Formal Verification of Arithmetic Masking in Hardware and Software

Barbara Gigerl[1], Robert Primas[1], and Stefan Mangard[1,2]

[1] Graz University of Technology, Graz, Austria
[2] Lamarr Security Research, Graz, Austria

Abstract. Masking is a popular countermeasure to protect cryptographic implementations against physical attacks like differential power analysis. So far, research focused on Boolean masking for symmetric algorithms like AES and Keccak. With the advent of post-quantum cryptography (PQC), arithmetic masking has received increasing attention because many PQC algorithms require a combination of arithmetic and Boolean masking and respective conversion algorithms (A2B/B2A), which represent an interesting but very challenging research topic. While there already exist formal verification concepts for Boolean masked implementations, the same cannot be said about arithmetic masking and accompanying mask conversion algorithms.

In this work, we demonstrate the first formal verification approach for (any-order) Boolean and arithmetic masking which can be applied to both hardware and software, while considering side-effects such as glitches and transitions. First, we show how a formal verification approach for Boolean masking can be used in the context of arithmetic masking such that we can verify A2B/B2A conversions for arbitrary masking orders. We investigate various conversion algorithms in hardware and software, and point out several new findings such as glitch-based issues for straightforward implementations of Coron et al.-A2B in hardware, transition-based leakage in Goubin-A2B in software, and more general implementation pitfalls when utilizing common optimization techniques in PQC. We provide the first formal analysis of table-based A2Bs from a probing security perspective and point out that they might not be easy to implement securely on processors that use of memory buffers or caches.

Keywords: Side-Channel Attacks · Arithmetic Masking · Formal Verification

1 Introduction

Passive side-channel attacks, including power or electromagnetic analysis, are among the most relevant attack vectors against cryptographic devices like smart cards, that are physically accessible by an attacker [47, 57]. A commonly used approach to protect against these attacks is to implement algorithmic countermeasures, for example masking [18, 40, 41, 44, 58]. Masking schemes split input and intermediate values of cryptographic computations into $d+1$ random shares such that observations of up to d shares do not reveal any information about the native (unmasked) value. Boolean masking, where native values correspond to the XOR-sum over its shares, have received much attention since such schemes are applicable to almost all symmetric cryptographic algorithms. On the other hand, arithmetic masking schemes have also gained increased importance, especially with the advent of PQC, for which they generally represent a better fit. In arithmetic masking, native values correspond to the arithmetic addition over their shares which allows to express operations like addition and subtraction much more efficiently than with Boolean masking. Since PQC algorithms often use symmetric building blocks, e.g. to achieve CCA2-security or for sampling random numbers, arithmetic masking often has to be combined with Boolean masking [14, 30, 60], which requires efficient and secure A2B/B2A conversion techniques. Many works have shown that hardware side-effects like glitches or transitions can violate the security of masking schemes in practice. Hence, the design of masked cryptography requires a detailed understanding of the targeted hardware platform, and is therefore a notoriously error-prone and time consuming task. Consequently, there is strong need for verification tooling that supports this effort to the highest possible extend.

While there already exists a vast amount of literature on the verification of Boolean masking, including formal verification approaches like Rebecca [13], maskVerif [3], Coco (Alma) [36,42], SILVER [46], or scVerif [6], the same cannot be said about arithmetic masking.

Limitations of Existing Approaches The first works on formal verification of arithmetic masking schemes were published with QMVerif by Gao et al. [33] and `LeakageVerif` by Meunier et al. [50]. These works already form a good foundation, but are limited in several ways.

QMVerif was published in 2019 by Gao et al. [31] for the verification of first-order Boolean and arithmetically masked software. QMVerif uses *type inference* to determine the distribution of every variable in the masked software, which is either uniform, independent of private inputs or dependent on private inputs. Due to the lack of completeness guarantees of type inference, deducing the distribution might not always be possible and might lead to false positives [50]. In that case, QMVerif uses *model-counting* to compute the exact distributions using a SAT solver, which is complete, but does not scale and consequently often requires significant computational resources such as GPU acceleration [32]. This scalability issue leads to the conclusion that model-counting-based masking verification is generally infeasible in the context of arithmetic masking. Besides that, QMVerif does not support masked hardware and heavily restricts supported masked software, e.g. by requiring a specific high-level syntax, not allowing branches, loops or functions and limiting variables to 8 bit. The leakage model of QMVerif hence also does not consider hardware side-effects like glitches and transitions, and it is unclear wether QMVerif can be applied to A2B/B2A conversions without a power-of-two-modulus. The same authors later propose HOME for higher orders following the same approach, but do not evaluate it for higher-order arithmetic masking. Since the tool is not (yet) open-source, it is not possible to investigate its further functionality.

Meunier et al. [50] propose LeakageVerif, a Python verification library based on *substitution* [4], which tries to show that an expression is leakage-free if it can be divided into sub-expresions, which can iteratively be substituted by fresh random variables. The evaluation shows that LeakageVerif is more efficient than QMVerif, but it is not complete and fails to verify common A2B/B2A conversions such as Goubin-A2B [39] and Coron et al.-B2A. LeakageVerif works for first-order implementations only, does not consider glitches, table lookups, or moduli which are not a power of two.

In general, both QMVerif and LeakageVerif are *sound* (leakages are never missed), but can sometimes only achieve *completeness* (leaks are only reported if they really exist) if they fall back to expensive and inefficient model-counting.

Other existing verification tools focus exclusively on Boolean masking and often perform exact model-counting, which are therefore unlikely to be applicable to arithmetic masking. For example, maskVerif has been shown infeasible in this context by several works [32,34,50]. In 2021, Bos et al. present scVerif [6], which was later modified for the verification of a first-order arithmetically masked software implementation of Kyber [14]. However, scVerif was not evaluated for other arithmetically masked programs, so no general statement about its efficiency or accuracy can be made. It does consider hardware side-effects but only if they have been identified in prior empirical experiments, which means the method is not sound and binds the evaluation stronger to the microarchitecture, while leaving no potential for masked hardware. We expect SILVER [46] to also not be able to deal with the complexity of arithmetic expressions since it exclusively tracks exact distributions with the help of binary decision diagrams.

In 2018, Bloem et al. suggest to approximate Fourier coefficients of Boolean functions [13] as a way to perform cheaper model counting that achieves soundness but not completeness. The resulting approach was evaluated for Boolean masked hardware (Rebecca), and later for software on CPU netlists (Coco) [36,42], and has shown to be efficient with a relatively low rate of false positives. However, it was not evaluated for arithmetic masking in terms of efficiency, accuracy, and general applicability for PQC-relevant use cases.

Our Contribution We improve this situation by demonstrating that the security of arithmetically masked software/hardware can be efficiently verified using verification approaches tailored to Boolean masking. More concretely, we provide the following contributions:

- We show how verification methods based on approximated Fourier coefficients of Boolean functions (as used by Rebecca/Coco) can be efficiently applied in the context of arithmetic masking. The resulting verification approach can successfully be applied to both masked hardware and software written in Assembly language. Its soundness is sufficient for many PQC/ARX applications. This approach is also the first to consider physical defaults (glitches, transitions) and the first to be evaluated for higher orders in the context of arithmetic masking (Section 3).

- In case of hardware implementations, we analyze different versions of the Coron et al.-A2B/B2A [23] conversion algorithms and identify potential weaknesses caused by glitches. We then present a proof-of-concept implementation that is secured against glitches and can be fully verified using our approach (Section 4).
- In the context of software implementations, we analyze various popular A2B/B2A conversion algorithms using power of two or prime moduli and provide new insights on implementation aspects that can reduce their protection order. More concretely, we report new findings of transition leakages in Goubin-A2B [39] and point out more general pitfalls when using lazy-reduction techniques in the context of masking. Additionally, we are the first to investigate architecture side-effects of table-based A2Bs and discuss why they might not be easy to implement securely on processors that make use of memory buffers or caches. Last but not least, we also show applicability of this approach in the context of symmetric cryptographic schemes by verifying the security of masked software implementations of one round of Speck and the ARX-box Alzette(Section 5).
- We plan to publish the software and hardware implementations on Github[3].

2 Background

In this section, we cover necessary background on masking, and A2B/B2A conversion techniques. Since our approach is based on Rebecca/Coco, we briefly describe the verification concept and the applied adversary model.

2.1 Masking Schemes and Applications

Masking is a prominent algorithmic countermeasure against Differential Power Analysis [47] that splits intermediate values of a computation into $d+1$ uniformly random shares [18, 40, 41, 44], such that an attacker who observes up to d shares cannot deduce information about native (unshared) intermediate values. Boolean masking is commonly used for symmetric cryptography, and uses the exclusive or ($\oplus$) operation to split a value b into $d+1$ uniformly random shares $b_0 \dots b_d$ such that $b = \bigoplus_i b_i = b_0 \oplus \cdots \oplus b_d$. In arithmetic masking schemes, the relation between shares of a value a is the modular addition, yielding $a = \sum_i a_i = a_0 + \cdots + a_d \mod q$. In both cases, masking linear functions is trivial since they can simply be computed for each share individually. Masking non-linear functions is more challenging since these functions operate on all shares of a native value and thus usually require additional fresh randomness to avoid unintended direct combination of shares. The concrete technique (Boolean or arithmetic) determines which operations are (non-)linear.

PQC algorithms often perform operations like matric/polynomial multiplication, which can be efficiently masked in the arithmetic domain when broken

[3] https://github.com/barbara-gigerl/arithmetic-masking-hw-sw

down into coefficient-wise modular addition/multiplications using e.g. the number theoretic transform (NTT). In practice, arithmetic masking often has to be combined with Boolean masking since building blocks including Gaussian samplers and lattice decoding, or constructions like the Fujisaki-Okamoto transform for achieving CCA2-security are more efficiently masked in the Boolean domain. Therefore, many masked implementations use dedicated conversion algorithms to transform shares from the arithmetic to the Boolean domain (A2B) and vice versa (B2A). Besides PQC, arithmetic masking is also applied to ARX-based implementations like SHA-256 or ChaCha, but comes with a significantly higher runtime overhead compared to Boolean masked variants of non-ARX symmetric algorithms.

2.2 Mask Conversion Techniques

Many cryptographic schemes require to switch between the Boolean and the arithmetic domain when respective masking techniques are applied. The performance of the protected scheme is mainly determined by the A2B and B2A conversions used, which is why there has been a lot of research in this direction [19, 20, 21, 22, 23, 59, 60]. Existing conversion algorithms either follow an *algebraic* or a *table-based* approach. An algebraic conversion algorithm performs the whole conversion at once, while table-based approaches first pre-compute a table which is later used during the actual conversion. B2A conversions can be done very efficiently following the algebraic approach, while A2B is less efficient, and therefore often apply a table-based approach.

In 2001, the first algebraic conversion algorithms were proposed by Goubin [39], and by Coron et al. [23] for higher orders. They propose the SecAdd algorithm, which allows to securely add Boolean shares at any order using a power-of-two modulus. Many follow-up works use Coron et al.-A2B/B2A as a basis, and suggest several performance improvements [11,19,22,43]. Since PQC applications often require a prime modulus, Barthe et al. [5], and later Schneider et al. [60] suggest how to adapt Coron et al.-B2A to work with prime moduli.

Table-based A2B conversion algorithms use pre-computed tables to reduce the computation effort during the actual conversion. In general, A2B conversions transform the shares together with the carry which is produced in an arithmetic addition. The pre-computed tables are used to handle the conversion of the carry, and prevent unintended unmasking of native values. The first table-based A2Bs were suggested by Coron-Tchulkine [25] and Neiße-Pulkus [54]. They were however shown to be incorrect and insecure by Debraize [26], who suggests several corrected and optimized versions of their algorithms. Recently, Beirendonck et al. [9] show that Debraize-A2B does also not fulfill its security claims, and propose two further table-based A2Bs.

