

Detecting Hardware Trojans Using Backside Optical Imaging of Embedded Watermarks

Boyou Zhou, Ronen Adato, Mahmoud Zangeneh, Tianyu Yang, Aydan Uyar,
Bennett Goldberg, Selim Unlu, Ajay Joshi
ECE Department, Boston University, Boston MA 02215
Email: bobzhou@bu.edu

ABSTRACT

Hardware Trojans are a critical security threat to integrated circuits. We propose an optical method to detect and localize Trojans inserted during the chip fabrication stage. We engineer the fill cells in a standard cell library to be highly reflective at near-IR wavelengths so that they can be readily observed in an optical image taken through the backside of the chip. The pattern produced by their locations produces an easily measured watermark of the circuit layout. Replacement, modification or re-arrangement of these cells to add a Trojan can therefore be detected through rapid post-fabrication backside imaging. We evaluate our approach using various hardware blocks where the Trojan circuit area is less than 0.1% of the total area and it consumes less than 2% leakage power of the entire chip. In addition, we evaluate the tolerance of our methodology to background measurement noise and process variation.

1. INTRODUCTION

CMOS integrated circuits (ICs) have become ubiquitous and essential elements in not only everyday consumer electronics, but also in critical defense technologies and core municipal support systems. This generates an enormous demand for high-quality and low-cost IC chips, and this demand is being increasingly supported through global supply chains. The IC industry has also experienced a trend away from vertical integration, and the IC chip development cycle today is progressively more fragmented and dependent on third party intellectual property (IP) and foreign fabs. A generic IC chip development cycle, consists of specification, design, fabrication, testing and packaging phases [1–3]. During these various phases the IC chip faces threats in the form of hardware Trojans (HTs), IP privacy and IC overbuilding, reverse engineering, side-channel analysis and IC counterfeiting [4]. HT has become one of the most significant problems in hardware security, since the Trojans can control, modify, disable, or monitor the key information in the IC chip [4]. Malicious IP core insertion, design modification

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

DAC '15 June 07 - 11, 2015, San Francisco, CA, USA
Copyright 2015 ACM 978-1-4503-3520-1/15/06 ...\$15.00.
<http://dx.doi.org/10.1145/2744769.2744822>.

and layout modification during fabrication process are the three most common ways of Trojan placement [4]. The first two types of insertions can be detected during design verification using functional simulations. However, HTs inserted after reverse-engineering the entire layout can be extremely hard to detect using post-fabrication functional tests. These Trojans are often designed to be extremely small in terms of area and power and could be triggered by signals that can be either internal or external to the IC chip. Therefore, there is a pressing need for the development of new technologies to rapidly, accurately and robustly detect any HTs inserted in CMOS ICs during the fabrication stage [5].

In this paper, we propose an optical method that can rapidly and accurately detect malicious tampering and the insertion of HTs at the fabrication stage. Our technique relies on imaging a large (of the order of $100 \mu m^2 - 1 mm^2$) area through the backside of the IC at near-IR wavelengths where silicon is transparent. We engineer the fill cells in the standard cell library used to generate the IC layout such that the fill cells are highly reflective and show up as bright spots in the optical image of the chip. The pattern produced by the location of these fill cells therefore acts as a watermark of the IC layout. This watermark is generated pre-fabrication, without any measured “golden chip” reference. Any replacement, modification or rearrangement of these fill cells to add HTs can be easily detected by backside imaging after fabrication. In our proposed approach, there is negligible overhead as the only modification is to the layout of the fill cells in a standard cell library. Hence the embedding of the optical watermarks can be seamlessly integrated into the standard IC design workflow. The main contributions of this paper are as follows:

– *We propose to engineer the layout of the fill cells in a standard cell library. We demonstrate that such fill cells yield reflectance signals well above the background reflectance of the functional circuitry. Any replacement, modification or rearrangement of these fill cells to add HTs can be easily detected by backside imaging.*

– *We designed Advanced Encryption Standard (AES) and PIC (a finite state machine) hardware blocks (with and without Trojans) using standard Cadence toolflows in 45 nm technology and we present an evaluation of the application of our proposed approach to correctly identify the tampered AES and PIC hardware blocks. We also demonstrate the robustness our technique to process variations and noise.*

2. RELATED WORK

HTs added during the fabrication stage are extremely difficult to detect using these approaches as they have very

small area and power.

