Towards Securing Low-Power Digital Circuits with Ultra-Low-Voltage Vdd Randomizers

Dina Kamel(✉), Guerric de Streel, Santos Merino Del Pozo, Kashif Nawaz, François-Xavier Standaert, Denis Flandre, and David Bol

ICTEAM/ELEN, Université catholique de Louvain, Louvain-la-neuve, Belgium
dina.kamel@uclouvain.be

Abstract. With the exploding number of connected objects and sensitive applications, security against side-channel attacks becomes critical in low-cost and low-power IoT applications. For this purpose, established mathematical countermeasures such as masking and shuffling always require a minimum amount of noise in the adversary's measurements, that may not be guaranteed by default because of good measurement setups and powerful signal processing. In this paper, we propose to improve the protection of sensitive digital circuits by operating them at a random ultra-low voltage (ULV) supplied by a V_{dd} randomizer. As the V_{dd} randomization modulates the switching current, it results in a multiplicative noise on both the current consumption amplitude and its time dependence. As ULV operation increases the sensitivity of the current on the supply voltage, it magnifies the generated noise while reducing the side-channel information signal thanks to the switching current reduction. As a proof-of-concept, we prototyped a simple V_{dd} randomizer based on a low-quiescent-current linear regulator with a digitally-controlled resistive feedback divider on which we apply a 4-bit random number stream. Using an information theoretic metric, the measurement results obtained in 65 nm low-power CMOS confirm that such randomizers can significantly improve the security of cryptographic implementations against standard side-channel attacks in case of low physical noise in the attacks' setups, hence enabling the use of mathematical countermeasures.

1 Introduction

With the increasing current trend of deploying billions of wireless Internet-of-Things (IoT) nodes, privacy and security concerns are raised [25]. However, due to strong power and area constraints, deploying cryptography for IoT systems is extremely challenging. Moreover, guaranteeing the physical security of these highly resource constrained applications against side-channel attacks is even more challenging. In a side-channel attack, the adversary exploits a physical signal (e.g. the supply current or electromagnetic field) to identify the secret key. Therefore, existing hardware countermeasures generally aim at reducing the side-channel signal-to-noise-ratio (SNR) [9] by decreasing the signal (e.g. [12,21,22]) or increasing the noise (e.g. [26]).

© Springer International Publishing AG 2016
C. Carlet et al. (Eds.): SPACE 2016, LNCS 10076, pp. 233–248, 2016.
DOI: 10.1007/978-3-319-49445-6_13

Following, the reduction of the SNR can be combined with mathematical countermeasures such as masking [18] and shuffling [24]. Nevertheless, for such mathematical countermeasures against side-channel attacks to be effective, it is strictly necessary that the original signal is hidden by a sufficient physical noise, i.e. that the original SNR is sufficiently small. Intuitively, this is because mathematical countermeasures can only amplify the impact of the physical noise (and therefore fall short if there is nothing to amplify). The usual approach for this purpose, that embeds sources of additive noise (algorithmic noise) in circuits [8], has an approximate cost that is linear with the noise level i.e. doubling the circuit size roughly doubles the noise variance, but also doubles the power consumption. A complementary approach would be to reduce the side-channel signal amplitude, for example by equalizing the power consumed with the design of custom logic gates (e.g. dual-rail pre-charged logic [21]). However, the power/area is approximately doubled and the design complexity is relatively high, which renders them unsuitable for resource-constrained applications. Other approaches to reduce the side-channel signal are the use of on-chip decoupling capacitors (current filtering) [12] and current equalization through switched capacitors [22].

This state-of-the-art raises new challenges regarding the design of advanced solutions that can be combined with mathematical countermeasures to increase the security regardless of the adversary's capabilities while keeping the cost and performance overheads limited. Therefore, we propose an ULV multiplicative source of noise in the form of a V_{dd} randomizer that embeds adequate noise due to the supply randomizataion in case the physical noise in the attack setup is insufficient, and at the same time reduces the side-channel signal due to its ULV operation ($< 0.55\,\mathrm{V}$). Our investigations show that this solution can be conveniently combined with mathematical countermeasures thanks to its low area ($1.8\times$) and low current consumption ($1.6\times$) overheads, and leads to a security improvement by a factor of 20 in case of low physical noise. Our results are based on real measurements of a fabricated chip using 65 nm low-power CMOS technology.

The rest of the paper is organized as follows. The related work and alternative approaches are discussed in Sect. 2, together with our contributions. We describe the V_{dd} randomizer implementation and the test setup in Sect. 3. We introduce our methodology for security evaluations in Sect. 4. This methodology is applied in Sect. 5, which details the security analysis of our side-channel signal reduction approach, and the embedding of the V_{dd} randomizer as a source of multiplicative noise. Finally, the design overheads are discussed in Sect. 6.

