

NuRedact: Non-Uniform eFPGA Architecture for Low-Overhead and Secure IP Redaction

Voktho Das, Kimia Azar, Hadi Kamali

Department of Electrical and Computer Engineering (ECE), University of Central Florida, Orlando, FL 32816, USA
{voktho.das, azar, kamali}@ucf.edu

Abstract—While logic locking has been extensively studied as a countermeasure against integrated circuit (IC) supply chain threats, recent research has shifted toward reconfigurable-based redaction techniques, e.g., LUT- and eFPGA-based schemes. While these approaches raise the bar against attacks, they incur substantial overhead, much of which arises not from genuine functional reconfigurability need, but from artificial complexity intended solely to frustrate reverse engineering (RE). As a result, fabrics are often underutilized, and security is achieved at disproportionate cost. This paper introduces *NuRedact*, the first full-custom eFPGA redaction framework that embraces architectural non-uniformity to balance security and efficiency. Built as an extension of the widely adopted OpenFPGA infrastructure, *NuRedact* introduces a three-stage methodology: (i) custom fabric generation with pin-mapping irregularity, (ii) VPR-level modifications to enable non-uniform placement guided by an automated Python-based optimizer, and (iii) redaction-aware reconfiguration and mapping of target IP modules. Experimental results show up to 9× area reduction compared to conventional uniform fabrics, achieving competitive efficiency with LUT-based and even transistor-level redaction techniques while retaining strong resilience. From a security perspective, *NuRedact* fabrics are evaluated against state-of-the-art attack models, including SAT-based, cyclic, and sequential variants, and show enhanced resilience while maintaining practical design overheads.

Index Terms—Embedded FPGA, Hardware Protection, IP Redaction, Non-uniform eFPGA

I. INTRODUCTION

Outsourcing critical stages of IC development to global parties has unlocked scale and cost advantages, but it has also exposed designs to IP piracy, RE, and unauthorized overproduction [1]. Logic locking emerged as a systematized countermeasure [2], yet attacks on logic locking (i.e., de-obfuscation attacks), both logical and physical [3]–[6], have eroded the practical security of many locking breeds under realistic threat models, motivating a reassessment of protection primitives [2]. Over the last few years, a natural pivot has been shift towards reconfigurable-based *redaction*, which isolates security-critical logic cones and IPs and replaces them with reconfigurable substrate [7]–[12], typically LUT or embedded FPGA (eFPGAs) fabrics, so that functionality is convoluted and only instantiated with a protected configuration (bitstream).

Early redaction flows have automated module selection and clustering for either LUT-based [13] and eFPGA-based [14] redaction. However, they have fixed the target fabric architecture/placement (particularly in eFPGA-based studies), leaving significant PPA slack [14]–[16]. The last two years have seen a sharper focus on fabric structure and granularity [10], [11], [16]–[19]. HIPR [17] inserts fine-grain configurable

blocks (including custom LUTs, sequential blocks, and interconnects) for Boolean logic, sequential logic, interconnect randomization, respectively, then compacts the configuration to curb PPA/bitstream cost while maintaining resistance to functional and structural attacks. ARIANNA [18] closes the loop between what to redact and how to build the fabric, by identifying a secure subset of eFPGA parameters up front and tailors bespoke fabrics per module-cluster. TRAP [19] moves below the gate level by a transistor-programmable fabric whose switch-level semantics fundamentally complicate attacks (e.g., the SAT attack [3], [20]–[24]).

While these approaches mitigate overhead, they still rely on uniform attributes that provision routing capacity, LUT structures, and tile regularity for generality rather than redaction specificity. As a result, when different logic cones are redacted, substantial configuration, logic, and routing resources remain underutilized [17]–[19]. Moreover, bitstream storage and configuration flip-flops incur nontrivial silicon area and timing costs. Such customizations may also leak structural cues, enabling oracle-less or structure-guided attacks.

Among these approaches, despite overhead concerns, eFPGAs remain attractive for integration and tooling, from open-source IP generators (e.g., OpenFPGA [25]), to commercial IP (Achronix Speedcore, Flex Logix EFLX, QuickLogic), being integrated in SoCs, with the value proposition of post-silicon patchability [25]. While eFPGAs are deployed to support future reconfiguration, in redaction (used for IP protection), these substrates are used to instantiate complexity (a large functional/search space) and to decouple revealed structure from protected behavior. This means the *fabric’s non-uniformity can be co-designed with the target IP*: redaction needs **programmability as a means** (security), not as a **product feature**.

Motivated by this fact, we propose *NuRedact*, an automated framework for non-uniform eFPGA redaction, realizing the need for programmability as a means of protection at the lowest overhead, while retaining security and maximizing integration and tooling. *NuRedact* replaces IP sub-components with irregular, design-specific fabrics generated atop OpenFPGA, modifies VPR to legalize non-uniform placement, and includes a redaction-aware mapper that exploits irregular pin-mapping and tailored resources. Compared to HIPR and TRAP, *NuRedact* preserves mainstream RTL-to-OpenFPGA automation and industry tool compatibility at comparable overhead while cutting area by up to 9x vs. baseline eFPGA-based redaction. The contributions of *NuRedact* are as follows:

TABLE I: Top-level Comparison of Redaction Approaches. (■□□□□ = low, ■■■■■■ = high). ✓ = present, ✗ = absent, ~ = partial/indirect).