State-of-the-art techniques for detecting HTs include both destructive and non-destructive approaches. Destructive approaches require costly reverse engineering and precise measurement. These techniques are mostly affordable to large semiconductor companies [3, 6]. Non-destructive testing approaches take the form of functionality tests and side-channel analyses [2]. Functionality tests excite the device under test using a variety of systematically chosen input combinations and compare the observed response with the expected response. These tests are already being applied in current designs [7]. Some HTs, however, are triggered using external inputs, and have extremely low triggering probability. As a result, detecting these types of HTs can be costly and time consuming [8]. Side-channel analyses detect the HT by analyzing the physical characteristics of the IC chip. Common analysis techniques include power analysis [9–12], timing-based analysis [13, 14], emission measurements [15, 16] and thermal analysis [17–19].

The key idea here is to use the physical characteristics as signature of the chip. During the testing process, any deviation in the challenge and response pairs (CRPs) [20] from the pristine or “golden” chip indicates tampering. These non-destructive approaches are easy to apply but they need a large number of CRPs to provide complete coverage of measurement space [2]. Moreover, side-channel analysis is also subject to noise stemming from fabrication variations and environment noise during measurements.

On the design front, the BISA technique [21], replaces all fill cells with functional cells to prevent the insertion of HT cells through fill cells. However, an intruder could replace or shift functional cells of the HT.

3. OPTICAL WATERMARKS: DESIGN AND MEASUREMENT

We propose a dramatically different approach to detect HT in IC chips. We propose to embed maximal amount of the M1 layer in the fill cells during the design phase. A digital ASIC designed using a library containing these fill cells will have optical watermarks encoded into its layout. Backside optical imaging of the fabricated chip will enable us to extract the full standard cell layout of the chip with the watermarks, which in turn can be validated to detect any HT. Due to the high resolution of the near-IR laser, it is very easy to tell the difference between the layouts using simple post processing. The high resolution characterization brings us an advantage that we do not need actual measurements of the “golden” chip, but only the simulation of the physical layout. We can use Euclidean distances or correlation to compare the simulation outputs against measurements to achieve high detection rate without any extensive testbenches from CRPs. Our technique is more tolerant of measurement and environment noise as high resolution image of the circuits provides detailed information of the circuits.

Compared to the electrical testing methods, our optical measurements are highly modular and independent of the multitude of connections in the full IC chip. In addition, the physical principles behind the implementation of our watermarking scheme are also highly distinct from previous approaches. Although, like the first PUF [22], it is an optical method that utilizes embedded scatters, their intended design and functionality are fundamentally different. The scattering in [22] was explicitly designed to be random, im-

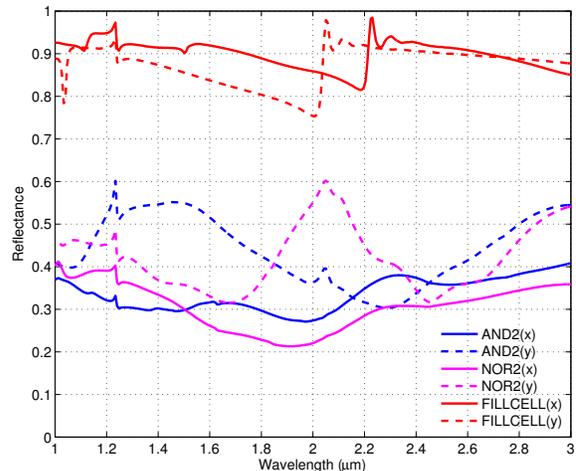


Figure 1: The reflectance spectrum of functional gates and fill cells, computed via FDTD simulations. The response is computed for both X and Y polarizations of the illuminating field (solid and dashed lines, respectively). For X polarization, the incident electric field is polarized along the VDD and VSS rails. For Y polarization, the polarization is perpendicular.

possible to predict and replicate. In contrast, determinism is essential to the functionality of our approach. The need for the optical watermark to be embedded at the design phase and its expected optical response is determined prior to fabrication means that its response must be predictable within the SNR of the measurement.

Backside imaging of integrated circuits is a well established technique for failure analysis [23–25]. Bright field images at near-infrared (IR) wavelengths (e.g. $\lambda \sim 1 - 2 \mu\text{m}$) can be used for passive measurements, such as inspecting the fidelity of the metal wires [26]. The active functionality of the circuit can also be probed via techniques such as thermal imaging [23] or laser-voltage imaging (LVI) [24]. These techniques record power dissipation via heat generation and the switching response of transistors respectively.