2 Related-Work and Contributions

2.1 Related Work

Countermeasures based on chip voltage regulation (VR) are currently gaining attention as they tend to pack the complexity into the regulator, which is a block present in almost all modern ICs, therefore reducing the power/energy and area

overheads. In this trend, we see two directions. One is to de-correlate the current traces (at the supply of the voltage regulator) from the load current (drawn by the crypto engine), thus reducing the side-channel signal amplitude. A lot of research focused on switched-capacitor VR (e.g. [20]) or switched-inductor VR (e.g. [7]). In [20], the authors describe a bi-channel power structure voltage regulator composed of a linear converter supplying the slowly varying part of the load current and a switched-capacitor converter clocked by a random digital signal. They also use a low dropout (LDO) regulation subsystem to further enhance the de-correlation of the external current traces from the load current. Their design occupies an area of $0.8\,\text{mm}^2$ using a $0.18\,\mu\text{m}$ CMOS technology. However, their security evaluation only relies on visible inspection of the current traces to show the effectiveness of the hiding performed by the system, which makes it difficult to assess in front of advanced side-channel attacks. Also, the authors do not use a real circuit load in their power/energy and area efficiency evaluations, which prevents quantifying the cost of the system relative to such loads. In [7], the authors use switched-inductors to de-correlate the observed current from the load current traces. But both security and performance evaluations are limited for similar reasons. Low-dropout (LDO) regulators are also investigated in [16] to attenuate the high frequency variations of the load current by reducing the LDO bandwidth. Here the authors mounted a correlation power analysis (CPA) attack against a protected AES engine which shows an 800× improvement in terms of the measurement to disclosure (MTD) metric (compared to an unprotected one) with 1.4 % area and 5 % power overheads. Yet, a possible shortcoming of their analysis is that the results they provided are based on simulation results without physical noise.

The second direction depends on creating multiple randomized observed current traces, thus introducing another source of noise (e.g. using multiphase switched-capacitor VR [26] or through random voltage scaling [1]). In [26], the authors propose to scramble the observed current by turning on and off the individual interleaved stages in a pseudo-random fashion. However, they do not provide the power/energy or area costs of their proposal and their security evaluation is uniquely based on power trace entropy, without indicating the amount of physical noise present in their evaluation system. The random voltage scaling technique proposed in [1] could lower the correlation coefficient by 10× when applied to the complete AES. Their approach aimed at FPGA designs and the authors suggested to alter the supply voltage once every 200 encryption rounds since in their settings, a successful DPA attack can be mounted against an AES engine after 2500 rounds and they need a changing supply rate much less than that.

So overall, and while these previous solutions are intuitively appealing and technically innovative, a more formal/comparative treatment of their pros and cons is still missing. As detailed next, this paper aims to make one step in this direction, by investigating the security of a current randomizer based on advanced side-channel security metrics together with its performance results in a comprehensive manner. Furthermore, since we believe the impact of such

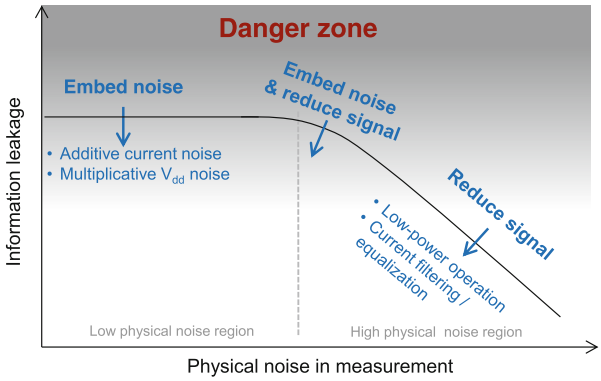

Fig. 1. Hardware countermeasures against side-channel analysis.

countermeasures are highly dependent on the actual level of physical noise found in the measurements, we propose to highlight trends in this respect, by considering noise as a parameter of our evaluations.

2.2 Contributions

First, we summarize the impact of these different countermeasures based on our understanding of their relevance in different physical noise regions as represented in Fig. 1. Generally, the side-channel information extracted from a circuit (precisely defined next) remains unchanged for low physical noise levels, and starts to decrease when increasing the physical noise. The gradient grey along the y-axis indicates the reduction of side-channel information as the color lightens. Therefore, the dark grey region is considered dangerous, since it corresponds to the case where mathematical countermeasures are ineffective. As clear from the figure, this typically happens in the low physical noise region. Fig. 1 also highlights the two main solutions for this purpose, namely embedding noise and reducing the side-channel signal. At the extreme, for extremely low physical noises, it is clear that the noise embedding approach is necessary. Similarly, in the large physical noise region, it is usually the signal reduction that brings the best benefits (since, e.g. additional additive noise could be small in front of the existing physical noise). Of course, most existing embedded devices fall inbetween these extremes and in this case, both approaches can be relevant. In the following, we pick up on this SNR reduction problem and investigate a new area- and power-efficient technique to generate hard-to-exploit noise that enables mathematical types of countermeasures. More precisely, our main contributions are threefold:

– We propose an ULV V_{dd} randomizer that modulates the switching current of a cryptographic implementation resulting in a multiplicative noise source. The V_{dd} randomizer employs an LDO regulator operating with sufficiently short

transition times between V_{dd} levels in order to prevent an adversary from easily profiling the V_{dd}'s at which each operation is performed. In addition, we propose to combine this V_{dd} randomization with an operation of the circuit-to-protect at ultra-low voltage (ULV) which reduces the side-channel signal amplitude. Both techniques help reduce the information leakage of the circuit-to-protect in low and high physical noise regions.
- We provide a security assessment of the proposed technique using an information theoretic metric described in [14,17] based on template attacks [3]. This allows us to demonstrate the effectiveness of the technique across the whole range of physical noise. In addition, the information theoretic metric we used is proven to be directly proportional to the success rate of a maximum likelihood adversary [5], thus justifying its use in this context.
- We show that our solution can be conveniently combined with mathematical countermeasures thanks to its low area ($\sim 1.8\times$) and low current consumption ($< 1.6\times$) overheads in addition to the security improvement by a factor of 20 in case of low physical noise against a standard template adversary doing Gaussian profiling. Our results are based on real measurements of a fabricated chip using 65 nm low-power CMOS technology.