A2B/B2A conversions are applied to masked implementations of various PQC and ARX schemes against side-channel attacks. For example, the SecAdd algorithm by Coron et al. [23] has been used as a cryptographic primitive in several software [1,5,14,22,35,60] and hardware implementations [16,29]. Debraize-A2B has also been applied recently in works on masking PQC [14,55].

2.3 Masking Verification with Rebecca/Coco

Rebecca [13] is a tool to formally verify Boolean-masked hardware implementations defined by gate-level netlists. In order to verify a circuit, a label is assigned to each circuit input. The label is either a share, fresh randomness or unimportant. During the verification process, these labels are propagated through the circuit and each gate is assigned a correlation set according to the propagation rules. In general, a correlation set contains information about the statistical dependence of the respective gate on the circuit inputs. Tools like Silver [46] compute these dependencies accurately, while Rebecca approximates statistical dependence with non-zero Fourier coefficients [13]. In general, the Fourier or Walsh expansion of a boolean function refers to the representation of the function as a multilinear polynomial [56]. A term in the polynomial, which is either a label or a combination of labels, with a non-zero Fourier coefficient indicates statistical dependence on the respective circuit input. The approximation is performed by not tracking the exact Fourier coefficient, but only whether a term has a non-zero coefficient or not. A correlation set contains all terms with non-zero coefficients.

Later, an optimized variant of this approach was implemented in Coco, a tool for the formal verification of (any-order) Boolean masked software implementations on concrete CPUs [36, 37]. The main purpose of Coco is to analyze the potential implications of hardware side-effects like glitches within a CPU on masked software implementations. Coco can additionally incorporate control flow logic, which is required for the verification of software and iterative hardware circuits. Before the verification, the CPU netlist is simulated together with a masked assembly implementation, in order to obtain a trace of the (constant) data-independent control signals like memory/register access patterns and branches. Next, similar to Rebecca, initial labels are assigned to registers and memory locations, which are further propagated through the netlist for multiple cycles to construct correlation sets, while considering software-specific control signals. The verification fails if there exists a gate in the netlist which directly correlates with a native value. In that case, Rebecca reports the leaking gate, while Coco additionally reports the exact clock cycle.

2.4 Adversary Model

The robust probing model for hardware [28, 44] allows an attacker to observe the values of up to d wires in a masked circuit using (g, t, c)-extended probing needles to optionally include glitches ($g = 1$), transitions ($t = 1$) or coupling ($c = 1$). The circuit is dth-order secure if the adversary is not able to learn anything about the native value by combining these observations. Accordingly, the standard probing model for software allows the adversary to probe intermediate program variables.

In this work we use the so-called *time-constrained probing model* [36], which is currently adopted by Coco for masked software implementations. The main difference to the robust probing model is the time restriction of each probe to one clock cycle, which is necessary to correctly model the execution of masked software on netlist level. More concretely, in the time-constrained probing model the attacker uses $(g, t, 0)$-extended probes to observe the value of any specific gate/wire in the CPU netlist for the duration of one clock cycle. The gate/wire and cycle can be chosen independently for each probe.

The time-constrained probing model can be applied to masked hardware circuits and allows to handle *iterative* circuits directly without the need to perform *unrolling* thanks to its time-awareness. In Appendix A we give an example of an iterative circuit and its unrolled version based on the suggestion of [12]. Verification approaches adopting the classic/robust probing model usually unroll the processed iterative circuit, which works well for simple circuits, but is more difficult for circuits with more complex control logic, such as state machines. Iterative circuits can be seen as a reduced version of a CPU, and therefore allows the direct application of the time-constrained probing model.

The original version of Coco provides two different verification modes in the time-constrained probing model. *Stable* verification focuses on pure algorithmic security. *Transient* verification uses $(g, t, 0)$-extended probes, and therefore considers algorithmic security and wire/register transitions and glitches within the hardware. For the purpose of this work, we add a third mode, the *Transitions* verification mode, working with $(0, t, 0)$-extended probes, which is convenient since it reports stable and transition leaks, but without the runtime overhead of the transient mode.

3 Verification of Arithmetic Masking in the Boolean Domain

In this section, we explain how one can perform verification of arithmetic masking using a method based on approximating Fourier coefficients of Boolean functions that was previously used by the tools Rebecca/Coco in the context of Boolean masking. In Section 3.1 we recall how arithmetic expressions can directly be broken down into equivalent Boolean expressions on bit-granularity. In Section 3.2 we discuss optimization strategies that can be used to reduce the complexity of the derived Boolean expressions for the initial labeling and to more efficiently propagate expressions through dedicated arithmetic addition circuits. We also comment on the soundness and completeness of our approach. Finally, we give a small self-contained example in Section 3.3.

Notation We denote with $a^{(i)}$ the i-th bit of variable a, with a^0 being the least significant bit (LSB). The j-th share of a native variable x is identified as x_j. Similar to Bloem et al. [13], we denote the correlation set of a gate/wire w by $\mathcal{C}(w) = \{...\}$. As introduced in [36], the $\otimes$-operator computes the element-wise multiplication of two correlation sets. We use small letters for symbolic expressions, while capital letters are used for wires in a circuit.

3.1 Modeling Arithmetic Expressions using Boolean Logic

A netlist represents a circuit design after logic synthesis that models gates as Boolean functions mapping 1-bit inputs to a 1-bit output, and indicates their interconnection. We aim at performing netlist-level verification of a circuit on bit granularity. In the end, a bitwise view on all terms computed by the circuit must still valid in the context of masking. This implies that the dependencies between the shares must be described using Boolean equations on bit granularity. Such a mapping can be obtained based on the definition of the Ripple-carry adder, which represents a cascade of 1-bit full adders, where each carry bit *ripples* to the next full adder. Each full adder takes two 1-bit summands and a 1-bit carry-in, and computes the arithmetic sum and respective carry-out [49].

Consider a sum s, which is computed from the summands u and v such that $s = u + v$. If u and v are n-bit values, s is represented by $n + 1$ bits, and hence, $n + 1$ full adders are needed to compute s. Each full adder takes two summand bits $u^{(i)}$ and $v^{(i)}$ together with the carry-in $c^{(i)}$, and computes $s^{(i)}$ as:

$$s^{(i)} = u^{(i)} \oplus v^{(i)} \oplus c^{(i)} {}_{\text{with } c^{(0)}=0} \tag{1}$$

The carry-out bit $c^{(i+1)}$ is then computed based on the carry-in $c^{(i)}$ by the following recursive formula:

$$c^{(i+1)} = (u^{(i)} \oplus v^{(i)}) \wedge c^{(i)} \vee (u^{(i)} \wedge v^{(i)}) \tag{2}$$

Equation 1 already gives a valid first-order Boolean sharing for s using the two shares $x_1 = u$ and $x_2 = v \oplus c$.

If a sum t is split into three summands u, v and w such that $t = u + v + w$, basically the same equations apply, and t can be computed in two steps. In the first step, the partial sum $s = u+v$ is computed, which yields the carry c. In the second step, t is computed by adding the partial sum to the remaining summand: $t = s + w$, which produces the carry e:

$$\begin{aligned} t^{(i)} &= \quad s^{(i)} \quad \oplus w^{(i)} \oplus e^{(i)} && (3) \\ &= u^{(i)} \oplus v^{(i)} \oplus c^{(i)} \quad \oplus w^{(i)} \oplus e^{(i)} && (4) \end{aligned}$$

Equation 4 gives a valid second-order Boolean sharing for t using three shares $x_1 = u, x_2 = v$ and $x_3 = w \oplus c \oplus e$. Formulas for more than three summands can be derived in a similar way, each resulting in a valid higher-order sharing. When working with $d + 1$ shares, the first d Boolean shares would always be equal to the first d arithmetic shares, while the last Boolean share needs to additionally include the carry.

3.2 Tailoring the Verification Approach

Arithmetically masked circuits process arithmetic input shares, while Rebecca/Coco expects Boolean input shares. The derived Boolean equations for arithmetic expressions in Section 3.1 can now be used to translate arithmetic shares to the Boolean domain, such that Rebecca/Coco could work with it. In the following, we describe how one can obtain such a translation in a correct and efficient way, how the resulting expressions can be propagated efficiently, and comment on soundness, completeness and scalability of the resulting approach.

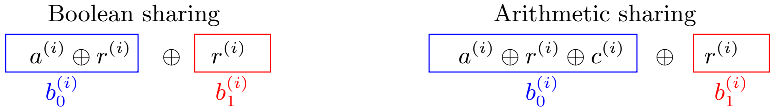

Fig. 1: Initial labeling for Boolean and arithmetic masking as given to the verifier

Initial Labeling Tools for the formal verification of masking require a set of initial labels that specify the location/dependency of shares on circuit inputs, registers or memory cells that are then further tracked throughout a circuit. In the case of (first-order) Boolean masking, each bit of a native value $a^{(i)}$ is initially masked with a random mask $r^{(i)}$. Therefore, the native value $a^{(i)}$ can simply be expressed as the XOR between the two shares $a^{(i)} \oplus r^{(i)}$ and $r^{(i)}$. As shown in Figure 1, the labels assigned prior to the verification would then be $b_0^{(i)} = a^{(i)} \oplus r^{(i)}$ and $b_1^{(i)} = r^{(i)}$.

In the case of (first-order) arithmetic masking, each bit of a native value $a^{(i)}$ is initially masked with a random mask $r^{(i)}$ using modular additions. According to Equation 1, the native value $a^{(i)}$ can be expressed as the XOR between the two shares $a^{(i)} \oplus r^{(i)} \oplus c^{(i)}$ and $r^{(i)}$. In contrast to Boolean masking, we also need to include the carry of the addition $c^{(i)}$, which depends on lower bits of $a^{(i)}$ and $r^{(i)}$. The first option to obtain a valid labeling for arithmetic shares is thus to resolve $c^{(i)}$ recursively according to Equation 2. The initial labels would then be given by $b_0^{(i)} = (a^{(i)} \oplus r^{(i)}) \oplus c^{(i)}$, and $b_1^{(i)} = r^{(i)}$. Here, the carry $c^{(i)}$ is computed recursively for each bit position, which adds already quite complex terms to the correlation set at the beginning of the verification, especially for the more significant bits of the arithmetic shares since the depend in a non-linear way on all lower bits.

It is however also possible to use a different initial labeling that incorporates additional information that is available at the beginning of the verification and significantly simplifies the resulting Boolean expressions. More concretely, with each $c^{(i)}$ being a non-linear combination of all lower bits (including their masks), this expression alone must never be observable by an attacker. Put differently, each bit of a fresh arithmetic share is only independent of any native values because the term $r^{(i)}$ is added in a linear way and does not occur in any of the lower bits (and thus also not $c^{(i)}$). It is hence sufficient to verify if the linear term $r^{(i)}$ in a certain bit of one arithmetic share ever gets in contact with the same $r^{(i)}$ in the corresponding bit of the other share, similarly as in the case of Boolean masking (c.f. Figure 1). This simplification leads to simpler expressions for the initial labels and thus improves verification runtime. Note that this simplification is only used for deriving initial labels but not during mask refresh operations throughout the masked computation where our assumptions on unique usage of fresh randomness does not necessarily hold anymore. This simplification also applies to initial labels of higher order arithmetic masking in a similar manner.

