Security on Plastics: Fake or Real?

Nele Mentens[1,2], Jan Genoe[3,2], Thomas Vandenabeele[1,2], Lynn Verschueren[3,4], Dirk Smets[2], Wim Dehaene[4] and Kris Myny[3] *

[1] KU Leuven, imec-COSIC, Leuven, Belgium
[2] KU Leuven, ES&S, Diepenbeek, Belgium
[3] IMEC, Large Area Electronics, Leuven, Belgium
[4] KU Leuven, MICAS, Leuven, Belgium
firstname.lastname@kuleuven.be

Abstract. Electronic devices on plastic foil, also referred to as flexible electronics, are making their way into mainstream applications. In the near future, flexible electronic labels can be embedded in smart blisters, but also used as mainstream technology for flexible medical patches. A key technology for flexible electronics is based on thin-film transistors, which have the potential to be manufactured at low cost, making them an ideal candidate for these applications. Yet, up to now, no-one is taking digital security into account in the design of flexible electronics.

In this paper, we present, to our knowledge, the first cryptographic core on plastic foil. Two main research challenges arise. The first challenge is related to the reliability of the circuit, which typically decreases when the circuit area increases. By integrating cryptographic modules, we explore the limits of the technology, since the smallest lightweight block ciphers feature a larger area than the largest digital circuit on flex foil reported up to now. The second challenge is related to key hiding. The relatively large features on the chip and the fact that electronic chips on plastics are used as bare dies, i.e. they are not packaged, make it easy to read out the value of the stored secret key. Because there is no dedicated non-volatile memory technology yet, existing methods for writing data to the flexible chip after fabrication are based on wire cutting with a laser or inkjet printing. With these techniques, however, it is extremely easy to "see" the value of the secret key under a microscope. We propose a novel solution that allows us to invisibly program the key after fabrication.

Keywords: flexible electronics · thin-film transistors (TFTs) on plastic foil · IGZO (indium-gallium-zinc-oxide) · KTANTAN · low-cost cryptographic hardware

1 Introduction

Flexible electronics refers to electronic circuits that are fabricated directly on plastic substrates. These plastic substrates are not able to deal with high process temperatures typically used in the silicon industry and, as a consequence of this lower temperature budget, a lower semiconductor performance is obtained, combined with lower tool cost. Because of the low dimensional stability of flexible substrates, high-end lithography is excluded. This results in devices on foil with gate lengths in the range of 5 µm to 2 µm. This again leads to lower overall circuit performance at lower cost due to lower-end lithography tools. There are several applications that do not require high-end performance operation: e.g., amongst others, our internet-of-things (IoT) age will need a broad range of tiny distributed sensors with limited data rate. These include activity sensors and

*This work received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement no. 716426 - FLICs project).

IACR Transactions on Cryptographic Hardware and Embedded Systems ISSN 2569-2925,
Vol. 2019, No. 4, pp. 1–16
DOI:10.13154/tches.v2019.i4.1-16

medical sensors. For those devices, low production costs are much more crucial than top performance. This is hence the market opportunity for flexible electronics [Fit18].

In this paper, we concentrate on radio-frequency identification tags on flex foil, implemented as passive tags, i.e. they are not battery-powered, but receive power wirelessly from another (active) device. Most passive electronic chips on plastics target Near Field Communication (NFC), which means that they communicate wirelessly within a range of about 10 cm with an active device. Examples of applications are medical patches, electronic labels on medicines and on groceries, and electronics in textile. While technological research concentrates on overcoming physical challenges to increase the reliability, it is also important to build in security features in these NFC systems; this challenge has not been tackled (up to now) by the research community.

In this work, we present, to our knowledge, the first cryptographic core implemented on plastic foil. Based on the realizations, we identify two main challenges:

1. Cryptographic algorithms need a significantly higher number of transistors than the largest digital circuit on flex foil reported up to now. Therefore, integrating security into flexible electronics pushes the limits of the technology.

2. The features on the flexible chip are relatively large, e.g. the transistor gate length in the technology we use is 5 µm, while silicon chips in the latest technology nodes have gate lengths below 10 nm. Moreover, chips on flex foil are not embedded in a package, which means that reverse engineering is possible without depackaging the chip. Both the large features and the absence of a package, make it extremely easy to inspect the chip under a microscope. Given that there is no dedicated memory technology yet, the secret key in a cryptographic circuit needs to be printed on the foil using inkjet printing or programmed by cutting wires with a laser. For both methods, it is straightforward to read out the key with an inexpensive microscope.

With respect to the first challenge, this paper serves as a driver for technology research, showing the importance of sufficiently large digital circuits on flex foil. In the light of the second challenge, this work proposes a novel solution for hiding the secret key based on laser injection. In summary, our contributions are the following:

- We report the first, to our knowledge, working cryptographic core on plastic foil.

- We provide experimental implementation results on the area, the execution time, and the power supply of the core.

- We identify the issue of key hiding and we propose a solution to invisibly program the key using a laser.

- We give an overview of the remaining challenges in the use of digital security on flex foil.

The paper is organized as follows. In Sect. 2, related work on flexible electronic circuits is given. Sect. 3 addresses the background knowledge that is necessary to understand the remainder of the paper. In Sect. 4, we present the chip and the measurement results. Sect. 5 introduces our proposed solution for key hiding. Finally, Sect. 6 addresses the remaining challenges and Sect. 7 concludes the paper.

2 Related Work

2.1 Circuits on Plastic Foil

The initiator behind the development of transistor technologies on flexible foil has been the display industry, performing research toward flexible, rollable, curved and unbreakable displays. Several thin film transistor (TFT) technologies on foil have therefore been developed, ranging from organic TFTs, amorphous silicon TFTs, Low-Temperature Polycrystalline Silicon (LTPS) and more recently a broad range of amorphous oxides, such as $InGaZnO_4$ (IGZO) [NOT+04].

