Nuclear Inst. and Methods in Physics Research, A 982 (2020) 164577

Contents lists available at ScienceDirect

Nuclear Inst. and Methods in Physics Research, A

journal homepage: www.elsevier.com/locate/nima

A secure measurement unit for an inspection system used in nuclear arms-control verification

Fred N. Buhler[1], David K. Wehe[1,*], Michael P. Flynn[1]

College of Engineering, University of Michigan, Ann Arbor, MI 48109, USA

ARTICLE INFO

Keywords:
ADC
Information barrier
Nonproliferation
Arms control

ABSTRACT

Verification of arms control treaties may require information barriers to protect sensitive data acquired during the verification process. These information barriers are commonly implemented in software, but must store and operate on data that are vulnerable to attack and tampering. We present a new Secure Measurement Unit based on a modified pipeline SAR ADC architecture. The idea is to move the information barrier back into the digitization process so that prohibited data is never created. A prototype 65 nm CMOS Secure Measurement Unit (SMU) was tested using U-235, U-238, Co-60, Cs-137 and Am-241 sources. The prototype accurately digitized allowable signals while being immune to side-channel attacks on signals within preset forbidden regions.

1. Introduction

Verifying compliance with the terms of future arms control agreements may involve radiation measurements of the objects, such as neutron counting, or more commonly, through gamma ray spectral measurements. As described in Ref. [1], the host party must certify that the measurement device meets not only the local facility security requirements, but also does not reveal sensitive information. On the other hand, the inspecting party must ensure that the measurement device produces results that will noninvasively authenticate the proffered object is as asserted. If a proposed authentication process acquires information that could be deemed by the host as sensitive, an information barrier must be included to block the inspecting party from that data. The security of the information barrier thus plays a rather critical role in treaty compliance negotiations.

Two general classes of authentication methods have evolved: attribute-based and template-based. In the attribute-based method, properties of the inspected object that will provide confidence in its authenticity are measured. Examples of such systems include the Controlled Intrusiveness Verification Technology (CIVET) [2], Trusted Radiation Attribute Demonstration System TRADS [3], Third Generation Attributes Measurement System 3G-AMS [4], Fieldable Nuclear Material Identification System FNMIS [5], and the UK–Norway Initiative UKNI [6].

In the template matching class, the object to be authenticated is compared to another device known to be authentic (i.e., the Master- or Golden-Copy). Systems such as Next Generation Trusted Radiation Identification System NG-TRIS [7] and CONFIDANTE [8], among others, follow this paradigm. In an ideal case, authentication can be accomplished without revealing sensitive data by only looking at the differences from the template. The reader is referred to Yan [9] or Hamel [10] and references therein for further details on instantiated verification systems.

Some of the proposed systems require a gamma ray spectrum, as produced from an HPGe detector, and embed an information barrier in internal software. An example is the warhead verification concept proposed Vavrek [11] that uses HPGe detectors to detect transmission nuclear resonance fluorescence gamma rays to achieve cryptographic warhead verification. The acquired gamma spectra may inadvertently reveal sensitive information, but an ADC that suppresses any secondary lines would harden this technique. Indeed, a natural concern of both the host and inspecting party is the vulnerability of any sensitive information residing internally. In our approach, we address these vulnerabilities by moving the information barrier into the digital measurement process. The idea is to avoid generating sensitive data, but retain only information needed for the verification process.

A typical implementation of an attribute system employing information barriers is depicted in Fig. 1. The detector converts ionizing radiation from the inspected material into weak electrical pulses, which are integrated, amplified, and shaped into Gaussian pulses, and then a peak-hold circuit captures the pulse peak amplitude. An Analog-to-Digital Converter (ADC) samples and digitizes the peak amplitude, which is proportional to the energy deposited in the detector. A computer or Field-Programmable-Gate-Array (FPGA) receives the digital

* Corresponding author.
E-mail addresses: fbuhler@umich.edu (F.N. Buhler), nima@umich.edu (D.K. Wehe), mpflynn@umich.edu (M.P. Flynn).
[1] All authors contributed to the work being reported.

https://doi.org/10.1016/j.nima.2020.164577
Received 22 October 2019; Received in revised form 18 August 2020; Accepted 19 August 2020
Available online 22 August 2020
0168-9002/

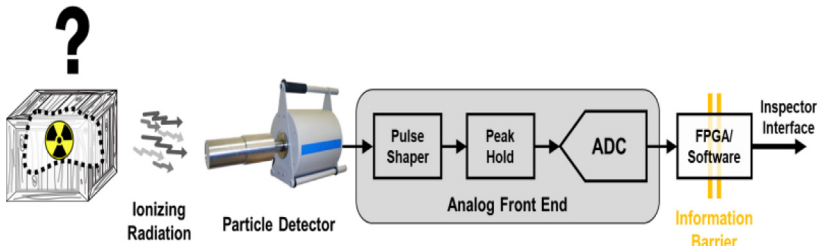

Fig. 1. Attribute verification scenario with FPGA/software information barrier.

Fig. 2. Example spectra of ionizing radiation with Information Barriers. In this scenario, U-235, Cs-137, and U-238 sources are permitted to be observed, while Am-241 and Co-60 are prohibited.

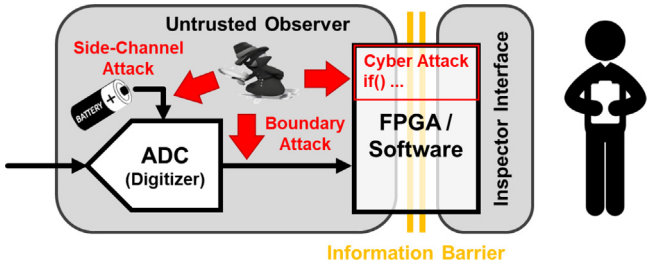

Fig. 3. Potential vulnerabilities of an information barrier implemented in FPGA/software.

code representing the peak amplitude, and implements the information barrier by selectively deleting sensitive data and only providing allowed information.

As an example, Fig. 2 shows a spectrum of ionizing radiation for a sample source containing U-235, U-238, Co-60, Cs-137, and Am-241. Suppose the sensitive information is contained within the yellow regions (at left), i.e., spectral data from these regions should not be revealed, then the final spectrum presented to the inspector (at right), after the information barrier, is devoid of information from the sensitive regions.

The sensitive information contained within the information barrier opens potential vulnerabilities in the FPGA/software. As shown in Fig. 3, an untrusted observer might obtain sensitive information before it is deleted by launching a boundary attack on the digital communication between the ADC and the FPGA/software. Information barrier software is difficult to verify and vulnerable to external cyber-attacks. Finally, commercial ADCs are vulnerable to relatively simple side-channel attacks, aided by machine-learning techniques, to extract restricted data by observing power supply activity. A better approach would prevent sensitive digital information from being created, thus eliminating the boundary and cyber-attack vulnerabilities. While the analog–digital interface provides an earlier step to establish the information barrier, it requires a hardware redesign of the standard ADC architecture, as described next.

2. Successive approximation (SAR) assisted pipeline ADC

We modified the common SAR-assisted pipeline ADC architecture [12,13] to both embed an information barrier and to secure the ADC from side-channel attack. To better understand the new architecture, we first review the operation of the conventional SAR ADC and SAR-assisted pipeline ADC architectures. We then explain how we implement the information barrier in the SAR-assisted pipeline and discuss how we harden the new ADC to side-channel attack.

2.1. Background: Conventional SAR ADC

A SAR ADC [12,13] digitizes an analog signal by sampling the analog voltage and then comparing it to a sequence of voltages. Typically, a SAR ADC applies a binary voltage search to find the closest equivalent digital code. In a switched-capacitor SAR ADC, charge redistribution on a binary-weighted capacitor array, often called a capacitor DAC or CDAC, generates the search voltages. In the example shown in Fig. 4, the CDAC array first samples the analog input voltage while the bottom plates of the CDAC capacitor array are grounded. A set of switches sequentially connects the bottom plates of the CDAC capacitors to the positive or the negative reference voltages. This action adds a scaled version of the reference voltage (depending on the capacitance value) to the top plate of the capacitor array. The search first adds enough voltage to move the top plate voltage by half of the full-scale voltage range of the ADC. The subsequent decision adds or removes one-quarter of the full-scale range, and so on. A finite state machine labeled “SAR Logic” determines the switching sequence based on the sign of the top-plate voltage as determined by the comparator. At the end of the decision chain, the voltage on the top plate of the capacitor array approaches zero. Since the voltages added to and subtracted from the top plate are known, the sampled analog input voltage is also known. An N-bit SAR ADC makes N decisions, which result in 2^N possible combinations of addition and subtraction of voltage. The SAR decision process is shown in tree form in Fig. 5 where each digital code represents a unique set of decisions.