Fourier Expansion of Arithmetic Addition One particularly challenging aspect of verifying arithmetic masking is scalability due to complex dependencies between shares on bit-level, introduced by the carry when an arithmetic addition is computed. In hardware, arithmetic additions are often performed by dedicated sub-circuits. For example, CPUs usually have such an adder circuit in their ALU (Arithmetic Logic Unit). In Equation 5 we propose the Fourier expansion W of arithmetic additions, which allows to directly obtain correlation sets for the result of an adder circuit, instead of computing an individual correlation set for every gate within the adder, and thus speeds up the verification runtime. The expansion of the sum is based on the Fourier expansion of the carry given in Equation 6.

$$W(s^{(j)}) = \frac{1}{2}u^{(j)} \cdot v^{(j)} \cdot c^{(j)} + \frac{1}{2}u^{(j)} + \frac{1}{2}v^{(j)} - \frac{1}{2}W(c^{(j)}) \tag{5}$$

$$W(c^{(j)}) = \frac{1}{2}W(c^{(j-1)}) + \frac{1}{2}v^{(j)} + \frac{1}{2}u^{(j)} - \frac{1}{2}u^{(j)} \cdot v^{(j)} \cdot W(c^{(j-1)}) \tag{6}$$

In Section 5.1 we give more details about how this can be used to increase the performance of software verification. More details on how we derived both expansions are given in Appendix B.

Soundness and Completeness While masking verification based on approximated Fourier coefficients of Boolean functions is sound (leakages are never missed), it is not complete (leaks might be reported although the implementation is secure). Throughout a masked computation it might happen that certain terms in the exact Fourier representation cancel out or evaluate to constants. Our verification approach might miss such situations since it only keeps track of whether a term occurs in a correlation set or not (for performance reasons), which ultimately results in an overapproximation of the exact Fourier representation. If a situation occurs in which e.g. multiple shares with a correlation coefficient of zero are combined, the verifier would report a leak that does not exist in practice (which implies non-completeness). Soundness is however guaranteed by the fact that the verifier always keeps track of an *over*approximation of all the terms that a register/wire could depend on, hence, a real leak can never be missed.

In case of sound but not complete masking verification approaches, the amount of false positive leakage reports in realistic scenarios plays an important for practicality. Simply speaking, the longer a computation becomes, the more likely a false positive occurs. Note however that after every mask refresh operation, the newly introduced randomness essentially eliminates possible future false-positive leaks caused by over-approximation that has happened thus far. In other words, as long as mask refreshing occurs somewhat frequently (which is generally the case) the occurrence of false positive leak reports will generally be quite low. Later, in Section 4 and Section 5, we show that the soundness of our approach is in fact sufficient to perform meaningful verification of masked SW/HW implementations in many typical PQC/ARX applications.

During our analysis in this work, we only really observe a single false positive when verifying Goubin-A2B [39] in software, as discussed in more detail in Section 5.2.

3.3 Example

Finally, we give an example about how correlation sets are constructed using our verification approach. Assume an example circuit which takes two 2-bit arithmetic shares $a+r$ (input signal A_0) and r (input signal A_1), and two bits of fresh randomness s (input signal S). The ultimate goal is to compute $(A_0+S)+A_1$ by using two *Full Adder*s. In order to verify the first-order security of this circuit, one first assigns the respective labels to the inputs which result in the following correlation sets:

$$\begin{aligned}
&\mathcal{C}(A_0^{(0)}) = \{\{b_0^{(0)}\}\}, && \mathcal{C}(A_1^{(0)}) = \{\{b_1^{(0)}\}\}, && \mathcal{C}(S^{(0)}) = \{\{s^{(0)}\}\},\\
&\mathcal{C}(A_0^{(1)}) = \{\{b_0^{(1)}\}\} && \mathcal{C}(A_1^{(1)}) = \{\{b_1^{(1)}\}\} && \mathcal{C}(S^{(1)}) = \{\{s^{(1)}\}\}
\end{aligned}$$

The input bits are propagated to the first adder, which computes $(A_0 + S)$. We obtain the following correlation sets at the output signals of the first adder:

$$\begin{aligned}
\mathcal{C}(\text{Adder1}_{sum}^{(0)}) &= \mathcal{C}(A_0^{(0)}) \otimes \mathcal{C}(S^{(0)}) = \{\{b_0^{(0)}, s^{(0)}\}\}\\
\mathcal{C}(\text{Adder1}_{sum}^{(1)}) &= \mathcal{C}(A_0^{(1)}) \otimes \mathcal{C}(S^{(1)}) \otimes \mathcal{C}(\text{Adder1}_{carry}^{(1)})\\
&= \{\{b_0^{(1)}, s^{(1)}\}\} \otimes \{\{1\}, \{b_0^{(0)}\}, \{s^{(0)}\}, \{b_0^{(0)}, s^{(0)}\}\}\\
&= \{\{b_0^{(1)}, s^{(1)}\}, \{b_0^{(1)}, s^{(1)}, b_0^{(0)}\}, \{b_0^{(1)}, s^{(1)}, s^{(0)}\}, \{b_0^{(1)}, s^{(1)}, b_0^{(0)}, s^{(0)}\}\}\\
\mathcal{C}(\text{Adder1}_{sum}^{(2)}) &= \mathcal{C}(\text{Adder1}_{carry}^{(2)})
\end{aligned}$$

Note that the second bit of the adder has to be labeled with the (recursively resolved) carry of the addition. These correlation sets are then propagated to the second adder:

$$\begin{aligned}
\mathcal{C}(\text{Adder2}_{sum}^{(0)}) &= \mathcal{C}(A_1^{(0)}) \otimes \mathcal{C}(\text{Adder1}_{sum}^{(0)}) = \{\{b_1^{(0)}, b_0^{(0)}, s^{(0)}\}\}\\
\mathcal{C}(\text{Adder2}_{sum}^{(1)}) &= \mathcal{C}(A_1^{(0)}) \otimes \mathcal{C}(\text{Adder1}_{sum}^{(1)}) \otimes \mathcal{C}(\text{Adder2}_{carry}^{(1)})\\
&= \{\{b_1^{(1)}, b_0^{(1)}, s^{(1)}\}, \{b_1^{(1)}, b_0^{(1)}, s^{(1)}, b_0^{(0)}\}, \{b_1^{(1)}, b_0^{(1)}, s^{(1)}, s^{(0)}\},\\
&\qquad \{b_1^{(1)}, b_0^{(1)}, s^{(1)}, b_0^{(0)}, s^{(0)}\}\}\\
\mathcal{C}(\text{Adder2}_{sum}^{(2)}) &= \mathcal{C}(\text{Adder1}_{sum}^{(2)}) \otimes \mathcal{C}(\text{Adder2}_{carry}^{(2)})\\
\mathcal{C}(\text{Adder2}_{sum}^{(3)}) &= \mathcal{C}(\text{Adder2}_{carry}^{(2)})
\end{aligned}$$

Obviously, $(A_0 + S) + A_1$ is a valid operation in the context of arithmetic masking and this is also visible on bit-level. The computation of the carry bits of the second adder combines shares in a non-linear way, which typically leads to a leak. However, the addition is still secure in the end since (A_0+S) adds randomness to each share bit linearly. When performing an addition of two operands we always conservatively label one bit more than the size of the largest operand to correctly capture bit width of the result independently on the concrete input values. Note that by performing modular reduction one can clear the carry residing in the most significant bit (MSB). This type of computation occurs very frequently in the beginning of A2B algorithms when two arithmetic shares should be added since the addition of fresh randomness is equivalent to a mask refreshing operation.

4 Application to Masked Hardware Implementations

In this section we apply our verification approach to hardware implementations of Coron et al.-A2B. While it has already been shown in the past that this algorithm is secure in the stable setting, which is also confirmed by our verifier, we want to put our focus mainly on settings where we also consider transition and glitch effects. We show, both via a formal analysis, and in empirical evaluations, that hardware side-effects can reduce the protection order of the implementation. While the straight-forward approach of adding additional register stages whenever needed can eliminate this problem, we also want to point out that this comes with a noticeable increase of latency.

Coron et al.-A2B/B2A In 2014, Coron et al. [23] have proposed the first higher-order mask conversion algorithm, which we refer to as Coron et al.-A2B/B2A in the following. This algorithm is based on the SecAdd function and allows to perform arithmetic additions in a Ripple-carry fashion on Boolean shares. More specifically, the algorithm converts the arithmetic shares a_0 and a_1 which correspond to the native value a into the Boolean shares b_0 and b_1. The conversion starts with the *initial remasking*, where the arithmetic input shares are refreshed by adding fresh randomness, followed by the *carry computation.* A single native carry bit is computed based on Equation 2, which can be rewritten as:

$$c^{(i+1)} = u^{(i)} \wedge v^{(i)} \oplus u^{(i)} \wedge c^{(i)} \oplus v^{(i)} \wedge c^{(i)} \text{ with } c^{(0)} = 0 \tag{7}$$

However, the algorithm operates on shared carries c_0 and c_1 instead on the native c, which are computed bit by bit using secure masked AND gadgets (SecAnd). In Appendix C we sketch the structure of Coron et al.-A2B when implemented in hardware. The corresponding B2A conversion chooses the first arithmetic share a_0 randomly, and computes $a_1 = (b_0 \oplus b_1) - a_0$ using SecAdd. The algorithm is very efficient for hardware implementations [29], since SecAdd can be used for both A2B and B2A, and both can also be applied to higher orders. In Section 5 we formally evaluate both Coron et al.-B2A, and a second-order masked software implementation of Coron et al.-A2B.

4.1 Formal Analysis

We implement Coron et al.-A2B with 16-bit shares in hardware. We store all inputs in registers, and implement the remaining parts as a pure combinatorial circuit, which takes a single cycle to finish and therefore does not require a state machine. The input shares as well as the necessary random values are stored in registers. The verifier confirms algorithmic security for this single-cycle implementation, while in the transient case under the consideration of glitches, first-order side-channel protection is not given. More concretely, glitches in the initial remasking phase and the SecAnd modules, which are part of the bigger SecAdd, may lead to a temporary combination of shares due to delayed addition of randomness.

Table 1: Verification of Coron et al.-A2B (broken and fixed) in hardware

Algorithm	Input shares	Runtime (cycles)	Verification result/runtime					
			Stable		Transitions		Transient	
Coron et al. [23]	16 bit	1	✓	11 s	✓	10 s	✗	1 s
Coron et al. [23]	16 bit	34	✓	56 s	✓	2 min	✓	3 min

Initial Remasking Two XOR gates at the circuit's inputs are used to perform the refreshing during the initial remasking phase, which each combines an arithmetic share with a random value. In the worst case, a glitch at the output of the XOR gate propagates the pure values of a_0 and a_1, for example, when the wire delay of the random values is bigger than the wire delay of the arithmetic shares. The input of SecAdd will then be the arithmetic shares without refreshing, and the circuit computes $a_0 + a_1 = a$ for a short time frame in the beginning of the clock cycle, until all wires stabilize and the randomness *arrives* at the gates. As a solution, we add a single additional register stage to store the result of these XOR computations. This ensures that the SecAdd module's input comes out of a register instead of combinatorial logic, and will therefore not glitch.

SecAnd Coron et al. suggest to use the masked AND gadget proposed by Ishai et al. [44], called ISW-AND, in the SecAnd-blocks of the conversion. Formal verification however reports a leak due to glitches in the SecAnd module because the ISW-AND is not glitch-resistant, and also does not fulfill the required composability properties. As a solution, we suggest to insert two register stages to the SecAnd component. Works like [24, 28, 48, 52] confirm our observation that these two register stages are indeed needed in this case. Combined with the register stage inserted for the initial remasking, this results in a high latency overhead, i.e., for n-bit input shares, the implementation now requires $34 = 2 + 2 \times n$ cycles to complete, and also utilizes a state machine in order to control the execution.