Researchers quickly realized that these low-cost upscalable TFT technologies had opportunities beyond display backplanes, such as, amongst others, low-cost flexible RFID tags [CGG+07], distributed sensor arrays on large area foils [MSVVH12] and flexible EMG front-end circuits [GAvdS+18]; even small organic 8-bit microprocessors have been implemented [MSR+14]. An overview of TFT technologies for low-cost, flexible integrated circuits is discussed in [Myn18]. The authors evaluate different non-silicon technologies and propose a roadmap similar to Moore's law for flexible electronics.

It is clear that the present state of TFT technologies on flexible foil does not yet enable circuits comprising billions of transistors, as is the case in silicon. However, displays on flex, comprising millions of TFTs have been shown. For digital circuitry based on unipolar technologies, the maximum integration density is presently around 3504 TFTs [MvVG+12]. Note that, in regular complementary metal oxide semiconductor (CMOS) technology, this would correspond to 876 equivalent NAND gates, since one equivalent NAND gate consists of four transistors. However, as explained in Sect. 3.1, in the standard cell library used in our work, one equivalent NAND gate consists of six transistors, which means the integration density reported in [MvVG+12] corresponds to only 584 equivalent NAND gates. Knowing that cryptographic ciphers, excluding a secure mode of operation, significantly exceed this amount of gates, it is clear that the ongoing progress in technological research is indispensable for the integration of security features in flexible chips. The technological challenges that are dealt with to increase the maximum integration density are intrinsic parameter variations, low noise margins and high power consumption [MGD16].

2.2 Non-Volatile Storage Mechanisms

Flexible chips make use of either inkjet printing or wire cutting with a laser to write non-volatile data into the chip after fabrication.

2.2.1 Inkjet Printing

In [MSR+14], a Write-Once-Read-Many (WORM) memory is proposed to store the instructions of an 8-bit microprocessor. The WORM memory consists of print-programmable memory cells, based on an interdigitated finger structure, as shown in Fig. 1 (left). The top and bottom connections are wires made of molybdenum (Mo). The fingers are situated in a well, created by an isolating ring. The ring is made of SU8, an epoxy-based negative photoresist. Silver (Ag) paste is used after fabrication to connect the fingers in the ring, resulting in the connection of the top and bottom wires. To construct a memory cell, the finger structure can be combined with a pull-up transistor load, that is always on. When the fingers are not connected, i.e. when no silver paste is applied, the current through the pull-up transistor connects the key bit to the supply voltage (V_{DD}), resulting in a logical 1. This is shown in Fig. 1 (middle). When the fingers are connected with silver paste, the current flows from the top to the bottom wire in the ring and the key bit is connected to the ground (GND), resulting in a logical 0. This works when the resistance of the finger structure, after the fingers have been connected with silver paste, is smaller than the resistance of the pull-up transistor load, as shown in Fig. 1 (right). It is clear that, when this memory cell is observed under a microscope, the key bit can be read out.

2.2.2 Wire Cutting with a Laser

Another way to program the key is through lasering [DRGK+19]. A simplified representation of the method proposed in [DRGK+19] can be obtained by replacing the interdigitated

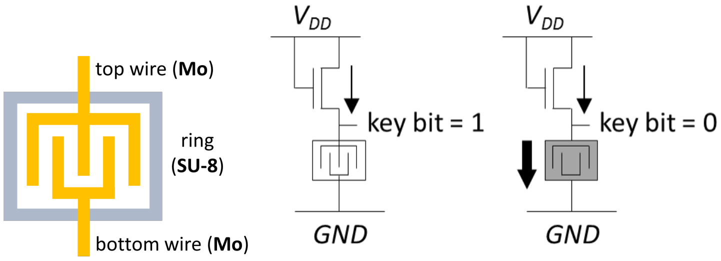

Figure 1: Interdigitated finger structure (left) and working principle of the WORM cells (middle and right) in [MSR+14].

finger structure in Fig. 1 by a connected wire, and to cut the wire with a laser to set the key bit to 1. The default key bit value when the wire is not cut is 0 in this mechanism. We use an even simpler structure to program the key bits: we connect each key bit to both V_{DD} and GND, and we break one of these connections with a laser. Obviously, the approach of using a laser for programming the key, is as vulnerable as the WORM memory when the key bits are observed under a microscope. One of our contributions is to propose a solution to this problem.

3 Preliminaries

3.1 Flex Foil Technology Used in This Paper

Fig. 2 shows the cross-section of a unipolar oxide n-type thin film transistor (n-TFT) based on a self-aligned IGZO (indium-gallium-zinc-oxide) semiconductor [NOF+14], which is used in this work. This technology is relatively stable and gives the best performance compared to other technologies with a similar cost. The semiconductor consists of amorphous IGZO (a-IGZO) and the transistor stack is fabricated on a plastic substrate. The metal used for the gate and source/drain layer is molybdenum (Mo). The gate dielectric is silicon dioxide (SiO_2) and the inter-metal dielectric is silicon nitride (SiN). The two metal layers used for the gate and source/drain layers also provide the wiring in the chip. While the best minimal channel length that has been published is 2 µm, the typically used minimal channel length is currently 5 µm. The charge carrier mobility of the TFTs is around 10 cm^2/Vs. The manufacturing only requires four photolithographic steps, which makes the manufacturing cost of the chips lower compared to other flexible TFT technologies with a comparable performance. A discussion on the yield is given in [GMSH10].

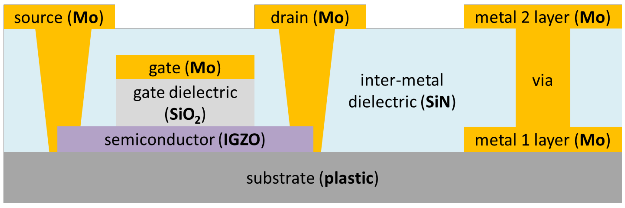

Figure 2: Cross-section of a unipolar oxide n-TFT on a plastic substrate.