2.2. Background: Conventional SAR assisted pipeline ADC

As an advancement over the SAR ADC, the SAR-assisted pipeline ADC introduced in Refs. [14,15], extends the energy-efficiency benefits of the SAR ADC to higher resolution. The SAR-assisted pipeline ADC is a two-stage pipeline that cascades two lower-resolution SAR sub-ADCs in a pipeline (Fig. 6) [15]. An advantage is that these lower-resolution SAR sub-ADCs are very compact and very energy efficient. The SAR-assisted pipeline operates in three phases. In the first phase (Phase 1), SAR sub-ADC 1 makes a low-resolution digital estimate of the analog input. In the second phase (Phase 2), a DAC creates the analog

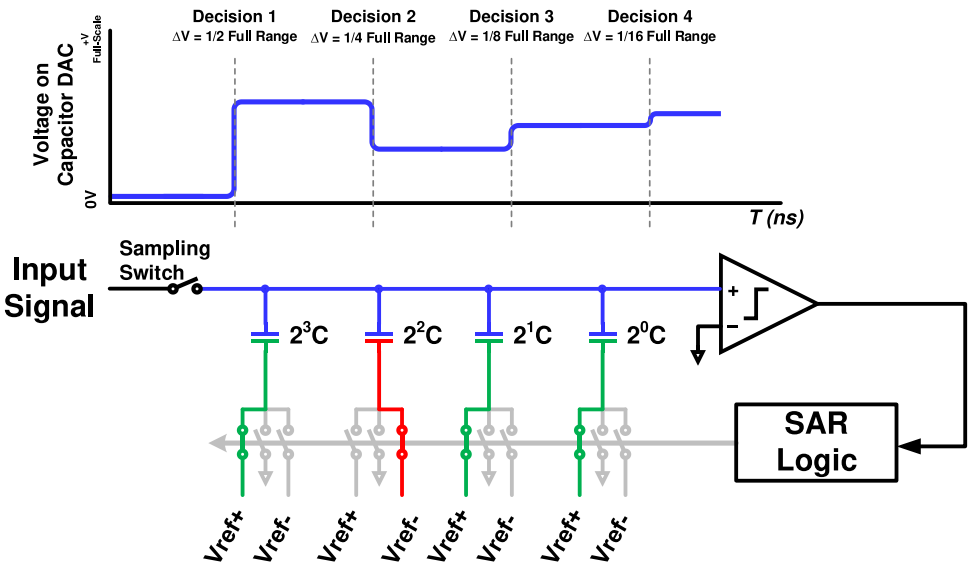

Fig. 4. Schematic of a SAR ADC with a 4-bit binary-weighted switched capacitor DAC (CDAC) array.

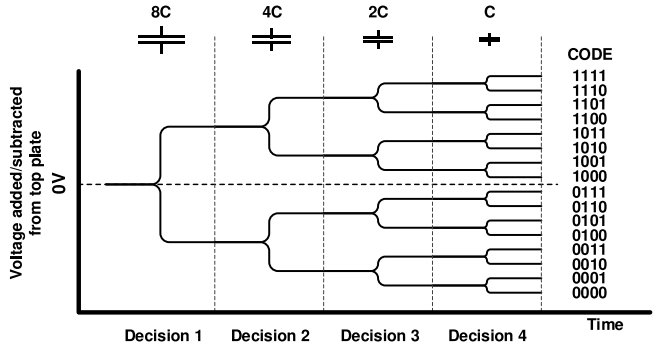

Fig. 5. Voltage added or subtracted to the top plate of the CDAC in a 4-bit SAR ADC versus time.

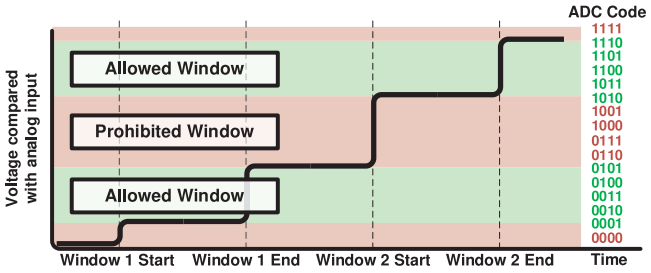

Fig. 7. Example of 4-bit sub-ADC 1 determining if the input is within two allowed code windows or three prohibited code windows. The trace represents the sequence of boundary voltage checks.

representation of this estimate, which is subtracted from the input signal to form the voltage residue of the first conversion. This residue is amplified to the full-scale resolution of the second stage sub-ADC. Finally, in the third phase (Phase 3), the sub-ADC 2 digitizes this residue. The digital outputs of the two sub-ADCs are combined to give the overall conversion result. The advantage of the SAR-assisted pipeline ADC architecture is that it relaxes the noise requirements on the second stage sub-ADC and enhances the overall ADC speed while maintaining excellent power efficiency [16–19].

3. Information barrier within the digitization process

To implement an information barrier within the digitization process, the two-stage SAR-assisted pipeline ADC architecture described above must be adapted. The first SAR sub-ADC is initially used to determine whether the analog input lies within an allowed code window (i.e., amplitude range) or within a prohibited code window. If the input is within an allowed range, then the full ADC hardware can be engaged to digitize the analog input. Conversely, if the input is not in an allowed region, then the ADC does not process the analog input. For illustration, Fig. 7 depicts an example with two allowed windows and three prohibited windows.

To implement an information barrier in sub-ADC 1, we developed an algorithm named the Consecutive SAR (C-SAR). In the initial step of C-SAR, sub-ADC 1 compares the analog input to the boundaries of the allowed and prohibited windows. After sampling the analog input voltage on the CDAC array, we consecutively check each boundary starting from the lowest boundary to the highest boundary. Each step in the C-SAR algorithm connects the CDAC capacitor bottom plates to a pre-determined combination of reference voltages to check a particular boundary. This consecutive code check is different from the traditional SAR algorithm which performs a binary search of all possible ADC codes. The example in Fig. 7 has two allowed windows, and therefore

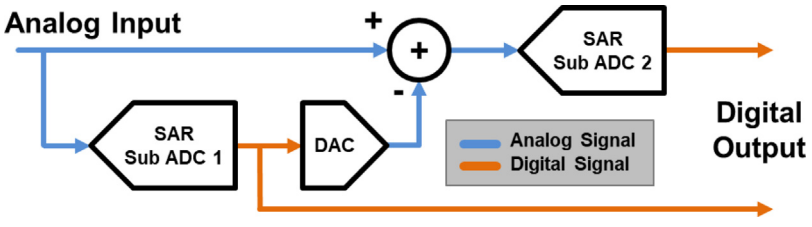

Fig. 6. Two-stage SAR-assisted pipeline ADC.

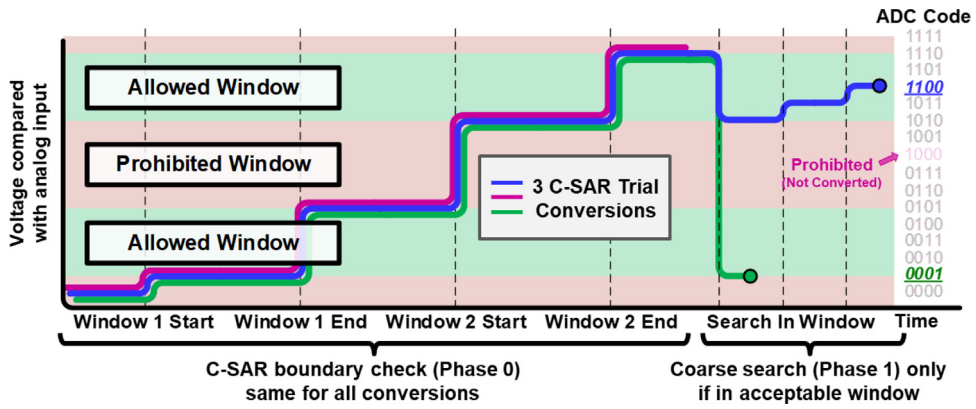

Fig. 8. C-SAR boundary check (Phase 0) followed by initial digitization by sub-ADC 1 for allowed windows (Phase 1). The traces are for three different analog input values, one of which is in a prohibited window.

