Side-Channel Attacks Using On-Chip Reliability Instruments

CHENXIN YANG
MASTER´S THESIS
DEPARTMENT OF ELECTRICAL AND INFORMATION TECHNOLOGY
FACULTY OF ENGINEERING | LTH | LUND UNIVERSITY

Side-Channel Attacks Using On-Chip Reliability Instruments

Chenxin Yang
ya3268ch-s@student.lu.se

Department of Electrical and Information Technology
Lund University

Supervisor: Erik Larsson

Examiner: Pietro Andreani

June 1 , 2026

Printed in Sweden
Department of Electrical and Information Technology, Lund

Abstract

Modern integrated circuits contain extensive on-chip instrumentation—Ring Oscillators (ROs), temperature sensors, voltage monitors, and delay-measurement circuits—originally intended for reliability monitoring and process calibration. The same physical sensitivity that makes these instruments effective sensors also makes them potential side-channel receivers. An adversary who can place or exploit such instrumentation near sensitive logic may be able to recover secret-dependent information without requiring external measurement equipment or physical probing. This threat is relevant across FPGA, ASIC, and SoC platforms, yet the boundary conditions under which on-chip instrumentation can reliably leak key information remain incompletely characterised.

This study uses a Xilinx Artix-7 FPGA as a controllable proof-of-concept platform to examine how key-bit observability degrades as the coupling path between a secret source and an on-chip RO listener is systematically weakened. Three completed experiment families span the coupling continuum from a direct electrical connection through a GPIO-level threshold-limited path to pure physical proximity without intentional logical wiring. The direct-connection baseline matches 255 of 256 streamed AES-256 key-bit decisions (99.61%) using on-chip digital capture alone. The GPIO-level path shows an abrupt threshold-limited failure: with a 400 Ω pull-down, agreement remains at 100% for 0–300 Ω series resistance and collapses to an all-zeros classifier at 400 Ω and above, where the returned voltage no longer crosses the FPGA input-buffer threshold. The proximity-only matrix—24 conditions covering three placement distances and eight RO lengths—shows that calibrated key-bit separability is strongly distance- and RO-length-dependent: near placement reaches 99.61% supervised separability at N=35 and 100% for N=45–75, while medium and far placement remain partially distinguishable or weak.

The central finding is that on-chip reliability instruments can act as effective side-channel receivers not only through explicit electrical paths but also, under favourable physical conditions, through unintended proximity coupling. The practical risk is therefore not limited to designs containing

deliberate Trojan wiring; it extends to any scenario in which an attacker can place or activate monitoring instrumentation near sensitive routing. The results motivate treating the placement and configuration of on-chip reliability instruments as security-relevant design decisions, enforcing physical separation between cryptographic key nets and untrusted instrumentation structures, and including routing-level proximity analysis in security verification flows.

Popular Science Summary

Modern computer chips contain thousands of tiny monitoring circuits that help engineers track temperature, voltage, and manufacturing quality. These circuits—including a common structure called the Ring Oscillator (RO)—are essential for keeping chips reliable. However, the same physical sensitivity that makes them good sensors also creates a security risk: if an attacker can place or activate such a circuit near sensitive information, it may inadvertently leak secret data.

This study investigates that risk by asking how much secret information an on-chip monitor can recover under different coupling conditions. To answer this question in a controlled setting, the researchers used an FPGA—a reconfigurable chip that allows precise experimental control—as a proof-of-concept platform. A fixed AES encryption key served as the secret data source, and an RO acted as the on-chip listener.

In the first and strongest scenario, the selected key bit was wired directly to the RO. The monitor almost perfectly distinguished whether the bit was 0 or 1, matching 255 of 256 key bits correctly. In the second scenario, the direct wire was replaced by an external path through resistors. The monitor continued to work perfectly as long as the electrical signal remained strong enough, but it failed abruptly once the voltage dropped below the chip's input recognition threshold. This shows that weakening a sensitive signal path does not necessarily remove the leakage risk.

The third and most challenging scenario removed any intentional wire connection entirely. The researchers asked whether physical proximity alone—through electrical cross-talk between nearby wires—could still allow the RO to distinguish key-bit values. The answer was nuanced: a short nearby RO could not reliably separate the bit values, but a longer nearby RO could. When the RO was sufficiently long and placed very close to the secret signal, it achieved perfect separability. At greater distances, the effect weakened and became unreliable.

The broader implication is that the security risk from on-chip monitoring circuits is not limited to designs containing deliberate malicious wiring. Even unintended physical proximity between a secret signal and a monitor-

ing structure can create measurable information leakage under favourable conditions. This suggests that chip designers should treat the placement of monitoring circuits as a security decision, not merely a reliability one, and should keep such structures physically separated from cryptographic key material.

Acknowledgements

I thank my supervisor Erik Larsson for his guidance, my examiner Pietro Andreani, and Fanny Lundberg for her foundational work. I also thank Joel Åhlund for his valuable assistance as a teaching assistant. I am grateful to my family and friends for their support.

Contents

List of Figures

List of Tables

List of Abbreviations

AES	Advanced Encryption Standard
ASIC	Application-Specific Integrated Circuit
CLB	Configurable Logic Block
CPA	Correlation Power Analysis
DRC	Design Rule Check
FPGA	Field-Programmable Gate Array
GPIO	General-Purpose Input/Output
HCI	Hot Carrier Injection
HT	Hardware Trojan
ILA	Integrated Logic Analyzer
IP	Intellectual Property
LUT	Look-Up Table
NBTI	Negative Bias Temperature Instability
RO	Ring Oscillator
RTL	Register-Transfer Level
SoC	System-on-Chip
XADC	Xilinx Analog-to-Digital Converter
XDC	Xilinx Design Constraints

Chapter 1

Introduction

1.1 Context and Motivation

1.1.1 The Dual-Use Dilemma of On-Chip Reliability Instruments

Modern integrated circuits contain extensive on-chip instrumentation originally designed for reliability monitoring, process calibration, and performance tuning. Ring Oscillators (ROs) measure local delay variations and detect ageing effects such as Negative Bias Temperature Instability (NBTI) and Hot Carrier Injection (HCI). Temperature sensors track thermal profiles. Voltage monitors observe supply droop. Timing slack detectors flag setup-time violations. These structures are now standard in both FPGA and ASIC design flows, embedded by the foundry, the IP vendor, or the system integrator to maintain operational health over the device lifetime.

The same physical sensitivity that makes these instruments effective reliability monitors also creates a security vulnerability. An RO, for example, oscillates at a frequency determined by the propagation delay of its inverter stages, which in turn depends on the local supply voltage and temperature. A small voltage fluctuation caused by neighbouring switching activity produces a measurable frequency shift. This effect is precisely what makes the RO useful for detecting silicon degradation—but it also means that an RO placed near a cryptographic core can act as an unintended side-channel receiver. If an attacker can observe the instrument output, either directly or through intermediate logic, fluctuations caused by secret-dependent activity may leak confidential information [8].

This dual-use property is not unique to ROs. Any on-chip sensor whose readout correlates with local electrical activity—whether delay, voltage, temperature, or leakage current—can in principle be repurposed as a side-channel observation point. The threat is therefore not limited to designs containing deliberately inserted malicious logic. Even legitimately placed reliability instrumentation can become an attack surface if its placement,

routing, or readout access is not subject to security-aware design constraints.

1.1.2 Hardware Security and the Supply-Chain Attack Surface

The risk is compounded by the disaggregated nature of modern hardware development. A System-on-Chip (SoC) may integrate IP cores from multiple vendors, be manufactured by an external foundry, and be deployed in environments where the end user cannot inspect the physical silicon. At each stage, a malicious actor could insert additional monitoring structures or exploit existing ones to leak information [1, 4, 5]. In multi-tenant FPGA environments, where independent designs share the same physical device, a malicious tenant may attempt to place instrumentation that observes neighbouring tenants' activity through shared routing, power distribution, or thermal resources [8, 12, 13].

The central security question raised by this dual-use dilemma is therefore not whether a specific instrument can be exploited in isolation, but rather how much secret information can leak through on-chip instrumentation under different physical coupling conditions, and what design practices can bound that leakage.

1.2 Problem Formulation

1.2.1 Side-Channel Attacks

Side-channel attacks exploit physical emanations during the execution of cryptographic algorithms. Mathematical cryptanalysis targets the algorithm's theoretical strength. By contrast, side-channel attacks use measurements of power consumption, electromagnetic radiation, or execution time to recover secret keys [16, 17].

Traditional side-channel attacks require physical access to the device and specialised measurement equipment such as high-bandwidth oscilloscopes. However, a newer class of internal side-channel attacks uses on-chip logic as the measurement instrument. In this model, the attacker does not need physical access—they only need to insert or exploit existing monitoring logic on the same die as the victim.

1.2.2 Prior Work and Its Limitations

The study of on-chip instrumentation as a side-channel vector has developed along several related lines. Understanding these lines is essential for

positioning the present work.

Traditional side-channel attacks. Kocher et al. [16] demonstrated that power consumption traces measured at the device package can reveal secret keys through Differential Power Analysis (DPA). Subsequent work extended this to Correlation Power Analysis (CPA) [2, 17], electromagnetic emanation [9], and template attacks [3]. These methods require physical access to the device and specialised measurement equipment such as high-bandwidth oscilloscopes or near-field probes. A newer class of *internal* side-channel attacks uses logic already present on the die as the measurement instrument, removing the need for external equipment and enabling attacks from software or malicious IP [16, 17].

On-chip instruments as attack tools. The repurposing of legitimate monitoring structures for attack has been demonstrated in multiple contexts. Zhao and Suh [12] showed that on-chip ROs can be used for remote power side-channel attacks in multi-tenant FPGAs, while Schellenberg et al. [13] demonstrated similar effects through internal power distribution network monitoring. Giechaskiel et al. [8] used RO-based sensors to detect and localise hardware Trojan activity, establishing that ROs can distinguish malicious from benign switching patterns. These studies confirm that on-chip instrumentation is a realistic attack surface, but they do not systematically characterise how observability changes as the coupling path between secret source and instrument is weakened.

Direct-connection RO attacks. Lundberg [18] demonstrated a functional internal side-channel attack on a Xilinx Artix-7 FPGA in which each bit of the AES key directly enabled or disabled a Ring Oscillator. When the key bit was 1, the RO oscillated and drew measurable current; when it was 0, the RO remained static. The resulting power difference was detectable on the FPGA supply lines using an external oscilloscope. This established an important proof of concept, but it relied on a direct key-to-RO electrical path that would be visible to design inspection, netlist analysis, or routing review. The open question is whether observability persists when this direct path is weakened or removed entirely.

Proximity and routing-level coupling. Long-wire leakage research has established that FPGA routing fabrics create measurable parasitic coupling between logically unrelated signals [10, 11]. Adjacent or parallel tracks on the same metal layer exhibit significant sidewall capacitance in deep-submicron processes, enabling both information leakage and covert communication [10]. Subsequent work demonstrated that on-chip sensors can detect malicious activity through these routing-level effects [8]. These studies primarily address activity detection, covert channels, or routing-resource leakage. They do not systematically ask whether a controlled secret-data source produces separable instrument readouts as the coupling mechanism

transitions from an explicit electrical connection to unintended physical proximity.

Countermeasures and design practice. Proposed defences against routing-level leakage include physical separation constraints, guard routing, shielding tracks, and placement policies that reject unauthorised oscillators in shared fabrics [14, 15]. These approaches treat the routing fabric as a shared resource requiring security-aware allocation. However, the effectiveness of separation under different coupling strengths, and the minimum safe distance between sensitive signals and on-chip instruments, remain incompletely quantified in the literature.