Since only n-type transistors can be built in this technology, it is not possible to use complementary logic to construct standard cells. Therefore, a number of topologies, that originate from nMOS logic, have been suggested to realize logic gates on flex foil [MGD16]. An example of a naive topology is shown in Fig. 3 (left). The logic gate consists of an always-on enhancement-mode transistor load and a pull-down network (PDN). When the PDN conducts, the output voltage (V_{out}) is pulled down to the ground voltage (GND). E.g., for a NAND gate, this would be when all inputs are equal to V_{DD}. When the PDN does not conduct, the transistor load pulls up the output to $V_{DD} - V_T$, where V_T is the threshold voltage of the transistor. However, for increased reliability, it is desirable that V_{out} is pulled up to V_{DD}.

The topology we use in this work is referred to as pseudo-CMOS [HFL+11]. In pseudo-CMOS logic, each logic gate needs to be powered by the supply voltage (V_{DD}) and an additional bias voltage (V_{bias}) that is higher than V_{DD}. This is shown in Fig. 3 (right), where the first stage corresponds to the logic gate in Fig. 3 (left). The output voltage of the first stage ranges from GND to $V_{bias} - V_T$. In order for the range of V_{out} to go from GND to V_{DD}, we need to make sure that $V_{bias} > V_{DD} + V_T$. In the experiments presented in Sect. 4.3, two combinations of V_{DD} and V_{bias} are used, namely $V_{DD} = 10V$ and $V_{bias} = 15V$, and $V_{DD} = 11V$ and $V_{bias} = 16.5V$.

From Fig. 3 (right), we can derive that, in pseudo-CMOS, a NAND gate, i.e. the basic gate that is used to express the area of a digital circuit, consists of 6 transistors compared to 4 transistors for standard CMOS logic. For digital circuit design, a limited standard cell library is proposed in [MSR+14] based on pseudo-CMOS gates with a gate length of 5 µm. The standard cells that are available in our library are a single-drive inverter, a triple-drive inverter, a 2-input NOR, a 3-input NOR, a 4-input NOR and a D-flipflop.

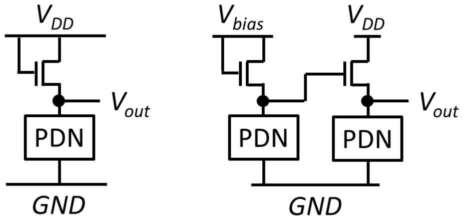

Figure 3: Diode-load logic gate with an always-on enhancement-mode transistor load (left) and pseudo-CMOS logic gate in which the pull-down network (PDN) is repeated twice (right).

In terms of digital design methodology, a strategy similar to standard CMOS design can be followed. However, since pseudo-CMOS logic gates only contain PDNs, NOR gates are favored over NAND gates, because they have a lower pull-down resistance when the same transistor sizes are used.

3.2 Selection of the Cryptographic Cipher

The implementation of digital logic circuits by n-type technologies yields a lower circuit robustness compared to complementary (n-type and p-type) circuits. Since the reliability of the circuit drops when the chip area increases, it is important to choose a cryptographic cipher that can be implemented with a minimal number of transistors. Of all properly analyzed block ciphers, the family of KATAN and KTANTAN ciphers [DCDK09] features the smallest number of gates and thus the smallest number of transistors. KATAN/KTANTAN is assumed to have a higher power consumption than other lightweight ciphers, due to the fact that the internal state is stored in shift registers. Nevertheless, KATAN/KTANTAN is preferred in our chip, where the first focus is on low area. All KATAN/KTANTAN variants are based on an 80-bit key, while the block size varies between 32, 48 and 64 bits. We implement the KTANTAN variant, which is smaller than KATAN, because it assumes that the key is burnt into the device, thus significantly reducing the gate count thanks to the avoidance of the 80-bit key register. In order to reduce the number of transistors to a minimum, we selected a block size of 32 bits. The security of KTANTAN is analyzed in [BR10, AL12, ZG14].

Block ciphers are preferred over stream ciphers in this work, since the security of block ciphers is much better understood by the cryptographic community. Therefore, it is much less risky and considered better practice to adopt a block cipher instead of a stream cipher.

4 Implementation and Measurement Results

4.1 The KTANTAN core

We describe the 32-bit KTANTAN core with an 80-bit key in Verilog. While the reference hardware implementation published by the inventors proposes the parallel loading and read-out of the plaintext and ciphertext, respectively, we implement a serial version to reduce the area and the input/output pads of the chip. This results in the digital hardware architecture for the 32-bit core in Fig. 4. It has three input bits, namely the plaintext (pt), the start signal (start) and the clock signal (clk). It has two output bits, namely the ciphertext (ct) and the ready signal (ready).

When the start signal is active, the pt bits are serially loaded into two shift registers, L1 (13 bits) and L2 (19 bits), and register T (8 bits) is initialized with the value 11111110. When the start signal becomes inactive, all registers (L1, L2 and T) work as linear shift registers. In each clock cycle, two new round key bits (ka and kb) are added to the feedback value of L2 and L1, respectively. After 254 clock cycles, T reaches the value 11111111 and the ready signal is activated. From then on, the bits in the registers L1 and L2 are serially shifted to the ciphertext output ct. Consequently, the number of clock cycles for one 32-bit KTANTAN encryption is equal to 32 (for shifting in pt) plus 254 (for the actual encryption) plus 32 (for shifting out ct). In total, this comes down to 318 clock cycles.

For testing purposes, we also route the outputs of the key schedule (ka and kb) and the outputs of the shift registers L1 and L2 (l1o and l2o) to the output. The programming of the 80-bit key is detailed in Sect. 4.2.