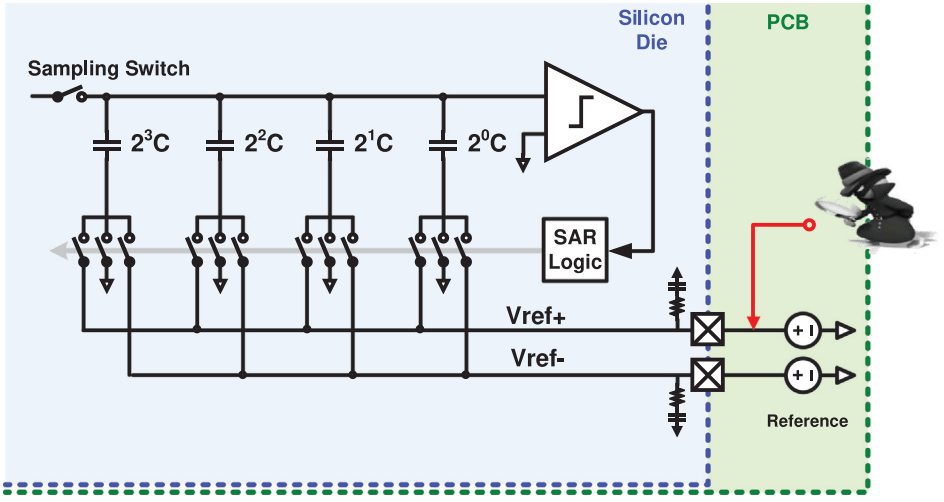

Fig. 9. An untrusted observer monitoring the current flowing through the external voltage reference of the CDAC in a SAR-ADC.

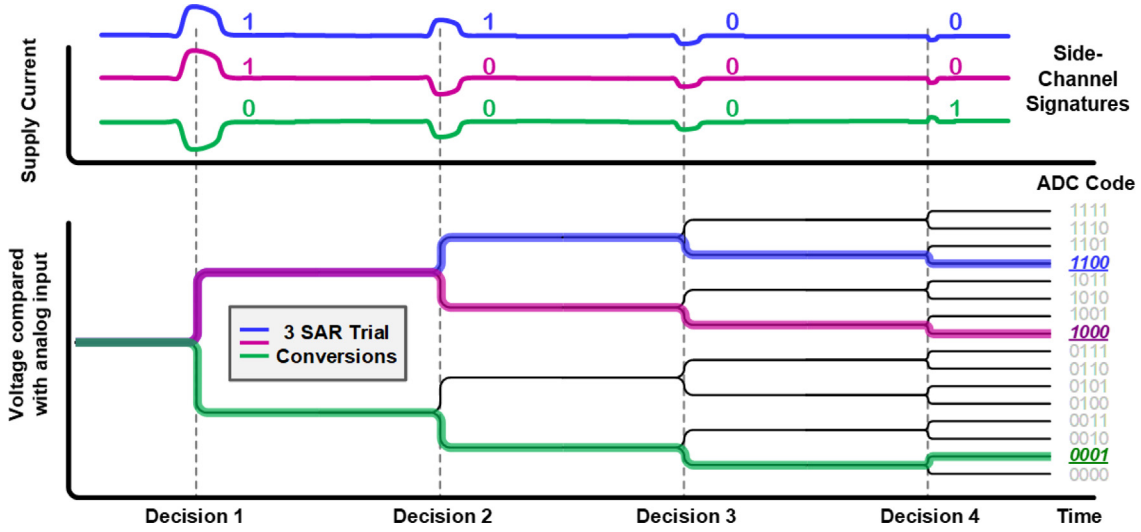

Fig. 10. (top) current flowing into voltage reference and (bottom) CDAC search voltage sequence for three different input voltages to a SAR ADC.

there are four boundary edges. The C-SAR algorithm consecutively checks the analog input against all four boundaries. If the sampled input is determined to be within an acceptable window, it will be allowed to be digitized to the full accuracy of the ADC, as described next. Alternatively, if the sampled input is not within an allowed window, the analog to conversion stops. To be consistent with the terminology of Section 2.2 above, we label this boundary check process as Phase 0 since it is a process to determine whether digitization is allowed.

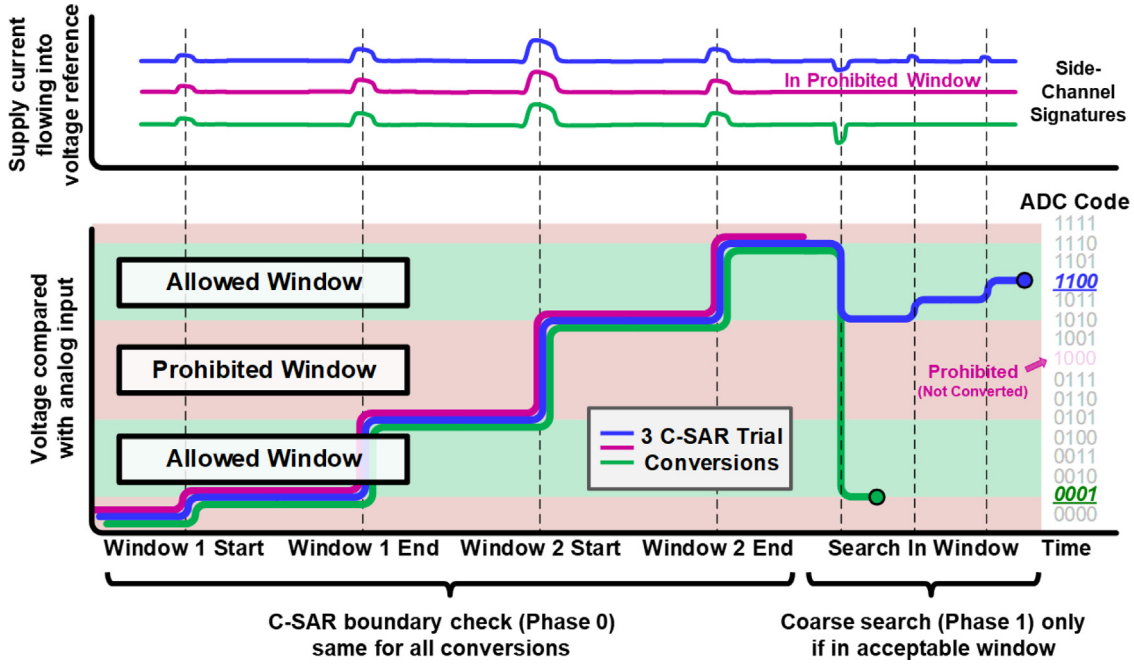

Fig. 11. Three sample C-SAR conversions and resulting reference current waveforms showing a constant side-channel signature for all three during the window detection phase (Phase 0). The coarse digitization phase (Phase 1) follows only if the analog input is within an acceptable window.

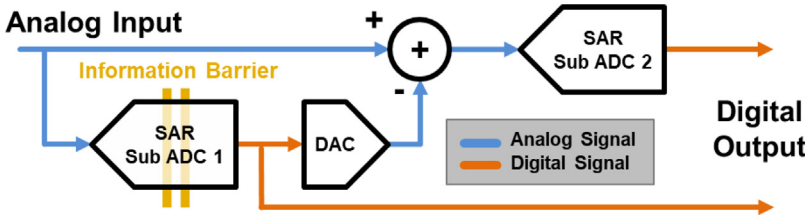

Fig. 12. The Secure Measurement Unit (SMU) is a 2-stage SAR Assisted Pipeline ADC, with the first sub-ADC using C-SAR logic to implement the Information Barrier. The C-SAR algorithm is immune to side-channel attacks and negates boundary attacks since the SMU only digitizes allowed information.