Using this case study, we evaluate our verification approach for masked hardware circuits in Table 1 by comparing the broken single-cycle implementation to the one which adapts our fixes. All experiments are run using a 64-bit Linux Operating System on an Intel Core i7-7600U CPU with a clock frequency of 2.70 GHz and 16 GB of RAM. The security on algorithmic level of both implementations can be shown in 11 seconds and respectively 56 seconds in the stable case. We need around a second to find the issues in the transient case, and about three minutes to prove that our fixes indeed provide first-order protection. Our implementation serves the purpose of a proof-of-concept and allows further extensive optimizations. However, we consider the discussion of these optimizations along with the evaluation of area and performance overhead out of scope for this paper.

4.2 Empirical analysis

In the last section we discussed the outcome of the formal analysis which indicates that glitches in the design are problematic in the context of masking. As a second step, we show practical evidence for the proposed statements.

Evaluation Setup We practically evaluate Coron et al.-A2B using a first-order t-test on the NewAE CW305 Artix-7 FPGA evaluation board connected to a

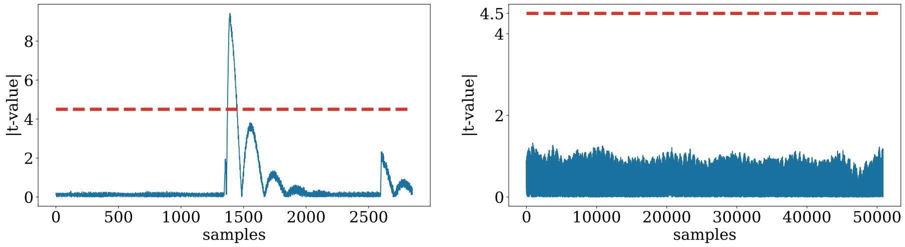

Fig. 2: T-test scores of the original (left) and the secured (right) implementation of Coron et al.-A2B using 400 000 power traces

PicoScope 6404C at 312.5 Ms/s sampling rate. The hardware design operates at a clock frequency of 1 MHz. In order to detect potential first-order leakage, we perform Welch's t-test following the guidelines of Goodwill et al. [38], which is a standard method to measure information leakage of masked implementations. The basic idea is to create two sets of measurements, one representing the power consumption of the design with random inputs (random set), and one with constant inputs (fixed set). For the fixed set, we assume the native value a is 0 and generate the respective shares a_0 and a_1 such that $a_0 + a_1 = a$. For the random set, we generate a_0 and a_1 completely at random. We use fresh random values for the random inputs in both cases. From these trace sets, one can compute Welch's t-score to measure the significance of the difference of means of the two distributions. The null-hypothesis is that both trace sets have equal means, which is rejected with a confidence greater than 99.999% if the absolute t-score does not exceed 4.5, and implies that the trace sets cannot be distinguished from each other.

Discussion Figure 2 shows our leakage assessment using 400 000 traces. The results for the original, unprotected single-cycle implementation are presented on the left. The t-test score shows significant peaks over the 4.5 border, indicating first-order leakage. On the right side, the leakage evaluation of our 34-cycle fixed implementation is shown, in which the t-score does not cross the significance boarder. Thus, these measurements confirm the security claim made by the formal tool. In Appendix D we show the functionality of the measurement setup by turning the random number generator off.

Note that Coco verifies ASIC netlists of masked implementations and identifies problematic wires where leaks might occur. The exact structure of this netlist must be reflected on the final FPGA layout to make concrete security statements, which is why we cannot simply synthesize the hardware design to the FPGA. The synthesis process will possibly merge multiple ASIC gates into a single lookup table (LUT) on the FPGA, and the original netlist structure will not be preserved. Consequently, one might see artifacts in the measurements stemming from this merging process, e.g. because the strict separation of shares is lost in the translation process [17]. Therefore, we must ensure to map each gate in the verified ASIC netlist to a functionally equivalent FPGA LUT, in order to preserve the original netlist structure as good as possible. We achieve this by mapping each ASIC gate to a LUT with 2 inputs and one output, by putting a `dont_touch = "true"` on every gate/wire in the netlist.

5 Application to Masked Software Implementations

In this section we discuss how Coco can be used to identify leaks in arithmetically masked RISC-V assembly implementations. In the beginning, we outline the software verification setup. First, we focus on algebraic conversions, including Coron et al. [23], Schneider et al. [60] for prime moduli and Goubin-A2B/B2A [39], for which we point out several register overwrite leaks. We discuss the table-based conversion algorithms of Debraize [26] and Beirendonck et al. [9], and explain how table lookups can be formally verified from a probing-security perspective. To conclude the section, we verify the masked ARX-based schemes Speck 32/64 and Alzette.

5.1 Software Verification Setup

Potential leaks in masked software are either caused by flaws in the algorithmic design, or due to microarchitectural side-effects of the processor's hardware. Flaws in the algorithmic design are mainly attributed to non-uniform sharings of intermediate variables, accidental combinations of masks, or transition leakage caused by register overwrites. However, even if such issues are taken into account there is still no guarantee that such an algorithm, once implemented for a specific processor, will be free of leaks. For example, a recent work by Gigerl et al. [36] has analyzed the RISC-V Ibex core in terms of architecture side-effects for masked software, and has pointed out multiple additional potential sources of leakage due to the design of the register file, the SRAM, the ALUs, and the load-store unit. They created a *secured* Ibex[4] that incorporates some relatively cheap hardware fixes that mostly eliminate glitch-related issues that are otherwise difficult to deal with purely on software-level.

For the purpose of this paper we are not so much interested into further netlist modifications, but rather focus on potential flaws in the algorithmic design of masked software implementations. We use their *secured* Ibex core as a reference platform that comes with a concrete list of hardware side-effects that do or do not need to be taken into consideration in software, thus allowing for an even playing field when evaluating and comparing different masked software implementations of A2B/B2A conversion algorithms. More specifically, the certain common microarchitectural leakages do not need to be addressed in software because the *secured* Ibex already has appropriate fixes on netlist-level. These fixes include:

- A glitch-resistant register file which allows to read and write shares without combination, as long as the respective software constraints are met
- No hidden registers or always-active computation units
- A glitch-resistant model of the SRAM (similar to the register file)

For more details on these fixes, we refer to the work of Gigerl et al. [36]. When a masked assembly implementation is executed by the *secured* Ibex and the software constraints are met, the leakages which are left are primarily register/memory overwrites and leaks caused by algorithmic flaws. The results of the

[4] https://github.com/IAIK/coco-ibex

Table 2: Verification results for masked software: ✓(no issues were found), (✓) (algorithmically secure, but potential problems like table-lookups), ✗(algorithmically insecure implementations), ✘(false positive).

Algorithm	Input shares	Runtime (cycles)	Verification result/runtime					
			Stable		Transitions		Transient	
A2Bs								
Coron et al. [23]	4 bit	225	✓	41 s	✓	67 s	✓	3 min
Coron et al. [23]	16 bit	984	✓	4 min	✓	5 min	✓	16 min
Coron et al. [23] (2nd order)	4 bit	1240	✓	3.8 min	✓	6 min	✓	20 min
Debraize [26] ⊞	4 bit ($n = 2$, $k = 2$)	140	✗	35 s	-		-	
Debraize [26] ⊞	16 bit ($n = 4$, $k = 4$)	450	✗	118 s	-		-	
Beirendonck et al.-fixed-Debraize [9] ⊞	4 bit ($n = 2$, $k = 2$)	180	(✓)	38 s	(✓)	48 s	-	
Beirendonck et al.-Dual-Lookup [9] ⊞	4 bit ($n = 2$, $k = 2$)	105	(✓)	28 s	(✓)	30 s	-	
Goubin [39]	16 bit	170	✘	37 s	✗		-	
B2As								
Goubin [39]	16 bit	23	✓	5 s	✓	8 s	✓	19 s
Coron et al. [23]	4 bit	650	✓	6 min	✓	4 min	✓	11 min
Coron et al. [23]	16 bit	2475	✓	11 min	✓	16 min	✓	38 min
Schneider et al. [60], without final mod instruction	4 bit, ($q = 257, \log_2 q = 9$)	400	(✓)	2 min	(✓)	21 min	-	
ARX-based schemes								
Speck 32/64 1 round	6 × 16 bit	1465	✓	6 min	✓	13 min	✓	5.13 h
Alzette 1 round	2 × 32 bit	3082	✓	29 min	✓	2.48 h	✓	27 h

following analysis can therefore be ported to any other microprocessor, as long as the respective device-specific fixes against these leaks, either in hardware or in software, are implemented.

The synthesis process will transform the adder, which lies in the ALU of the *secured* Ibex, to a set of logic gates. Theoretically, each gate is each assigned a correlation set during verification, which is very time-consuming. We wrap up the addition into a custom `adder` "gate" instead of splitting it up, which means only the output wires of the adder must be assigned correlation sets. In order to achieve this, we identify the addition in the CPU design before synthesis (which is trivial), move it into a distinct module, and apply `keep_hierarchy` on this module, which results in a single `adder` gate on netlist level. In case of the *secured* Ibex core, the `adder` gate is represented by 2 × 32-bit inputs, and creates a 33-bit output, for which we can compute the correlation set quite efficiently using Equation 5 and Equation 6. Without this optimization, the synthesizer would split the adder up into individual logic gates and one would check the correlation of each of these gates individually. Consequently, especially verification in transient mode would then not be possible in a feasible time frame [5]. It is important to note that this does not affect the soundness guarantees of our approach because the correlation sets computed for the outputs of the `adder` gates are identical to the correlation sets of an adder which is split up.

5.2 Verification of Algebraic Share Conversions

Coron et al.-A2B/B2A [23] In Section 4 we discuss the verification of Coron et al. [23] conversion algorithms in hardware. When verified on a CPU netlist, the algorithm in general behaves very similar. As shown in Table 1, we implement

[5] Runtime of a few hours for a single 32-bit addition

Coron et al.-A2B and -B2A in software and verify it successfully. We provide 16-bit A2B and B2A implementations which we verify in all three modes. Additionally, we implement 4-bit first- and second-order implementations, which can also successfully be verified with our approach. Compared to the results of Section 4, where we verify a 34-cycle implementation in 3 min, we can verify the respective software implementation (~1000 cycles) in 20 min, which shows the efficiency of our tool. Interestingly, both QMVerif and `LeakageVerif` have to fall back to exhaustive enumeration when verifying Coron et al.-B2A, while the other direction (A2B) is possible [50].

Schneider et al.-B2A [60] Various A2Bs/B2As work with power-of-two moduli exclusively, while many lattice-based constructions require a prime modulus. To address this issue, one can first transform the shares from $\mathbb{F}_q$ to $\mathbb{F}_{2^k}$, and then apply conversion algorithms working with power-of-two moduli [55]. Another possibility is to directly adapt a $\mathbb{F}_{2^k}$-conversion algorithm to work in $\mathbb{F}_q$. We investigate Schneider et al.-B2A [60], which is an adaption of Coron et al.-B2A, and was initially proposed to build a masked binomial sampler. We sketch the algorithm in Appendix E. We construct a first-order implementation of Schneider et al.-B2A with $q = 257$ and 4-bit input shares.