A key challenge in all of these techniques is to obtain sufficient spatial resolution to resolve the nanometer scale circuit elements associated with the current advanced technology nodes. The diffraction limit imposes a fundamental restriction on the maximum spatial frequency that can propagate macroscopic distances and therefore be imaged using conventional optical systems. This is given by $k_{max} = \frac{2\pi n}{\lambda}$ [27], where n is the refractive index of the material in which light propagates and λ is its wavelength. In a conventional microscope, not all spatial frequencies can be collected by the objective lens. Its numerical aperture ($NA = n \sin \theta$) dictates the maximum spatial frequencies that are used to form the image. The result is that the impulse response function of an image generally takes the form of an Airy function, [27]

$$I(\rho) \propto \left[2 \frac{J_1(2\pi\rho)}{2\pi\rho} \right]^2 \quad (1)$$

where $\rho = NA r / (M\lambda)$ is the image space coordinate. The size of the impulse response is

$$\Delta x = 0.61 \frac{\lambda}{NA} \quad (2)$$

While complex solid immersion lenses are used for the nanometer scale resolution required for failure analysis [25], we show here that relatively low resolution imaging at near-

IR wavelengths combined with a judicious engineering of the fill cell geometry can be used to rapidly and accurately detect malicious tampering and the presence of a HT.

Our approach is based on the fact that for low NA 's in the near-IR we can achieve impulse response functions with widths on the order of the gate size in 45 nm or below technology nodes (see Equation (2)). NA 's of 0.14, 0.42 and 0.5 correspond to spot sizes of approximately $4.6 \mu\text{m}$, $1.5 \mu\text{m}$ and $1.3 \mu\text{m}$, respectively at $\lambda = 1.064 \mu\text{m}$. These correspond to common near-IR commercial objective capable of imaging over $0.1 - 1 \text{ mm}$ fields of view (several thousand to half a million gates simultaneously). An image collected in this manner (i.e. at low NA , without a solid immersion lens) would, rather than resolve the detailed substructure of individual gates, produce a slowly varying image that tracks the average reflectance of each gate over its area. Although the individual gates are comprised of unique layouts of metal lines, Figure 1 shows that the variation in the response over the near-IR is characterized by subtle fluctuations in the reflected light about a baseline of $\sim 50\%$.

In contrast, the response of the fill cells can be engineered to explicitly reflect the majority of the incident light. Figure 1 demonstrates the results achieved by designing the fill cells containing the maximal amount of metal allowable within the design rule constraints. As a result, the cells produce signal well above the background fluctuations, which can then be easily observed in a low resolution image. Because the positions and sizes of the fill cells are indicative of the gate layout produced during the place-and-route stage of the design stage, they are directly influenced by the surrounding circuitry.

Therefore, if the fill cells in a standard cell library are explicitly designed as described, the clear pattern produced by their layout can serve as a robust, easily recordable optical watermark of the given circuit layout. Any changes in the placement or number of the fill cells (to tamper the IC) will result in a change in the watermark that can be measured with high fidelity. The potential to perform these measurements at low resolution enables large fields of view to be measured and therefore the simultaneous measurement of a large number of gates. It also has the added benefit of considerably simplifying the required optical setup in comparison with commercial failure analysis tools. Our approach represents a simple, rapid test for the insertion of hardware Trojans or tampering at the fabrication stage. Our approach requires negligible overhead and can be seamlessly integrated into a typical IC design test.

While we demonstrate our approach at the 45 nm technology node level, because our technique *does not* require high resolution imaging we expect it to scale well with the move to smaller technology nodes. We can estimate the efficacy of our proposed technique as gate sizes are reduced by considering the impulse response of our optical system as described by Equation (1). As the gate sizes are reduced and become progressively smaller in comparison with the impulse response, two effects will limit our ability to (i) detect the presence of a highly reflective fill cell and; (ii) detect a shift in the position of a highly reflective fill cell, and therefore the performance our our technique.

The first issue is due to the fact that a given point in the optical image will consist of the average of a progressively larger number of different gates. As a conservative estimate, consider an average gate size, D , such that $N =$

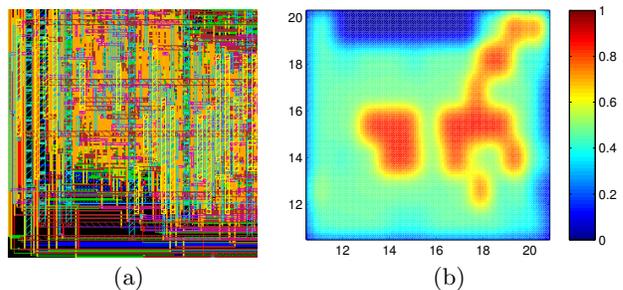


Figure 2: (a) Physical layout of a $10\mu\text{m} \times 10\mu\text{m}$ region of the AEST100 hardware block. (b) Backside image (reflectance value) of the $10\mu\text{m} \times 10\mu\text{m}$ region. The fill cells have the highest reflectance.

