Insights into the Mind of a Trojan Designer

The Challenge to Integrate a Trojan into the Bitstream

Maik Ender*
Ruhr-Universität Bochum
Horst Görtz Institute for IT Security
Bochum, Germany
maik.ender@rub.de

Pawel Swierczynski*
Digital Society Institute, ESMT
Berlin, Germany
pawel.swierczynski@esmt.org

Sebastian Wallat*
University of Massachusetts
Amherst, MA, USA
swallat@umass.edu

Matthias Wilhelm
Ruhr-Universität Bochum
Horst Görtz Institute for IT Security
Bochum, Germany
matthias.wilhelm@rub.de

Paul Martin Knopp
Ruhr-Universität Bochum
Horst Görtz Institute for IT Security
Bochum, Germany
paul.knopp@rub.de

Christof Paar
Ruhr-Universität Bochum
Horst Görtz Institute for IT Security
Bochum, Germany
Christof.Paar@rub.de

ABSTRACT

The threat of inserting hardware Trojans during the design, production, or in-field poses a danger for integrated circuits in real-world applications. A particular critical case of hardware Trojans is the malicious manipulation of third-party FPGA configurations. In addition to attack vectors during the design process, FPGAs can be infiltrated in a non-invasive manner after shipment through alterations of the bitstream. First, we present an improved methodology for bitstream file format reversing. Second, we introduce a novel idea for Trojan insertion.

CCS CONCEPTS

• Security and privacy → Hardware reverse engineering; Malicious design modifications; Embedded systems security;

ACM Reference Format:

Maik Ender, Pawel Swierczynski, Sebastian Wallat, Matthias Wilhelm, Paul Martin Knopp, and Christof Paar. 2019. Insights into the Mind of a Trojan Designer: The Challenge to Integrate a Trojan into the Bitstream. In *Proceedings of ACM Asia South Pacific Design Automation Conference (ASP-DAC 2019)*, Jennifer B. Sartor, Theo D'Hondt, and Wolfgang De Meuter (Eds.). ACM, New York, NY, USA, Article 4, 8 pages.
<https://doi.org/10.1145/3287624.3288742>

1 INTRODUCTION

The threat of Intellectual Property (IP) theft, imposed by hardware reverse engineering, has been historically considered as the main practical security issue. The move from on-site fabrication to a globally distributed supply-chain and the arising threats of interdiction changed this perspective significantly for all kind of applications.

*All three authors contributed equally to the paper

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the Owner/Author.
ASPDAC '19, January 21–24, 2019, Tokyo, Japan
© 2019 Copyright is held by the owner/author(s).
ACM ISBN 978-1-4503-6007-4/19/01.
<https://doi.org/10.1145/3287624.3288742>

Since the Snowden's surveillance revelations, malicious hardware manipulations became an increasing concern, including SRAM-based Field Programmable Gate Arrays (FPGAs). Due to the volatile nature of SRAM-based FPGAs, new attack vectors arise such as bitstream interception and manipulation. A prerequisite for those kinds of attacks is to reverse engineer the bitstream file formats. For this reason and to support highly customisable bitstream generation tools, various research works [3, 5, 7, 8, 14–18, 24, 28] aimed at reverse engineering the proprietary bitstream file format of SRAM-based FPGAs, which mainly focused on Xilinx FPGAs.

However, today it is not possible to fully reverse the entire bitstream format of Xilinx FPGAs which reveals all the details of a specific hardware configuration. Hence, there is no official support for developing open source bitstream generation tools similar to Project IceStorm [27], which reversed the Lattice iCE40 FPGAs. Such a tool improves the flexibility for designers and researchers, i.e., it could extend (security) frameworks like *HAL* [6], *Torc* [19], or *RapidSmith* [9].

Knowing the entire bitstream file format, the security of cryptographic hardware configuration can be appropriately analyzed. Thus quick-and-easy malicious bitstream manipulation attacks [1, 4, 20–22], leading to a potential security breach, can be pentested beforehand and accordingly addressed by a security analyst. Defending of FPGA designs is even more crucial since most FPGA bitstream encryption schemes of older FPGA generations are vulnerable to side-channel attacks [10–13, 23] or do not offer any bitstream encryption/authentication at all. Hence the hardware layout reverse engineering and manipulation of bitstreams are a real threat. Notably, many old systems used in large infrastructures deploy hundreds of (older) FPGA devices. Newer hardware modules cannot simply replace them due to high costs or environmental reasons. Hence, considering the long life span [2] of (older) deployed SRAM-based FPGAs, it is always worthwhile to explore the practical doability of bitstream reverse engineering, hardware design reverse engineering, and the corresponding potential malicious hardware manipulations, cf. Wallat *et al.* [26]. All those methods need to be understood well, as they are crucial for improving the security of critical systems. Summarizing bitstream reverse engineering can be used for illegitimate and legitimate purposes.