Technique	Granularity	Overhead	Bitstream	Interconnect obfuscation	Custom fabric	Tooling & integration	Highlight
LUT-Lock [7]	LUT obfuscation	■■■□□	■■■□□	✗	✗	✓	SAT-hardening
Full-Lock [8]	Configurable Routing + LUT	■■■□□	■■■□□	✓	✗	✗	SAT-hardening + Proof
ALICE [15]	eFPGA (fixed fabric)	■■■■■	■■■□□	✗	✗	✓	End-to-end Flow
SheLL [10]	eFPGA (tile shrinking)	■■■□□	■■■□□	~	✓	✓	↓~55% vs. ALICE [15]
HIPR [17]	LUT + seq + interconnect	■■□□□	■■□□□	✓	✗	✓	Security-aware compaction
ARIANNA [18]	eFPGA (bespoke, uniform tiles)	■■□□□	■■■□□	~	✓	✓	3.3× lower vs. fixed fabric
SOAR [12]	eFPGA + Monitoring	■■■■■	■■■□□	✓	✗	✓	Dynamic Key
TRAP [19]	Transistor-level fabric	■□□□□	■■■□□	✗	✗	~	SAT-hardening
NuRedact (proposed)	eFPGA (non-uniform)	■■□□□	■■■□□	~	✓	✓	Low underutilization Low Black-box Learnability up to 9x smaller vs. fixed

Ratings are qualitative and synthesized from reported results in prior work.

(i) An automated framework for generating design-specific, **non-uniform** eFPGA fabrics that reduce area overhead while preserving full functionality of the redacted logic.

(ii) A Python-based package for layout transformation and VPR architecture customization on a SkyWater 130nm OpenFPGA architecture, for toolchain-compatible fabric generation.

(iii) A *security-driven overhead-aware selective redaction*, replacing critical logic maximizing overhead reduction with irregular, non-uniform fabrics.

(iv) Evaluation of overhead using Cadence Genus for synthesis and Synopsys Design Compiler (DC) to produce attack-friendly netlists for evaluation (evaluated by IcySAT unrolling and the KC2 SAT toolchain).

II. BACKGROUND AND PRIOR WORK

A. Hardware IP Protection through Redaction

Prior art in redaction can be divided into three main categories: (i) fine-grain LUT redaction, (ii) coarse-grain eFPGA redaction, and (iii) sub-gate transistor-level redaction:

(i) *Fine-grain LUT redaction*. These techniques redact logic by replacing selected gates/cones with small LUTs [7], [13], [17], [26], [27]. These techniques control overhead by compacting LUTs [7], using attack-resistive insertion policies [7], [13], [17], and bitstream compression [17]. Advantages of these techniques are layout friendliness, medium-effort integration into standard netlists. This is while structure may leak if compaction/placement is careless [17], [27].

(ii) *Fine-grain LUT+Routing redaction*. Several studies, including HIPR [17], extend LUT-based redaction by configuring cascaded twisted multiplexers to realize both LUT functionality and routing crossbars [8], [28]. This design closely resembles FPGA architectures, in which part of the configuration enables logic operations while the remainder supports routing recovery. Despite its robustness, the approach is vulnerable to link-prediction-based machine-learning attacks [29].

(iii) *Coarse-grain eFPGA redaction*. RTL modules (or IPs) are moved behind an eFPGA macro, where early end-to-end flows automated what to redact [15]. More recent frameworks tailor the fabric to each cluster (bespoke K/N/IO/tile grids) after pre-filtering to secure configurations, improving utilization and

cutting overhead [18]. These techniques benefit from strong decoupling between revealed structure and hidden behavior and mature tool flows (OpenFPGA/VPR) for full automation. However, they are often over-provisioned for redaction, wasting logic/routing and driving overhead. Additionally, configuration registers and I/O pins can feed attack inferences if, leading to function/bitstream leakage [30]–[33].

(iv) *Sub-gate transistor-level redaction*. Switch-level fabrics purposefully break gate-level assumptions for overhead mitigation [19]. Similar approaches have been used in the past for only routing-based obfuscation [34], where more recent ones shift to full configurable substrate at transistor-level. These approaches eliminate hierarchical LUT cues. However, tooling (automation) is less turnkey than LUT/eFPGA flows.

Table I shows a detailed comparison of key studies in these three categories. *NuRedact* keeps the integration/automation of eFPGA redaction but discards uniformity. By co-designing non-uniform fabrics with the target cones, it retains the redaction security model while sharply reducing under-utilization and area—precisely the gap revealed by the prior art above.

B. Custom eFPGA Design Frameworks

An eFPGA redaction flow couples (i) an architecture description (e.g., VTR/VPR XML) that defines grid layout, tiles, logic clusters, routing, etc. (ii) a fabric generator (e.g., OpenFPGA) that binds architectural primitives to RTL and bitstream tooling, and (iii) an ASIC integration path that compiles eFPGA macro with the vendor’s compiler (e.g., Cadence Genus) for timing, floorplaning, DFT, etc. [25], [35]. Various open-source and commercial frameworks have been introduced for the integration of eFPGA into ASICs, i.e., eFPGA-based system-on-chips (SoCs), where VTR/VPR [35] + OpenFPGA [25] are widely used for prototyping. This tooling stacks provide key knobs for configurations, including logic clustering (LUT size and fracturability, carry chain, and hard blocks), routing topology (e.g., channel width, switch-box types, and crossbar topology), clocking, and configuration protocol (frame-based or memory-based) [10], [25].

Crucially, although fabrics are often presented as perfectly tileable and uniform, both VTR/VPR and OpenFPGA natively

support heterogeneity and irregularity [25], [35], including mixing tile types, using irregular pin maps, depopulating connection/switch boxes asymmetrically, and varying channel width. Legality is preserved through pin-equivalence classes, placement regions, and explicit resource constraints, so the standard pack/place/route remains correct even when the fabric deviates from a checkerboard (mesh model) [35], [36]. For redaction, these capabilities matter because the objective is not end-user flexibility but a large functional/search space at low silicon cost and with minimal structural regularity. *NuRedact* implements redaction by co-designing the fabric with the target cones using irregular tilings, pin-map asymmetry, and selective logic depopulation, then enforces legality through explicit placer/router constraints and binding metadata, thus retaining security while lowering overhead and integrating cleanly with OpenFPGA/VTR and standard ASIC flows.