1.2.3 Research Gap

Taken together, prior work establishes three points: on-chip instrumentation can leak information when directly wired to a secret source; routing-level proximity can create measurable coupling between logically isolated circuits; and countermeasures exist in principle but lack quantitative coupling-distance guidance. The missing element is a systematic characterisation of the boundary conditions that determine whether an on-chip instrument can reliably distinguish secret bit values as the coupling path is weakened from an explicit electrical connection toward less controlled physical interaction.

It is not sufficient to know that instruments can sense activity or that coupling exists in the routing fabric. For a key-leakage attack, the measurement must resolve the value of a selected secret bit with enough reliability to reconstruct information. The present study addresses this gap by examining observability across a continuous spectrum of coupling conditions—from direct electrical control through a threshold-limited return path to pure physical proximity—using a controlled proof-of-concept platform that isolates the coupling mechanism from other variables.

1.3 Goals and Research Questions

The overarching goal is to characterise how secret-key observability through on-chip reliability instrumentation changes as the physical coupling path between the secret source and the instrument is progressively weakened. To examine this, the study uses an FPGA-based proof-of-concept platform that provides controllable placement, routing, and measurement access. The platform choice enables precise experimental control, but the physical mechanisms studied—direct electrical coupling, input-threshold behaviour, and routing-level cross-talk—are relevant beyond FPGAs to any integrated circuit containing sensitive logic and on-chip instrumentation.

The work is guided by three research questions. First, under a direct electrical connection between a selected key bit and an RO listener, what level of bit-value observability can be achieved using only on-chip digital measurement, and how does this compare with prior external-oscilloscope methods? Second, when the direct path is replaced by a threshold-limited external return path, at what point does key-bit observability fail, and does the failure follow a gradual analog degradation or an abrupt digital-threshold pattern? Third, when no intentional logical connection exists, can physical proximity alone—through routing-level cross-talk, power-network coupling, or substrate coupling—produce measurable key-bit separability, and how do instrument geometry and placement distance affect the result?

These questions are investigated through three completed experiment families: a direct-connection baseline that establishes the high-SNR reference; a GPIO-level threshold-limited path experiment that tests the digital input boundary; and a proximity-based pure cross-talk matrix that measures 24 distance-and-instrument-length conditions.

1.4 Controlled Experiment Scope

The experimental design employs a controlled key-scan methodology in which each key bit is sequentially selected and driven onto a defined coupling path. This approach isolates the physical coupling mechanism for systematic study, separating the question of whether a selected bit remains observable from the broader problem of recovering a full key from arbitrary cryptographic switching activity. The scope is therefore physical observability under controlled conditions, not a universal end-to-end cryptanalytic attack on arbitrary AES implementations.

The FPGA platform provides the controllability needed for this isolation, but the results are interpreted in terms of physical mechanisms that generalise beyond the specific device. Reported findings are based on hardware captures and processed summaries. For the proximity-coupling study, the 24-condition matrix is treated as measured evidence of distance- and instrument-length-dependent separability, while post-route adjacency proof, parasitic extraction, and repeated captures are acknowledged as remaining steps needed for full physical characterisation.

1.5 Contributions

The study makes four main contributions to the understanding of on-chip instrumentation as a side-channel vector. The first is a direct-connection baseline that reproduces and extends prior RO-based key-leakage work us-

ing only on-chip digital measurement, establishing that an external oscilloscope is not required for observing the leakage mechanism in this controlled setting. The second is a GPIO-level threshold-limited path experiment that reveals an abrupt input-threshold-limited failure boundary, showing that attenuating a sensitive signal path does not gradually reduce leakage but rather preserves full observability until the digital input receiver can no longer resolve the signal. The third is a 24-condition proximity-coupling matrix across three placement distances and eight RO lengths, providing the first systematic measurement of how key-bit separability degrades with distance and instrument geometry under pure cross-talk conditions. The fourth is an analysis of the cumulative attack risk posed by weak but repeatable observability, arguing that even partially separable measurements—when combined across multiple observations, multiple instruments, or multiple side channels—can pose a realistic security threat that is not captured by single-measurement supervised accuracy alone.

Chapter 2

Background

2.1 Hardware Trojans and the Instrumentation Threat Model

A Hardware Trojan (HT) is a malicious, intentional modification of an electronic circuit inserted at any stage of its lifecycle [5]. HTs are typically composed of a trigger and a payload: the trigger monitors for a specific condition, and when met, activates the payload to leak information, modify functionality, or cause denial of service [4]. In FPGA design flows, Trojans can be inserted through compromised third-party IP cores, malicious EDA tool plugins, or intercepted bitstream updates [1]. In multi-tenant environments, a malicious tenant may insert or exploit monitoring logic that observes neighbouring tenants' activity through shared routing, power, or thermal resources [8, 10, 12, 13].

The present study adopts a broader threat model that is not limited to deliberately inserted Trojans. Any on-chip instrument whose readout correlates with local electrical activity—whether originally placed for reliability monitoring, process calibration, or debug—can act as a side-channel receiver if an adversary gains access to its output. Ring Oscillators are particularly relevant examples because they have legitimate uses in reliability monitoring, clock generation, and process calibration, and a small RO added to a design may pass scrutiny as a benign structure while covertly leaking information. The experiments therefore treat the RO as a representative on-chip instrument and ask how much key-dependent information it can recover under different coupling conditions.

2.2 AES and the Observability Framework

2.2.1 AES as a Controlled Secret Source

The Advanced Encryption Standard (AES) is a symmetric block cipher adopted by NIST in 2001 [19]. AES-256 uses a 256-bit key and 14 rounds

of SubBytes, ShiftRows, MixColumns, and AddRoundKey operations. In conventional side-channel analysis, the Hamming weight and Hamming distance of intermediate registers correlate with secret-dependent values, modulating power consumption in a way that can be statistically recovered from many traces [16, 17].

The present study does not perform Correlation Power Analysis on AES round intermediates. Instead, the AES-256 key register is treated as a controlled sensitive data source: each key bit is sequentially selected and driven onto a defined coupling path, and the instrument readout is measured as a function of that single bit value. This *controlled key-scan* methodology isolates the physical coupling mechanism from the complexity of full cryptographic switching, making it possible to study how observability degrades as the coupling path is weakened. The focus is therefore on *key-bit observability*—the ability of an on-chip instrument to resolve whether a selected secret bit is 0 or 1—rather than on end-to-end key recovery from arbitrary AES operation.

2.2.2 Separability and Supervised Thresholding

Two key concepts recur throughout the analysis. *Class separation* (ΔC) is the difference between the mean instrument readouts associated with the two bit values: $\Delta C = \mu_1 - \mu_0$, where μ_1 is the mean readout when the actual bit is 1 and μ_0 is the mean readout when the actual bit is 0. Large positive ΔC indicates that the two bit classes produce distinguishable instrument responses.

Supervised separability quantifies how well the two labelled classes can be separated by a threshold after a calibration pass. For conditions where both bit classes produce non-zero instrument readouts, a fixed low threshold is inappropriate. Instead, a midpoint threshold $T = (\mu_0 + \mu_1)/2$ is computed from labelled calibration data, and the fraction of bits correctly classified by this threshold is reported as supervised separability. It is important to distinguish this metric from *unsupervised key recovery*: supervised separability measures whether the physical coupling produces separable distributions when the class labels are known, not whether an attacker without labels could discover the threshold independently. This distinction is central to the conservative interpretation of the proximity-coupling results.

2.3 Ring Oscillators as Representative On-Chip Sensors

2.3.1 Instrument Physics and Voltage Sensitivity

A Ring Oscillator formed from N odd-numbered inverters oscillates at

$$f_{osc} = \frac{1}{2 \cdot N \cdot t_{delay}} \tag{2.1}$$

where the per-stage delay t_{delay} depends on the supply voltage through the Alpha-Power Law

$$t_{delay} \propto \frac{C_L V_{dd}}{(V_{dd} - V_{th})^{\gamma}} \tag{2.2}$$

This voltage dependence is the physical basis for both the RO's utility as a reliability monitor and its vulnerability as a side-channel receiver. Local voltage fluctuations caused by neighbouring switching activity produce measurable frequency shifts that can be accumulated into a count over a fixed measurement window. The same principle applies, with appropriate transduction mechanisms, to other on-chip instruments: temperature sensors exploit the temperature dependence of threshold voltage and carrier mobility; delay monitors measure path slack under varying supply and temperature conditions; and voltage monitors directly sample the local supply rail.

The experiments use the RO as a *representative* instrument because it is simple to implement, requires no analogue reference, and produces a digital count that is easy to capture and analyse. The findings are not limited to ROs: any instrument whose readout correlates with local electrical perturbations can in principle leak secret-dependent information under similar coupling conditions.

2.3.2 FPGA Implementation and Measurement Method

On the Xilinx Artix-7 proof-of-concept platform, an inverter is implemented using a 6-input LUT configured as a NOT gate. The RO is enabled by replacing the first inverter with a NAND (enable='1' allows oscillation; '0' holds the output static). Odd N is required to prevent latch-up. Routing delays between LUTs contribute to t_{delay} and depend on placement.

The RO output is a high-speed asynchronous signal sampled by the 100 MHz system clock through two flip-flop synchroniser stages. The counter increments on observed rising transitions and is frozen at the end of a 2^{20}-cycle measurement window (approximately 10.49 ms). The resulting

count is a *relative sampled observable*, not a calibrated oscillator frequency; conclusions are based on within-scenario class separation. Asynchronous sampling introduces edge ambiguity—a transition near the window boundary may or may not be counted depending on the phase relationship between the RO and system clock—which is the most likely cause of the single count=1 false positive observed in the baseline experiment. This ambiguity is negligible for large class separations but limits interpretation of the raw count as an absolute frequency.

Synthesis tools attempt to eliminate combinational loops, so attributes `(* KEEP = "TRUE" *)` and `(* DONT_TOUCH = "TRUE" *)` are required on every inverter instance and the feedback net [20]. Jitter, temperature drift, and supply variation introduce low-frequency noise, motivating relative comparisons within a single capture session.

2.4 Routing-Level Coupling and Its Generalisation Beyond FPGAs

2.4.1 FPGA Routing Fabric and Proximity Effects

The Artix-7 routing fabric is built from hierarchical wire resources connected through programmable switch matrices. Adjacent or parallel routing resources on the same metal layer have significant sidewall capacitance due to high aspect ratios in 28 nm technology, creating measurable coupling between logically unrelated circuits [8, 10, 11]. A switching signal on an aggressor track injects displacement current into an adjacent victim track through mutual capacitance:

$$I_{crosstalk} = C_{mutual} \frac{dV_{aggressor}}{dt} \tag{2.3}$$

Equation 2.3 captures the capacitive component, but the measured effects in any real chip may combine capacitive coupling through the routing fabric, resistive coupling through the shared power distribution network, and substrate coupling through the silicon bulk. The experimental design does not isolate these individually.

2.4.2 Relevance to ASIC and SoC Platforms

While the proof-of-concept experiments are implemented on an FPGA for controllability, the physical coupling mechanisms are not FPGA-specific. In custom ASICs and SoCs, adjacent metal layers and parallel wire segments create the same sidewall capacitance effects. The power distribution

network is shared across all blocks, enabling resistive coupling from cryptographic switching to nearby sensors. The silicon substrate provides a common ground plane that can mediate substrate coupling between physically close but logically isolated circuits. The key difference is that FPGA fabrics offer regular, predictable routing structures that make coupling easier to study experimentally, whereas ASIC coupling depends on custom layout decisions that vary from design to design. The qualitative conclusions—that proximity matters, that coupling strength decays with distance, and that instrument geometry affects sensitivity—are expected to transfer to any integrated circuit containing sensitive logic and on-chip instrumentation.

2.4.3 Routing-Level Defences

Routing-level countermeasures aim to prevent untrusted measurement circuits from sharing sensitive physical resources. Proposed approaches include separation constraints, routing audits, shielding or guard routing, and placement policies that reject unauthorised oscillators in shared fabrics [14, 15]. These defences are relevant because the proximity-only measurements remain dependent on implementation-specific placement and routing. The experiments use them as design guidance and interpret the measured distance trends as motivating stricter separation policies, not as validated countermeasures.

