

A Switched-Capacitor Converter with Switch-Mode Time-Domain Locking and HCI-Based Self-Destruction

Heejoon Yoon Heejoon Yoon

Northeastern University

Utkarsh Kumar Utkarsh Kumar

kumar.ut@northeastern.edu

Northeastern University

Aatmesh Shrivastava Aatmesh Shrivastava

Northeastern University

Research Article

Keywords: Switched-Capacitor Converter, Switch-Mode Time-Domain Locking, Hardware Security, Physically Unclonable Function, Hot-Carrier Injection

Posted Date: July 8th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10230711/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

A Switched-Capacitor Converter with Switch-Mode Time-Domain Locking and HCI-Based Self-Destruction

Heejoon Yoon ¹, Utkarsh Kumar ^{1*}, Aatmesh Shrivastava ¹

¹Dept. of Electrical & Computer Engineering, Northeastern University, Boston, MA, USA.

*Corresponding author(s). E-mail(s): kumar.ut@northeastern.edu;
Contributing authors: yoons.he@northeastern.edu; aatmesh@northeastern.edu;

Abstract

Switched-capacitor converters (SCCs) are widely used in power-management ICs, conventionally driven by periodic clocks. This paper presents a key-driven SCC in which a 128-bit non-periodic digital key from a physically unclonable function (PUF) replaces the clock, enabling hardware authentication through temporal switching behavior. Stable output voltage and prompt startup are achieved only when the supplied key matches the on-chip reference. Statistical analysis of 100 key sets yields a Shannon entropy of 0.993 bit/position and a Hamming-weight distribution consistent with an ideal Bernoulli(0.5) source, giving an effective key space of 89 bits after run-length correction. Long-term robustness is modeled by treating the hot-carrier-injection (HCI)-induced threshold-voltage shift (ΔV_{th}) during shoot-through events as a positive-feedback process. With 64 incorrect bits, power conversion efficiency drops from 0.86 to 0.29, output voltage from 2.18 V to 1.96 V, and startup time from 3.8 to 7.7 μ s. After a 100 mV HCI-induced ΔV_{th} , output voltage falls to 1.84 V and startup time reaches 23.7 μ s. A startup-time detector disables downstream circuitry under unauthorized access. The proposed design achieves lower power and area overhead than prior work, with negligible degradation under authorized operation.

Keywords: Switched-Capacitor Converter, Switch-Mode Time-Domain Locking, Hardware Security, Physically Unclonable Function, Hot-Carrier Injection

1 Introduction

Protection of digital hardware has been extensively studied, ranging from established cryptographic techniques [1] to design-level approaches such as logic locking [2, 3], which remains an active area of research. Security techniques for analog and mixed-signal circuits, by contrast, remain comparatively limited [4, 5]. This asymmetry reflects the historical trajectory of semiconductor

research, which has been driven predominantly by digital computing: substantial progress has been made in digital IP protection, while analog and mixed-signal hardware has received comparatively little attention.

The rapid growth of artificial intelligence (AI) and other data-intensive applications is expanding the role of analog and mixed-signal circuits in modern system-on-chip (SoC) designs [6]. Analog blocks are no longer confined to supporting digital circuitry but increasingly perform critical computational and signal-processing functions in their own right. At the same time, the widespread

This work was supported in part by the National Science Foundation (NSF) under Grant Nos. CCSS-2225368 and CCSS-2144703.

reliance on third-party foundries has exposed these designs to a range of supply-chain threats, including intellectual-property piracy [7], unauthorized overproduction [8], counterfeiting [9], and reverse engineering [10]. Protecting analog hardware has therefore become an increasingly pressing design challenge.

1.1 Conventional Analog Locking Techniques

Several techniques have been proposed to improve the security of analog and mixed-signal circuits. Bias-based obfuscation [4] replaces biasing transistors with key-controlled width-selectable transistor banks. However, because analog circuits operate over a continuous range of acceptable bias conditions rather than a single unique operating point, multiple key combinations can produce similar circuit behavior, resulting in repeated correct keys (CKs) across chips. Consequently, the reported incorrect-key (IK) variation can be as low as 3% [11], while the effective key size is ultimately limited by the analog noise floor regardless of the nominal key length.

Calibration-based locking techniques [5, 12] utilize process-calibration bits as the authentication key, achieving negligible hardware overhead in the lock-less case [12] and up to 50% IK variation in logic-locked implementations [5]. Nevertheless, since fabricated dies converge toward similar calibration targets, the resulting keys exhibit limited inter-chip entropy and frequently repeat across devices.

Design-based locking techniques exploit layout-dependent effects to obfuscate circuit functionality. Although these methods cause complete loss of amplifier functionality for incorrect keys in an OTA [13], they provide only a single valid key for a given design and incur area overheads of up to 175%. ANN-based bias generation [14] generates the correct bias voltages only for a valid analog key, improving resistance against model-approximation attacks. However, floating-gate non-idealities limit the effective key space to approximately 2^{14} combinations for a three-input implementation, while the reported IK variation remains around 20%.

Authors in [11] recently proposed the switch-mode time-domain locking (SMDL) technique, which overcomes many of these limitations by




Fig. 1 Conceptual comparison of hardware locking strategies. (a) Conventional digital logic locking (b) Proposed PSU-level locking.

embedding authentication into the temporal operation of switching analog circuits. Instead of modifying static bias conditions or calibration states, SMDL employs a long PUF-derived key to control the switching sequence, thereby providing high key entropy with minimal hardware overhead. Experimental results demonstrate near-100% IK variation in both a folded-cascode amplifier (FCA) and a switched-capacitor bandgap reference, establishing SMDL as a scalable foundation for securing switching-based analog circuits.

1.2 Proposed SMDL Framework for SCCs

Among the analog blocks in a typical SoC, the on-chip SCC sits at the top of the power delivery hierarchy. Every downstream digital and analog block draws its operating voltage from the power supply unit (PSU), so a disabled PSU cascades into a complete system shutdown. SCCs are particularly popular for implementing PSUs in ULP IoT and edge-AI SoCs, owing to their fully integrated nature, high efficiency across large




Fig. 2 (a) Charging phase and (b) Discharging phase of a single-phase SCC.

load variation, and compatibility with low supply voltages [15–17].

Fig. 1 illustrates the difference between conventional digital logic locking and the proposed SMDL-SCC locking method. In Fig. 1(a), the comparator and the protected logic are powered by the on-chip converter, so the locking mechanism remains fully energized and accessible to repeated probing whenever the chip is powered. This makes the system vulnerable to brute-force attacks, since IKs result only in transient failures and allow the adversary to promptly initiate an unlimited number of IK attempts. Embedding key verification directly into the SCC, rather than delegating it to a separate digital comparator as in Fig. 1(b), offers a protection leverage that locking a single functional block cannot provide: a single secured SCC implicitly secures every block it supplies power to. Moreover, the SCC is a particularly suitable target for SMDL, since its switching network directly handles the full charge-transfer current.

