266)	96.9 (50)	100 (15)	66.7 (50)	66.7 (4)
0 – 3.0	99.0 (300)	100 (50)	100 (52)	50.0 (50)	62.5 (8)
0 – 3.5	92.7 (282)	100 (50)	100 (41)	50.0 (50)	83.3 (7)
0 – 4.0	96.9 (287)	100 (50)	100 (8)	50.0 (50)	82.3 (25)
0 – 4.5	96.9 (290)	100 (50)	100 (7)	50.0 (50)	83.3 (7)
0 – 5.0	93.8 (267)	86.5 (50)	100 (4)	66.7 (50)	83.3 (16)

Table A4: K-means and HDBSCAN clustering accuracy using PCA models fitted to PMA spectra from the HTB study, captured using the BTE setup.

$V_{\min} - V_{\max}$	K-means Accuracy		HDBSCAN Accuracy	
	Expected	Maximum	Expected	Maximum
0 – 0.5	41.2 (2)	58.8 (5)	18.8 (2)	26.2 (4)
0 – 1.0	47.5 (4)	58.8 (2)	23.8 (4)	31.2 (7)
0 – 1.5	70.0 (3)	76.2 (25)	30.0 (3)	52.5 (5)
0 – 2.0	28.8 (5)	28.8 (2)	28.8 (5)	36.2 (9)
0 – 3.5	57.5 (5)	65.0 (2)	33.7 (5)	37.5 (20)
0 – 5.0	56.2 (5)	72.5 (2)	26.2 (5)	42.5 (6)

Table A5: K-means and HDBSCAN clustering accuracy using PCA models fitted to PMA spectra with interharmonic components excluded from the HTB study, captured using the BTE setup.

$V_{\min} - V_{\max}$	K-means Accuracy		HDBSCAN Accuracy	
	Expected	Maximum	Expected	Maximum
0 – 0.5	43.8 (3)	52.5 (13)	21.1 (3)	38.8 (13)
0 – 1.0	58.8 (4)	58.8 (2)	30.0 (4)	37.5 (9)
0 – 1.5	77.5 (5)	77.5 (3)	28.8 (5)	62.5 (8)
0 – 2.0	28.8 (6)	28.8 (2)	31.2 (6)	35.0 (7)
0 – 3.5	72.5 (4)	90.0 (9)	28.8 (4)	31.2 (5)
0 – 5.0	66.2 (4)	76.2 (2)	20.0 (4)	30.0 (9)

Table A6: K-means and HDBSCAN clustering accuracy using engineered features extracted from PMA spectra from the HTB study, captured using the BTE setup.

$V_{\min} - V_{\max}$	K-means Accuracy			HDBSCAN Accuracy	
	FS-SCC	FS-PCA	FS-LDA	FS-PCA	FS-LDA
0 – 0.5	60.0 (286)	46.3 (50)	82.5 (23)	33.8 (50)	42.5 (10)
0 – 1.0	55.0 (297)	45.0 (50)	76.2 (47)	31.3 (50)	40.0 (6)
0 – 1.5	52.5 (299)	51.3 (50)	96.2 (8)	42.5 (50)	51.2 (8)
0 – 2.0	70.0 (292)	52.5 (50)	76.2 (66)	40.0 (50)	40.0 (8)
0 – 3.5	100 (292)	70.0 (50)	100 (38)	57.5 (50)	52.5 (15)
0 – 5.0	78.8 (263)	96.3 (50)	100 (5)	75.0 (50)	86.2 (65)




Figure A5: Autoencoder reconstructor error, trained on PMA spectra (0–3.5 V) of control devices, computed for devices at different stages of HTB degradation. Vertical dashed lines indicate the median of each degradation stage.




Figure A6: Impedance sweep measurements using 0–1.0 V input of HTB degraded TC74HC02 devices, showing impedance (left) and phase (right). Shaded region indicates ± 1 standard deviation.




Figure A7: Impedance sweep measurements using 0–2.0 V input of HTB degraded TC74HC02 devices, showing impedance (left) and phase (right). Shaded region indicates ± 1 standard deviation.

Appendix B.

Additional TID Study Results

B.1 Additional DMP2104V Results

Table B1: K-means and HDBSCAN clustering accuracy using PCA models fitted to PMA spectra of unbiased DMP2104V devices from the TID study, captured at CALCE.

$V_{\min} - V_{\max}$	K-means Accuracy		HDBSCAN Accuracy	
	Expected	Maximum	Expected	Maximum
–0.05 – 0.05	62.5 (3)	68.8 (6)	50.0 (3)	50.0 (3)
0.0 – 0.10	56.2 (3)	68.8 (5)	53.1 (3)	59.4 (5)
0.0 – 0.25	56.2 (4)	65.6 (13)	43.8 (4)	56.2 (3)
0 – 0.5	50.0 (5)	59.4 (10)	53.1 (5)	53.1 (5)
0 – 1.0	65.6 (5)	68.8 (16)	56.2 (5)	56.2 (4)
0 – 1.5	75.0 (4)	81.2 (7)	68.8 (4)	68.8 (3)
0 – 2.0	78.1 (4)	81.2 (17)	68.8 (4)	78.1 (7)
0 – 2.5	75.0 (4)	84.4 (11)	71.9 (4)	75.0 (7)
0 – 3.0	75.0 (4)	78.1 (8)	81.2 (4)	81.2 (4)
0 – 3.5	50.0 (5)	50.0 (5)	40.6 (5)	43.8 (12)
0 – 4.0	62.5 (5)	65.6 (21)	62.5 (5)	78.1 (3)

Table B2: K-means and HDBSCAN clustering accuracy using engineered features extracted from PMA spectra of unbiased DMP2104V devices from the TID study, captured at CALCE.

$V_{\min} - V_{\max}$	K-means Accuracy			HDBSCAN Accuracy	
	FS-SCC	FS-PCA	FS-LDA	FS-PCA	FS-LDA
-0.05 – 0.05	65.6 (304)	62.5 (50)	84.4 (36)	18.78 (50)	65.6 (8)
0.0 – 0.10	56.3 (316)	68.8 (50)	90.6 (35)	53.1 (50)	65.6 (19)
0.0 – 0.25	53.1 (315)	50.0 (50)	75.0 (9)	15.6 (50)	59.4 (5)
0 – 0.5	56.2 (287)	43.8 (50)	78.1 (43)	25.0 (50)	68.8 (28)
0 – 1.0	59.4 (279)	59.4 (50)	81.2 (59)	28.1 (50)	46.9 (2)
0 – 1.5	68.8 (289)	71.9 (50)	84.4 (36)	28.1 (50)	28.1 (3)
0 – 2.0	78.1 (277)	81.2 (50)	90.6 (17)	28.1 (50)	81.2 (12)
0 – 2.5	71.9 (279)	81.2 (50)	90.6 (7)	28.1 (50)	87.5 (6)
0 – 3.0	84.4 (281)	71.9 (50)	87.5 (9)	28.1 (50)	81.2 (8)
0 – 3.5	46.9 (256)	43.8 (50)	68.8 (24)	28.1 (50)	53.1 (2)
0 – 4.0	59.4 (286)	56.2 (50)	78.1 (16)	28.1 (50)	50.0 (2)

Table B3: K-means and HDBSCAN clustering accuracy using PCA models fitted to PMA spectra of biased DMP2104V devices from the TID study, captured at CALCE.

$V_{\min} - V_{\max}$	K-means Accuracy		HDBSCAN Accuracy	
	Expected	Maximum	Expected	Maximum
-0.05 – 0.05	50.0 (3)	56.2 (22)	53.1 (3)	53.1 (3)
0.0 – 0.10	59.4 (3)	62.5 (5)	40.6 (3)	46.9 (4)
0.0 – 0.25	56.2 (3)	65.6 (4)	40.6 (3)	46.9 (5)
0 – 0.5	56.2 (3)	68.8 (12)	56.2 (3)	68.8 (5)
0 – 1.0	75.0 (3)	78.1 (6)	68.8 (3)	81.2 (5)
0 – 1.5	78.1 (4)	78.1 (4)	62.5 (4)	71.9 (6)
0 – 2.0	78.1 (5)	78.1 (5)	65.5 (5)	78.1 (2)
0 – 2.5	71.9 (5)	81.2 (6)	75.0 (5)	75.0 (4)
0 – 3.0	78.1 (4)	81.2 (6)	68.8 (4)	81.2 (12)
0 – 3.5	53.1 (5)	53.1 (2)	50.0 (5)	56.2 (7)
0 – 4.0	56.2 (5)	62.5 (23)	53.1 (5)	62.5 (8)

Table B4: K-means and HDBSCAN clustering accuracy using engineered features extracted from PMA spectra of biased DMP2104V devices from the TID study, at CALCE.

$V_{\min} - V_{\max}$	K-means Accuracy			HDBSCAN Accuracy	
	FS-SCC	FS-PCA	FS-LDA	FS-PCA	FS-LDA
-0.05 – 0.05	53.1 (319)	56.2 (50)	78.1 (30)	56.2 (50)	65.6 (63)
0.0 – 0.10	53.1 (312)	62.5 (50)	75.0 (16)	50.0 (50)	53.1 (2)
0.0 – 0.25	53.1 (306)	53.1 (50)	71.9 (12)	31.2 (50)	59.4 (9)
0 – 0.5	53.1 (285)	62.5 (50)	71.9 (2)	37.5 (50)	56.2 (2)
0 – 1.0	59.4 (280)	65.6 (50)	81.2 (13)	28.1 (50)	59.4 (13)
0 – 1.5	71.9 (286)	71.9 (50)	78.1 (33)	28.1 (50)	56.2 (2)
0 – 2.0	81.2 (278)	75.0 (50)	96.9 (71)	28.1 (50)	75.0 (11)
0 – 2.5	75.0 (279)	78.1 (50)	87.5 (15)	28.1 (50)	78.1 (15)
0 – 3.0	65.6 (283)	78.1 (50)	96.9 (29)	28.1 (50)	78.1 (12)
0 – 3.5	59.4 (257)	50.0 (50)	68.8 (59)	46.9 (50)	53.1 (4)
0 – 4.0	62.5 (295)	59.4 (50)	78.1 (22)	34.4 (50)	65.6 (16)

Table B5: K-means and HDBSCAN clustering accuracy using PCA models fitted to PMA spectra of biased and unbiased DMP2104V devices from the TID study, captured at CALCE.

$V_{\min} - V_{\max}$	K-means Accuracy		HDBSCAN Accuracy	
	Expected	Maximum	Expected	Maximum
-0.05 – 0.05	40.7 (3)	50.8 (5)	35.6 (3)	37.3 (2)
0.0 – 0.10	50.8 (3)	50.8 (2)	39.0 (3)	42.4 (9)
0.0 – 0.25	47.5 (4)	49.2 (13)	32.2 (4)	37.3 (2)
0 – 0.5	49.2 (3)	59.3 (13)	42.4 (3)	45.8 (5)
0 – 1.0	64.4 (4)	69.5 (15)	62.7 (4)	67.8 (5)
0 – 1.5	69.5 (5)	76.3 (8)	64.4 (5)	71.2 (7)
0 – 2.0	67.8 (6)	74.6 (13)	67.8 (6)	69.5 (9)
0 – 2.5	64.4 (5)	72.9 (11)	67.8 (5)	72.9 (9)
0 – 3.0	71.2 (5)	74.6 (7)	71.2 (5)	74.6 (7)
0 – 3.5	33.9 (5)	42.4 (27)	35.6 (5)	39.0 (6)
0 – 4.0	59.3 (5)	59.3 (5)	59.3 (5)	59.3 (5)

Table B6: K-means and HDBSCAN clustering accuracy using engineered features extracted from PMA spectra of biased and unbiased DMP2104V devices from the TID study, captured at CALCE.

$V_{\min} - V_{\max}$	K-means Accuracy			HDBSCAN Accuracy	
	FS-SCC	FS-PCA	FS-LDA	FS-PCA	FS-LDA
-0.05 – 0.05	39.0 (311)	42.4 (50)	54.2 (12)	10.2 (50)	44.1 (4)
0.0 – 0.10	50.8 (311)	45.8 (50)	61.0 (26)	40.7 (50)	49.2 (3)
0.0 – 0.25	45.8 (309)	47.5 (50)	55.9 (48)	8.5 (50)	44.1 (2)
0 – 0.5	50.8 (286)	39.0 (50)	59.3 (25)	27.1 (50)	47.5 (2)
0 – 1.0	61.0 (277)	49.2 (50)	69.5 (12)	22.0 (50)	39.0 (3)
0 – 1.5	66.1 (389)	59.3 (50)	71.2 (36)	22.0 (50)	44.1 (2)
0 – 2.0	62.7 (280)	66.1 (50)	79.7 (30)	18.6 (50)	50.8 (2)
0 – 2.5	59.3 (279)	52.5 (50)	78.0 (60)	22.0 (50)	54.2 (3)
0 – 3.0	61.0 (291)	55.9 (50)	83.1 (43)	25.4 (50)	57.6 (4)
0 – 3.5	42.4 (218)	35.6 (50)	55.9 (13)	32.2 (50)	40.7 (3)
0 – 4.0	52.5 (291)	47.5 (50)	67.8 (23)	23.7 (50)	50.8 (3)

Table B7: K-means and HDBSCAN clustering accuracy using PCA models fitted to PMA spectra of unbiased DMP2104V devices from the TID study, captured at Sandia National Laboratories.