C. Attacks and Security Analysis of eFPGA Redaction

Attacks on redaction aim to (i) recover functionality (partial or full) [30], [32] or (ii) shrink the search space until the remaining configuration bits are solvable. [16], [33]. For LUT-based schemes, early random LUT insertion fell quickly to SAT de-obfuscation [37]. Later heuristics improved hardness but leaked structure, where ML-based attacks target missing k-cuts infer LUT placement/keys [29]. For eFPGA-based redaction, the uniformity of the fabric blunts direct structural inference, but black-box functional attacks remain viable, e.g., FuncTeller that queries the eFPGA I/O, identifies on-set minterms, expands them to prime implicants, and synthesizes an approximate netlist [30]. In parallel, SAT derivations, e.g., sequential/cyclic SAT (e.g., unrolling), can collapse the effective configuration space once cycles are handled, enabling oracle-guided recovery [23], [24], [38]. More recently, data-driven approaches such as MANTIS learn approximate bitstreams from observed I/O behavior without an explicit SAT model, underscoring that the attack surface extends beyond classic logic-level formulations [32]. These trends point to architectural knobs that reduce attack success (targeted by *NuRedact*): (i) limit structural regularity that ML models and oracle-less methods exploit and (ii) raise entropy density by avoiding low-utilization, over-provisioned eFPGA fabrics.

III. THREAT MODEL

In line with prior redaction studies, the assets (attack target) would be (i) the functionality of the redacted cones, (ii) the eFPGA configuration (bitstream) that realizes them, and (iii) the mapping between architectural bits and configuration addresses. Attacker goal is to replicate the protected logic by recovering a correct (or approximate) configuration, or by reconstructing the redacted functionality. The attacker has the redacted gate-level netlist of the ASIC and can apply input–output queries to a working device (oracle). Attack surfaces we evaluate is Oracle-guided SAT family, black-box functional recovery, and Oracle-less/structural guidance. We assume bitstreams are not available in plaintext, and configuration interfaces are placed within the system’s security perimeter. Additionally, invasive physical extraction and advanced side-channel attacks on the

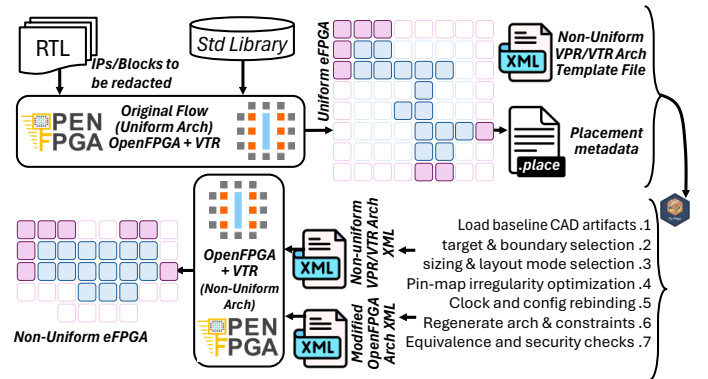


Fig. 1: *NuRedact* Framework: Automated Generation of Custom Non-Uniform eFPGA Fabrics by Leveraging Placement Data from Uniform Flow and Integrating Python-Based VPR/VTR XML Customization within OpenFPGA.

bitstream loader are important but excluded here to align with prior redaction work.

IV. PROPOSED REDACTION SCHEME: *NuRedact*

NuRedact is a fully automated flow that synthesizes design-specific, a first-of-its-kind non-uniform eFPGA fabrics for redaction. To guarantee reliable automation (keep everything running on widely-used frameworks, e.g., OpenFPGA and VPR/VTR), the key idea is to extract structured metadata from the baseline VPR/OpenFPGA placement/routing (from the baseline uniform run), then regenerate a custom layout that preserves function and CAD legality while removing under-utilized logic/routing and injecting architectural irregularity (pin-maps, logic depopulation, regional segments). The result is a compact fabric that matches the redacted cone’s footprint, integrates with the unmodified toolchain, and maintains security at much lower overhead. The overall implementation flow of *NuRedact* is shown in Figure 1, highlighting the integration of our custom package for layout transformation, customized VPR/VTR architecture generation, and OpenFPGA run.

A. Uniform Fabric Analysis and Logic Utilization Extraction

NuRedact starts the process from a regular (uniform) eFPGA fabric, composed of a regular array of homogeneous logic tiles and I/O pins. The design is packed/placed/routed using the standard OpenFPGA+VPR flow, yielding the canonical artifacts (including `.place` that records the spatial coordinates of all placed logic blocks and I/O elements). By processing `.place` log, *NuRedact* extracts a utilization map (occupied tiles and I/Os), a routing map, and timing criticality per net. This analysis enables the elimination of unused logic and I/O resources, initiating customized non-uniform layout that preserves only the actively used tiles and I/Os.