Eventually, we conclude by discussing the limitations of our randomizer against adversaries able to access the V_{dd} levels during profiling. We highlight that increasing the number of V_{dd} levels (limited to 16 in our work) would be necessary to prevent such worst-case attacks, which we leave as an interesting scope for further research.

3 Vdd Randomizer Design

3.1 Circuit Implementation

Intuitively, randomizing the supply voltage of the circuit to protect modulates the switching current or I_{on} in a multiplicative fashion. When the gates are switching at nominal V_{dd} the transistors mostly operate in saturation regime:

$$I_{on} \sim (V_{dd} - V_t)^{\alpha}. \tag{1}$$

where V_t is the transistor threshold voltage and α is a factor between 1 and 2 [15].

When the transistors operate in the subthreshold regime:

$$I_{on} \sim 10(V_{dd}/S). \tag{2}$$

where S is the subthreshold swing between 60 and 100 mV/decade [2]. At the supply voltage operating range of the V_{dd} randomizer in 65 nm LP CMOS, the transistors are in the near-threshold regime, which results in an I_{on} dependence on V_{dd} between Eqs. 1 and 2. Therefore, adding voltage noise on the supply voltage results in a multiplicative noise on the dynamic supply current as I_{on} is modulated with a dependence between linear-to-quadratic in saturation regime and exponential in subthreshold regime.

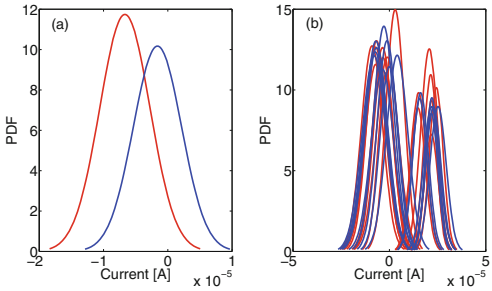

Fig. 2. Probability density functions of two (S-box) computations, for a prototype circuit operating under (a) a fixed supply (0.5 V) and (b) sixteen randomized supplies (from 0.45 V to 0.55 V).

Now, the main goal of a multiplicative noise source is to avoid the simple Gaussian leakage functions of unprotected implementations that easily allow distinguishing different events (e.g. S-box computations) happening in a target chip, as represented in Fig. 2(a)[1]. Instead, the V_{dd} randomizer turns the simple Gaussian leakage into a Gaussian mixture where every mode of the distribution represents one possible supply voltage, as represented in Fig. 2(b), thus increasing the overlap between these two events.

In this work, we use a simple V_{dd} randomizer to evaluate the security of the proposed approach. Its architecture is based on a conventional linear regulator with an error amplifier driving a power stage and a digitally-controlled 4-bit feedback resistive divider on which we apply a 4-bit random number stream generated off chip[2], as shown in Fig. 3. The error amplifier is a folded cascode amplifier and the power stage is an NMOS device with body connected to source [4]. The amplifier and its bias dissipates only 280 nA, whereas the feedback resistor network consumes about 1 µA, allowing the V_{dd} randomizer to limit the power overhead compared to the protected circuit. The feedback resistive divider is designed to provide a supply voltage to the digital circuit to protect ranging from 0.45 V to 0.55 V. This voltage range is carefully chosen in order to reduce the side-channel signal amplitude while operating the digital circuits at a maximum frequency of 1 MHz, suitable for low-speed IoT applications, and at the same time provide sufficient variations in the current traces as can be seen in Fig. 2. Current state of the art design of IoT sensor nodes uses ULV operation to minimize the energy consumption, e.g. in [19]. Therefore our analysis focuses only on the ULV region of operation which ranges from 0.45 V to 0.55 V in this case.

[1] The sign of the current is not preserved due to the clock coupling on the printed circuit board (PCB).

[2] We considered an off-chip implementation of the 4-bit random number stream generation as a proof of concept for demonstration purpose only. Of course, a full implementation of the V_{dd} randomizer would consider designing the random number stream on-chip to deny the adversary access.

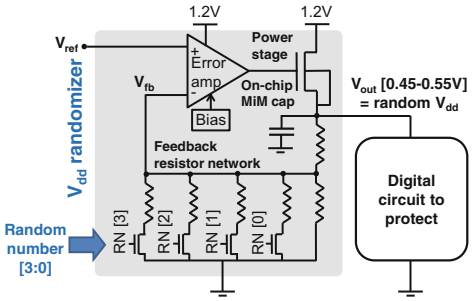

Fig. 3. Circuit architecture.