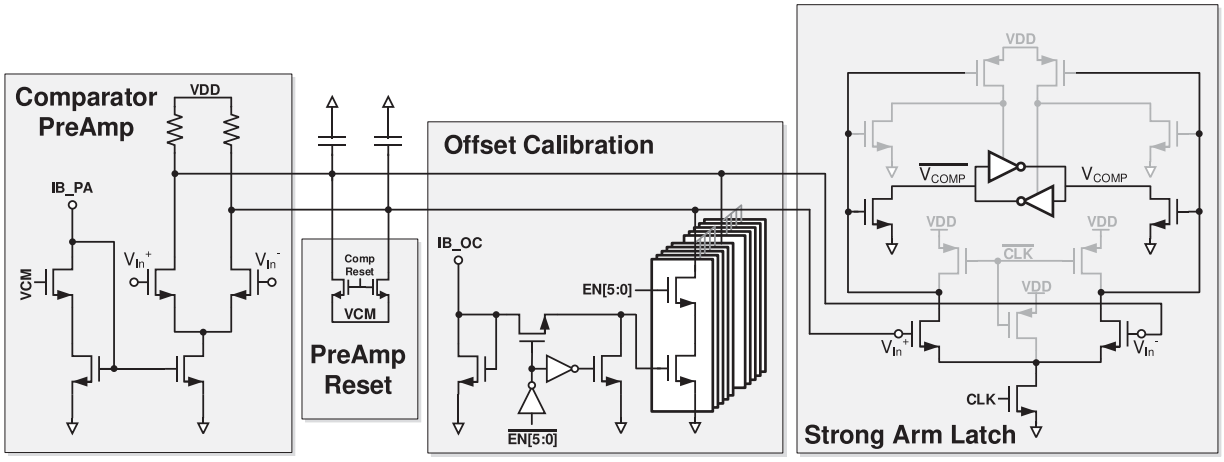

Fig. 13. Schematic of comparator showing preamplifier and offset calibration.

If the analog input passes the Phase 0 test of acceptability, digitization occurs in three phases. The first phase (Phase 1) again makes use of the hardware in sub-ADC 1, reusing its CDAC to conduct a coarse search within the allowed window. As shown in Fig. 8, when the input is found to be in an allowed window, we compare against a stepped ramp of voltages to provide a coarse (i.e., 4 bits in the example) digitization of the input. This search is only within the allowed window and therefore avoids partially digitizing the forbidden regions. This Phase 1 search differs from the first phase in a conventional SAR-assisted pipeline (cf. Section 2.2) where sub-ADC performs a binary search of the entire signal range. Phase 2 and Phase 3 make standard use of the SAR Assisted Pipeline architecture to complete the high-resolution digitization of the signal. But this digitization process is only activated if the analog input lies in an allowed window.

4. Hardening against side channel attacks

Any verification system acquiring protected data should not only implement an information barrier but also be resilient to indirect methods of unmasking the data. These indirect methods of discovery are referred to as side-channel attacks. With a conventional ADC, a side-channel attack could observe the external power supply current to determine protected information. As our modified SAR-assisted ADC relies on two SAR sub-ADCs, we first explain the vulnerabilities of the conventional SAR ADC. We then explain how our C-SAR-based information barrier avoids the side-channel attack vulnerabilities of the traditional SAR operation.

Each digital code generated by a SAR ADC represents a unique sequence of decisions, and thus produces a unique sequence of directions of charge flowing to and from the CDAC references. If these reference currents can be observed (Fig. 9), then information can be obtained about the ADC digitization process. In particular, if the direction of the charge can be determined, then the analog input to the ADC can be inferred. As an example, Fig. 10 shows side-channel signatures for three different SAR conversion values. The externally observed sequences of charges directions are directly correlated with the reported ADC digital code. Thus, an embedded information barrier in a conventional SAR ADC that digitizes, but does not report the digital code, is highly vulnerable to side-channel attacks since an untrusted observer could determine the digital code based solely on the side-channel signature.

The C-SAR algorithm avoids the side-channel attack vulnerability of the traditional SAR ADC. The example in Fig. 11 shows the side channel signatures for the same analog inputs as for the conventional SAR ADC example in Fig. 10. Since the C-SAR algorithm checks all window edges sequentially regardless of the comparator decisions, the reference current polarity sequence is the same for all input values. Thus, the side-channel signature is the same for all conversions during the window detection phase (i.e., Phase 0).

As discussed above, if the input is in an allowed range then the digitization by the modified SAR-pipeline SAR continues through the three phases (i.e., Phases 1 to 3). As shown in Fig. 11, during Phase 1 while sub-ADC 1 is performing a coarse digitization it produces a unique, and thus vulnerable, side-channel signature. However, this coarse search only processes allowed information, so side-channel attack exploits are of no concern. Similarly, Phase 2 and Phase 3 operate only if the input is in an allowed window and therefore cannot disclose secure information through a side channel.

5. Circuit implementation

The prototype Secure Measurement Unit (SMU) is implemented as a 14-bit pipeline that cascades a 6-bit C-SAR sub-ADC 1 and a traditional 8-bit SAR sub-ADC 2 with an inter-stage amplifier to achieve the high accuracy needed to measure ionizing radiation spectra (Fig. 12). The inter-stage amplifier is implemented as a 3-stage ring-amplifier [16]. The advantages of the ring amplifier are its efficiency and its ability to handle a much larger signal range than a conventional CMOS amplifier. The output of the C-SAR ADC and the 8-bit SAR sub-ADC are then combined to produce the overall 14-bit ADC output. The boundary search sequence and the subsequent digitization are directed by on-chip digital logic.

The nature of the SMU means that some conventional pipeline-ADC techniques cannot be applied. The boundary search must be at the full accuracy (i.e., 14 bits) of the ADC, otherwise the ADC might inadvertently digitize prohibited regions. Conventional pipeline ADCs use redundancy [20] to relax the accuracy of the first sub-ADC. Through redundancy, errors in the first stage are corrected by the second stage. With redundancy, a first-stage sub-ADC error causes a change in the residue which can easily be accounted for by the second stage. While redundancy is very effective, it presents the possibility of information in a prohibited window leaking to the ADC output (i.e., through the residue). To counteract this possibility, we do not use redundancy, and instead we ensure that the first stage sub-ADC operates at the full accuracy of the system (i.e., 14 bits).

A high accuracy comparator is essential for sub-ADC 1 to operate with 14-bit accuracy. The first stage sub-ADC uses the comparator shown in Fig. 13 to determine whether the analog voltage exceeds a particular value. This design uses a pre-amplifier with offset calibration and averaging to reduce errors in the C-SAR-generated comparisons with window-boundary voltages. The pre-amplifier is a resistor-load common source amplifier. A large capacitive load on the pre-amplifier greatly reduces the comparator's noise bandwidth. The preamplifier is reset between comparisons to eliminate memory effects. Running the comparator seven times and taking the average further reduces noise. The simulated input-referred comparator noise is 30 μVrms.

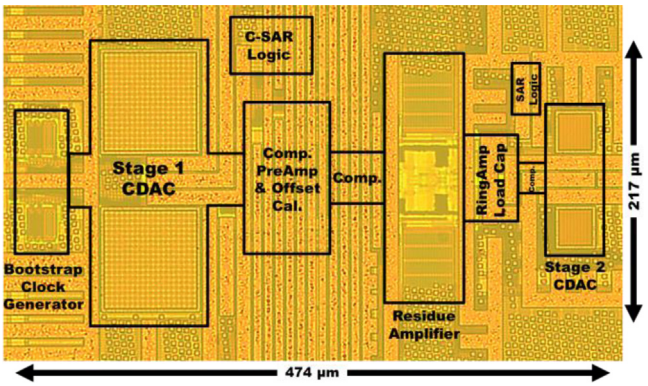

Fig. 14. Die photograph of Secure Measurement Unit.