In this work, we present a key-driven SCC implementation that also includes an HCI-induced self-destruction mechanism, extending SMDL [11] to an SCC target while addressing the iterative

key-guessing attack surface inherent to recoverable analog locks. The on-chip clocks are replaced by a 128-bit non-periodic key generated by an embedded current-mode PUF [18] and conditioned by a digital chain consisting of a serial-in parallel-out (SIPO) shift register, a level shifter, and a maximum-run-length controller (MRLC). With the CK, the converter operates within a negligible margin of the conventional clock-driven baseline. Under an IK attack, however, persistent shoot-through events accumulate HCI stress on the switching transistors and progressively shift their V_{th} , until the converter can no longer reach its target output even when the CK is later reapplied. The locking mechanism is therefore physically irreversible, as shown in Fig. 1(b): a sustained IK attack permanently disables the chip rather than merely interrupting it, eliminating the repeated-attack advantage that adversaries exploit against recoverable analog locks.

The remainder of this paper is organized as follows. Section 2 describes the cross-coupled SCC topology, the SMDL operating principle, and the definition of QSS. Section 3 presents the FCA-based PUF and the digital conditioning chain. Section 4 develops the HCI model used in this work and the corresponding ΔV_{th} calculation. Section 5 provides a statistical evaluation of the 128-bit PUF keys against the i.i.d. Bernoulli reference and quantifies the impact of MRLC on bit-balance, run-length, autocorrelation, and per-bit entropy. Section 6 reports the SCC performance under CK and IK operation as a function of bit-error count n_{err} and HCI-induced ΔV_{th} , along with a startup-time violation detection scheme. Section 7 concludes the paper.

2 Cross-Coupled SCC

2.1 Conventional SCC Operation

A voltage doubler [19–21] is a switched-capacitor converter (SCC) that ideally generates an output voltage equal to twice its input by alternately charging and discharging a flying capacitor under periodic steady-state (PSS) operation.

As shown in Fig. 2, a conventional single-phase SCC consists of four power switches: two NMOS and two PMOS devices. Its operation comprises two alternating phases. During the charging phase (Fig. 2(a)), the NMOS switches are turned on to




Fig. 3 Schematic of a cross-coupled SCC.

charge the flying capacitor from the input supply while the load is supported by the output capacitor. During the discharging phase (Fig. 2(b)), the PMOS switches connect the charged flying capacitor to the output, transferring its stored charge to the load. Switches M1 and M4 operate in the high-voltage domain (HVD), whereas M2 and M3 operate in the low-voltage domain (LVD). Because the gate voltages of the HVD switches exceed the input supply voltage, dedicated level shifters are required to generate their gate-drive signals, resulting in additional static power consumption. To eliminate this level-shifter overhead, a cross-coupled topology [22] (Fig. 3) is employed, in which the HVD switches are driven through internal cross-coupling between two interleaved phases. During the first phase, the LVD switches are controlled by CLK, with negligible overlap between their conduction intervals. The second phase is driven by the complementary clock, CLKb, so that one phase charges the flying capacitor while the other discharges it. The gate of each HVD switch is driven by the top-plate voltage of the flying capacitor in the opposite phase, thereby eliminating the need for a level shifter. This interleaved operation also reduces the output-voltage ripple and improves the power conversion efficiency (PCE).

For ultra-low-power (ULP) switched-capacitor converters operating at low input voltages and light loads, the dominant efficiency loss arises from charging and discharging the bottom-plate parasitic capacitance of the flying capacitor (bottom-plate loss). Several techniques have been proposed to reduce this loss, including bottom-plate charge sharing during the dead time between switching phases [23] and process-level optimization to minimize the parasitic capacitance [24].

The PCE of an SCC is determined by several loss mechanisms, including bottom-plate parasitic loss, switching loss, gate-driving loss, conduction loss, and charge-redistribution loss [16]. Switching and gate-driving losses increase with operating frequency, whereas charge-redistribution loss decreases with frequency; consequently, the optimum operating frequency is obtained when these competing losses are balanced, corresponding to the maximum-power-point operating frequency of the converter.

2.2 SMDL SCC

Conventional SCCs are driven by periodic clock pulses that guarantee PSS and uniform switching. In this work, we instead drive the SCC with a non-periodic 128-bit digital key generated from a PUF. The intent is to bring the converter into the hardware-security domain: a stable supply and prompt startup are obtained only when INPUT is identical to CK. To ensure that the cross-coupled SCC operates correctly only when the two input keys are identical, the LVD switches of the two stages are driven as illustrated in Fig. 4(a). CLK and CLKb drive the LVD NMOS switch of each phase, while INPUT and its complement (INPUTb) drive the LVD PMOS switch of each phase. Under an IK trial, two additional phases arise that do not exist in the matched-key case; these are illustrated in Fig. 4(b). These additional phases are the main reasons why the SCC cannot maintain reliable performance under IK operation.

- (CK, INPUT) = (HIGH, LOW): a direct conduction path is created from V_{IN} to ground through the two LVD switches of one phase. The shoot-through current is set by V_{IN} and R_{on} of the two LVD switches, and dissipates energy without delivering charge to the load.
- (CK, INPUT) = (LOW, HIGH): all four switches in the affected phase are off and the SCC is momentarily inactive. The output node continues to be discharged by the load current, producing a deeper droop than a normal discharge phase.

The primary concern when driving an SCC with a non-periodic key is its low-frequency content, which is governed by the key's maximum run length (the longest sequence of identical consecutive bits). Even when INPUT is identical to




Fig. 4 (a) Schematic of the proposed cross-coupled SCC with digital input keys. (b) Four stages of operation of the digital key-driven SCC.

CK, a key with a large run length causes long intervals during which the SCC does not switch, producing a large transient droop on the output. This behavior is fundamentally different from the clock-driven case, where switching remains uniform regardless of the underlying data.

To optimize the power conversion efficiency of a key-driven SCC, its effective switching frequency must therefore be analyzed explicitly. In this work, the 128-bit PUF key is constrained to a maximum run length of 2, so that every run in the key has a length of either 1 or 2. Let N_1 and N_2 denote the number of runs of length 1 and length 2, respectively.

$$\text{Total number of bits: } 128 = N_1 + 2N_2 \quad (1)$$

$$\text{Total number of transitions: } N_{tr} = N_1 + N_2 \quad (2)$$

A uniformly random sequence with maximum run length 2 has, on average, an equal number of length-1 and length-2 runs, i.e.,

$$N_1 \approx N_2 \approx 128/3 \approx 42.7 \quad (3)$$




Fig. 5 (a) Output voltage of a clock-driven SCC and digital key-driven SCC. (b) Enlarged view of PSS and QSS.

so the expected number of transitions per 128-bit cycle is

$$N_{tr} = N_1 + N_2 \approx 2 \times 128/3 \approx 85.3 \quad (4)$$

The equivalent switching frequency is defined as

$$f_{eq} = N_{tr}/(128 \times T_{CLK}) \quad (5)$$

Since each complete switching cycle requires two edges, this reduces to

$$f_{eq} = (N_{tr}/128) \times f_{CLK} \approx \frac{2}{3} f_{CLK} \approx 0.67 f_{CLK} \quad (6)$$

which lies between the two natural bounds $0.5f_{CLK}$ (every run is length 2) and f_{CLK} (every run is length 1). The SCC should therefore be optimized at f_{eq} rather than at f_{CLK} .