Notably, the randomness in the output voltage should come from the random number stream and not from the impact of the load current, which is correlated to the computation and might therefore leak side-channel information. Stability and load regulation can thus not be compromised in the Vdd randomizer, just as in conventional voltage regulators. In order to ensure stability over a wide range of loading currents, a current-mode capacitance multiplication in the bias of the error amplifier is used for pole splitting and an on-chip filtering MiM capacitor is added on the supply voltage output. As a result, stability is ensured for load currents up to 0.5 mA without an off-chip capacitor.

The slew rate of the randomizer is determined by the sizing of the power stage, and is in a direct trade-off with the area through the sizing of the stabilization capacitances. The slew rate specification is that the regulator has to render a transition time comparable with the clock period used in the digital circuit. This ensures sufficiently short transition times between V_{dd} levels in order to prevent an adversary from easily profiling the V_{dd}'s at which each operation is performed. In [1], which aimed at FPGA designs, the supply voltage is modified at a rate of one change per 200 encryptions. While this may be sufficient to improve security against certain types of Differential Power Analyses (DPA) attacks, it is still insufficient against advanced adversaries exploiting multiple samples/intermediate computations per encryption [10,23]. Therefore we target a slew rate of 10^5 V/s allowing the voltage to ramp from 0.45 V to 0.55 V in 1 µs compatible with the 1 MHz clock frequency.

3.2 Performance Benchmark and Test Setup

In order to characterize the V_{dd} randomizer, we designed a test chip implementing an 8-bit AES S-box as benchmark for the circuit to protect[3]. The input signal

[3] In our test chip we only implemented the 8-bit AES S-box instead of the whole AES as a proof of concept. Of course when used with the full AES, the V_{dd} randomizer should be able to drive the whole AES circuit. The results we later provide in Sect. 6 are for the measured V_{dd} randomizer with the AES S-box and also an estimation in case the randomizer operates with the full AES.

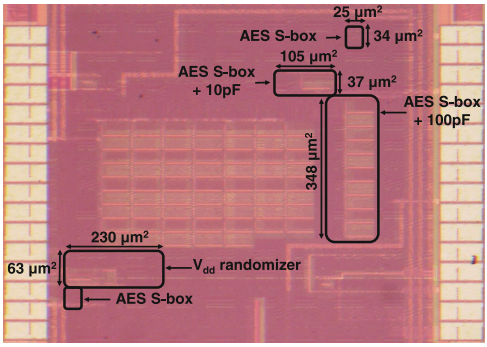

Fig. 4. Die microphotograph of the ultra-low-voltage V_{dd} randomizer and other structures of the AES S-box with decoupling capacitors.

has 256 possible values whose transitions are chosen between 0 and an arbitrary input.[4] In order to hide the impact of the supply voltage transients on the captured traces, the random number stream is clocked synchronously with the S-box input signal, i.e. the V_{dd} randomizer is synchronized with the operation of the circuit to protect. The V_{dd} randomizer was manufactured in 65 nm LP CMOS, and its area is $0.0145\,\mathrm{mm}^2$, as represented in Fig. 4. Our dies also contain several versions of an unprotected S-box with various levels of decoupling capacitances for comparison. All input signals were generated externally using a National instrument PXI 6552 waveform generator. The clock frequency of all circuits under test is 1 MHz. Current traces for security analysis were captured with a differential probe over a resistor with a 2 GS/s oscilloscope.

4 Methodology

4.1 Evaluation Settings

Capital letters are assigned to random variables, while lower case letters refer to samples of these random variables. The leakage function in case the implementation uses a fixed supply voltage has two input arguments: the discrete random variable X, which denotes the value of the processed data under investigation, and the continuous random variable N, which represents the physical noise in the measurements. When the V_{dd} randomizer is used, we also consider the discrete random variable V, which denotes the supply voltage. The leakage function variable denoted by $L(\,,\,)$ contains either random variable arguments or fixed arguments. We denote the t^{th} time sample in a leakage trace as $L_t(\,,\,)$. We consider

[4] We only considered 256 input transitions for the S-box in order to limit the time of our measurement campaigns. Since our security evaluation will essentially reflect the improved overlap of Gaussian mixture models such as in Fig. 2, this should not impact our comparisons between fixed and randomized power supplies. Yet, a more expensive profiling of 256^2 transitions should admittedly allow adversaries to extract slightly more information from their traces.

two types of traces in our analysis. First, the real measurements with actual physical noise are denoted as $L_t^1(X,N) = L_t^{meas}(X,N)$. Second, "hybrid" traces, in which the average measurement traces $\overline{L_t^{meas}}(X) = \hat{\mathbf{E}}_n L(X,n)$ (where $\hat{\mathbf{E}}$ denotes the sample mean operator) are combined with simulated Gaussian noise. The leakage function in this context is denoted as $L_t^2(X,N) = \overline{L_t^{meas}}(X) + N$. These hybrid traces allow us to quantify the impact of a change of physical noise level in our different experiments.

