ARL-TR-9652 ● FEB 2023

DEVCOM
ARMY RESEARCH
LABORATORY

Revisiting Obfuscation Metrics for Trusted Fabrication of Designs

by Theodros Nigussie

NOTICES

Disclaimers

The findings in this report are not to be construed as an official Department of the Army position unless so designated by other authorized documents.

Citation of manufacturer's or trade names does not constitute an official endorsement or approval of the use thereof.

Destroy this report when it is no longer needed. Do not return it to the originator.

ARL-TR-9652 ● Feb 2023

Revisiting Obfuscation Metrics for Trusted Fabrication of Designs

Theodros Nigussie
DEVCOM Army Research Laboratory

REPORT DOCUMENTATION PAGE

1. REPORT DATE	**2. REPORT TYPE**	**3. DATES COVERED**	
		START DATE	**END DATE**
February 2023	Technical Report	July 2022	January 2023

4. TITLE AND SUBTITLE
Revisiting Obfuscation Metrics for Trusted Fabrication of Designs

5a. CONTRACT NUMBER	**5b. GRANT NUMBER**	**5c. PROGRAM ELEMENT NUMBER**
5d. PROJECT NUMBER	**5e. TASK NUMBER**	**5f. WORK UNIT NUMBER**

6. AUTHOR(S)
Theodros Nigussie

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) DEVCOM Army Research Laboratory ATTN: FCDD-RLA-PE Adelphi, MD 20783-1138	**8. PERFORMING ORGANIZATION REPORT NUMBER** ARL-TR-9652

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	**10. SPONSOR/MONITOR'S ACRONYM(S)**	**11. SPONSOR/MONITOR'S REPORT NUMBER(S)**

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release: distribution unlimited.

13. SUPPLEMENTARY NOTES
ORCID ID: Theodros Nigussie, 0000-0003-0096-5059

14. ABSTRACT

There are several published papers on obfuscation techniques to prevent reverse engineering of an integrated circuit design. Among these, the split fabrication approach is the most promising to provide strong obfuscation. In this report, proposed improvements and additional analyses are included for two of the metrics that can be used to measure the level of obfuscation for the split fabrication approach. First, the partitioning depth metric provides guidance to the partitioning tool with respect to the extent of partitioning needed to sufficiently hide the design. Second, the connection possibility metric provides a statistical analysis of the potential to reconstruct the original design from a partitioned design. A detail is revealed that shows more precisely how difficult reverse engineering will be based on this metric.

15. SUBJECT TERMS
design obfuscation; trusted fabrication; obfuscation metrics; partitioning; Photonics, Electronics, and Quantum Sciences

16. SECURITY CLASSIFICATION OF:			**17. LIMITATION OF ABSTRACT**	**18. NUMBER OF PAGES**
a. REPORT	**b. ABSTRACT**	**C. THIS PAGE**		
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	UU	23

19a. NAME OF RESPONSIBLE PERSON	**19b. PHONE NUMBER** *(Include area code)*
Theodros Nigussie	(802) 310-4133

STANDARD FORM 298 (REV. 5/2020)
Prescribed by ANSI Std. Z39.18

Contents

List of Figures

List of Tables

1. Introduction

Today, more than 90% of the world's semiconductor foundry production is controlled by non-US companies that are located overseas; the share of the US-based companies is diminishing from year to year.[1] Also, the foreign foundries are demonstrating faster transition of leading-edge technology nodes to commercial production.[2] Such semiconductor industry trends present a challenge to the traditional threat prevention model in which DOD relies on domestically owned and operated trusted suppliers for the acquisition of integrated circuits (ICs).[3] However, the current trusted model does not allow access to the latest nodes and restricts DOD access to technologies that are at least two generations behind what is available in the advanced (untrusted) foundries. The advanced foundries offer chips with 10 times reduction in chip size and 5 times better energy efficiency than what is provided by the current trusted partnership. Thus, it is imperative to have an alternative framework that allows the DOD secure access to the most advanced commercial processes. Such a framework involves implementation of a split-manufacturing technique where a portion of the design, which includes the critical intellectual property and the integration of the whole, is fabricated in a trusted environment for security while the rest of the design is fabricated at more advanced untrusted foundries for optimal power–performance–area (PPA). Comprehensive summaries of the research state on protecting designs using split fabrication techniques[4–8] and the proposed obfuscation metrics[7–11] for measuring the level of security provided by the techniques are found in the literature.

This report presents a brief survey of a few of the representative obfuscation metrics and analyses and proposed improvements for the partitioning depth (P_{depth}) and connection possibility (C_p) metrics discussed in the literature.[8]

2. Brief Survey of Obfuscation Metrics

Several metrics that are used to evaluate the security provided by obfuscation techniques have been discussed in the literature. A brief summary of these metrics follows.