During conversion, Schneider et al.-B2A heavily uses reductions mod q, which are usually not implemented using the processor's `mod` instruction due to the instruction's large runtime overhead. Instead, many practical implementations use efficient reduction methods like Montgomery [51] or Barret [2] in combination with lazy reduction, i.e., skipping reductions as long as intermediate values are guaranteed to fit inside 32-bit words (on 32-bit architectures) [15]. For our implementation, we eliminate all reductions except the very last at the end of the algorithm, where we stick to Barret reduction. These tricks not only significantly improve runtime but also reduce the verification runtime drastically, since mod operations create very complex dependencies between individual bits of a share. In this setting, we want to point out and interesting pitfall that should be avoided when using lazy reduction techniques in the context of masking. For example, in order to convert the Boolean shares b_0 and b_1, Schneider et al.-B2A first generates a random number $E_0 \in \mathbb{F}_q$, and then computes $E_1 = ((b_1 - E_0 \mod q) - 2 \cdot ((b_1 - E_0 \mod q) \cdot b_0)) \mod q$. If one now lazily skips the reductions, the upper bits of E_1 will not be masked anymore due to the smaller bit width of E_0 . To mitigate this potential pitfall on 32-bit architectures, one could simply always use 32-bit words of randomness whenever mask refreshing is required. Other than that, the verification points out no issues in the algorithm.

Goubin-A2B/B2A [39] Goubin's algorithms [39] fix one output share, while the other is computed accordingly in order to derive a correct arithmetic or Boolean sharing. Goubin-B2A fixes $a_1 = b_1$ and then applies the recursive rule $a_0 = (b \oplus b_1) - b_1$ to compute the second share, while the respective A2B fixes $b_1 = a_1$ and computes b_0 by recursion instead. Goubin-B2A remains popular due to its efficiency ($\mathcal{O}(1)$), while the A2B conversion is more costly ($\mathcal{O}(n)$ for n bits [23]). In Appendix G we outline both algorithms. We can successfully verify the security of the B2A conversion in the stable, transition, and transient case. However, we encounter several issues with the A2B conversion that we now describe in more detail. To the best of our knowledge, these findings have not been reported yet.

First, Goubin-A2B introduces several problems regarding insecure register overwrites even if we ensure that our implementation uses dedicated registers for each of the variables proposed in the original algorithm. The algorithm uses an intermediate Y, which is initialized with a random variable and overwritten several times during the computation. Each of these overwrites leaks the XOR between the old (Y_{old}) and the new (Y_{new}) value. The verifier points out two situations during the computation in which $Y_{\text{old}} \oplus Y_{\text{new}}$ reveals information about the native value a. This issue can however be fixed easily be ensuring that different registers are used for every re-assignment of Y. In Appendix G we give a detailed calculation of the issue.

Second, the verifier indicates another leak during the computation, which is however not a practical problem, but a false positive as already mentioned in Section 3.2. In the following, we want to briefly highlight the circumstances and give the exact calculation in Appendix G. During the computation, an attacker can probe the following 1-bit expression: $(Y^{(0)} \wedge (Y^{(0)} \oplus a_1^{(0)})) \oplus (a_1^{(0)} \wedge (Y^{(0)} \oplus a_0^{(0)}))$, with $Y^{(0)}$ being random. The exact Fourier expansion of this expression does not contain a single term which depends on both $a_0^{(0)}$ and $a_1^{(0)}$ alone, but only in connection with $Y^{(0)}$, and is therefore properly masked. However, the verifier works with approximated correlation sets, which contain a subset $\{a_0^{(0)}, a_1^{(0)}\}$ where the random value $Y^{(0)}$ is not contained, and therefore represents a leak. According to [50], both QMVerif and `LeakageVerif` also fail to verify Goubin-A2B correctly because their tools produce false positives. Unfortunately, they do not discuss the exact issue, and therefore we were not able to make further investigations.

5.3 Verification of Table-based Share Conversions

Besides algebraic approaches, several A2Bs utilize lookups into pre-computed tables, such as the ones from Debraize [26] and Beirendonck et al. [9]. A table lookup represents a data-dependent memory access, i.e., an operation that loads data from a memory address that is data-dependent. Coco was mostly intended to verify symmetric cryptography, where table lookups are not common and have therefore not been considered. However, our study shows that the verification approach can be successfully applied under specific conditions, which are fortunately fulfilled by all table-based A2Bs that we are aware of.

First, it must be possible to compute all entries in the table with a single unique function $f(i)$, which depends only on the table index i and constants. This ensures that every table entry is assigned the same label during the verification independently of the address. For example, Debraize-A2B uses $f(i) = i + r + p \oplus (p||r)$ for initially generated random values ("constants") r and p. Second, the evaluation platform must guarantee constant-time memory accesses, i.e., memory accesses always require the same amount of cycles independently of the memory address. For example, the original Ibex core fetches multiple memory locations in case of a misaligned memory access, and therefore requires more cycles compared to an aligned memory access. Therefore, we simply disable the *secured* Ibex core's ability to perform misaligned memory accesses, which represents a quite reasonable modification for verification purposes since constant-time is anyway a desired property of cryptographic implementations.

5.4 Application to Table-based Conversion Algorithms

Table-based A2Bs usually take over one Boolean share from the arithmetic domain and derive the second by computing $b_0 = (a_0 + a_1) \oplus a_1$. From a masking perspective, $(a_0 + a_1)$ leaks the native value a, which is prevented using a precomputed look-up table which stores $(a_0 + r) \oplus r$ for a fixed r [9]. However, generating the table for each possible value of a_0 is not efficient.

Debraize-A2B [26] In 2012, Debraize [26] suggests to split up a_0 into n parts of k bits each, and precompute a table entry for each of the 2^k possible values. The actual conversion is performed by iterating over the n parts of a_0 and converting each part individually by performing a table lookup to the precomputed table. The table returns (a) the transformed part of a_0 into the Boolean domain, and (b) the one-bit carry that is Boolean masked and has to be considered in the next iteration. We sketch the algorithm in Appendix F. We implement Debraize-A2B with $n = 2, k = 2$ as well as $n = 4, k = 4$ and verify its execution as shown in Table 2. Two leaks are already reported in the stable verification mode (indicated by ✖), which points towards algorithmic errors.

First, the verifier reports a leak when performing the table lookup due to a combination of the (share-dependent) address bits and the memory content. On gate level, a table lookup using 32-bit addresses is realized using an equality comparator, which itself consists of 32 XNOR gates, whose output is combined by a single AND gate [49]. In case of equality, the AND gate outputs 1, and 0 otherwise. This information is finally used to decide whether to read data from a specific location. We give an example of an equality comparator in Appendix H. When performing the table lookup in Debraize-A2B in the first iteration, the address bits depend on both arithmetic shares, a random value r and a random bit β, which represents an intermediate result of the transformation. The content of the precomputed lookup-table is determined by r and β. Combining both values cancels out the random values r and β, and the attacker can probe an expression depending on the native value a. One can argue that an SRAM module is constructed in a way such that the address and the memory cell content will never be combined. However, in bigger CPUs, the memory access logic is much more complicated and might contain buffers or caches, which employ such an addressing mechanism. For example, data caches usually require the computation of a tag based on the address, and compare this tag to the one in the cache.

Second, the verifier points out that the value obtained from the lookup-table in the first iteration is not uniformly distributed, although used as a mask in the algorithm. Beirendonck et al. [9] already report the issue in their work, which was found by empirical measurements, and provide a theoretical analysis afterwards. We want to emphasize that another advantage of our verification approach is the fast discovery of such bugs, which happens in 35 s and 118 s according to Table 2 in this case, which is much quicker than empirical/theoretical evaluations. Coco reports the leaking cycle and netlist gate immediately and therefore one does not need to carry out a laborious empirical analysis.

Beirendonck et al.-A2Bs [9] In their work, Beirendonck et al. [9] propose two new secure table-based A2Bs, Beirendonck et al.-fixed-Debraize A2B (a secured version of Debraize-A2B) and Beirendonck et al.-Dual-Lookup A2B (an efficient version of Beirendonck et al.-fixed-Debraize A2B). We verify both algorithms by choosing parameters $n = 2, k = 2$. As shown in Table 2, table lookups cause a similar leak as we already discussed for Debraize-A2B. Since the issue however strongly depends on the underlying microarchitecture, and no further issues were found, we mark it with (✔) in the table.

5.5 Application to ARX-based Constructions

The ARX (Addition-Rotation-XOR) design principle has been used for several well-known symmetric cryptographic constructions like the block cipher Speck [7], the stream cipher ChaCha [10], or the hash function SHA-256 [53]. We focus on first-order implementations of a single round of Speck 32/64 [7], and the 64-bit ARX-based S-box Alzette [8]. Alzette is a central building block of Sparkle, which is currently one of the finalists of the NIST LWC Standardization Process [61]. Masking these implementations requires both Boolean masking (for the Rotation and XOR) and arithmetic masking (for the addition). One option is to apply an algorithm like SecAdd, which implements modular addition directly on Boolean shares [23, 27, 45, 59]. Another possibility is to first convert the Boolean shares to arithmetic shares, then perform the addition on arithmetic shares, and convert the shares back to the Boolean domain. In our implementation, we choose the second option using Goubin-B2A before each addition, perform the addition on arithmetic shares, and switch back to the Boolean domain using Coron et al.-A2B. We are able to verify algorithmic security in under 30 minutes for both schemes (stable mode). For the transient mode, the verification requires several hours, which is mostly spent by solving the SAT equation, and therefore offers several possibilities for further optimization.

6 Conclusion

In this paper, we presented an approach for the formal verification of masked software and hardware implementations, which supports both arithmetic and Boolean masking schemes of any order. On the hardware side, we show that glitches may cause issues in the context of masking for a straightforward implementation of Coron et al.-A2B. We demonstrate that this issue exists in practice using empirical measurements. On the software side, we first analyze algebraic share conversions, report a previously unknown register transition issue in

Goubin-A2B and provide new insights on the security of lazy reduction, a popular optimization technique in PQC. Second, we discuss table-based conversions and demonstrate that table lookups might not be secure due to architectural side-effects. Last but not least, we underline the scalability of our approach by applying it to entire round functions of masked ARX-based ciphers.

Acknowledgements This work was supported by the TU Graz LEAD project "Dependable Internet of Things in Adverse Environments". Additionally, this project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 681402).

References

1. A. Adomnicai, J. J. A. Fournier, and L. Masson. Bricklayer attack: A side-channel analysis on the chacha quarter round. In A. Patra and N. P. Smart, editors, *Progress in Cryptology - INDOCRYPT 2017 - 18th International Conference on Cryptology in India, Chennai, India, December 10-13, 2017, Proceedings*, volume 10698 of *Lecture Notes in Computer Science*, pages 65–84. Springer, 2017.
2. P. Barrett. Implementing the rivest shamir and adleman public key encryption algorithm on a standard digital signal processor. In A. M. Odlyzko, editor, *Advances in Cryptology - CRYPTO '86, Santa Barbara, California, USA, 1986, Proceedings*, volume 263 of *Lecture Notes in Computer Science*, pages 311–323. Springer, 1986.
3. G. Barthe, S. Belaïd, G. Cassiers, P. Fouque, B. Grégoire, and F. Standaert. maskverif: Automated verification of higher-order masking in presence of physical defaults. In *Computer Security - ESORICS 2019 - 24th European Symposium on Research in Computer Security, Luxembourg, September 23-27, 2019, Proceedings, Part I*, volume 11735 of *Lecture Notes in Computer Science*, pages 300–318. Springer, 2019.
4. G. Barthe, S. Belaïd, F. Dupressoir, P. Fouque, B. Grégoire, and P. Strub. Verified proofs of higher-order masking. In E. Oswald and M. Fischlin, editors, *Advances in Cryptology - EUROCRYPT 2015 - 34th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Sofia, Bulgaria, April 26-30, 2015, Proceedings, Part I*, volume 9056 of *Lecture Notes in Computer Science*, pages 457–485. Springer, 2015.
5. G. Barthe, S. Belaïd, T. Espitau, P. Fouque, B. Grégoire, M. Rossi, and M. Tibouchi. Masking the GLP lattice-based signature scheme at any order. In J. B. Nielsen and V. Rijmen, editors, *Advances in Cryptology - EUROCRYPT 2018 - 37th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tel Aviv, Israel, April 29 - May 3, 2018 Proceedings, Part II*, volume 10821 of *Lecture Notes in Computer Science*, pages 354–384. Springer, 2018.
6. G. Barthe, M. Gourjon, B. Grégoire, M. Orlt, C. Paglialonga, and L. Porth. Masking in fine-grained leakage models: Construction, implementation and verification. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2021(2):189–228, 2021.
7. R. Beaulieu, D. Shors, J. Smith, S. Treatman-Clark, B. Weeks, and L. Wingers. The SIMON and SPECK families of lightweight block ciphers. *IACR Cryptol. ePrint Arch.*, page 404, 2013.