B. Automated Generation of Non-Uniform Layouts

Given the placement data, a Python-based layout transformation framework is implemented in *NuRedact* that constructs a customized non-uniform architecture (w.r.t. structured metadata from `.place` log). Based on two inputs, (i) structured metadata from `.place` log, and (ii) parameterized VPR architecture template XML, this Python-based framework generates a new architecture XML file with an updated layout and routing

Algorithm 1 *NuRedact*: Non-Uniform Fabric Generation

```

1: Input:  $P$ : placement file,  $A_{temp}$ : template arch.
2: Output:  $A_{nu}$ : non-uniform arch.
3:  $C_{io} \leftarrow$  IO capacity from  $A_{temp}$ 
4:  $(N_{clb}, N_{io}) \leftarrow \text{parse}(P)$ 
5:  $B_{io} \leftarrow \left\lceil \frac{N_{io}}{C_{io}} \right\rceil$ 
6:  $(W, H) \leftarrow \min_{w,h} (w+2)(h+2)$  s.t.  $wh \geq N_{clb}$ ,  $2(w+h) \geq B_{io}$ 
7: Place  $N_{clb}$  CLBs  $\rightarrow \mathcal{C}$ 
8: Boundary IOs  $\rightarrow \mathcal{I}$ 
9: while  $|\mathcal{I}| < B_{io}$  do
10:    $\mathcal{I} \leftarrow \mathcal{I} \cup \mathcal{P}$ 
11: end while
12:  $A_{nu} \leftarrow \mathcal{C} \cup \mathcal{I}$ 
13: Export  $A_{nu} \rightarrow$  VPR XML
  
```

configuration that reflects a reduced and design-specific layout while maintaining compatibility with VPR’s input requirements, as illustrated in Algorithm 1.

Unlike conventional VPR architectures that utilize `auto_layout` to define tile arrangement, the customized non-uniform architectures rely on a *fixed_layout* specification. This allows for the explicit definition of tile locations, enabling precise control over the spatial organization of logic tiles and I/O pins. The overall array size in the fixed layout is not static but instead dynamically determined by the spatial extent of the utilized tiles in the original placement. As a result, the final fabric dimensions directly reflect the design’s actual resource demands. Our framework supports two layout strategies:

(i) *Conservative Reduction*: It preserves the original spatial coordinates of all utilized logic tiles and removes all unused tiles. I/O tiles are retained only up to the required number, with at least one I/O per side to maintain routing accessibility. This approach ensures minimal disruption to routing and timing while achieving resource reduction. An example of this strategy is shown in Figure 2(b).

(ii) *Layout Compaction*: It applies aggressive compaction by relocating all utilized logic tiles into a contiguous, tightly packed region. This relocation minimizes the overall fabric dimensions and improves spatial locality. I/O tiles are selected based on adjacency to used logic tiles, and the number of I/O tiles is reduced by adjusting per-tile I/O capacity where applicable. An example of this strategy is shown in Figure 2(c). To support flexible routing in the resulting compact layouts, *NuRedact* uses an additional routing segment into the customized VPR XML. In particular, a short-length segment (L1) is introduced alongside the default long-length segment (L4). While L4 spans four logic blocks and eases longer communication across the fabric, the newly added L1 segment is limited to a single block length. This finer routing improves local connectivity and is particularly beneficial in densely packed, non-uniform layouts where short interconnects dominate. The combination of L1 and L4 routing resources enhances routability and mitigates congestion that may arise due to irregular tile spacing.

C. Fabric Regeneration and Flow Compatibility

The custom architecture file produced by the transformation script is used to regenerate the fabric through a second invocation of the OpenFPGA flow. This stage produces the non-uniform eFPGA fabric, including routing infrastructure and




Fig. 2: Uniform and Non-Uniform eFPGAs in *NuRedact* using (b) Conservative Reduction and (c) Aggressive Layout Compaction.

configuration bitstream generation support. The regenerated architecture remains compatible with the original RTL design and synthesis results. The process preserves all logical mappings and ensures consistent functionality, while delivering a reduced-area implementation. Integration into OpenFPGA’s task-based infrastructure allows the entire process, starting from a uniform run and ending with a synthesized non-uniform fabric, to be executed automatically without manual intervention. *NuRedact* enables the construction of logic-agnostic, general-purpose non-uniform eFPGA fabrics that are tightly matched to actual design requirements. By eliminating unused tiles and compacting placement, the resulting architectures demonstrate improved spatial efficiency and reduced overhead without sacrificing correctness or tool compatibility.

V. SECURE IP REDACTION USING *NuRedact* FRAMEWORK

From the redaction standpoint, *NuRedact* introduces spatially constrained intra-IP redaction that eases selective redaction of critical sub-components with low external observability, e.g., arithmetic blocks of encryption cores, control logic, or decision-making units. As *NuRedact* maximizes the utilization ratio per logic tile (using logic depopulation), the ratio of logic to I/O would be increased drastically (denser logic in a smaller footprint and fewer I/O pins). Accordingly, this non-uniform redaction, coupled with pin-map asymmetry, minimizes black-box learnability by algorithmic attacks, e.g., SAT-derived attacks. Our experiments (see Section VI-B) demonstrate that this dense logic in a smaller footprint increases SAT variables per bitstream size by one order of magnitude, showcasing the security enhancement at lower PPA overhead.

Additionally, in *NuRedact*, ASIC to/from eFPGA boundary is chosen by a simple multi-objective policy: minimize interface width (I/O cut) and exposed controllability/observability, subject to timing slack and floorplan constraints. In practice, we rank candidate cuts by (i) cut-signal count and toggle/observe metrics, (ii) uniqueness/rarity of the logic, and (iii) a detour budget for top-critical nets. Architectural irregularity is introduced by pin-map asymmetry (permuting equivalent pins to break port regularity) and logic/routing depopulation to further strengthen security by eliminating predictable architectural patterns. Compared to square, uniform overlays, this (i) removes under-utilized silicon, (ii) reduces visual and graph-structural cues that aid oracle-less inference, and (iii) raises configuration entropy per logic, all while running through standard OpenFPGA/VTR and sign-off flows.

TABLE II: Specifications of the selected Benchmark Circuits.