$\pi(S/2)^2/(D^2)$ gates are present (i.e. averaged together) in a spot size of diameter S . Using Figure 1 as a reference, we aim to detect a fill cell with reflectance of ~ 0.9 over a background with mean ~ 0.5 and fluctuations of $\sim \pm 0.1$. Therefore we need,

$$\frac{0.9 + 0.5(N - 1)}{N} \geq 0.6 \quad (3)$$

which implies $N \leq 4$ or $D \geq 330 \text{ nm}$. Given that a gate dimension, in the worst case, scales linearly with technology node, this estimate implies the potential to reach the 11 nm node while keeping a modest NA of 0.5 at $\lambda = 1.064 \mu\text{m}$. Increases in NA of up to 0.8 should be feasible if increased resolution or longer wavelengths are required.

At the same time, our technique also scales well to measuring large areas since we rely on low resolution imaging. This enables measurements over a large field of view and therefore a large number of gates can be imaged simultaneously. Considering an IC with $\sim 1 \text{ B}$ transistors, an average of 4 transistors per gate corresponds to approximately 250 M gates. For a gate size of $1.3 \mu\text{m}^2$ corresponding to the 45 nm technology node we examine here, imaging fields of view of $100 \mu\text{m}^2 - 1 \text{ mm}^2$ correspond to $10^4 - 10^6$ gates probed per image, and hence 250 - 25,000 images are required to scan a full IC. At 1-2 seconds per image, the entire IC area can be inspected by our technique in a matter of a few hours. It should be noted that to detect HTs inserted during the fabrication step we need to measure only a few chips.

4. HT DETECTION PROCESS

We explain our proposed approach of HT detection using AEST100 as an example hardware block [28]. We used the Cadence RTL Compiler and Cadence Encounter tool for synthesis, floor planning, place and route of the AEST100 hardware block. From the physical layout solution files (Figure 2(a)), we extract information to the DEF file (Gate Level Geometry Description File). We identify each gate location for imaging. The optical image can be generated by first mapping each gate type and location to the reflectance (see Figure 1) that is computed from detailed FDTD simulation. We approximate each gate as having uniform reflectance over its area. This is reasonable to the first order given that our spot size is larger than the gate area. To form the optical image, we convolve the reflectance map with the impulse response of our imaging system described by Equation 1. For the images in Figures 2 and 3, we used a wavelength of $1.064 \mu\text{m}$ and an NA of 0.6.

To insert HTs into the AEST100 hardware block we considered the following three changes to the physical layout:

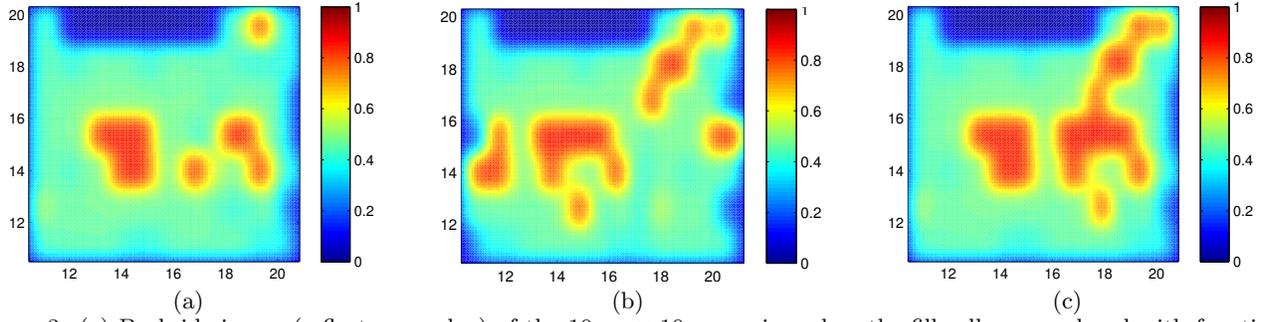


Figure 3: (a) Backside image (reflectance value) of the $10\mu\text{m} \times 10\mu\text{m}$ region when the fill cells are replaced with functional gates that constitute the Hardware Trojans. (b) Backside image (reflectance value) of the $10\mu\text{m} \times 10\mu\text{m}$ region when the bottom 3 rows are shifted by $5\mu\text{m}$ to the left make room for cells that constitute the HT. (c) Backside image (reflectance value) of the $10\mu\text{m} \times 10\mu\text{m}$ region when the functional cells are replaced by a different set of cells that constitute the HT.