8. C. Beierle, A. Biryukov, L. C. dos Santos, J. Großschädl, L. Perrin, A. Udovenko, V. Velichkov, and Q. Wang. Alzette: A 64-bit arx-box - (feat. CRAX and TRAX). In D. Micciancio and T. Ristenpart, editors, *Advances in Cryptology - CRYPTO 2020 - 40th Annual International Cryptology Conference, CRYPTO 2020, Santa Barbara, CA, USA, August 17-21, 2020, Proceedings, Part III*, volume 12172 of *Lecture Notes in Computer Science*, pages 419–448. Springer, 2020.
9. M. V. Beirendonck, J. D'Anvers, and I. Verbauwhede. Analysis and comparison of table-based arithmetic to boolean masking. *IACR Cryptol. ePrint Arch.*, 2021:67, 2021.
10. D. J. Bernstein. ChaCha, a variant of Salsa20. In *Workshop record of SASC*, volume 8, pages 3–5, 2008.
11. L. Bettale, J. Coron, and R. Zeitoun. Improved high-order conversion from boolean to arithmetic masking. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2018(2):22–45, 2018.
12. S. Bhasin, S. Guilley, L. Sauvage, and J. Danger. Unrolling cryptographic circuits: A simple countermeasure against side-channel attacks. In J. Pieprzyk, editor, *Topics in Cryptology - CT-RSA 2010, The Cryptographers' Track at the RSA Conference 2010, San Francisco, CA, USA, March 1-5, 2010. Proceedings*, volume 5985 of *Lecture Notes in Computer Science*, pages 195–207. Springer, 2010.
13. R. Bloem, H. Groß, R. Iusupov, B. Könighofer, S. Mangard, and J. Winter. Formal verification of masked hardware implementations in the presence of glitches. In *Advances in Cryptology - EUROCRYPT 2018 - 37th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tel Aviv, Israel, April 29 - May 3, 2018 Proceedings, Part II*, volume 10821 of *Lecture Notes in Computer Science*, pages 321–353. Springer, 2018.
14. J. W. Bos, M. Gourjon, J. Renes, T. Schneider, and C. van Vredendaal. Masking kyber: First- and higher-order implementations. *IACR Cryptol. ePrint Arch.*, 2021:483, 2021.
15. L. Botros, M. J. Kannwischer, and P. Schwabe. Memory-efficient high-speed implementation of kyber on cortex-m4. In J. Buchmann, A. Nitaj, and T. Rachidi, editors, *Progress in Cryptology - AFRICACRYPT 2019 - 11th International Conference on Cryptology in Africa, Rabat, Morocco, July 9-11, 2019, Proceedings*, volume 11627 of *Lecture Notes in Computer Science*, pages 209–228. Springer, 2019.
16. C. Chen, T. Eisenbarth, I. von Maurich, and R. Steinwandt. Masking large keys in hardware: A masked implementation of mceliece. In O. Dunkelman and L. Keliher, editors, *Selected Areas in Cryptography - SAC 2015 - 22nd International Conference, Sackville, NB, Canada, August 12-14, 2015, Revised Selected Papers*, volume 9566 of *Lecture Notes in Computer Science*, pages 293–309. Springer, 2015.
17. T. D. Cnudde, B. Bilgin, B. Gierlichs, V. Nikov, S. Nikova, and V. Rijmen. Does coupling affect the security of masked implementations? In S. Guilley, editor, *Constructive Side-Channel Analysis and Secure Design - 8th International Workshop, COSADE 2017, Paris, France, April 13-14, 2017, Revised Selected Papers*, volume 10348 of *Lecture Notes in Computer Science*, pages 1–18. Springer, 2017.
18. T. D. Cnudde, O. Reparaz, B. Bilgin, S. Nikova, V. Nikov, and V. Rijmen. Masking AES with d+1 shares in hardware. In *Cryptographic Hardware and Embedded Systems - CHES 2016 - 18th International Conference, Santa Barbara, CA, USA, August 17-19, 2016, Proceedings*, volume 9813 of *Lecture Notes in Computer Science*, pages 194–212. Springer, 2016.

19. J. Coron. High-order conversion from boolean to arithmetic masking. In W. Fischer and N. Homma, editors, *Cryptographic Hardware and Embedded Systems - CHES 2017 - 19th International Conference, Taipei, Taiwan, September 25-28, 2017, Proceedings*, volume 10529 of *Lecture Notes in Computer Science*, pages 93–114. Springer, 2017.
20. J. Coron, F. Gérard, S. Montoya, and R. Zeitoun. High-order table-based conversion algorithms and masking lattice-based encryption. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2022(2):1–40, 2022.
21. J. Coron, C. Giraud, E. Prouff, S. Renner, M. Rivain, and P. K. Vadnala. Conversion of security proofs from one leakage model to another: A new issue. In *Constructive Side-Channel Analysis and Secure Design - Third International Workshop, COSADE 2012, Darmstadt, Germany, May 3-4, 2012. Proceedings*, volume 7275 of *Lecture Notes in Computer Science*, pages 69–81. Springer, 2012.
22. J. Coron, J. Großschädl, M. Tibouchi, and P. K. Vadnala. Conversion from arithmetic to boolean masking with logarithmic complexity. In G. Leander, editor, *Fast Software Encryption - 22nd International Workshop, FSE 2015, Istanbul, Turkey, March 8-11, 2015, Revised Selected Papers*, volume 9054 of *Lecture Notes in Computer Science*, pages 130–149. Springer, 2015.
23. J. Coron, J. Großschädl, and P. K. Vadnala. Secure conversion between boolean and arithmetic masking of any order. In L. Batina and M. Robshaw, editors, *Cryptographic Hardware and Embedded Systems - CHES 2014 - 16th International Workshop, Busan, South Korea, September 23-26, 2014. Proceedings*, volume 8731 of *Lecture Notes in Computer Science*, pages 188–205. Springer, 2014.
24. J. Coron, E. Prouff, M. Rivain, and T. Roche. Higher-order side channel security and mask refreshing. In S. Moriai, editor, *Fast Software Encryption - 20th International Workshop, FSE 2013, Singapore, March 11-13, 2013. Revised Selected Papers*, volume 8424 of *Lecture Notes in Computer Science*, pages 410–424. Springer, 2013.
25. J. Coron and A. Tchulkine. A new algorithm for switching from arithmetic to boolean masking. In C. D. Walter, Ç. K. Koç, and C. Paar, editors, *Cryptographic Hardware and Embedded Systems - CHES 2003, 5th International Workshop, Cologne, Germany, September 8-10, 2003, Proceedings*, volume 2779 of *Lecture Notes in Computer Science*, pages 89–97. Springer, 2003.
26. B. Debraize. Efficient and provably secure methods for switching from arithmetic to boolean masking. In E. Prouff and P. Schaumont, editors, *Cryptographic Hardware and Embedded Systems - CHES 2012 - 14th International Workshop, Leuven, Belgium, September 9-12, 2012. Proceedings*, volume 7428 of *Lecture Notes in Computer Science*, pages 107–121. Springer, 2012.
27. D. Dinu, J. Großschädl, and Y. L. Corre. Efficient masking of arx-based block ciphers using carry-save addition on boolean shares. In P. Q. Nguyen and J. Zhou, editors, *Information Security - 20th International Conference, ISC 2017, Ho Chi Minh City, Vietnam, November 22-24, 2017, Proceedings*, volume 10599 of *Lecture Notes in Computer Science*, pages 39–57. Springer, 2017.
28. S. Faust, V. Grosso, S. M. D. Pozo, C. Paglialonga, and F. Standaert. Composable masking schemes in the presence of physical defaults & the robust probing model. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2018(3):89–120, 2018.
29. T. Fritzmann, M. V. Beirendonck, D. B. Roy, P. Karl, T. Schamberger, I. Verbauwhede, and G. Sigl. Masked accelerators and instruction set extensions for post-quantum cryptography. *IACR Cryptol. ePrint Arch.*, 2021:479, 2021.

30. E. Fujisaki and T. Okamoto. Secure integration of asymmetric and symmetric encryption schemes. In M. J. Wiener, editor, *Advances in Cryptology - CRYPTO '99, 19th Annual International Cryptology Conference, Santa Barbara, California, USA, August 15-19, 1999, Proceedings*, volume 1666 of *Lecture Notes in Computer Science*, pages 537–554. Springer, 1999.
31. P. Gao. Formal verification of masking countermeasures for arithmetic programs. In *35th IEEE/ACM International Conference on Automated Software Engineering, ASE 2020, Melbourne, Australia, September 21-25, 2020*, pages 1385–1387. IEEE, 2020.
32. P. Gao, H. Xie, F. Song, and T. Chen. A hybrid approach to formal verification of higher-order masked arithmetic programs. *CoRR*, abs/2006.09171, 2020.
33. P. Gao, H. Xie, J. Zhang, F. Song, and T. Chen. Quantitative verification of masked arithmetic programs against side-channel attacks. In T. Vojnar and L. Zhang, editors, *Tools and Algorithms for the Construction and Analysis of Systems - 25th International Conference, TACAS 2019, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2019, Prague, Czech Republic, April 6-11, 2019, Proceedings, Part I*, volume 11427 of *Lecture Notes in Computer Science*, pages 155–173. Springer, 2019.
34. P. Gao, J. Zhang, F. Song, and C. Wang. Verifying and quantifying side-channel resistance of masked software implementations. *ACM Trans. Softw. Eng. Methodol.*, 28(3):16:1–16:32, 2019.
35. F. Gérard and M. Rossi. An efficient and provable masked implementation of qtesla. In S. Belaïd and T. Güneysu, editors, *Smart Card Research and Advanced Applications - 18th International Conference, CARDIS 2019, Prague, Czech Republic, November 11-13, 2019, Revised Selected Papers*, volume 11833 of *Lecture Notes in Computer Science*, pages 74–91. Springer, 2019.
36. B. Gigerl, V. Hadzic, R. Primas, S. Mangard, and R. Bloem. Coco: Co-Design and Co-Verification of Masked Software Implementations on CPUs. *30th USENIX Security Symposium, USENIX Security 2021*, 2021.
37. B. Gigerl, R. Primas, and S. Mangard. Secure and efficient software masking on superscalar pipelined processors. In *Advances in Cryptology - ASIACRYPT 2021*, 2021.
38. G. Goodwill, B. Jun, J. Jaffe, and P. Rohatgi. A testing methodology for side-channel resistance validation. In *NIST Non-Invasive Attack Testing Workshop*, 2011.
39. L. Goubin. A sound method for switching between boolean and arithmetic masking. In Ç. K. Koç, D. Naccache, and C. Paar, editors, *Cryptographic Hardware and Embedded Systems - CHES 2001, Third International Workshop, Paris, France, May 14-16, 2001, Proceedings*, volume 2162 of *Lecture Notes in Computer Science*, pages 3–15. Springer, 2001.
40. H. Groß, R. Iusupov, and R. Bloem. Generic low-latency masking in hardware. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2018(2):1–21, 2018.
41. H. Groß, S. Mangard, and T. Korak. Domain-oriented masking: Compact masked hardware implementations with arbitrary protection order. In *Proceedings of the ACM Workshop on Theory of Implementation Security, TIS@CCS 2016 Vienna, Austria, October, 2016*, page 3. ACM, 2016.
42. V. Hadzic and R. Bloem. COCOALMA: A versatile masking verifier. In *Formal Methods in Computer Aided Design, FMCAD 2021, New Haven, CT, USA, October 19-22, 2021*, pages 1–10. IEEE, 2021.
43. M. Hutter and M. Tunstall. Constant-time higher-order boolean-to-arithmetic masking. *J. Cryptogr. Eng.*, 9(2):173–184, 2019.

