High-Throughput Side-Channel-Protected Stream Cipher Hardware for 6G Systems

Yuluan Cao[1,2†][0009−0000−2390−1280], Cankun Zhao[1,2†][0000−0002−6875−3557], Bohan Yang[1,2][0000−0002−5204−1707], Wenping Zhu[1,2][0000−0002−3276−4019], Hanning Wang[1,2][0000−0003−1117−2740], Min Zhu[3][0009−0005−1807−0810], and Leibo Liu[1,2*][0000−0001−7548−4116]

[1] Beijing National Research Center for Information Science and Technology, School of Integrated Circuits, Tsinghua University.
[2] State Key Laboratory of Cryptography and Digital Economy Security, Tsinghua University, Beijing, 100084, China.
cao-yl25@mails.tsinghua.edu.cn
{zhaock,bohanyang,zhuwp,wanghn,liulb}@tsinghua.edu.cn
[3] Wuxi Micro Innovation Integrated Circuit Design Co., Ltd, Jiangsu Wuxi, China.
zhumin@mucse.com
† These authors contributed equally to this work. * Corresponding authors.

Abstract. Emerging 6G communication systems impose unprecedented requirements on cryptographic primitives, demanding ultra-high throughput, low latency, and strong resistance to implementation-level attacks. LOL2.0 is a recently proposed stream cipher framework that achieves high software efficiency and strong security in post-quantum settings. While several stream ciphers have been proposed to address performance demands, side-channel-protected hardware implementations capable of sustaining 6G-class throughput remain largely unexplored.

In this work, we present the first side-channel-protected hardware implementation that meets the 6G-class throughput demand. Focusing on the LOL2.0 stream cipher framework, we leverage Time Sharing Masking to achieve first-order security under the glitch-extended probing model. This design realizes full-phase protection covering initialization, keystream generation, and tag generation. To address diverse deployment requirements, we design two masked architectures: a compact variant optimized for area and randomness efficiency, and a fast variant targeting the maximum achievable throughput.

The proposed fast implementations achieve peak throughputs of 183 Gbps and 142 Gbps for the unmasked and masked configurations, respectively. Meanwhile, the compact architecture reduces hardware cost by achieving areas as low as 23.25 kGE and 141.98 kGE in unmasked and masked designs, respectively, while still maintaining competitive throughput. Security is validated through practical side-channel evaluations using Test Vector Leakage Assessment on FPGA platforms. Across up to 100 million measured power traces, no statistically significant first-order leakage is observed for any protected configuration.

Overall, this work realizes side-channel-protected stream cipher hardware that sustains ultra-high throughput, providing a concrete path toward secure cryptographic deployment in future 6G communication systems.

Keywords: Stream Cipher · Side-Channel · Countermeasure · Hardware Security · Hardware Implementation.

1 Introduction

Nowadays, fifth-generation (5G) mobile networks have entered large-scale deployment, enabling the interconnection of massive numbers of devices. Looking beyond 5G, sixth-generation (6G) communication systems are currently being standardized under the IMT-2030 framework. IMT-2030 [24], established by the International Telecommunication Union (ITU), defines the overarching vision and objectives for the development of International Mobile Telecommunications (IMT) systems toward 2030 and beyond.

Compared with 5G, 6G systems are expected to support peak throughputs ranging from 100 Gbps to 1 Tbps [29]. Such Tbps-class throughput is widely regarded as essential for enabling next-generation applications, such as virtual and augmented reality [2]. Beyond high performance, security is also explicitly identified by the ITU as a core capability of IMT-2030 systems [24]. In this context, security extends beyond data confidentiality to encompass data integrity and authentication, which are essential for end-to-end trustworthy communication. As a result, Authenticated Encryption with Associated Data (AEAD) has become a fundamental cryptographic primitive for such application scenarios.

Moreover, the advent of quantum attacks threatens the security of cryptographic primitives [10, 26]. In particular, Kaplan et al. [26] showed that several structured symmetric-key constructions can be attacked using Simon's algorithm [41], indicating that, for such constructions, doubling the key length alone may be insufficient to restore the expected security level. Therefore, in the post-quantum setting, the design of symmetric primitives should take structural security into account rather than relying solely on parameter-level adjustments.

Under the dual requirements of data rate and security, LOL is a general stream cipher framework proposed by Feng et al. to meet the performance demands of 6G systems while providing strong security for post-quantum cryptography [13]. LOL2.0 is an enhanced version of LOL that significantly improves its security guarantees [14]. In addition, the authors proposed SCMAC, a generic construction that enables the transformation of stream ciphers into AEAD schemes.

Although modern cryptographic algorithms are widely regarded as mathematically secure in theory, security in practical deployments is severely threatened by implementation-level vulnerabilities. Among these, side-channel attacks (SCAs) represent one of the most powerful attacks. 6G application scenarios significantly amplify the practical threat posed by side-channel attacks. The diversity of 6G devices—ranging from resource-constrained end nodes to physically accessible edge devices—makes it easier for adversaries to observe physical leakages.

Consequently, in cryptographic implementations for 6G systems, side-channel resistance must be regarded as a core design objective on par with throughput.

Achieving verifiable side-channel security while simultaneously meeting data-rate requirements on the order of hundreds of Gbps has emerged as a critical challenge for cryptographic hardware implementations in 6G communication systems.

1.1 Related Work

Cryptographic Primitives. Several cryptographic primitives have been proposed or improved in the context of 6G systems, including Rocca, SNOW-V, AEGIS, and LOL/LOL2.0. Based on the LOL2.0 framework, two concrete cipher instances, namely LOL2.0-Mini and LOL2.0-Double, were instantiated, both of which support stream cipher (SC) and AEAD modes.

SNOW-Vi achieves a peak throughput of up to 92 Gbps [11], while Rocca reaches 154.03 Gbps [14, 37]. AEGIS achieves a throughput of approximately 90.71 Gbps [14, 44]. By contrast, LOL2.0-Mini and LOL2.0-Double achieve throughputs of approximately 90 Gbps and 144 Gbps, respectively.

Although SNOW-V/Vi demonstrates strong software performance, it has been shown to be vulnerable to correlation attacks [40], and its GCM-based AEAD construction exhibits limited efficiency. While Rocca outperforms LOL2.0-Double in terms of throughput, its reliance on AES round functions exposes it to existing key-recovery attacks [22]. AEGIS has also been shown to be vulnerable to linear distinguishers [32]. Moreover, despite claims of 256-bit key support by SNOW-V/Vi, AEGIS, and Rocca, existing cryptanalytic results indicate that these constructions may fall short of the theoretical 256-bit security level [14]. In contrast, the design of LOL2.0 avoids these security risks by construction. It enhances resistance against differential attacks, division attacks, and fault correlation analysis [14].

Side-Channel-Protected Implementation. Side-channel protection for cryptographic algorithms has been extensively studied in recent years. Among the various countermeasures, masking, first introduced by Goubin et al. [16], is one of the most effective defenses, as it can provide provable security under leakage models. While extensive efforts have been devoted to hardware implementations of cryptographic algorithms, particularly block ciphers, comparatively few works address complete circuit-level protection for SC or AEAD schemes, particularly in high-throughput settings.

Saurabh et al. [39] presented the first power side-channel attack on SNOW-V and further proposed Boolean masking countermeasures, which were validated through experimental measurements. For AES-round-based encryption primitives such as Rocca and AEGIS, the prevailing approach to side-channel protection is to apply secure masking schemes to the AES round function or its S-boxes. Representative techniques include threshold implementations (TI) [35] and domain-oriented masking (DOM) [19]. For lightweight AEAD schemes such as Ascon, Adomnicai et al. [1] studied masked software implementations of ACORN and Ascon under the assumption that the secret key cannot be recovered from the internal state after initialization, thereby protecting only the

initialization phase to reduce implementation overhead. More recently, Prasad et al. proposed a low-latency masking scheme for Ascon that leverages a changing of the guards strategy to eliminate the need for fresh randomness [36]. They also protect and perform leakage assessment solely on the initialization stage.

However, side-channel-protected implementations for high-throughput 6G-oriented ciphers remain notably scarce. Mentens et al. [33] implemented a TI-protected AES-GCM design on a Virtex-7 FPGA, achieving a maximum throughput of 15.24 Gbps at 119 MHz. This result highlights the substantial performance gap between state-of-the-art masked implementations and the hundreds-of-Gbps throughput targets envisioned for 6G systems.