The main issue of bitstream file format reverse engineering and a meaningful hardware configuration manipulation is a seemingly complicated and time-consuming practical task. In general, it is unclear to what extent an attacker can use a non-perfect converted netlist from a bitstream to inject a hardware Trojan into it.

In this work, we provide insights into bitstream reverse engineering techniques and Trojan insertion strategies at the hardware configuration level. Our contribution is as follows.

- (1) We present an improved methodology for bitstream file format reversing targeting its routing encoding. Moreover, it is capable of extracting the bitstream encoding rules for Programmable Interconnect Points (PIPs), Look-Up Tables (LUTs), and Flip Flops (FFs), which we exemplarily conducted for a Xilinx Spartan 6 FPGA. Additionally, our framework manipulates bitstreams, e.g., it can replace LUT configurations and set/unset single PIPs.
- (2) We introduce a new method for a hardware Trojan insertion into a self-test-protected AES IP core at the hardware configuration level. This gives an idea of how advanced attacker may perform malicious hardware configuration manipulations.

2 BACKGROUND

In this section, we briefly introduce the needed background on SRAM-based FPGAs, which is necessary to fully comprehend the bitstream reversing and manipulation methods presented within this paper. For an in-depth description of FPGAs, we refer the interested reader to [25].

The Spartan-6 has a two-dimensional array structure, where SLICES encounter the LUTs and FFs. Two SLICES and an adjacent switch matrix form one Configurable Logic Block (CLB), as this is depicted in Figure 1. The switch matrix realizes the interconnection logic of the FPGA by configuring so-called PIPs.

Figure 1: Part of configured FPGA internals showing a switch matrix, a NET, and a SLICE.

As indicated by Figure 1, PIPs are configurable wires within a switch matrix that connect static wires with other static wires. Hence, they allow building specific bridges between wires. Since the PIPs are reprogrammable, the information must be stored in

the proprietary bitstream files. Generally speaking, the PIPs realize the actual routing functionality of an FPGA.

Furthermore, a dedicated signal with a defined source and sink is called a NET. Sources and sinks are for example the LUT's input or FF's output. The PIPs realize the NETs' routing from a dedicated source to its sinks, i.e., an FF's output is routed via several PIPs to a LUT's input. Thus, when knowing all potential PIP configurations, the NETs can be reconstructed.

Finally, we shortly note that the bitstream is an encoded version of the placed and routed XDL netlist file. We refer to this information as hardware configuration.

3 BITSTREAM REVERSE ENGINEERING

We divide the reverse engineering process into two phases; the (1) bitstream reverse engineering and (2) bitstream conversion phase.

The first phase determines the relation between all bits in the bitstream and the associated hardware configuration of a specific FPGA model. The result of the first step is a database containing the mapping of a single bitstream's bit and its impact on the hardware configuration of a hardware primitive such as FFs, multiplexers, PIPs, or LUTs.

In the second phase, a targeted bitstream is converted back into a human-readable netlist representation by using the previously created database. The resulting gate-level netlist (or: XDL file) can then be processed by officially supported tools such as the ISE suite and FPGA Editor from Xilinx or by unofficially supported tools like *RapidSmith* [9]. With this step, the before unknown targeted bitstream's hardware configuration is revealed. Thus, it can be for example further processed by a reverse engineer.

The reverse engineering of the information that encodes the routing within a bitstream was already depicted in [5], but we provide an improved methodology for the Xilinx Spartan-6 Series. In particular, we were successful for the 6slx16csg324 and 6slx75csg484 FPGAs. Our approach simplifies and speeds up the entire bitstream reverse-engineering process. For reverse-engineering and verification purposes, we used the Xilinx ISE Suite 14.7.

3.1 Phase 1: Bitstream Reversing

Our tool for reversing the bitstream creates a database containing the mapping between most configurable FPGA resources and its configuration bits in the bitstream. In the following, we focus on the PIP reversing, but the methodology is also applicable to other FPGA components. Note that we reversed LUTs, flip-flops, and multiplexers as well, but we do not provide further details. Our approach works as follows.

- (1) Creation of a minimalistic template (XDL netlist) that either configures a single hardware elements, e.g., a PIP.
- (2) Generation of a reference bitstream which encodes supportive instances regarding the template of Step 1.
- (3) Conducting template alterations followed by bitstream generation and bit toggle observation.
- (4) Database creation.

To correlate the configuration of FPGA resources with the bitstream bits, we created specific crafted XDL netlist templates (Step 1) that are translated by the vendors tool to a bitstream (Step 2). Note that for each FPGA element, an individual template is needed, e.g.,

PIPs require a different template than LUTs. During Step 3, the configuration of the templates are slightly modified and another bitstream is generated, again using the vendors tool. The resulting change of both generated bitstream files are recorded into a database (Step 4). For bitstream creation, we used the Xilinx tools *xdl* and *bitgen* as follows.