This work targets an SCC operating from a 1.2-V input supply, where the average output voltage depends directly on the switching frequency. Fig. 5(a) and (b) compare the output voltage under conventional clock-driven operation and CK-driven operation, using a CK with a maximum run length of 2. As shown in Fig. 5(a), both




Fig. 6 (a) Schematic of the PUF generation system. (b) Timing diagram. (c) Waveform of V_{OUT} and $V_{SAMPLING}$.

implementations reach steady operation within approximately 5 μs . The enlarged waveform in Fig. 5(b) highlights the resulting difference in output ripple: the clock-driven implementation exhibits a periodic ripple of 51 mV peak-to-peak around 2.24 V, whereas the CK-driven implementation produces a non-periodic ripple of 114 mV peak-to-peak around 2.19 V – approximately $2.2\times$ larger. Because the ripple amplitude and duration vary with the local run length of the key, the converter settles into a quasi-steady-state (QSS) operating condition rather than a fixed periodic one.

3 PUF System for Digital Key Generation

The PUF generation system employed in this work uses transistor mismatch as its entropy source to generate a random digital key. A differential-input, single-ended FCA is used: the random mismatch between the two PMOS input transistors of the differential pair is amplified and converted into a polarity decision at V_{out} . The schematic of the

amplifier is shown in Fig. 6(a). Unlike prior work on FCA-based PUF generation [18], the proposed structure employs a low-voltage cascode biasing scheme such that V_{out} settles near 0.7 V, close to half of the 1.2 V supply rail. Centering V_{out} near mid-rail simplifies the design of the downstream inverter-chain comparator and shortens the settling time in the next input pair’s compensation phase.

Fig. 6(b) shows the time-domain operation of the proposed PUF generation system. Generation of the i th bit involves one compensation phase followed by the i th comparison phase, with sampling performed at the end of the comparison phase. A finite-state machine controls the N differential input stages and their associated transmission gates required to generate the N -bit key. The compensation phase is intentionally made longer than the comparison phase for two reasons. First, the compensation phase must be relatively long to accommodate sufficient transient settling for the auxiliary amplifier to equalize the cascode branch currents. Second, the sampling phase can be relatively short, since sampling V_{out} before it reaches

either supply rail both shortens the per-bit cycle and improves transient stability – settling a node that has not yet hit a rail is more reliable than re-settling one that has already saturated. The corresponding waveforms of V_{out} and $V_{sampling}$ during a representative bit-generation cycle are shown in Fig. 6(c).

The voltage sampled from the FCA must be conditioned into a digital key with a bounded maximum run length to guarantee the QSS of the downstream SCC. The first stage of this back-end chain is an inverter chain, which converts the sampled analog voltage into a clean digital signal. Because V_{out} settles near 0.7 V during the compensation phase, the trip point of the first inverter is also designed at 0.7 V, so that the resulting bit assumes HIGH or LOW with equal probability and the key is statistically balanced (50:50). The output of the inverter chain is then passed through a level shifter that translates the 1.2 V digital signal to the 0.8 V domain. While the FCA requires a 1.2 V supply to achieve sufficient gain to resolve the input-pair mismatch, the digital back-end does not, and operating it from a lower rail reduces both the static and dynamic power of the chain and downstream digital blocks. A second motivation for this voltage down-scaling is related to the HCI mechanism described in Section 4.

A SIPO shift register collects the serial bit stream and assembles it into a 128-bit parallel key. The key is then passed to an MRLC, which deterministically flips the minimum number of bits required to enforce a maximum run length of 2. Because the FCA generates bits at a rate of a few kHz, the SIPO shift register is clocked at the same frequency. The MRLC, by contrast, is a parallel-in serial-out (PISO) block and is clocked at the much higher CLK frequency of the SCC, so that its serial output can drive the LVD switches at the required switching rate.

4 Hot Carrier Injection

4.1 Modeling of HCI

HCI is a reliability degradation mechanism in which carriers in the MOSFET channel gain sufficient kinetic energy from the lateral electric field to overcome the $Si - SiO_2$ energy barrier and become trapped in the gate oxide or at the interface. The resulting fixed charge shifts up

the threshold voltage V_{th} and lowers the carrier mobility, leaving a permanent shift in the device characteristics. The damage is particularly severe in NMOS devices because the electron mobility is roughly three times higher than the hole mobility in silicon, allowing electrons to reach the injection-threshold energy more easily than holes. Unlike NBTI, which exhibits partial recovery once the gate stress is removed, HCI damage involves irreversible structural trap formation that does not anneal at operating temperatures.

HCI is most severe when V_{GS} and V_{DS} are simultaneously high: a high V_{GS} supplies a large channel current that provides carriers, while a high V_{DS} establishes the strong lateral electric field that heats them. This condition arises mainly during switching transitions, when the transistor traverses the saturation region. The time evolution of ΔV_{th} follows a power law,

$$\Delta V_{th} \propto t^{0.45} \quad (7)$$

The modeling of HCI-induced ΔV_{th} in the SCC follows the analytical formulation of [25],

$$\Delta V_{th} = \frac{q}{C_{ox}} K_2 \sqrt{Q_i} \exp\left(\frac{E_{ox}}{E_{o2}}\right) \exp\left(-\frac{\varphi_{it}}{q\lambda E_m}\right) t^{n'} \quad (8)$$

$$E_{ox} = \frac{(V_{gs} - V_{th})}{t_{ox}}, \quad E_m = \frac{V_{ds} - V_{dsat}}{l} \quad (9)$$

$$Q_i = C_{ox}(V_{gs} - V_{th}) \quad (10)$$

which characterizes HCI degradation in a 65 nm CMOS technology. Although the present work targets the same node, the foundry differs from that of [25], and the device-level parameters, most notably the oxide thickness, are adjusted to match the TSMC 65-nm PDK used here. The calibrated formulation is then combined with the operating point extracted from Cadence Spectre simulation, and the resulting ΔV_{th} is evaluated numerically through a Python implementation. While the majority of prior work on HCI seeks to mitigate its impact in order to preserve long-term reliability [26, 27], this work instead exploits HCI as a physical security mechanism. To maximize the stress accumulated in the LVD NMOS during IK operation (Fig. 7(a)), the two LVD switches of the SCC are sized highly asymmetrically. As shown in Fig. 7(b), nearly the entire 1.2 V supply that bridges V_{IN} and ground during the presence of shoot-through current is dropped




Fig. 7 (a) Shoot-through current inducing HCI in the LVD NMOS during IK operation. (b) Voltage division between the on-resistances of the switches.

across the NMOS drain–source terminals, while the PMOS sustains only a few tens of millivolts. This asymmetric voltage division concentrates the lateral field, so that the HCI-induced degradation is almost exclusively confined to the LVD NMOS, transforming it into a dedicated “sacrificing” device that records each unauthorized access attempt as a permanent V_{th} shift.

The formulation of [25] provides the ΔV_{th} expression but evaluates it under the assumption of a fixed operating point: V_{th} and the associated conditions (V_{ds} , V_{dsat} , E_m) are held at their initial values throughout the projected lifetime. While this assumption is adequate for conventional reliability analysis of digital and analog circuits, where ΔV_{th} remains a small fraction of V_{th} , it underestimates the degradation in the present scenario, which intentionally maximizes the effect of HCI: the asymmetric stress drives the NMOS toward a regime in which ΔV_{th} becomes large enough to alter the amount of stress itself.