2.1 Hamming Distance (HD)

The HD metric performs a bit-by-bit comparison of two bitstreams: 1) one from the output of the unobfuscated design and 2) another from the output of the obfuscated version of the same design, which uses a percentage figure to describe how many bits are different between the bitstreams.[7] Each bitstream stands for the output of a circuit and an average HD of 50% implies that the responses between the two circuits are completely different. The metric has been used for evaluating gate-level obfuscation techniques such as logic encryption and camouflaging.[9] As this metric requires logic simulation, it will not scale well with an increase in the number of inputs and size of circuit.[7]

2.2 Verification Failure

The verification failure metric uses formal verification tools such as Synopsys Formality to perform logical equivalence checking between the obfuscated design and the original.[7] The equivalence checking is done on the logic cones of the two designs to compare their output ports and flip-flop outputs. The metric is expressed as a percentage of compare points that failed equivalence checking to the total number of compare points. This metric was used to evaluate security of obfuscation techniques discussed in Chakraborty et al.[10] Compared to the simulation-based HD approach, this metric is much faster and does not suffer coverage issues.[7]

2.3 Entropy

Entropy refers to the amount of information an adversary can extract by observing the obfuscated version of the design based on the distribution of gate types.[7] Jagasivamani et al.[11] proposed this metric advocating for the synthesis of a design to a larger number of gate types to minimize the entropy as opposed to a synthesis to one or two types of gates that will make it easier for the attacker. They also proposed a complementary metric that is termed standard cell composition bias.[11] The metric analyzes the proportion of standard cells in the design. The idea is that a design with high bias (i.e., the dominant presence of XORs as compared to other gates) might indicate a cryptographic core revealing useful information about the design. Therefore, the design should be synthesized with low bias (i.e., with equal proportion of different types of standard cells) so that it will not be easier for the attacker to make a generalized guess about the functionality of the design.

2.4 Partitioning Depth (P_{depth})

This obfuscation metric is based on the success rate of the partitioning tool in disguising the connection of cells so that reconstruction of the netlist will not be easy.[8] The analysis of this metric starts with identifying the logic paths between input ports and an output port that form a logic cone. P_{depth} for a logic path is calculated by dividing the number of partitioned nets in a logic path by the total number of nets in the path. A logic cone can have multiple logic paths that may stretch from multiple inputs to a single output port. The P_{depth} value for a logic cone is calculated by averaging the P_{depth} values of the logic paths that constitute the logic cone. The calculated P_{depth} value for the entire design is the average of the P_{depth} values of all of the logic cones and ranges between 0 and 1. The values approach 1 as the number of partitioned nets increases—implying a high level of obfuscation—whereas a value closer to 0 suggests a minimal number of cuts and a low level of obfuscation.

2.5 Connection Possibility (C_p)

The C_p metric was proposed by Nigussie et al.[8] and measures the security of an obfuscated design based on the number of possible connections that must be tested to find the correct connection between the ports of two design splits. This metric is especially applicable to an application specific integrated circuit (ASIC)-on-ASIC stacking approach using three-dimensional IC bonding technologies. The connection points between the design splits are hidden in an interposer where interconnected wires are routed to connect two different locations in the two stacked chips. Also, the routing from point, p1, on chip1 is only allowed to be made within a predefined radius to point, p2, on chip2 so that the impact to PPA is minimized. The C_p value can grow exponentially with the increase in the number of partitioned nets complicating the reverse engineering work of an attacker.

3. Analysis and Proposed Improvements

In this section, two of the metrics that are more relevant to the partitioning-based obfuscation are revisited for further analysis and improvements.

3.1 Partitioning Depth

The P_{depth} metric is based on averaging the number of cuts in logic paths that form a logic block.[8] In formal verification terms, these logic blocks are known as logic cones where the output port is considered a compare point for logical equivalence checks. A logic cone refers to a combinational logic originating from a compare point (e.g., primary outputs, internal registers, and black box input pins) fanning backward to another compare point (e.g., primary input ports, black box output pins, and register outputs). The adder circuit shown in Fig. 1 is a good example of a logic cone. A logic path refers to the path that starts from a compare point, in this case an input port and terminating at an output port by propagating through a cascade of gates.

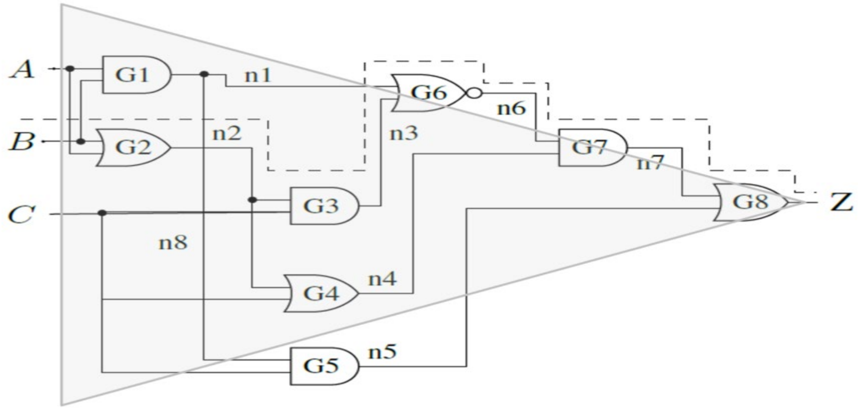

Fig. 1 Logic cone representing a full adder circuit. The dashed line shows the logic path from the input, *B*, to the output.