- `xdl -xdl2ncd -force empty.xdl // creates an NCD file`
- `bitgen -d empty.ncd // creates the corresponding bitstream`

By using the *-force* and *-d* option, this enables us to generate bitstreams even though the hardware configuration itself is an illegitimate design.

Note that some targeted hardware elements require the configuration of other hardware elements. Otherwise, the vendor's tools do not manipulate the bitstream files accordingly. This is an issue that we faced ourselves when we tried to reverse engineer the bitstream encoding for one PIP configuration. It is not sufficient to only configure a PIP within a NET, but one needs to attach the NET to other instances. Listing 1 exemplary shows an XDL template that generates a valid bitstream containing information about the PIP configuration.

```

1 design "basic_pip" xc6slx16csg324-3 v3.2;
2 inst "q" "IOB",placed BIOB_X11Y0 T8;
3 inst "q_OBUF" "OLOGIC2",placed BIOI_OUTER_X11Y0 OLOGIC_X8Y0;
4 net "q_OBUF" ,
5   outpin "q_OBUF" OQ ,
6   inpin "q" O ,
7   pip INT_X10Y6 WW2E2 -> NL1B2 ;

```

Listing 1: XDL template for one PIP configuration (WW2E2 → NL1B2) of the switch matrix INT_X10Y16

We first reverse engineer the routing encoding of one switch matrix and later on simplify the process for all other available ones. In our example (cf. Listing 1), we target the switch matrix labeled as INT_X10Y16. In line 4, a NET is instantiated containing the examined PIP (WW2E2 → NL1N2) in line 7. To our surprise, we found out that it is just necessary to specify two *arbitrary chosen* instances (line 2 and 3) that are supposed to be connected with the NET's OUTPIN and INPIN, even if the NET does not route such connectivity. Hence, there is no need to worry about finding and connecting the NET to valid instances, which eases the entire reverse engineering task significantly.

Additionally, thanks to the *-force* and *-d* parameters, the bitstream generation tools nevertheless encode the PIP's information into the resulting bitstream even though our specified instance and NET connectivity are not necessarily correct. If the two arbitrary chosen instances are missing in the XDL template, the PIP's configuration is not encoded into the bitstream file even though the specified NET itself is valid.

Since we instantiate two additional instances which are not associated with the PIP, its configuration is undesirably encoded into every generated bitstream file. This information can be easily removed by letting the vendors tool generate a reference bitstream file (only once) that does only encode both instances (line 2 and 3) without the NET in Listing 1, i.e., Step 2.

The difference between any generated bitstream (instances plus NET with individual PIP) and the one-time generated reference bitstream (with instances only) reveals the correct bitstream encoding of an individual PIP configuration. Hence, both bitstreams simply need to be XORed.

Which PIPs are available for a switch-matrix configuration, can be figured out quickly by parsing the information from so-called XDL report files that can be easily generated, cf. [3]. For every possible PIP the template is changed, e.g., Line 7 is replaced with another examined PIP and the corresponding bitstream is generated.

Further note that in cases where a PIP configuration does not lead to any bit toggle in the bitstream, we mark it as a default configuration. During our reverse engineering efforts, this only happened for a fraction of all PIPs (under 1% of all PIPs).

The presented PIP reversing approach of Ding *et al.* [5] relies on *i)* creating NETs with multiple PIPs and on *ii)* conducting further pre- and postprocessing steps. As opposed to that, our approach minimizes the required pre- and postprocessing steps and hence can be carried out faster and in a less complicated manner. It also does require less knowledge about the FPGA internals.

Note that by applying the described approach of [5], we can also reduce the time for reversing the bitstream and the file size of the database. This works by exploiting the repetitiveness of hardware elements distributed over the FPGA's grid structure.

Most switch matrices are of the same type, i.e., they contain the same labeled PIPs. Moreover, the byte distances among the configuration bits from different PIPs are always equally distributed within one switch matrix type. Exemplarily, the configuration bits of the first and second PIP could always be separated by k byte positions, which would apply to (most) switch matrices. This was described by Ding *et al.* [5].

We were able to verify this, as we have also reversed various switch matrices with different locations on the FPGA grid for verification purposes. As expected, the bitstream encoding yielded the same distances. Additionally, the routing of all our tested FPGA designs could be later on correctly recovered. A sample design recovery for the routing of an AES IP core is depicted in Figure 2. Our recovered netlist only lacks the clock tree information which is work in progress. Otherwise, the recovered routing information is complete.

Once all PIPs of a single switch matrix are reverse engineered (revealing all distances), it is sufficient to only generate one bitstream for every remaining unreversed switch matrix in the FPGA. Each of those bitstreams encodes the configuration of a chosen fixed reference PIP.