2.5 Positioning Relative to Prior Work

The present work sits at the intersection of three research lines: direct-connection RO attacks [18], routing-level proximity leakage [8, 10, 11], and on-chip instrumentation as a side-channel vector [12, 13]. It advances beyond these lines by systematically characterising the boundary conditions under which key-bit observability persists as the coupling path is weakened from a direct electrical connection through a threshold-limited path to pure physical proximity. The contribution is not a new claim that coupling exists—that is established—but a quantitative measurement of how key-specific separability degrades across the coupling continuum, together with an analysis of the cumulative attack risk posed by weak but repeatable observability.

Chapter 3

Method and Experimental Setup

The experimental design translates the three research questions into a controlled sequence of coupling conditions. A direct-connection baseline establishes the observable upper bound when the selected key bit fully controls the RO. A GPIO-level threshold path then weakens that path while keeping the signal electrically connected through the FPGA I/O boundary. Finally, a proximity-only setup removes the intentional logical connection and tests whether placement alone can produce measurable class separation. All three experiment families are reported as hardware measurements; the proximity-only results are interpreted with additional physical-design caveats.

3.1 Working Hypotheses

The hypotheses below connect the background mechanisms in Chapter 2 to the measurable outcomes used in the experiments. They are working hypotheses for interpreting the controlled measurements, not assumptions that the results are expected to confirm.

3.1.1 H1: GPIO-Level Threshold-Limited Path

When the key-dependent control path is routed out through a GPIO-level series/pull-down network and then back into the FPGA, streamed-bit agreement is expected to remain reliable while the returned input voltage exceeds the FPGA input-buffer logic-high threshold. It is expected to fail abruptly once the voltage falls below that threshold. The transition should therefore be digital-threshold-limited rather than a smooth analog attenuation, because the FPGA input buffer acts as a comparator with a sharp switching boundary.

3.1.2 H2: Proximity-Dependent Observability

Under controlled placement, a key-dependent aggressor routing path placed near an RO-based receiver path may produce measurable RO-count separation even when no explicit RTL-level or netlist-level logical connection exists between them. For H2 to be supported, the measured separation must exceed the fixed-key baselines and the bit-class ranges should separate or narrow toward separation under stronger coupling conditions.

3.1.3 H3: RO Length Amplification and Saturation

In weak-coupling conditions, RO length N is expected to affect observability through a trade-off. Short ROs (small N) may not accumulate sufficient perturbation to produce reliable separation. Intermediate ROs (moderate N) may improve sensitivity by accumulating delay variations across more stages. Very long ROs (large N) may saturate or decline because fewer sampled transitions are observed within the fixed measurement window, while accumulated jitter or routing detours may add noise. The pure cross-talk matrix therefore uses N as an explicit experimental factor rather than treating RO length as a fixed implementation detail.

3.2 Experimental Platform

All experiments were implemented on a Digilent Nexys 4 board with a Xilinx Artix-7 FPGA, device `xc7a100tcsg324-1` (28 nm process, 101,440 logic cells) [7, 23]. The FPGA was chosen as a proof-of-concept platform because its reconfigurable fabric provides precise experimental control over placement, routing, and instrumentation access. The physical mechanisms studied—direct electrical coupling, input-threshold behaviour, and routing-level cross-talk—are not FPGA-specific; they occur in any integrated circuit containing sensitive logic and on-chip instruments. The FPGA platform makes these mechanisms easier to isolate and measure systematically.

The implementation flow used Xilinx Vivado 2024.2 with experiment-specific Tcl scripts and XDC constraint files. The main clock source was the on-board 100 MHz oscillator.

3.2.1 Measurement Philosophy

Measurements used an Integrated Logic Analyzer (ILA) IP core to export compact result rows after the on-chip scan had accumulated the instrument counts. In the final slice-recovery captures, the ILA data depth was 4,096 samples and the trigger was `stream_frame_start`. This buffer depth is

not the duration of the key scan. For each selected bit, an internal 32-bit counter accumulates sampled RO rising transitions over a fixed 2^{20}-cycle window of the 100 MHz system clock (approximately 10.49 ms). After all 256 bits have been measured, the controller streams stored result rows containing the bit index, known bit value, online recovered bit, and RO count. The ILA therefore records the compact stream interface rather than the raw continuous RO waveform over the approximately 2.7 s scan. The use of on-chip digital capture rather than an external oscilloscope demonstrates that the leakage mechanism can be observed using only instrumentation already present or insertable on the die.

3.3 Design Overview

3.3.1 Module Architecture

The design comprises five functional blocks. An AES-256 key register stores the fixed key, and a four-state Moore FSM selects one bit at a time. The selected bit then passes through the scenario-specific coupling path before it reaches an enable-controlled RO listener with a 32-bit counter. Finally, an ILA stream interface exports the compact labelled result rows.

3.3.2 Key-Scan FSM

The FSM has four states. IDLE loads the key and waits for the trigger. SELECT drives the selected bit onto the coupling path and allows a short settling interval. MEASURE enables the RO counter for 2^{20} system-clock cycles and stores the final count in an on-chip result array. After the last bit has been measured, STREAM emits one compact row per bit. The full 256-bit scan completes in approximately 2.7 s, excluding small control overheads, whereas the streaming phase occupies only a small ILA capture window. Each measurement result comprises bit index (8-bit), actual bit (1-bit), online recovered bit (1-bit), and RO count (32-bit).

3.3.3 ILA Capture and Reduction

The ILA is triggered by the first `stream_frame_start` pulse, which marks the first streamed result row after the 256-bit scan has completed. No hardware capture qualification was used to filter the ILA buffer to valid rows only. Instead, `stream_valid` is recorded as a probe, and offline Python scripts retain the first 256 valid streamed rows for the summary. This explains why the ILA buffer depth and the scan duration are different quantities: the counter performs the long per-bit accumulation, while the ILA exports the compact post-measurement stream.

3.3.4 Thresholding Strategy

Two thresholding strategies are distinguished. For the direct-connection and GPIO-level threshold-limited scenarios, bit=0 counts are zero and bit=1 counts exceed 30,000 in the above-threshold measurements; a fixed threshold of 10 counts separates the classes in these measurements. For weak-coupling cross-talk measurements, both classes may produce large non-zero counts, so a fixed low threshold is inappropriate. A midpoint threshold $T = (\mu_0 + \mu_1)/2$ quantifies labelled class separability after a calibration pass, but it is interpreted as a supervised observability metric rather than as an unsupervised attack procedure.

3.4 Evaluation Metrics

The analysis computes the same set of metrics for each 256-row summary:

- μ_1: mean RO count for bits with actual value 1; μ_0: mean RO count for bits with actual value 0.
- $\Delta C = \mu_1 - \mu_0$: raw-count class separation.
- **Accuracy**: fraction correctly classified, $N_{correct}/256$.
- **Majority-class baseline**: $145/256 = 56.64\%$ (always predict 1 for this fixed key).
- **Random expectation**: 50%.
- **All-zeros constant classifier**: $111/256 = 43.36\%$.
- **Overlap**: whether the per-class min–max ranges intersect.

Because the fixed key is imbalanced (145 ones, 111 zeros), raw accuracy alone is misleading. An accuracy of 43.36% indicates constant-output behaviour (always predicting 0), not performance below random. The appropriate reference is therefore the majority-class baseline of 56.64%.

3.5 Experimental Controls and Validity Boundaries

The experiments keep the FPGA board, 100 MHz system clock, AES-256 key, key-scan order, RO-count window, and ILA export method fixed across the main scenarios. These controls make the coupling path the primary experimental variable. The direct and GPIO-level experiments therefore form a controlled progression from a full digital connection to a threshold-limited external return path.

This control also limits the generality of the results. The fixed key enables direct comparison across scenarios, but its class imbalance requires

the majority-class and all-zeros baselines defined in Section 3.4. The key-scan method isolates one selected bit at a time, which makes the physical coupling mechanism easier to study, but it is not the same as observing a naturally switching AES datapath. The results should therefore be interpreted as controlled observability measurements, not as a complete end-to-end cryptanalytic attack.

For the proximity-only study, the strongest validity limitation is physical-design uncertainty. Placement constraints define approximate CLB locations, but they do not by themselves prove that the routed aggressor and receiver wires are adjacent for a known length. This is why the pure crosstalk matrix supports a distance- and RO-length-dependent observability conclusion, but not a fully characterised physical coupling model.

3.6 Direct-Connection Baseline Setup

The selected key bit directly drives the RO enable, as shown in Figure 3.1. If the bit is 1, the RO oscillates and the counter accumulates edges; if 0, the RO is disabled and the counter output is zero. This scenario provides the high-SNR reference for all subsequent weakened-coupling comparisons.

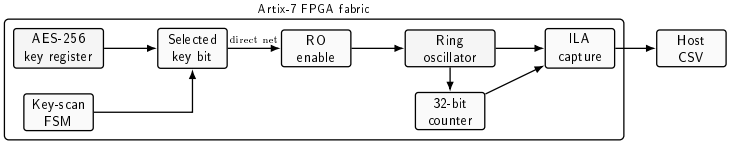

Figure 3.1: Direct-connection baseline setup. The key-scan FSM selects one AES-256 key bit at a time and drives it directly onto the RO enable net. The RO count is captured internally by the ILA and exported for offline analysis.

3.7 GPIO-Level Threshold-Limited Path Setup

The selected bit drives FPGA output pin JA3. As summarised in Figure 3.2, the external path is: JA3 $\rightarrow$ R_{series} $\rightarrow$ JA4 $\rightarrow$ R_{pull} to ground $\rightarrow$ FPGA input buffer $\rightarrow$ synchroniser $\rightarrow$ RO enable. This setup is a GPIO-level threshold experiment. It does not model distributed on-chip interconnect resistance; instead, it tests whether the returned digital signal still crosses the FPGA input-buffer threshold. The ideal sampled voltage follows:

$$V_{in} = V_{drv} \frac{R_{pull}}{R_{series} + R_{pull}} \tag{3.1}$$

The main completed sweep fixed R_{pull} = 400 Ω and varied R_{series} from 0 Ω to 1 kΩ plus open circuit. Additional partial controls with other pull-

down values were performed, but they are not treated as a complete four-by-series matrix. The experiment models how a progressively weakened external digital return path affects the RO enable at the FPGA's input boundary.

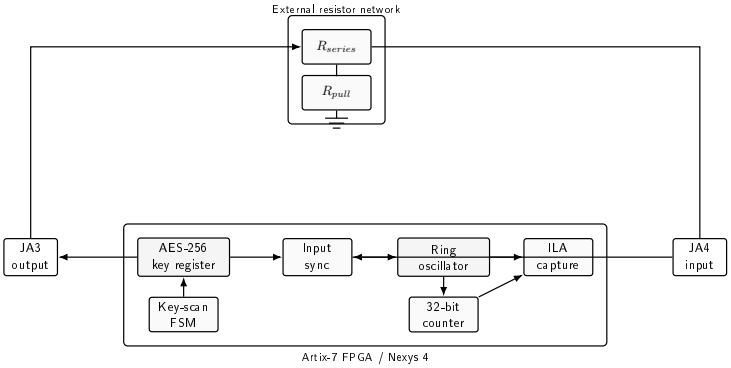

Figure 3.2: GPIO-level threshold-limited path setup. The selected key bit leaves the FPGA through JA3, passes through the external series resistor and pull-down network, and re-enters through JA4 before controlling the RO listener. The measured transition is interpreted at the input-buffer threshold, not as distributed on-chip interconnect resistance.

3.8 Proximity-Based Pure Cross-Talk Setup

3.8.1 Concept

The proximity-only concept is shown in Figure 3.3. No intentional logical connection should exist between the aggressor and receiver nets. The selected key bit drives a dedicated routing track (the aggressor), while the RO includes an explicit receiver stub—a named wire segment within the oscillation loop—intended to be placed near the aggressor. Any distinguishable effect would have to be mediated by physical coupling rather than by a designed digital path.