44. Y. Ishai, A. Sahai, and D. A. Wagner. Private circuits: Securing hardware against probing attacks. In *Advances in Cryptology - CRYPTO 2003, 23rd Annual International Cryptology Conference, Santa Barbara, California, USA, August 17-21, 2003, Proceedings*, volume 2729 of *Lecture Notes in Computer Science*, pages 463–481. Springer, 2003.
45. M. Karroumi, B. Richard, and M. Joye. Addition with blinded operands. In E. Prouff, editor, *Constructive Side-Channel Analysis and Secure Design - 5th International Workshop, COSADE 2014, Paris, France, April 13-15, 2014. Revised Selected Papers*, volume 8622 of *Lecture Notes in Computer Science*, pages 41–55. Springer, 2014.
46. D. Knichel, P. Sasdrich, and A. Moradi. SILVER - statistical independence and leakage verification. In S. Moriai and H. Wang, editors, *Advances in Cryptology - ASIACRYPT 2020 - 26th International Conference on the Theory and Application of Cryptology and Information Security, Daejeon, South Korea, December 7-11, 2020, Proceedings, Part I*, volume 12491 of *Lecture Notes in Computer Science*, pages 787–816. Springer, 2020.
47. P. C. Kocher, J. Jaffe, and B. Jun. Differential power analysis. In *CRYPTO*, volume 1666 of *Lecture Notes in Computer Science*, pages 388–397. Springer, 1999.
48. S. Mangard, T. Popp, and B. M. Gammel. Side-channel leakage of masked CMOS gates. In A. Menezes, editor, *Topics in Cryptology - CT-RSA 2005, The Cryptographers' Track at the RSA Conference 2005, San Francisco, CA, USA, February 14-18, 2005, Proceedings*, volume 3376 of *Lecture Notes in Computer Science*, pages 351–365. Springer, 2005.
49. M. M. Mano. *Computer system architecture*. Prentice Hall, 1982.
50. Q. L. Meunier, E. Pons, and K. Heydemann. Leakageverif: Scalable and efficient leakage verification in symbolic expressions. *IACR Cryptol. ePrint Arch.*, page 1468, 2021.
51. P. L. Montgomery. Modular multiplication without trial division. *Mathematics of computation*, 44(170):519–521, 1985.
52. T. Moos, A. Moradi, T. Schneider, and F. Standaert. Glitch-resistant masking revisited or why proofs in the robust probing model are needed. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2019(2):256–292, 2019.
53. National Institute of Standards and Technology (NIST). FIPS-180-2: Secure Hash Standard, 2002.
54. O. Neiße and J. Pulkus. Switching blindings with a view towards IDEA. In M. Joye and J. Quisquater, editors, *Cryptographic Hardware and Embedded Systems - CHES 2004: 6th International Workshop Cambridge, MA, USA, August 11-13, 2004. Proceedings*, volume 3156 of *Lecture Notes in Computer Science*, pages 230–239. Springer, 2004.
55. T. Oder, T. Schneider, T. Pöppelmann, and T. Güneysu. Practical cca2-secure and masked ring-lwe implementation. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2018(1):142–174, 2018.
56. R. O'Donnell. *Analysis of Boolean Functions*. Cambridge University Press, 2014.
57. J. Quisquater and D. Samyde. Electromagnetic analysis (EMA): measures and counter-measures for smart cards. In *E-smart*, volume 2140 of *Lecture Notes in Computer Science*, pages 200–210. Springer, 2001.
58. O. Reparaz, B. Bilgin, S. Nikova, B. Gierlichs, and I. Verbauwhede. Consolidating masking schemes. In *Advances in Cryptology - CRYPTO 2015 - 35th Annual Cryptology Conference, Santa Barbara, CA, USA, August 16-20, 2015, Proceedings, Part I*, volume 9215 of *Lecture Notes in Computer Science*, pages 764–783. Springer, 2015.

59. T. Schneider, A. Moradi, and T. Güneysu. Arithmetic addition over boolean masking - towards first- and second-order resistance in hardware. In T. Malkin, V. Kolesnikov, A. B. Lewko, and M. Polychronakis, editors, *Applied Cryptography and Network Security - 13th International Conference, ACNS 2015, New York, NY, USA, June 2-5, 2015, Revised Selected Papers*, volume 9092 of *Lecture Notes in Computer Science*, pages 559–578. Springer, 2015.

60. T. Schneider, C. Paglialonga, T. Oder, and T. Güneysu. Efficiently masking binomial sampling at arbitrary orders for lattice-based crypto. In D. Lin and K. Sako, editors, *Public-Key Cryptography - PKC 2019 - 22nd IACR International Conference on Practice and Theory of Public-Key Cryptography, Beijing, China, April 14-17, 2019, Proceedings, Part II*, volume 11443 of *Lecture Notes in Computer Science*, pages 534–564. Springer, 2019.

61. M. S. Turan, K. McKay, D. Chang, Ç. Çalık, L. Bassham, J. Kang, and J. Kelsey. Status report on the second round of the nist lightweight cryptography standardization process. Technical report, Tech. rep. https://doi. org/10.6028/NIST. IR. 8369. Gaithersburg, MD, USA . . . , 2021.

A Iterative and unrolled circuits

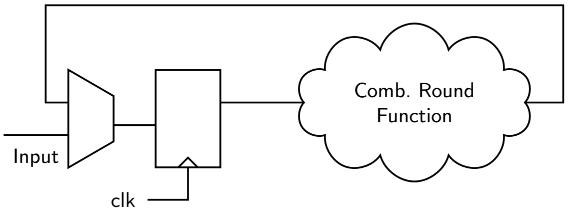

Fig. 3: Iterative circuit [12]

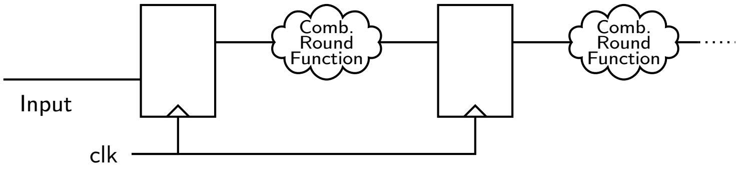

Fig. 4: Unrolled circuit [12]

B Fourier Expansion of the Arithmetic Addition

Recall the Fourier expansion of the AND, OR and XOR functions:

$$\begin{aligned}
&\textsc{And} \quad W(a \wedge b) = \frac{1}{2} + \frac{1}{2}a + \frac{1}{2}b - \frac{1}{2}ab \\
&\textsc{Or} \quad W(a \vee b) = -\frac{1}{2} + \frac{1}{2}a + \frac{1}{2}b + \frac{1}{2}ab \\
&\textsc{Xor} \quad W(a \oplus b) = ab
\end{aligned}$$

Additionally, note that Fourier expansions represent Boolean functions as a polynomial over the real domain $\{1, -1\}$, where 1 represents FALSE and -1 represents TRUE. Consequently, monomials x^c with even exponents c evaluate to 1 in Fourier expansions. The Fourier expansion of the carry and sum can hence be expressed as:

$$
\begin{aligned}
\textsc{Carry}\quad & W(c^{(j)}) = W((u^{(j)} \oplus u^{(j)}) \wedge c^{(j-1)}) \vee (u^{(j)} \wedge u^{(j)})) \\
& = -(0.25u^{(j)})^2(u^{(j)})^2c^{(j-1)} - 0.25(u^{(j)})^2(u^{(j)})^2 - 0.25(u^{(j)})^2u^{(j)}c^{(j-1)} - 0.25u^{(j)}(u^{(j)})^2c^{(j-1)} \\
& + (0.25u^{(j)})^2u^{(j)} + 0.25u^{(j)}(u^{(j)})^2 - 0.5u^{(j)}u^{(j)}c^{(j-1)} + 0.25u^{(j)}c^{(j-1)} + 0.25u^{(j)}c^{(j-1)} \\
& + 0.25u^{(j)} + 0.25u^{(j)} + 0.25c^{(j-1)} + 0.25 \\
& = 0.25c^{(j-1)} - 0.25 - 0.25u^{(j)}c^{(j-1)} - 0.25u^{(j)}c^{(j-1)} + 0.25u^{(j)} + 0.25u^{(j)} \\
& - 0.5u^{(j)}u^{(j)}c^{(j-1)} + 0.25u^{(j)}c^{(j-1)} + 0.25u^{(j)}c^{(j-1)} \\
& + 0.25u^{(j)} + 0.25u^{(j)} + 0.25c^{(j-1)} + 0.25 \\
& = 0.5c^{(j-1)} + 0.5u^{(j)} + 0.5u^{(j)} - 0.5u^{(j)}u^{(j)}c^{(j-1)} \\
& W(c[0]) = 1 \\
\textsc{Sum}\quad & W(sum^{(j)}) = W(W(u^{(j)} \oplus u^{(j)}) \oplus c^{(j)}) \\
& = W(u^{(j)}u^{(j)} \oplus c^{(j)}) \\
& = u^{(j)}u^{(j)}W(c^{(j)}) \\
& = 0.5u^{(j)}u^{(j)}c^{(j)} + 0.5u^{(j)} + 0.5u^{(j)} - 0.5c^{(j)}
\end{aligned}
$$

C Coron et al.-A2B

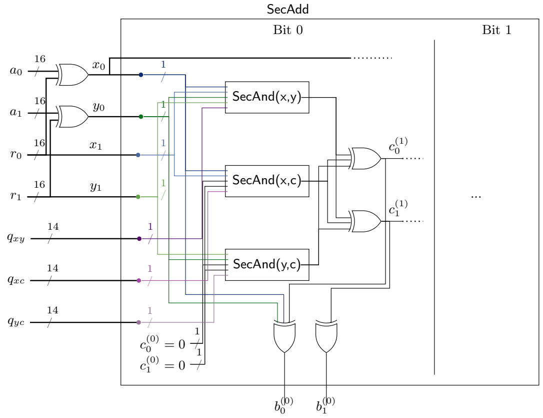

Fig. 5: Schematic image of Coron et al.-A2B [23] when implemented in hardware. The arithmetic input shares a_0, a_1 are transformed into Boolean shares b_0, b_1. The carry computation happens in the SecAdd module, from which we draw the first part responsible for bits 0 of the final result.

D Sanity Check Measurement Setup (RNG Off)

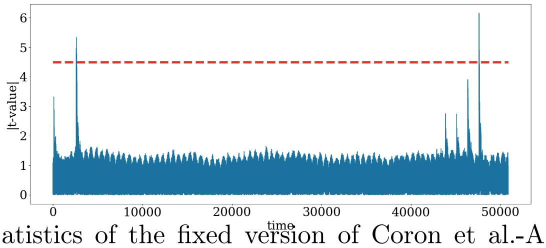

Fig. 6: T-test statistics of the fixed version of Coron et al.-A2B with 400 000 traces and RNG off.