1.2 Motivation

Compared with other stream ciphers designed for 6G scenarios, LOL2.0 demonstrates competitive performance and stands out as a promising candidate capable of simultaneously satisfying the stringent security and throughput requirements of future 6G systems. Since LOL2.0 comprises two cipher instances, LOL2.0-Mini and LOL2.0-Double, both of which follow the same design principles and employ identical nonlinear components, the masking architectures for LOL2.0-Mini can be straightforwardly extended to the double mode. Thus, we focus on LOL2.0-Mini, which captures the core update structure of LOL2.0, and leave the extension to LOL2.0-Double for future work. Unless otherwise stated, all references to LOL2.0 in the remainder of this paper refer to LOL2.0-Mini.

Meanwhile, existing works have made substantial progress in side-channel protection for SC and AEAD schemes. However, achieving provably secure and low-latency side-channel protection at ultra-high throughput remains an open challenge, particularly for designs targeting 6G-scale performance.

Traditionally, the keystream generated after initialization is assumed to behave like a pseudo-random sequence, leading to the belief that the initialization phase is the most vulnerable to side-channel attacks [21]. Consequently, many prior works apply masking only to the initialization phase. Recent studies, however, indicate that the keystream generation phase can also be successfully attacked using machine-learning-based and other AI-assisted techniques [28].

Therefore, in this work, we adopt a full-phase side-channel protection strategy for LOL2.0 and present a low-latency hardware implementation targeting the ultra-high-throughput requirements of 6G systems.

1.3 Our Contributions

Motivated by the throughput and security requirements of 6G systems, we present a comprehensive hardware design space exploration of LOL2.0, covering both unmasked and masked implementations and spanning performance-oriented and resource-constrained architectures to support diverse 6G application scenarios. Our main contributions are summarized as follows:

1. **First high-throughput side-channel-protected stream cipher hardware targeting 6G scenarios.** We present the first side-channel-protected hardware implementation of a stream cipher that achieves throughput exceeding the 100 Gbps target envisioned for 6G systems. While the unmasked implementation reaches a peak throughput of 182.86 Gbps, the proposed first-order masked implementation sustains up to 142.22 Gbps, delivering ultra-high throughput together with strong resistance to side-channel attacks.
2. **First-order full-phase masked stream cipher hardware implementation based on Time Sharing Masking (TSM).** We present first-order masked hardware implementation of the LOL2.0 framework using TSM [43]. The masked architectures provide full-phase protection across initialization, keystream generation, and authentication tag generation. To the best of our knowledge, this is the first application of TSM to stream cipher and AEAD hardware design, achieving glitch-extended PINI security while maintaining low latency and ultra-high throughput.
3. **Compact and fast hardware architectures for diverse deployment scenarios.** We design multiple hardware variants of LOL2.0, covering both unmasked and first-order masked implementations and targeting distinct application scenarios through compact and fast architectures. The compact architecture is optimized for resource-constrained environments by prioritizing implementation cost, while the fast architecture targets high-performance applications by maximizing throughput through parallelism and pipelining. All designs support both SC and AEAD modes. Together, these variants form an extensible hardware design space, enabling secure deployment across diverse 6G application scenarios.

2 Preliminaries

2.1 System Requirements of 6G

According to the ITU, the 6G mobile communication system corresponds to the next generation of International Mobile Telecommunications standardized under the IMT-2030 framework. IMT-2020 defines the performance targets for 5G systems, specifying a peak downlink data rate of 20 Gbit/s and a peak uplink data rate of 10 Gbit/s [23]. Building upon 5G, IMT-2030 further expands system capabilities and introduces a diverse set of usage scenarios, including integrated sensing and communication, immersive communication, AI-native communication, massive connectivity, ubiquitous connectivity, and hyper-reliable low-latency communication (HRLLC). These emerging scenarios impose increasingly stringent requirements on system performance, with target data rates exceeding 100 Gbps. In addition, these end-to-end application scenarios also impose heightened security requirements. Reliable communication must be ensured not only within the network core, but also at the network edge and end devices, where attack surfaces are substantially expanded [9].

2.2 Masking

Masking is a well-established countermeasure against side-channel attacks, in which a secret value v is split into n statistically independent random shares $v_1, \ldots, v_n$, such that all computations are carried out on these shares rather than on the secret itself [25]. Formally, the secret is reconstructed as $v = v_1 \circ v_2 \circ \cdots \circ v_n$, where $\circ$ denotes the share-combining operator. The exact meaning of $\circ$ depends on the masking scheme. Boolean masking is commonly adopted in symmetric cryptographic algorithms and is particularly well-suited for securing linear operations. It represents a secret value v as the bitwise exclusive-OR of n uniformly random shares, i.e., $v = v_1 \oplus v_2 \oplus \cdots \oplus v_n$, where $\oplus$ denotes the bitwise XOR operation. In contrast, masking nonlinear operations is considerably more challenging. Nonlinear transformations generally introduce statistical dependencies among shares, which may lead to information leakage. As a result, secure masked implementations of nonlinear operations typically require additional fresh randomness to preserve the independence of the shares and maintain the desired security guarantees.

2.3 d-Probing Model

We adopt the d-probing model to characterize the side-channel adversary. The standard probing model, first proposed by Ishai et al. [25], provides a formal abstraction for analyzing the resistance of masked circuits under idealized leakage assumptions. A probing order d means that the adversary can observe the internal signals of up to d nodes in the circuit, assuming ideal probes. A circuit is said to be d-probing secure if any combination of up to d intermediate variables is statistically independent of the secret.

However, because real circuits exhibit physical imperfections such as glitches, Faust et al. [12] extended it to the glitch-extended probing model. In this model, a single probe may observe not only the stable output of its target gate but also combinations of upstream transient signals until the last register boundary.

Directly constructing large circuits that satisfy d-probing security is extremely challenging, and the composition of d-probing-secure gadgets does not necessarily yield a globally d-probing-secure circuit [8]. To address this issue, Barthe et al. [4] introduced the notion of Non-Interference (NI) to ensure that security is preserved when composing masking gadgets into larger circuits. They later proposed Strong Non-Interference (SNI) [5], which further restricts probe propagation but incurs substantial implementation overhead due to its conservative assumptions.

Probe-Isolating Non-Interference (PINI), proposed by Cassiers et al. [7], offers a more efficient security notion. PINI isolates probe propagation within a single share domain, reducing area overhead compared to SNI while still guaranteeing composability. A particularly useful property of PINI is that any composition of PINI gadgets remains PINI, making it highly suitable for constructing large masked circuits.

2.4 Low-Latency Masking

For high-throughput stream ciphers, latency is the dominant performance metric, as it directly determines the critical path and limits the achievable keystream generation rate. All state-of-the-art masking schemes inevitably incur additional latency when protecting high-degree nonlinear functions. In practice, this often results in one or more additional clock cycles along the critical nonlinear paths, significantly degrading throughput. Consequently, low-latency masking schemes are essential in the context of stream cipher implementations. We briefly summarize representative low-latency masking paradigms and compare their properties to motivate our choice of TSM.

Sasdrich et al. [38] presented a single-cycle masked AES implementation based on LUT-based Masked Dual-Rail with Pre-charge Logic (LMDPL). This design achieves first-order security under the d-glitch-extended probing model and exhibits a certain degree of resistance against higher-order attacks. However, no formal higher-order masking security guarantees are provided and the approach relies on a dedicated pre-charge phase. The masked AES S-box proposed by Gross et al. [17] follows a circuit-level domain separation strategy within the probing model. Although effective in reducing latency, the construction does not support composable security, which limits its applicability in large-scale masked designs. Knichel et al. [27] introduced GHPC_{LL}, which supports the glitch probing model and achieves PINI security. However, this scheme incurs a substantial overhead in terms of both randomness consumption and area, making it less attractive for resource-constrained or high-throughput implementations. LLTI [3] is a low-latency masking scheme based on threshold implementations. It requires no fresh randomness but does not provide composability guarantees.