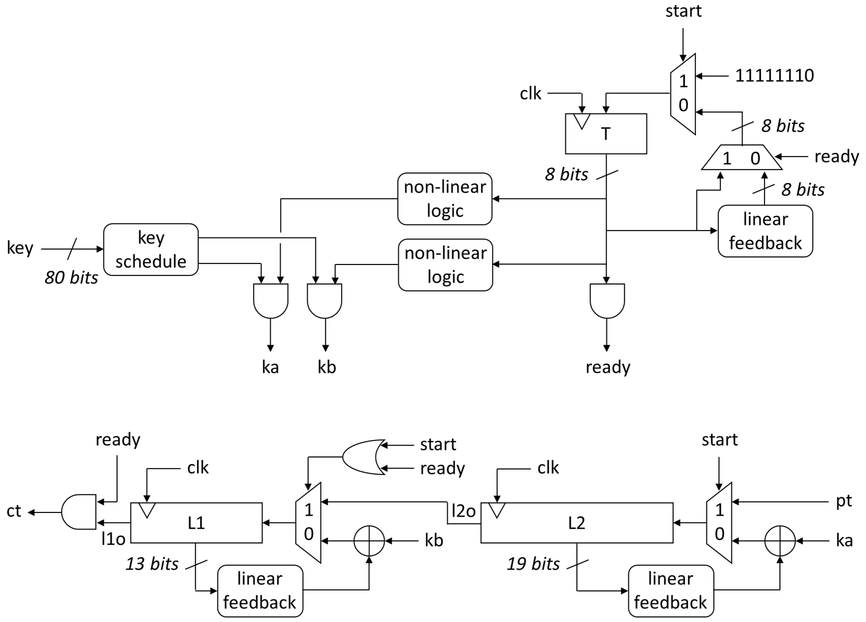

Figure 4: Hardware architecture of the KTANTAN32 implementation.

The physical lay-out of the KTANTAN32 core is shown in Fig. 5. It consists of 4044 transistors and measures a total area of 331.5 mm^2. The 80-bit key is applied on the top side of the core and the input/output pads are placed at the bottom side of the core. Note that there are 48 input/output pads, while we only have three inputs and two outputs. The reason is that our measurement setup contains a probe card that is compatible with this structure. Each core has three power rails, namely V_{DD}, V_{bias} and GND, as required for the pseudo-CMOS logic gates, explained in Sect. 3.1. The power rails are also driven by inputs that are part of these 48 pads.

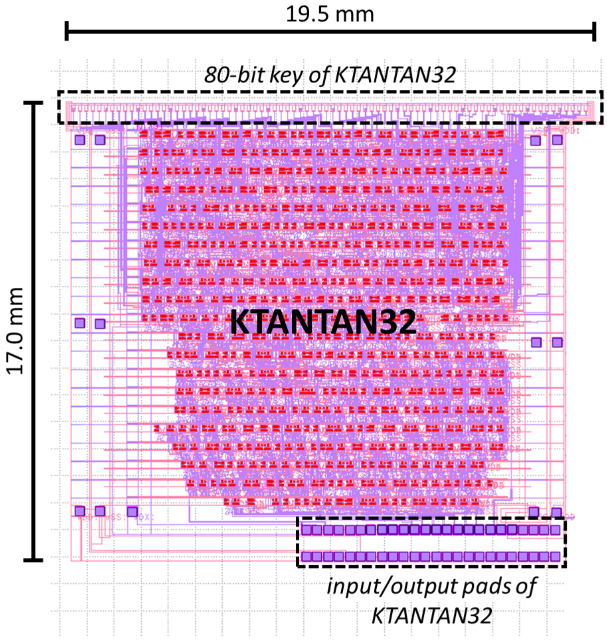

Figure 5: Physical lay-out of the KTANTAN32 core.

4.2 The one-time programmable key

Fig. 6 and Fig. 7 are microscope images that indicate how the 80-bit key is programmed. In Fig. 6, a subset of the key bits is shown, namely bits 26-36. After fabrication, each key signal is connected to both the power supply (V_{DD}) and the ground line (GND). When the key is programmed, one of both connections is broken with a laser; this is the state that can be seen in Fig. 6. The laser is an S-EZLI of Signatone, producing a green (532 nm) beam. Underneath the key bits, the figure also shows part of the digital logic gates.

When we zoom into two key bits, we get the image in Fig. 7, in which a clear visible distinction can be made between a connected wire and a wire that is cut by the laser. The left key signal in the figure is connected to V_{DD} (the connection with GND is broken). The right key signal in the figure is connected to GND (the connection with V_{DD} is broken).

Based on these microscope images, we can conclude that the applied method for programming the secret key clearly reveals the value of each bit, just like the WORM-based approach explained in Sect. 2.2. Therefore, we propose a novel approach, also using a laser, to visually hide the secret key bits. Sect. 5 elaborates on the novel key hiding mechanism.

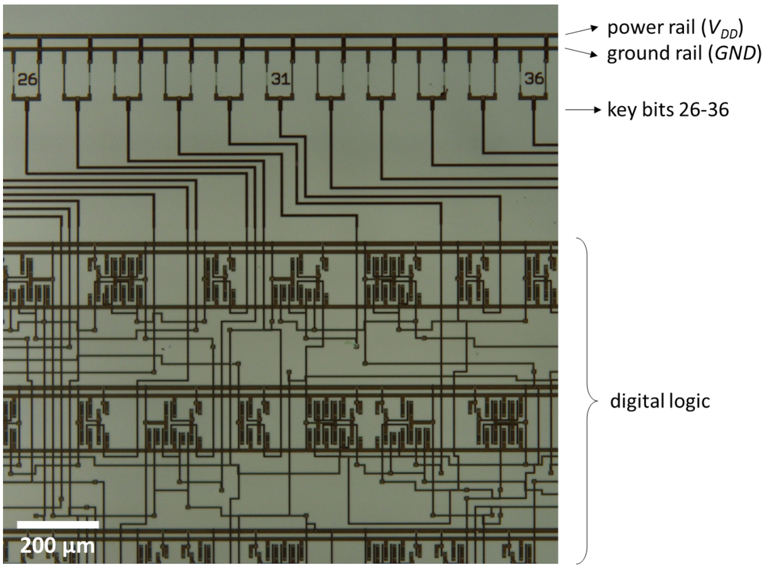

Figure 6: Microscope image of key bits 26-36 and part of the logic gates.