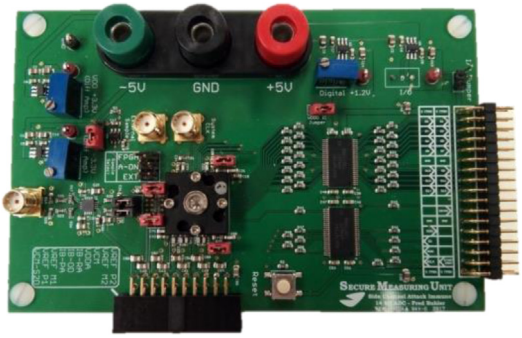

Fig. 15. SMU test board.

Fig. 16. Test board for side-channel attacks.

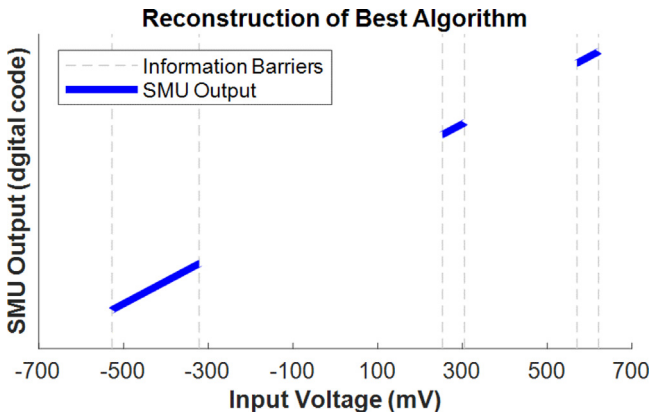

Fig. 17. Digitization output for a ramp voltage input with the SMU operating in secure mode configured with three allowable windows.

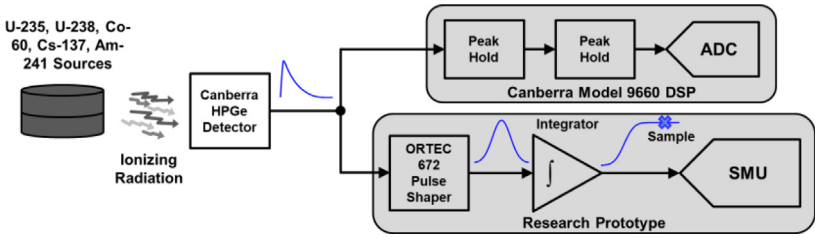

Fig. 18. Radioisotope measurement setup using both the SMU and conventional commercial units for comparison.

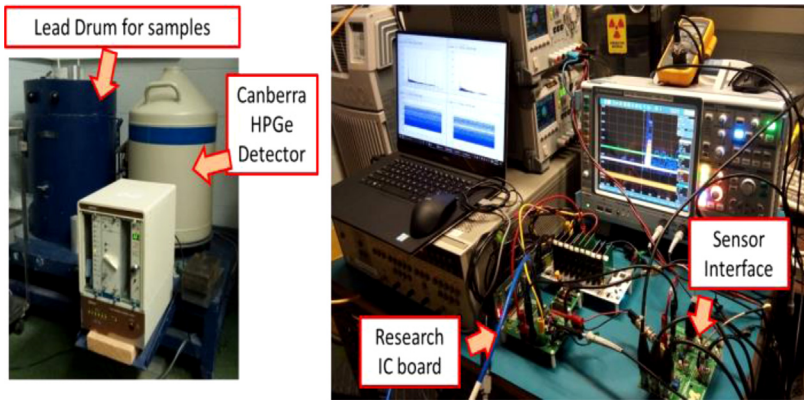

Fig. 19. Measurement lab with commercial and research IC equipment.

A large (8.7 pF) first stage CDAC limits ADC sampling noise (i.e., kT/C noise). The first stage sampling switches are boot-strapped for high linearity. Unlike some designs, the entire first stage capacitance is applied to the first stage sub-ADC (i.e., the C-SAR sub-ADC). While this increases first-stage power consumption, it ensures more accurate sub-ADC decisions. We also make use of the rail-to-rail operation of the ring inter-stage amplifier to enhance the security of the information barrier. Rail-to-rail operation means that out-of-range residues cannot be generated. This natural residue signal clipping means that larger residues due to first stage decision errors cannot be propagated to the second stage.

6. Prototype and measurements

The prototype 14-bit SMU is implemented in 65 nm CMOS and occupies 0.103 mm^2 (Fig. 14). The prototype is packaged in a 56-pin 8 × 8 mm QFN package. A bootstrapped sampling circuit lowers sampling distortion. The first stage CDAC, C-SAR logic, comparator preamp, and the first comparator comprise the first sub-ADC that implements the information barrier. Stage two CDAC, SAR logic, and the second comparator comprise the second sub-ADC. Two custom-designed printed circuit boards support the use and evaluation of the SMU. The board, shown in Fig. 15, connects to the analog source and also interfaces to the controller. The specialized board shown in Fig. 16 is used to test the device against side channel attacks.

6.1. Information barrier

To evaluate the functionality of the information barrier, we configured the SMU to have three allowable input regions and applied a ramp voltage to its analog input. As can be seen in Fig. 17, the SMU only digitized the analog inputs which fell within the allowed regions and generated no output otherwise. The overall measured input-referred noise is 38.4 μVrms, or 0.26 bits at 14-bit level.

6.2. Radioisotope testing

We tested the SMU prototype using combination of radioisotopes to check its ability to measure different spectra. This test was run

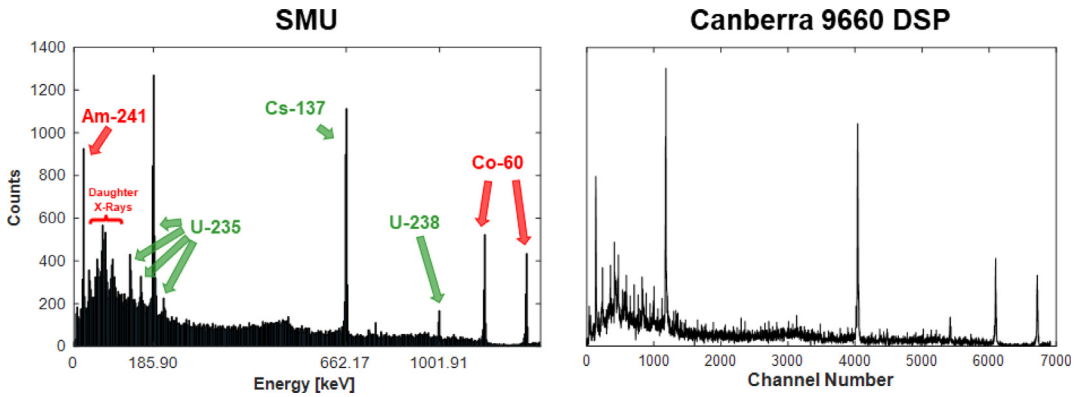

Fig. 20. The energy spectrum of a sample with U-235 (185.90 keV), Cs-137 (662.16 keV), and U-238 (1001.91 keV) (left) measured by SMU measuring over a 5-minute period in non-secure mode (without Information Barrier) and (right) measured with Canberra 9660.

without any restrictive information barrier. The U-235, U-238, Co-60, Cs-137 and Am-241 sources provide energy spectra with unique peaks. The sources were placed within the lead housing that contained a Canberra HPGe detector, and a Canberra 20002C preamplifier provided the analog input for our testing.

To evaluate the accuracy of the research prototype SMU, as shown in Fig. 18 the preamplifier output signal was split into two paths, one path feeding the prototype SMU and the other to a commercial Canberra 9660 unit. An Ortec 972 Pulse Shaper followed by a custom-designed integrator printed circuit board fed the prototype SMU integrated circuit. Fig. 19 shows a picture of the lab setup.

Fig. 20 shows the measured spectra from both the prototype SMU and the commercial Canberra 9660 DSP system. The SMU spectra is comparable to its commercial counterpart in that all peaks from the U-235, U-238, Co-60, Cs-137, and Am-241 are sources clearly visible. A Gaussian fit to the Cs-137 peak counts yielded a Full Width at Half-Maximum (FWHM) of 2.16 keV, slightly worse than the commercial system's FWHM of 1.41 keV.