3.8.2 Synthesis Constraints

Verilog attributes and XDC constraints prevent optimisation of the receiver stub and fix approximate placement. The implementation used `KEEP` and `DONT_TOUCH` attributes for preservation, `LOC/BEL` constraints for the relevant aggressor and receiver LUTs, and `PRESERVE` on the named receiver-stub net. Appendix A records the representative constraint syntax. These constraints preserve logic and place cells, but they do not guarantee exact post-route wire adjacency. The distance groups (near, medium, far)

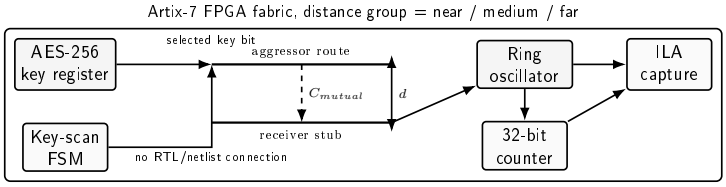

Figure 3.3: Proximity-based pure cross-talk setup. The selected key bit drives an aggressor route, while the RO listener contains a separate receiver stub. The two routes are placed with controlled near, medium, or far spacing, without an intentional RTL-level or netlist-level connection.

are therefore implementation-controlled approximations based on CLB locations, not verified post-route track spacing.

3.8.3 Experimental Matrix

The measured matrix used three distance groups—near (adjacent CLBs), medium (2–3 CLBs apart), and far (opposite device ends, approximately 30 CLBs apart)—and eight RO lengths: N = 5, 15, 25, 35, 45, 55, 65, and 75. This gives 24 pure cross-talk conditions. Each condition is reduced to a 256-row key-bit summary using the same bit-scan pipeline and the supervised midpoint threshold described in Section 3.4.

3.8.4 Mechanism Attribution

All circuits share the same FPGA substrate, power distribution network, and clocking infrastructure. The measured effect is attributed to proximity-dependent coupling with caution; the design cannot isolate capacitive coupling as the sole physical contributor.

3.9 Post-Implementation Verification

The following checks were performed during the proximity-only study: RTL inspection, Vivado schematic inspection for unexpected logical connectivity, Design Rule Check (DRC), and placement inspection in the device view. These checks are informative but not sufficient for a fully characterised cross-talk mechanism. Post-route track adjacency, parallel run length, parasitic extraction, and repeated captures were not completed; exact coupling capacitance values are unavailable.

The metrics in Section 3.4 are then applied to the available 256-row summaries. For the proximity-only study, the overlap indicator is especially important because an accuracy near the majority-class baseline can occur even when the classifier carries little key-dependent information.

Chapter 4

Experimental Results

The results are organised according to the progression of coupling conditions defined in Chapter 3. The direct-connection baseline establishes the reference behaviour when the key bit fully controls the RO enable. The GPIO-level threshold-limited path then weakens the signal through an external return network. The proximity-coupling study removes the intentional logical connection and evaluates the measured distance/RO-length matrix. Because the fixed AES-256 key contains 145 ones and 111 zeros, all accuracy values are interpreted relative to both the all-ones majority-class baseline (56.64%) and the all-zeros constant classifier (43.36%).

4.1 Direct-Connection Baseline

The direct-connection baseline validates the experimental infrastructure and provides the reference signal level for later comparisons. Figure 4.1 and Table 4.1 show per-bit RO counts: bit=0 samples produce zero or near-zero counts because the RO is disabled, while bit=1 samples produce 30,000–42,000 counts.

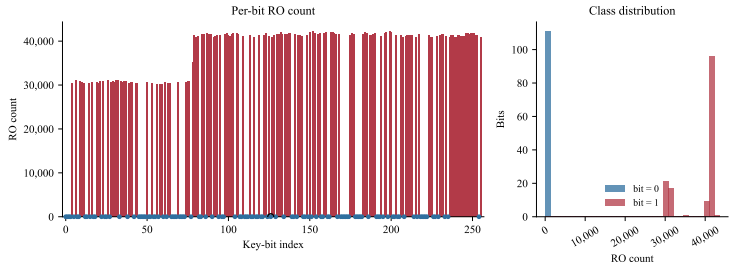

Figure 4.1: Measured direct-connection baseline. Left: per-bit RO count, with bit=1 samples shown in red and bit=0 samples in blue. Right: class distributions showing the large count separation between the two bit classes.

The bit=1 count variation (30,000–42,000) arises from differences in routing delays and local process variation as the scan index advances through

Table 4.1: Baseline performance. The direct connection matches 255 of 256 streamed bit decisions, with a large count separation between enabled and disabled RO states.

Metric	Value
Matched streamed decisions	255 / 256
Accuracy	99.61%
μ_1	38,592
μ_0	0.009
ΔC	38,592

two floorplan regions, not from the key-bit value itself. Bit 126 (actual=0, count=1) is a false positive consistent with a measurement-window edge effect; a threshold margin of 10 would eliminate it.

Figure 4.2 shows the class-level distribution of RO counts for the direct-connection baseline. The two populations are completely separated: bit=0 samples cluster near zero, while bit=1 samples form a distinct distribution centred at approximately 38,600 counts. This large separation is the physical basis for the near-perfect accuracy observed.

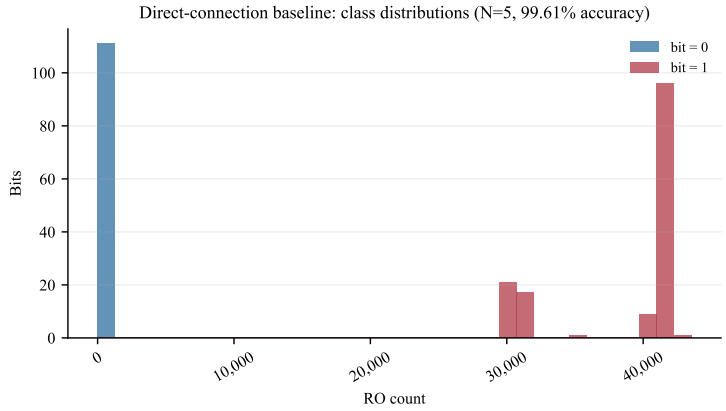

Figure 4.2: Class distributions for the direct-connection baseline. The bit=0 and bit=1 populations are completely non-overlapping, enabling reliable threshold-based classification.

4.2 Direct-Connection RO Sweep

The direct-connection RO sweep (N = 5, 7, 9, 11, 13) is an implementation and capture diagnostic, not the primary weak-coupling sensitivity test. Table 4.2 and Figure 4.3 report two different quantities. The mean count gap ΔC measures count-level class separation between enabled and

disabled RO states. The streamed accuracy reports agreement between the online recovered-bit field exported by the design and the known bit label. In this full-swing direct connection, large ΔC values confirm that the bit classes are separable, but the streamed accuracy can still be affected by build-specific implementation, stream timing, threshold logic, or capture/reduction details.

Each RO length in this diagnostic sweep was implemented as a separate FPGA build. Therefore, the N=5 raw count in the sweep is not expected to exactly match the single-bitstream baseline in Table 4.1, and the sweep should not be used as a reliable ranking of RO lengths.

Table 4.2: Direct-connection RO sweep. The values are diagnostic because each RO length was implemented as a separate FPGA build.

N	Accuracy	μ_1	μ_0	ΔC
5	99.61%	39,707	∼0	39,706
7	87.50%	510,507	∼0	510,506
9	94.53%	187,331	∼0	187,331
11	90.23%	494,262	∼0	494,262
13	95.31%	207,525	∼0	207,525

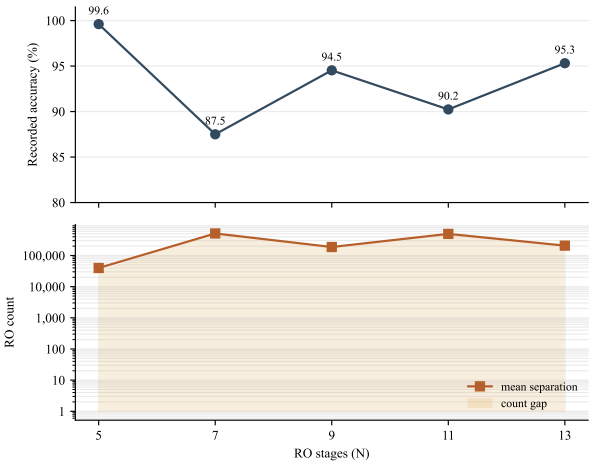

Figure 4.3: Direct-connection RO sweep. The upper panel reports streamed recovered-bit agreement for each separately implemented RO length, while the lower panel reports count-level class separation. The two diagnostics should not be collapsed into a single RO-length ranking.

The non-monotonic ΔC values are implementation-specific. For example, N=7 gives 510,506 whereas N=9 gives 187,331. The streamed recovered-bit agreement also varies across builds, from 87.50% to 99.61%, despite the large count-level separation in every case. This mismatch does not

imply that N=7 is a better attack configuration than N=5 or N=9. It indicates that the sweep is useful for diagnosing implementation and export behaviour under full-swing conditions. The narrower conclusion is that a short RO (N=5) suffices for the final strong-signal direct-connection baseline, while RO-length optimisation for weak coupling is evaluated separately in the proximity-only matrix.

4.3 GPIO-Level Threshold-Limited Path Results

4.3.1 Fixed Pull-Down 400 Ω

The completed series-resistance sweep used $R_{pull} = 400\,\Omega$. Table 4.3 reports the measured streamed-agreement statistics, and Figure 4.4 visualises the abrupt transition.

Table 4.3: GPIO-level threshold-limited path, $R_{pull} = 400\ \Omega$.

R_{series}	Accuracy	μ_1	μ_0	ΔC	V_{in} (est.)
0 Ω	100.00%	112,792	0	112,792	3.30 V
100 Ω	100.00%	113,221	0	113,221	2.64 V
200 Ω	100.00%	112,321	0	112,321	2.20 V
300 Ω	100.00%	112,916	0	112,916	1.89 V
400 Ω	43.36%	0	0	0	1.65 V
500 Ω	43.36%	0	0	0	1.47 V
600 Ω	43.36%	0	0	0	1.32 V
700 Ω	43.36%	0	0	0	1.20 V
800 Ω	43.36%	0	0	0	1.10 V
900 Ω	43.36%	0	0	0	1.01 V
1,000 Ω	43.36%	0	0	0	0.94 V
Open	43.36%	0	0	0	floating

The transition is abrupt. For $R_{series} = 0$–$300\ \Omega$, accuracy is 100% with $\Delta C \approx 113,000$. At 400 Ω, accuracy drops to 43.36%. This value corresponds to an all-zeros constant classifier, *not* random guessing. Since an all-ones majority-class classifier would achieve 56.64%, the drop to 43.36% indicates that the digital return path no longer carries distinguishable key-dependent information.

Figures 4.5 and 4.6 show the class distributions for the above-threshold and below-threshold regimes, respectively. At R_{series}=100 Ω, the bit=0 and bit=1 populations are completely separated, with bit=0 counts at zero and bit=1 counts clustered near 113,000. At R_{series}=400 Ω, both populations collapse to zero counts, confirming that the RO enable is held static and no oscillation occurs regardless of the key-bit value.