4.2 Information Theoretic Metric

We evaluate the leakage information of the traces with the information theoretic metric described in [17] and refined in [14]. Namely, the Perceived Information (PI) corresponds to the amount of information that can be exploited by a side-channel adversary given a certain leakage model:

$$\hat{\mathrm{PI}}(X;L) = \mathrm{H}[X] - \sum_{x \in X} \Pr[x] \sum_{l \in L} \Pr_{chip}[l|x] \; . \; \log_2 \left(\hat{\Pr}_{model}[x|l] \right).$$

In case the true (unknown) leakage distribution of an implementation (denoted as $\Pr_{chip}[l|x]$) and the adversary's leakage model estimate (given by $\hat{\Pr}_{model}[x|l]$) are identical (e.g. in a simulated environment), then a perfect evaluation is achieved. That is, the PI is equivalent to the standard definition of mutual information and it captures the worst-case information leakages. By contrast, if these distributions deviate (because of practical limitations which lead to bad profiling, or because there exists significant inter-chip variability, or because the adversary's model is simplified), then the PI is the best available estimate of the implementation's leakage. Compared to the previous analyses in [1], using such an information theoretic metric allows our conclusions to be closer to those of a worst-case security evaluation. Indeed, such a PI metric is directly proportional to the success rate of a maximum likelihood adversary (as proven in [5]).[5]

Note that the PI can be viewed as a generalization of the SNR metric discussed in introduction [5]. It is even proportional to the SNR in case of Gaussian leakages. We next use the PI (rather than the SNR) as evaluation metric since it can capture other types of leakage distributions, in particular the Gaussian mixtures that are relevant in our experiments.

4.3 Information Extraction Tools

In order to evaluate the previous information theoretic metric, one essentially requires a good model, aka estimation of the leakage probability function. For this purpose, our strategy will follow the one already established, e.g. in [13,14], and consider a univariate setting as a starting point. That is, models will be built

[5] If positive, otherwise it indicates that the model exploited by the adversary does not guarantees successful key recoveries.

exhaustively for all the time samples of our leakage traces, and the PI value for the most informative time sample will be kept.[6] Concretely, building models for the fixed supply voltage case can directly exploit the Gaussian template attacks described in [3]. That is, in this case we start by building 256 templates of the form:

$$\hat{\Pr}_{model}[l|x] = \mathcal{N}(l|\mu_{x,N}, \sigma^2_{x,N}). \tag{3}$$

$\hat{\Pr}_{model}[x|l]$ is then obtained by applying Bayes' rule. Eventually, the PI metric is directly estimated according to its equation, by sampling the true distribution $\Pr_{chip}[l|x]$ (i.e. by measuring the chip) and estimating the conditional probabilities of the 256 x values based on these measurements.

By contrast, the procedure can be slightly more involved in the case of randomized power supplies. We will consider two types of adversaries for this purpose: a standard one and powerful one. In the first case, the adversary is assumed incapable of identifying the 16 V_{dd} values during profiling. Therefore, the power supply randomizations are (wrongly) considered as a part of the measurement physical noise when building the templates and estimating the PI. In practice, such a setting would typically correspond to a context where the random numbers are unknown during profiling. As a result, the profiling phase exactly corresponds to the previous Gaussian templates building, but with $\sigma'^2_{x,N}$ made of a truly physical part $\sigma^2_{x,N_{meas}}$ to which we add a randomization part $\sigma^2_{x,N_{Vdd}}$. We call this scenario *Gaussian profiling*.

Next, the more powerful adversary is assumed capable of identifying the 16 random supply voltages during profiling. In this case we build 256×16 templates corresponding to the 256 S-box inputs and the 16 supply voltages:

$$\hat{\Pr}_{model}[x|l, v] = \mathcal{N}(l|\mu_{x,v,N}, \sigma^2_{x,v,N}). \tag{4}$$

Quite naturally, the random numbers selecting V_{dd} remain unknown during the PI estimation phase:

$$\hat{\mathrm{PI}}(X; L) = \mathrm{H}[X] - \sum_{x \in X} \Pr[x] \sum_{v \in V} \Pr[v] \cdot \sum_{l \in L} \Pr_{chip}[l|x, v] \,.\, \log_2(\hat{\Pr}_{model}[x|l]),$$

where the conditional probability of the events x given the leakages l is computed by summing over all possible v's, namely: $\hat{\Pr}_{model}[x|l] = \sum_{v \in V} \hat{\Pr}_{model}[x|l, v]$. In the following, we call this scenario *Gaussian mixture profiling*.

Note that the more powerful adversary could additionally target the leakage of the 4-bit random values controlling the randomizer. This is an interesting scope for further research. Yet, as the following results already show that our

[6] Extending this analysis towards multivariate attacks, possibly including a dimensionality reduction phase, is an interesting scope for further research. As for Footnote 1, it should not impact our comparisons between fixed and randomized supplies, but allow more efficient attacks.

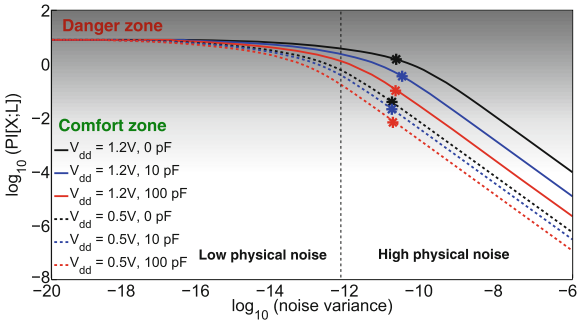

Fig. 5. Perceived information of the AES S-box with different decoupling capacitor values at a supply voltage of 1.2 V (solid lines) and 0.5 V (dashed lines). Curves correspond to the hybrid case with simulated Gaussian noise. The stars indicate the actual physical noise measured on chip.

instance of randomizer is not sufficient to prevent such powerful adversaries without this additional leakage, results in this direction will not affect our conclusions.