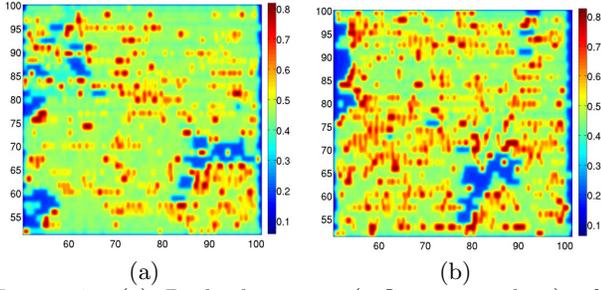


Figure 4: (a) Backside image (reflectance values) of a Trojan-free $50\mu\text{m} \times 50\mu\text{m}$ region of the AEST100 hardware block. (b) Backside image (reflectance values) of the same $50\mu\text{m} \times 50\mu\text{m}$ region of the AEST100 hardware block with CDMA private key disclosure type of HT inserted in it.

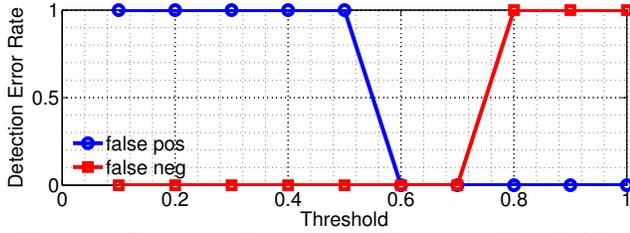


Figure 5: Detection Error rate vs Detection Threshold.

- *Replacement_Type1* - Changing fill cells into functional gates that are used as part of a Hardware Trojan.
- *Replacement_Type2* - Replacing functional gates with other types of functional gates that are used as part of a Hardware Trojan.
- *Shift* - Shifting several rows of logic gates to accommodate Hardware Trojans.

Figures 2 and 3 show the actual layout and the backside image (intensity of the reflected light) subject to various changes to the physical layout of a $10\mu\text{m} \times 10\mu\text{m}$ region in the bottom left hand corner of the AEST100 hardware block. Here, the color scale indicates the reflectance, which is the fraction of the incident power that is reflected. Figure 2(b) shows the ‘golden’ reference that is determined using simulations. In Figure 2(b) the high reflectance value regions correspond to the fill cells with metal structures. Figure 3(a) shows the backside image of *Replacement_Type1* case. We can see that this changes the reflectance map. This change can be easily detected, and it indicates the tampering of the IC physical layout. Figure 3(b) corresponds to the case where the bottom three rows of cells are shifted by $5\mu\text{m}$ to the left to make room for HT. Similar to the earlier

case, we can easily see a difference in the reflectance map in Figure 2(b) and 3(b). This difference indicates tampering of the layout.

Figure 3(c) shows the backside image of the case where functional logic gates are replaced with functional logic gates that constitute the HTs. In this case the reflectance map looks similar to the reflectance map of the original untampered layout (see Figure 2(b)). To determine if there is any difference between the two maps, we use the 2D-correlation method. Here, using the normalized values in the backside imaging matrix of two images, the correlation co-efficient can be used to quantify the displacement and deformation of the two images [29–31]. If the matrix of the measured backside image (post-fabrication) is $\mathbf{M}_{i \times j}$ and the matrix for the ‘golden’ reference (determined through simulation) is $\mathbf{N}_{i \times j}$ in Equation (4), then the similarity between the two images can be quantified using the correlation coefficient (r) between two matrices M and N . This correlation coefficient can be calculated as:

$$r = \frac{\sum_i \sum_j (M_{ij} - \bar{M})(N_{ij} - \bar{N})}{\sqrt{(\sum_i \sum_j (M_{ij} - \bar{M})^2)(\sum_i \sum_j (N_{ij} - \bar{N})^2)}} \quad (4)$$

We use a thresholding mechanism to determine if the two images match with each other. As we may have noise from various sources like process variation and measurement noise, we need to determine a threshold to determine whether the chip has been tampered or not. If r is greater than a $r_{threshold}$ value, it indicates there is a match between the two images. This means that the layout has not been tampered. On the other hand, if r is lower than the $r_{threshold}$ value then it indicates a mismatch, which corresponds to tampering of the IC layout. Here, the $r_{threshold}$ value needs to be carefully chosen such that we do not incorrectly detect a Trojan (false positive) and at the same time we do not miss a Trojan (false negative). Figure 4 shows the variation in the false positive and false negative detection rate of HT for various threshold values. For small threshold values we have a non-zero probability of false positive detection. Beyond a threshold value of 0.6, the probability of a false positive detection is zero. On the other hand, we may have a false negative detection for large values of threshold. As we decrease the threshold value, the probability of false negative detection goes to zero. Hence, we choose a threshold value of 0.65, which has a zero probability of false positive detection and false negative detection.