$V_{\min} - V_{\max}$	K-means Accuracy		HDBSCAN Accuracy	
	Expected	Maximum	Expected	Maximum
-0.05 – 0.05	78.1 (4)	81.2 (21)	78.1 (4)	78.1 (2)
0.0 – 0.10	75.0 (4)	90.6 (2)	75.0 (4)	81.2 (6)
0.0 – 0.25	65.6 (4)	71.9 (7)	68.8 (4)	81.2 (26)
0 – 0.5	71.9 (3)	84.4 (2)	59.4 (3)	81.2 (11)
0 – 1.0	78.1 (3)	90.6 (19)	71.9 (3)	81.2 (9)
0 – 1.5	87.5 (6)	90.6 (3)	78.1 (6)	78.1 (6)
0 – 2.0	87.5 (6)	93.8 (2)	71.9 (6)	81.2 (2)
0 – 2.5	93.8 (5)	93.8 (5)	81.2 (5)	84.4 (6)
0 – 3.0	68.8 (5)	81.2 (2)	84.4 (5)	84.4 (4)
0 – 3.5	71.9 (5)	71.9 (5)	62.5 (5)	68.8 (25)
0 – 4.0	71.9 (4)	81.2 (2)	75.0 (4)	78.1 (2)

Table B8: K-means and HDBSCAN clustering accuracy using engineered features extracted from PMA spectra of unbiased DMP2104V devices from the TID study, captured at Sandia National Laboratories.

$V_{\min} - V_{\max}$	K-means Accuracy			HDBSCAN Accuracy	
	FS-SCC	FS-PCA	FS-LDA	FS-PCA	FS-LDA
-0.05 – 0.05	65.6 (257)	62.5 (50)	90.6 (12)	50.0 (50)	84.4 (45)
0.0 – 0.10	68.8 (256)	59.4 (50)	87.5 (33)	25.0 (50)	75.0 (6)
0.0 – 0.25	71.9 (263)	68.8 (50)	90.6 (47)	21.9 (50)	81.2 (5)
0 – 0.5	56.3 (264)	62.5 (50)	90.6 (21)	68.8 (50)	87.5 (21)
0 – 1.0	75.0 (254)	75.0 (50)	87.5 (8)	53.1 (50)	65.6 (3)
0 – 1.5	71.9 (254)	62.5 (50)	93.8 (27)	62.5 (50)	87.5 (8)
0 – 2.0	71.9 (273)	71.9 (50)	93.8 (13)	62.5 (50)	90.6 (28)
0 – 2.5	68.8 (264)	68.8 (50)	93.8 (8)	56.3 (50)	87.5 (2)
0 – 3.0	53.1 (252)	68.8 (50)	90.6 (4)	68.8 (50)	90.6 (17)
0 – 3.5	62.5 (265)	59.4 (50)	84.4 (2)	50.0 (50)	75.0 (5)
0 – 4.0	53.1 (258)	53.3 (50)	87.5 (3)	53.1 (50)	81.2 (12)

Table B9: K-means and HDBSCAN clustering accuracy using PCA models fitted to PMA spectra of biased DMP2104V devices from the TID study, captured at Sandia National Laboratories.

$V_{\min} - V_{\max}$	K-means Accuracy		HDBSCAN Accuracy	
	Expected	Maximum	Expected	Maximum
-0.05 – 0.05	65.6 (4)	75.0 (2)	59.4 (4)	62.5 (3)
0.0 – 0.10	62.5 (4)	75.0 (28)	56.2 (4)	56.2 (4)
0.0 – 0.25	62.5 (4)	75.0 (14)	56.2 (4)	56.2 (4)
0 – 0.5	68.8 (6)	75.0 (5)	62.5 (6)	65.6 (7)
0 – 1.0	62.5 (4)	78.1 (3)	78.1 (4)	78.1 (4)
0 – 1.5	65.6 (3)	84.4 (4)	75.0 (3)	78.1 (4)
0 – 2.0	75.0 (6)	84.4 (3)	81.2 (6)	81.2 (5)
0 – 2.5	84.4 (5)	84.4 (5)	71.9 (5)	81.2 (6)
0 – 3.0	53.1 (5)	62.5 (7)	75.0 (5)	75.0 (4)
0 – 3.5	59.4 (5)	65.6 (6)	53.1 (5)	62.5 (22)
0 – 4.0	65.6 (6)	68.8 (8)	65.6 (6)	65.6 (5)

Table B10: K-means and HDBSCAN clustering accuracy using engineered features extracted from PMA spectra of biased DMP2104V devices from the TID study, captured at Sandia National Laboratories.

$V_{\min} - V_{\max}$	K-means Accuracy			HDBSCAN Accuracy	
	FS-SCC	FS-PCA	FS-LDA	FS-PCA	FS-LDA
-0.05 – 0.05	59.4 (296)	65.6 (50)	81.2 (55)	53.1 (50)	68.8 (4)
0.0 – 0.10	53.1 (285)	65.6 (50)	75.0 (12)	53.1 (50)	56.2 (2)
0.0 – 0.25	59.4 (304)	68.8 (50)	81.2 (48)	40.6 (50)	68.8 (2)
0 – 0.5	59.4 (284)	75.0 (50)	78.1 (15)	62.5 (50)	65.6 (4)
0 – 1.0	59.4 (278)	71.9 (50)	84.4 (12)	62.5 (50)	84.4 (9)
0 – 1.5	78.1 (285)	71.9 (50)	84.4 (34)	62.5 (50)	71.9 (29)
0 – 2.0	87.5 (267)	81.3 (50)	90.6 (13)	78.1 (50)	87.5 (18)
0 – 2.5	68.8 (265)	68.8 (50)	84.4 (61)	62.5 (50)	78.1 (4)
0 – 3.0	53.1 (276)	53.3 (50)	81.2 (13)	28.1 (50)	75.0 (19)
0 – 3.5	68.8 (245)	65.6 (50)	78.1 (17)	56.3 (50)	75.0 (9)
0 – 4.0	68.8 (255)	75.0 (50)	81.2 (28)	40.6 (50)	68.8 (20)

Table B11: Classification model performance for TID-degraded biased and unbiased DMP2104V devices using different input datasets, including features extracted from the PMA spectra (0–2.0 V) captured at CALCE.

Method	Input	Accuracy	F1
MLP	FS-SCC Features	0.19 ± 0.06	0.10 ± 0.04
	PMA Spectra	0.16 ± 0.05	0.08 ± 0.03
	PCA	0.10 ± 0.03	0.04 ± 0.01
	FS-PCA Features	0.07 ± 0.02	0.01 ± 0.0
	Harmonic Spectra	0.25 ± 0.10	0.18 ± 0.07
	Curve Trace	0.11 ± 0.04	0.04 ± 0.02
RF	FS-SCC Features	0.65 ± 0.06	0.59 ± 0.10
	PMA Spectra	0.54 ± 0.07	0.49 ± 0.09
	PCA	0.65 ± 0.06	0.59 ± 0.10
	FS-PCA Features	0.54 ± 0.10	0.51 ± 0.10
	Harmonic Spectra	0.60 ± 0.08	0.55 ± 0.09
	Curve Trace	0.50 ± 0.06	0.47 ± 0.06
SVC	FS-SCC Features	0.56 ± 0.08	0.50 ± 0.10
	PMA Spectra	0.63 ± 0.08	0.57 ± 0.11
	PCA	0.66 ± 0.10	0.60 ± 0.10
	FS-PCA Features	0.63 ± 0.09	0.57 ± 0.10
	Harmonic Spectra	0.63 ± 0.11	0.58 ± 0.12
	Curve Trace	0.58 ± 0.06	0.50 ± 0.06

Table B12: Classification model performance for TID-degraded unbiased DMP2104V devices using different input datasets, including features extracted from the PMA spectra (0–2.0 V) captured at Sandia National Laboratories.

Method	Input	Accuracy	F1
MLP	FS-SCC Features	0.28 ± 0.09	0.16 ± 0.08
	PMA Spectra	0.28 ± 0.10	0.19 ± 0.07
	PCA	0.08 ± 0.08	0.04 ± 0.04
	FS-PCA Features	0.26 ± 0.11	0.17 ± 0.10
	Harmonic Spectra	0.36 ± 0.10	0.25 ± 0.06
	Curve Trace	0.28 ± 0.08	0.13 ± 0.04
	FS-SCC Features	0.80 ± 0.13	0.80 ± 0.13
RF	PMA Spectra	0.79 ± 0.15	0.78 ± 0.15
	PCA	0.85 ± 0.10	0.85 ± 0.09
	FS-PCA Features	0.81 ± 0.09	0.82 ± 0.09
	Harmonic Spectra	0.85 ± 0.14	0.86 ± 0.12
	Curve Trace	0.82 ± 0.08	0.80 ± 0.07
	FS-SCC Features	0.82 ± 0.14	0.81 ± 0.14
	PMA Spectra	0.72 ± 0.13	0.67 ± 0.12
SVC	PCA	0.77 ± 0.13	0.75 ± 0.12
	PCA Features	0.80 ± 0.12	0.78 ± 0.13
	Harmonic Spectra	0.82 ± 0.10	0.79 ± 0.09
	Curve Trace	0.79 ± 0.09	0.77 ± 0.09

Table B13: Classification model performance for TID-degraded biased DMP2104V devices using different input datasets, including features extracted from the PMA spectra (0–2.0 V) captured at Sandia National Laboratories.

Method	Input	Accuracy	F1
MLP	FS-SCC Features	0.08 ± 0.06	0.02 ± 0.02
	PMA Spectra	0.33 ± 0.17	0.25 ± 0.12
	PCA	0.03 ± 0.04	0.01 ± 0.01
	FS-PCA Features	0.06 ± 0.06	0.03 ± 0.03
	Harmonics Spectra	0.38 ± 0.13	0.30 ± 0.11
	Curve Trace	0.34 ± 0.04	0.20 ± 0.02
	FS-SCC Features	0.74 ± 0.08	0.71 ± 0.10
RF	PMA Spectra	0.68 ± 0.12	0.63 ± 0.13
	PCA	0.79 ± 0.08	0.76 ± 0.08
	FS-PCA Features	0.81 ± 0.08	0.77 ± 0.10
	Harmonic Spectra	0.74 ± 0.08	0.70 ± 0.08
	Curve Trace	0.80 ± 0.08	0.77 ± 0.12
	FS-SCC Features	0.82 ± 0.09	0.78 ± 0.09
	PMA Spectra	0.75 ± 0.08	0.71 ± 0.08
SVC	PCA	0.78 ± 0.11	0.73 ± 0.11
	FS-PCA Features	0.81 ± 0.10	0.77 ± 0.11
	Harmonic Spectra	0.80 ± 0.08	0.74 ± 0.11
	Curve Trace	0.77 ± 0.07	0.74 ± 0.07

B.2 Additional TC74HC02 Results

Table B14: K-means and HDBSCAN clustering accuracy using PCA models fitted to PMA spectra of TC74HC02 devices from the TID study.

$V_{\min} - V_{\max}$	K-means Accuracy		HDBSCAN Accuracy	
	Expected	Maximum	Expected	Maximum
-0.05 – 0.05	62.2 (2)	70.3 (13)	67.6 (2)	67.6 (2)
0.0 – 0.10	59.5 (2)	64.9 (3)	64.9 (2)	64.9 (2)
0.0 – 0.25	59.5 (2)	64.9 (4)	62.2 (2)	62.2 (2)
0 – 0.5	73.0 (2)	73.0 (2)	67.6 (2)	73.0 (3)
0 – 1.0	100 (4)	100 (4)	100 (4)	100 (4)
0 – 1.5	100 (3)	100 (2)	100 (3)	100 (2)
0 – 2.0	78.4 (3)	78.4 (3)	89.2 (3)	89.2 (3)
0 – 2.5	86.5 (4)	86.5 (2)	91.9 (4)	97.3 (5)
0 – 3.0	86.5 (4)	86.5 (4)	100 (4)	100 (3)
0 – 3.5	56.8 (3)	86.5 (6)	83.8 (3)	97.3 (6)
0 – 4.0	70.3 (5)	70.3 (2)	91.9 (5)	91.9 (4)

Table B15: K-means and HDBSCAN clustering accuracy using engineered features extracted from PMA spectra of TC74HC02 devices from the TID study.