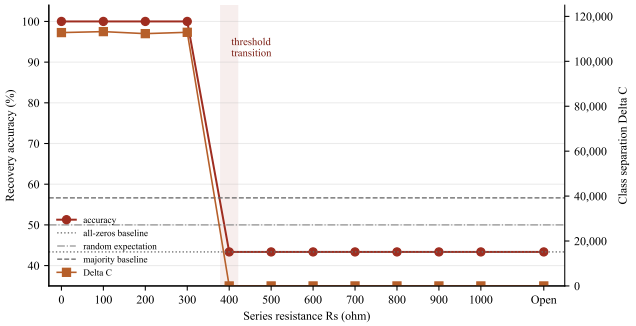

Figure 4.4: GPIO-level threshold-limited path (R_{pull} = 400 Ω). Accuracy and ΔC are plotted against R_{series} on separate axes. The transition is interpreted as FPGA input-buffer threshold behaviour, not as distributed on-chip interconnect resistance. The reference lines mark the all-zeros constant output (43.36%), random expectation (50%), and the all-ones majority-class baseline (56.64%).

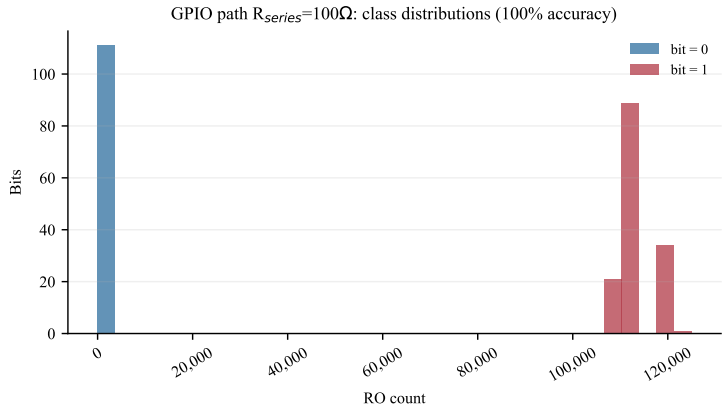

Figure 4.5: Class distributions for the GPIO-level path with R_{series}=100 Ω (above threshold). The bit=1 population is centred near 113,000 counts, while bit=0 remains at zero, enabling perfect separation.

4.3.2 Model Agreement and Deviation

The voltage-divider model (Eq. 3.1) is consistent with the threshold transition for the completed 400 Ω sweep. Figure 4.7 visualises the ideal divider trend. Failure begins when the ideal V_{in} falls to approximately 1.65 V, close to the expected Artix-7 LVCMOS33 switching region. Partial controls with other pull-down values are summarised in Table 4.4 and indicate that the model alone is insufficient when the output driver is heavily loaded.

The partial 100 Ω and 200 Ω pull-down controls are best interpreted as output-drive-limited cases. A 100 Ω pull-down can demand approximately 33 mA from the FPGA output buffer, exceeding the configured drive-strength range normally used for LVCMOS33 outputs in 7-series FPGA

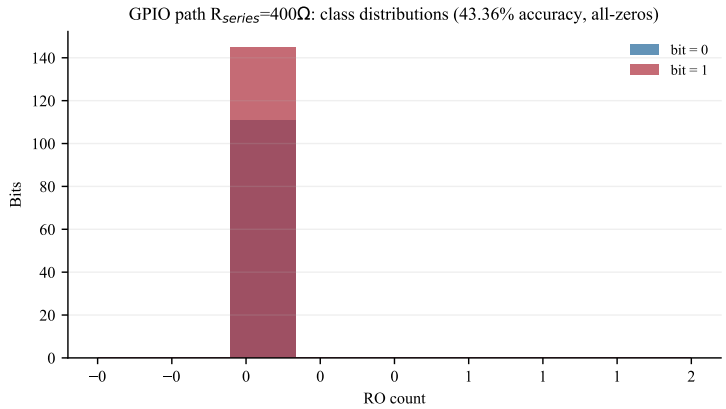

Figure 4.6: Class distributions for the GPIO-level path with R_{series}=400 Ω (below threshold). Both bit classes collapse to zero counts, confirming that the input buffer no longer recognises a logic-high signal.

Table 4.4: Measured and partial-control accuracy matrix across pull-down and series resistance. Blank entries were not completed in the available measurements. The matrix is a GPIO-level threshold-path study, not an on-chip interconnect-resistance model.

R_{pull}	Accuracy at R_{series} =						
	0 Ω	100 Ω	200 Ω	300 Ω	400 Ω	1 kΩ	Open
100 Ω	43.4%	—	—	—	—	—	—
200 Ω	100%	43.4%	—	—	—	—	—
400 Ω	100%	100%	100%	100%	43.4%	43.4%	43.4%
1 kΩ	100%	100%	100%	100%	100%	—	43.4%

I/O banks [24]. These controls therefore should not be merged with the 400 Ω series sweep as if they were a complete resistance matrix. Direct analog voltage measurements at JA4 were not performed, so the voltage values in Table 4.3 are model-based estimates. The threshold interpretation is derived from the digital recovery accuracy transition.

4.3.3 Metrics within the Above-Threshold Regime

Within the completed 400 Ω pull-down sweep, the above-threshold regime (R_{series} = 0–300 Ω) keeps ΔC approximately constant at 112,000–113,000. Bit=0 counts are zero because the RO is fully disabled.

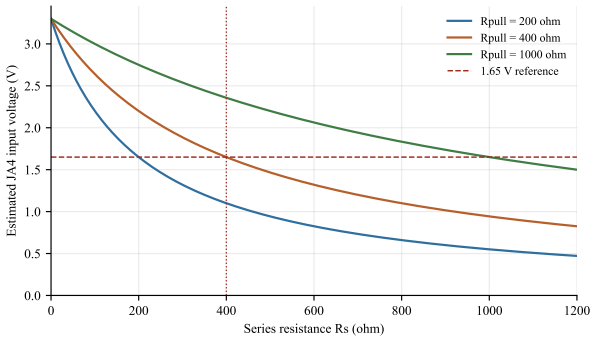

Figure 4.7: Voltage-divider model. The dashed line marks 1.65 V. The 100 Ω and 200 Ω pull-downs are limited by output-drive capability, not the divider alone.

4.4 Pure Cross-Talk Distance/RO Matrix

4.4.1 Measurement Coverage

The proximity-only experiment was designed to test whether physical placement alone can create measurable key-dependent RO-count separation. Unlike the direct and GPIO-level experiments, there is no intentional logical connection between the selected key bit and the RO listener. The measured matrix covers three placement groups (near, medium, and far) and eight RO lengths (N=5–75), giving 24 conditions in total. Each condition contains a 256-bit key scan.

The reported cross-talk accuracy uses the supervised midpoint threshold described in Chapter 3. It should therefore be read as calibrated class separability, not as an unsupervised attack that discovers its own threshold from an unknown key. The overlap column in Table 4.5 records whether the min–max ranges of the bit=0 and bit=1 count distributions intersect. Figures 4.8 and 4.9 separate supervised separability from count-level class separation so that accuracy-like and count-level diagnostics are not conflated.

4.4.2 Distance and RO-Length Trends

The near-distance result changes the interpretation of the pure cross-talk experiment. The near/N=5 condition alone is weak (54.30%, $\Delta C = 149.3$) and would not support a key-recovery claim. However, it is not representative of the full matrix. Supervised separability rises to 87.11% at N=25, reaches 99.61% with non-overlapping class ranges at N=35, and reaches 100% for N=45, 55, 65, and 75. The count separation peaks at N=55 ($\Delta C = 6{,}636.3$) and then declines slightly, which is consistent with the amplification-and-saturation behaviour hypothesised in Section 3.1.3.

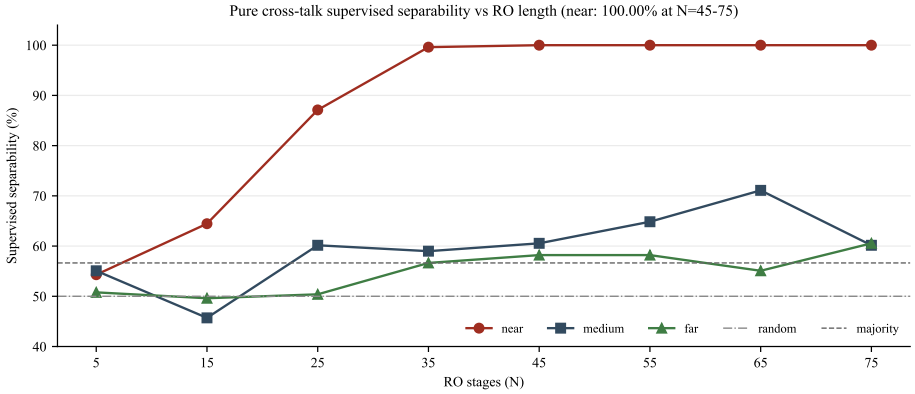

Figure 4.8: Pure cross-talk supervised separability across distance and RO length. Near placement reaches 100.00% for N=45–75, whereas medium and far placement remain overlapping and substantially weaker.

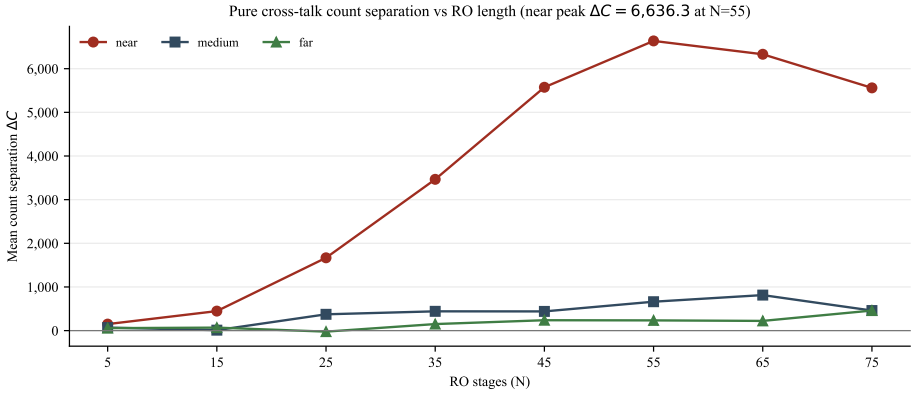

Figure 4.9: Pure cross-talk count-level separation across distance and RO length. The near-placement count gap peaks at $\Delta C = 6{,}636.3$ for N=55 and then declines slightly, while medium and far placement retain much smaller gaps.

The medium-distance group shows partial but unreliable distinguishability. Its best condition is N=65 with 71.09% supervised separability and $\Delta C = 813.9$, but the class ranges still overlap in every measured medium-distance condition. The far-distance group remains close to the random or majority-class baselines. Its best condition is N=75 with 60.55% supervised separability and $\Delta C = 459.5$, again with overlapping class ranges. These two groups do not support unsupervised or reliable full-key recovery in the present setup.

Figure 4.11 shows representative per-condition distributions for the four selected conditions. The near/N=5 case exhibits substantial overlap, while near/N=55 shows clean separation. The medium/N=65 and far/N=75 conditions retain significant overlap, confirming that distance and instrument

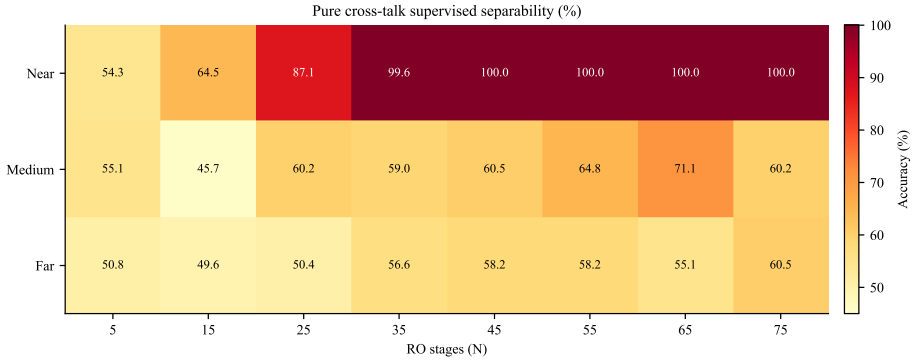

Figure 4.10: Supervised-separability heatmap for the 24-condition pure cross-talk matrix. The strongest calibrated separability is concentrated in the near-placement, larger-RO corner of the matrix.

geometry strongly affect observability under pure cross-talk.