To calculate the positions of the remaining unreversed PIPs of a switch matrix, the offset of the fixed reference PIP is added with the corresponding previously derived PIP distance. Each PIP position can be hence simply computed with

$$PIP_position = reference_PIP_position + PIP_distance \quad (1)$$

Given that we have to reverse-engineer one switch matrix with 3461 possible PIPs (distance PIPs) and there are only 2278 switch matrices (reference pips), we just need to generate $5739 = 3461 + 2278$ bitstreams to derive a complete list of most available PIPs.

A straight-forward attempt to reverse most PIPs of all switch matrices individually would make the required reverse-engineering

Figure 2: Routing comparison between the original and the converted AES IP core.

time practically infeasible: when considering that a bitstream generation plus its processing (single threaded) takes ~ 7 seconds on an average machine (i7-7700HQ, 3.8GHz, 4 cores, 8 threads) and that there are approx. ~ 5.5 million possible PIPs on a Spartan-6 (6slx16), the sequential bitstream generation of all 5.5 million bitstreams would take approximately ~ 381 days.

As opposed to that, it took us only approximately 21.5 hours to reverse engineer the entire routing encoding by following the distance approach of Ding. We also implemented a parallelized version with 8 threads and could further reduce the reverse engineering time to ~ 2.6 hours.

Once we have generated all bitstreams, where each bitstream encodes the configuration of a chosen fixed reference PIP for one switch-matrix, we derive a database (Step 4) that stores all those reference byte positions. Additionally, we store the distance patterns for all PIPs that can later on be used to reconstruct the exact byte and bit positions (Equation 1) for any queried switch matrix.

3.2 Phase 2: Bitstream Conversion

For Phase 2, we have created a bitstream converter and manipulation framework which uses the previously generated database. It is capable of converting a targeted bitstream back to its partial XDL netlist representation, e.g., routing, flip-flops, MUXs, and LUTs. Additionally, it is capable of modifying LUTs and PIPs directly in the bitstream, e.g., it can manipulate the Boolean LUT equations and set/unset PIPs in any arbitrary switch matrix. Figure 3 shows the architecture of our developed tool, which is written in C++. During development, we used a modular architecture allowing to support

new FPGA devices and features, e.g., to support the conversion of BRAM or IO blocks.

Figure 3: Architecture of our bitstream framework.

We describe all modules as follows.

- The database wrapper parses the previously created database and provides an interface to it. It stores all information regarding the mapping of bitstream bits to a hardware configuration.
- Similarly, the bitstream wrapper parses the targeted - to be converted - bitstream. The targeted bitstream is XORed with a generated reference bitstream that encodes an empty hardware configuration. This eliminates unwanted default configuration bits.
- The hardware configuration information, e.g., a list of PIP objects, is stored into the hardware configuration wrapper, which can be accessed by the XDL writer.

- The XDL writer is responsible for generating valid XDL code from the available information of the hardware configuration wrapper.

In the following, we discuss the bitstream reverser class in greater detail, as it is mainly responsible for correctly extracting the hardware configuration from a targeted bitstream. Algorithm 1 describes the entire extraction workflow.

```

Data: bitstreamWrapper* bs, database* db, configurationWrapper* c
begin
  /* Iterate over every bit of the bitstream */
  for bitPos = bs->getSyncWordPos() to bs->bitstreamLength() do
    /* Check if the bit is set and marked */
    if bs->getBit(bitPos) == 1 and !marked then
      /* Search for a suitable object in the db */
      DataObject* obj ←
        db->GetElementFromDatabase(bitPos);
      if obj != nullptr then
        /* If an object is found, process it
           further depending on its type */
        type ← obj->getType();
        if type == PIP then
          /* Reverse the PIP */
          mostBitsToggled ← 0;
          usedPIP ← nullptr;
          for PIP p : db->getLinkedPIPs(bitPos) do
            if bs->bitsToggled(p->getToggledBits())
              then
                if mostBitsToggled <
                  p->getNumToggledBits() then
                  mostBitsToggled ←
                    p->getNumToggledBits();
                  usedPIP ← p;
                end
              end
            end
          end
          if usedPIP != nullptr then
            /* set usedPIP in the config */
            c->addPIP(p);
          end
        else if type == lutBit then
          /* Reverse the LUT Bit */
        end
      end
    end
  end
end

```

Algorithm 1: Bitstream conversion algorithm needed for extracting PIP configurations from a targeted bitstream.

As can be seen, the algorithm iterates over every single set bit from the targeted bitstream for which the bitstream wrapper removed the disturbing default bits in a previous processing step.

Depending on the bit's position, our database assigns it to the correct bits' object type, e.g., PIPs, MUXs, or LUTs as well as to the correct location on the two-dimensional FPGA grid. Then it stores the returned information in the hardware configuration wrapper.

Since the reconstruction of LUTs is for example already explained in [16, 22, 28], we only describe the reconstruction of PIPs in further detail, as its conversion to a correct PIPs' object list requires appropriate processing.