$V_{\min} - V_{\max}$	K-means Accuracy			HDBSCAN Accuracy	
	FS-SCC	FS-PCA	FS-LDA	FS-PCA	FS-LDA
-0.05 – 0.05	51.4 (302)	67.6 (50)	86.5 (4)	35.1 (50)	81.1 (2)
0.0 – 0.10	51.4 (294)	54.1 (50)	78.4 (38)	29.7 (50)	73.0 (3)
0.0 – 0.25	64.9 (301)	83.8 (50)	86.5 (27)	35.1 (50)	83.8 (3)
0 – 0.5	62.2 (299)	64.9 (50)	86.5 (8)	37.8 (50)	89.2 (8)
0 – 1.0	59.5 (249)	97.3 (50)	100 (9)	78.4 (50)	100 (10)
0 – 1.5	75.7 (234)	100 (50)	100 (2)	89.2 (50)	100 (3)
0 – 2.0	83.8 (244)	78.4 (50)	100 (3)	81.1 (50)	100 (2)
0 – 2.5	73.0 (236)	86.5 (50)	100 (3)	97.3 (50)	100 (5)
0 – 3.0	78.4 (243)	86.5 (50)	100 (3)	67.6 (50)	100 (10)
0 – 3.5	70.3 (266)	73.0 (50)	100 (10)	89.2 (50)	100 (12)
0 – 4.0	70.3 (265)	59.5 (50)	86.5 (22)	51.4 (50)	91.9 (54)

B.3 Additional LM741 Results

Table B16: K-means and HDBSCAN clustering accuracy using PCA models fitted to PMA spectra of LM741 devices from the TID study.

$V_{\min} - V_{\max}$	K-means Accuracy		HDBSCAN Accuracy	
	Expected	Maximum	Expected	Maximum
-0.05 – 0.05	37.8 (2)	51.4 (11)	45.5 (2)	45.5 (2)
0.0 – 0.10	40.5 (2)	45.9 (3)	36.4 (2)	39.4 (3)
0.0 – 0.25	37.8 (2)	48.6 (3)	42.4 (2)	45.5 (3)
0 – 0.5	32.4 (2)	45.9 (8)	39.4 (2)	39.4 (2)
0 – 1.0	43.2 (3)	51.4 (6)	36.4 (3)	51.5 (4)
0 – 1.5	37.8 (2)	45.9 (9)	39.4 (2)	39.4 (2)
0 – 2.0	37.8 (3)	40.5 (2)	27.3 (3)	51.5 (2)
0 – 2.5	29.7 (3)	43.2 (5)	48.5 (3)	48.5 (3)
0 – 3.0	32.4 (3)	45.9 (30)	45.5 (3)	51.5 (2)
0 – 3.5	43.2 (4)	54.1 (7)	51.5 (2)	51.5 (2)
0 – 4.0	40.5 (3)	54.1 (11)	48.5 (2)	48.5 (2)

Table B17: K-means and HDBSCAN clustering accuracy using engineered features extracted from PMA spectra of LM741 devices from the TID study.

$V_{\min} - V_{\max}$	K-means Accuracy			HDBSCAN Accuracy	
	FS-SCC	FS-PCA	FS-LDA	FS-PCA	FS-LDA
-0.05 – 0.05	37.8 (312)	35.1 (50)	59.5 (16)	21.6 (50)	45.9 (3)
0.0 – 0.10	43.2 (313)	40.5 (50)	59.5 (8)	18.9 (50)	48.6 (3)
0.0 – 0.25	40.5 (313)	32.4 (50)	73.0 (58)	18.9 (50)	43.2 (4)
0 – 0.5	37.8 (312)	27.0 (50)	70.3 (17)	21.6 (50)	54.1 (12)
0 – 1.0	51.4 (306)	35.1 (50)	64.9 (16)	24.3 (50)	62.2 (2)
0 – 1.5	46.0 (301)	40.5 (50)	73.0 (41)	24.3 (50)	51.4 (3)
0 – 2.0	35.1 (296)	32.4 (50)	70.3 (5)	29.7 (50)	48.6 (3)
0 – 2.5	35.1 (306)	37.8 (50)	78.4 (48)	24.3 (50)	51.4 (3)
0 – 3.0	37.8 (296)	32.4 (50)	62.2 (41)	24.3 (50)	40.5 (3)
0 – 3.5	46.0 (286)	43.2 (50)	75.7 (34)	32.4 (50)	56.8 (3)
0 – 4.0	40.5 (287)	32.4 (50)	70.3 (51)	32.4 (50)	45.9 (2)

Appendix C:

Additional Counterfeit Screening Study Results

C.1 Additional NOR Logic Gate Device Results

Table C1: K-means and HDBSCAN clustering accuracy using PCA models fitted to PMA spectra of NOR logic gate devices from the counterfeit screening study.

$V_{\min} - V_{\max}$	K-means Accuracy		HDBSCAN Accuracy	
	Expected	Maximum	Expected	Maximum
0.0 – 0.10	74.7 (2)	84.0 (57)	22.7 (2)	64.0 (5)
0.0 – 0.25	66.7 (2)	86.7 (53)	22.7 (2)	66.7 (5)
0 – 0.5	68.0 (2)	96.0 (55)	18.7 (2)	66.7 (8)
0 – 1.0	100 (3)	100 (2)	20.0 (3)	68.0 (5)
0 – 1.5	98.7 (3)	98.7 (2)	46.7 (3)	77.3 (8)
0 – 2.0	80.0 (3)	80.0 (2)	25.3 (3)	53.3 (10)
0 – 2.5	100 (3)	100 (2)	76.0 (3)	100 (6)
0 – 3.0	100 (3)	100 (2)	24.0 (3)	100 (4)
0 – 3.5	100 (3)	100 (2)	26.7 (3)	100 (5)
0 – 4.0	100 (3)	100 (2)	70.7 (3)	100 (4)

Table C2: K-means and HDBSCAN clustering accuracy using PCA models fitted to PMA spectra with interharmonic components excluded of NOR logic gate devices from the counterfeit screening study.

$V_{\min} - V_{\max}$	K-means Accuracy		HDBSCAN Accuracy	
	Expected	Maximum	Expected	Maximum
0.0 – 0.10	68.0 (2)	80.0 (36)	18.7 (2)	66.7 (5)
0.0 – 0.25	66.7 (2)	80.0 (42)	24.0 (2)	66.7 (4)
0 – 0.5	68.0 (3)	68.0 (2)	20.0 (3)	68.0 (8)
0 – 1.0	100 (3)	100 (2)	21.3 (3)	100 (10)
0 – 1.5	100 (3)	100 (2)	26.7 (3)	100 (8)
0 – 2.0	80.0 (3)	80.0 (2)	48.0 (3)	50.7 (43)
0 – 2.5	100 (3)	100 (2)	73.3 (3)	100 (5)
0 – 3.0	100 (3)	100 (2)	50.7 (3)	100 (6)
0 – 3.5	100 (3)	100 (2)	21.3 (3)	100 (5)
0 – 4.0	100 (3)	100 (2)	72.0 (3)	100 (5)

Table C3: K-means and HDBSCAN clustering accuracy using engineered features extracted from PMA spectra of NOR logic gate devices from the counterfeit screening study.

$V_{\min} - V_{\max}$	K-means Accuracy			HDBSCAN Accuracy	
	FS-SCC	FS-PCA	FS-LDA	FS-PCA	FS-LDA
0.0 – 0.10	73.3 (314)	89.3 (50)	100 (2)	66.7 (50)	66.7 (4)
0.0 – 0.25	78.7 (296)	81.3 (50)	100 (2)	66.7 (50)	74.7 (3)
0 – 0.5	69.3 (294)	100 (50)	100 (2)	66.7 (50)	100 (4)
0 – 1.0	100 (246)	100 (50)	100 (2)	100 (50)	100 (6)
0 – 1.5	96.0 (243)	96.0 (50)	100 (2)	66.7 (50)	100 (18)
0 – 2.0	60.0 (254)	61.3 (50)	92.0 (5)	58.7 (50)	69.3 (20)
0 – 2.5	100 (258)	100 (50)	100 (7)	100 (50)	100 (14)
0 – 3.0	100 (263)	100 (50)	100 (2)	100 (50)	100 (4)
0 – 3.5	100 (260)	100 (50)	100 (2)	100 (50)	100 (9)
0 – 4.0	100 (230)	100 (50)	100 (2)	100 (50)	100 (9)

C.2 Additional NAND Logic Gate Device Results

Table C4: K-means and HDBSCAN clustering accuracy using PCA models fitted to PMA spectra of NAND logic gate devices from the counterfeit screening study.

$V_{\min} - V_{\max}$	K-means Accuracy		HDBSCAN Accuracy	
	Expected	Maximum	Expected	Maximum
0.0 – 0.10	100 (2)	100 (2)	38.0 (2)	100 (4)
0.0 – 0.25	100 (2)	100 (2)	18.0 (2)	100 (3)
0 – 0.5	100 (2)	100 (2)	18.0 (2)	100 (3)
0 – 1.0	100 (2)	100 (2)	16.0 (2)	100 (4)
0 – 1.5	100 (2)	100 (2)	20.0 (2)	100 (3)
0 – 2.0	100 (2)	100 (2)	18.0 (2)	100 (4)
0 – 2.5	100 (2)	100 (2)	18.0 (2)	100 (4)
0 – 3.0	100 (2)	100 (2)	26.0 (2)	100 (4)
0 – 3.5	100 (2)	100 (2)	18.0 (2)	100 (4)
0 – 4.0	100 (2)	100 (2)	22.0 (2)	100 (4)

Table C5: K-means and HDBSCAN clustering accuracy using PCA models fitted to PMA spectra with interharmonic components excluded of NAND logic gate devices from the counterfeit screening study.

$V_{\min} - V_{\max}$	K-means Accuracy		HDBSCAN Accuracy	
	Expected	Maximum	Expected	Maximum
0.0 – 0.10	100 (2)	100 (2)	16.0 (2)	100 (4)
0.0 – 0.25	100 (2)	100 (2)	22.0 (2)	100 (4)
0 – 0.5	100 (2)	100 (2)	22.0 (2)	100 (5)
0 – 1.0	100 (2)	100 (2)	24.0 (2)	100 (6)
0 – 1.5	100 (2)	100 (2)	24.0 (2)	100 (4)
0 – 2.0	100 (2)	100 (2)	26.0 (2)	100 (7)
0 – 2.5	100 (2)	100 (2)	18.0 (2)	100 (8)
0 – 3.0	100 (2)	100 (2)	24.0 (2)	100 (6)
0 – 3.5	100 (2)	100 (2)	20.0 (2)	100 (7)
0 – 4.0	100 (2)	100 (2)	14.0 (2)	100 (7)

Table C6: K-means and HDBSCAN clustering accuracy using engineered features extracted from PMA spectra of NAND logic gate devices from the counterfeit screening study.

$V_{\min} - V_{\max}$	K-means Accuracy			HDBSCAN Accuracy	
	FS-SCC	FS-PCA	FS-LDA	FS-PCA	FS-LDA
0.0 – 0.10	100 (293)	100 (50)	100 (2)	100 (50)	100 (4)
0.0 – 0.25	100 (281)	100 (50)	100(2)	100 (50)	100 (4)
0 – 0.5	100 (295)	100 (50)	100 (2)	100 (50)	100 (5)
0 – 1.0	100 (295)	100 (50)	100 (2)	100 (50)	100 (4)
0 – 1.5	100 (256)	100 (50)	100 (2)	100 (50)	100 (4)
0 – 2.0	100 (258)	100 (50)	100 (2)	100 (50)	100 (3)
0 – 2.5	100 (270)	100 (50)	100 (2)	100 (50)	100 (4)
0 – 3.0	100 (258)	100 (50)	100 (2)	100 (50)	100 (4)
0 – 3.5	100 (254)	100 (50)	100 (2)	100 (50)	100 (6)
0 – 4.0	100 (242)	100 (50)	100 (2)	100 (50)	100 (6)

C.3 Additional NOR, OR, and NAND Logic Gate Device Results

Table C7: K-means and HDBSCAN clustering accuracy using PCA models fitted to PMA spectra of NOR, OR, and NAND logic gate devices from the counterfeit screening study.

$V_{\min} - V_{\max}$	K-means Accuracy		HDBSCAN Accuracy	
	Expected	Maximum	Expected	Maximum
0.0 – 0.10	80.0 (2)	88.0 (83)	27.3 (2)	66.0 (6)
0.0 – 0.25	83.3 (3)	95.3 (2)	19.3 (3)	66.7 (5)
0 – 0.5	81.3 (3)	89.3 (2)	29.3 (3)	64.0 (7)
0 – 1.0	100 (4)	100 (2)	75.3 (4)	85.3 (5)
0 – 1.5	98.0 (5)	99.3 (4)	86.0 (5)	100 (11)
0 – 2.0	46.7 (3)	66.0 (4)	22.7 (3)	76.7 (12)
0 – 2.5	100 (4)	100 (3)	86.0 (4)	100 (5)
0 – 3.0	100 (5)	100 (2)	78.7 (5)	100 (7)
0 – 3.5	100 (4)	100 (3)	65.3 (4)	100 (6)
0 – 4.0	100 (5)	100 (2)	76.7 (5)	100 (8)

Table C8: K-means and HDBSCAN clustering accuracy using PCA models fitted to PMA spectra with interharmonic components excluded of NOR, OR, and NAND logic gate devices from the counterfeit screening study.