For example, there are 10 logic paths in the adder circuit of Fig. 1, the longest of which is *B*$_{\text{(in port)}}$→G2→G3→G6→G7→G8→Z$_{\text{(out port)}}$. Decomposing the logic cone into logic paths will aid automated computation of cuts in a given path and average it over the total number of paths. This can be done by converting the cascade of gates in the logic path into a bitstring. If a gate is assigned to a Tier1, it will be assigned a 0 value and a Tier2 gate will have a value of 1. Suppose G3 and G7 are assigned to Tier1 and the rest are assigned to Tier2, the bitstring value for the logic path will be 10101, and the number of cuts along the path can be calculated by counting the number of 0→1 and 1→0 transitions in the bitstring.

The calculated P$_{\text{depth}}$ value for the example circuits shown in Fig. 1 would range between 0 for the unpartitioned circuit and 1 if all nets n1–n7 are cut. This metric will yield a value between 0 and 1 for any partitioning less than full partition. Having discussed how the number of cuts is calculated, the following paragraphs examine how this approach can be used for a large netlist that has combinational gates, registers, and memory macros.

Figure 2 shows an example generic netlist that may have a combination of these gates. This generic netlist is a union of a large number of logic cones, registers, and black box units such as memory macros. A logic path in this case includes all the combinational gates, registers, and memory macros between input and output ports. One example follows: I$_{0\text{(in port)}}$ → CL$_1$ → R1 → CL$_2$ → BBOX → . . . → R$_n$ → CL$_n$→ O$_{0\text{(out port)}}$. The number of logic paths analyzed will increase with the size of the design netlist. Large design netlists can have millions of logic paths; therefore, it is imperative to consider only representative samples of the logic paths for these metric analyses to reduce run time. The samples are chosen based on a fixed number of cascaded standard cells or macros (N$_{\text{cascaded\_cells}}$) that form a logic path. For example, let us consider the logic cone in Fig. 1. Suppose the value for N$_{\text{cascaded\_cells}}$

is selected to be 5, then there are two logic paths that qualify the criteria: $A_{\text{(in port)}}$ →G2→G3→G6→G7→G8→Z $_{\text{(out port)}}$, and $B_{\text{(in port)}}$ →G2→G3→G6→G7→G8 →Z $_{\text{(out port)}}$. The P_{depth} value for the logic cone will be the average of the partitioning ratios of the two paths. In Section 3.1.1, analyses data are provided based on a study of partitioning different designs.

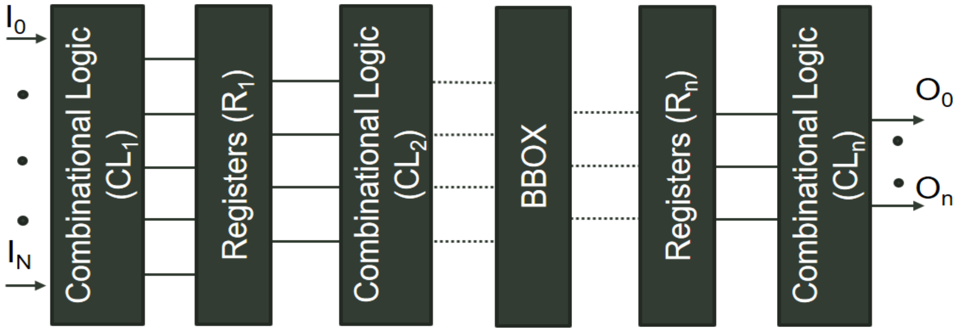
I_0
I_N
Combinational Logic (CL_1)
Registers (R_1)
Combinational Logic (CL_2)
BBOX
Registers (R_n)
Combinational Logic (CL_n)
O_0
O_n

Fig. 2 Generic netlist example

3.1.1 P_{depth} Analysis

The implementation and results of the P_{depth} metric are presented in this section. Eight designs of different sizes representing a range of application areas were analyzed. Also, the P_{depth} data was correlated with the mismatch rate obtained from comparing the original netlist with each of the partitioned netlists.

3.1.1.1 Designs

A summary of the designs and their characteristics is included in Table 1.

All designs except the field programmable neural array (FPNA) were synthesized using the open-source tool OpenROAD-Yosys[12] to the *nangate45* free process design kit. The FPNA was synthesized to GlobalFoundries' 12-nm library.

Table 1 Design characteristics

Design	Standard cells	Nets	Memory macros[a]	Description
GCD	365	436	0	Greatest common denominator
AES	16,478	17,273	0	Advanced encryption standards cryptography
IBEX	15,568	18,327	. . .	32-bit RISC-V core: 2-stage pipeline
JPEG	57,137	73,898	0	Image compression algorithm
TinyRocket	23,917	29,209	2: [64 × 32] I&D	RISC-V core with I&D caches
SweRV	85,457	100,354	. . .	RISC-V core (39 bit)
FPNA	1,507,459	1,843,156	64: [4096 × 16]	Field programmable neural array

BlackParrot	165,824	195,058	24: [1: 64 × 7;
			1:64 × 15;
			4:64 × 96;
			1:256 × 95;
			17:512 × 64]
				Full 64-bit RISC-V core with cache coherence

[a] Please note “2: [64 × 32]” refers to two memory macros with 64-word lines (rows) and 32-bit lines (columns) each.