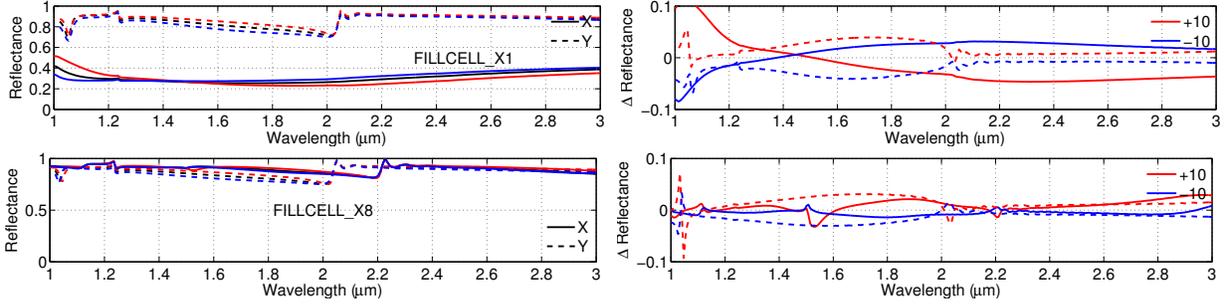


Figure 6: Impact of process variations on the reflectance signal for two different types of fill cells for various wavelengths. Black lines correspond to the nominal value, Red lines correspond to +10% variation and Blue lines correspond to -10% variation. X and Y correspond to X polarization and Y polarization, respectively.

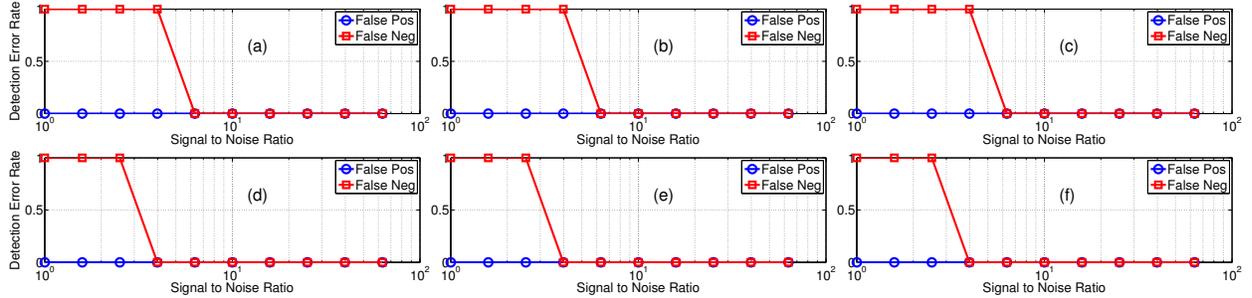


Figure 7: Testbench AES and PIC Trojan Detection Rate under different Signal-to-Noise Ratio (a) is AES100 (b) is AES200 (c) is AES1000 (d) is PIC100 (e) is PIC200 (f) is PIC300

Testbench	Trojan Leakage Percentage (% of Total leakage)	Trojan Total Power Percentage (% of Total power)
AES100	1.64	1.05
AES200	0.8	1.0
AES1000	1.64	1.05
PIC100	1.0	29
PIC200	1.21	34
PIC300	16	60

Table 1: Trojan power overhead in various hardware blocks. The baseline designs of AES100, AES200 and AES1000 are the same, the inserted HTs are different. Similarly, the baseline design of PIC100, PIC200 and PIC300 are the same but the inserted HTs are different. The power calculation was done using *Nangate45nm* high performance library.

We applied this 2D-correlation approach to images in Figures 2(b) and 3(c). Here the correlation co-efficient evaluates to 0.11 which is less than the $r_{threshold}$. This indicates tampering of the physical layout. The approach described above for Trojan detection scales well with image sizes. The optical test setup can easily capture images as large as $50\mu m \times 50\mu m$ from the backside of the IC. Figure 4(a) shows the backside image of a Trojan-free $50\mu m \times 50\mu m$ region of the AEST100 hardware block. The positions of the fill cells can be easily identified as the bright spots with reflectance of around 0.8. The functional gates have an average reflectance value of 0.5 with variations of 0.1. After inserting the Trojans into the circuits, Figure 4(b) shows the backside image of the tampered layout. In this case, the correlation co-efficient evaluates to 0.18, which is lower than the $r_{threshold}$. This indicates tampering of the layout.

Testbench	Area without Trojans	Trojan Area	Trojan Area Percentage (%)
AES100	274177.6	253.2	0.0923
AES200	274177.6	169.5	0.0618
AES1000	274177.6	251.1	0.0915
PIC100	4215.0	351	8.33
PIC200	4215.0	89.6	2.13
PIC300	4215.0	253.2	6.01

Table 2: Area (in μm^2) occupied by various hardware blocks.