$V_{\min} - V_{\max}$	K-means Accuracy		HDBSCAN Accuracy	
	Expected	Maximum	Expected	Maximum
0.0 – 0.10	79.3 (2)	86.0 (14)	20.7 (2)	66.7 (4)
0.0 – 0.25	64.0 (3)	88.7 (87)	36.7 (3)	66.7 (4)
0 – 0.5	60.0 (3)	78.7 (12)	32.0 (3)	67.3 (13)
0 – 1.0	100 (4)	100 (2)	82.0 (4)	100 (12)
0 – 1.5	100 (5)	100 (4)	62.7 (5)	100 (8)
0 – 2.0	90.0 (5)	90.0 (5)	66.7 (5)	75.3 (45)
0 – 2.5	100 (4)	100 (3)	74.0 (4)	100 (16)
0 – 3.0	100 (5)	100 (2)	50.0 (5)	100 (16)
0 – 3.5	100 (4)	100 (3)	52.7 (4)	100 (14)
0 – 4.0	100 (5)	100 (2)	74.7 (5)	100 (12)

Table C9: K-means and HDBSCAN clustering accuracy using engineered features extracted from PMA spectra of NOR, OR, and NAND logic gate devices from the counterfeit screening study.

$V_{\min} - V_{\max}$	K-means Accuracy			HDBSCAN Accuracy	
	FS-SCC	FS-PCA	FS-LDA	FS-PCA	FS-LDA
0.0 – 0.10	82.7 (269)	86.7 (50)	92.7 (4)	33.3 (50)	66.7 (6)
0.0 – 0.25	70.0 (269)	88.7 (50)	97.3 (10)	66.7 (50)	67.3 (9)
0 – 0.5	66.0 (260)	96.7 (50)	98.0 (16)	66.7 (50)	72.7 (6)
0 – 1.0	85.3 (224)	100 (50)	100 (2)	100 (50)	100 (9)
0 – 1.5	94.7 (232)	94.7 (50)	100 (8)	83.3 (50)	100 (10)
0 – 2.0	81.3 (257)	62.7 (50)	98.7 (5)	85.3 (50)	78.0 (20)
0 – 2.5	88.7 (238)	100 (50)	100 (7)	100 (50)	100 (6)
0 – 3.0	86.0 (244)	100 (50)	100 (5)	83.3 (50)	100 (9)
0 – 3.5	84.0 (232)	100 (50)	100 (3)	83.3 (50)	100 (9)
0 – 4.0	100 (235)	100 (50)	100 (5)	98.7 (50)	100 (6)

C.4 Additional Flip-Flop Device Results

Table C10: K-means and HDBSCAN clustering accuracy using PCA models fitted to PMA spectra of flip-flop devices from the counterfeit screening study.

$V_{\min} - V_{\max}$	K-means Accuracy		HDBSCAN Accuracy	
	Expected	Maximum	Expected	Maximum
0.0 – 0.10	100 (2)	100 (2)	18.0 (2)	100 (5)
0.0 – 0.25	100 (2)	100 (2)	26.0 (2)	100 (4)
0 – 0.5	100 (3)	100 (2)	30.0 (3)	100 (6)
0 – 1.0	100 (2)	100 (2)	36.0 (2)	96.0 (8)
0 – 1.5	96.0 (4)	96.0 (2)	54.0 (4)	84.0 (10)
0 – 2.0	100 (3)	100 (2)	58.0 (3)	96.0 (9)
0 – 2.5	100 (3)	100 (2)	72.0 (3)	96.0 (7)
0 – 3.0	100 (3)	100 (2)	56.0 (3)	96.0 (8)
0 – 3.5	100 (3)	100 (2)	50.0 (3)	92.0 (8)
0 – 4.0	96.0 (2)	96.0 (2)	52.0 (2)	88.0 (42)

Table C11: K-means and HDBSCAN clustering accuracy using PCA models fitted to PMA spectra with interharmonic components excluded of flip-flop devices from the counterfeit screening study.

$V_{\min} - V_{\max}$	K-means Accuracy		HDBSCAN Accuracy	
	Expected	Maximum	Expected	Maximum
0.0 – 0.10	100 (2)	100 (2)	20.0 (2)	100 (4)
0.0 – 0.25	100 (2)	100 (2)	20.0 (2)	100 (4)
0 – 0.5	100 (3)	100 (2)	30.0 (3)	100 (18)
0 – 1.0	100 (2)	100 (2)	22.0 (2)	92.0 (9)
0 – 1.5	96.0 (4)	96.0 (2)	14.0 (4)	84.0 (17)
0 – 2.0	100 (3)	100 (2)	18.0 (3)	96.0 (9)
0 – 2.5	100 (3)	100 (2)	18.0 (3)	96.0 (11)
0 – 3.0	100 (3)	100 (2)	80.0 (3)	96.0 (9)
0 – 3.5	100 (3)	100 (2)	20.0 (3)	92.0 (8)
0 – 4.0	96.0 (2)	96.0 (2)	28.0 (2)	76.0 (6)

Table C12: K-means and HDBSCAN clustering accuracy using engineered features extracted from PMA spectra of flip-flop devices from the counterfeit screening study.

$V_{\min} - V_{\max}$	K-means Accuracy			HDBSCAN Accuracy	
	FS-SCC	FS-PCA	FS-LDA	FS-PCA	FS-LDA
0.0 – 0.10	100 (301)	100 (50)	100 (2)	100 (50)	100 (3)
0.0 – 0.25	100 (294)	100 (50)	100 (2)	100 (50)	100 (4)
0 – 0.5	100 (294)	100 (50)	100 (2)	100 (50)	100 (5)
0 – 1.0	100 (254)	100 (50)	100 (2)	92.0 (50)	100 (6)
0 – 1.5	92.0 (269)	92.0 (50)	100 (2)	84.0 (50)	98.0 (10)
0 – 2.0	100 (265)	100 (50)	100 (2)	96.0 (50)	100 (7)
0 – 2.5	100 (270)	100 (50)	100 (2)	96.0 (50)	96.0 (6)
0 – 3.0	100 (260)	100 (50)	100 (2)	96.0 (50)	100 (71)
0 – 3.5	100 (251)	100 (50)	100 (2)	92.0 (50)	100 (9)
0 – 4.0	92.0 (259)	92.0 (50)	100 (2)	76.0 (50)	96.0 (17)

C.5 Additional OpAmp Results

Table C13: K-means and HDBSCAN clustering accuracy using PCA models fitted to PMA spectra of 072 and 4558 OpAmp devices from the counterfeit screening study.

$V_{\min} - V_{\max}$	K-means Accuracy		HDBSCAN Accuracy	
	Expected	Maximum	Expected	Maximum
0.0 – 0.10	76.0 (2)	76.0 (2)	20.0 (2)	52.0 (29)
0.0 – 0.25	81.0 (2)	85.0 (18)	18.0 (2)	50.0 (5)
0 – 0.5	81.0 (2)	85.0 (49)	16.0 (2)	50.0 (6)
0 – 1.0	75.0 (3)	81.0 (70)	32.0 (3)	75.0 (7)
0 – 1.5	82.0 (3)	84.0 (23)	17.0 (3)	75.0 (6)
0 – 2.0	77.0 (2)	77.0 (2)	19.0 (2)	75.0 (14)
0 – 2.5	46.0 (5)	46.0 (2)	57.0 (5)	69.0 (17)
0 – 3.0	76.0 (2)	76.0 (2)	14.0 (2)	75.0 (11)
0 – 3.5	77.0 (2)	78.0 (3)	18.0 (2)	77.0 (17)
0 – 4.0	80.0 (4)	81.0 (57)	58.0 (4)	77.0 (6)

Table C14: K-means and HDBSCAN clustering accuracy using PCA models fitted to PMA spectra with interharmonic components excluded of 072 and 4558 OpAmp devices from the counterfeit screening study.

$V_{\min} - V_{\max}$	K-means Accuracy		HDBSCAN Accuracy	
	Expected	Maximum	Expected	Maximum
0.0 – 0.10	81.0 (2)	81.0 (2)	18.0 (2)	50.0 (3)
0.0 – 0.25	83.0 (2)	87.0 (46)	23.0 (2)	50.0 (5)
0 – 0.5	78.0 (3)	78.0 (2)	52.0 (3)	65.0 (6)
0 – 1.0	75.0 (4)	75.0 (3)	25.0 (4)	75.0 (17)
0 – 1.5	77.0 (3)	77.0 (3)	24.0 (3)	75.0 (6)
0 – 2.0	77.0 (2)	77.0 (2)	16.0 (2)	71.0 (7)
0 – 2.5	46.0 (5)	46.0 (2)	41.0 (5)	69.0 (61)
0 – 3.0	76.0 (2)	76.0 (2)	18.0 (2)	77.0 (27)
0 – 3.5	77.0 (2)	77.0 (2)	16.0 (2)	76.0 (16)
0 – 4.0	80.0 (4)	80.0 (3)	18.0 (4)	77.0 (9)

Table C15: K-means and HDBSCAN clustering accuracy using engineered features extracted from PMA spectra of 072 and 4558 OpAmp devices from the counterfeit screening study.

$V_{\min} - V_{\max}$	K-means Accuracy			HDBSCAN Accuracy	
	FS-SCC	FS-PCA	FS-LDA	FS-PCA	FS-LDA
0.0 – 0.10	57.0 (297)	75.0 (50)	81.0 (33)	50.0 (50)	58.0 (4)
0.0 – 0.25	53.0 (292)	74.0 (50)	83.0 (48)	50.0 (50)	75.0 (3)
0 – 0.5	81.0 (293)	82.0 (50)	85.0 (64)	75.0 (50)	75.0 (4)
0 – 1.0	79.0 (277)	68.0 (50)	90.0 (3)	75.0 (50)	75.0 (5)
0 – 1.5	77.0 (275)	76.0 (50)	93.0 (2)	75.0 (50)	75.0 (6)
0 – 2.0	77.0 (245)	77.0 (50)	80.0 (11)	74.0 (50)	75.0 (7)
0 – 2.5	78.0 (223)	77.0 (50)	78.0 (13)	69.0 (50)	77.0 (22)
0 – 3.0	77.0 (235)	76.0 (50)	81.0 (45)	75.0 (50)	77.0 (60)
0 – 3.5	77.0 (238)	77.0 (50)	84.0 (6)	77.0 (50)	77.0 (24)
0 – 4.0	77.0 (226)	76.0 (50)	88.0 (3)	75.0 (50)	77.0 (12)

Table C16: K-means and HDBSCAN clustering accuracy using PCA models fitted to PMA spectra of 4558 OpAmp devices from the counterfeit screening study.

$V_{\min} - V_{\max}$	K-means Accuracy		HDBSCAN Accuracy	
	Expected	Maximum	Expected	Maximum
0.0 – 0.10	100 (2)	100 (2)	34.0 (2)	100 (7)
0.0 – 0.25	100 (2)	100 (2)	24.0 (2)	100 (4)
0 – 0.5	100 (2)	100 (2)	18.0 (2)	100 (7)
0 – 1.0	100 (3)	100 (2)	28.0 (3)	100 (6)
0 – 1.5	100 (2)	100 (2)	22.0 (2)	100 (6)
0 – 2.0	100 (2)	100 (2)	26.0 (2)	100 (5)
0 – 2.5	100 (2)	100 (2)	56.0 (2)	100 (4)
0 – 3.0	100 (2)	100 (2)	24.0 (2)	100 (4)
0 – 3.5	100 (2)	100 (2)	20.0 (2)	100 (5)
0 – 4.0	100 (2)	100 (2)	18.0 (2)	100 (4)

Table C17: K-means and HDBSCAN clustering accuracy using PCA models fitted to PMA spectra with interharmonic components excluded of 4558 OpAmp devices from the counterfeit screening study.

$V_{\min} - V_{\max}$	K-means Accuracy		HDBSCAN Accuracy	
	Expected	Maximum	Expected	Maximum
0.0 – 0.10	100 (2)	100 (2)	20.0 (2)	100 (5)
0.0 – 0.25	100 (2)	100 (2)	20.0 (2)	100 (4)
0 – 0.5	100 (3)	100 (2)	40.0 (3)	100 (5)
0 – 1.0	100 (3)	100 (2)	20.0 (3)	100 (11)
0 – 1.5	100 (2)	100 (2)	22.0 (2)	100 (4)
0 – 2.0	100 (2)	100 (2)	26.0 (2)	100 (4)
0 – 2.5	100 (2)	100 (2)	16.0 (2)	100 (5)
0 – 3.0	100 (2)	100 (2)	20.0 (2)	100 (6)
0 – 3.5	100 (2)	100 (2)	16.0 (2)	100 (7)
0 – 4.0	100 (2)	100 (2)	20.0 (2)	100 (6)

Table C18: K-means and HDBSCAN clustering accuracy using engineered features extracted from PMA spectra of 4558 OpAmp devices from the counterfeit screening study.