3.1.1.2 Partitioners

The P_{depth} analysis was conducted for the designs listed in Table 1 across five different partitioners and six different $N_{cascaded\_cells}$ values. The partitioners are summarized in Table 2. The $N_{cascaded\_cells}$ values used in the analysis include 7, 9, 11, 14, 17, and 19. These values were selected so that a sufficient number of the logic paths are represented, and the extraction of the logic paths will complete within a reasonable time. The extraction time will grow exponentially as the value of $N_{cascaded\_cells}$ increases.

Table 2 Partitioners used in the analysis

Partitioner	Algorithm
MLPART[13]	Minimum cut
hMETIS[14]	Minimum cut
Obf1[8]	Assigning driver and receiver cells in separate tiers
Obf2[8]	Assign every other cell randomly
Obf3[8]	Assign every fourth cell randomly

Also, logic-equivalence check was done using the Synopsys Formality tool between the original netlist and the design splits for each partitioner. Therefore, a total of 16 verification checks (2 per design) were completed. The calculated failure (mismatch) rate along with the P_{depth} value are presented in Section 3.1.2.

3.1.2 Results and Discussion

3.1.2.1 Number of Logic Paths Versus the Number of Cascaded Cells

Figure 3 summarizes the number of cascaded cells that were selected for the analysis and the associated number of logic paths used in calculating the P_{depth} metric for each design.

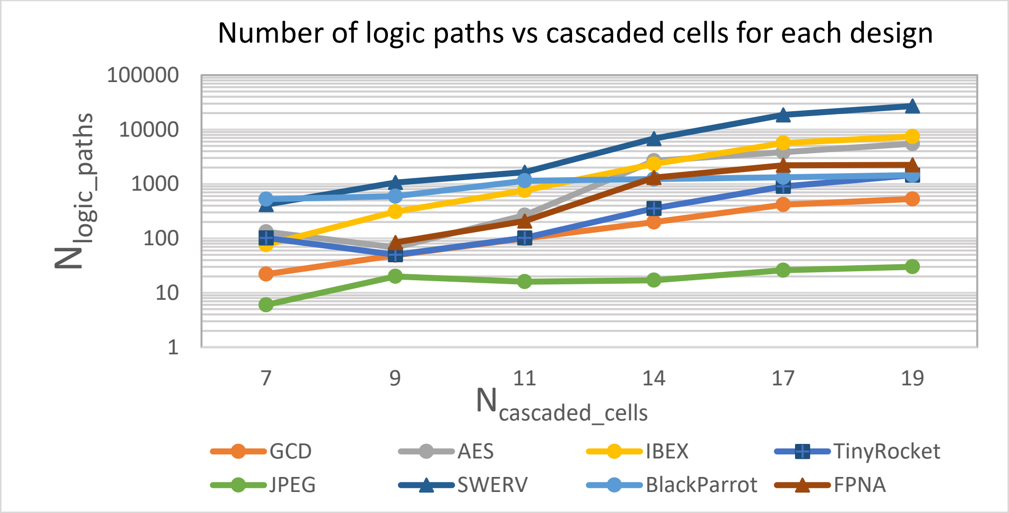

Fig. 3 Selected number of cascaded cells and the associated number of logic paths for each design

A data point is missing for $N_{cascaded\_cells} = 7$ for the FPNA design; the minimum number of cascaded cells that connect an input port to an output port is greater than seven.

3.1.2.2 Mismatch Rate Versus P_{depth} Grouped by $N_{cascaded\_cells}$

The effectiveness of the P_{depth} metric can be measured by correlating it with the logical equivalency test between the obfuscated and original designs. The tests for each of the designs used Synopsis Formality. The mismatch rate between identical netlists will be 0. The mismatch between a fully partitioned netlist and an unpartitioned netlist will be 1. The P_{depth} data collected for the designs were combined and are shown in Fig. 4 for each of the $N_{cascaded\_cells}$. Each chart is broken into subgroups by the partitioned tiers.