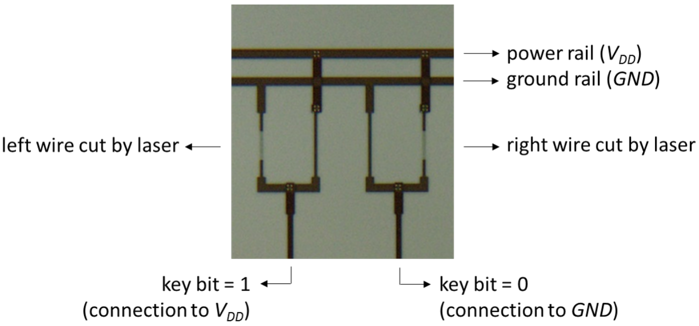

Figure 7: Microscope image of two key bits that are programmed to 1 and 0, respectively. The wires that are cut by the laser are visible.

4.3 Measurement Setup and Results

Before powering up the cores, the 80-bit keys are programmed to the fixed hexadecimal value 07C1F07C1F07C1F07C1F. The measurement setup contains a field-programmable gate array (FPGA) that generates 1000 different plaintexts and verifies the result. The FPGA keeps track of the number of correct ciphertexts and outputs the result on three LEDs: the first LED lights up when all 1000 ciphertext are correct; the second LED lights up when at least one but less than 1000 ciphertexts are correct; the third LED lights up when all ciphertexts are incorrect.

The inputs and outputs of the FPGA are connected to a printed circuit board (PCB) that contains level shifters to make the conversion between the 3.3 V pins of the FPGA and the voltage level of the flexible chip's inputs/outputs, which is equal to the power supply (V_{DD}). In between the flexible chip and the PCB, there is a probe card that connects to the 48 input/output pads that are described in Sect. 4.1.

Adaptive power supplies are used to drive V_{DD} and V_{bias}, such that we can explore which V_{DD} and V_{bias} voltages lead to a working cryptographic core. For the experiments, the flexible foil is not yet removed from the glass carrier on which it is fabricated. When the chip is used in a real-life application, it needs to be removed from the glass carrier.

We explore a number of V_{DD}-V_{bias} combinations, maintaining a fixed ratio of 1.5 for V_{bias}/V_{DD}. The 32-bit core functions correctly at a clock frequency of 10 kHz for V_{DD} = 10 V and V_{bias} = 15 V as well as for V_{DD} = 11 V and V_{bias} = 16.5 V. Given the measured clock frequency of 10 kHz and the cycle count of 318 cycles (cfr. Sect. 4.1), the total latency of one 32-bit KTANTAN encryption is equal to 31.8 ms. The settings and implementation results for the working 32-bit core are summarized in Table 1.

Table 1: Settings and implementation results for the KTANTAN32 core.

V_{DD}, V_{bias}	10 V, 15 V or 11 V, 16.5 V
Frequency	10 kHz
Cycle count	318 cycles
Latency	31.8 ms
Number of transistors	4044
Area	331.5 mm^2

5 Novel Approach for Key Hiding

5.1 Proposed Concept

This section describes our novel solution for programming the cryptographic key such that the key bits cannot be read out when the chip is inspected under a microscope. We explain how we change the state of a transistor through lasering and we propose two possible circuits for programming the key bits.

5.1.1 Changing the state of the transistor through lasering

Fig. 8 shows the effect that we expect to obtain. The graph shows the I_d current, i.e. the current that flows through the transistor, as a function of the V_g voltage, i.e. the voltage at the input of the transistor. The full line shows the I_d-V_g graph before lasering. The transistor starts conducting when $V_g > V_T$, where V_T is the threshold voltage. Although the threshold voltage is drawn as a negative voltage in Fig. 8, it can be positive or negative in reality. Lasering the source and the drain of the transistor causes a V_T shift. We refer to the new threshold voltage as V_T'. As a result of this V_T shift, we can modify the state of the transistor from non-conducting to conducting when we apply a well-chosen voltage V_g at the input, like V_{neg} in Fig. 8.

The physical effect that causes this phenomenon can be explained as follows. The amorphous structure of the IGZO semiconductor guarantees that good uniformity can be obtained over large area on foil, but also makes that the background doping and hence the threshold voltage V_T is strongly dependent on the overall temperature budget of the semiconductor layer during the anneal process. We cannot take the overall temperature of the anneal process too high, as this would have a detrimental impact on the foil and moreover would increase the overall circuit costs. This is important as these circuits are usually cost sensitive. However, although a longer overall system anneal is not preferred, a local anneal step of a single transistor channel can still be done by laser. Careful control of the laser pulse shape allows us to obtain a specific amount of V_T shift.

5.1.2 Programming the key bits

Each key bit is programmed by a circuit that uses the effect explained in Sect. 5.1.1. We propose two options for the circuit.

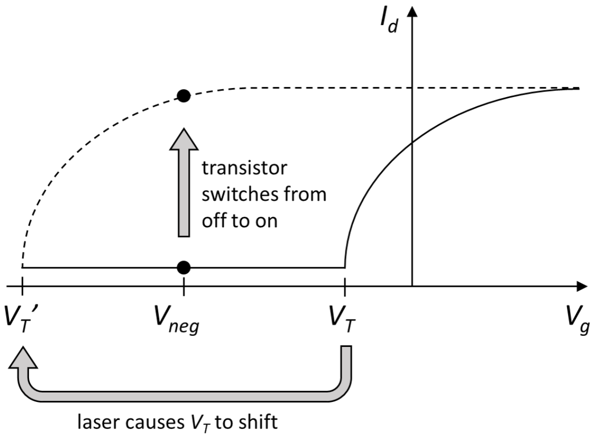

Figure 8: I_d-V_g graph before (full line) and after lasering (dotted line). The threshold voltage changes from V_T into V_T'. When V_{neg} is applied at the input, the transistor goes from a non-conducting to a conducting state.