5. EVALUATION

In this section we provide an evaluation of our proposed optical watermarking technique when subject to process variations as well as when applied to other hardware blocks. As described in the previous section, the measured backside optical image can be correlated with the simulated ‘golden’ version to detect any insertion of HTs. The measured reflectance signal can be affected by the process variations. To determine the fidelity of our approach in presence of process variations, we used Monte Carlo simulations to determine the reflectance signal of the fill cells in presence of $\pm 10\%$ process variations. Figure 6 shows the absolute reflectance values and the change in reflectance value with $\pm 10\%$ variations in the metal structure inserted in two different types of fill cells over a range of wavelengths. On average there is less than 5% change in reflectance for $\pm 10\%$ process variations. This change in reflectance value is below the measurement noise level and can be easily tolerated.

To test our proposed technique on other hardware blocks, we used the standard digital ASIC toolflow described earlier to generate the detailed place-and-route solutions for each hardware block. These detailed solutions were used to calculate the leakage power and area of the Trojan circuits within each hardware block (see Table 1 and 2). Here we consid-

ered leakage power instead of dynamic power because the HTs have extremely low triggering rates and so we may see variations in the dynamic power only when the circuit is operated under specific conditions. The power overhead of the HT circuits for the AES hardware blocks is within the standard deviation of leakage power due to process variation. At the same time, the HTs are added such that they are not part of the critical path. As a result, it would be very difficult to detect these HTs through power or delay analysis. Our proposed optical watermarking technique though was able to detect these HTs as it is agnostic of these power/area overheads.

To further evaluate the capabilities of our approach to detect HTs, we ran Monte Carlo simulations to determine the rate of incorrectly detecting (false positive) and incorrectly not detecting (false negative) HTs as a function of signal-to-noise ratio. Figure 7 shows the plots for the rate of false positive detections and false negative detections for varying signal-to-noise ratio for the three different AES circuits. On average for all three AES and PIC circuits, detection error rate is zero if the SNR is above 7. On the other hand, the rate of incorrectly detecting a HT is zero for all SNR values. In case of the PIC circuit, the leakage power and area overhead of the HTs is non-trivial (see Table 1 and 2). So power or delay analysis approach could be used for detecting HTs in these hardware blocks. Our proposed optical watermarking technique can also readily detect HTs in these hardware blocks. On average for all three PIC circuits, detection error rate is zero if the SNR is above 4.

6. CONCLUSION

Hardware Trojans added during the fabrication stage are hard to detect using current HT detection techniques. In this paper, we have proposed a new technique that uses backside imaging for detecting HTs. We embed maximal amount of the M1 layer in the fill cells during the design phase so that any HTs inserted by replacement, modification or shifting of the fill cells during the fabrication stage can be easily detected. We evaluate our proposed approach using a variety of AES and PIC blocks. We show that we are able to detect HTs that have power consumption that is less than 2% of the total chip and an area that is less than 0.1% of the total area (which makes it very difficult to detect these HTs using power and delay analysis). We also show that our approach is robust to measurement noise and $\pm 10\%$ process variations. Moving forward, we are going to work on classifications of the functional cells under near-IR photonic responses. This will provide protection against single cell modification HTs.