E Schneider et al.-B2A [60]

Algorithm 1 Schneider et al. B2A [60] (simplified for 1st order)

```
Input: k-bit shares $b_0, b_1$ such that $b = b_0 \oplus b_1$
Output: Shares $a_0, a_1 \in \mathbb{F}_q$ such that $b = a_0 + a_1 \mod q$
$b'_0 \leftarrow b_0^{(k-1)}$
$b'_1 \leftarrow b_1^{(k-1)}$
$a_0, a_1 \leftarrow \text{B2A\_Bit}(b'_0, b'_1)$
$R \overset{\$}{\leftarrow} \mathcal{R}_q$
$a_0 \leftarrow (a_0 + R) \mod q$
$a_1 \leftarrow (a_1 - R) \mod q$
for $j = 2$ to $k - 1$ do
    $b'_0 \leftarrow b_0^{(k-j)}$
    $b'_1 \leftarrow b_1^{(k-j)}$
    $C_0, C_1 \leftarrow \text{B2A\_Bit}(b'_0, b'_1)$
    $R \overset{\$}{\leftarrow} \mathcal{R}_q$
    $C_0 \leftarrow (C_0 + R) \mod q$
    $a_0 \leftarrow ((a_0 << 1) + C_0) \mod q$
    $C_1 \leftarrow (C_1 - R) \mod q$
    $a_1 \leftarrow ((a_1 << 1) + C_1) \mod q$
end for
return $a_0, a_1$
```

Algorithm 2 Schneider et al. B2A_Bit [60] (simplified for 1st order)

```
Input: 1-bit shares $b'_0, b'_1$ such that $b = b'_0 \oplus b'_1$
Output: $E_0, E_1$ such that $E_0 + E_1 = b \mod q$
$E_0 \overset{\$}{\leftarrow} \mathcal{R}_q$
$E_1 \leftarrow b'_1 - E_0 \mod q$
$E_1 \leftarrow E_1 - 2 \cdot (E_1 \cdot b'_0) \mod q$
$E_0 \leftarrow E_0 - 2 \cdot (E_0 \cdot b'_1) \mod q$
$E_1 \leftarrow E_1 + b'_1 \mod q$
return$E_1, E_0$
```

F Debraize-A2B

Algorithm 3 Table T generation [26]

```
Input: $k$
Output: Conversion table $T$, random variables $r, \rho$
$r \leftarrow \mathcal{U}(0,1)^k$
$\rho \leftarrow \mathcal{U}(0,1)$
for $i \leftarrow 0$ to $2^k - 1$ do
    $T[\rho||i] \leftarrow (i + r) \oplus (\rho||r)$
    $T[(\rho \oplus 1)||i] \leftarrow (i + r + 1) \oplus (\rho||r)$
end for
return $T, r, \rho$
```

Algorithm 4 Debraize-A2B [26]

```
Input: $(n \cdot k)$-bit shares $a_0, a_1$ such that $a = a_0 + a_1 \mod 2^{(n \cdot k)}, T, r, \rho$
Output: $(n \cdot k)$-bit shares $b_0, b_1$ such that $a = b_0 \oplus b_1$
$a_0 \leftarrow a_0 - (r||...||r||...||r) \mod 2^{n \cdot k}$
$\beta \leftarrow \rho$
for $i \leftarrow 0$ to $n - 1$ do
    Split $a_0$ into $(a_{0h}||a_{0l})$, split $a_1$ into $(a_{1h}||a_{1l})$
    $a_0 \leftarrow a_0 + a_{1l} \mod 2^{(n-i) \cdot k}$
    $\beta||x'_i \leftarrow T[\beta||a_{0l}]$
    $x'_i \leftarrow x'_i \oplus a_{1l}$
    $a_0 \leftarrow a_{0h}, a_1 \leftarrow a_{1h}$
end for
$b_0 = (x'_0||...||x'_i||...||x'_{n-1}) \oplus (r||...||r||...||r)$
$b_1 = a_1$
return $b_0, b_1$
```

G Goubin [39]

Algorithm 5 Goubin-A2B [39]

```
Input: n-bit shares $a_0, a_1$ such that $a = a_0 + a_1 \bmod 2^n$
Output: n-bit shares $b_0, b_1$ such that $a = b_0 \oplus b_1$
 1: $Y \leftarrow \mathcal{U}(0,1)^n$
 2: $T \leftarrow 2Y$
 3: $b_0 \leftarrow Y \oplus a_1$
 4: $\Omega \leftarrow Y \wedge b_0$
 5: $b_0 \leftarrow T \oplus a_0$
 6: $Y \leftarrow Y \oplus b_0$
 7: $Y \leftarrow Y \wedge a_1$
 8: $\Omega \leftarrow \Omega \oplus Y$
 9: $Y \leftarrow T \wedge a_0$
10: $\Omega \leftarrow \Omega \oplus Y$
11: for $i \leftarrow 0$ to $n-1$ do
12:     $Y \leftarrow T \wedge a_1$
13:     $Y \leftarrow Y \oplus \Omega$
14:     $T \leftarrow T \wedge a_0$
15:     $Y \leftarrow Y \oplus T$
16:     $T \leftarrow 2Y$
17: end for
18: $b_0 \leftarrow b_0 \oplus T$
19: $b_1 \leftarrow a_1$
20: return $b_0, b_1$
```

Algorithm 6 Goubin-B2A [39]

```
Input: n-bit shares $b_0, b_1$ such that $a = b_0 \oplus b_1$
Output: n-bit shares $a_0, a_1$ such that $a = a_0 + a_1 \bmod 2^n$
 1: $Y \leftarrow \mathcal{U}(0,1)^n$
 2: $T \leftarrow b_0 \oplus Y$
 3: $T \leftarrow T - Y$
 4: $T \leftarrow T \oplus b_0$
 5: $Y \leftarrow Y \oplus b_1$
 6: $a_0 \leftarrow b_0 \oplus Y$
 7: $a_0 \leftarrow a_0 - Y$
 8: $a_0 \leftarrow a_0 \oplus T$
 9: $a_1 \leftarrow b_1$
10: return $a_0, a_1$
```

Overwrite leakages In line 9 of the algorithm, the attacker probes the reassignment of Y:

$$\begin{aligned}
Y_{\text{old}} &= Y_{\text{line 6}} \wedge a_1 \\
&= (Y_{\text{line 1}} \oplus b_{0\text{line 5}}) \wedge a_1 \\
&= (Y_{\text{line 1}} \oplus (T \oplus a_0)) \wedge a_1 \\
Y_{\text{new}} &= T \wedge a_0 \\
Y_{\text{old}} \oplus Y_{\text{new}} &= ((Y_{\text{line 1}} \oplus (T \oplus a_0)) \wedge a_1) \oplus (T \wedge a_0) \\
&= (a_0 \wedge a_1) \oplus (a_0 \wedge T) \oplus (a_1 \wedge Y)
\end{aligned}$$

Hence, for every bit $>= 0$, this expression will correlate with native value a.

Another similar situation occurs in Figure 12 where $Y_{\text{old}} = T \wedge a_0$ is overwritten by $Y_{\text{new}} = T \wedge a_1$ in the first loop iteration.

False positive in Goubin-A2B Assume the attacker probes the expression $\Omega \oplus Y_{\text{line 9}}$ in line 10, which is $(Y^{(0)} \wedge (Y^{(0)} \oplus a_1^{(0)})) \oplus (a_1^{(0)} \wedge (Y^{(0)} \oplus a_0^{(0)}))$. For reasons of readability, we omit to indicate that we always refer to the LSB, i.e., skip $^{(0)}$.

Exact Fourier expansion

$$\begin{aligned}
W((Y \wedge (Y \oplus a_1)) \oplus ((Y \oplus a_0) \wedge a_1)) &=? \\
W(Y \oplus a_0) &= Y a_0 \\
W(Y \oplus a_1) &= Y a_1 \\
W(Y \wedge (Y \oplus a_1)) &= -0.5\, Y^2 a_1 + 0.5\, Y a_1 + 0.5\, Y + 0.5 \\
&= -0.5\, a_1 + 0.5\, Y a_1 + 0.5\, Y + 0.5 \\
W((Y \oplus a_0) \wedge a_1)) &= -0.5\, Y a_0 a_1 + 0.5\, Y a_0 + 0.5\, a_1 + 0.5 \\
W((Y \wedge (Y \oplus a_1)) \oplus ((Y \oplus a_0) \wedge a_1)) &= -0.25\, Y^2 a_0 a_1^2 + 0.25\, Y a_0 a_1^2 + 0.25\, Y^2 a_0 - 0.5\, Y a_0 a_1 \\
&\quad + 0.25\, Y a_1^2 + 0.25\, Y a_0 + 0.50\, Y a_1 - 0.25\, a_1^2 + 0.25\, Y + 0.25 \\
&= -0.25\, a_0 + 0.25\, Y a_0 + 0.25\, a_0 - 0.5\, Y a_0 a_1 \\
&\quad + 0.25\, Y + 0.25\, Y a_0 + 0.50\, Y a_1 - 0.25 + 0.25\, Y + 0.25
\end{aligned}$$

Approximated Fourier expansion

$$\begin{aligned}
\mathcal{C}((Y \wedge (Y \oplus a_1)) \oplus ((Y \oplus a_0) \wedge a_1)) &=? \\
\mathcal{C}(Y \oplus a_0) &= \{\{Y, a_0\}\} \\
\mathcal{C}(Y \oplus a_1) &= \{\{Y, a_1\}\} \\
\mathcal{C}(Y \wedge (Y \oplus a_1)) &= \{\{1\}, \{Y\}, \{Y, a_1\}, \{a_1\}\} \\
\mathcal{C}((Y \oplus a_0) \wedge a_1)) &= \{\{1\}, \{Y, a_0\}, \{a_1\}, \{Y, a_0, a_1\}\} \\
\mathcal{C}((Y \wedge (Y \oplus a_1)) \oplus ((Y \oplus a_0) \wedge a_1)) &= \mathcal{C}((Y \oplus a_0) \wedge a_1)) \otimes \mathcal{C}(Y \wedge (Y \oplus a_1))\} \\
&= \{\{1\}, ...\{Y^2, a_0, a_1\}, ...\}
\end{aligned}$$

Note: $Y^2 = 1$ because in Fourier expression each element is either 1 (False) or -1 (True).

H Table lookup on gate-level

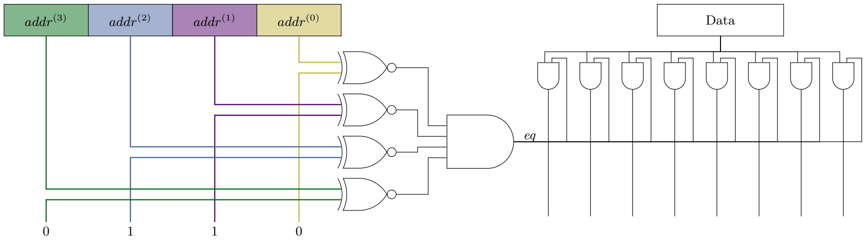

Fig. 7: Example of table lookup including equality comparator on gate-level with 4-bit addresses and 8-bit data words. The address *addr* is compared to the constant address of the SRAM cell $((0110)_b)$. If both values are equal, the resulting 1-bit signal *eq* is 1, and 0 otherwise. *eq* is further used to decide whether the respective data word should be read or not.