To test the effectiveness of the information barrier, a spectrum from the U-235, U-238, Co-60, Cs-137, and Am-241 sources was acquired with restricting the input signals to lie within small windows around the U-235 (185.90 keV), Cs-137 (662.16 keV), and U-238 (1001.91 keV) peaks. Fig. 21 plots the measured spectrum over a 300 s acquisition. As intended, the SMU only digitized signals within the allowed range. Note that some Compton background counts from higher energy sources are captured within the allowable range.

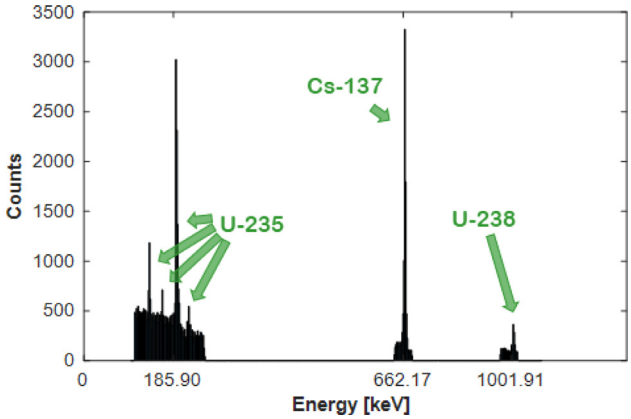

Fig. 21. The energy spectrum of a sample with U-235 (185.90 keV), Cs-137 (662.16 keV), and U-238 (1001.91 keV) measured by the SMU over a 5 min while operating in secure mode, implementing an information barrier using the C-SAR algorithm.

6.3. Side-channel attack setup

We tested the prototype to determine its resilience against side-channel attacks by measuring the reference inputs to infer information about the analog input signal. We tested both with analog signals from a signal generator and signals from a radiation detector.

The test setup is shown in Fig. 22. Two ultra-low-noise voltage sources provided the positive and negative reference voltages for the capacitor DACs (CDACs) in the first and second stages of the ADC. Each reference voltage is heavily decoupled with an on-silicon-die RC filter to reduce noise (Fig. 22). A de-Q resistor on the Printed Circuit Board (PCB) dampens transient voltages caused by the parasitic inductor of the chip package. As a side attack, we measured the differential current flowing through the de-Q resistors to both ADC stages. These current measurements were made by measuring the voltage drop across the de-Q resistors with instrumentation amplifiers.

Fig. 23 shows examples of measured reference currents for a fixed input voltage both in the secure and non-secure modes. Note that in the non-secure mode, the first six distinct peaks (one peak per decision of the first-stage 6-bit sub-ADC) are clearly visible, as expected. Fig. 23 also shows the measured reference current in secure mode. In this test we configured the SMU with three allowed windows; thus, six window edge decisions appear as the first six peaks in the measured current. The fixed-input voltage for this test was in an allowed window, and the subsequent peaks are related to the coarse and fine digitization in Phases 1–3.

6.3.1. Side-channel attack on unsecured ADC

To highlight the vulnerability of conventional ADCs, we tested the resilience of the SMU in non-secure mode. When operating in non-secure mode, the SMU uses the conventional SAR algorithm instead of the secure C-SAR algorithm. We trained machine learning algorithms to predict the analog input voltage based on the reference current signatures. (The Appendix provides more detail about the machine learning algorithms and how they were applied). In non-secure mode, these side-channel signatures are highly correlated to the analog input. We performed training by recording the reference current signatures for 45,000 different input voltages uniformly spanning the input signal range (Fig. 24).

Our experiment shows that machine learning provided a very accurate reconstruction of the analog input from the externally measured reference current. This is shown in Fig. 25, which plots the best reconstruction (from a Rational Quadratic Gaussian Process Regression) from 47 different machine learning algorithms. The blue line represents the ground truth which was the analog voltage supplied to the SMU. The machine-learning prediction of the analog input (red line) is accurate to 9 bits, highlighting that a conventional SAR ADC is highly susceptible to side-channel attacks.

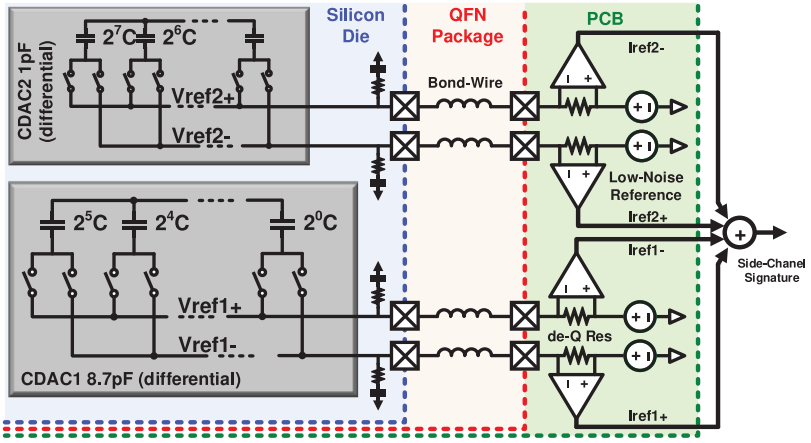

Fig. 22. Side-channel attack measuring the reference current flowing the first and second stage sub-ADCs.

6.3.2. Side-channel attack on secured SMU

We repeat this test with the SMU operating in secure mode (using the C-SAR algorithm) to show the robustness of the proposed scheme to side channel attacks. As before, we trained machine learning algorithms from recorded reference current signatures for 45,000 different input voltages uniformly spanning the input signal range. As shown in Fig. 26, the first six peaks of the side-channel signature are identical for all analog input voltages. Furthermore, no reference current activity is apparent if the SMU finds that the input lies within a prohibited window.

Fig. 27 plots the best reconstruction of the analog input from 47 different machine learning algorithms. The blue line represents the analog voltage supplied to the SMU, which is the ground truth. The red line represents the machine learning predicted (or reconstructed) input voltage. As can be seen, the machine learning algorithm reasonably reconstructed data in allowed windows (this is acceptable since the data are not restricted). However, in the prohibited regions where measurement is not allowed, the machine learning algorithms were unable to reconstruct the analog input.

6.3.3. Side-channel attack of radioisotope measurements in secure mode

To demonstrate the robustness of the SMU in an inspection scenario, we evaluated the SMU against side-channel attacks during the measurement of radioactive sources. Again, we used trained machine-learning algorithms (see Appendix) to predict the actual input signal based on the measurement of external reference currents. For this test, the sources were the same combination of U-235, U-238, Co-60, Cs-137, and Am-241 that produced the measured spectrum shown in Fig. 21.

Fig. 28 shows the most accurate reconstructed spectrum produced by the machine-learning algorithms using the measured side-channel reference currents. The peak at 185 keV represents U-235, which was allowed in this scenario. The algorithm produced a uniform distribution across the restricted windows, i.e., the Am-241 and Co-60 peaks were not revealed. This proved that an untrusted observer could gain no useful information within the restricted energies by observing these external system inputs.

We quantitatively evaluated the machine-learning reconstruction accuracy (defined as effective bits in the signal-to-noise ratio) for energies within allowed windows and energies within prohibited windows and the results are shown in Fig. 29. Reconstruction accuracy within an allowed window is of no concern since it is not sensitive information and will be reported (digitized) regardless. Reconstruction accuracy within a prohibited window represents how susceptible the system is to side-channel attacks. The best machine learning algorithm reconstructed the external signals from within secure windows to ~3 effective bits, which is far too crude to obtain useful information from within the prohibited energy regions.