Benchmark	# Modules	# Inputs	# Outputs	Redacted Modules
GPS	12	6-128	1-256	Cacode
RISC-V	18	3-130	1-79	CTRL Arbiter Ld/St
DES3	11	6-240	4-64	SBOX_8
AES	9	10-128	8-128	AES_ShR
ADDER	2	3-4	2-3	FA_ARRAY HA_ARRAY
GCD	8	8-45	1-18	Comparator

VI. EXPERIMENTAL RESULTS AND ANALYSIS

To evaluate the effectiveness of the proposed *NuRedact* framework, we conducted experiments on a set of benchmarks, including a RISC-V-based PULPino SoC and additional IP cores listed in Table II. For each benchmark, as shown, specific modules were selected as targets for eFPGA-based redaction using the OpenFPGA toolchain¹. *NuRedact* is implemented in Python and leverages the PyVerilog for (lexical) parsing and generation of non-uniform eFPGAs, all integrated into the OpenFPGA flow. Logic synthesis was performed using Cadence Genus 21.18 with the Skywater 130nm technology library to evaluate area overheads, on a server with an Intel Xeon E5-2640 v4 2.4 GHz processor (40 cores, 48 GB RAM). To assess the security of the redacted fabrics, netlists compatible with IcySAT and KC2 SAT attack tools were generated using Synopsys DC on Nangate 15nm technology node. Attacks were launched on the resulting non-uniform eFPGA fabrics using an Ubuntu environment on an Intel® Core™ i7-14700 system (28 cores, 32 GB RAM).

A. Hardware Overhead Analysis

The primary goal of *NuRedact* framework is to minimize eFPGA fabric area while preserving functionality and robustness against attacks. For each benchmark, both uniform (baseline) and non-uniform fabrics are generated by the *NuRedact* flow. Fabric layouts were obtained with VPR, leveraging a custom VPR architecture file together with the BLIF representation of

TABLE III: Area Efficiency of eFPGA Fabrics Generated via the *NuRedact* Framework Compared to Traditional Uniform Architectures.

IP	Module	Fabric Size	# Total Tiles	# Used Tiles	Area(μm^2)		ΔArea (%)
					Uniform	Non-Uni	
GPS	Cacode	3x3	9	7	513,361	357,251	-30.41
AES	AES_ShR	2x2	4	3	212,716	158,113	-25.67
Seq_Comb	Seq_Comb	3x3	9	5	560,845	257,653.5	-54.06
RISC-V	CTRL	3x3	9	5	529,113	306,957	-41.99
	Arbiter	6x6	36	5	2026,867	353,273.5	-82.57
	Ld/St	1x1	1	1	64,195	57,833	-9.91
Logic7	Logic7	3x3	9	7	513,361	357,251	-30.41
ADDER	FA_Array	1x1	1	1	64,195	57,833	-9.91
DES3	SBOX_8	1x1	1	1	64,195	57,833	-9.91

Negative values of ΔArea (%) indicate area reduction relative to the uniform fabric.

¹ Amongst the Benchmarks, smaller ranges are selected to visualize the compactness effectiveness of *NuRedact* over the fabrics.

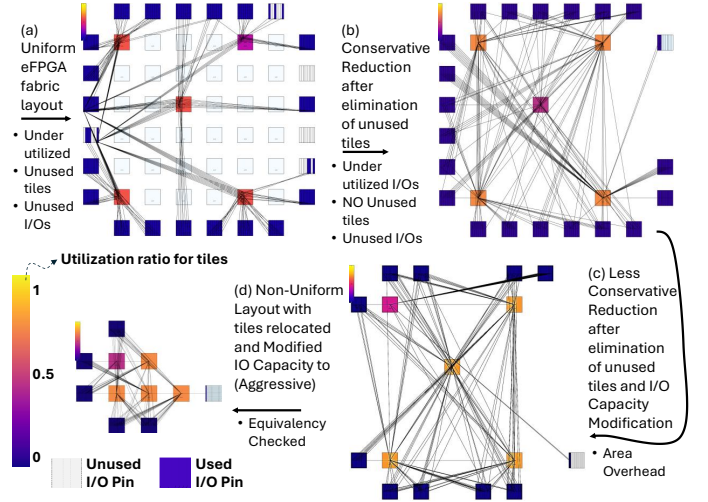


Fig. 3: Comparison of uniform and non-uniform eFPGA fabric layouts for Arbiter module of the RISC-V benchmark, generated using the *NuRedact* framework (using both Conservative and Aggressive).

the target modules. Figures 3 and 4 illustrate the transformation pipeline on two RISC-V modules, from uniform to compacted non-uniform, all extracted from OpenFPGA framework. For Arbiter (Figure 3(a)) is a 6x6 fabric, while after removing unused CLBs, a reduced non-uniform layout is obtained (Figure 3(b)), further refined by the conservative reduction (Figure 3(c)), where unused CLBs are pruned and the I/O capacity is adjusted. This is followed by compact layout generation (Figure 3(d)), where the I/O capacity is increased and the used tiles are relocated into a contiguous region to achieve a more compact fabric. This sequence shrinks the footprint of Arbiter from 6x6 tiles to a five-tile design ($\sim 86\%$ or 7.14x area reduction), while keeping the standard tool flow unchanged.

The CTRL shows the same trend (Figure 4), albeit with a smaller compaction ratio due to a more fit logic footprint. Heatmaps by OpenFPGA confirm that *NuRedact* increases per-tile utilization: for the Arbiter, the average tile utilization rises from 0.68 (uniform) to 0.94 (aggressive compaction). Higher utilization yields fewer idle configuration bits and a denser configuration entropy per unit area, which reduces structural cues and limits black-box learnability, while preserving CAD legality and compatibility with OpenFPGA/VTR.

We synthesized both the uniform and *NuRedact* non-uniform fabrics with Cadence Genus under identical libraries and timing constraints to compare silicon area. Table III summarizes results across all redacted modules. As an instance, for Seq_Comb (3x3 fabric with 5 used tiles), *NuRedact* cuts area by 54.06% (2.18x) relative to the uniform baseline, and for the RISC-V Arbiter (baseline 6x6), the reduction reaches 82.57%² (5.73x), highlighting the efficiency of tailoring the fabric footprint to actual utilization.