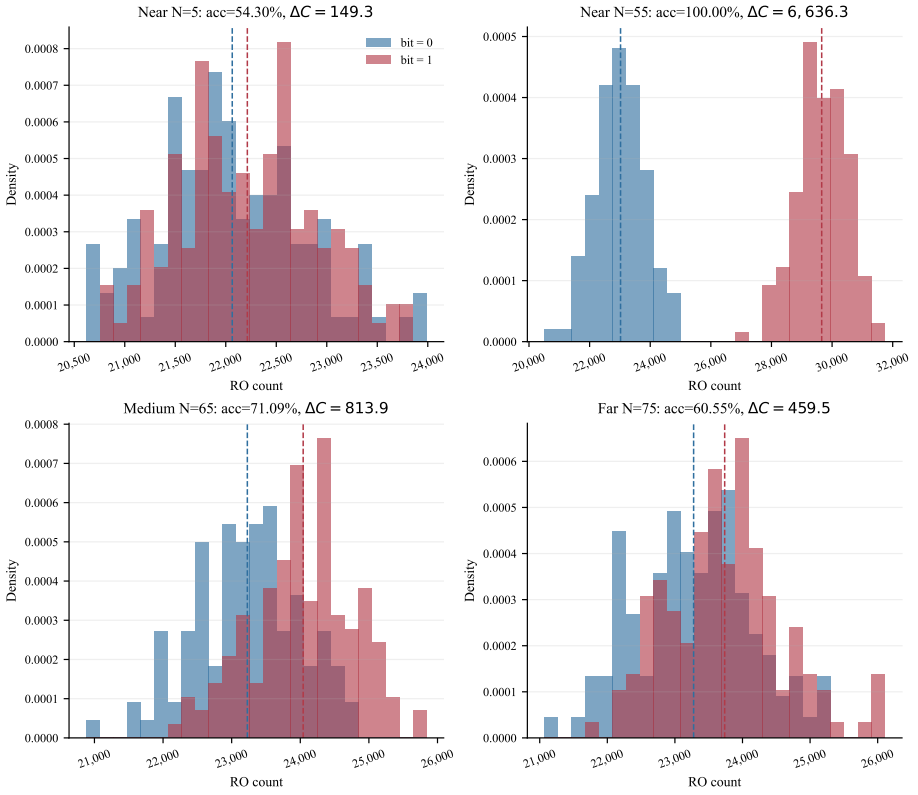

Figure 4.11: Representative count distributions for weak, strong, and distant pure cross-talk conditions. The titles report the supervised accuracy and ΔC for each condition. Near/N=55 separates the bit classes cleanly under calibration, while medium/N=65 and far/N=75 retain substantial overlap.

Table 4.5: Measured pure cross-talk matrix. Accuracy is calibrated supervised separability using the per-condition midpoint threshold. "Overlap" indicates whether the bit-class count ranges intersect.

Distance	RO stages N	Supervised acc.	ΔC	Overlap
near	5	54.30%	149.3	Yes
near	15	64.45%	447.1	Yes
near	25	87.11%	1,668.1	Yes
near	35	99.61%	3,465.5	No
near	45	100.00%	5,574.5	No
near	55	100.00%	6,636.3	No
near	65	100.00%	6,330.4	No
near	75	100.00%	5,560.7	No
medium	5	55.08%	69.7	Yes
medium	15	45.70%	10.6	Yes
medium	25	60.16%	373.5	Yes
medium	35	58.98%	441.8	Yes
medium	45	60.55%	440.0	Yes
medium	55	64.84%	662.2	Yes
medium	65	71.09%	813.9	Yes
medium	75	60.16%	459.9	Yes
far	5	50.78%	56.6	Yes
far	15	49.61%	68.7	Yes
far	25	50.39%	-23.2	Yes
far	35	56.64%	150.0	Yes
far	45	58.20%	238.1	Yes
far	55	58.20%	234.3	Yes
far	65	55.08%	222.5	Yes
far	75	60.55%	459.5	Yes

4.5 Cross-Scenario Comparison

Table 4.6 summarises representative completed scenarios. The comparison separates direct electrical observability, GPIO-level threshold behaviour, and the distance-dependent pure cross-talk separability result. A visual version of the same comparison is provided in Appendix A.6.

Table 4.6: Cross-scenario comparison of the completed measurements.

Scenario	Reported acc.	ΔC	Conditions	Interpretation
Direct connection	99.61%	38,592	N=5	Direct electrical control
GPIO threshold path	100%	113,221	$R_{pull} = 400\ \Omega$, $R_{series} = 100\ \Omega$	Above input threshold
Pure cross-talk, near	100%	6,636	Near, N=55	Strong calibrated separability
Pure cross-talk, medium	71.09%	814	Medium, N=65	Partial calibrated separability
Pure cross-talk, far	60.55%	459	Far, N=75	Weak and overlapping

Chapter 5

Discussion

The measurements support a more nuanced picture than a simple connected-versus-unconnected boundary. Direct and GPIO-level electrical paths produce strong observability, while the pure cross-talk matrix shows that calibrated key-bit separability is possible under close placement with a sufficiently long instrument. The discussion below interprets these results in depth, examines the cumulative attack risk that arises when weak observability is exploited repeatedly, relates the findings to prior work and practical threat models, and derives design recommendations.

5.1 Interpretation of the Direct-Connection Result

The direct-connection baseline validates the experimental infrastructure and establishes the high-SNR reference for all weakened-coupling comparisons. Matching 255 of 256 streamed bit decisions with a 5-stage RO confirms that a directly wired instrument can recover key information with near-perfect reliability using only on-chip digital capture. The single false positive (bit 126, count=1) is consistent with the edge ambiguity of asynchronous sampling and does not alter the conclusion. This result extends prior work by showing that an external oscilloscope is not required: the leakage mechanism can be observed using instrumentation already present or insertable on the die. In practical terms, any adversary who can establish a direct electrical path from a secret bit to an on-chip instrument—whether through malicious wiring, compromised debug logic, or a bitstream modification—faces a trivially easy observation task.

5.2 GPIO-Level Threshold Behaviour

The GPIO-level threshold-path experiment does not model true on-chip distributed interconnect resistance. It is a GPIO-level emulation in which the

RO enable must cross an FPGA I/O boundary. The abrupt threshold behaviour is a consequence of the I/O buffer's digital input stage. On-chip interconnect degradation, such as a poorly driven net within the fabric, could behave differently because there is no input buffer with a sharp threshold.

Despite this limitation, the experiment has evidential value within its stated scope. It shows that, for this GPIO-level threshold-limited path, moderate attenuation did not reduce leakage until the sampled input collapsed to a constant zero.

5.3 Interpretation of the Proximity-Coupling Results

The proximity-coupling study exposes both a measurable separability condition and important methodological boundaries. The key result is that near placement is not adequately represented by the N=5 case; increasing RO length turns a weak, overlapping measurement into clean calibrated separation. This suggests that instrument geometry can amplify weak physical interaction enough for labelled key-bit observability, but the effect decays strongly with distance. Medium placement reaches only partial separability (best 71.09% at N=65), and far placement remains weak and overlapping (best 60.55% at N=75).

These numbers should not be interpreted as proof that pure cross-talk is harmless at medium and far distances. The accuracy values report *single-measurement* supervised separability under ideal laboratory conditions: a fixed key, a quiet device, no concurrent workload noise, and a threshold calibrated from the same data. An actual adversary is not constrained to a single measurement. The implications of repeated observation are examined in the next section.

At the same time, placement constraints alone do not prove a specific adjacent aggressor–receiver track pair. Without post-route route inspection or parasitic extraction, the physical coupling geometry is unknown. The single-capture design also leaves repeatability and environmental robustness untested. These limitations motivate the conservative interpretation of the proximity results as evidence of *implementation-controlled* observability rather than as a precise physical model.

5.4 Cumulative Attack Risk

A central insight that emerges from the proximity-coupling matrix is that single-measurement accuracy understates the practical security risk. Even when a single observation achieves only 60–70% supervised separability—

barely above random guessing—an adversary can improve recovery probability through three complementary strategies: repeated measurement, spatial aggregation, and side-channel fusion.

Repeated measurement. If the physical coupling mechanism is stable over time, an attacker who observes the same key bit multiple times can average or vote across measurements to reduce noise. A simple majority vote across n independent and identically distributed (i.i.d.) observations with per-observation accuracy p yields a collective accuracy given by the binomial cumulative distribution:

$$P_{\text{maj}}(n, p) = \sum_{k=\lceil n/2 \rceil}^{n} \binom{n}{k} p^k (1-p)^{n-k} \tag{5.1}$$

Equation 5.1 quantifies the probability that at least half of the n independent observations agree with the true bit value. For example, at $p = 0.60$ (close to the far-placement best case), Equation 5.1 gives $P_{\text{maj}}(11, 0.60) \approx 75.4\%$ and $P_{\text{maj}}(101, 0.60) \approx 97.9\%$. The increase is not linear: the collective accuracy grows sigmoidally with n, crossing 90% at $n \approx 35$ for $p = 0.60$ and at $n \approx 15$ for $p = 0.70$.

This model assumes statistical independence between observations, which is an idealisation. In practice, temporal correlation in supply noise, temperature drift, or measurement-setup jitter can reduce the effective number of independent samples. Even with moderate correlation, however, the principle that repeated observation amplifies weak signals is well established in side-channel analysis, where averaging across traces is a standard noise-reduction technique [17]. The key point is that a single-measurement accuracy of 60–70% is not a reliable safety indicator when the attacker can repeat the observation many times.

Spatial aggregation. A sophisticated adversary is not limited to a single instrument. If multiple ROs or other sensors can be placed near the same secret signal, their readouts can be combined. The near-placement matrix shows that a sufficiently long RO achieves perfect separability; an attacker who can instantiate several ROs of different lengths and select the best-performing one effectively bypasses the weak short-RO case. In multi-tenant FPGA environments, this corresponds to a malicious tenant flooding the fabric with instrument variants and using the one that happens to fall closest to the victim routing.

Side-channel fusion. The proximity-coupling measurement is not the only source of information available to an attacker. Power consumption traces, electromagnetic emanations, timing measurements, and other on-chip sensors all leak correlated information about the same secret. An attacker who combines weak proximity-coupling observations with even

weaker power-analysis traces may achieve a combined signal that exceeds any individual channel. The information-theoretic principle is straightforward: independent noisy observations of the same secret carry additive mutual information.

These considerations change the practical interpretation of the proximity-coupling matrix. A condition that achieves only 60.55% supervised separability in a single capture is not *safe*; it is merely *harder*. The security boundary is not the point at which single-measurement accuracy drops to 50%, but rather the point at which the coupling is so weak that no feasible amount of repetition, aggregation, or fusion can reconstruct useful information. Determining that boundary requires a formal information-theoretic treatment that is beyond the scope of the present controlled measurements, but the experiments provide the empirical foundation on which such an analysis could be built.

5.5 Relation to Prior Work

The direct-connection result supports Lundberg's observation that a small RO can be used as an effective leakage mechanism when it is directly controlled by a secret-dependent signal. The measurement method differs: an internal counter and ILA capture replace the external oscilloscope used in the earlier work. This shows that the leakage mechanism can be observed using only on-chip digital instrumentation, which is relevant for scenarios involving malicious IP, debug logic, or built-in monitors.

The GPIO-level threshold-path experiment adds a boundary case that is not present in the direct-connection attack. It shows that a weakened external path may remain fully observable as long as the digital input receiver still interprets the signal correctly. This supports the broader lesson that attenuating or externalising a sensitive signal does not necessarily remove key-bit observability.