TSM is a recently proposed low-latency masking scheme introduced by Kumar et al. [42]. TSM provides first-order security under the glitch-extended probing model and satisfies the PINI composability requirement. The central idea of TSM is to temporally separate the processing of $share_0$ and $share_1$ using a pipeline register. $share_0$ is processed in the first stage while $share_1$ is only processed in the second stage. The intermediate register layer enforces temporal isolation between the shares, effectively preventing glitch propagation across domains. The authors later introduced higher-order TSM in [43], which strictly generalizes TSM to any probing order d while preserving the same architectural principles.

3 Hardware Implementations

3.1 SCMAC and LOL2.0

SCMAC represents an intermediate design paradigm between Encrypt-then-MAC and single-pass AEAD schemes. It partially integrates encryption and authentication mechanisms while mitigating state-leakage risks associated with immediate data absorption and squeezing. SCMAC adopts a dual-state structure, in which the internal encryption state **S** is responsible for keystream generation, while the authentication state **E** accumulates authentication information.

These two states are updated in parallel but remain isolated to prevent unintended information leakage.

The overall data flow of the encryption and authentication process is illustrated in Fig. 1. After initialization, **E** first absorbs the associated data AD. Then, **S** and **E** are updated in parallel: **S** generates the keystream Z to encrypt the plaintext M, while **E** absorbs M. After encryption completes, **E** absorbs Θ, which encodes the lengths of M and AD. In tag generation phase, **E** feeds back into **S** and produces the authentication tag.

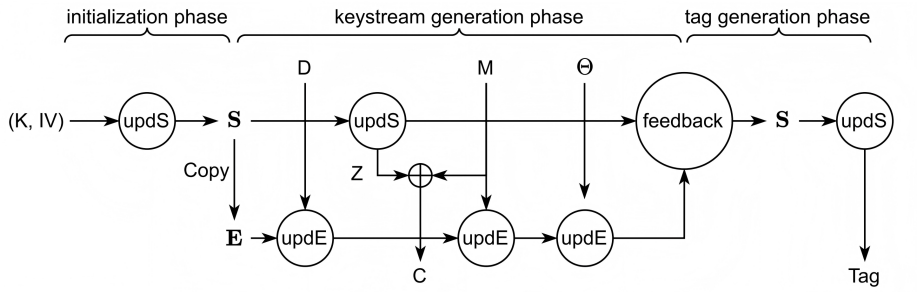

Fig. 1: SCMAC structure of LOL2.0.

LOL2.0 supports multiple ciphers, namely LOL2.0-Mini and LOL2.0-Double. In this work, we focus on LOL2.0-Mini encryption, which uses a 256-bit key and a 128-bit IV and produces a 128-bit keystream block per iteration. This allows us to clearly analyze architectural and security trade-offs without loss of generality. Our proposed hardware architectures are generic and can be readily extended to LOL2.0-Double.

Internal State Structure. The stream-cipher component of LOL2.0 updates the internal state $\mathbf{S} = (L, H, N, S_0, S_1, S_2)$, where all elements are 128-bit registers. The update structure of **S** is illustrated in Fig. 2. These six registers are involved in both the initialization and keystream generation phases. The pair (L, H) forms a linear feedback shift register (LFSR), N is a nonlinear feedback register, and (S_0, S_1, S_2) constitute a finite-state machine. The LFSR update function is defined as $f(H, L) = \lambda(H) \oplus \sigma(L)$, where the linear transformation $\lambda(\cdot)$ and the permutation $\sigma(\cdot)$ are defined in Equations 1 and 2.

$$\begin{aligned}\lambda(x) = (&h_7 \ll 5 \oplus h_6 \gg 11,\ h_6 \ll 5,\ h_5 \ll 5 \oplus h_4 \gg 11,\\ &h_4 \gg 6, h_3 \gg 6,\ h_2 \ll 5,\ h_1 \ll 5 \oplus h_0 \gg 11,\ h_0 \gg 6)\end{aligned} \tag{1}$$

$$\sigma(x_7, \ldots, x_0) = (x_5, x_0, x_3, x_6, x_4, x_7, x_2, x_1). \tag{2}$$

The AEAD component maintains a separate authentication state $\mathbf{E} = (E_0,\ E_1,\ E_2,\ E_3)$, where each element is also a 128-bit register. These registers are initialized as copies of **S** after the initialization phase, after which **E** and **S** evolve

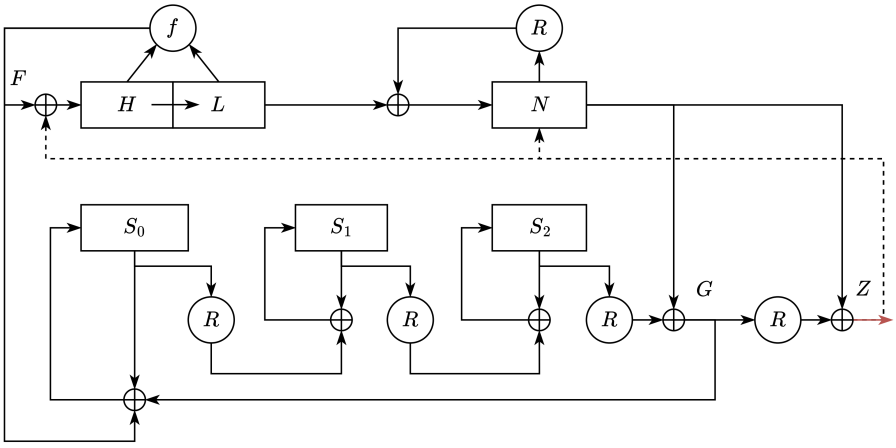

Fig. 2: Update structure of the internal state **S** in LOL2.0. Dashed connections are active only during initialization, while the brown path indicates keystream generation.

independently. The update equation of **E** is given in Equation 3.

$$(E_0 \oplus R(E_3) \oplus AD, E_1 \oplus R(E_0), E_2 \oplus R(E_1), E_3 \oplus R(E_2)) \rightarrow \mathbf{E}' \tag{3}$$

All nonlinear operations in LOL2.0 are confined to a single function R, which corresponds to the AES round function without the AddRoundKey operation, i.e., $R = SubBytes \circ ShiftRows \circ MixColumns$. This property is critical for hardware protection, as it allows side-channel countermeasures to focus exclusively on the implementation of R.

3.2 Masking Scheme

Throughout this work, we target composable first-order security under the glitch-extended probing model, formalized through the PINI semantics. This model captures the effects of glitches in combinational logic and enables secure composition of masked hardware gadgets. We adopt Boolean masking to protect sensitive internal variables by decomposing them into multiple shares. For any w-bit sensitive value X, first-order Boolean masking represents it as $X = X^{(0)} \oplus X^{(1)}$, where the two shares $X^{(0)}$ and $X^{(1)}$ are stored, processed, and propagated independently in separate registers and signal paths.

From an architectural perspective, the entire circuit can be decomposed into two classes of operations: (i) linear operations and (ii) nonlinear operations implemented as masked S-box gadgets. Since linear operations do not mix different shares, they can be implemented independently per share without introducing additional leakage. In contrast, nonlinear operations must be implemented using masking gadgets that satisfy the required composability properties. Under

this construction paradigm, end-to-end composable security is ensured as long as the nonlinear gadgets provide the desired PINI guarantees and the interfaces between modules preserve strict share separation and proper register isolation, thereby preventing unintended share recombination.

In LOL2.0, all nonlinearity is confined to the AES S-box embedded in the R function. We employ an 8-bit masked S-box gadget based on the TSM scheme proposed by Kumar et al. [43] as the fundamental nonlinear building block. The S-box contains two register stages and is constructed by decomposing the algebraic normal form (ANF) of the AES S-box into Boolean monomials. This construction consumes fresh randomness across the evaluation of the monomials and across the internal pipeline stages. Each invocation of the 8-bit gadget consumes a fixed amount of 46 bits of fresh randomness. This randomness breaks deterministic relationships between shares across clock cycles and prevents exploitable leakage caused by glitches and signal recombination in the nonlinear logic.

Moreover, the TSM-based S-box achieves low latency and composable security under the glitch-extended probing model, facilitating efficient integration into the iterative and scheduling-intensive structure of LOL2.0. To support the 128-bit internal state, sixteen such 8-bit masked S-boxes are instantiated in parallel to form a 128-bit masked S-box layer used to protect all nonlinear operations in the design. The R core serves as the fundamental building block of the LOL2.0 datapath and consists of a TSM-protected nonlinear S-box combined with linear operations. Since PINI security guarantees composability of masked gadgets without requiring additional register-based isolation [42], this masking scheme ensures that the entire circuit achieves first-order side-channel security.