$V_{\min} - V_{\max}$	K-means Accuracy			HDBSCAN Accuracy	
	FS-SCC	FS-PCA	FS-LDA	FS-PCA	FS-LDA
0.0 – 0.10	100 (293)	100 (50)	100 (2)	100 (50)	100 (5)
0.0 – 0.25	100 (289)	100 (50)	100 (2)	100 (50)	100 (4)
0 – 0.5	100 (286)	100 (50)	100 (2)	100 (50)	100 (5)
0 – 1.0	100 (280)	100 (50)	100 (2)	100 (50)	100 (4)
0 – 1.5	100 (291)	100 (50)	100 (2)	100 (50)	100 (4)
0 – 2.0	100 (267)	100 (50)	100 (2)	100 (50)	100 (6)
0 – 2.5	100 (280)	100 (50)	100 (2)	100 (50)	100 (5)
0 – 3.0	100 (271)	100 (50)	100 (2)	100 (50)	100 (4)
0 – 3.5	100 (266)	100 (50)	100 (2)	100 (50)	100 (4)
0 – 4.0	100 (275)	100 (50)	100 (2)	100 (50)	100 (3)

Table C19: K-means and HDBSCAN clustering accuracy using PCA models fitted to PMA spectra of 072 OpAmp devices from the counterfeit screening study.

$V_{\min} - V_{\max}$	K-means Accuracy		HDBSCAN Accuracy	
	Expected	Maximum	Expected	Maximum
0.0 – 0.10	58.0 (2)	66.0 (43)	10.0 (2)	46.0 (13)
0.0 – 0.25	70.0 (2)	76.0 (17)	18.0 (2)	46.0 (22)
0 – 0.5	62.0 (2)	72.0 (17)	20.0 (2)	48.0 (6)
0 – 1.0	50.0 (2)	58.0 (44)	14.0 (2)	52.0 (7)
0 – 1.5	64.0 (2)	66.0 (25)	12.0 (2)	46.0 (18)
0 – 2.0	54.0 (3)	54.0 (2)	42.0 (3)	50.0 (6)
0 – 2.5	56.0 (4)	56.0 (2)	12.0 (4)	38.0 (13)
0 – 3.0	52.0 (4)	52.0 (2)	14.0 (4)	50.0 (11)
0 – 3.5	54.0 (2)	54.0 (2)	12.0 (2)	54.0 (13)
0 – 4.0	60.0 (2)	62.0 (8)	16.0 (2)	54.0 (4)

Table C20: K-means and HDBSCAN clustering accuracy using PCA models fitted to PMA spectra with interharmonic components excluded of 072 OpAmp devices from the counterfeit screening study.

$V_{\min} - V_{\max}$	K-means Accuracy		HDBSCAN Accuracy	
	Expected	Maximum	Expected	Maximum
0.0 – 0.10	62.0 (2)	68.0 (10)	12.0 (2)	50.0 (13)
0.0 – 0.25	60.0 (2)	68.0 (9)	12.0 (2)	42.0 (17)
0 – 0.5	56.0 (2)	56.0 (2)	18.0 (2)	52.0 (19)
0 – 1.0	50.0 (2)	50.0 (2)	12.0 (2)	50.0 (27)
0 – 1.5	54.0 (3)	54.0 (2)	12.0 (3)	50.0 (15)
0 – 2.0	54.0 (3)	54.0 (2)	24.0 (3)	42.0 (6)
0 – 2.5	56.0 (4)	56.0 (2)	16.0 (4)	38.0 (42)
0 – 3.0	52.0 (5)	52.0 (2)	14.0 (5)	54.0 (26)
0 – 3.5	54.0 (4)	54.0 (2)	12.0 (4)	52.0 (5)
0 – 4.0	60.0 (3)	60.0 (2)	16.0 (3)	54.0 (7)

Table C21: K-means and HDBSCAN clustering accuracy using engineered features extracted from PMA spectra of 072 OpAmp devices from the counterfeit screening study.

$V_{\min} - V_{\max}$	K-means Accuracy			HDBSCAN Accuracy	
	FS-SCC	FS-PCA	FS-LDA	FS-PCA	FS-LDA
0.0 – 0.10	66.0 (326)	52.0 (50)	68.0 (60)	52.0 (50)	44.0 (24)
0.0 – 0.25	52.0 (311)	52.0 (50)	74.0 (45)	52.0 (50)	56.0 (10)
0 – 0.5	58.0 (320)	58.0 (50)	76.0 (30)	52.0 (50)	48.0 (27)
0 – 1.0	54.0 (313)	52.0 (50)	90.0 (8)	38.0 (50)	54.0 (5)
0 – 1.5	52.0 (313)	52.0 (50)	98.0 (3)	34.0 (50)	60.0 (5)
0 – 2.0	54.0 (273)	54.0 (50)	66.0 (4)	42.0 (50)	48.0 (23)
0 – 2.5	56.0 (250)	54.0 (50)	74.0 (2)	32.0 (50)	48.0 (8)
0 – 3.0	52.0 (283)	54.0 (50)	74.0 (16)	50.0 (50)	54.0 (10)
0 – 3.5	54.0 (291)	54.0 (50)	74.0 (36)	52.0 (50)	54.0 (44)
0 – 4.0	52.0 (246)	54.0 (50)	86.0 (37)	52.0 (50)	54.0 (34)

References

- [1] M. Pecht and M. Kang, *Prognostics and Health Management of Electronics: Fundamentals, Machine Learning, and Internet of Things*, 2nd ed. Hoboken: Wiley-IEEE press, 2018.
- [2] J. W. McPherson, *Reliability Physics and Engineering: Time-To-Failure Modeling*. Cham, Switzerland: Springer International Publishing, 2019. doi: 10.1007/978-3-319-93683-3.
- [3] K. C. Kapur, *Reliability Engineering*. in Wiley series in systems engineering and management. Hoboken, N.J: J. Wiley & Sons, 2014.
- [4] V. Khemani, M. H. Azarian, and M. Pecht, “WavePHMNet: A comprehensive diagnosis and prognosis approach for analog circuits,” *Advanced Engineering Informatics*, vol. 59, p. 102323, Jan. 2024, doi: 10.1016/j.aei.2023.102323.
- [5] J. Lutz, H. Schlangenotto, U. Scheuermann, and R. De Doncker, *Semiconductor Power Devices*. Cham: Springer International Publishing, 2018. doi: 10.1007/978-3-319-70917-8.
- [6] M. Randolph and W. Diehl, “Power side-channel attack analysis: a review of 20 years of study for the layman,” *Cryptography*, vol. 4, no. 2, p. 15, May 2020, doi: 10.3390/cryptography4020015.
- [7] D. R. Richman, M. H. Azarian, D. Das, and M. Pecht, “Analysis of standards-based counterfeit microelectronics detection methods,” *IEEE Access*, vol. 13, pp. 7691–7704, 2025, doi: 10.1109/ACCESS.2025.3526518.

- [8] P. Tangyunyong *et al.*, “Power spectrum analysis,” in *International Symposium for Testing and Failure Analysis*, Pasadena, California, USA: ASM International, Nov. 2017, pp. 73–78. doi: 10.31399/asm.cp.istfa2017p0073.
- [9] S. DeJong, R. Multari, K. Wilson, and P. Tangyunyong, “Evaluation of COTS electronics by power spectrum analysis and multivariate data analysis,” Sandia National Laboratories (SNL), Albuquerque, NM, and Livermore, CA (United States), Sep. 2022. doi: 10.2172/1890397.
- [10] P. Socha, V. Miškovský, and M. Novotný, “A comprehensive survey on the non-invasive passive side-channel analysis,” *Sensors*, vol. 22, no. 21, p. 8096, Oct. 2022, doi: 10.3390/s22218096.
- [11] M. Joye and F. Olivier, “Side-Channel Analysis,” in *Encyclopedia of Cryptography and Security*, H. C. A. Van Tilborg and S. Jajodia, Eds., Boston, MA: Springer US, 2011, pp. 1198–1204. doi: 10.1007/978-1-4419-5906-5_516.
- [12] S. Bhunia and M. H. Tehranipoor, “Side-Channel Attacks,” in *Hardware Security: A Hands-On Learning Approach*, Cambridge, Mass: Morgan Kaufmann, 2019, pp. 193–218. [Online]. Available: <https://shop.elsevier.com/books/hardware-security/bhunias/978-0-12-812477-2>
- [13] T. Popp, S. Mangard, and E. Oswald, “Power analysis attacks and countermeasures,” *IEEE Des. Test. Comput.*, vol. 24, no. 6, pp. 535–543, Nov. 2007, doi: 10.1109/MDT.2007.200.
- [14] P. Kocher, J. Jaffe, and B. Jun, “Differential Power Analysis,” in *Advances in Cryptology — CRYPTO’ 99*, vol. 1666, in Lecture Notes in Computer Science,

- vol. 1666. , Berlin, Heidelberg: Springer Berlin Heidelberg, 1999, pp. 388–397.
doi: 10.1007/3-540-48405-1_25.
- [15] P. Kocher, J. Jaffe, B. Jun, and P. Rohatgi, “Introduction to differential power analysis,” *J Cryptogr Eng*, vol. 1, no. 1, pp. 5–27, Apr. 2011, doi: 10.1007/s13389-011-0006-y.
- [16] U. Guin, D. DiMase, and M. Tehranipoor, “Counterfeit integrated circuits: detection, avoidance, and the challenges ahead,” *Journal of Electronic Testing*, vol. 30, pp. 9–23, Feb. 2014, doi: 10.1007/s10836-013-5430-8.
- [17] U. Guin, K. Huang, D. DiMase, J. M. Carulli, M. Tehranipoor, and Y. Makris, “Counterfeit integrated circuits: a rising threat in the global semiconductor supply chain,” *Proceedings of the IEEE*, vol. 102, no. 8, pp. 1207–1228, Aug. 2014, doi: 10.1109/JPROC.2014.2332291.
- [18] *Test Methods Standard; General Requirements, Suspect/Counterfeit, Electrical, Electronic, and Electromechanical Parts*, AS6171A, SAE International, 2018.
- [19] S. Bhunia, M. S. Hsiao, M. Banga, and S. Narasimhan, “Hardware Trojan Attacks: Threat Analysis and Countermeasures,” *Proc. IEEE*, vol. 102, no. 8, pp. 1229–1247, Aug. 2014, doi: 10.1109/JPROC.2014.2334493.
- [20] S. Kutzner, A. Y. Poschmann, and M. Stöttinger, “Hardware Trojan design and detection: a practical evaluation,” in *Proceedings of the Workshop on Embedded Systems Security*, Montreal Quebec Canada: ACM, Sep. 2013, pp. 1–9. doi: 10.1145/2527317.2527318.
- [21] J. Roberts and M. Riley, “The Big Hack: How China Used a Tiny Chip to Infiltrate U.S. Companies,” *Bloomberg Businessweek*, Oct. 04, 2018.

- [22] GIDEP, “Government-Industry Data Exchange Program (GIDEP).” [Online].
Available: <https://www.gidep.org/>
- [23] ERAI, “Electronics Reseller Association International (ERAI).” [Online].
Available: <https://www.eraí.com/>
- [24] D. Akhiundov, “2024 annual report,” ERAI. Accessed: Jun. 24, 2024. [Online].
Available:
https://www.eraí.com/eraí_blog/3187/_2024_annual_report#:~:text=In%202024%2C%20ERAI%20reported%20a,reported%20by%20ERAI%20since%202015.
- [25] S. Agarwal, L. Risse, C. Barnes, and R. Menke, “Dangerous counterfeit electronics,” *EEE Parts Bulletin*, vol. 3, no. 1, May 2011.
- [26] A. Varas, R. Varadarajan, J. Goodrich, and F. Yinug, “Government incentives and US competitiveness in semiconductor manufacturing,” Boston Consulting Group (BCC) and Semiconductor Industry Association (SIA), Jul. 2020. [Online].
Available: <https://www.semiconductors.org/wp-content/uploads/2020/09/Government-Incentives-and-US-Competitiveness-in-Semiconductor-Manufacturing-Sep-2020.pdf>
- [27] *Acceptability of Electronic Components Distributed in the Open Market*, IDEA-STD-1010-B, Independent Distributors of Electronics Association (IDEA), 2011.
- [28] *Techniques for Suspect/Counterfeit EEE Parts Detection by External Visual Inspection, Remarking and Resurfacing, and Surface Texture Analysis Using SEM Test Methods*, AS6171/2A, SAE International, 2017.
- [29] *Techniques for Suspect/Counterfeit EEE Parts Detection by Radiological Test Methods*, AS6171/5, SAE International, 2016.