A switch matrix of an FPGA contains multiple sinks and sources. From now on, we only consider one sink, e.g., a PIP. Depending on the type, N configurable sources are wired with the PIP, cf. Fig. 4, but only one valid PIP connection can be routed through the sink at the same time, cf. Figure 4.

Figure 4: Upper part of a switch matrix

Roughly speaking, a PIP acts like a multiplexer. The configuration of which source is routed through the PIP is stored in the bitstream. Such an N -to-1-multiplexer requires at least $\lceil \log_2(N) \rceil$ control bits, but we observed that the bitstream stores more than $\lceil \log_2(N) \rceil$ bits (without further explanation). The fact that a single PIP share multiple configuration bits and these bits are not stored continuously complicates the bitstream conversion task. To address this, we have developed Algorithm 1.

Figure 5: Possible distribution of PIP's configuration bits within a bitstream. If all arrows of a PIP point to only set bits and if it has the most bits set, the corresponding PIP is configured. If any outgoing arrow of a PIP points to at least one cleared bit, then the PIP is obviously not configured.

As indicated before, our tool iterates over each bit in the targeted pre-processed bitstream. If the algorithm encounters a PIP configuration bit, it yields a set of possible PIPs and iterates over each PIP candidate. It then successfully figures out, which PIP is the correct candidate, cf. the caption of Figure 5 or Algorithm 1.

In the shown example of Figure 5, our tool processes the most-left set PIP configuration bit. In this case, the database will return PIP A, PIP B, and PIP C as potential configured PIPs. Since one bit for PIP C is cleared, our algorithm discards this candidate. As a valid PIP requires to have set all configuration bits. Even though this is the case for PIP B, it will be discarded as well, since all configuration bits of PIP A are set and the Hamming weight of PIP A is the largest one. Therefore, in this toy example PIP A is the correct PIP candidate.

Additionally, once we have reconstructed all configured PIP bits within a bitstream, we convert back all set NETs. For this purpose, we have re-implemented the described ICG algorithm from Ding *et al.* [5].

3.3 Bitstream Manipulator

Note that we started to implement a bitstream manipulator to avoid the need for a complete bitstream conversion. So far, our framework is capable of setting or unsetting PIPs and changing Boolean functions in LUTs. This allows us to carry out targeted bitstream manipulations for which we do not possess the fully reverse engineered XDL netlist. Therefore, we do not rely on the vendor's tools, e.g., xdl and bitgen. They usually require a complete XDL netlist to ensure the intended functionality of a hardware configuration. In the future, we plan to enhance our bitstream manipulation capabilities further.

Having introduced our capabilities, we now describe how a Trojan may be inserted into a targeted selftest-protected AES IP core.

4 CASE STUDY: TROJAN INSERTION

Only little is known about how a Trojan designer may exactly proceed to tamper a third-party hardware configuration. This case study demonstrates a possible Trojan insertion methodology into an Advanced Encryption Standard (AES) IP core at hardware configuration level. Note that in practice such Trojan requires to either develop or use a tool that can

- convert the bitstream file format to a gate-level netlist, cf. Section 1 and 3
- manipulate the underlying hardware primitives without violating any timing constraints.
- perform correct bitstream file patching

In the field, a bitstream manipulation is possible during various life-cycle phases of a device, i.e., it can be intercepted during shipment, during selling, or even during operation from the customer himself. Even though we did not test the following manipulations directly on the bitstream, we provide a new idea of how an attacker can proceed to manipulate a placed and routed hardware configuration where only partial information is available, i.e., we assume that a third-party bitstream was partially converted. This is a realistic assumption¹ for most deployed embedded devices, where the bitstream is usually stored in plaintext on the same PCB along with the targeted FPGA. We used the *HAL* framework [6] to identify and manipulate the relevant netlist components. From our result, we are confident that a hardware configuration must not be fully

¹It is likely that state actors are able to obtain all data from a bitstream, but we nevertheless try to explore the actual practical capabilities of weaker attackers who only possess partial information

known to also accomplish more complicated Trojan insertions at a very late stage.

4.1 Target and Notation

Our targeted AES-128 IP core provides an interface to set a key for data decryption and encryption. To interact with the circuit we integrated an Universal Asynchronous Receiver Transmitter (UART)/RS-232 interface. Further, we use the following notations: p - Plaintext (16 bytes), k - Key (16 bytes), $c = \text{AES}_k(p)$ - Ciphertext (16 bytes), $(p_{\text{ref}}, c_{\text{ref}})$ - Plaintext/ciphertext pair for the self-test, k_{st} - Key for the self-test, k_u - Key for user data.