5 Security Analysis

5.1 ULV Operation and Decoupling Capacitors

In our analysis, the AES S-box is first operated at two different constant supply voltages: the nominal 1.2 V and a ULV (near-threshold) supply which is 0.5V. In each case, the impact of adding different values of on-chip decoupling capacitors is explored as well. Figure 5 exploits both the actual measured traces ($L_t^1(.)$) denoted by the stars and the hybrid traces ($L_t^2(.)$) explained in Sect. 4.1. It demonstrates how the reduction of the supply voltage and the addition of on-chip decoupling capacitors are effective in case the physical noise in the attack setup is high enough. Both techniques reduce the side-channel signal. This is clearly seen as the stars' horizontal positions in Fig. 5 remain nearly the same (corresponding to physical noise in the attack setup), whereas their vertical positions (corresponding to the perceived information) decreases while operating at ULV or using on-chip decoupling capacitors. However, if the physical noise can be reduced (e.g. thanks to a better measurement setup, or signal processing) as in the left part of the figure, neither lowering the supply voltage, nor using on-chip decoupling capacitors can help to escape the danger zone.

5.2 Vdd Randomizer

In Fig. 6 we compare the security of the V_{dd} randomizer implementation to the unprotected S-box at 0.5 V (without decoupling capacitors) again exploiting both the actual measured traces ($L_t^1(.)$) denoted by the stars and the hybrid traces

($L_t^2(.)$) explained in Sect. 4.1. Furthermore, we have considered different settings for the actual measured traces to explore various physical noise values. The bandwidth of the oscilloscope was configured to full (600 MHz) and to 20 MHz in addition to using the singular spectrum analysis (SSA) post-processing tool introduced in [11] to reduce the physical noise in the attack setup[7]. First we consider Gaussian profiling that correspond to a standard adversary who (wrongly) considers the power supply randomizations as a part of the measurement physical noise. In the low physical noise region, the perceived information of the V_{dd} randomizer, using Gaussian profiling, is 20× better than the unprotected S-box at 0.5 V, thus approaching the comfort zone. Note that once in the comfort zone, typically corresponding to a PI below 0.1, a factor 20 for the PI reduction implies a multiplication of the attack's data complexity by the same factor in case of unprotected devices, and a factor 20^d if masking with d shares is exploited [5]. Meanwhile, the security of the V_{dd} randomizer is bounded by the unprotected S-box at 0.5 V in the high physical noise region which naturally lies in the comfort zone. This is expected as the physical noise dominates in this region. Consequently, these results prove the importance of combining the V_{dd} randomization technique with the ultra-low voltage operation to sustain sufficient security for the whole physical-noise range.

We insist that the concrete noise level of our experiments is in general less relevant than the trends indicated by our PI curves. In particular, since we target a combinatorial circuit, the SNR of these measurements is lower than what would be expected for sequential circuits and complete systems. Besides, it is interesting to see that the physical noise in the attack setup can be reduced by lowering the bandwidth of the oscilloscope to 20 MHz (acting as a low-pass filter) and by employing the SSA tool as shown by the symbols in Fig. 6. In general, the goal of the V_{dd} randomizer is indeed to mitigate the risk of a strong physical noise reduction.

On the other hand, if the Gaussian mixtures are considered, where the adversary is assumed capable of accurately identifying the 16 random supply voltages used, then there is obviously no security gain compared to the unprotected S-box at $0.5\,V$ in the low physical noise region. This is expected, since for the V_{dd} randomizer to be effective in this context, we need a noise such that the modes of the distributions in Fig. 2 start to overlap. In this respect, it is important to stress that this observation does not invalidate the interest of the randomizer. First, and very concretely, such a powerful profiling may be difficult to be performed by practical adversaries, since the internal randomness of the randomizer is not supposed to leave the chip. Yet, it is an interesting conceptual challenge to prevent even those adversaries (and their possible extension towards non-parametric pdf estimation techniques that would not require the knowledge of the masks during profiling, at the cost of higher sampling requirements). Second, and more importantly, V_{dd} randomizers can in principle enforce the modes of their Gaussian mixtures to be arbitrarily close, by increasing the range of

[7] SSA can be viewed as a type of filtering. Details are not necessary for the understanding of our results.

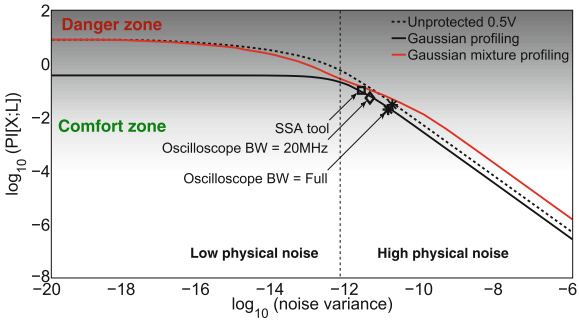

Fig. 6. Perceived information of the unprotected AES S-box at 0.5 V and the one with variable supply voltages using the Gaussian profiling and the perfect profiling scenarios. The symbols indicate the actual physical noise on chip with different settings.

the power supplies. Hence, our results show that our simple V_{dd} randomizer is already a good solution to prevent most state-of-the-art side-channel attacks, and that their generalization towards a wider range of V_{dd} levels to face even more powerful adversaries is an interesting research track. Note that by "most state-of-the-art attacks" we mean in particular all the CPA-like attacks that were used to assess the security of the solutions mentioned in Sect. 2.