Table IV compares the two non-uniform modes, conservative reduction vs. layout compaction, while sweeping the per-tile I/O capacity (e.g., 4, 8, 16, 32; and 48 for Arbiter). To

² Per tile area reduction estimated from OpenFPGA layout counts is $\approx 86\%$, whereas post-synthesis area from Cadence Genus reports 82.57%; the close agreement validates the per-tile comparison as a quick proxy for silicon area.

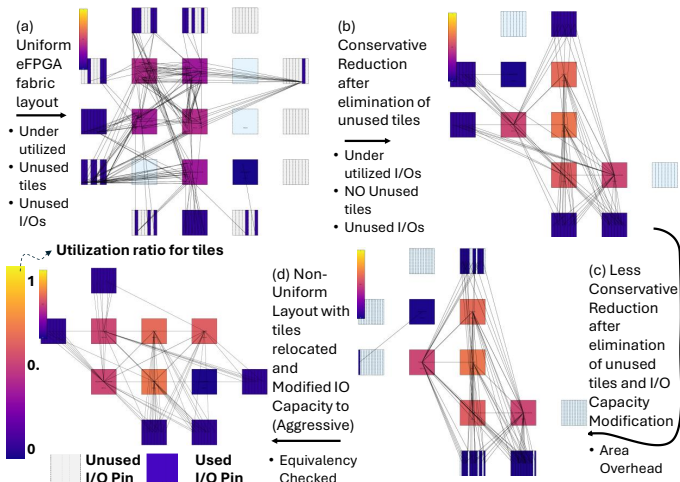


Fig. 4: Comparison of uniform and non-uniform eFPGA fabric layouts for CTRL module of the RISC-V benchmark, generated using the *NuRedact* framework (using both Conservative and Aggressive)

preserve routability, we kept at least one I/O tile on each fabric edge even when demand fell below four. The general trend is consistent: as per-tile I/O capacity increases, *NuRedact*'s layout compaction yields larger area savings over conservative reduction. For example, for larger modules like *Arbiter*, with higher I/Os, aggressive mode achieves 37.88% area reduction for compaction vs. conservative (Total of 89.17% (9.23x reduction) vs. uniform. Across all benchmarks, layout compaction produces the most compact fabrics (smallest area), while conservative reduction provides a lower-risk fallback when timing/routability are tight.

B. De-obfuscation Analysis

For the security analysis, we redacted each target module onto two fabrics: (i) the standard SkyWater-130 OpenFPGA architecture, and (ii) the non-uniform VPR architecture generated by *NuRedact*. Both fabrics use scan-chain configuration FFs and fracturable 4-input LUTs. Similar to prior state-of-the-art eFPGA-based (and other reconfigurable-based) studies [17]–[19], [31], our security analysis of *NuRedact* will focus exclusively on derived SAT attacks. This is because eFPGA-based redaction is inherently resilient against other categories

TABLE IV: Area Optimization in *NuArch*: Conservative vs. Aggressive.

Module	#IO/Tile	Conservative		Aggressive		Δ Area (%)
		#I/O Tile	Area (μm^2)	#I/O Required (Allocated)	Area (μm^2)	
Seq_Comb	4	6	264,708	6	255,618	-3.43
	8 (Default)	6	270,975	3(4)	257,653.5	-4.92
	16	6	280,724	2(4)	264,199	-5.89
CTRL	4	10	315,753.5	10	307,256.5	-2.69
	8 (Default)	7	319,507	5	306,957	-3.93
	16	7	331,092.5	2(4)	311,303	-5.98
Arbiter	8 (Default)	21	384,266	21	353,273.5	-8.06
	16	11	381,784.5	11	321,403	-15.82
	32	6	410,700	6	325,971	-20.63
	48	6	527,432.5	4	327,645.5	-37.88
Logic7	4	10	363,347.5	6	355,217.5	-2.24
	8 (Default)	10	373,705.5	4	357,251	-4.40
	16	10	389,769.5	2(4)	363,775.55	-6.67

Negative values of Δ Area (%) indicate area reduction.

TABLE V: Results of SAT-Attack on Non-uniform eFPGA Fabrics Generated via the *NuRedact* Framework.

Design	Bitstream	Unroll Factor	# Clauses	# Variables	Time (s)	Key (S/F)
HA_ARRAY	915	98	6165189	2318300	8102.2	S
CTRL	8215	434	-	-	TO	F
FA_ARRAY	937	254	19595603	7327291	40818.5	S
Cacode	9582	534	-	-	TO	F
Ld/St	937	254	-	-	TO	F
Comparator	861	303	-	-	TO	F
AES_ShR	2252	813	-	-	TO	F
Arbiter	8278	-	-	-	TO	F
SBOX_8	2001	-	-	-	TO	F
Logic7	4787	-	-	-	TO	F

of de-obfuscation attacks, such as structural- or ML-based approaches [31], [33]. To enable SAT-style analysis (for both cyclic and sequential), we convert the redacted fabrics to .bench by exposing all configuration FFs as primary key inputs. The configuration FFs are stitched as a scan chain driven by `prog_clk` (we recover the correct bit ordering via a depth-first walk from `scan_in_head` to `scan_in_tail`), then validate the mapping by re-simulation against the OpenFPGA bitstream. For sequential/cyclic behavior, we use *IcySAT* to unroll hard combinational loops and instantiate an oracle by fixing the key inputs to the ground-truth configuration from the OpenFPGA bitstream. We then launch *KC2* on the key-exposed netlist to assess functional recovery. (Time/memory budgets are fixed across all runs; uniform vs. non-uniform are evaluated under identical conditions.)