4.2 System Model

We assume that an FPGA and an external entity, e.g., a dedicated microprocessor, are integrated on the same embedded system along with an FPGA-based AES accelerator in a bitstream. The workflow is as follows:

- (1) After the FPGA is configured with the AES IP core, the external entity conducts a self-test with the FPGA by setting the self-test key k_{st} and the reference plaintext p_{ref} .
- (2) The external entity analyzes the FPGA's computed ciphertext $c = \text{AES}_{k_{\text{st}}}(p_{\text{ref}})$ and verifies the integrity by comparing if c is equal to the reference ciphertext c_{ref} .
- (3) In case the self-test is successful, an AES key k_u is derived after a user legitimates to unblock his encrypted device. A key derivation is computed and the user key is automatically passed to the FPGA.

Note that such an embedded system is not only a theoretical assumption, since it is similar to the FIPS-140-2 level 2 certified USB flash drive from Kingston, cf. the work of Swierczynski *et al.* [21].

4.3 Adversary's Goal and Trojan Idea

The high-level goal of the adversary is to permanently burn the key k_{st} into the FPGA hardware configuration, so that the user-data is always encrypted/decrypted with k_{st} regardless of which key k_u is set by the user. This way, the adversary can, later on, decrypt all seemingly securely encrypted user data. To accomplish this goal, the adversary targets to attach a payload circuit to the existing set-key circuit to permanently write the known key bits of k_{st} into the FFs that process k . Thus, the external self-test falsely confirms the integrity of the FPGA AES IP core. To sum up, the described Trojan can trick an external self-test in cases where the self-test key is known like in [21].

4.4 Detection of Set-Key Circuitry

For our targeted AES core, we noticed that eight 16-bit shift-registers were integrated into the hardware configuration. Its task is to temporarily store the key received sequentially via UART and is later on forwarded to the AES IP core, cf. Figure 6. For an attacker, the key registers are an appealing target, as its alteration can, for example, enable key extraction or Trojan insertion. Since our goal is to burn a fixed key into the design, we first extract the relation between a FF and the AES key bit and second override the corresponding parts of the hardware configuration that delivers the (usually correct) key bit values. For that purpose, we developed

a plugin to algorithmically detect deployed shift-registers and its corresponding length, i.e., how many FFs are cascaded, within a netlist.

Figure 6: 8 observed parallel 16-bit shift-registers sampling one key byte from a repeatedly updated 8-bit buffer.

In order to reverse-engineer the relation between each FF and key bit, we manipulate/clear the first pass-through LUT (in this case it was always present, not shown in Figure 6) of a found shift-register in the hardware configuration. By means of simulation², we let the hardware configuration compute a ciphertext for a known plaintext and an AES key for which all 128 key bits are set. By comparing the resulting FPGA ciphertext output with various pre-computed known plaintext/ciphertexts pairs for suitable keys³, we can conclude the relation of each shift-register and the key bit position of a key byte it processes. Note that we similarly determined whether the first or last flip-flop of a shift-register store the key bit of the first or last key byte. Having figured out which hardware primitive is responsible for processing one segment of the AES key, we can now describe the Trojan payload.

4.5 Manipulating a Set-Key Circuitry

The key idea is to detach any signal from the data input pin D of each key bit FF and to attach the output signal from the so-called *payload LUTs*, cf. Figure 6, so that all key-based shift-registers are independent from the key buffer.

For the communication with our device an RS232 interface is used, which receives the key byte by byte. Eight shift-registers are used to store the bits of the current key byte during the loading phase. Each shift-register itself can hold up to 16 values to store all 8-bit chunks of a 128-bit key (*key_byte* in Figure 6). In total, we added 128 payload LUTs and routed them to their target FFs, using *HAL*. Of course alternative approaches may be used like for example attaching GND or VCC signals to the corresponding flip-flops inputs, but so far we did not reverse-engineer the corresponding bitstream encoding for those hardware elements. Each payload LUT is programmed individually to output either 0 or 1 (depending on the

key bit k_{st}). Overall, the functionality of the $AES_k(\cdot)$ is substituted by the malicious one, i.e., by $AES_{k=k_{st}}(\cdot)$.

Figure 7: Key-based Shift-register after IP manipulation.

We verified the correctness of our Trojan by setting a random key $k \neq k_{st}$ and testing that the manipulated AES core nevertheless computes $AES_{k_{st}}(\cdot)$ when encrypting one plaintext.

4.6 Stealthiness

Considering this Trojan, the self-test is rendered useless if only one key k_{st} is tested by the external entity. Furthermore, all user data is encrypted with the known key k_{st} instead of the derived key k_u making the decryption entirely possible. Note that in case of an embedded device such as the Universal Serial Bus (USB) flash drive [21], the user usually has no direct access to k_u unless he attempts to derive k_u by himself, which requires exact knowledge of the key derivation function. Certainly, once the user has access to k_u and obtains one plaintext/ciphertext pair, this Trojan can be easily detected.