The proximity-coupling result extends the long-wire leakage literature in a narrower direction. Prior work established that physically close routes can reveal activity or support covert communication [8, 10, 11]. The near-placement matrix here shows controlled key-bit separability under pure cross-talk conditions when the RO is long enough. The extension remains bounded: it is a labelled key-scan experiment with calibrated thresholds, not a demonstrated recovery of arbitrary AES state from naturally switching logic.

5.6 Threat Model and Scope

The demonstrated scenarios require the ability to drive or place a selected key-dependent signal near the measurement structure, the ability to read the RO count, and controlled sequential exposure of each key bit. This is narrower than a universal AES attack. It is most directly relevant to malicious IP, compromised bitstreams, or debug/instrumentation logic that can influence placement or access sensitive control paths. The results do not demonstrate a multi-tenant attack in which an untrusted tenant recovers a neighbouring tenant's key from arbitrary AES switching without calibration.

5.7 Countermeasures

The experimental findings motivate several design-time countermeasures. These are informed recommendations rather than experimentally validated defences.

Physical separation. Prior long-wire leakage work and the proximity-coupling matrix both motivate treating placement as security-relevant [8, 10, 11]. Sensitive key-dependent routing paths should be kept away from untrusted RO structures through floorplanning directives such as Vivado Pblock constraints. The measured near/medium/far trend suggests that separation can substantially reduce observability, although the required distance is technology- and routing-dependent and must be characterised experimentally.

Routing-aware verification. Security verification flows should include a proximity analysis pass that identifies routing tracks carrying static key-dependent signals that run parallel to RO structures or combinational loops for significant distances. This could be implemented as a post-route DRC extension that flags potential aggressor–victim pairs based on parallel run length and inter-track spacing. Routing-level defences proposed for crosstalk and voltage-based FPGA side channels support this type of post-route policy rather than relying only on RTL-level isolation [14, 15].

Guard routing. Grounded shielding tracks placed between sensitive nets and potential RO receiver structures would shunt coupled displacement current to ground rather than allowing it to perturb the receiver node. This technique is standard in analog and mixed-signal design and has conceptual overlap with FPGA routing defences against crosstalk leakage [15], though practical adoption would consume routing resources and require tool support.

RO monitoring and policy. Designs with multiple or atypically placed

ROs near cryptographic cores warrant manual scrutiny. In multi-tenant FPGA environments, unauthorised combinational loops should be detected and rejected during bitstream verification because RO-based monitors have been repeatedly used as sensors in remote or co-located FPGA side channels [12, 13]. The legitimate use of ROs for reliability monitoring complicates this approach, but policy can require that reliability ROs be placed in designated regions away from security-sensitive logic.

Key-storage design. Static key bits exposed on long routing tracks contribute to leakage risk. Techniques such as register-level key splitting (where no single path carries the full key) or key bit masking (where bits are XORed with a random mask before routing) may reduce the signal-to-noise ratio available to a proximity observer. These approaches come with area and latency overheads that must be evaluated per application.

Treatment of reliability sensors as security assets. A central organisational recommendation is that the placement and configuration of on-chip reliability instruments should be treated as a security-relevant design decision, not solely a reliability decision. An RO placed near a cryptographic core for legitimate temperature monitoring may inadvertently act as a proximity-based side-channel receiver. Design reviews should consider both the intended function and the potential dual use of every RO instance.

5.8 Limitations

Table 5.1 summarises the main limitations, their effect on interpretation, and why the corresponding result remains useful within the controlled thesis scope.

Table 5.1: Threats to validity and interpretation boundaries.

Limitation	Impact on interpretation	Why the result remains useful
Single FPGA device	Device-to-device and family-to-family variation are not characterised.	The results provide a concrete Artix-7 case study and identify measurements that should be repeated on additional devices.
Single fixed AES key	The key is imbalanced (145 ones, 111 zeros), so raw accuracy must be compared with constant-class baselines.	The fixed key enables controlled cross-scenario comparison and makes majority-baseline effects explicit.

Limitation	Impact on interpretation	Why the result remains useful
Single capture per measured condition	No confidence intervals or repeatability statistics are available.	The measurements still reveal large qualitative differences between direct, GPIO-threshold, near, medium, and far conditions.
Controlled key-scan model	Results do not directly generalise to naturally switching AES datapaths or CPA-style attacks.	The scan isolates physical observability of selected key bits before broader cryptanalytic claims are attempted.
Supervised cross-talk thresholding	Pure cross-talk accuracy is calibrated separability, not unsupervised key recovery.	It quantifies whether the two labelled classes are physically separable under the implemented placement and RO length.
No post-route adjacency quantification	The near, medium, and far groups are placement-controlled but not reduced to exact routed spacing or parallel run length.	The distance trend remains useful as an implementation-level warning and motivates route-level verification.
No parasitic extraction	Exact mutual capacitance and coupling-current values are unavailable.	The thesis reports measured RO-count effects without claiming a calibrated physical coupling model.
Mechanism attribution	Capacitive coupling, PDN coupling, substrate coupling, and routing artefacts cannot be isolated.	The security question is whether a separable observation exists; physical decomposition is identified as future work.
Coarse distance granularity	Three placement groups do not resolve the coupling decay function.	The near/medium/far comparison is sufficient to show that placement distance materially affects separability.
No concurrent workload noise	Robustness against unrelated activity, supply noise, and temperature drift is untested.	Variable isolation makes the first-order coupling trend easier to interpret.
Single-bit coupling	Only one selected key bit is driven at a time; multi-bit coupling is not investigated.	The setup provides a controlled lower-level measurement before studying realistic parallel switching.

Limitation	Impact on interpretation	Why the result remains useful
Fixed clock and measurement window	Frequency dependence and alternative window lengths are uncharacterised.	Keeping the 100 MHz clock and 2^{20}-cycle window fixed makes the coupling path the main variable.
No direct analog JA4 measurement	GPIO voltage values are model-based estimates rather than oscilloscope measurements at the pin.	The digital threshold transition is still directly observed in the recovered stream and RO counts.

5.9 Implications for Future Experiments

The most immediate continuation is repeatability and physical-design verification for the completed pure cross-talk matrix. Each distance group and RO length should be captured repeatedly and reported with confidence intervals. After place-and-route, the routed aggressor and receiver tracks should be inspected to quantify parallel run length, spacing, and, where possible, route-level capacitance. This would replace cell-placement intent with explicit coupling evidence.

The measurement method should then be strengthened through matched differential ROs, longer capture windows, repeated captures under controlled temperature and voltage conditions, and thresholds fixed from calibration data independent of the evaluated key-bit scan.

Finally, the scope should be broadened beyond one Artix-7 board and one fixed AES key. Repeating the study across FPGA families, key distributions, multi-bit coupling patterns, and realistic AES switching activity would clarify how far the controlled key-scan results generalise.

Chapter 6

Conclusion

6.1 Summary of Findings

This study examined whether on-chip reliability instrumentation can act as an effective side-channel receiver as the coupling path between a secret source and the instrument is systematically weakened. Using an FPGA as a controllable proof-of-concept platform, three experiment families spanned the coupling continuum from direct electrical connection through a threshold-limited return path to pure physical proximity without intentional logical wiring.

The direct-connection baseline demonstrates that a directly wired instrument can recover key information with near-perfect reliability (99.61% accuracy, $\Delta C \approx 38,600$) using only on-chip digital capture. The GPIO-level threshold-limited path reveals an abrupt failure boundary: observability remains at 100% until the returned signal falls below the digital input-buffer threshold, after which the instrument readout collapses to a constant-zero classifier. This shows that attenuating a sensitive signal path does not gradually reduce leakage but rather preserves full observability until the digital receiver can no longer resolve the signal.

The proximity-coupling matrix provides the first systematic measurement of how key-bit separability degrades with placement distance and instrument geometry under pure cross-talk conditions. Near placement achieves perfect supervised separability when the instrument is sufficiently long (N=45–75), while medium and far placement remain partially distinguishable or weak. The most important practical insight is that single-measurement accuracy understates the true security risk: even weak separability—when combined across repeated measurements, multiple instruments, or fused with other side channels—can pose a realistic threat that is not captured by supervised accuracy alone.

6.2 Recommendations

The findings motivate five design-time recommendations. First, the placement and configuration of on-chip reliability instruments should be treated as a security-relevant design decision, not solely a reliability decision. Second, sensitive key-dependent routing paths should be physically separated from untrusted instrumentation structures through floorplanning and placement constraints. Third, security verification flows should include proximity analysis that flags routing tracks carrying static key-dependent signals running parallel to instrument structures. Fourth, static key bits should not be exposed on debug, GPIO, or instrumentation paths that can drive instrument-like listener structures. Fifth, key-storage designs should consider register-level splitting or masking to reduce the signal-to-noise ratio available to a proximity observer.

6.3 Future Work

The immediate next step is to strengthen the physical-design evidence for the proximity-coupling matrix. Each distance group and instrument length should be inspected after place-and-route to quantify parallel run length and spacing, captured repeatedly under controlled conditions, and evaluated with thresholds fixed from independent calibration data. The cumulative attack risk analysis presented here should be formalised through an information-theoretic treatment that quantifies how much mutual information weak single-measurement separability provides when aggregated over time and space.

Beyond the proof-of-concept platform, the experiments should be repeated across FPGA families, ASIC test chips, and SoC platforms to verify that the qualitative conclusions transfer to different technologies and design styles. Repeating the study with multiple keys, realistic AES switching activity, concurrent workload noise, and variable environmental conditions would clarify how far the controlled key-scan observations generalise to practical attack settings. Ultimately, the goal is to develop quantitative design rules—minimum separation distances, maximum instrument densities, and verification checklists—that chip designers can use to bound the side-channel risk posed by on-chip reliability instrumentation.

Bibliography

[1] 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, 2014.

[2] E. Brier, C. Clavier, and F. Olivier, "Correlation power analysis with a leakage model," in *Proc. Cryptographic Hardware and Embedded Systems (CHES)*, ser. LNCS, vol. 3156, 2004, pp. 16–29.

[3] S. Chari, J. R. Rao, and P. Rohatgi, "Template attacks," in *Proc. Cryptographic Hardware and Embedded Systems (CHES)*, ser. LNCS, vol. 2523, 2003, pp. 13–28.

[4] B. Shakya, T. He, H. Salmani, D. Forte, S. Bhunia, and M. Tehranipoor, "Benchmarking of hardware Trojans and maliciously affected circuits," *J. Hardware and Systems Security*, vol. 1, no. 1, pp. 85–102, 2017.

[5] M. Tehranipoor and F. Koushanfar, "A survey of hardware Trojan taxonomy and detection," *IEEE Design & Test of Computers*, vol. 27, no. 1, pp. 10–25, 2010.

[6] J. Szefer, "A survey of FPGA security," in *Proc. Int. Conf. Field-Programmable Technology (FPT)*, 2013, pp. 1–8.

[7] Xilinx Inc., *7 Series FPGAs Data Sheet: Overview* (DS180), v2.6.1, 2020.

[8] I. Giechaskiel, K. B. Rasmussen, and J. Szefer, "Read between the lines: Detecting and locating hardware Trojans through on-chip activity," in *Proc. IEEE Symp. Security and Privacy (SP)*, 2020, pp. 1235–1252.

[9] K. Gandolfi, C. Mourtel, and F. Olivier, "Electromagnetic analysis: Concrete results," in *Proc. Cryptographic Hardware and Embedded Systems (CHES)*, ser. LNCS, vol. 2162, 2001, pp. 251–261.

[10] I. Giechaskiel, K. B. Rasmussen, and K. Eguro, "Leaky wires: Information leakage and covert communication between FPGA long wires," in *Proc. ACM Asia Conf. Computer and Communications Security (ASIACCS)*, 2018, pp. 15–27.

[11] G. Provelengios, C. Ramesh, S. B. Patil, K. Eguro, R. Tessier, and D. Holcomb, "Characterization of long wire data leakage in deep submicron FPGAs," in *Proc. ACM/SIGDA Int. Symp. Field-Programmable Gate Arrays (FPGA)*, 2019, pp. 292–297.