Table V summarizes SAT outcomes on *NuRedacted* fabrics. For a small design (2-bit adder, bitstream size 915), *IcySAT* unrolled 98 loops and *KC2* recovered the key in 2+ hours. Compared to its uniform counterpart, the non-uniform fabrics introduced $\geq 10\times$ SAT variables on *NuRedacted* modules, substantially increasing solve time, which is due to reducing structural cues and black-box learnability. For larger designs and bitstreams (e.g., the RISC-V CTRL and GPS Cacode), *IcySAT* completed unrolling after a long runtime, but *KC2* did not recover a key within the available RAM and time budgets.

VII. CONCLUSION

This paper introduced *NuRedact*, an automated eFPGA-based redaction flow that replaces security-critical logic with design-specific, non-uniform eFPGA fabrics. *NuRedact* collapses eFPGA under-utilization by injecting pin-map asymmetry, selective logic depopulation, local routing, and compact fixed layout, while remaining CAD-legal compatible with OpenFPGA/VTR and standard sign-off. On RISC-V and other benchmarks, *NuRedact* delivers substantial efficiency; up to 9x area reduction, while per-tile utilization improves by 38% on average, which helps increase configuration-entropy density per unit area for reducing structural cues and black-box learnability. In (cyclic and sequential) attack evaluations, very small fabrics remain solvable, but non-uniform fabrics with ≥ 5 logic tiles provide practical resistance against de-obfuscation attacks (variables to resolve is increased by 10x).