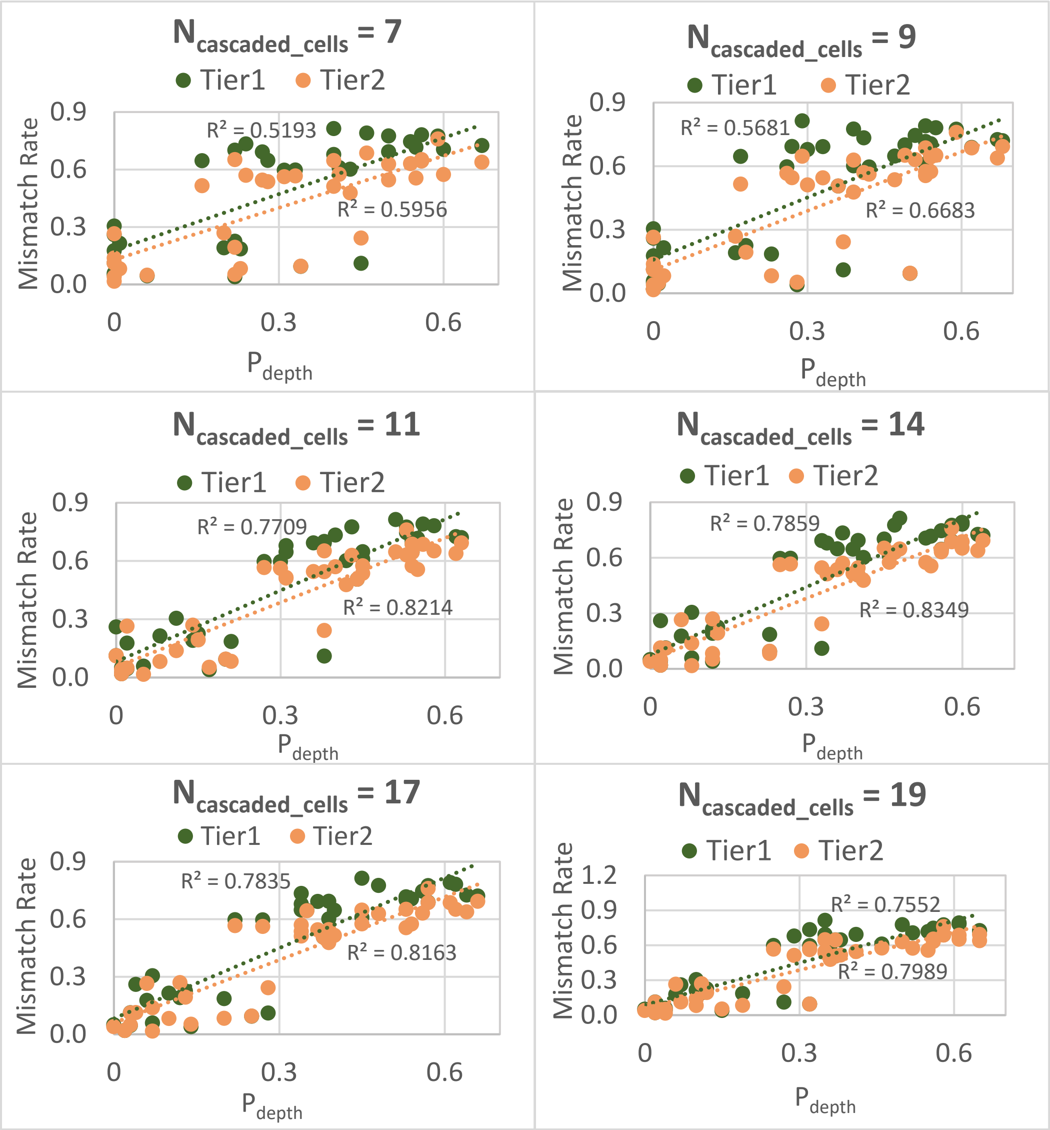

Fig. 4 Mismatch rate vs. P_{depth} using combined data for all designs

The regression lines for Tier1 and Tier2 are different because the partitioning flow assigns cells connected to primary input/output (I/O) ports to Tier2 only. This results in Tier2 having a lower number of mismatches with respect to the reference design. Among the six charts, the one for $N_{cascaded\_cells}$ = 14 shows the best correlation, which suggests the logic paths represented by this value can be sampled for P_{depth} calculation.

3.1.2.3 Mismatch Rate Versus P_{depth} Subgrouped by Design Name

The chart in Fig. 5 shows the correlation of mismatch rate and P_{depth} for each design. The $N_{cascaded\_cells}$ value with the best correlation was chosen to show how the trend for each design looks. There are three regression lines shown in the chart. The black line is for BlackParrot, which shows the least correlation among the designs. This is attributed to the large number of cells connected to primary I/O ports that are assigned to only one of the tiers. The brown line is for the FPNA design, which shows the best correlation. This is attributed to the fact that the number of cells connected to I/O ports is minimal compared to the size of the design. The blue line that stands for the combined data of the designs is added as a reference.