3.3 Implementation Structure

This section presents comprehensive hardware implementation designs of LOL2.0, with a particular focus on the architectural trade-offs among throughput, area, and side-channel security. We first develop two unmasked architectures that target different implementation objectives. The unmasked-fast architecture is optimized for high-throughput scenarios, as required by emerging 6G communication systems, while the unmasked-compact architecture aims to minimize hardware cost for resource-constrained deployments.

For side-channel resistance, we further design first-order secure implementations based on the masking scheme described in Section 3.2. Two protected variants are implemented, referred to as the masked-fast and masked-compact versions, following the same design philosophy as their unmasked counterparts. To support heterogeneous security requirements for confidentiality and authentication, all architectures are designed to support both SC and AEAD modes.

Unmasked Implementation. Before introducing masked implementations, we first describe the unmasked architectures, which serve both as performance baselines and as references for evaluating the overhead introduced by side-channel countermeasures.

Unmasked-compact Design. The unmasked-compact architecture prioritizes reduced implementation cost. In LOL2.0, the R function constitutes the dominant contributor to hardware area. Instead of directly instantiating a full 128-bit R function, the compact design decomposes it into smaller 32-bit units to reduce logic resource consumption. Specifically, the architecture instantiates three 32-bit R functions, each composed of four parallel 8-bit AES S-boxes, resulting in a total of 12 S-box instances. Due to this limited nonlinear resource budget, a time-multiplexed update strategy is employed, in which subsets of state registers are sequentially processed by the available R functions according to a predefined and deterministic schedule. This time-multiplexed design significantly reduces area overhead at the cost of increased per-round latency.

In SC mode, only the encryption state $\mathbf{S}$ is updated, and one full round requires eight clock cycles. In contrast, AEAD mode with the authentication state $\mathbf{E}$ increases the total number of required R function invocations and consequently requires twelve clock cycles per round.

Unmasked-fast Design. The fast architecture is designed to establish an upper bound on achievable throughput. In contrast to the compact design, it directly instantiates a complete 128-bit R function, allowing the full transformation to be completed within a single clock cycle. To avoid scheduling constraints caused by resource sharing, each internal state register is updated with its own dedicated R function. This fully parallel design eliminates scheduling dependencies among state updates and achieves both the minimum number of update cycles and the highest throughput. As a result, all state registers can be updated in parallel without structural dependencies between different update paths.

This fully parallel update strategy enables the minimum possible number of update cycles per round and achieves the highest throughput among all unmasked design variants. The performance gain is obtained at the cost of increased hardware resource consumption, as the number of instantiated R functions scales with the number of state registers.

Masked Implementation. In LOL2.0, the initialization and tag generation phases involve feedback paths in which the keystream output Z is injected into several internal state registers. In unprotected implementations, these dependencies can be resolved within a single cycle by freely composing multiple invocations of the nonlinear function. However, such unrestricted composition is no longer feasible once side-channel protection is applied, since masking inevitably introduces additional latency. As a result, these two phases cannot be treated as a simple extension of the keystream generation datapath, and careful scheduling of state updates is required to preserve functional correctness.

Fig. 3 illustrates the overall masked architecture of LOL2.0. The design centers around a TSM-protected nonlinear S-box, which is accessed through input/output selection logic, enabling flexible scheduling of state updates. The highlighted green block corresponds to the pipelined masked R function, while the yellow block represents the linear update network. This modular organiza-

tion allows both time-multiplexed and fully parallel instantiations of the masked architecture, forming the basis for the compact and fast architectures.

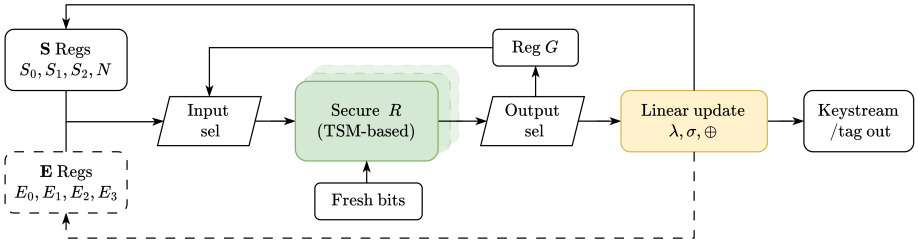

Fig. 3: Overview of the masked LOL2.0 architecture. Dashed block and datapath are enabled only in AEAD mode.

All masked implementations employ a fully pipelined realization of the protected R core. Although masking introduces a fixed latency, the pipeline is designed with an initiation interval of one cycle, allowing a new input to be issued to each R instance every clock cycle and yielding one valid output per cycle after pipeline fill.

Accordingly, as formalized in Algorithm 1, a unified scheduling strategy is employed to dispatch registers in successive cycles, while the corresponding results are written back in order for state updates. The concrete datapath realization—including the number of parallel R cores, pipeline depth, and masking configuration—is implementation-dependent and detailed in the following subsections.

Masked-compact Design. The masked-compact architecture aims to minimize hardware cost and randomness consumption while maintaining acceptable throughput. Accordingly, only a single 128-bit masked R core is instantiated and time-multiplexed across all state updates, maximizing the reuse of nonlinear resources. Because the computation of the intermediate variable G depends on two consecutive invocations of the R function, the register S_2 is always scheduled first. This ordering avoids data hazards and ensures that the masked output of S_2 is available before being consumed in the update of G.

During the initialization and tag generation phases, the encryption state **S** completes one full update every five clock cycles. In AEAD mode, both the encryption state **S** and the authentication state **E** share the same masked R core, which introduces additional resource contention. As a result, one full keystream-generation round is completed in nine clock cycles, producing one 128-bit keystream block per traversal.

By reusing a single masked R core, the compact architecture significantly reduces both area overhead and the number of masked S-boxes per unit time, leading to substantially lower fresh randomness consumption compared to the

Algorithm 1 Unified scheduling algorithm for LOL2.0 masked implementations.

Require: mode $m \in \{\mathrm{SC}, \mathrm{AEAD}\}$; architecture $a \in \{\mathrm{compact}, \mathrm{fast}\}$;
Require: initialization iterations N_{init} ; tag generation iterations (AEAD only) N_{fin}
Require: number of keystream blocks B
Schedule templates: ▷ one traversal completes one full-state update
$\mathcal{L}_{\mathrm{init}} \leftarrow [S_2, S_1, S_0, N, G]$
$\mathcal{L}_{\mathrm{fin}} \leftarrow [S_2, S_1, S_0, N, G]$
if $a = \mathrm{compact}$ **then**
  $\mathcal{L}_{\mathrm{SC}} \leftarrow [S_2, S_1, S_0, N, G]$
  $\mathcal{L}_{\mathrm{AEAD}} \leftarrow [E_3, E_0, E_1, E_2, S_2, S_1, S_0, N, G]$
else ▷ $a = \mathrm{fast}$
  $\mathcal{L}_{\mathrm{SC}} \leftarrow [(S_0, N), (S_1, G), S_2]$
  $\mathcal{L}_{\mathrm{AEAD}} \leftarrow [(E_3, S_1), (E_2, N), (E_1, S_2), (E_0, S_0), G]$
end if
function RUN_PHASE($\mathcal{L}, T, m$) ▷ T traversals of schedule list $\mathcal{L}$
  for $t = 1$ to T **do**
    for each schedule item u in $\mathcal{L}$ **do** ▷ e.g., $u = x$ or $u = (x, y)$
      ISSUE_R(u, m) ▷ issue u to pipelined R, and update state regs
    end for
  end for
end function

Initialization phase:
RUN_PHASE($\mathcal{L}_{\mathrm{init}}$, N_{init}, m)
Keystream generation phase:
for $blk = 1$ to B **do**
  RUN_PHASE($\mathcal{L}_m$, 1, m) ▷ output 128-bit keystream / ciphertext
end for
if $m = \mathrm{AEAD}$ **then**
Tag generation phase (AEAD only):
  RUN_PHASE($\mathcal{L}_{\mathrm{fin}}$, N_{fin}, m) ▷ output 128-bit authentication tag
end if

fast variant. This design is therefore well suited for resource-constrained environments.