To capture this, this work extends the analytical framework into an iterative form: after each IK event, the updated V_{th} is fed back into the voltage-divider model to recompute V_{DS} and E_m , and the next increment of ΔV_{th} is evaluated at the new operating point. The resulting positive-feedback loop (Fig. 8) yields a degradation trajectory that progresses substantially faster




Fig. 8 Positive feedback of HCI.

than the constant- V_{th} prediction, and more accurately reflects the actual aging behavior under repeated IK attempts.

4.2 Effect of HCI on the SCC

Fig. 9 shows the HCI-induced degradation of the LVD NMOS in the SCC. The x-axis represents the cumulative stress time, defined as the total duration for which shoot-through current is present. Since a shoot-through event occurs only when the CK and INPUT correspond to one of four equally likely combinations, the probability of stress during each clock cycle is 25%. Consequently, one year of continuous brute-force probing corresponds to approximately three months of cumulative stress on the NMOS. The y-axis shows the resulting threshold-voltage shift, ΔV_{th} . The proposed iterative model, which accounts for the positive-feedback effect, is compared with a constant- V_{th} baseline. The two curves begin to diverge after approximately one year of cumulative stress, indicating the onset of positive-feedback-induced degradation. The effect of ΔV_{th} on converter performance is discussed in Section 6.

The proposed SCC employs separate analog and digital supply domains. The analog circuitry, including the SCC and the PUF, operates from a 1.2-V supply, while the digital circuitry that generates the maximum-run-length-controlled key for the LVD switches operates from a 0.8-V supply. Since the key directly drives the LVD switches, the digital supply voltage determines the available gate-drive margin before HCI-induced ΔV_{th} affects converter operation. Using a 0.8-V key




Fig. 9 HCI-induced ΔV_{th} of the LVD NMOS switch versus cumulative shoot-through stress time.

reduces this margin by 400 mV compared with a conventional SCC, where both the analog and digital circuits operate from a common 1.2-V supply. This reduction does not affect correct-key (CK) operation, because the LVD PMOS threshold voltage remains greater than 400 mV under all operating conditions, and the resulting increase in switch on-resistance causes less than a 1% reduction in power conversion efficiency.

Reducing the digital supply below 0.8 V would cause the SCC to fail earlier under HCI stress; however, two practical limitations arise. First, the LVD PMOS threshold voltage sets the minimum supply required for reliable switching. Second, as $|V_{GS} - V_{TH}|$ decreases, the switch on-resistance increases rapidly, leading to a noticeable reduction in power conversion efficiency even under CK operation.

5 Key Space and Entropy Analysis

5.1 Theoretical Analysis of Key Space

The bit length of the PUF-generated key in this work is 128, which can be modified flexibly, since an N -bit extension of the key only requires N additional input transistor pairs and a few milliseconds. The magnitude of the total key space can be computed through a recurrence formula, since the maximum run length is limited to 2.

$$a_n = a_{n-1} + a_{n-2}, \quad (n \geq 3, a_1 = 2, a_2 = 4) \quad (11)$$

The recurrence follows the Fibonacci form, since any valid n -bit key can be partitioned into two disjoint groups according to its terminal bits: those ending in 11 or 00, and those ending in 01 or 10. Solving this recurrence yields a key space of $a_{128} \approx 8.15 \times 10^{26}$, which corresponds to an effective key length of $\log_2(8.15 \times 10^{26}) \approx 89.4$ bits, or approximately 89 bits. The resulting key space constitutes only 8.15% of the full 128-bit domain 2^{128} .

$$t_{eval} = N_{bit} \times T_{SCC,CLK} = \frac{N_{bit}}{f_{SCC,CLK}} \quad (12)$$

Here, t_{eval} denotes the evaluation time required per IK, N_{bit} is the total key length, and $f_{SCC,CLK}$ is the operating frequency of the SCC. For the proposed SCC operating at 3.5 MHz, a single IK evaluation requires $36.6 \mu s$, limiting the attacker to at most 27,344 IK trials per second. Even under the assumption that the attacker has identified the maximum-run-length constraint of 2, an exhaustive search of the entire key space would require 9.46×10^{14} years. In the absence of such knowledge, the corresponding search time extends to 4×10^{26} years for the full 128-bit space. These estimates assume that the attacker waits for each IK evaluation to complete. Although early termination is, in principle, possible when the first 12 bits fail to drive the SCC into QSS, the resulting reduction is negligible compared to the overall magnitude of the required search time.

5.2 Effect of MRLC on Key Space and Entropy

5.2.1 Raw PUF Key Statistics

The PUF generator described consists of 128 differential input pairs, one per output bit. The MRLC algorithm is a deterministic, purely digital greedy scan that flips the minimum number of bits required to enforce a circular maximum run length of two. Its behavior depends only on the input bit pattern and the maximum-run-length parameter.

The 100 raw 128-bit keys generated by the PUF are evaluated against the ideal of an unbiased, memoryless Bernoulli(0.5) binary source, using two complementary criteria: bit-balance statistics (Hamming weight) and run-length distribution. The same statistics are then computed on the MRLC-processed keys, and the effect of




Fig. 10 Maximum run length distribution across 100 raw 128-bit PUF keys, compared with the theoretical distribution for an i.i.d. Bernoulli(0.5) source.

MRLC is quantified in terms of the bit-flip rate and the corresponding shifts in these two statistics.

Across 100 raw 128-bit keys (12,800 bits aggregate), the PUF generates HIGH and LOW bits in a measured ratio of 50.24% / 49.76%, well within the 95% confidence interval [49.38%, 51.11%] of an ideal unbiased Bernoulli source. The per-key HIGH count has mean $\mu = 64.31$ and standard deviation $\sigma = 4.93$, in close agreement with the binomial expectation ($\mu = 64$, $\sigma = \sqrt{32} \approx 5.66$) for $n = 128$ independent Bernoulli(0.5) draws. The inter-key Hamming distance (HD), averaged over 4,950 pairs, is 49.73%, confirming inter-key uniqueness consistent with the same i.i.d. assumption.

Fig. 10 shows the distribution of the maximum (circular) run length L_{max} observed in each of the 100 raw keys, together with the theoretical distribution for a 128-bit i.i.d. Bernoulli(0.5) sequence, computed by dynamic programming over run-state transitions. The measured distribution peaks at $L_{max} = 7$ (27/100 keys), bracketed by $L_{max} = 6$ (23/100) and $L_{max} = 8$ (12/100), and closely matches the theoretical mode and right tail. The measured mean $\langle L_{max} \rangle = 7.52$ agrees with the analytical expectation of 7.33. A χ^2 goodness-of-fit test against the theoretical distribution, with bins $\{5, 6, 7, 8, 9-10, 11 \leq\}$, yields $\chi^2 = 8.30$ (df = 5, $p = 0.14$). The measured distribution is therefore statistically indistinguishable from that of an ideal i.i.d. binary source at $\alpha = 0.05$.