Fig. 9 shows the first option, which consists of a pull-up transistor load that is always on, and a pull-down transistor that is driven by a fixed negative voltage (V_{neg}). In the default state (left side of Fig. 9), this voltage does not induce a current in the pull-down transistor. As a result, the pull-up current causes the key bit to be equal to a logical 1. When the source and the drain of the pull-down transistor are lasered, the pull-down transistor starts conducting (right side of Fig. 9) and the key bit becomes equal to a logical 0. The disadvantage of this circuit is the static power consumption of the always-on pull-up transistor.

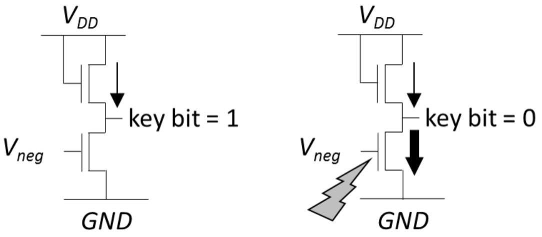

Figure 9: First option for the key programming circuit.

Fig. 10 solves this issue by using a pull-up and a pull-down transistor that are both switched off in the default state - they both have $V_g = V_{neg}$. The laser pulse is then applied to one of these transistors in order to pull the output either down or up.

5.2 Experimental Validation

To validate the solution proposed in Sect. 5.1, we explore different settings of the Signatone S-EZLI laser in order to find a setting that causes the effect explained in Sect. 5.1.1. We apply two different energy settings (“low” and “high”), 16 different attenuation settings (ranging from 1 dB to 45 dB), and a repetition of one or two pulses for lasering individual transistors. The first thing we notice, as shown in Fig. 11, is that there is a visual difference between a transistor that is lasered (left side of the figure) and a transistor that is not (right side of the figure). The edges of the spots that were lasered, are indicated in the

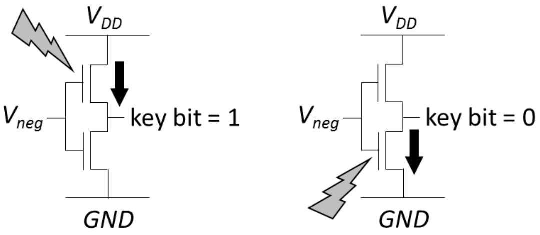

Figure 10: Second option for the key programming circuit.

figure. On both microscope images, we see the wiring of the source and the drain of the transistor at the top and the bottom. The gate extends to the right.

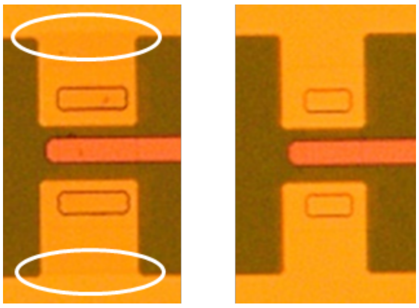

Figure 11: Microscope image of a transistor that is lasered (left) and a transistor that is not lasered (right).

Due to the observed difference in Fig. 11, we cannot directly use the concept explained Sect. 5.1. However, a slight modification of the approach allows us to reach the desired goal of changing the state of a transistor without visually revealing this. Our modified approach consists of finding two laser settings that are visually indistinguishable but lead to a different transistor state. The result of two settings that are particularly interesting, are shown in Fig. 12 and Fig. 13. From both images, we can detect the edge of the laser mark in the source and the drain, as indicated with white ovals. However, because different laser settings are applied, the lasering has a different effect on the transistor in Fig. 12 and Fig. 13. For the first transistor, the V_T shift causes the I_d-V_g curve to significantly shift to the left, switching the transistor from off to on when -5 V is applied at the input. For the second transistor, there is also a V_T shift resulting in a shift of the I_d-V_g curve, but it is not enough to make the transistor conduct at an input voltage of -5 V. For the transistor in Fig. 12, the attenuation of the Signatone S-EZLI laser is 45 dB in low energy mode, and one pulse is applied (we call this Setting 1). For the transistor in Fig. 13, the attenuation of the laser is 35 dB in low energy mode, and two pulses are applied (we call this Setting 2).

Consequently, our technique can be applied to the key programming circuit in Fig. 9 to change the state of the pull-down transistor from off to on by using a laser with Setting 1, programming the key bit to a logical 0. When we use Setting 2, the state of the pull-down transistor in Fig. 9 does not change, and the key bit remains at a logical 0.

Lasering the pull-down and pull-up transistors of the key programming circuit in Fig. 10 with Setting 1 and Setting 2, respectively, programs the key bit to a logical 1. Alternatively, swapping the settings for both transistors programs the key bit to a logical 0. Both key programming techniques (Fig. 9 and 10) lead to an overhead of two transistors per key bit, resulting in 4204 transistors for the entire circuit.

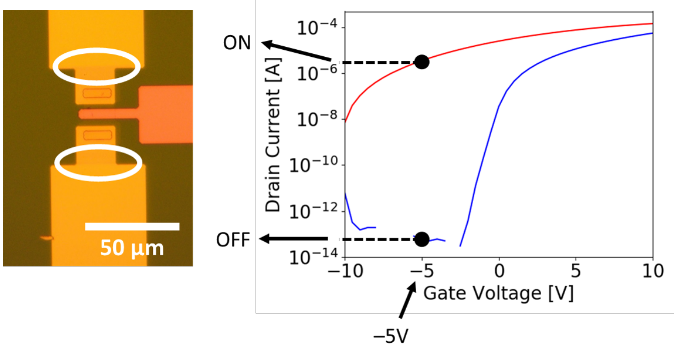

Figure 12: Microscope image and I_d-V_g graph of a transistor, corresponding to a laser configuration that changes the state of the transistor from off to on when the gate voltage is -5 V. The blue line indicates the current before lasering. The red line indicates the current after lasering. The laser is configured to an attenuation of 45 dB in low energy mode; one pulse is applied (Setting 1).