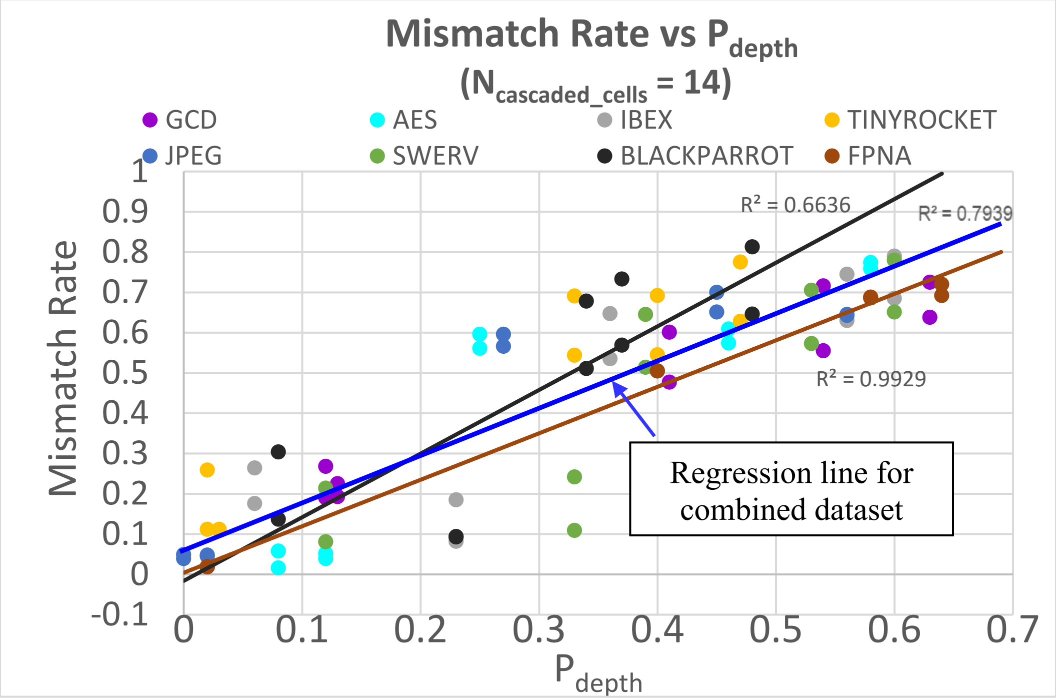

Fig. 5 Mismatch rate vs. P_{depth} for each design; the blue line is for all designs (combined)

3.1.3 Proposed P_{depth} Values for Obfuscation Efforts

Although partitioning all the nets in a netlist maximizes obfuscation, the overhead PPA is untenable. Therefore, a lower P_{depth} value is desirable. The P_{depth} to mismatch curves consistently have a significant increase near a value of 0.3. Choosing a P_{depth} value of 0.41 will result in an approximately 50% mismatch rate, which is proposed as an initial target for obfuscation (Fig. 6). Partitioning that results in a 50% mismatch rate is expected to result in sufficiently broken logic cones to prevent discovery of the original design. The precise P_{depth} to choose will depend on the impact to PPA.

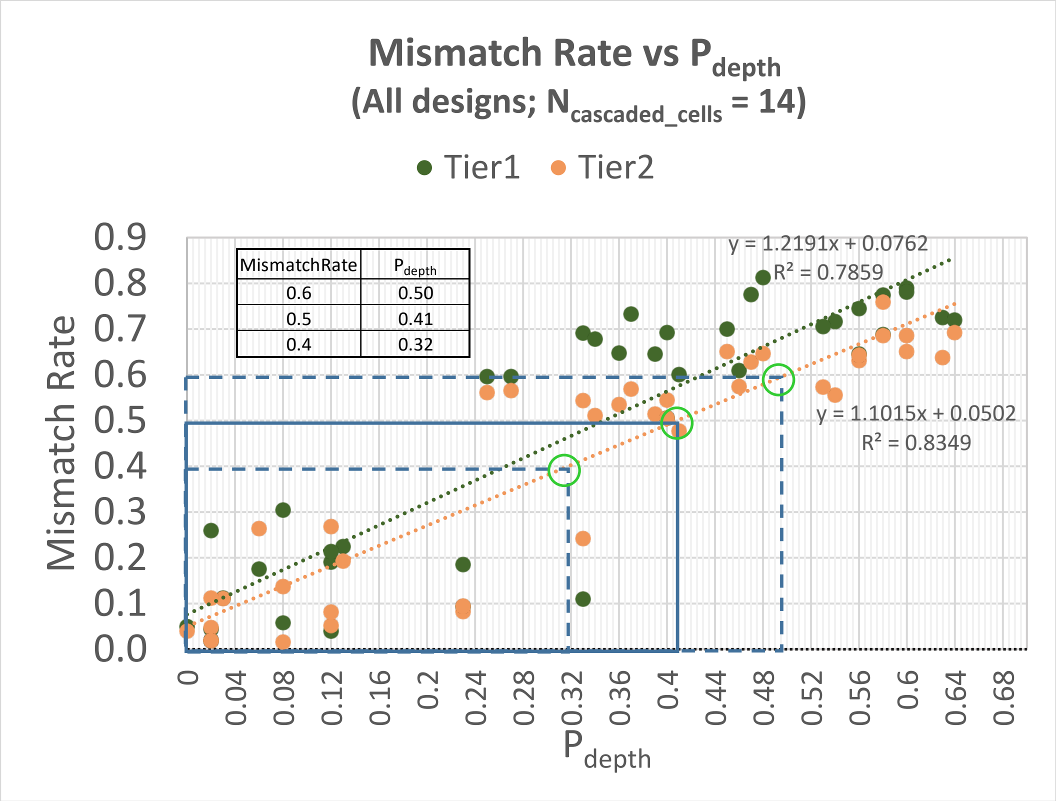

MismatchRate	P_{depth}
0.6	0.50
0.5	0.41
0.4	0.32

Fig. 6 Mismatch rate vs. P_{depth} by design split for $N_{cascaded\_cells}$ = 14

3.2 Connection Possibility

The C_p is illustrated in Fig. 7. It is based on the idea that a potential adversary having access to the design of only one partition will be able to identify the I/O ports between the two partitions depending on how they are connected to the cells in the design layout. If this is discovered, the split manufacturing obfuscation would be largely overcome. This metric provides a statistical analysis of the likelihood of this being discovered. The ports are divided into I/O groups so that an output port in a tier is connected to an input port in the opposite tier.

Suppose N stands for the number of partitioned nets to be vertically connected between the tiers and M is the number of output ports (B_1, B_2, . . . B_M) in Tier1 that are connected to a group of input ports (T_1, T_2, . . .T_M) in Tier2. Similarly, let us consider T_{M+1}, T_{M+2}, . . . T_N to be output ports of Tier2 that are driving the input ports B_{M+1}, B_{M+2}, . . . B_N in Tier1.