More technically, it is worth mentioning that in the Gaussian mixture profiling we notice the "waved" shape of the information theoretic curve for the intermediate noise levels that is typical from masking [18]. It indicates that several moments of the statistical distribution are actually exploited for such noise levels. Besides, for the worst-case Gaussian mixture profiling, the V_{dd} randomizer actually leaks (slightly) more information than the unprotected chip running at 0.5 V in the high noise region. This is explained by the fact that the randomized supplies also lead to computations at (more informative) higher supplies in this case.

6 Cost Comparison

Table 1 summarizes the costs of the techniques in this paper. First, reducing the supply voltage of the unprotected S-box from 1.2 V to 0.5 V decreases both the current consumption and the PI at actual measured physical noise by 2.3× and 34×, respectively (for the standard side-channel adversary doing Gaussian profiling). Next, when the V_{dd} randomizer is used with the S-box, we gain a factor of 20 in PI at low physical noise for a similar increase of 17× in area, while maintaining the security gain of nearly 50× at the actual measured physical noise. This is more or less what additive noise would cost. But quite naturally, the performance gains are significantly amplified if the V_{dd} randomizer was used for a full AES design, since we could then amortize its cost (e.g. the area is expected to increase only by a factor of 1.8 compared to the unprotected AES reported in [6], still leading to the same security gain 20× at low physical noise).

The current consumption overheads in this case are even smaller: the full AES with the V_{dd} randomizer would consume $< 1.6\times$ higher current than the unprotected one. Finally, decoupling capacitances are only effective in the high physical noise region at a large area cost.

Table 1. Security versus cost (area and current consumption at $1\,MHz$) for a standard adversary.

Implementation	Area [GE]	Current [μA]	PI @ low noise	PI @ actual noise
S-box (1.2 V)	220	0.74	8	1
S-box (0.5 V)	220	0.32	8	0.029
Full AES[a]	4,721	2.12	8	NA
S-box + V_{dd} rand	3,753	1.64	0.36	0.019
Full AES + V_{dd} rand	8,255	3.48	0.36	NA
S-box + 10 pF (0.5 V)	1,011	0.32	8	0.021
S-box + 100 pF (0.5 V)	6,849	0.32	8	0.007

[a] The power consumption of the unprotected full AES reported in [6] is at 0.4 V (890 kHz).

7 Conclusions

Noise is always assumed as the basic ingredient to prevent side-channel attacks. Confirming previous works in this direction, this paper shows that designing secure and efficient noise engines is not a trivial task, and certainly deserves more attention. In particular, while trying to hide the side-channel signal in a sufficient amount of physical noise with signal reduction techniques (as done with decaps in this paper) or mathematical countermeasures is well understood, how to generate hard-to-exploit noise in the low physical noise region is very challenging, especially in front of powerful adversaries able to perform Gaussian mixture profiling.

As a first step towards the better understanding of these issues, we analyzed the security improvements offered by a V_{dd} randomizer prototype to supply the digital circuits to protect at ULV. It shows good results against standard DPA adversaries usually considered in the literature (and evaluation laboratories), at a low die area cost. This confirms that randomizing the supplies can be used to make sure that the (possibly small) physical noise in an adversary's attack setup creates confusion when trying to distinguish cryptographic computations. Mathematical countermeasures such as masking can then be used to amplify this confusion.

But interestingly, our results also show that the impact of such randomizers may be limited in front of powerful adversaries able to profile the leakage distributions with full access to the chip's randomness (which is not advisable from

a design point-of-view, but is interesting to reflect worst-case security levels). Our discussion (in Sect. 5.2) suggests that preventing such powerful adversaries is conceptually feasible, e.g. with supplies covering a wider range of V_{dd} levels, with a more granular randomization. So, our results raise new research challenges. Namely, how to design efficient noise engines that guarantee low information leakage (in the comfort zone) across the whole range of physical noise and against adversaries exploiting non-Gaussian profiling methods (either Gaussian mixtures, as in this paper, or non-parametric ones).

Acknowledgements. This work has been funded in parts by the ARC Project NANOSEC. François-Xavier Standaert is a research associate of the Belgian Fund for Scientific Research.