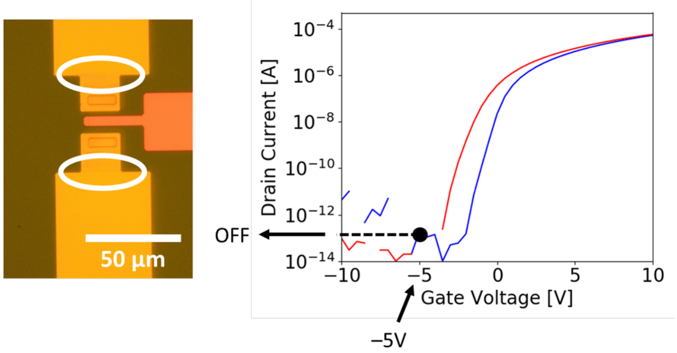

Figure 13: Microscope image and I_d-V_g graph of a transistor, corresponding to a laser configuration that does not change the state of the transistor when the gate voltage is -5 V. The blue line indicates the current before lasering. The red line indicates the current after lasering. The laser is configured to an attenuation of 35 dB in low energy mode; two pulses are applied (Setting 2).

6 Remaining Challenges and Opportunities

Since this is the first work that considers security of chips on flex foil, we elaborate on the challenges and opportunities in this section. Issues related to the reliability of digital circuits and key hiding have already been treated in this paper and will therefore not be repeated here.

6.1 Reverse Engineering and Intellectual Property Theft

Due to the much larger features sizes and the fact that flexible chips are unpackaged, it is much easier and cheaper to inspect a circuit on flex foil than on silicon. On the other hand, exposing the different layers in the chip is much harder thanks to the fragile nature of the foil. A possible solution to prevent an adversary from inspecting the chip, is to use printed ink as an additional layer at the top and the bottom of the foil. Note that we could have used this for key hiding as well. Nevertheless, for many types of ink, solvents exist, which would undo the protection method, while the lasering technique we propose, cannot be undone. Another solution to prevent reverse engineering and intellectual property theft is to obfuscate the presence of a via between two crossing wires or to re-use the key hiding technique that we propose in this work.

6.2 Side-channel Analysis

One of the strongest protection mechanisms against side-channel analysis attacks is performed at the logic level, using dual-rail, dynamic logic styles. All these protection mechanisms assume that the underlying technology is based on CMOS gates. In n-type technologies, like the one we consider in this work, novel logic styles need to be designed and experimentally validated.

6.3 True Random Number Generators

The slope of the input-output transfer characteristic of pseudo-CMOS gates in the a-IGZO technology is much less steep compared to CMOS gates, resulting in much lower noise margins. The noise margin is illustrated by the opening of the "eyes" in Fig. 14, which shows the transfer characteristic of an inverter in pseudo-CMOS logic (on the left) and of a CMOS inverter (on the right). The figure also indicates that, in pseudo-CMOS, it is harder to position the transition voltage at $V_{DD}/2$ than in CMOS, which leads to a biased output. This should be taken into account in the design of True Random Number Generators (TRNGs) in the presented technology. For example, TRNGs based on ring oscillators [CFAF13, RLG14] should be reinvestigated to maximize the entropy and to minimize the bias.

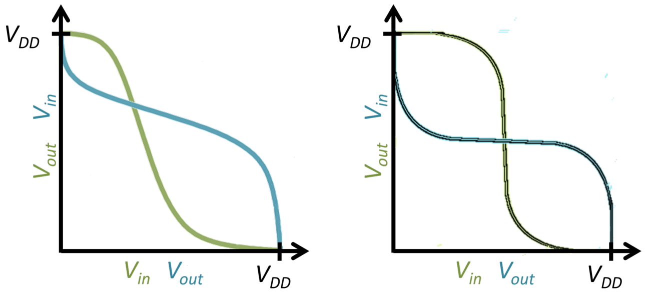

Figure 14: Voltage transfer characteristic for a pseudo-CMOS inverter (on the left) and a CMOS inverter (on the right).

6.4 Physical Unclonable Functions

Whereas temperature and supply voltage are the most important parameters to take into account in the reliability evaluation of silicon Physical Unclonable Functions (PUFs), we expect that an additional parameter will influence the physical behavior of the circuit in a flexible chip, namely the bending radius and the bending direction. This should be investigated for all types of PUFs, since the behavior of a PUF is influenced by very small changes in the transistor parameters, which might be caused by bending. Note that the functionality of deterministic digital and analog circuits does not suffer from bending, as discussed in [TMH$^+$15].

7 Conclusions

This paper addresses the security of chips on flexible plastic foil. The technology is rapidly gaining importance, especially in the design of medical patches and smart textiles, facilitated by the low-cost fabrication of flexible thin-film transistors. We present, to our knowledge, the first working cryptographic core on flex foil. An important issue that arises, is related to the hiding of the secret key. The lack of a dedicated memory technology, the large features on the chip and the fact that flexible chips are used as unpackaged bare dies, make it extremely easy to read out the secret key. We propose and validate a novel mechanism to "invisibly" program the bits of the secret key. The first results look promising, and further analysis will be conducted to investigate the stability of the proposed key hiding technique. Further, since this is the first work that considers the security of flexible chips, we give an outlook to the expected challenges and limitations related to reverse engineering, side-channel security, true random number generators and physical unclonable functions.

References

[AL12] Martin R Albrecht and Gregor Leander. An all-in-one approach to differential cryptanalysis for small block ciphers. In *International Conference on Selected Areas in Cryptography*, pages 1–15. Springer, 2012.

[BR10] Andrey Bogdanov and Christian Rechberger. A 3-subset meet-in-the-middle attack: cryptanalysis of the lightweight block cipher KTANTAN. In *Workshop on Selected Areas in Cryptography*, pages 229–240. Springer, 2010.