5 CONCLUSION AND FUTURE WORK

In this paper, we demonstrated further attackers' capabilities based on the example of a malicious manipulation of a self-test protected AES core. We improved the bitstream reverse engineering methods by simplifying known routing extraction mechanisms. Thus we conclude that the efforts for a successful bitstream conversion are even lower than commonly assumed. Note that bitstream reversing is also an essential step for legitimate purposes such as Trojan detection, IP theft exposure, bitstream verification, or advanced bitstream tooling. Further, we demonstrated the implication of partial bitstream exposure. Combined with our Trojan case study, we emphasized that key-based shift-registers can be exploited. The identification of such points of interest is usually one of the various practical hurdles during reverse engineering. For future projects, one may also analyze how to algorithmically identify key registers with more complicated hardware structures and whether partial bitstream reverse engineering is sufficient to carry out similar attacks. Consequently, future research should explore defense mechanisms for such vulnerable hardware structures, e.g., hardware obfuscation methods.

²Note that one may also manipulate the corresponding change in the bitstream and observe the resulting ciphertexts through querying the AES module with several known plaintexts. This way one can also conclude the relation between flip-flops and key bits without conducting any simulation. For the sake of simplicity, we used simulation.

³For example, we computed look-up tables with plaintext/ciphertext pairs for a key, where the last bit of each of 16 key bytes was cleared, while all other key bits were set.

6 ACKNOWLEDGEMENT

Part of this work was supported by the European Research Council (ERC) under the European Union's Horizon 2020 Research and Innovation programme (ERC Advanced Grant No. 695022 (EPoCH)) and by NSF grant CNS-1563829.