- [30] *Techniques for Suspect/Counterfeit EEE Parts Detection by X-ray Fluorescence Test Methods*, AS6171/3, SAE International, 2016.
- [31] *Techniques for Suspect/Counterfeit EEE Parts Detection by Delid/Decapsulation Physical Analysis Test Methods*, AS6171/4, SAE International, 2016.
- [32] *Techniques for Suspect/Counterfeit EEE Parts Detection by Electrical Test Methods*, AS6171/7, SAE International, 2016.
- [33] M. Tehranipoor, U. Guin, and D. Forte, *Counterfeit Integrated Circuits: Detection and Avoidance*. Cham: Springer International Publishing, 2015. doi: 10.1007/978-3-319-11824-6.
- [34] S. Foster and S. Walters, “A proficiency scheme for counterfeit part testing effectiveness,” presented at the Parts and Material Management Conference, Aurora, Colorado, Mar. 09, 2022. [Online]. Available: <http://www.arctosmeetings.com/agenda/Agendas.asp?sid=7239&ID=pmmc202286473184>
- [35] G. Mura, R. Murru, and G. Martines, “Analysis of counterfeit electronics,” *Microelectronics Reliability*, vol. 114, Nov. 2020, doi: 10.1016/j.microrel.2020.113793.
- [36] G. Mura, R. Murru, and G. Martines, “Analysis of fake amplifiers,” in *IEEE 32nd International Conference on Microelectronics (MIEL)*, Nis, Serbia, Sep. 2021, pp. 131–134. doi: 10.1109/MIEL52794.2021.9569080.
- [37] Y. L. Wang, X. Kuang, C. Huang, and S. P. Li, “Case studied of failure threat caused by counterfeit plastic encapsulated microcircuits,” in *Proceedings of the 20th IEEE International Symposium on the Physical and Failure Analysis of*

- Integrated Circuits (IPFA)*, Suzhou, China, Jul. 2013, pp. 574–577. doi: 10.1109/IPFA.2013.6599226.
- [38] A. Shrivastava and M. Pecht, “Counterfeit capacitors in the supply chain,” *Journal of Materials Science: Materials in Electronics*, vol. 25, no. 2, pp. 645–652, Feb. 2014, doi: 10.1007/s10854-013-1147-9.
- [39] A. Shrivastava, M. H. Azarian, C. Morillo, B. Sood, and M. Pecht, “Detection and reliability risks of counterfeit electrolytic capacitors,” *IEEE Transactions on Reliability*, vol. 63, no. 2, pp. 468–479, Jun. 2014, doi: 10.1109/TR.2014.2315914.
- [40] A. Kanovsky, P. Spanik, and M. Frivaldsky, “Detection of electronic counterfeit components,” in *16th International Scientific Conference on Electric Power Engineering (EPE)*, Kouty nad Desnou, Czech Republic, May 2015, pp. 701–705. doi: 10.1109/EPE.2015.7161140.
- [41] M. H. Azarian, “Analysis of counterfeit ceramic capacitors,” presented at the Symposium on Avoiding, Detecting, and Preventing Counterfeit Electronic Parts, College Park, Maryland, Nov. 07, 2007.
- [42] K. Xiao, D. Forte, Y. Jin, R. Karri, S. Bhunia, and M. Tehranipoor, “Hardware Trojans: lessons learned after one decade of research,” *ACM Trans. Des. Autom. Electron. Syst.*, vol. 22, no. 1, pp. 1–23, Jan. 2017, doi: 10.1145/2906147.
- [43] F. T. Werner, M. Prvulovic, and A. Zajic, “Detection of recycled ICs using backscattering side-channel analysis,” *IEEE Trans. VLSI Syst.*, vol. 30, no. 9, pp. 1244–1255, Sep. 2022, doi: 10.1109/TVLSI.2022.3190236.

- [44] A. Vakil, F. Niknia, A. Mirzaeian, A. Sasan, and N. Karimi, “Learning assisted side channel delay test for detection of recycled ICs,” in *Proceedings of the 26th Asia and South Pacific Design Automation Conference*, Tokyo, Japan, Jan. 2021, pp. 455–462. doi: 10.1145/3394885.3431640.
- [45] Y. Zheng, X. Wang, and S. Bhunia, “SACCI: scan-based characterization through clock phase sweep for counterfeit chip detection,” *IEEE Trans. VLSI Syst.*, vol. 23, no. 5, pp. 831–841, May 2015, doi: 10.1109/TVLSI.2014.2326556.
- [46] P. Samarin and K. Lemke-Rust, “Detection of counterfeit ICs using public identification sequences and side-channel leakage,” *IEEE Trans. Inform. Forensic Secur.*, vol. 14, no. 3, pp. 803–813, Mar. 2019, doi: 10.1109/TIFS.2018.2866333.
- [47] H. Huang, A. Boyer, and S. Ben Dhia, “Electronic counterfeit detection based on the measurement of electromagnetic fingerprint,” *Microelectronics Reliability*, vol. 55, no. 9–10, pp. 2050–2054, Aug. 2015, doi: 10.1016/j.microrel.2015.07.008.
- [48] A. Stern, U. Botero, F. Rahman, D. Forte, and M. Tehranipoor, “EMFORCED: EM-based fingerprinting framework for remarked and cloned counterfeit IC detection using machine learning classification,” *IEEE Trans. VLSI Syst.*, vol. 28, no. 2, pp. 363–375, Feb. 2020, doi: 10.1109/TVLSI.2019.2949733.
- [49] Y. Zheng, A. Basak, and S. Bhunia, “CACI: dynamic current analysis towards robust recycled chip identification,” in *Proceedings of the 51st Annual Design Automation Conference*, San Francisco, CA, USA, Jun. 2014, pp. 1–6. doi: 10.1145/2593069.2593102.

- [50] H. Liu, L. Tinsley, S. Addepalli, X. Liu, A. Starr, and Y. Zhao, “Detectability evaluation of attributes anomaly for electronic components using pulsed thermography,” *Infrared Physics & Technology*, vol. 111, p. 103513, Dec. 2020, doi: 10.1016/j.infrared.2020.103513.
- [51] H. Liu *et al.*, “A novel inspection technique for electronic components using thermography (NITECT),” *Sensors*, vol. 20, no. 17, p. 5013, Sep. 2020, doi: 10.3390/s20175013.
- [52] L. Tinsley, H. Liu, S. Addepalli, W. Lam, and Y. Zhao, “Inspection of electronic component using pulsed thermography,” *Procedia Manufacturing*, vol. 49, pp. 132–138, 2020, doi: 10.1016/j.promfg.2020.07.008.
- [53] M. S. Safa, T. Mosavirik, and S. Tajik, “Counterfeit chip detection using scattering parameter analysis,” in *2023 26th International Symposium on Design and Diagnostics of Electronic Circuits and Systems (DDECS)*, Tallinn, Estonia, May 2023, pp. 99–104. doi: 10.1109/DDECS57882.2023.10139623.
- [54] Y. Yang, L. Wu, X. Zhang, and J. He, “A novel Hardware Trojan detection with chip ID based on relative time delays,” in *2017 11th IEEE International Conference on Anti-counterfeiting, Security, and Identification (ASID)*, Xiamen, Oct. 2017, pp. 163–167. doi: 10.1109/ICASID.2017.8285766.
- [55] Y. Jin and Y. Makris, “Hardware Trojan detection using path delay fingerprint,” in *2008 IEEE International Workshop on Hardware-Oriented Security and Trust*, Anaheim, CA, USA: IEEE, Jun. 2008, pp. 51–57. doi: 10.1109/HST.2008.4559049.

- [56] Jie Li and J. Lach, "At-speed delay characterization for IC authentication and Trojan Horse detection," in *2008 IEEE International Workshop on Hardware-Oriented Security and Trust*, Anaheim, CA, USA: IEEE, Jun. 2008, pp. 8–14. doi: 10.1109/HST.2008.4559038.
- [57] J. Balasch, B. Gierlichs, and I. Verbauwhede, "Electromagnetic circuit fingerprints for Hardware Trojan detection," in *2015 IEEE International Symposium on Electromagnetic Compatibility (EMC)*, Dresden, Germany, Aug. 2015, pp. 246–251. doi: 10.1109/ISEMC.2015.7256167.
- [58] C. He, B. Hou, L. Wang, Y. En, and S. Xie, "A novel hardware Trojan detection method based on side-channel analysis and PCA algorithm," in *2014 10th International Conference on Reliability, Maintainability and Safety (ICRMS)*, Guangzhou, China, Aug. 2014, pp. 1043–1046. doi: 10.1109/ICRMS.2014.7107362.
- [59] J. He, Y. Zhao, X. Guo, and Y. Jin, "Hardware Trojan detection through chip-free electromagnetic side-channel statistical analysis," *IEEE Trans. VLSI Syst.*, vol. 25, no. 10, pp. 2939–2948, Oct. 2017, doi: 10.1109/TVLSI.2017.2727985.
- [60] R. Shende and D. D. Ambawade, "A side channel based power analysis technique for Hardware Trojan detection using statistical learning approach," in *2016 Thirteenth International Conference on Wireless and Optical Communications Networks (WOCN)*, Hyderabad, Telangana State, India, Jul. 2016, pp. 1–4. doi: 10.1109/WOCN.2016.7759894.

- [61] M. Banga and M. S. Hsiao, “A novel sustained vector technique for the detection of Hardware Trojans,” in *2009 22nd International Conference on VLSI Design*, New Delhi: IEEE, Jan. 2009, pp. 327–332. doi: 10.1109/VLSI.Design.2009.22.
- [62] M. Cozzi, J.-M. Galliere, and P. Maurine, “Exploiting phase information in thermal scans for stealthy Trojan detection,” in *2018 21st Euromicro Conference on Digital System Design (DSD)*, Prague, Aug. 2018, pp. 573–576. doi: 10.1109/DSD.2018.00100.
- [63] R. Thummala, K. Winton, L. Flores, E. Redmond, and G. Falco, “Out-of-band power side-channel detection for semiconductor supply chain integrity at scale,” 2026, *arXiv*. doi: 10.48550/ARXIV.2601.01054.
- [64] A. Portune, “Advanced detection of electronic counterfeits (ADEC) for enhanced supply chain assurance against sophisticated counterfeits,” presented at the CALCE/SMTA Symposium on Counterfeit Parts and Materials, College Park, Maryland, Apr. 19, 2013.
- [65] A. Portune, “Counterfeit electronic part avoidance and mitigation through advanced detection of electronic counterfeits (ADEC),” presented at the Diminishing Manufacturing Sources and Material Shortages (DMSMS), 2018.
- [66] A. Nechiyl, J. J. McCue, R. Lee, and G. Chapman, “Classification and novelty detection of tampered ICs using ResCav,” *J Fail. Anal. and Preven.*, vol. 24, no. 5, pp. 2105–2112, Oct. 2024, doi: 10.1007/s11668-024-01998-4.
- [67] K. Liszewski and T. Bergman, “Battelle barricade—microelectronic device authentication and counterfeit detection utilizing power analysis,” presented at the

- ISTFA 2017, Pasadena, California, USA, Nov. 2017, pp. 79–87. doi: 10.31399/asm.cp.istfa2017p0079.
- [68] T. D. Bergman, C. P. Manager, and K. T. Liszewski, “Battelle barricade: a nondestructive electronic component authentication and counterfeit detection technology,” in *2016 IEEE Symposium on Technologies for Homeland Security (HST)*, Waltham, MA, USA, May 2016, pp. 1–6. doi: 10.1109/THS.2016.7568901.
- [69] PFP Cybersecurity, “Supply chain protection a white paper on counterfeit detection.” 2015.
- [70] E. Koziel *et al.*, “Side channel authenticity discriminant analysis for device class identification,” presented at the 2016 IEEE Int. Conf. on Image Processing, Sep. 2016.
- [71] K. Liszewski and M. Brewer, “System and method for generating test vectors,” U.S. Patent 10789550, Sep. 29, 2020
- [72] P. Tangyunyong, J. Buetler, E. Cole, and G. Loubriel, “Defect screening method for electronic circuits and circuit components using power spectrum analysis,” U.S. Patent 10145894, Dec. 04, 2018
- [73] J. T. Christiansen, J. A. Walraven, F. A. Coleman, E. H. Smith, and E. I. Cole, “Heterogeneous integrated failure analysis using radio frequency signal detection/injection with power spectrum analysis,” presented at the ISTFA 2024, San Diego, California, USA, Oct. 2024, pp. 501–508. doi: 10.31399/asm.cp.istfa2024p0501.