5.2.2 MRLC-Processed Key Statistics

Fig. 11 shows the histogram of per-key HIGH counts before and after MRLC, overlaid with




Fig. 11 Histogram of the per-key HIGH count for the 100 PUF keys, before and after MRLC processing.

the ideal Binomial(128, 0.5) distribution. Applying MRLC to each raw key produces 100 corrected keys whose per-key HIGH count has mean $\mu_{MRLC} = 64.05$ and standard deviation $\sigma_{MRLC} = 2.02$, with an aggregate HIGH/LOW ratio of 50.04% / 49.96%. The mean is essentially unchanged from the raw case, while σ contracts by a factor of approximately 2.5: the MRLC, by replacing locally repeated bits with their complements, removes the rare keys with strong Hamming-weight excursions and concentrates the per-key weight tightly around 64. The inter-key HD, averaged over the same 4,950 pairs after MRLC, is 49.75%, still statistically indistinguishable from the raw value.

By construction, the MRLC enforces $L_{max} \leq 2$ in every key. Across all 100 corrected keys, the empirical distribution collapses to a delta function at $L_{max} = 2$ (100/100 keys); no key exhibits $L_{max} = 1$, consistent with the fact that a length-1-only sequence would require strict alternation and is statistically improbable in 128 bits.

5.2.3 Quantifying the Effect of MRLC

The statistical characterization [28] of the PUF keys before and after MRLC processing is summarized in Table 1. The MRLC modifies, on average, $\langle n_{flip} \rangle = 21.60$ bits per 128-bit key ($16.88\% \pm 2.62\%$), with a minimum of 14 and a maximum of 30 flips across the 100 key samples. The flip count correlates strongly with the raw L_{max} : keys whose raw L_{max} exceeds 10 require 25–30 flips, whereas keys with raw $L_{max} = 5$ typically require 15 flips. The lower bound is set by the number of length-3-or-greater runs in the raw key; every such run forces at least $\lfloor (L-1)/(L_{max}+1) \rfloor$ flips to break it into sub-runs of length ≤ 2 .

Table 1 Statistical Characterization of the 128-bit PUF Keys ($N = 100$)

Metric	Raw	After MRLC	Theoretical
Mean HIGH ratio (uniformity)	50.24%	50.04%	50.00%
Std. dev. of per-key HIGH ratio	3.85%	1.58%	4.42%
Mean max run length, $\langle L_{\max} \rangle$	7.52	2.00	7.33
Mean inter-key HD (uniqueness)	49.73%	49.75%	50.00%
Mean bit-flip ratio (MRLC overhead)	—	16.88%	—




Fig. 12 Circular autocorrelation function (ACF) of the 100 PUF keys before (blue) and after (red) MRLC processing.

MRLC preserves the mean Hamming weight ($\Delta\mu = -0.26$ bits) but reduces its variance ($\sigma : 4.93 \rightarrow 2.02$). This is an expected consequence of replacing the most “uneven” runs of identical bits with alternating patterns: long runs of zeros (or ones) produce extreme Hamming weights, and shortening them moves each key closer to the 50:50 mean. The inter-key HD changes by less than 0.1 percentage points ($49.73\% \rightarrow 49.75\%$), and its standard deviation increases modestly ($5.98 \rightarrow 7.61$). The MRLC therefore does not degrade keyspace uniqueness; each key remains approximately as distinguishable from every other key as in the raw case. Fig. 12 shows the circular autocorrelation function (ACF) of the 100 PUF keys before and after MRLC processing. For the raw keys (blue), the mean ACF remains within the ± 0.1 band, which is the empirical randomness criterion used in prior SMDL work [11]. The 5–95 percentile envelope across the 100 keys is tightly centered on zero, confirming the absence of any measurable long-range correlation in the PUF output, consistent with an i.i.d. binary source.

The MRLC-processed keys (red), by contrast, exhibit a sharply localized short-range structure: the mean ACF reaches approximately -0.33 at




Fig. 13 Per-bit Shannon entropy H_j across the 128 bit positions of the MRLC-processed keys.

lags $k = \pm 1, \pm 2$ and $+0.34$ at $k = \pm 3$, returning to within the ± 0.1 band for $|k| \geq 5$. This is a mathematical consequence of the enforced maximum run length $L_{\max} = 2$: after MRLC, every run consists of either one or two identical bits with nearly equal frequency (50.3% and 49.7% in the present dataset), producing a bounded oscillatory pattern at the length-1-to-2 scale. This structure is fully deterministic given the raw input and therefore introduces no information leakage beyond the entropy of the raw key itself; the cryptographic strength of the key against brute-force attack is instead governed by its total keyspace entropy, which we analyze next. Fig. 13 plots the per-bit Shannon entropy H_j across the 128 bit positions of the MRLC-processed keys. The mean per-bit entropy is $\langle H \rangle = 0.9933$ bit/position, with all positions exceeding 0.95 bit/position. This result approaches the theoretical limit of $H = 1$ bit/position, which holds exactly for any source where $P(\text{bit} = 1) = 0.5$ at every position. The small gap between the measured $\langle H \rangle = 0.9933$ and the ideal value of 1 is due to finite-sample noise at $N = 100$, rather than to any structural bias introduced by the MRLC. The observed per-position minimum of 0.958 falls within the 5–95 percentile range




Fig. 14 Power conversion efficiency versus operating frequency.

expected under the i.i.d. null hypothesis at $N = 100$ ($[0.904, 0.965]$, obtained by Monte Carlo), confirming the absence of position-dependent bias.

6 Simulation Results

6.1 Maximum Power Point of SMDL SCC

As introduced earlier, $f_{eq} \approx 0.67f_{CLK}$ in the SMDL SCC, implying that the optimum frequency at which maximum power conversion efficiency is achieved will change when the SCC is driven with non-periodic digital keys.

Fig. 14 shows the simulated power conversion efficiency as a function of operating frequency for the clock-pulse-driven and digital-key-driven SCCs, swept from 1 MHz to 4 MHz. The clock-pulse-driven SCC exhibits a well-defined efficiency peak of 83.9% at 2.0 MHz, with a smooth and monotonic roll-off on both sides. Efficiency rises steadily from 1 MHz, since switching losses are smaller relative to conduction losses; above the peak, it declines gradually to 80.1% at 4 MHz as dynamic losses begin to dominate. The flat, symmetric shape of the curve reflects the deterministic, fixed-duty-cycle nature of the clock drive signal. The digital-key-driven SCC displays markedly different behavior. At low frequencies, efficiency is significantly degraded, dropping to 73.6% at 1 MHz. This is due to the data-dependent switching pattern of the key signal, which introduces irregular charge-transfer intervals that are particularly pronounced at longer periods. As frequency increases, efficiency improves, and the two curves converge near 2.3–2.6 MHz, where their difference




Fig. 15 DC gain and phase margin of the PUF control loop under temperature variation.

narrows to within 1%. Beyond this crossover, the key-driven SCC exceeds the clock-driven counterpart, reaching a peak efficiency of 84.6% at 3.7 MHz. As discussed in Section 2, the equivalent switching frequency of the SMDL SCC is approximated as $f_{eq} \approx 0.67f_{CLK}$; consequently, the frequency at which maximum power conversion efficiency is achieved shifts relative to the clock-driven case, as shown in the simulation results of Fig. 14.

6.2 Process and Temperature Variation of the PUF System

The PUF system must guarantee reliable operation under external variations such as process and temperature. Monte Carlo simulations were performed to evaluate the controllability of the compensation-phase feedback loop under these variations, quantifying how reliably the loop equalizes the cascode branch currents across the expected operating corners.