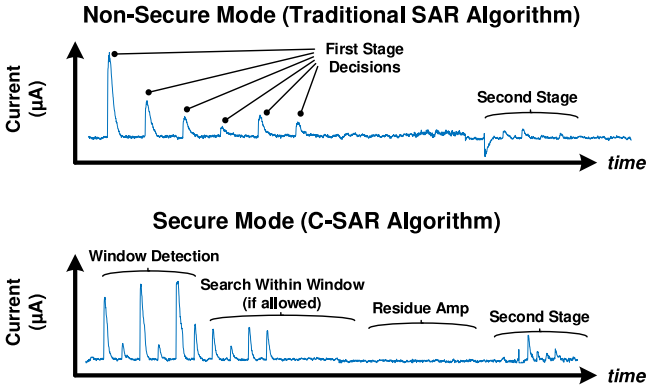

Fig. 23. Measured side channel signatures during a conversion cycle for (Top) non-secure and (Bottom) secure mode.

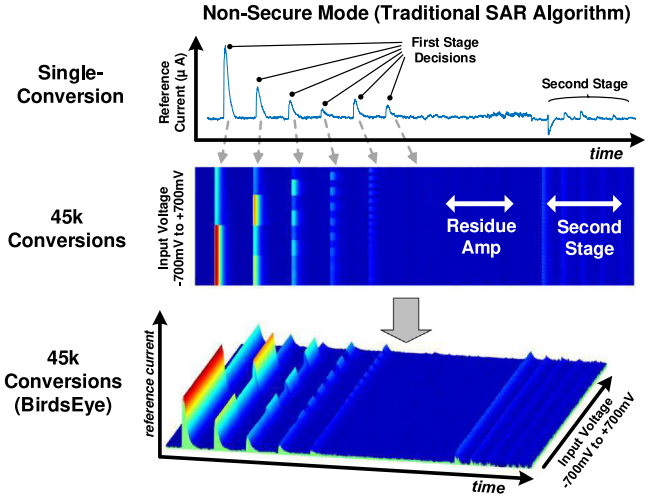

Fig. 24. (Color online) Measured differential reference current waveforms for the SMU in non-secure mode, (i.e., using the traditional SAR ADC algorithm) used for training the machine-learning algorithms. (Top) A single conversion, or digitization, showing a unique peak for each SAR-decision in the first and second stage sub-ADC. (Middle) 45,000 conversions covering the full dynamic range of the SMU showing unique waveforms for each analog input. (Bottom) 3D visualization of the middle figure.

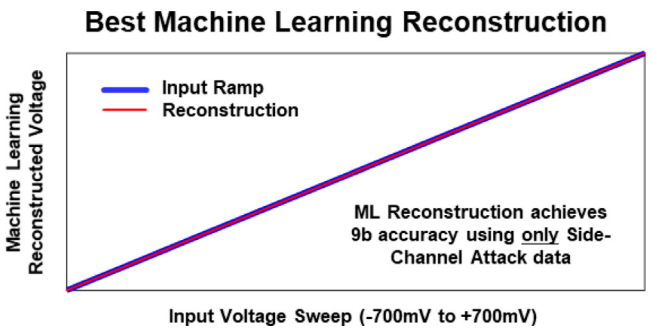

Fig. 25. Machine learning reconstructed voltage versus actual analog input voltage (i.e., ground truth) for the SMU operating in non-secure mode. The best performing algorithm was a Rational Quadratic Gaussian Process Regression.. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

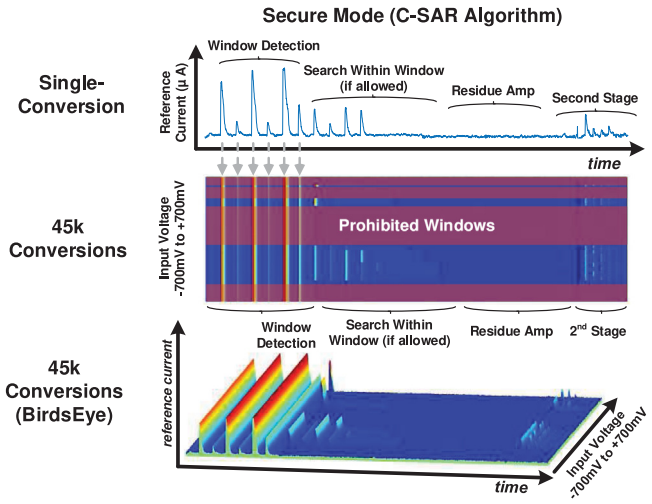

Fig. 26. (Color online) Measured differential reference current versus time for the SMU using the C-SAR ADC algorithm. (Top) A single conversion, or digitization, showing the first six peaks resulting from each of the three window boundaries. (Middle) 45k conversions covering the full dynamic range of the SMU with three allowable windows, showing a constant side-channel signature for all conversions within a prohibited window. (Bottom) 3D visualization of the middle figure.

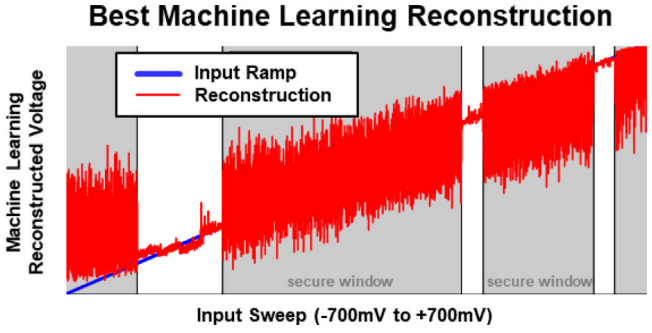

Fig. 27. Machine learning reconstructed voltage versus supplied voltage for SMU in secure mode with three allowable windows.. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

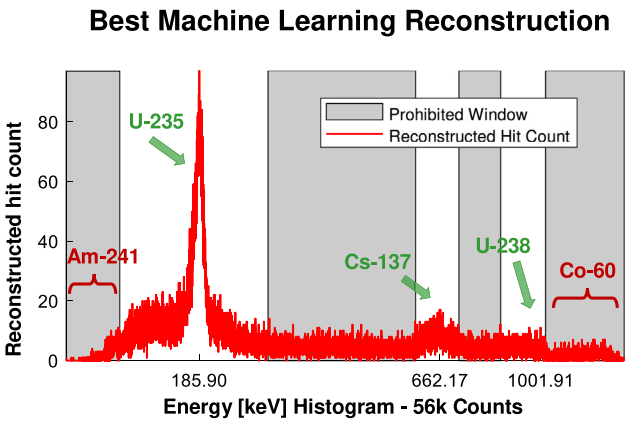

Fig. 28. (Color online) Histogram reconstruction of the most accurate machine learning algorithm.

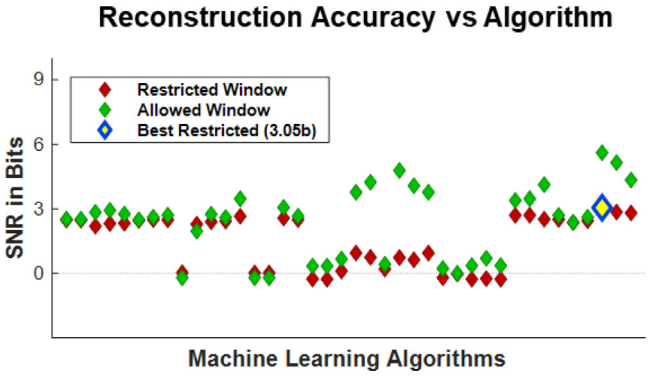

Fig. 29. (Color online). Reconstruction accuracy versus machine learning algorithm.

7. Conclusion

In order to mitigate security vulnerabilities during verification tasks and to enhance the possibility of more intrusive treaty inspections, we moved the information barrier into the digitization process so that sensitive data was never created. This eliminates vulnerabilities present in other verification approaches that acquire the data and then delete it later in an information barrier. A 14-bit prototype, secure, custom-ADC provided HPGe energy resolution comparable to an analog commercial setup but adds an intrinsic information barrier and resilience against side-channel attacks.

Note that moving the information barrier within the digitization process may not be possible for verification scenarios in which sensitive data is required in order to calculate derived quantities, such as ratios (e.g., Pu-240/Pu-239), that are not sensitive. Similarly, the measured counts in an allowable region may contain Compton scattered events from within prohibited regions. For the complex spectra expected in practice, we do not envision this to be a significant security issue, but it should be considered. Finally, one area that deserves further exploration, unrelated to the passive inspection scenario, is the ability to enhance signal throughput by adaptively culling uninteresting data during acquisition. This could be helpful for interrogation scenarios in which the interesting data are a small fraction of the total.