- [74] C. Bohman *et al.*, “Multi-tone analysis for the authentication of electronic devices,” in *NAECON 2023 - IEEE National Aerospace and Electronics Conference*, Dayton, OH, USA: IEEE, Aug. 2023, pp. 54–58. doi: 10.1109/naecon58068.2023.10365951.
- [75] M. H. Azarian, D. Das, and P. Campbell, “Machine vision pilot (MVP) and microelectronic authenticity and security, evaluation and research (MASER),” 2020. [Online]. Available: <https://calce.umd.edu/news/story/final-report-of-calce-counterfeit-microelectronics-study-for-dmea-released-for-public-access>
- [76] P. Tangyunyong and G. M. Loubriel, “Power spectrum analysis (PSA) for counterfeit electronics,” presented at the CALCE/SMTA Symposium on Counterfeit Parts and Materials, College Park, Maryland, 2017.
- [77] P. Tangyunyong and G. M. Loubriel, “Power spectrum analysis (PSA) for counterfeit and aging detection,” Feb. 2018. [Online]. Available: <https://www.osti.gov/biblio/1806510>
- [78] P. Tangyunyong *et al.*, “Power spectrum analysis (PSA),” presented at the 43rd International Symposium for Testing and Failure Analysis, 2017. [Online]. Available: <https://www.osti.gov/servlets/purl/1467997>
- [79] L. House, “Battelle barricade electronic component authentication technology,” 2014.
- [80] L. Milechin, E. Koziel, M. Vai, R. Khazan, K. Bergevin, and P. Comer, “Automatic feature selection and improved classification in SICADA counterfeit electronic detection,” 2017.

- [81] B. Copos, “Modeling systems using side channel information,” Ph.D. dissertation, University of California Davis, California, 2017.
- [82] C. Gonzalez, “Power fingerprinting for integrity assessment of embedded systems dissertation,” Ph.D. dissertation, Virginia Polytechnic Institute and State University, Virginia, 2011.
- [83] I. Jolliffe, *Principal Component Analysis*. in Springer Series in Statistics. New York, USA: Springer-Verlag, 2002. doi: 10.1007/b98835.
- [84] W. E. Cobb, E. D. Laspe, R. O. Baldwin, M. A. Temple, and Y. C. Kim, “Intrinsic physical-layer authentication of integrated circuits,” *IEEE Trans.Inform.Forensic Secur.*, vol. 7, no. 1, pp. 14–24, Feb. 2012, doi: 10.1109/TIFS.2011.2160170.
- [85] L. House and D. Engelhart, “Electronic component classification,” U.S. Patent 9759757, Sep. 12, 2017
- [86] *Failure Mechanisms and Models for Semiconductor Devices*, JEP122H, JEDEC, 2016.
- [87] S. J. Babu, F. Hu, L. Zhu, S. Singhal, and X. Guo, “Extending silicon lifetime: a review of design techniques for reliable integrated circuits,” 2025, *arXiv preprint*. doi: 10.48550/ARXIV.2503.21165.
- [88] P. Viswanadham and P. Singh, *Failure Modes and Mechanisms in Electronic Packages*. Boston, MA: Springer US, 1998. doi: 10.1007/978-1-4615-6029-6.
- [89] Y. Ren *et al.*, “Overview on radiation damage effects and protection techniques in microelectronic devices,” *Science and Technology of Nuclear Installations*, vol. 2024, no. 1, p. 3616902, Jan. 2024, doi: 10.1155/2024/3616902.

- [90] R. Velazco, P. Fouillat, and R. Reis, Eds., *Radiation Effects on Embedded Systems*. Dordrecht: Springer Netherlands, 2007. doi: 10.1007/978-1-4020-5646-8.
- [91] C. Ostrom, A. Manis, M. Matney, and J. Liou, “Orbital Debris Quarterly News,” National Aeronautics and Space Administration (NASA), Houston, TX, USA, 20250008945, Sep. 2025. Accessed: Dec. 04, 2025. [Online]. Available: <https://ntrs.nasa.gov/citations/20250008945?>
- [92] G. Brunetti, G. Campiti, M. Tagliente, and C. Ciminelli, “COTS devices for space missions in LEO,” *IEEE Access*, vol. 12, pp. 76478–76514, 2024, doi: 10.1109/ACCESS.2024.3405373.
- [93] I. Jun *et al.*, “A review on radiation environment pathways to impacts: radiation effects, relevant empirical environment models, and future needs,” *Advances in Space Research*, p. S0273117724003296, Apr. 2024, doi: 10.1016/j.asr.2024.03.079.
- [94] J. W. Evans and J. Y. Evans, Eds., *Product Integrity and Reliability in Design*. London: Springer London, 2001. doi: 10.1007/978-1-4471-0253-3.
- [95] M. I. Bâzu and T. I. Băjenescu, *Failure Analysis: A Practical Guide for Manufacturers of Electronic Components and Systems*. Chichester, West Sussex, U.K: Wiley, 2011. doi: 10.1002/9781119990093.
- [96] *Test Method Standard Environmental Test Methods for Microcircuits Part I: Test Methods 1000-1999*, MIL-STD-883-1 Method 1005.11 Steady-State Life.

- [97] E. Bender, J. B. Bernstein, and D. S. Boning, “Modern trends in microelectronics packaging reliability testing,” *Micromachines*, vol. 15, no. 3, p. 398, Mar. 2024, doi: 10.3390/mi15030398.
- [98] L. Lanzieri, L. Butkowski, J. Kral, G. Fey, H. Schlarb, and T. C. Schmidt, “Studying the degradation of propagation delay on FPGAs at the European XFEL,” Jul. 09, 2024, *arXiv*: arXiv:2407.06643. doi: 10.1109/DSD64264.2024.00018.
- [99] K. Sutherlin, “Investigation of electromagnetic signatures of a FPGA using an April EM-Isight system,” M.S. thesis, Air Force Institute of Technology, Ohio, 2015. [Online]. Available: https://scholar.afit.edu/etd/239/?utm_source=scholar.afit.edu%2Fetd%2F239&utm_medium=PDF&utm_campaign=PDFCoverPages
- [100] N. Karimi, T. Moos, and A. Moradi, “Exploring the effect of device aging on static power analysis attacks,” *TCHES*, pp. 233–256, May 2019, doi: 10.46586/tches.v2019.i3.233-256.
- [101] X. Guo, N. Karimi, F. Regazzoni, C. Jin, and R. Karri, “Simulation and analysis of negative-bias temperature instability aging on power analysis attacks,” in *2015 IEEE International Symposium on Hardware Oriented Security and Trust (HOST)*, Washington, DC, USA: IEEE, May 2015, pp. 124–129. doi: 10.1109/HST.2015.7140250.
- [102] Y. Kaneko, Y. Hayashi, and N. Homma, “Experimental evaluation for detecting aging effect on microcontrollers based on side-channel analysis,” in *2024 14th International Workshop on the Electromagnetic Compatibility of Integrated*

- Circuits (EMC Compo)*, Torino, Italy: IEEE, Oct. 2024, pp. 1–5. doi: 10.1109/EMCCompo61192.2024.10742051.
- [103] K. Takasaki, K. Hasegawa, R. Kida, and N. Togawa, “An anomalous behavior detection method for IoT devices by extracting application-specific power behaviors,” in *2020 IEEE 26th International Symposium on On-Line Testing and Robust System Design (IOLTS)*, Napoli, Italy: IEEE, Jul. 2020, pp. 1–4. doi: 10.1109/IOLTS50870.2020.9159732.
- [104] K. Hasegawa, K. Chikamatsu, and N. Togawa, “Empirical evaluation on anomaly behavior detection for low-cost micro-controllers utilizing accurate power analysis,” in *2019 IEEE 25th International Symposium on On-Line Testing and Robust System Design (IOLTS)*, Rhodes, Greece: IEEE, Jul. 2019, pp. 54–57. doi: 10.1109/IOLTS.2019.8854456.
- [105] T. Moos, A. Moradi, and B. Richter, “Static power side-channel analysis—an investigation of measurement factors,” *IEEE Trans. VLSI Syst.*, vol. 28, no. 2, pp. 376–389, Feb. 2020, doi: 10.1109/TVLSI.2019.2948141.
- [106] J. E. Jackson and G. S. Mudholkar, “Control procedures for residuals associated with principal component analysis,” *Technometrics*, vol. 21, no. 3, pp. 341–349, Aug. 1979, doi: 10.1080/00401706.1979.10489779.
- [107] H. Hotelling, “The generalization of student’s ratio,” *Ann. Math. Statist.*, vol. 2, no. 3, pp. 360–378, Aug. 1931, doi: 10.1214/aoms/1177732979.
- [108] “EDU33210 series 20 MHz function/arbitrary waveform generators.” Keysight.
- [109] “InfiniiVision 7000B series oscilloscopes.” Keysight.

- [110] P. Tangyunyong, “CTAP REPORT. Commercialization of power spectrum analysis (PSA) technology,” Sandia National Laboratories (SNL), Albuquerque, NM, and Livermore, CA (United States), Jun. 2022. doi: 10.2172/1874426.
- [111] “Analog discovery 3 specifications - digilent reference.” Accessed: May 20, 2025. [Online]. Available: <https://digilent.com/reference/test-and-measurement/analog-discovery-3/specifications?srsId=AfmBOorv6dv6TWo2P7fCTJcIGuMOxSAn7O7xmQNwhKXoQTz1g1vxQb0P>
- [112] A. V. Oppenheim, A. S. Willsky, and S. H. Nawab, *Signals & Systems*, 2nd ed. in Prentice-Hall signal processing series. Upper Saddle River, N.J: Prentice Hall, 1997.
- [113] R. G. Lyons, *Understanding Digital Signal Processing*, 3rd ed. Upper Saddle River, NJ: Prentice Hall, 2011.
- [114] L. G. Nonato and M. Aupetit, “Multidimensional projection for visual analytics: linking techniques with distortions, tasks, and layout enrichment,” *IEEE Trans. Visual. Comput. Graphics*, vol. 25, no. 8, pp. 2650–2673, Aug. 2019, doi: 10.1109/TVCG.2018.2846735.
- [115] M. L. Duc, P. Bilik, and R. Martinek, “Harmonics signal feature extraction techniques: a review,” *Mathematics*, vol. 11, no. 8, p. 1877, Apr. 2023, doi: 10.3390/math11081877.
- [116] X. Zhang and Z. W. Ras, “Analysis of sound features for music timbre recognition,” in *2007 International Conference on Multimedia and Ubiquitous*

- Engineering (MUE'07)*, Seoul: IEEE, Apr. 2007, pp. 3–8. doi: 10.1109/MUE.2007.85.
- [117] R. B. Randall, *Vibration-Based Condition Monitoring: Industrial, Automotive and Aerospace Applications*, Second edition. Hoboken, NJ: Wiley, 2022. doi: 10.1002/9781119477631.
- [118] A. Patil, B. Mishra, and S. Harsha, “Fault diagnosis of rolling element bearing using autonomous harmonic product spectrum method,” *Proceedings of the Institution of Mechanical Engineers, Part K: Journal of Multi-body Dynamics*, vol. 235, no. 3, pp. 396–411, Sep. 2021, doi: 10.1177/1464419321994986.
- [119] G. Peeters, B. L. Giordano, P. Susini, N. Misdariis, and S. McAdams, “The timbre toolbox: extracting audio descriptors from musical signals,” *The Journal of the Acoustical Society of America*, vol. 130, no. 5, pp. 2902–2916, Nov. 2011, doi: 10.1121/1.3642604.
- [120] M. Christ, N. Braun, J. Neuffer, and A. W. Kempa-Liehr, “Time series Feature extraction on basis of scalable hypothesis tests (tsfresh – a python package),” *Neurocomputing*, vol. 307, pp. 72–77, Sep. 2018, doi: 10.1016/j.neucom.2018.03.067.
- [121] M. Barandas *et al.*, “TSFEL: time series feature extraction library,” *SoftwareX*, vol. 11, p. 100456, Jan. 2020, doi: 10.1016/j.softx.2020.100456.
- [122] A. Zytek, I. Arnaldo, D. Liu, L. Berti-Equille, and K. Veeramachaneni, “The need for interpretable features: motivation and taxonomy,” *SIGKDD Explor. Newsl.*, vol. 24, no. 1, pp. 1–13, Jun. 2022, doi: 10.1145/3544903.3544905.

- [123] S. Kazazis, P. Depalle, and S. McAdams, “Interval and ratio scaling of spectral audio descriptors,” *Front. Psychol.*, vol. 13, p. 835401, Mar. 2022, doi: 10.3389/fpsyg.2022.835401.
- [124] *IEEE Standard for Harmonic Control in Electric Power Systems*. doi: 10.1109/IEEESTD.2022.9848440.
- [125] G. Peeters, “A large set of audio features for sound description (similarity and classification) in the CUIDADO project,” Jan. 2004, [Online]. Available: http://recherche.ircam.fr/anasyne/peeters/ARTICLES/Peeters_2003_cuidadoaudiofeatures.pdf
- [126] L. R. Rabiner and R. W. Schafer, *Digital Processing of Speech Signals*, 13. [print.]. in Prentice-Hall signal processing series. Englewood Cliffs, NJ: Prentice-Hall, 1978.
- [127] G. E. P. Box, G. M. Jenkins, G. C. Reinsel, and G. M. Ljung, *Time Series Analysis: Forecasting and Control*, Fifth edition. in Wiley series in probability and statistics. Hoboken, New Jersey: John Wiley & Sons, Inc, 2016.
- [128] S. Mallat, *A Wavelet Tour of Signal Processing: The Sparse Way*, 3rd ed. Amsterdam Boston: Elsevier/Academic Press, 2009.
- [129] A. V. Oppenheim and R. W. Schafer, “Dsp history - from frequency to quefrency: a history of the cepstrum,” *IEEE Signal Process. Mag.*, vol. 21, no. 5, pp. 95–106, Sep. 2004, doi: 10.1109/MSP.2004.1328092.
- [130] D. G. Childers, D. P. Skinner, and R. C. Kemerait, “The cepstrum: a guide to processing,” *Proc. IEEE*, vol. 65, no. 10, pp. 1428–1443, 1977, doi: 10.1109/PROC.1977.10747.