REFERENCES

- [1] M. Rostami, F. Koushanfar, and R. Karri, "A primer on hardware security: Models, methods, and metrics," *Proceedings of the IEEE*, vol. 102, no. 8, pp. 1283–1295, 2014.
- [2] H. M. Kamali, K. Z. Azar, F. Farahmandi, and M. Tehranipoor, "Advances in logic locking: Past, present, and prospects," *Cryptology ePrint Archive*, 2022.
- [3] P. Subramanyan, S. Ray, and S. Malik, "Evaluating the security of logic encryption algorithms," in *2015 IEEE International Symposium on Hardware Oriented Security and Trust (HOST)*. IEEE, 2015, pp. 137–143.
- [4] K. Z. Azar, H. M. Kamali, H. Homayoun, and A. Sasan, "Smt attack: Next generation attack on obfuscated circuits with capabilities and performance beyond the sat attacks," *IACR Transactions on Cryptographic Hardware and Embedded Systems*, pp. 97–122, 2019.
- [5] A. Barengi, L. Breveglieri, I. Koren, and D. Naccache, "Fault injection attacks on cryptographic devices: Theory, practice, and countermeasures," *Proceedings of the IEEE*, vol. 100, no. 11, pp. 3056–3076, 2012.
- [6] P. Kocher, J. Jaffe, and B. Jun, "Differential power analysis," in *Annual international cryptology conference*. Springer, 1999, pp. 388–397.
- [7] H. M. Kamali, K. Z. Azar, K. Gaj, H. Homayoun, and A. Sasan, "Lut-lock: A novel lut-based logic obfuscation for fpga-bitstream and asic-hardware protection," in *2018 IEEE Computer Society Annual Symposium on VLSI (ISVLSI)*. IEEE, 2018, pp. 405–410.
- [8] H. M. Kamali, K. Z. Azar, H. Homayoun, and A. Sasan, "Full-lock: Hard distributions of sat instances for obfuscating circuits using fully configurable logic and routing blocks," in *Proceedings of the 56th Annual Design Automation Conference 2019*, 2019, pp. 1–6.
- [9] L. Collini, B. Tan, C. Pilato, and R. Karri, "Reconfigurable logic for hardware ip protection: Opportunities and challenges," in *Proceedings of the 41st IEEE/ACM International Conference on Computer-Aided Design*, 2022, pp. 1–7.
- [10] H. M. Kamali, K. Z. Azar, F. Farahmandi, and M. Tehranipoor, "Shell: Shrinking efpga fabrics for logic locking," in *2023 Design, Automation & Test in Europe Conf. & Exhibition (DATE)*. IEEE, 2023, pp. 1–6.
- [11] Z. U. Abideen, S. Gokulanathan, M. J. Aljafar, and S. Pagliarini, "An overview of fpga-inspired obfuscation techniques," *ACM Computing Surveys*, vol. 56, no. 12, pp. 1–35, 2024.
- [12] V. Das, K. Azar, and H. Kamali, "Soar: Secure once, adapt at runtime with efpga-based redaction for ip protection," *IEEE Access*, 2025.
- [13] R. Guo, M. S. Rahman, H. M. Kamali, F. Rahman, F. Farahmandi, and M. Tehranipoor, "Evolute: evaluation of look-up-table-based fine-grained ip redaction," in *2023 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. IEEE, 2023, pp. 1–6.
- [14] J. Bhandari, A. K. T. Moosa, B. Tan, C. Pilato, G. Gore, X. Tang, S. Temple, P.-E. Gaillardon, and R. Karri, "Exploring efpga-based redaction for ip protection," in *2021 IEEE/ACM International Conference On Computer Aided Design (ICCAD)*. IEEE, 2021, pp. 1–9.
- [15] C. M. Tomajoli, L. Collini, J. Bhandari, A. K. T. Moosa, B. Tan, X. Tang, P.-E. Gaillardon, R. Karri, and C. Pilato, "Alice: An automatic design flow for efpga redaction," in *Proceedings of the 59th ACM/IEEE Design Automation Conference*, 2022, pp. 781–786.
- [16] J. Bhandari, A. K. T. Moosa, B. Tan, C. Pilato, G. Gore, X. Tang, S. Temple, P.-E. Gaillardon, and R. Karri, "Not all fabrics are created equal: Exploring efpga parameters for ip redaction," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 31, no. 10, pp. 1459–1471, 2023.
- [17] A. Dasgupta, S. Paria, and S. Bhunia, "Hipr: Hardware ip protection through low-overhead fine-grain redaction," *IACR Transactions on Cryptographic Hardware and Embedded Systems*, vol. 2025, no. 3, pp. 781–805, 2025.
- [18] L. Collini, J. Bhandari, C. Muscari Tomajoli, A. Moosa, B. Tan, X. Tang, P.-E. Gaillardon, R. Karri, and C. Pilato, "Arianna: An automatic design flow for fabric customization and efpga redaction," *ACM Transactions on Design Automation of Electronic Systems*, 2025.
- [19] A. Fowler, S. Mohammed, M. Shihab, T. Broadfoot, P. Beerel, C. Sechen, and Y. Makris, "A trap for sat: On the imperviousness of a transistor-level programmable fabric to satisfiability-based attacks," *IACR Transactions on Cryptographic Hardware and Embedded Systems*, vol. 2025, no. 2, pp. 579–603, 2025.
- [20] M. El Massad, S. Garg, and M. V. Tripunitara, "The sat attack on ic camouflaging: Impact and potential countermeasures," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 39, no. 8, pp. 1577–1590, 2019.
- [21] K. Shamsi, M. Li, T. Meade, Z. Zhao, D. Z. Pan, and Y. Jin, "Appsat: Approximately deobfuscating integrated circuits," in *2017 IEEE International Symposium on Hardware Oriented Security and Trust (HOST)*. IEEE, 2017, pp. 95–100.
- [22] Y. Shen and H. Zhou, "Double dip: Re-evaluating security of logic encryption algorithms," in *Proceedings of the Great Lakes Symposium on VLSI 2017*, 2017, pp. 179–184.
- [23] H. Zhou, R. Jiang, and S. Kong, "Cycsat: Sat-based attack on cyclic logic encryptions," in *2017 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*. IEEE, 2017, pp. 49–56.
- [24] K. Shamsi, D. Z. Pan, and Y. Jin, "Icysat: Improved sat-based attacks on cyclic locked circuits," in *2019 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*. IEEE, 2019, pp. 1–7.
- [25] X. Tang, E. Giacomini, A. Alacchi, B. Chauviere, and P.-E. Gaillardon, "Openfpga: An opensource framework enabling rapid prototyping of customizable fpgas," in *2019 29th Int'l Conference on Field Programmable Logic and Applications (FPL)*. IEEE, 2019, pp. 367–374.
- [26] A. Baumgarten, A. Tyagi, and J. Zambreno, "Preventing ic piracy using reconfigurable logic barriers," *IEEE design & Test of computers*, vol. 27, no. 1, pp. 66–75, 2010.
- [27] G. Kolhe, S. M. PD, S. Rafatirad, H. Mahmoodi, A. Sasan, and H. Homayoun, "On custom lut-based obfuscation," in *Proceedings of the 2019 Great Lakes Symposium on VLSI*, 2019, pp. 477–482.
- [28] H. M. Kamali, K. Z. Azar, H. Homayoun, and A. Sasan, "Interlock: An intercorrelated logic and routing locking," in *Proceedings of the 39th International Conference on Computer-Aided Design*, 2020, pp. 1–9.
- [29] L. Alrahis, S. Patnaik, M. A. Hanif, M. Shafique, and O. Sinanoglu, "Untangle: Unlocking routing and logic obfuscation using graph neural networks-based link prediction," in *2021 IEEE/ACM Int'l Conference On Computer Aided Design (ICCAD)*. IEEE, 2021, pp. 1–9.
- [30] Z. Han, M. Shayan, A. Dixit, M. Shihab, Y. Makris, and J. Rajendran, "{FuncTeller}: How well does {eFPGA} hide functionality?" in *32nd USENIX Security Symposium*, 2023, pp. 5809–5826.
- [31] P. Karmakar, M. Bharani, and C. Karfa, "Evaluating the robustness of large scale efpga-based hardware redaction," in *2024 37th International Conference on VLSI Design and 2024 23rd International Conference on Embedded Systems (VLSID)*. IEEE, 2024, pp. 517–522.
- [32] C. G. Sathe, Y. Makris, and B. C. Schafer, "Mantis: Machine learning-based approximate modeling of redacted integrated circuits," in *2023 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. IEEE, 2023, pp. 1–6.
- [33] A. Rezaei, R. Afsharmazayejani, and J. Maynard, "Evaluating the security of efpga-based redaction algorithms," in *Proceedings of the 41st IEEE/ACM Int'l Conf. on Computer-Aided Design*, 2022, pp. 1–7.
- [34] K. Shamsi, M. Li, D. Z. Pan, and Y. Jin, "Cross-lock: Dense layout-level interconnect locking using cross-bar architectures," in *Proceedings of the 2018 Great Lakes Symposium on VLSI*, 2018, pp. 147–152.
- [35] J. Luu, J. Goeders, M. Wainberg, A. Somerville, T. Yu, K. Nasartschuk, M. Nasr, S. Wang, T. Liu, N. Ahmed *et al.*, "Vtr 7.0: Next generation architecture and cad system for fpgas," *ACM Transactions on Reconfigurable Technology and Systems (TRETS)*, vol. 7, no. 2, pp. 1–30, 2014.
- [36] L. McMurchie and C. Ebeling, "Pathfinder: A negotiation-based performance-driven router for fpgas," in *Proceedings of the 1995 ACM third international symposium on Field-programmable gate arrays*, 1995, pp. 111–117.
- [37] J. Sweeney, M. J. Heule, and L. Pileggi, "Modeling techniques for logic locking," in *Proceedings of the 39th International Conference on Computer-Aided Design*, 2020, pp. 1–9.
- [38] K. Shamsi, M. Li, D. Z. Pan, and Y. Jin, "Kc2: Key-condition crunching for fast sequential circuit deobfuscation," in *2019 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. IEEE, 2019, pp. 534–539.