REFERENCES

- [1] A. C. Aldaya, A. J. C. Sarmiento, S. Sánchez-Solano. 2015. AES T-Box tampering attack. *Journal of Cryptographic Engineering* (2015), 1–18. <https://doi.org/10.1007/s13389-015-0103-4>
- [2] Altera. 2015. ALTERA Annual Report for Form 10-K for 2014. <https://www.sec.gov/Archives/edgar/data/768251/000076825115000008/altera10k12312014.htm>.
- [3] F. Benz, A. Seffrin, and S.A. Huss. 2012. BIL: A tool-chain for bitstream reverse-engineering. In *Field Programmable Logic and Applications (FPL)*. IEEE, 735–738.
- [4] R.S. Chakraborty, I. Saha, A. Palchoudhuri, and G.K. Naik. 2013. Hardware Trojan Insertion by Direct Modification of FPGA Configuration Bitstream. *Design Test, IEEE* 30, 2 (April 2013), 45–54.
- [5] Zheng Ding, Qiang Wu, Yizhong Zhang, and Linjie Zhu. 2013. Deriving an NCD file from an FPGA bitstream: Methodology, architecture and evaluation. *Microprocessors and Microsystems - Embedded Hardware Design* 37, 3 (2013), 299–312.
- [6] M. Fyrbiak, S. Wallat, P. Swierczynski, M. Hoffmann, S. Hoppach, M. Wilhelm, T. Weidlich, R. Tessier, and C. Paar. 2018. HAL- The Missing Piece of the Puzzle for Hardware Reverse Engineering, Trojan Detection and Insertion. *IEEE Transactions on Dependable and Secure Computing* (2018), 1–1. <https://doi.org/10.1109/TDSC.2018.2812183>
- [7] Steve Guccione, Delon Levi, and Prasanna Sundararajan. 2011. JBits: Java based interface for reconfigurable computing. In *CCS 2011*. ACM.
- [8] Edson L. Horta, John W. Lockwood, and Sérgio T. Kofuji. 2002. *Using PARBIT to Implement Partial Run-Time Reconfigurable Systems*. Springer Berlin Heidelberg, Berlin, Heidelberg, 182–191. https://doi.org/10.1007/3-540-46117-5_20
- [9] Christopher Lavin, Marc Padilla, Philip Lundrigan, Brent E. Nelson, and Brad L. Hutchings. 2010. Rapid prototyping tools for FPGA designs: RapidSmith. In *Proceedings of the International Conference on Field-Programmable Technology, FPT 2010, 8-10 December 2010, Tsinghua University, Beijing, China*. 353–356. <https://doi.org/10.1109/FPT.2010.5681429>
- [10] A. Moradi, A. Barengi, T. Kasper, and C. Paar. 2011. On the vulnerability of FPGA bitstream encryption against power analysis attacks: Extracting keys from Xilinx Virtex-II FPGAs. In *CCS 2011*. ACM, 111–124.
- [11] A. Moradi, M. Kasper, and C. Paar. 2012. Black-Box Side-Channel Attacks Highlight the Importance of Countermeasures - An Analysis of the Xilinx Virtex-4 and Virtex-5 Bitstream Encryption Mechanism. In *CT-RSA 2012 (LNCS)*, Vol. 7178. Springer, 1–18.
- [12] Amir Moradi, David Oswald, Christof Paar, and Pawel Swierczynski. 2013. Side-channel Attacks on the Bitstream Encryption Mechanism of Altera Stratix II: Facilitating Black-box Analysis Using Software Reverse-engineering. In *ACM/SIGDA International Symposium on Field Programmable Gate Arrays*. ACM, 91–100. <https://doi.org/10.1145/2435264.2435282>
- [13] A. Moradi and T. Schneider. 2016. Improved Side-Channel Analysis Attacks on Xilinx Bitstream Encryption of 5, 6, and 7 Series. In *Workshop on Constructive Side-Channel Analysis and Secure Design*. Springer.
- [14] Jean-Francois Nguyen. 2016. Analysing the Bitstream of Altera's MAX-V CPLDs. https://lse.epita.fr/lse-summer-week-2016/slides/lse-summer-week-2016-07-maxv_cpld.pdf
- [15] Jean-Baptiste Note. 2008. debit. <https://github.com/djn3m0/debit/tree/master/altera>
- [16] Jean-Baptiste Note and Éric Rannaud. 2008. From the bitstream to the netlist. In *Proceedings of the ACM/SIGDA 16th International Symposium on Field Programmable Gate Arrays, FPGA 2008, Monterey, California, USA, February 24-26, 2008*. 264. <https://doi.org/10.1145/1344671.1344729>
- [17] K. Dang Pham, E. Horta, and D. Koch. 2017. BITMAN: A Tool and API for FPGA Bitstream Manipulations. In *Design, Automation Test in Europe Conference Exhibition (DATE)*, 2017. 894–897. <https://doi.org/10.23919/DATE.2017.7927114>
- [18] A. K. Raghavan and P. Sutton. 2002. JPG - a partial bitstream generation tool to support partial reconfiguration in virtex FPGAs. In *Proceedings 16th International Parallel and Distributed Processing Symposium*. 6 pp–. <https://doi.org/10.1109/IPDPS.2002.1016552>
- [19] Neil Steiner, Aaron Wood, Hamid Shojaei, Jacob Couch, Peter Athanas, and Matthew French. 2011. Torc: towards an open-source tool flow. In *Proceedings of the ACM/SIGDA 19th International Symposium on Field Programmable Gate Arrays, FPGA 2011, Monterey, California, USA, February 27, March 1, 2011*. 41–44. <https://doi.org/10.1145/1950413.1950425>
- [20] Pawel Swierczynski, Georg T. Becker, Amir Moradi, and Christof Paar. 2017. Bitstream Fault Injections (BiFI) - Automated Fault Attacks against SRAM-based FPGAs. *IEEE Trans. Comput. PP*, 99 (January 2017), 1–1. <https://doi.org/10.1109/TC.2016.2646367>
- [21] Pawel Swierczynski, Marc Fyrbiak, Philipp Koppe, Amir Moradi, and Christof Paar. 2016. Interdiction in Practice - Hardware Trojan Against a High-Security USB Flash Drive. *Journal of Cryptographic Engineering* (June 2016), 1–13. <https://doi.org/10.1007/s13389-016-0132-7>
- [22] Pawel Swierczynski, Marc Fyrbiak, Philipp Koppe, and Christof Paar. 2015. FPGA Trojans Through Detecting and Weakening of Cryptographic Primitives. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 34, 8 (August 2015), 1236–1249. <https://doi.org/10.1109/TCAD.2015.2399455>
- [23] Pawel Swierczynski, Amir Moradi, David Oswald, and Christof Paar. 2014. Physical Security Evaluation of the Bitstream Encryption Mechanism of Altera Stratix II and Stratix III FPGAs. *ACM Trans. Reconfigurable Technol. Syst.* 7, 4, Article 34 (January 2014), 34:1–34:23 pages. <https://doi.org/10.1145/2629462>
- [24] SymbiFlow. 2017. Project X-Ray. <https://github.com/SymbiFlow/prjxray>
- [25] Russell Tessier, Kenneth L. Pocek, and André DeHon. 2015. Reconfigurable Computing Architectures. *Proc. IEEE* 103, 3 (2015), 332–354. <https://doi.org/10.1109/JPROC.2014.2386883>
- [26] S. Wallat, M. Fyrbiak, M. Schlögel, and C. Paar. 2017. A look at the dark side of hardware reverse engineering - a case study. In *2017 IEEE 2nd International Verification and Security Workshop (IVSW)*. 95–100. <https://doi.org/10.1109/IVSW.2017.8031551>
- [27] Clifford Wolf and Mathias Lasser. [n. d.]. Project IceStorm. <http://www.clifford.at/icestorm/>.
- [28] D. Ziener, S. Assmus, and J. Teich. 2006. Identifying FPGA IP-Cores Based on Lookup Table Content Analysis. In *Field Programmable Logic and Applications, 2006. FPL '06. International Conference on*. 1–6. <https://doi.org/10.1109/FPL.2006.311255>