- [131] G. Tzanetakis and P. Cook, “Musical genre classification of audio signals,” *IEEE Trans. Speech Audio Process.*, vol. 10, no. 5, pp. 293–302, Jul. 2002, doi: 10.1109/TSA.2002.800560.
- [132] A. K. Jain, M. N. Murty, and P. J. Flynn, “Data clustering: a review,” *ACM Comput. Surv.*, vol. 31, no. 3, pp. 264–323, Sep. 1999, doi: 10.1145/331499.331504.
- [133] J. MacQueen, “Some methods for classification and analysis of multivariate observations.” 1967.
- [134] U. von Luxburg, “A tutorial on spectral clustering,” 2007, doi: 10.1007/s11222-007-9033-z.
- [135] E. Schubert, J. Sander, M. Ester, H. P. Kriegel, and X. Xu, “DBSCAN revisited, revisited: why and how you should (still) use DBSCAN,” *ACM Trans. Database Syst.*, vol. 42, no. 3, pp. 1–21, Sep. 2017, doi: 10.1145/3068335.
- [136] D. Müllner, “Modern hierarchical, agglomerative clustering algorithms,” Sep. 12, 2011, *arXiv*: arXiv:1109.2378. doi: 10.48550/arXiv.1109.2378.
- [137] R. J. G. B. Campello, D. Moulavi, and J. Sander, “Density-Based Clustering Based on Hierarchical Density Estimates,” in *Advances in Knowledge Discovery and Data Mining*, vol. 7819, J. Pei, V. S. Tseng, L. Cao, H. Motoda, and G. Xu, Eds., in Lecture Notes in Computer Science, vol. 7819. , Berlin, Heidelberg: Springer Berlin Heidelberg, 2013, pp. 160–172. doi: 10.1007/978-3-642-37456-2_14.

- [138] X. Ran, Y. Xi, Y. Lu, X. Wang, and Z. Lu, “Comprehensive survey on hierarchical clustering algorithms and the recent developments,” *Artif Intell Rev*, vol. 56, no. 8, pp. 8219–8264, Aug. 2023, doi: 10.1007/s10462-022-10366-3.
- [139] B. J. Frey and D. Dueck, “Clustering by passing messages between data points,” *Science*, vol. 315, no. 5814, pp. 972–976, Feb. 2007, doi: 10.1126/science.1136800.
- [140] E. Schubert, “Stop using the elbow criterion for k-means and how to choose the number of clusters instead,” *SIGKDD Explor. Newsl.*, vol. 25, no. 1, pp. 36–42, Jun. 2023, doi: 10.1145/3606274.3606278.
- [141] R. Tibshirani, G. Walther, and T. Hastie, “Estimating the number of clusters in a data set via the gap statistic,” *Journal of the Royal Statistical Society Series B: Statistical Methodology*, vol. 63, no. 2, pp. 411–423, Jul. 2001, doi: 10.1111/1467-9868.00293.
- [142] P. J. Rousseeuw, “Silhouettes: a graphical aid to the interpretation and validation of cluster analysis,” *Journal of Computational and Applied Mathematics*, vol. 20, pp. 53–65, Nov. 1987, doi: 10.1016/0377-0427(87)90125-7.
- [143] D. L. Davies and D. W. Bouldin, “A cluster separation measure,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. PAMI-1, no. 2, pp. 224–227, Apr. 1979, doi: 10.1109/TPAMI.1979.4766909.
- [144] T. Calinski and J. Harabasz, “A dendrite method for cluster analysis,” *Comm. in Stats. - Theory & Methods*, vol. 3, no. 1, pp. 1–27, 1974, doi: 10.1080/03610927408827101.

- [145] J. C. Dunn, “A fuzzy relative of the ISODATA process and its use in detecting compact well-separated clusters,” *Journal of Cybernetics*, vol. 3, no. 3, pp. 32–57, Jan. 1973, doi: 10.1080/01969727308546046.
- [146] D. Moulavi, P. A. Jaskowiak, R. J. G. B. Campello, A. Zimek, and J. Sander, “Density-based clustering validation,” in *Proceedings of the 2014 SIAM International Conference on Data Mining*, Society for Industrial and Applied Mathematics, Apr. 2014, pp. 839–847. doi: 10.1137/1.9781611973440.96.
- [147] G. McLachlan and D. Peel, *Finite Mixture Models*, 1st ed. in Wiley Series in Probability and Statistics. Wiley, 2000. doi: 10.1002/0471721182.
- [148] H. Akaike, “A new look at the statistical model identification,” *IEEE Trans. Automat. Contr.*, vol. 19, no. 6, pp. 716–723, Dec. 1974, doi: 10.1109/TAC.1974.1100705.
- [149] M. Ebrahimi Warkiani and M. H. Moattar, “A comprehensive survey on recent feature selection methods for mixed data: challenges, solutions and future directions,” *Neurocomputing*, vol. 623, p. 129372, Mar. 2025, doi: 10.1016/j.neucom.2025.129372.
- [150] R. Dwivedi, A. Tiwari, N. Bharill, M. Ratnaparkhe, and A. K. Tiwari, “A taxonomy of unsupervised feature selection methods including their pros, cons, and challenges,” *J Supercomput*, vol. 80, no. 16, pp. 24212–24240, Nov. 2024, doi: 10.1007/s11227-024-06368-3.
- [151] T. Hastie, R. Tibshirani, and J. H. Friedman, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, 2nd ed. in Springer series in statistics. New York, NY: Springer, 2009.

- [152] R. De Maesschalck, D. Jouan-Rimbaud, and D. L. Massart, “The Mahalanobis distance,” *Chemometrics and Intelligent Laboratory Systems*, vol. 50, no. 1, pp. 1–18, Jan. 2000, doi: 10.1016/s0169-7439(99)00047-7.
- [153] R. Ott, M. Skouson, A. Jennings, C. Bohman, and C. Eakins, “A separability metric for second order effect system performance evaluation,” presented at the Government Microcircuit Applications & Critical Technology (GOMACTech), San Diego, California, USA, Mar. 2020.
- [154] I. Goodfellow, A. Courville, and Y. Bengio, “Chapter 14 Autoencoders,” in *Deep learning*, in Adaptive computation and machine learning. , Cambridge, Massachusetts: The MIT Press, 2016.
- [155] B. Schölkopf and A. J. Smola, *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. in Adaptive computation and machine learning. Cambridge, Mass: MIT Press, 2002.
- [156] B. Schölkopf, R. C. Williamson, A. Smola, J. Shawe-Taylor, and J. Platt, “Support vector method for novelty detection,” in *Advances in Neural Information Processing Systems*, S. Solla, T. Leen, and K. Müller, Eds., MIT Press, 1999. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/1999/file/8725fb777f25776ffa9076e44fcfd776-Paper.pdf
- [157] M. M. Breunig, H.-P. Kriegel, R. T. Ng, and J. Sander, “LOF: identifying density-based local outliers,” *SIGMOD Rec.*, vol. 29, no. 2, pp. 93–104, Jun. 2000, doi: 10.1145/335191.335388.

- [158] J. Cohen, *Statistical Power Analysis for the Behavioral Sciences*, 0 ed. Routledge, 2013. doi: 10.4324/9780203771587.
- [159] K. A. Ashman, C. M. Bird, and S. E. Zepf, “Detecting bimodality in astronomical datasets,” *The Astronomical Journal*, vol. 108, p. 2348, Dec. 1994, doi: 10.1086/117248.
- [160] Toshiba, “Quad 2-input NOR gate TC74HC02AP, TC74HC02AF, TC74HC02AFN.” 2007.
- [161] *Steady-State Temperature-Humidity Bias Life Test*, JESD22-A101D.01, JEDEC Solid State Technology Association, 2021.
- [162] *Temperature, Bias, and Operating Life*, JESD22-A108G, JEDEC Solid State Technology Association, 2022.
- [163] “2.4. Gauge R & R Studies,” in *NIST/SEMATECH e-Handbook of Statistical Methods*. Accessed: Jan. 18, 2024. [Online]. Available: <https://www.itl.nist.gov/div898/handbook/mpc/section4/mpc4.htm>
- [164] O. Senvar and S. U. O. Firat, “An overview of capability evaluation of measurement systems and gauge repeatability and reproducibility studies,” *Int. J. Metrol. Qual. Eng.*, vol. 1, no. 2, pp. 121–127, 2010, doi: 10.1051/ijmqe/2010022.
- [165] L. Čepová, A. Kovacikova, R. Cep, P. Klaput, and O. Mizera, “Measurement system analyses - gauge repeatability and reproducibility methods,” *Measurement Science Review*, vol. 18, pp. 20–27, Feb. 2018, doi: 10.1515/msr-2018-0004.
- [166] M. Naouss and F. Marc, “FPGA LUT delay degradation due to HCI: experiment and simulation results,” *Microelectronics Reliability*, vol. 64, pp. 31–35, Sep. 2016, doi: 10.1016/j.microrel.2016.07.048.

- [167] K. Kang, S. Gangwal, S. P. Park, and R. Kaushik, “NBTI induced performance degradation in logic and memory circuits: how effectively can we approach a reliability solution?,” in *2008 Asia and South Pacific Design Automation Conference*, Seoul: IEEE, Jan. 2008, pp. 726–731. doi: 10.1109/ASPDAC.2008.4484047.
- [168] J. R. Schwank, M. R. Shaneyfelt, and P. E. Dodd, “Radiation hardness assurance testing of microelectronic devices and integrated circuits: radiation environments, physical mechanisms, and foundations for hardness assurance,” *IEEE Trans. Nucl. Sci.*, vol. 60, no. 3, pp. 2074–2100, Jun. 2013, doi: 10.1109/TNS.2013.2254722.
- [169] R. H. Maurer, M. E. Fraeman, M. N. Martin, and D. R. Roth, “Harsh environments: space radiation environment, effects, and mitigation,” *Johns Hopkins APL Technical Digest*, vol. 28, no. 1, 2008.
- [170] J. R. Schwank *et al.*, “Radiation effects in MOS oxides,” *IEEE Trans. Nucl. Sci.*, vol. 55, no. 4, pp. 1833–1853, Aug. 2008, doi: 10.1109/TNS.2008.2001040.
- [171] M. V. O’Byrne *et al.*, “NASA Goddard space flight center’s current radiation effects test results,” in *2023 IEEE Radiation Effects Data Workshop (REDW) (in conjunction with 2023 NSREC)*, Kansas City, MO, USA: IEEE, Jul. 2023, pp. 1–13. doi: 10.1109/REDW61050.2023.10265859.
- [172] L. Finazzi *et al.*, “Total ionizing dose measurements in small satellites in LEO using LabOSat-01,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 1064, p. 169344, Jul. 2024, doi: 10.1016/j.nima.2024.169344.

- [173] O. Gutiérrez, M. Prieto, A. Perales-Eceiza, A. Ravanbakhsh, M. Basile, and D. Guzmán, “Toward the use of electronic commercial off-the-shelf devices in space: assessment of the true radiation environment in low earth orbit (LEO),” *Electronics*, vol. 12, no. 19, p. 4058, Sep. 2023, doi: 10.3390/electronics12194058.
- [174] T. R. Oldham and F. B. McLean, “Total ionizing dose effects in MOS oxides and devices,” *IEEE Trans. Nucl. Sci.*, vol. 50, no. 3, pp. 483–499, Jun. 2003, doi: 10.1109/TNS.2003.812927.
- [175] M. Obrenovic, M. Pejovic, D. Lazarevic, and N. Kartalovic, “The effects induced by the gamma-ray responsible for the threshold voltage shift of commercial p-channel power VDMOSFET,” *Nucl Technol Radiat Prot*, vol. 33, no. 1, pp. 81–86, 2018, doi: 10.2298/NTRP1801081O.
- [176] S. Rathod, S. Dasgupta, and A. Saxena, “Radiation effects in MOS-based devices and circuits: a review,” *IETE Tech Rev*, vol. 28, no. 6, p. 451, 2011, doi: 10.4103/0256-4602.90747.
- [177] H. J. Barnaby, “Total dose effects in linear bipolar integrated circuits,” *Int. J. Hi. Spe. Ele. Syst.*, vol. 14, no. 02, pp. 519–541, Jun. 2004, doi: 10.1142/S0129156404002491.
- [178] E. KendiR and Ş. Yaltkaya, “Gamma ray effect on op-amp 741 and timer 555 ic,” *UUJFE*, pp. 193–206, Apr. 2023, doi: 10.17482/uumfd.1094384.