Masked-fast Design. The masked-fast architecture targets maximum throughput by exploiting parallel masked computation. Since each 128-bit masked R core has a fixed two-cycle latency, the design constructs a lightweight pipeline and instantiates multiple masked cores.

In AEAD mode, five parallel cores alternately update the registers of **S** and **E** in each pipeline stage, enabling all required state updates to be completed within two clock cycles. In SC mode, three masked R cores are sufficient to update the encryption state **S**. As a result, the masked-fast design produces one 128-bit keystream block every two clock cycles during the keystream generation phase, achieving substantially higher throughput than the compact variant. This increased throughput, however, comes at the cost of a proportional increase in

area and randomness consumption, as the number of masked R evaluations scales with the degree of parallelism.

Together, the masked-compact and masked-fast architectures expose a clear trade-off among area, randomness consumption, and throughput. The compact variant minimizes nonlinear resources and randomness usage through aggressive reuse and careful scheduling, while the fast variant exploits parallelism and pipelining to approach the throughput limits imposed by masked nonlinear latency.

4 Experiments and Evaluations

4.1 Performance Evaluation

We provide a comprehensive hardware performance evaluation of all proposed LOL2.0 implementations. All designs were synthesized using Synopsys Design Compiler with the TSMC 28-nm technology node. To ensure a consistent comparison across different architectural variants, the same compilation options, i.e., `compile_ultra -exact_map -no_autoungroup`, were applied. For each implementation, the clock period was individually optimized to achieve its maximum operating frequency. Throughput is computed as $\left[\frac{\text{clock frequency} \times \text{block size}}{\text{cycles per block}}\right]$, and represents the peak throughput measured during the keystream generation phase, i.e., the steady-state rate at which 128-bit keystream blocks are produced and output. The resulting metrics for all design variants are summarized in Table 1.

Table 1: Post-synthesis hardware implementation results of LOL2.0 in SC and AEAD modes using the TSMC 28-nm technology.

Variant	Mode	Area (kGE)	Rand (bits)	Frequency (GHz)	Latency (cycles)	Throughput (Gbps)
unmasked-compact	SC	23.25	–	4	8	64.00
	AEAD	34.57	–	4	12	42.67
unmasked-fast	SC	90.36	–	1.43	1	182.86
	AEAD	123.20	–	1.43	1	182.86
masked-compact	SC	141.98	736	2.22	5	56.89
	AEAD	163.25	736	2.22	9	31.60
masked-fast	SC	371.44	2208	2.22	2	142.22
	AEAD	626.14	3680	2.22	2	142.22

The results in Table 1 highlight clear architectural trade-offs among the proposed design variants. Most notably, both fast architectures comfortably satisfy the ultra-high-throughput requirements envisioned for 6G systems. Even the masked-fast variant sustains throughput beyond 100 Gbps, demonstrating that

Table 2: Throughput performance (Gbps) comparison in SC mode under different message bytes (Software vs. Hardware). SW results are measured on an Intel i7-11800H CPU (2.30 GHz, Turbo up to 4.6 GHz), implemented in C++ using AVX2 intrinsics in a single-thread manner [14].

Impl.	Variant	Arch.	32B	128B	256B	2048B	16384B
SW [14]	unmasked	–	4.77	26.70	41.68	78.59	90.31
HW	unmasked	compact	9.14	25.60	36.57	58.51	63.26
		fast	26.12	73.14	104.49	167.18	180.74
	masked	compact	7.90	22.31	32.05	51.86	56.21
		fast	8.13	27.75	46.44	113.07	137.78

in our design implementation-level security can be achieved without sacrificing performance.

In contrast, the compact architecture targets resource-constrained deployment scenarios. The masked-compact variant requires substantially less area and fewer fresh random bits than the masked-fast variant, thereby providing a practical trade-off among security, performance, and implementation cost. For the unmasked implementation, the additional area-reduction technique described in Section 3.3 is further employed, at the cost of more serialized execution and higher latency. As a result, the unmasked-compact variant occupies only 23.25 kGE while delivering a throughput of 64 Gbps in SC mode. Although this rate does not reach the 6G-class target, it remains sufficiently high for many practical high-speed applications.

Practical deployment of masked cryptographic accelerators requires considering system-level support for fresh randomness. Nevertheless, the maximum randomness rate required here, 8 Tbps, is achievable with several practical PRNG architectures reported in [6].

Tables 2 and 3 compare the throughput of the proposed hardware implementations with the unprotected software implementation of LOL2.0 under varying message lengths. For large plaintext sizes (e.g., 16,384 bytes), the impact of initialization overhead becomes negligible, and the measured throughput closely reflects steady-state keystream generation performance.

In SC mode, the unmasked hardware implementation achieves a throughput of 180.74 Gbps, corresponding to a 2× speedup over the software implementation. The throughput of the masked implementation remains as high as 137.78 Gbps, still outperforming the software implementation by 52%.

In AEAD mode, the unprotected hardware implementation reaches a throughput of 178.33 Gbps, corresponding to a 3× speedup over the software implementation. With countermeasures, the throughput remains at 133.30 Gbps, achieving a 126% improvement over the software implementation and approaching the throughput requirements expected for future 6G systems. It is worth noting that the software baseline is evaluated on a high-performance CPU platform with a higher operating frequency and more advanced resources than the hardware de-

Table 3: Throughput performance (Gbps) comparison in AEAD mode under different message lengths (Software vs. Hardware). Same SW setup as Table 2.

Impl.	Variant	Arch.	32B	128B	256B	2048B	16384B
SW [14]	unmasked	–	4.48	15.12	24.87	50.36	59.06
HW	unmasked	compact	4.41	13.47	20.48	37.58	41.96
		fast	13.06	43.03	69.66	151.99	178.33
	masked	compact	3.82	11.21	16.55	28.38	31.16
		fast	4.03	14.87	26.93	92.64	133.30

signs. Despite this, the proposed hardware implementations achieve substantially higher throughput, highlighting the efficiency of the proposed architectures.

Table 4: Comparison with state-of-the-art protected hardware implementations of stream ciphers and AEAD designs.

Design	Scheme	Standard-cell library	Area (kGE)	Throughput (Gbps)	Efficiency (Mbps/GE)
SC					
masked-fast	TSM	TSMC 28-nm	371.44	142.22	0.38
Trivium [31]	DOM	TSMC 130-nm	5.15	1.92	0.37
AEAD					
masked-fast	TSM	TSMC 28-nm	626.14	142.22	0.23
Grain-128AEADv2 [30]	DOM	STM 65-nm	22.72	8.33	0.37
Spook v2[a b] [34]	HPC2	TSMC 65-nm	64.60	3.12	0.05
Ascon-128 [17]	DOM	UMC 90-nm	42.75	2.77	0.06
Ascon-128 [18]	UMA	UMC 90-nm	27.18	2.25	0.08
Ascon-fast-TI [20]	TI	UMC 90-nm	123.52	9.02	0.07

[a] This implementation provides second-order security; all others are first-order secure.
[b] This implementation provides only partial phase protection; all others provide full-phase protection.

Side-channel-protected stream cipher implementations targeting 6G-scale performance remain very limited in the open literature. To provide a more comprehensive assessment, Table 4 compares our masked-fast version with representative protected hardware implementations of stream ciphers and AEAD designs. In SC mode, our masked-fast design delivers approximately 74× the throughput of the protected Trivium implementation, while also achieving the best throughput efficiency. In AEAD mode, although its throughput efficiency is slightly lower than that of Grain-128AEADv2, its throughput is about 17× higher, and both its throughput and efficiency remain substantially better than those of the other protected AEAD implementations. Overall, these results indicate that

the proposed design achieves the highest throughput and competitive efficiency among representative protected implementations, making it particularly suitable for ultra-high-throughput secure applications.

4.2 Security Evaluation

In this section, we evaluate the side-channel security of the proposed LOL2.0 hardware implementations. Since the masked-compact architecture in AEAD mode strictly subsumes the SC mode, as the AEAD datapath activates all components required in the SC mode as well as additional authentication logic, we focus our evaluation on three implementations: masked-compact-AEAD, masked-fast-SC, and masked-fast-AEAD. The assessment is carried out at two complementary levels: formal leakage verification under the glitch-extended probing model and practical side-channel evaluation based on power measurements on FPGA platforms. Together, these evaluations provide validation of first-order security for the proposed designs.