[12] M. Zhao and G. E. Suh, "FPGA-based remote power side-channel attacks," in *Proc. IEEE Symp. Security and Privacy (SP)*, 2018, pp. 229–244.

[13] F. Schellenberg, D. R. E. Gnad, A. Moradi, and M. B. Tahoori, "An inside job: Remote power analysis attacks on FPGAs," in *Proc. Design, Automation & Test in Europe Conf. (DATE)*, 2018, pp. 1111–1116.

[14] J. Krautter, D. R. E. Gnad, F. Schellenberg, A. Moradi, and M. B. Tahoori, "Active fences against voltage-based side channels in multi-tenant FPGAs," in *Proc. IEEE/ACM Int. Conf. Computer-Aided Design (ICCAD)*, 2019.

[15] Z. Seifoori, S. S. Mirzargar, and M. Stojilović, "Closing leaks: Routing against crosstalk side-channel attacks," in *Proc. ACM/SIGDA Int. Symp. Field-Programmable Gate Arrays (FPGA)*, 2020, pp. 197–203.

[16] P. Kocher, J. Jaffe, and B. Jun, "Differential power analysis," in *Proc. Advances in Cryptology (CRYPTO)*, 1999, pp. 388–397.

[17] S. Mangard, E. Oswald, and T. Popp, *Power Analysis Attacks: Revealing the Secrets of Smart Cards*. Springer, 2007.

[18] F. Lundberg, *Exploiting Ring Oscillators to Leak Secret Information*, Master's thesis, Dept. of Electrical and Information Technology, Lund University, Sweden, 2025.

[19] National Institute of Standards and Technology, "Advanced Encryption Standard (AES)," FIPS PUB 197, 2001.

[20] Xilinx Inc., *Vivado Design Suite User Guide: Using Constraints* (UG903), 2020.

[21] J. Zhang, G. Qu, and Q. Zhou, "Detecting and preventing side-channel attacks on FPGA," *IEEE Trans. VLSI*, vol. 23, no. 12, pp. 2969–2982, 2015.

[22] A. Moradi, A. Barenghi, T. Kasper, and C. Paar, "On the vulnerability of FPGA bitstream encryption against power analysis attacks," in *Proc. ACM CCS*, 2012, pp. 111–122.

[23] Digilent Inc., *Nexys 4 Reference Manual*, 2016.

[24] Xilinx Inc., *7 Series FPGAs SelectIO Resources User Guide* (UG471), v1.10, 2018. Available: https://docs.amd.com/v/u/en-US/ug471_7Series_SelectIO

Chapter A

Supplementary Implementation Details

A.1 Key-Scan Controller FSM

A.1.1 State Table

Table A.1 records the four-state controller used to expose, measure, and stream each selected key bit.

Table A.1: Key-scan FSM states used to expose one key bit, measure the RO count, store the result, and later stream one labelled output row per scan step.

State	Cycles	Action
IDLE	Until trigger	Load key, reset pointer, wait for start
SELECT	Short settling interval	Drive selected bit onto coupling path, settle
MEASURE	2^{20}	Enable sampled RO-transition counter, accumulate, store result
STREAM	Post-scan stream	Emit compact stored result rows to ILA probes

A.1.2 Signal List

- `key_shift_reg[255:0]`: AES-256 key loaded at IDLE.
- `bit_pointer[7:0]`: Current scan index (0–255).
- `selected_bit`: The key bit at the current pointer.
- `coupling_path_out`: Signal driven onto the coupling medium (direct wire, GPIO, or aggressor track).
- `ro_enable`: Enables the RO counter during MEASURE.
- `ro_count[31:0]`: Accumulated oscillation edges.

- `stream_valid`: Indicates that the stream bus is carrying post-scan result rows.
- `stream_bit_index`: Bit index associated with the streamed row.
- `stream_actual_bit`: Ground-truth key bit used for evaluation.
- `stream_recovered_bit`: Online bit decision inferred from the RO-count threshold.
- `stream_ro_count`: Captured RO counter value.

A.2 RO Implementation Excerpt

The following RTL excerpt documents the RO structure used to describe the implemented architecture. It is included as an architectural record of the design pattern; synthesis was performed from the repository RTL rather than from this abbreviated listing.

```
// Parameterised N-stage ring oscillator
// N must be odd. KEEP and DONT_TOUCH are applied to all instances.
module ring_oscillator #(parameter N = 5) (
    input  wire        enable,
    output wire        ring_out
);
    wire [N-1:0] stage;
    genvar i;

    // Stage 0: NAND-based enable control
    LUT6 #(.INIT(64'hFFFFFFFFFFFDFFFF)) stage0 (
        .O(stage[0]),
        .I0(enable), .I1(stage[N-1]),
        .I2(1'b0), .I3(1'b0), .I4(1'b0), .I5(1'b0)
    );

    // Stages 1..N-1: LUT-based inverters
    for (i = 1; i < N; i = i + 1) begin : inv
        LUT6 #(.INIT(64'h00000000000000FF)) stageN (
            .O(stage[i]), .I0(stage[i-1]),
            .I1(1'b0), .I2(1'b0), .I3(1'b0),
            .I4(1'b0), .I5(1'b0)
        );
    end

    assign ring_out = stage[N-1];
endmodule
```

Consequently, the listing should be read as an architectural excerpt. The compiled project source remains authoritative for exact LUT INIT values, receiver-stub insertion, and generate/endgenerate syntax.

A.3 XDC Constraint Strategy

Table A.2 summarises the main synthesis and placement constraints used to preserve the RO structures and define approximate placement groups.

Table A.2: Constraint summary, including the distinction between preserving cells and proving post-route wire adjacency.

Constraint	Object	Purpose	Limitation
`KEEP=TRUE`	Net, cell	Preserve through synthesis	No route control
`DONT_TOUCH`	Cell	Disable re-synthesis	Increases compile time
`LOC`	Cell	Fix slice location	Does not guarantee wire adjacency
`BEL`	Cell	Fix LUT within slice	Intra-slice routing not fully controlled
`PRESERVE`	Net	Prevent net removal	Does not fix routing path

`LOC` and `BEL` fix cell placement but do not guarantee that the routing between two constrained cells follows a specific track or maintains a specific edge-to-edge distance. The distance groups (near, medium, far) are therefore interpreted as CLB-location separation groups rather than post-route wire-spacing measurements.

A.4 ILA Export and Data Reduction Pipeline

A.4.1 ILA Export Mechanism

The ILA export should be distinguished from the per-bit measurement interval. The hardware counter accumulates sampled RO transitions during each 2^{20}-cycle MEASURE state. The count is then stored in an on-chip array. Only after all 256 key bits have been measured does the controller assert `stream_frame_start` and stream compact result rows to the ILA probes.

The exported ILA CSV therefore does not contain a raw 2.7 s waveform of `ring_out`. It contains a 4,096-sample capture window around the compact stream interface. The trigger is the first `stream_frame_start` pulse. No hardware capture qualification was used to keep only valid rows in the ILA buffer. Instead, `stream_valid` is recorded as a probe, and the offline reducer keeps the first 256 valid streamed rows. Some CSV exports split the 32-bit count bus into high and low fields; the reducer recombines those fields before computing summary metrics.

A.4.2 Pseudocode

```
THRESHOLD_CROSS = 0  # set per-condition; see below

raw_rows = read_ila_csv(capture.csv)
valid_stream_rows = [r for r in raw_rows if r.stream_valid == '1']

# The stream bus may remain valid after the first 256 rows.
# The evaluated key scan uses the first complete 256-row frame.
rows = []
for each row in valid_stream_rows[0:256]:
    bit_idx  = int(row.stream_bit_index)
    actual   = int(row.stream_actual_bit)
    ro_count = decode_stream_count(row)

    # Scenario-specific thresholding
    if scenario == 'direct' or scenario == 'gpio_threshold':
        # bit=0 gives count=0, bit=1 gives count>30000
        recovered = 1 if ro_count > 10 else 0
    elif scenario == 'crosstalk':
        # Supervised separability metric only;
        # not an unsupervised attack threshold.
        recovered = 1 if ro_count > THRESHOLD_CROSS else 0

    rows.append((bit_idx, actual, recovered, ro_count))

# For cross-talk, THRESHOLD_CROSS is set per-condition as:
# THRESHOLD_CROSS = (mean_bit0_counts + mean_bit1_counts) / 2
# estimated from the data after a calibration scan

n_correct = sum(1 for r in rows if r.actual == r.recovered)
accuracy = n_correct / 256
mu1 = mean(r.ro_count for r in rows if r.actual == 1)
mu0 = mean(r.ro_count for r in rows if r.actual == 0)
delta_c = mu1 - mu0
majority_baseline = 145/256  # for this fixed key
overlap = (min_bit1 <= max_bit0) or (min_bit0 <= max_bit1)
```

A.4.3 Threshold Strategy for Cross-Talk

For measured pure cross-talk experiments, both bit classes may produce large non-zero RO counts, so a fixed low threshold (e.g., 10) does not separate them. A per-condition midpoint threshold derived from labelled class means, $T = (\mu_0 + \mu_1)/2$, is used to report supervised separability after a calibration pass. It must not be presented as an unsupervised recovery method.

A.5 Notes on Pure Cross-Talk Matrix Data

The pure cross-talk dataset contains 24 measured conditions: three placement groups (near, medium, and far) crossed with eight RO lengths (N=5, 15, 25, 35, 45, 55, 65, and 75). The matrix is reported in Chapter 4 and is summarised as follows:

- Near placement becomes separable when the RO is sufficiently long: N=35 reaches 99.61%, and N=45–75 reaches 100% supervised separability.
- The strongest near-placement separation occurs at N=55, with $\Delta C \approx 6{,}636$.
- Medium placement is partially distinguishable but overlapping, with a best result of 71.09% at N=65.
- Far placement remains weak and overlapping, with a maximum accuracy of 60.55% at N=75.
- Post-route adjacency proof, parasitic extraction, repeated captures, and unsupervised threshold selection remain future work.

A.6 Cross-Scenario Visual Summary

Figure A.1 provides a visual companion to Table 4.6. It is placed in the appendix to keep the main results section focused on the numerical comparison table.

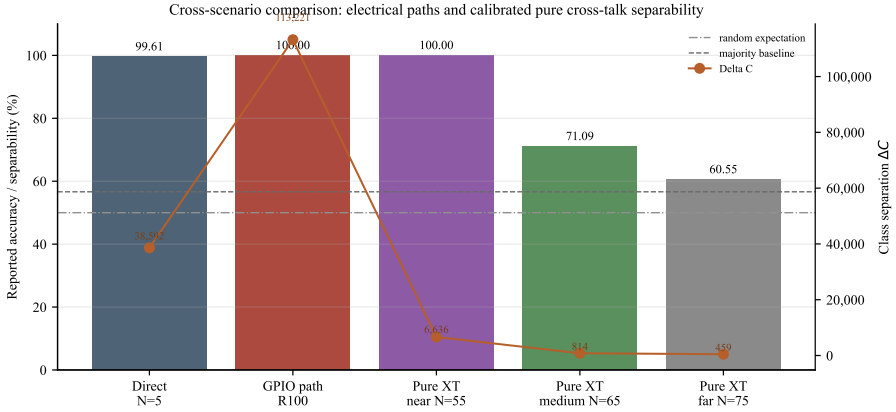

Figure A.1: Visual comparison of representative completed scenarios. Direct and GPIO-level threshold paths give the largest separations; pure cross-talk is reported as calibrated separability and degrades sharply at medium and far distances.

Printed by Tryckeriet i E-huset, Lund 2026

LUND
UNIVERSITY

Series of Master´s theses
Department of Electrical and Information Technology
LU/LTH-EIT 2026-1163
http://www.eit.lth.se