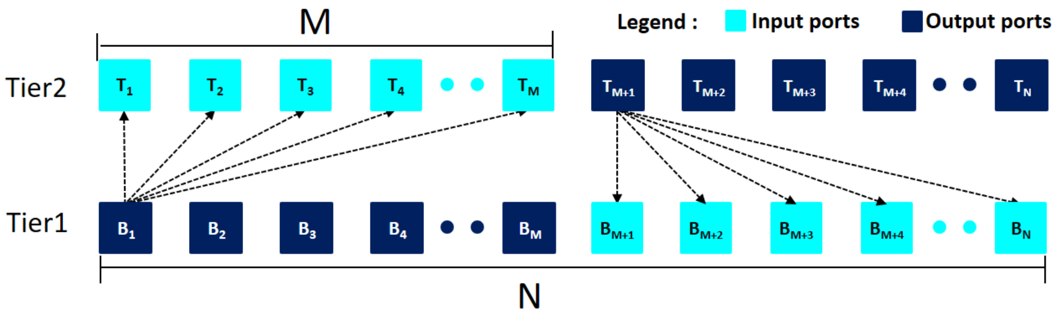

Fig. 7 Possible vertical connections between output ports in Tier1 and input ports in Tier2 and vice versa. Only connections from B_1 and from T_{M+1} are shown for clarity.

The possible connection that can be made between the Tier1 output ports and Tier2 input ports is M!. Similarly, Tier2 output ports can be linked to Tier1 ports in (N-M)! ways. Thus, the total C_P between the tiers will be expressed as

$$Cp = M! + (N - M)! \tag{1}$$

Adversaries will not stop at this. They will work to estimate the worst-case speed of the design by conducting critical path analysis on the unpartitioned portion of the netlist in order to narrow down the connection possibilities. Therefore, the assertion in Fig. 7 where any output port in Tier1 can drive any input port in Tier2 should be reassessed.

Figure 8 shows a subset of the ports included in Fig. 8 in the ascending order of propagation delay, which is a portion of the delay in the timing path up to/from the partitioning point.

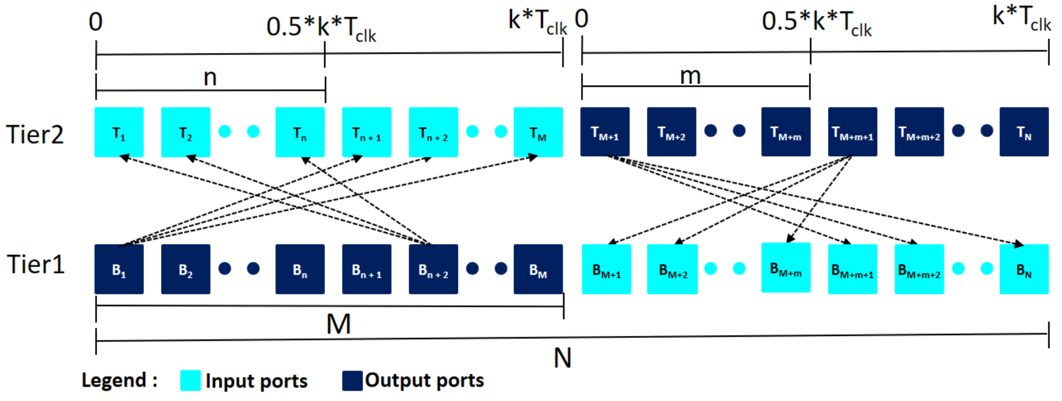

Fig. 8 Subset of the vertical connection points listed in ascending order of propagation delay

Suppose n and m are the number of output ports of Tier1 and Tier2, respectively, that have propagation delay of less than or equal to 0.5 × k × T_{clk} where T_{clk} is the worst-case speed or in other terms the value of the critical path delay in the unpartitioned timing paths and k is a T_{clk} factor. The output ports in Tier1 that include B_1, B_2, . . . B_n, can only be connected to Tier2's input ports: T_{n+1}, T_{n+2}, . . .T_M. Similarly, output ports B_{n+1}, B_{n+2}, and B_M, will be restricted to connect to T_1, T_2, . . . T_n. This restriction is made because the sum of the propagation delays of the connected ports cannot exceed k × T_{clk}, which is the critical path delay for the entire design. Therefore, Eq. 1 can be rewritten as follows:

$$\mathrm{Cp} = n! + (\mathrm{M} - \mathrm{n})! + (\mathrm{N} - \mathrm{M} - \mathrm{m})! + \mathrm{m}! \quad (2)$$

The minimum value for C_p is obtained when M = N/2 and n = m = N/4. Thus, $C_{p(min)}$ can be formulated as follows:

$$\mathrm{Cp(min)} = 4 * \left(\frac{\mathrm{N}}{4}\right)! \quad (3)$$

Note that k will take a value of 1 if the critical path of the entire design is in the unpartitioned timing paths and a value greater than 1 if the critical path is in the partitioned timing paths. This suggests that it will add more guesswork for the potential attacker to determine the value of k if the obfuscation strategy involves putting the critical path in the partitioned timing paths.