7. REFERENCES

- [1] DARPA. (2007) Trust in integrated circuits (tic) - proposer information pamphlet. [Online]. Available: <http://www.darpa.mil/MTO/solicitations/baa07-24/index.html>
- [2] R. Chakraborty *et al.*, "Hardware trojan: Threats and emerging solutions," in *Proc. HLDVT*, Nov 2009, pp. 166–171.
- [3] R. Karri *et al.*, "Trustworthy hardware: Identifying and classifying hardware trojans," *Computer*, vol. 43, no. 10, pp. 39–46, Oct 2010.
- [4] M. Rostami *et al.*, "A primer on hardware security: Models, methods, and metrics," *Proceedings of the IEEE*, vol. 102, no. 8, pp. 1283–1295, 2014.
- [5] N. Tsoutsos and M. Maniatakos, "Fabrication attacks: Zero-overhead malicious modifications enabling modern microprocessor privilege escalation," *Emerging Topics in Computing, IEEE Transactions on*, vol. 2, no. 1, pp. 81–93, March 2014.
- [6] M. Tehranipoor and F. Koushanfar, "A survey of hardware trojan taxonomy and detection," *Design Test of Computers, IEEE*, vol. 27, no. 1, pp. 10–25, Jan 2010.
- [7] L. Lin *et al.*, "Trojan side-channels: Lightweight hardware trojans through side-channel engineering," in *Proc. CHES*, 2009, pp. 382–395.
- [8] S. Wei *et al.*, "Hardware trojan horse benchmark via optimal creation and placement of malicious circuitry," in *Proc. DAC*, 2012, pp. 90–95.
- [9] R. Rad *et al.*, "Power supply signal calibration techniques for improving detection resolution to hardware trojans," in *Proc. ICCAD*, Nov 2008, pp. 632–639.
- [10] Y. Alkabani and F. Koushanfar, "Consistency-based characterization for ic trojan detection," in *Proc. ICCAD*, Nov 2009, pp. 123–127.
- [11] M. Potkonjak *et al.*, "Hardware trojan horse detection using gate-level characterization," in *Proc. DAC*, July 2009, pp. 688–693.
- [12] S. Wei and M. Potkonjak, "Scalable hardware trojan diagnosis," *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, vol. 20, no. 6, pp. 1049–1057, June 2012.
- [13] J. Li and J. Lach, "At-speed delay characterization for ic authentication and trojan horse detection," in *Proc. HST*, 2008, pp. 8–14.
- [14] Y. Jin and Y. Makris, "Hardware trojan detection using path delay fingerprint," in *Proc. HOST*, June 2008, pp. 51–57.
- [15] P. Song, F. Stellari, D. Pfeiffer, J. Culp, A. Weger, A. Bonnoit, B. Wisnieff, and M. Taubenblatt, "Marvel: Malicious alteration recognition and verification by emission of light," in *Hardware-Oriented Security and Trust (HOST), 2011 IEEE International Symposium on*. IEEE, 2011, pp. 117–121.
- [16] F. Stellari, P. Song, and H. A. Ainspan, "Functional block extraction for hardware security detection using time-integrated and time-resolved emission measurements," in *VLSI Test Symposium (VTS), 2014 IEEE 32nd*. IEEE, 2014, pp. 1–6.
- [17] A. N. Nowroz, K. Hu, F. Koushanfar, and S. Reda, "Novel techniques for high-sensitivity hardware trojan detection using thermal and power maps," *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on*, vol. 33, no. 12, pp. 1792–1805, 2014.
- [18] D. Forte *et al.*, "Temperature tracking: An innovative run-time approach for hardware trojan detection," in *Proc. ICCAD*, 2013, pp. 532–539.
- [19] K. Hu *et al.*, "High-sensitivity hardware trojan detection using multimodal characterization," in *Proc. DATE*, March 2013, pp. 1271–1276.
- [20] J. Kong *et al.*, "Pufatt: Embedded platform attestation based on novel processor-based pufs," in *Proc. DAC*, 2014, pp. 1–6.
- [21] K. Xiao and M. Tehranipoor, "Bisa: Built-in self-authentication for preventing hardware trojan insertion," in *Hardware-Oriented Security and Trust (HOST), 2013 IEEE International Symposium on*. IEEE, 2013, pp. 45–50.
- [22] R. Pappu, B. Recht, J. Taylor, and N. Gershenfeld, "Physical one-way functions," *Science*, vol. 297, no. 5589, pp. 2026–2030, 2002.
- [23] S. B. Ippolito *et al.*, "High spatial resolution subsurface thermal emission microscopy," *Applied Physics Letters*, vol. 84, no. 22, p. 4529, 2004.
- [24] U. Kindereit *et al.*, "Quantitative Investigation of Laser Beam Modulation in Electrically Active Devices as Used in Laser Voltage Probing," *IEEE Transactions on Device and Materials Reliability*, vol. 7, no. 1, pp. 19–30, Mar. 2007.
- [25] F. H. Köklü and M. S. Ünlü, "Subsurface microscopy of interconnect layers of an integrated circuit," *Optics letters*, vol. 35, no. 2, pp. 184–6, Jan. 2010.
- [26] S. B. Ippolito, B. B. Goldberg, and M. S. Ünlü, "Theoretical analysis of numerical aperture increasing lens microscopy," *Journal of Applied Physics*, vol. 97, no. 5, p. 053105, 2005.
- [27] L. Novotny and B. Hecht, *Principles of Nano-Optics*. Cambridge, UK: Cambridge University Press, 2006.
- [28] "Trust-hub website," <https://www.trust-hub.org/>, accessed: 2014-11-30.
- [29] H. Pohl *et al.*, "Arrangement for control of aerial cameras," Dec. 14 1976, uS Patent 3,997,795.
- [30] T. Keating, P. Wolf, and F. Scarpace, "An improved method of digital image correlation," *Photogrammetric Engineering and Remote Sensing*, vol. 41, no. 8, 1975.
- [31] L. Sorgi and K. Daniilidis, "Normalized cross-correlation for spherical images," in *Computer Vision-ECCV 2004*. Springer, 2004, pp. 542–553.