Formal Verification. We evaluate the security of the masked implementations using PROLEAD, a hardware leakage detection tool that supports analysis under the glitch-extended probing model. Unlike conventional verification tools that are typically limited to small masking gadgets, PROLEAD is capable of analyzing large-scale masked cryptographic circuits within a reasonable time, making it suitable for full-design verification. In addition, to provide an explicit insecure reference for comparison, we evaluate a masked-compact-AEAD implementation without fresh random bits to illustrate detectable leakage behavior. Due to the large gate count of the proposed masked architectures, performing formal leakage verification over the entire encryption process would require prohibitive memory resources, exceeding the terabyte scale. Consequently, we restrict the PROLEAD analysis to two consecutive rounds of the keystream generation phase, which is sufficient to capture the relevant leakage behavior while keeping the verification tractable.

The leakage assessment results are illustrated in Fig. 4. We adopt a detection threshold of 10^{-6}, indicated by the horizontal dashed line, to mitigate the impact of false positives. For the masked implementations, the assessment was conducted with up to 70 million simulations. Although a small number of leakage estimates temporarily exceed the threshold, they converge back below the threshold as the number of simulations increases, indicating that the observed excursions correspond to statistical fluctuations rather than genuine leakage. By contrast, for the PRNG-off reference design shown in Fig. 4a, the leakage diverges rapidly and grows unbounded after only 1,280 simulations. Thus, the results confirm the absence of first-order leakage in the proposed masked implementations under the glitch-extended probing model.

Practical Evaluation. To evaluate practical side-channel security, we applied the first-order Test Vector Leakage Assessment (TVLA) methodology [15] on an

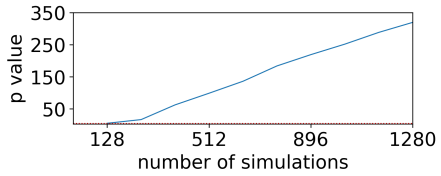

(a) Masked-compact design in AEAD mode, PRNG off

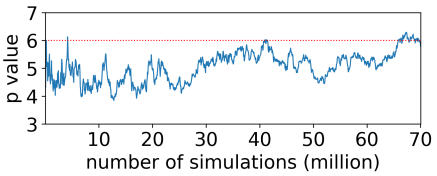

(b) Masked-compact design in AEAD mode, PRNG on

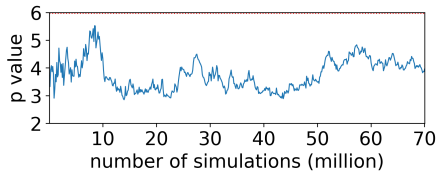

(c) Masked-fast design in SC mode, PRNG on

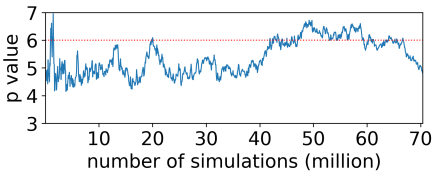

(d) Masked-fast design in AEAD mode, PRNG on

Fig. 4: First-order leakage assessment results of the proposed masked implementations using PROLEAD.

FPGA platform. The experimental setup is based on the CW310 Bergen Board equipped with a Xilinx Kintex-7 FPGA. All design variants are implemented on the same FPGA device and evaluated under identical measurement conditions. Power measurements are obtained via the on-board SMA interface using an amplified shunt-based measurement setup with a PicoScope 3418E oscilloscope. An overview of the measurement setup is illustrated in Fig. 5.

Fresh randomness is generated using a Trivium-based PRNG. Since the evaluated designs consume a large amount of randomness during a single execution, generating random bits on the fly would introduce additional switching activity, potentially affecting the noise characteristics and thus biasing the TVLA results. To avoid this effect, we adopt a pre-generated randomness supply strategy to better isolate the protected core and assess whether the core itself exhibits leakage. Specifically, a dedicated BRAM region on the FPGA is used as a randomness reservoir. The random bits required for one execution are pre-generated and stored in the BRAM, while the PRNG is disabled during cryptographic processing. In a practical deployment, enabling the PRNG would introduce additional switching noise, which generally reduces the signal-to-noise ratio available to an attacker, thereby making side-channel attacks more difficult.

For all evaluated designs, the measurement configuration is chosen to ensure an average of approximately 50 samples per clock cycle, providing sufficient temporal resolution for TVLA. This sampling density provides sufficient temporal resolution to capture intra-cycle switching activity, thereby reducing the risk of undersampling-induced artifacts in the TVLA statistic. As the raw power traces

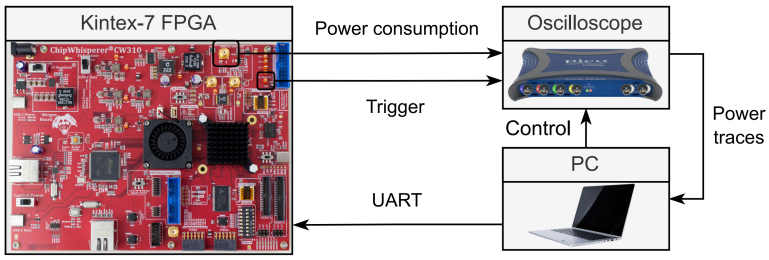

Fig. 5: Experimental setup for practical TVLA-based side-channel evaluation.

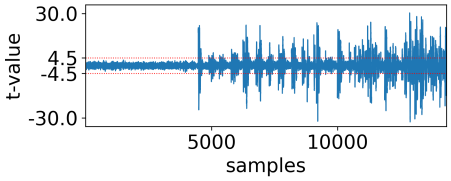

(a) Masked-compact design in AEAD mode, PRNG off (60k traces)

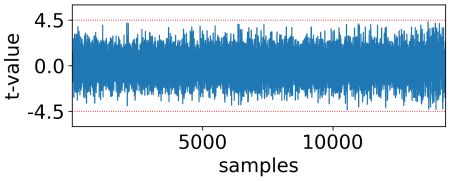

(b) Masked-compact design in AEAD mode, PRNG on (100M traces)

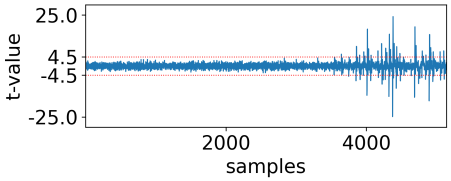

(c) Masked-fast design in SC mode, PRNG off (80k traces)

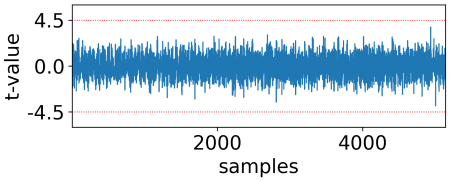

(d) Masked-fast design in SC mode, PRNG on (100M traces)

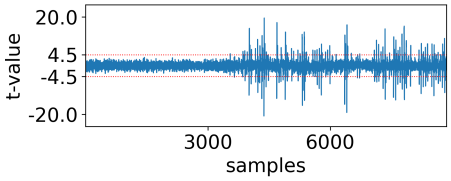

(e) Masked-fast design in AEAD mode, PRNG off (60k traces)

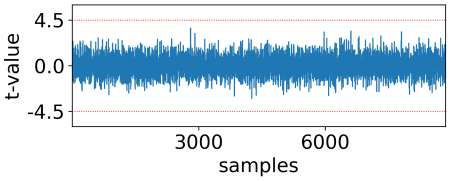

(f) Masked-fast design in AEAD mode, PRNG on (100M traces)

Fig. 6: First-order TVLA results for the masked LOL2.0 implementations under different architectural configurations and operating modes.

exhibit a slowly varying envelope, a high-pass filter is applied prior to TVLA to suppress low-frequency components.

The first-order TVLA results are presented in Fig. 6. For all PRNG-off configurations, significant leakage is observed within fewer than 80k traces. In contrast, for the masked implementations with randomness enabled, up to 100 million traces are analyzed, and the test statistic never exceeded the ± 4.5 threshold. These results provide strong evidence for the first-order side-channel security of the proposed masked implementations.