Fig. 15 shows the DC gain and phase margin of the FCA compensation loop over the temperature range -40° to 100°C . The loop maintains a DC gain above 60 dB and a phase margin above 45° across the entire range, ensuring that the auxiliary amplifier can reliably equalize the two cascode branch currents under all expected operating temperatures.

Fig. 16 shows the DC gain and phase margin of the loop under process variation at a nominal temperature of 27°C . Similarly, a DC gain above 60 dB and a phase margin above 45° are maintained in most cases. These criteria on DC gain and phase margin ensure stable closed-loop settling of the auxiliary amplifier. Although the 25 nA input-pair tail current limits the available slewing current, the compensation capacitor is sized such that V_{out} can complete its transient within the allotted compensation phase, balancing




Fig. 16 DC gain and phase margin of the PUF control loop under process variation.

slew-limited settling speed against the area cost of a larger capacitor.

6.3 Effect of Bit Errors and Threshold Voltage Shift on SCC Performance

The performance of the SCC is evaluated using three figures-of-merit: power efficiency, average output voltage $V_{out,avg}$, and startup time t_{2V} , defined as the time required for the output voltage to reach 2 V from initialization. The analysis assumes that the load circuit powered by the SCC operates correctly only when supplied with a stable dc voltage of at least 2 V. This 2 V threshold corresponds to 83% of the ideal 2.4 V doubled output for $V_{IN} = 1.2$ V, and serves as a practical indicator that the converter has reached its nominal operating range. In practice, the exact $V_{out,min}$ and output-ripple limits required for reliable load operation are system-specific and depend on the downstream circuitry supplied by the SCC.

Fig. 17 shows the average output voltage. All curves decrease monotonically with ΔV_{th} , and the ideal 2.4 V doubled output degrades progressively as either n_{err} or ΔV_{th} increases; for example, even in the absence of ΔV_{th} , an IK with 64 errors already fails to reach 2 V at the output.

The asymmetric effect of HCI degradation on authorized versus unauthorized operation is one of the key findings of this work. For an authorized user supplying the CK, the converter remains within specification even at $\Delta V_{th} = 150$ mV of the LVD NMOS, with $V_{out,avg}$ staying above 2 V. In contrast, under unauthorized operation with $n_{err} = 32$, the same $\Delta V_{th} = 150$ mV is already sufficient to push $V_{out,avg}$ below the 2 V threshold, and the droop deepens further as n_{err} grows. The HCI mechanism therefore degrades the converter selectively along the attack direction, preserving normal functionality while progressively disabling its response to incorrect keys.




Fig. 17 Average output voltage versus ΔV_{th} and n_{err} .




Fig. 18 Power conversion efficiency versus ΔV_{th} and n_{err} .

Fig. 18 shows the combined effect of ΔV_{th} and bit-error count n_{err} on power conversion efficiency. Efficiency degrades substantially as n_{err} increases, as expected. A more subtle feature distinguishes the correct-key (CK) and incorrect-key (IK) regimes: every IK curve exhibits an interior maximum in ΔV_{th} , whereas the CK curve decreases monotonically. This arises from two competing loss mechanisms whose ΔV_{th} dependences point in opposite directions. The first is the shoot-through loss of the (CK, INPUT) = (HIGH, LOW) phases introduced in Section 2. The energy dissipated per shoot-through phase scales as $V_{IN}^2 t_{on} / R_{on}$, so increasing ΔV_{th} raises R_{on} and thereby reduces the shoot-through loss. This mechanism dominates on the left side of the peak, and is the reason a small ΔV_{th} can actually improve the efficiency of IK operation.

The second is the conduction loss of the 128 – n_{err} correct-switching cycles per period, which carry the actual load charge through the LVD




Fig. 19 Startup time versus ΔV_{th} and n_{err} .

switches in triode operation. The energy dissipated by these cycles scales as $I^2 R_{on}$ per transferred packet, so increasing ΔV_{th} instead raises the conduction loss:

$$R_{on} = \frac{1}{\mu C_{ox} \frac{W}{L} (V_{gs} - V_{th})} \quad (13)$$

Because R_{on} grows hyperbolically as $V_{gs} - V_{th}$ decreases, the conduction loss becomes steep at moderate-to-large ΔV_{th} and dominates on the right side of the peak.

Fig. 19 shows the startup-time violation induced by error bits and HCI degradation, where t_{2V} is the startup time required for V_{OUT} to reach 2 V and enter the QSS; operating points for which V_{OUT} never reaches 2 V are omitted. This plot exposes the functional death point of the SCC: once accumulated HCI has shifted V_{th} by approximately 160 mV, even CK operation can no longer reach 2 V. IK operation reaches this boundary much earlier – the $n_{err} = 64$ curve, for instance, fails to reach 2 V after a 180 mV shift, corresponding to approximately 2.22 years of accumulated stress.

A startup-time detector, built from two D-flip-flops (DFFs), two AND gates, and a voltage comparator, is driven by the SCC to detect startup-time extension under IK operation (Fig. 20). The voltage comparator, with hysteresis, compares the SCC output voltage against a reference voltage; if the output exceeds the reference, the comparator output propagates through the downstream DFF. DFF1 is clocked by a single-pulse deadline (DL) signal, whose timing marks the startup-time limit: the SCC output must exceed V_{REF} before DL and hold that state, so that DFF1 can capture the comparator output. Under IK operation,




Fig. 20 Block diagram of a startup-time detector.




Fig. 21 Timing diagram of the startup-time detector.

if the startup time exceeds DL, DFF1 latches a LOW at the clock's rising edge and feeds it to an AND gate and to DFF2. DFF2's output is initialized HIGH, so as to keep the AND gate active under CK operation; however, once DFF2 receives a LOW, it is fed back to the AND gate, producing an irreversible lock. A final AND gate ensures that EN initializes LOW. Once startup has been verified, the detector can be powered down to reduce static consumption.

Fig. 21 plots the SCC output voltage, comparator output, DL, and EN during the startup stage. The selected IK contains only 8 error bits – a statistically improbable sample – yet the startup-time detector still distinguishes it through deliberate DL timing and comparator-voltage selection. The comparator reference voltage is set to 2.0 V with a ± 0.05 V hysteresis band. An error bit near $t = 3 \mu s$ introduces a voltage difference that delays the comparator output and ultimately prevents EN from activating. Realistic IK attempts typically exhibit error counts well

above 8 bits, making them substantially easier for the detector to identify.


6.4 Self-Destruction and Hardware Security

According to the analysis above, a brute-force attacker must identify the CK before ΔV_{th} reaches 180 mV, since the SCC cannot produce a 2 V output once $\Delta V_{th} > 180$ mV. Reaching this point requires approximately 3.5 years of cumulative shoot-through current, which, under the 25% duty cycle of brute-force probing, corresponds to roughly 14 years of unauthorized access attempts, or 1.19×10^{13} IK trials. This attempt window is 14 orders of magnitude shorter than the time required to exhaust even the entropy-reduced 89-bit key space.

$$P_{success} = \frac{N_{trials}}{N_{keyspace}} = \frac{1.19 \times 10^{13}}{8.15 \times 10^{26}} \approx 1.15 \times 10^{-14} \quad (14)$$

The attacker therefore faces a power supply unit that ages itself out of legitimate operation before any meaningful fraction of the key space can be searched, yet degrades far too slowly for that degradation itself to constitute a practical attack vector.