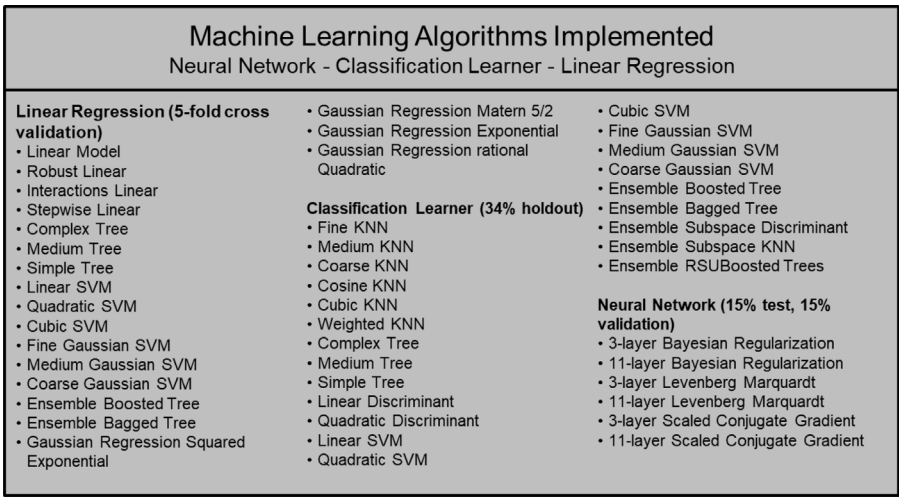

Fig. 30. Machine learning algorithms used for predicting the analog input.

CRediT authorship contribution statement

Fred N. Buhler: Execution of R&D, Manuscript composition. **David K. Wehe:** Conceptualization, Guidance, Manuscript composition. **Michael P. Flynn:** R&D guidance, Manuscript edits.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work was performed within the Consortium for Verification Technology and funded by the Department of Energy National Nuclear Security Administration award number DE-NA0002534.

Appendix

Machine learning was implemented using MATLAB's (R2018b) Machine Learning Toolbox using the standard methods shown in Fig. 30. The programmed input voltage is used as the response variable for all algorithms. The data shown in Figs. 24 and 26 are used as predictors. These side-channel signatures shown in Figs. 24 and 26 are aligned in time from the first movement of charge to/from the Capacitor DAC. Several of the standard Machine Learning Toolbox algorithms failed to converge and are not included in Fig. 29.

References

[1] M. Kütt, M. Göttsche, A. Glaser, Information barrier experimental: Toward a trusted and open-source computing platform for nuclear warhead verification, Meas.: J. Int. Meas. Confed. 114 (2018) 185–190.

[2] Peter B. Zuhoski, et al., Building a Dedicated Information Barrier System for Warhead and Sensitive Item Verification, BNL-66214, http://www.bnl.gov/isd/documents/19942.pdf;
Peter E. Vanier, et al., Study of CIVET Design of a Trusted Processor for Non-Intrusive Measurements, INMM, 2001.

[3] Jacob Benz, Paul Booker, Benjamin McDonald, Verification challenges at low numbers, in: Stephanie Spies, Sarah Weiner (Eds.), Project on Nuclear Issues: A Collection of Papers from the 2012 Conference Series, Rowman and Littlefield, Lanham, 2013, pp. 14–27;
Also see, D.J. Mitchell, K.M. Tolk, Trusted Radiation Attribute Demonstration System, SAND2000-1481C, 2000.

[4] Glen Warren, et al., Concepts for the measurement subsystems of the third generation attributes measurement system, PNNL-SA-89171, in: Paper Presented at the INMM 53rd Annual Meeting, Orlando, FL, July 10–14, 2012.

[5] L.G. Chiang, et al., Nuclear Materials Identification System Operations Manual, Oak Ridge National Laboratory, 2001, (ORNL/TM-2001/65) Rev. 2;
For FNMIS, see, T.A. Wellington, B.A. Palles, J.A. Mullens, J.T. Mihalczo, D.E. Archer, T. Thompson, C.L. Britton, N.D. Bull Ezell, M.N. Ericson, E. Farquhar, R. Lind, J. Cartera, Recent fast neutron imaging measurements with the fieldable nuclear materials identification system, Phys. Procedia 66 (2015) 432–438.

[6] David M. Chambers, et al., UK-Norway initiative: research into information barriers to allow warhead attribute verification without release of sensitive or proliferative information, in: Paper Presented at the INMM 51 Annual INMM Meeting, Baltimore, MD, July 11–15, 2010. Also see: https://ukni.info/project/information-barrier/ and references therein.

[7] Adam J. Flynn, et al., Next Generation Trusted Radiation Identification System (NG-TRIS), SAND2010-3502C, 2010. INMM Annual Meeting, July 11–16, 2010 in Baltimore, MD.

[8] P. Marleau, R. Krentz-Wee, Investigation into Practical Implementations of a Zero Knowledge Protocol, SAND2017-1649, 2017, https://prod-ng.sandia.gov/techlib-noauth/access-control.cgi/2017/171649.pdf.

[9] J. Yan, A. Glaser, Nuclear warhead verification: A review of attribute and template systems, Sci. Glob. Secur. 23 (2015) 157–170.

[10] M. Hamel, Next-generation arms-control agreements based on emerging radiation detection technologies, SAND2018-6183C, in: Institute of Nuclear Materials Management 59th Annual Meeting Baltimore, MD, USA, July 22–26, 2018, https://www.osti.gov/servlets/purl/1526354.

[11] J.R. Vavrek, B.S. Henderson, A. Danagoulian, Experimental demonstration of an isotope-sensitive warhead verification technique using nuclear resonance fluorescence, Proc. Natl. Acad. Sci. 115 (17) (2018) 4363–4368, https://www.pnas.org/content/115/17/4363.

[12] H. Russell, An improved successive-approximation register design for use in A/D converters, IEEE Trans. Circuits Syst. 25 (7) (1978) 550–554.

[13] J. Johnston, New design techniques yield low power, high resolution delta-sigma and SAR ADCS for process control, medical, seismic and battery-powered applications, in: 1991 International Conference on Analogue to Digital and Digital to Analogue Conversion, Swansea, UK, 1991, pp. 118–123.

[14] C.C. Lee, M.P. Flynn, A SAR-assisted two-stage pipeline ADC, IEEE J. Solid-State Circuits 46 (4) (2011) 859–869.

[15] C.C. Lee, M.P. Flynns, A 12b 50MS/s 3.5mW SAR assisted 2-stage pipeline ADC, in: VLSI Circ. Symp. Dig. Tech. Papers, 2010, pp. 230–239.

[16] H.-Y. Lee, et al., A 31.3fJ/conversion-step 70.4dB SNDR 30MS/s 1.2V two-step pipelined ADC in 0.13μm CMOS, in: ISSCC Dig. Tech. Papers, Feb. 2012, pp. 474–475.

[17] F. van der Goes, et al., A 1.5mW 68dB SNDR 80MS/s 2x interleaved SAR-assisted pipelined ADC in 28 nm CMOS, in: ISSCC Dig. Tech. Papers, Feb. 2014, pp. 200–201.

[18] B. Verbruggen, et al., A 70 dB SNDR 200 MS/s 2.3 mW dynamic pipelined SAR adc in 28 nm digital CMOS, in: VLSI Circ. Symp. Dig. Tech. Papers, 2014, pp. 242–243.

[19] Y. Lim, M.P. Flynn, 26.1 A 1mW 71.5dB SNDR 50MS/S 13b fully differential ring-amplifier-based SAR-assisted pipeline ADC, in: 2015 IEEE International Solid-State Circuits Conference - (ISSCC) Digest of Technical Papers, San Francisco, CA, 2015, pp. 1–3.

[20] J. Guerber, M. Gande, U. Moon, The analysis and application of redundant multistage ADC resolution improvements through PDF residue shaping, IEEE Trans. Circuits Syst. I. Regul. Pap. 59 (8) (2012) 1733–1742.