Table 3 includes estimated evaluation time for the $C_{p(min)}$ value formulated in Eq. 3 as a function of N. The analysis assumes the adversary can deploy one billion of the current fastest GPUs[15] in parallel at 12 TFLOPS each. This is 11,000 times faster than the fastest supercomputer at the time of this report. As N is chosen by the designer, it can be made arbitrarily large to ensure that no classical computer will be able to decipher the I/O ports within a meaningful time frame.

Table 3 Estimated evaluation time

N	C_p	Estimated evaluation time
10	8	6.67E – 22 s
20	480	4E – 20 s
40	14,515,200	1.21E – 15 s
80	9.7316E + 18	0.811 ms
100	6.2E + 25	1.4 h
120	1E + 33	2799 years
140	4.13E + 40	1.09E + 11 years

4. Conclusion

In this report, two of the obfuscation metrics proposed to be used for split manufacturing-based obfuscation techniques have been revisited and improvements were proposed. The metrics are complementary and should be used together. The P_{depth} metric measures the level of partitioning conducted on a logic path connecting an input port with an output port and helps to show how well logic cones in a given design's netlist are broken down in such a way that the logic structure and Boolean expression of the design are significantly changed. The C_p metric shows that the possible connections to be tested will grow exponentially with the increase in the number of partitioned nets. Future work will include working with a Red Team to get feedback on the robustness of our obfuscation effort and determine how well the metrics track with practice. Also, a metric that will combine obfuscation and the associated impact to PPA will be formulated.

5. References

1. Statista. [accessed 2023 Jan 31]. https://www.statista.com/statistics/867223/worldwide-semiconductor-foundries-by-market-share/.

2. Clark P. TSMC ramps revenue per wafer, other foundries dip. European Business Press SA; 2020 Feb 20 [accessed 2023 Jan 31]. https://www.eenewsanalog.com/news/tsmc-ramps-revenue-wafer-other-foundries-dip.

3. Department of Defense (US). DOD announces $117 million defense production act title III agreement with GlobalFoundries to strengthen the domestic microelectronics industrial base. Defense Media Activity; 2022 May 2 [accessed 2023 Feb 1]. https://www.defense.gov/News/Releases/Release/Article/3016070/dod-announces-117-million-defense-production-act-title-iii-agreement-with-globa/.

4. Colombier B, Bossuet L. Survey of hardware protection of design data for integrated circuits and intellectual properties. Computers & Digital Techniques, IET. 2014 Nov;8(6):274–287.

5. Perez T, Pagliarini S. A survery on split manufacturing: attacks, defenses, and challenges. IEEE Access. 2020 Oct;8.

6. Bhunia S. Hardware security: a hands-on learning approach. Morgan Kaufmann, an imprint of Elsevier, 2019.

7. Forte D. Hardware protection through obfuscation. Cham: Springer, 2017.

8. Nigussie T, Schabel J, Lipa S, McIlrath L, Patti R, Franzon P. Design obfuscation through 3D split fabrication with smart partitioning. IEEE TVLSI Systems. 2022 June.

9. Rajendran J, Pino Y, Sinanoglu O, Karri R. Logic encryption: a fault analysis perspective. Proceedings of the Conference on Design, Automation and Test in Europe. EDA Consortium; 2012. p. 953–958.

10. Chakraborty RS, Bhunia S. Harpoon: an obfuscation-based SOC design methodology for hardware protection. IEEE Trans Computer Aided Design Integration Circuits System. 2009;28(10):1493–1502.

11. Jagasivamani M, Gadfort P, Sika M, Bajura M, Fritze M. Split-fabrication obfuscation: metrics and techniques. In: 2014 IEEE International Symposium on Hardware-Oriented Security and Trust (HOST); 2014. p. 7–12.

12. OpenROAD. GitHub, Inc [accessed 2023 Feb 10]. https://github.com/The-OpenROAD-Project/OpenROAD.

13. Caldwell AE, Kahang AB, Kennings AA, Markov IL. Hypergraph partitioning for VLSI CAD: methodology for heuristic development, experimentation and reporting. In: Proc. 1999 Design Automation Conference; 1999 June.

14. Karypis G, Aggarwal R, Kumar V, Shekhar S. Multilevel hypergraph partitioning: application in VLSI domain. In: Proc. 34th Design Automation Conference; 1997. p. 526–529. doi: 10.1109/DAC.1997.597203.

15. Microsoft. Xbox Series X [accessed 2023 Jan 31]. https://www.xbox.com/en-US/consoles/xbox-series-x#overview.

List of Symbols, Abbreviations, and Acronyms

AES	advanced encryption standard
ASIC	application specific integrated circuit
C_p	connection possibility
DOD	Department of Defense
FPNA	field programmable neural array
GCD	greatest common denominator
GPU	graphics processing unit
HD	Hamming Distance
IC	integrated circuit
I&D	instruction and data
I/O	input/output
$N_{cascaded\_cells}$	cascaded standard cells or macros
P_{depth}	partitioning depth
PPA	power–performance–area
RISC	reduced instruction set computer
TFLOP	one trillion floating-point operations per second

1 DEFENSE TECHNICAL
(PDF) INFORMATION CTR
DTIC OCA

1 DEVCOM ARL
(PDF) FCDD RLB CI
TECH LIB

1 DEVCOM ARL
(PDF) FCDD RLA PE
T NIGUSSIE