As shown in the simulation results, IKs with a Hamming distance of 32 or greater from the CK produce noticeable degradation across all three performance metrics, most prominently in efficiency. In contrast, provided that HCI-induced degradation remains within acceptable bounds, IKs with a Hamming distance below 32 can yield performance comparable to that of the CK in every metric except power conversion efficiency. Nevertheless, as summarized in Table 1, the probability that a brute-force trial falls within a Hamming distance of 16 from the CK is sufficiently small to be considered negligible. For an attacker aware of the MRLC, the expected time required to encounter an IK within a Hamming distance of 16 is approximately 1.25 years; for an attacker unaware of this constraint, that expected time extends to 1.27×10^6 years. Furthermore, since the MRLC is implemented as on-chip digital circuitry, its existence is inherently difficult for an external attacker to identify or exploit.



Total Active Area: 0.0550 mm²
Security Overhead: 18.4% (relative to baseline)

Fig. 22 Area overhead of the proposed SMDL-SCC.

6.5 Performance Comparison

Table 2 compares the performance of the proposed SMDL-SCC with prior work, including the SMDL-FCA and SMDL-bandgap reference (BGR) [11]. Compared to this prior work, the IK variation is improved because the startup-time detector can distinguish small startup-time violations caused by bit errors and digitize the result, so that the power supply to downstream logic or circuitry is blocked under IK access.

Fig. 22 shows the estimated area overhead of the proposed SMDL-SCC after layout. The SCC, with two MIM flying capacitors and a MOS output capacitor, consumes 0.0464 mm², while the PUF generation system and digital circuitry consume only 0.0051 mm² and 0.0035 mm², respectively. Since the flying and output capacitors dominate the area of the entire system, the digital back-end and PUF generation blocks together occupy a relatively small area compared to the SCC. The resulting security area overhead is therefore 18.4%, smaller than both cases reported in [11]. Multiple CKs do exist in this SMDL-SCC, since IKs with no error bits during the startup period can pass the startup-time detection phase; however, this probability is statistically negligible, and any bit errors occurring after startup detection induce severe voltage droops that prevent the SCC from maintaining QSS.

Table 2 Comparison of the proposed analog locking technique with other analog locking methods.

	ISCAS'19 [29]	TETC'20 [5]	TVLSI'21 [12]	ICCD'18 [30]	ASHES'22 [31]	VTS'19 [14]	TVLSI25 [11]	This Work
Technique	Bias Control	Calibration Control		Design-Based		ANN	SMDL	
Technology	180 nm	180 nm	65 nm	14 nm	65 nm	130 nm	65 nm	65 nm
Analog Circuit	PLL	BPF/LNA/LDO	$\Sigma\Delta$ ADC	$(I/V)_{\text{REF}}/\text{AMP}$	OTA	LNA	FCA/BGR	SCC
Key Size (bit)	24	80/84/122	194	75-eq. [#]	36-eq. [#]	14-eq. [#]	128	128
Key Entropy	Weak	Weak	Weak	–	–	Weak	High	High
Multiple CK	Yes	No	No	No	No	Yes	No	Negligible
Repeated Keys	Yes	Yes	Yes	No	No	No	No	No
IK Variation	3%	27%	96.8%	92%	55%	20%	$\approx 100\%$	100%
Perf. Over.	7.6%	NA	NA	< 1%	NA	< 1%	< 0.1%	0%
Area Over.	145%	8.8/2.6/3.1%	0	3%	175%	High	116%/45%*	18%
Power Over.	14%	0	0	NA	70%	sub- μ W	1%/76%*	5 μ W

[#] Binary equivalent key size.

* Overhead includes the overhead of the PUF, which lies outside the area of the analog circuit.

7 Conclusion

This paper presented an SCC driven by a 128-bit non-periodic digital key generated from an on-chip PUF. Unlike conventional clock-driven SCCs, the proposed architecture performs correct power conversion only when the user-supplied key matches the on-chip reference key. Statistical evaluation confirmed high key entropy, with an effective key space of 8.15×10^{26} max-run-constrained sequences. Under IK operation, the SCC exhibits substantial degradation in power conversion efficiency, output voltage, and startup time. Moreover, repeated shoot-through current during unauthorized operation accelerates HCI, progressively increasing the threshold voltage of the switching devices and ultimately rendering the converter inoperable, thereby introducing a hardware-level constraint on brute-force attacks by limiting the number of incorrect authentication attempts a device can sustain. Future work will focus on experimental validation of this HCI-induced degradation on fabricated silicon, extending the security analysis to representative analog loads, and addressing the residual attack surface posed by timing side-channels and low-error incorrect-key conditions.

Statements and Declarations

Competing Interests

The authors declare no competing interests.

References

- [1] Rajendran, J., Zhang, H., Zhang, C., Rose, G.S., Pino, Y., Sinanoglu, O., Karri, R.: Fault analysis-based logic encryption. *IEEE Transactions on Computers* **64**(2), 410–424 (2015) <https://doi.org/10.1109/TC.2013.193>
- [2] Roy, J.A., Koushanfar, F., Markov, I.L.: Ending piracy of integrated circuits. *Computer* **43**, 30–38 (2010)
- [3] Subramanyan, P., Ray, S., Malik, S.: Evaluating the security of logic encryption algorithms. In: 2015 IEEE International Symposium on Hardware Oriented Security and Trust (HOST), pp. 137–143 (2015). <https://doi.org/10.1109/HST.2015.7140252>
- [4] Rao, V.V., Savidis, I.: Parameter biasing obfuscation for analog ip protection. In: 2017 IEEE International Symposium on Hardware Oriented Security and Trust (HOST), pp. 161–161 (2017). <https://doi.org/10.1109/HST.2017.7951825>
- [5] Jayasankaran, N.G., BÓrbon, A.S., SÁnchez-Sinencio, E., Hu, J., Rajendran, J.: Towards provably-secure analog and mixed-signal locking against overproduction. *IEEE Transactions on Emerging Topics in Computing* **10**(1), 386–403 (2022) <https://doi.org/10.1109/TETC.2020.3025561>
- [6] Leroux, N., Manea, P.-P., Sudarshan, C., Finkbeiner, J., Siegel, S., Strachan, J.P.,