[CFAF13] Abdelkarim Cherkaoui, Viktor Fischer, Alain Aubert, and Laurent Fesquet. A self-timed ring based true random number generator. In *2013 IEEE 19th international symposium on asynchronous circuits and systems*, pages 99–106. IEEE, 2013.

[CGG$^+$07] Eugenio Cantatore, Thomas Geuns, Gerwin Gelinck, Erik van Veenendaal, Arnold Gruijthuijsen, Laurens Schrijnemakers, Steffen Drews, and Dago de Leeuw. A 13.56-MHz RFID system based on organic transponders. *IEEE Journal of Solid-State Circuits*, 42(1):84–92, 2007.

[DCDK09] Christophe De Canniere, Orr Dunkelman, and Miroslav Knežević. KATAN and KTANTAN—a family of small and efficient hardware-oriented block ciphers. In *Cryptographic Hardware and Embedded Systems-CHES 2009*, pages 272–288. Springer, 2009.

[DRGK$^+$19] Florian De Roose, Jan Genoe, Auke J Kronemeijer, Kris Myny, and Wim Dehaene. Memory solutions for flexible thin-film logic: up to 8kb,$>$ 105.9 kb/s lprom and sram with integrated timing generation meeting the iso nfc standard. In *IEEE International Solid-State Circuits Conference-(ISSCC)*, pages 206–208. IEEE, 2019.

[Fit18] Alissa M. Fitzgerald. The internet of disposable things: Throwaway paper and plastic sensors will connect everyday items. *IEEE Spectrum*, 55(12):30–35, 2018.

[GAvdS$^+$18] Carmine Garripoli, Sahel Abdinia, Jan-Laurens van der Steen, Gerwin Gelinck, and Eugenio Cantatore. A fully integrated 11.2mm^2 a-IGZO EMG front-end circuit on flexible substrate achieving up to 41db SNR and 29MΩ input impedance. *IEEE Solid-State Circuits Letters*, 1(6):142–145, 2018.

[GMSH10] Jan Genoe, Kris Myny, Soeren Steudel, and Paul Heremans. Design and manufacturing of organic RFID circuits: Coping with intrinsic parameter variations in organic devices by circuit design. In *2010 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*, pages 496–499, 2010.

[HFL$^+$11] Tsung-Ching Huang, Kenjiro Fukuda, Chun-Ming Lo, Yung-Hui Yeh, Tsuyoshi Sekitani, Takao Somey, and Kwang-Ting Cheng. Pseudo-CMOS: A design style for low-cost and robust flexible electronics. *IEEE Trans. Electron. Dev.*, 58(1):141–150, 2011.

[MGD16] Kris Myny, Jan Genoe, and Wim Dehaene. *Robust Design of Digital Circuits on Foil.* Cambridge University Press, 2016.

[MSR$^+$14] Kris Myny, Steve Smout, Maarten Rockelé, Ajay Bhoolokam, Tung Huei Ke, Soeren Steudel, Brian Cobb, Aashini Gulati, Francisco Gonzalez Rodriguez, Koji Obata, Marko Marinkovic, Duy-Vu Pham, Arne Hoppe, Gerwin Gelinck, Jan Genoe, Wim Dehaene, and Paul Heremans. A thin-film microprocessor with inkjet print-programmable memory. *Scientific reports*, 4:7398, 2014.

[MSVVH12] Hagen Marien, Michiel Steyaert, Erik Van Veenendael, and Paul Heremans. 1D and 2D analog 1.5 kHz air-stable organic capacitive touch sensors on plastic foil. In *IEEE International Solid-State Circuits Conference (ISSCC)*, pages 310–312. IEEE, 2012.

[MvVG$^+$12] Kris Myny, Erik van Veenendaal, Gerwin Gelinck, Jan Genoe, Wim Dehaene, and Paul Heremans. An 8-bit, 40-instructions-per-second organic microprocessor on plastic foil. *IEEE Journal of Solid-State Circuits*, 47(1):284–291, 2012.

[Myn18] Kris Myny. The development of flexible integrated circuits based on thin-film transistors. *Nature Electronics*, 1(1):30–39, 2018.

[NOF$^+$14] Manoj Nag, Koji Obata, Yusuke Fukui, Kris Myny, Sarah Schols, Peter Vicca, Tung Huei Ke, Steve Smout, Myriam Willegems, Marc Ameys, Ajay Bhoolokam, Robert Muller, Brian Cobb, Abhishek Kumar, Jan-Laurens van der Steen, Tim Ellis, Gerwin Gelinck, Jan Genoe, Paul Heremans, and Soeren Steudel. Flexible AMOLED display and gate-driver with self-aligned IGZO TFT on plastic foil. In *SID Symposium Digest of Technical Papers*, volume 45, pages 248–251. Wiley Online Library, 2014.

[NOT$^+$04] Kenji Nomura, Hiromichi Ohta, Akihiro Takagi, Toshio Kamiya, Masahiro Hirano, and Hideo Hosono. Room-temperature fabrication of transparent flexible thin-film transistors using amorphous oxide semiconductors. *Nature*, 432(7016):488, 2004.

[RLG14] Stewart Robson, Bosco Leung, and Guang Gong. Truly random number generator based on a ring oscillator utilizing last passage time. *IEEE Transactions on Circuits and Systems II: Express Briefs*, 61(12):937–941, 2014.

[TMH$^+$15] Ashutosh Kumar Tripathi, Kris Myny, Bo Hou, Kimberley Wezenberg, and Gerwin H Gelinck. Electrical Characterization of Flexible InGaZnO Transistors and 8-b Transponder Chip Down to a Bending Radius of 2 mm. *IEEE Transactions on Electron Devices*, 62(12):4063–4068, 2015.

[ZG14] Bo Zhu and Guang Gong. Multidimensional meet-in-the-middle attack and its applications to KATAN32/48/64. *Cryptography and Communications*, 6(4):313–333, 2014.