To evaluate the resistance of the proposed designs against higher-order attacks, we further performed second-order univariate TVLA tests. The results are shown in Fig. 7. With up to 100 million traces, the masked-compact variant exhibits slight second-order leakage, while no leakage is observed for the masked-fast variant. This suggests that, although the protected architectures target first-order security, they still offer a certain level of resistance to higher-order leakage.

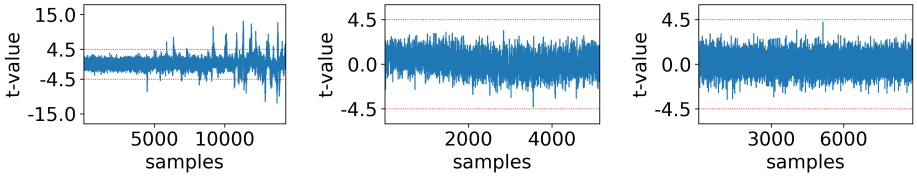

(a) Masked-compact design in AEAD mode (b) Masked-fast design in SC mode (c) Masked-fast design in AEAD mode

Fig. 7: Second-order univariate TVLA results for the masked LOL2.0 implementations with PRNG enabled over 100 million traces.

5 Conclusions

Achieving ultra-high throughput while maintaining rigorous resistance to side-channel attacks remains a fundamental challenge for stream cipher hardware in emerging 6G communication systems. Although recent years have witnessed significant advances in high-performance stream cipher design, practical hardware implementations that simultaneously provide verifiable side-channel protection and sustain 6G-class throughput remain largely unexplored.

In this work, we present the first stream cipher hardware implementation that meets the throughput demands of 6G systems under first-order full-phase side-channel protection. Focusing on LOL2.0, we leverage TSM to achieve first-order security under the glitch-extended probing model with PINI composability, while preserving the low latency required for high-throughput operation. Two masked architectures are developed to address different deployment scenarios:

a compact variant that prioritizes area and randomness efficiency, and a fast variant that exploits parallelism and pipelining to maximize throughput.

All designs were synthesized using the TSMC 28-nm technology node. The masked-fast architecture sustains a throughput of 142 Gbps in both SC and AEAD modes, meeting the throughput requirements for 6G systems even under full first-order side-channel protection. The unprotected fast implementation reaches a peak throughput of 183 Gbps, delivering approximately a twofold performance improvement over the software implementation under the same plaintext configuration. The masked-compact architecture significantly reduces implementation cost, achieving substantially lower area and randomness consumption than the fast variant. Despite these reductions, it still delivers 56.89 Gbps in SC mode and 31.60 Gbps in AEAD mode, providing a well-balanced trade-off between performance, hardware cost, and side-channel security.

Importantly, both proposed architectures preserve first-order security, and this security claim is substantiated through formal and practical evaluations. Formal leakage analysis using PROLEAD, with up to 70 million simulations, confirms the absence of first-order leakage under the glitch-extended probing model. FPGA-based measurements using TVLA, covering up to 100 million power traces, likewise reveal no statistically significant first-order leakage for the masked implementations, while clear leakage is observed when masking is disabled. Although the first-order masking is inherently limited against higher-order adversaries, the second-order TVLA results show that our protected designs exhibit a certain degree of resistance to higher-order leakage. Taken together, these results validate the security of the proposed designs throughout the entire encryption and authentication process.

By enabling ultra-high-throughput encryption and authentication under strict side-channel protection, the proposed designs directly support emerging 6G use cases such as massive connectivity, high-data-rate services, and HRLLC. As such, this work provides a concrete hardware foundation for integrating secure and efficient cryptographic primitives into next-generation 6G communication systems.

Future work includes extending the proposed designs to support higher-order protection and reducing randomness consumption for resource-constrained platforms. In addition, the proposed design methodology can be generalized to high-throughput, side-channel-resistant hardware implementations of other stream ciphers and AEAD constructions targeting future 6G systems.

Acknowledgments. This work was supported in part by the National Key R&D Program of China (No. 2024YFB3108103) and in part by the National Natural Science Foundation of China (Grant No. 62504134).

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Adomnicai, A., Fournier, J.J.A., Masson, L.: Masking the lightweight authenticated ciphers acorn and ascon in software (2018), https://eprint.iacr.org/2018/708

2. Akyildiz, I.F., Kak, A., Nie, S.: 6g and beyond: The future of wireless communications systems. IEEE Access **8**, 133995–134030 (2020). https://doi.org/10.1109/ACCESS.2020.3010896
3. Arribas, V., Zhang, Z., Nikova, S.: Llti: Low-latency threshold implementations. IEEE Transactions on Information Forensics and Security **16**, 5108–5123 (2021). https://doi.org/10.1109/TIFS.2021.3123527
4. Barthe, G., Belaïd, S., Dupressoir, F., Fouque, P.A., Grégoire, B., Strub, P.Y.: Verified proofs of higher-order masking. In: Oswald, E., Fischlin, M. (eds.) Advances in Cryptology – EUROCRYPT 2015. pp. 457–485. Springer (2015). https://doi.org/10.1007/978-3-662-46800-5_18
5. Barthe, G., Belaïd, S., Dupressoir, F., Fouque, P.A., Grégoire, B., Strub, P.Y., Zucchini, R.: Strong non-interference and type-directed higher-order masking. In: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security. pp. 116–129. CCS '16, Association for Computing Machinery (Oct 2016). https://doi.org/10.1145/2976749.2978427
6. Cassiers, G., Masure, L., Momin, C., Moos, T., Moradi, A., Standaert, F.X.: Randomness generation for secure hardware masking – unrolled trivium to the rescue. IACR Communications in Cryptology **1**(2) (Jul 2024). https://doi.org/10.62056/akdkp2fgx
7. Cassiers, G., Standaert, F.X.: Trivially and efficiently composing masked gadgets with probe isolating non-interference. IEEE Transactions on Information Forensics and Security **15**, 2542–2555 (2020). https://doi.org/10.1109/TIFS.2020.2971153
8. Coron, J.S., Prouff, E., Rivain, M., Roche, T.: Higher-order side channel security and mask refreshing. In: Moriai, S. (ed.) Fast Software Encryption. pp. 410–424. Springer (2014). https://doi.org/10.1007/978-3-662-43933-3_21
9. Dang, S., Amin, O., Shihada, B., Alouini, M.S.: What should 6g be? Nature Electronics **3**(1), 20–29 (Jan 2020). https://doi.org/10.1038/s41928-019-0355-6
10. Dong, X., Dong, B., Wang, X.: Quantum attacks on some feistel block ciphers. Designs, Codes and Cryptography **88**(6), 1179–1203 (Jun 2020). https://doi.org/10.1007/s10623-020-00741-y
11. Ekdahl, P., Maximov, A., Johansson, T., Yang, J.: Snow-vi: An extreme performance variant of snow-v for lower grade cpus. In: Proceedings of the 14th ACM Conference on Security and Privacy in Wireless and Mobile Networks. pp. 261–272. WiSec '21, Association for Computing Machinery (Jun 2021). https://doi.org/10.1145/3448300.3467829
12. Faust, S., Grosso, V., Pozo, S.M.D., Paglialonga, C., Standaert, F.X.: Composable masking schemes in the presence of physical defaults & the robust probing model. IACR Transactions on Cryptographic Hardware and Embedded Systems pp. 89–120 (Aug 2018). https://doi.org/10.13154/tches.v2018.i3.89-120
13. Feng, D., Jiao, L., Hao, Y., Zheng, Q., Wu, W., Qi, W., Zhang, L., Zhang, L., Sun, S., Tian, T.: Lol: A highly flexible framework for designing stream ciphers. Science China Information Sciences **67**(9), 192101 (Aug 2024). https://doi.org/10.1007/s11432-023-3901-0
14. Feng, D., Jiao, L., Hao, Y., Zheng, Q., Wu, W., Qi, W., Zhang, L., Zhang, L., Sun, S., Tian, T.: Scmac and lol2.0: An aead design framework and a new version of lol stream cipher design framework (2025), https://eprint.iacr.org/2025/925
15. Gilbert Goodwill, B.J., Jaffe, J., Rohatgi, P., et al.: A testing methodology for side-channel resistance validation. In: NIST Non-Invasive Attack Testing Workshop. vol. 7, pp. 115–136 (2011)