- Neftci, E.: Analog In-Memory Computing Attention Mechanism for Fast and Energy-Efficient Large Language Models (2024). <https://arxiv.org/abs/2409.19315>
- [7] Rostami, M., Koushanfar, F., Karri, R.: A primer on hardware security: Models, methods, and metrics. *Proceedings of the IEEE* **102**(8), 1283–1295 (2014) <https://doi.org/10.1109/JPROC.2014.2335155>
- [8] Alkabani, Y.M., Koushanfar, F.: Active hardware metering for intellectual property protection and security. In: 16th USENIX Security Symposium (USENIX Security 07). USENIX Association, Boston, MA (2007). <https://www.usenix.org/conference/16th-usenix-security-symposium/active-hardware-metering-intellectual-property-protection>
- [9] Guin, U., Huang, K., DiMase, D., Carulli, J.M., Tehranipoor, M., Makris, Y.: Counterfeit integrated circuits: A rising threat in the global semiconductor supply chain. *Proceedings of the IEEE* **102**(8), 1207–1228 (2014) <https://doi.org/10.1109/JPROC.2014.2332291>
- [10] Acharya, R.Y., Chowdhury, S., Ganji, F., Forte, D.: Attack of the genes: Finding keys and parameters of locked analog ics using genetic algorithm. In: 2020 IEEE International Symposium on Hardware Oriented Security and Trust (HOST), pp. 284–294 (2020). <https://doi.org/10.1109/HOST45689.2020.9300277>
- [11] Kumar, U., Khanna, S., Mittal, A., Shrivastava, A.: Protecting analog circuits using switch mode time domain locking. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems* **33**(4), 916–928 (2025) <https://doi.org/10.1109/TVLSI.2025.3528320>
- [12] Elshamy, M., Sayed, A., Louërat, M.-M., Aboushady, H., Stratigopoulos, H.-G.: Locking by untuning: A lock-less approach for analog and mixed-signal ic security. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems* **29**(12), 2130–2142 (2021) <https://doi.org/10.1109/TVLSI.2021.3117584>
- [13] Aljafar, M.J., Azais, F., Flottes, M.-L., Pagliarini, S.: Leveraging Layout-based Effects for Locking Analog ICs (2022). <https://arxiv.org/abs/2209.01856>
- [14] Volanis, G., Lu, Y., Rao Nimmalapudi, S.G., Antonopoulos, A., Marshall, A., Makris, Y.: Analog performance locking through neural network-based biasing. In: 2019 IEEE 37th VLSI Test Symposium (VTS), pp. 1–6 (2019). <https://doi.org/10.1109/VTS.2019.8758657>
- [15] Ramadass, Y.K., Chandrakasan, A.P.: Voltage scalable switched capacitor dc-dc converter for ultra-low-power on-chip applications. In: 2007 IEEE Power Electronics Specialists Conference, pp. 2353–2359 (2007). <https://doi.org/10.1109/PESC.2007.4342378>
- [16] Le, H.-P., Sanders, S.R., Alon, E.: Design techniques for fully integrated switched-capacitor dc-dc converters. *IEEE Journal of Solid-State Circuits* **46**(9), 2120–2131 (2011) <https://doi.org/10.1109/JSSC.2011.2159054>
- [17] Butzen, N., Krishnamurthy, H.K., Yu, J., Ahmed, Z.K., Weng, S., Ravichandran, K., Hosseinian Ahangharnejhad, R., Waldemer, J., Peltó, C., Tschanz, J.W.: A monolithic 12.7 w/mm², 92% peak-efficiency switched-capacitor dc-dc converter using cscr-first topology. *IEEE Journal of Solid-State Circuits* **59**(12), 4114–4123 (2024) <https://doi.org/10.1109/JSSC.2024.3465388>
- [18] Mirchandani, N., Shafiee, N., Fei, Y., Shrivastava, A.: An ultra-low power and lower area current-mode based physically unclonable function with less than 100nw power consumption and a native instability of 0.6875% for iot applications. In: 2020 IEEE 63rd International Midwest Symposium on Circuits and Systems (MWSCAS), pp. 273–276 (2020). <https://doi.org/10.1109/MWSCAS48704.2020.9184628>
- [19] Dickson, J.F.: On-chip high-voltage generation in mnos integrated circuits using an improved voltage multiplier technique. *IEEE*

- Journal of Solid-State Circuits **11**(3), 374–378 (1976) <https://doi.org/10.1109/JSSC.1976.1050739>
- [20] Favrat, P., Deval, P., Declercq, M.J.: A high-efficiency cmos voltage doubler. IEEE Journal of Solid-State Circuits **33**(3), 410–416 (1998) <https://doi.org/10.1109/4.661206>
 - [21] Kumar, U., Mittal, A., Muncuk, U., Shrivastava, A.: Non-iterative square-wave modeling of multistage cmos dickson rectifiers for rf energy harvesting. IEEE Open Journal of Circuits and Systems **7**, 82–94 (2026) <https://doi.org/10.1109/OJCAS.2026.3666784>
 - [22] Maksimovic, D., Dhar, S.: Switched-capacitor dc-dc converters for low-power on-chip applications. In: 30th Annual IEEE Power Electronics Specialists Conference. Record. (Cat. No.99CH36321), vol. 1, pp. 54–591 (1999). <https://doi.org/10.1109/PESC.1999.788980>
 - [23] Butzen, N., Steyaert, M.S.J.: Scalable parasitic charge redistribution: Design of high-efficiency fully integrated switched-capacitor dc-dc converters. IEEE Journal of Solid-State Circuits **51**(12), 2843–2853 (2016) <https://doi.org/10.1109/JSSC.2016.2608349>
 - [24] Deng, S., Chou, H.-C., Zhang, C., Kwak, J., Asokan, A.D., Quezada, E., Sarkar, E., Lee, J., Shin, J., Chakraborty, D., Sathe, V.S., Yu, S., Datta, S.: High-efficiency beol-compatible on-chip switched capacitor dc-dc converter using planar and vertical oxide transistors with trench capacitors. In: 2025 IEEE International Electron Devices Meeting (IEDM), pp. 1–4 (2025). <https://doi.org/10.1109/IEDM50572.2025.11353815>
 - [25] Wang, W., Reddy, V., Krishnan, A.T., Vattikonda, R., Krishnan, S., Cao, Y.: Compact modeling and simulation of circuit reliability for 65-nm cmos technology. IEEE Transactions on Device and Materials Reliability **7**(4), 509–517 (2007) <https://doi.org/10.1109/TDMR.2007.910130>
 - [26] Babu, S.J., Hu, F., Zhu, L., Singhal, S., Guo, X.: Extending Silicon Lifetime: A Review of Design Techniques for Reliable Integrated Circuits (2025). <https://arxiv.org/abs/2503.21165>
 - [27] Bravaix, A., Cacho, F., Federspiel, X., Mora, P.: Hot-carrier injection degradation in advanced cmos nodes: A bottom-up approach to circuit and system reliability, 401–448 (2015) https://doi.org/10.1007/978-3-319-08994-2_14
 - [28] Maiti, A., Gunreddy, V., Schaumont, P.: A Systematic Method to Evaluate and Compare the Performance of Physical Unclonable Functions. Cryptology ePrint Archive, Paper 2011/657 (2011). <https://eprint.iacr.org/2011/657>
 - [29] Rao, V.V., Savidis, I.: Mesh based obfuscation of analog circuit properties. In: 2019 IEEE International Symposium on Circuits and Systems (ISCAS), pp. 1–5 (2019). <https://doi.org/10.1109/ISCAS.2019.8702671>
 - [30] Saki, A.A.-, Ghosh, S.: How multi-threshold designs can protect analog ips. In: 2018 IEEE 36th International Conference on Computer Design (ICCD), pp. 464–471 (2018). <https://doi.org/10.1109/ICCD.2018.00075>
 - [31] Aljafar, M.J., Azais, F., Flottes, M.-L., Pagliarini, S.: Leveraging layout-based effects for locking analog ics. In: Proceedings of the 2022 Workshop on Attacks and Solutions in Hardware Security. ASHES’22, pp. 5–13. Association for Computing Machinery, New York, NY, USA (2022). <https://doi.org/10.1145/3560834.3563826> . <https://doi.org/10.1145/3560834.3563826>