16. Goubin, L., Patarin, J.: Des and differential power analysis (the "duplication" method). In: Proceedings of the First International Workshop on Cryptographic Hardware and Embedded Systems. pp. 158–172. CHES '99, Springer-Verlag (Aug 1999)
17. Gross, H., Iusupov, R., Bloem, R.: Generic low-latency masking in hardware. IACR Transactions on Cryptographic Hardware and Embedded Systems pp. 1–21 (May 2018). https://doi.org/10.13154/tches.v2018.i2.1-21
18. Gross, H., Mangard, S.: Reconciling d+1 masking in hardware and software. In: Fischer, W., Homma, N. (eds.) Cryptographic Hardware and Embedded Systems – CHES 2017. pp. 115–136. Springer International Publishing (2017). https://doi.org/10.1007/978-3-319-66787-4_6
19. Gross, H., Mangard, S., Korak, T.: Domain-oriented masking: Compact masked hardware implementations with arbitrary protection order. In: Proceedings of the 2016 ACM Workshop on Theory of Implementation Security. p. 3. TIS '16, Association for Computing Machinery (Oct 2016). https://doi.org/10.1145/2996366.2996426
20. Groß, H., Wenger, E., Dobraunig, C., Ehrenhöfer, C.: Suit up! – made-to-measure hardware implementations of ascon. In: 2015 Euromicro Conference on Digital System Design. pp. 645–652 (Aug 2015). https://doi.org/10.1109/DSD.2015.14
21. Gross, H., Wenger, E., Dobraunig, C., Ehrenhöfer, C.: Ascon hardware implementations and side-channel evaluation. Microprocessors and Microsystems **52**, 470–479 (Jul 2017). https://doi.org/10.1016/j.micpro.2016.10.006
22. Hosoyamada, A., Inoue, A., Ito, R., Iwata, T., Minematsu, K., Sibleyras, F., Todo, Y.: Cryptanalysis of rocca and feasibility of its security claim. IACR Transactions on Symmetric Cryptology pp. 123–151 (2022)
23. International Telecommunication Union - Radiocommunication Sector (ITU-R): Minimum requirements related to technical performance for imt-2020 radio interface(s). Report itu-r m.2410-0, International Telecommunication Union (2017), https://www.itu.int/pub/R-REP-M.2410-2017
24. International Telecommunication Union - Radiocommunication Sector (ITU-R): Framework and overall objectives of the future development of imt for 2030 and beyond. Recommendation M.2160-0, ITU-R (Nov 2023), https://www.itu.int/rec/R-REC-M.2160-0-202311-I/en
25. Ishai, Y., Sahai, A., Wagner, D.: Private circuits: Securing hardware against probing attacks. In: Boneh, D. (ed.) Advances in Cryptology - CRYPTO 2003. pp. 463–481. Springer (2003). https://doi.org/10.1007/978-3-540-45146-4_27
26. Kaplan, M., Leurent, G., Leverrier, A., Naya-Plasencia, M.: Breaking symmetric cryptosystems using quantum period finding. In: Robshaw, M., Katz, J. (eds.) Advances in Cryptology – CRYPTO 2016. pp. 207–237. Springer (2016). https://doi.org/10.1007/978-3-662-53008-5_8
27. Knichel, D., Sasdrich, P., Moradi, A.: Generic hardware private circuits: Towards automated generation of composable secure gadgets. IACR Transactions on Cryptographic Hardware and Embedded Systems pp. 323–344 (2022). https://doi.org/10.46586/tches.v2022.i1.323-344
28. Kumar, S., Dasu, V.A., Baksi, A., Sarkar, S., Jap, D., Breier, J., Bhasin, S.: Side channel attack on stream ciphers: A three-step approach to state/key recovery. IACR Transactions on Cryptographic Hardware and Embedded Systems pp. 166–191 (Feb 2022). https://doi.org/10.46586/tches.v2022.i2.166-191
29. Latva-aho, M., Leppänen, K., University of Oulu, G.F.: Key drivers and research challenges for 6g ubiquitous wireless intelligence (Sep 2019), https://oulurepo.oulu.fi/handle/10024/36430

30. Li, B., Zhang, H., Lin, D.: Efficient (masked) hardware implementation of grain-128aeadv2. Security and Communication Networks **2023**(1), 8044164 (2023). https://doi.org/10.1155/2023/8044164
31. Li, B., Zhang, H., Lin, D.: Higher-order masking scheme for trivium hardware implementation. In: Deng, Y., Yung, M. (eds.) Information Security and Cryptology. pp. 337–356. Springer Nature Switzerland (2023). https://doi.org/10.1007/978-3-031-26553-2_18
32. Liu, F., Isobe, T., Meier, W., Sakamoto, K.: Weak keys in reduced aegis and tiaoxin. IACR Transactions on Symmetric Cryptology pp. 104–139 (Jun 2021). https://doi.org/10.46586/tosc.v2021.i2.104-139
33. Mentens, N., Miskovsky, V., Novotny, M., Vliegen, J.: High-speed side-channel-protected encryption and authentication in hardware (2018), https://eprint.iacr.org/2018/1088
34. Momin, C., Cassiers, G., Standaert, F.X.: Unprotected and masked hardware implementations of spook v2 (2022), https://eprint.iacr.org/2022/254
35. Moradi, A., Poschmann, A., Ling, S., Paar, C., Wang, H.: Pushing the limits: A very compact and a threshold implementation of aes. In: Paterson, K.G. (ed.) Advances in Cryptology – EUROCRYPT 2011. pp. 69–88. Springer (2011). https://doi.org/10.1007/978-3-642-20465-4_6
36. Prasad, S.H., Mendel, F., Schläffer, M., Nagpal, R.: Efficient low-latency masking of ascon without fresh randomness (2023), https://eprint.iacr.org/2023/1914
37. Sakamoto, K., Liu, F., Nakano, Y., Kiyomoto, S., Isobe, T.: Rocca: An efficient aes-based encryption scheme for beyond 5g. IACR Transactions on Symmetric Cryptology pp. 1–30 (Jun 2021). https://doi.org/10.46586/tosc.v2021.i2.1-30
38. Sasdrich, P., Bilgin, B., Hutter, M., Marson, M.E.: Low-latency hardware masking with application to aes. IACR Transactions on Cryptographic Hardware and Embedded Systems pp. 300–326 (Mar 2020). https://doi.org/10.13154/tches.v2020.i2.300-326
39. Saurabh, H., Golder, A., Titti, S.S., Kundu, S., Li, C., Karmakar, A., Das, D.: Snow-sca: Ml-assisted side-channel attack on snow-v. In: 2024 IEEE International Symposium on Hardware Oriented Security and Trust (HOST). pp. 139–149 (May 2024). https://doi.org/10.1109/HOST55342.2024.10545384
40. Shi, Z., Jin, C., Zhang, J., Cui, T., Ding, L., Jin, Y.: A correlation attack on full snow-v and snow-vi. In: Dunkelman, O., Dziembowski, S. (eds.) Advances in Cryptology – EUROCRYPT 2022. pp. 34–56. Springer International Publishing (2022). https://doi.org/10.1007/978-3-031-07082-2_2
41. Simon, D.R.: On the power of quantum computation. SIAM Journal on Computing **26**(5), 1474–1483 (Oct 1997). https://doi.org/10.1137/S0097539796298637
42. V, D.K.S., Dhooghe, S., Balasch, J., Gierlichs, B., Verbauwhede, I.: Time sharing - a novel approach to low-latency masking. IACR Transactions on Cryptographic Hardware and Embedded Systems **2024**(3), 249–272 (Jul 2024). https://doi.org/10.46586/tches.v2024.i3.249-272
43. V, D.K.S., Dhooghe, S., Balasch, J., Gierlichs, B., Verbauwhede, I.: Higher-order time sharing masking. IACR Transactions on Cryptographic Hardware and Embedded Systems **2025**(2), 235–267 (Mar 2025). https://doi.org/10.46586/tches.v2025.i2.235-267
44. Wu, H., Preneel, B.: Aegis: A fast authenticated encryption algorithm. In: Lange, T., Lauter, K., Lisoněk, P. (eds.) Selected Areas in Cryptography – SAC 2013. pp. 185–201. Springer (2014). https://doi.org/10.1007/978-3-662-43414-7_10