References

1. Baddam, K., Zwolinski, M.: Evaluation of dynamic voltage and frequency scaling as a differential power analysis countermeasure. In: VLSI Design, pp. 854–862. IEEE (2007)
2. Bol, D., Ambroise, R., Flandre, D., Legat, J.-D.: Interests and limitations of technology scaling for subthreshold logic. IEEE Trans. VLSI Syst. **17**(10), 1508–1519 (2009)
3. Chari, S., Rao, J.R., Rohatgi, P.: Template attacks. In: Kaliski, B.S., Koç, K., Paar, C. (eds.) CHES 2002. LNCS, vol. 2523, pp. 13–28. Springer, Heidelberg (2003). doi:10.1007/3-540-36400-5_3
4. de Streel, G., De Vos, J., Flandre, D., Bol, D.: A 65 nm 1 V to 0.5 V linear regulator with ultra low quiescent current for mixed-signal ULV SoCs. In: FTFC, pp. 1–4. IEEE (2014)
5. Duc, A., Faust, S., Standaert, F.-X.: Making masking security proofs concrete. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015. LNCS, vol. 9056, pp. 401–429. Springer, Heidelberg (2015). doi:10.1007/978-3-662-46800-5_16
6. Hocquet, C., Kamel, D., Regazzoni, F., Legat, J.-D., Flandre, D., Bol, D., Standaert, F.-X.: Harvesting the potential of nano-CMOS for lightweight cryptography: an ultra-low-voltage 65 nm AES coprocessor for passive RFID tags. J. Cryptographic Eng. **1**(1), 79–86 (2011)
7. Kar, M., Lie, D., Wolf, M., De, V., Mukhopadhyay, S.: Impact of inductive integrated voltage regulator on the power attack vulnerability of encryption engines: a simulation study. In: CICC, pp. 1–4. IEEE (2014)
8. Mangard, S.: Hardware Countermeasures against DPA – a statistical analysis of their effectiveness. In: Okamoto, T. (ed.) CT-RSA 2004. LNCS, vol. 2964, pp. 222–235. Springer, Heidelberg (2004). doi:10.1007/978-3-540-24660-2_18
9. Mangard, S., Oswald, E., Popp, T.: Power Analysis Attacks: Revealing the Secrets of Smart Cards. Advances in Information Security. Springer, Secaucus (2007)
10. Mather, L., Oswald, E., Whitnall, C.: Multi-target DPA Attacks: pushing DPA beyond the limits of a desktop computer. In: Sarkar, P., Iwata, T. (eds.) ASIACRYPT 2014. LNCS, vol. 8873, pp. 243–261. Springer, Heidelberg (2014). doi:10.1007/978-3-662-45611-8_13

11. Merino Del Pozo, S., Standaert, F.-X.: Blind source separation from single measurements using singular spectrum analysis. In: Güneysu, T., Handschuh, H. (eds.) CHES 2015. LNCS, vol. 9293, pp. 42–59. Springer, Heidelberg (2015). doi:10.1007/978-3-662-48324-4_3
12. Nakai, T., Shiozaki, M., Kubota, T., Fujino, T.: Evaluation of on-chip decoupling capacitors effect on AES cryptographic circuit. In: SASIMI (2013)
13. Renauld, M., Kamel, D., Standaert, F.-X., Flandre, D.: Information theoretic and security analysis of a 65-nanometer DDSLL AES S-Box. In: CHES, pp. 223–239 (2011)
14. Renauld, M., Standaert, F.-X., Veyrat-Charvillon, N., Kamel, D., Flandre, D.: A formal study of power variability issues and side-channel attacks for nanoscale devices. In: Paterson, K.G. (ed.) EUROCRYPT 2011. LNCS, vol. 6632, pp. 109–128. Springer, Heidelberg (2011). doi:10.1007/978-3-642-20465-4_8
15. Sakurai, T., Newton, A.R.: Alpha-power law MOSFET model and its applications to CMOS inverter delay and other formulas. IEEE J. Solid-State Circuits **25**(2), 584–594 (1990)
16. Singh, A., Kar, M., Ko, J. H., Mukhopadhyay, S.: Exploring power attack protection of resource constrained encryption engines using integrated low-drop-out regulators. In: ISLPED, pp. 134–139. IEEE/ACM (2015)
17. Standaert, F.-X., Malkin, T.G., Yung, M.: A unified framework for the analysis of side-channel key recovery attacks. In: Joux, A. (ed.) EUROCRYPT 2009. LNCS, vol. 5479, pp. 443–461. Springer, Heidelberg (2009). doi:10.1007/978-3-642-01001-9_26
18. Standaert, F.-X., Veyrat-Charvillon, N., Oswald, E., Gierlichs, B., Medwed, M., Kasper, M., Mangard, S.: The World Is Not Enough: another look on second-order DPA. In: Abe, M. (ed.) ASIACRYPT 2010. LNCS, vol. 6477, pp. 112–129. Springer, Heidelberg (2010). doi:10.1007/978-3-642-17373-8_7
19. Takamiya, M.: Energy efficient design and energy harvesting for energy autonomous systems. In: VLSI Design, Automation and Test, VLSI-DAT 2015, Hsinchu, Taiwan, 27–29 April 2015, pp. 1–3 (2015)
20. Telandro, V., Kussener, E., Malherbe, A., Barthelemy, H.: On-chip voltage regulator protecting against power analysis attacks. In: MWSCAS, pp. 507–511 (2006)
21. Tiri, K., Verbauwhede, I.: Securing encryption algorithms against DPA at the logic level: next generation smart card technology. In: Walter, C.D., Koç, Ç.K., Paar, C. (eds.) CHES 2003. LNCS, vol. 2779, pp. 125–136. Springer, Heidelberg (2003). doi:10.1007/978-3-540-45238-6_11
22. Tokunaga, C., Blaauw, D.: Secure AES engine with a local switched-capacitor current equalizer, In: ISSCC, pp. 64–65. IEEE (2009)
23. Veyrat-Charvillon, N., Gérard, B., Standaert, F.-X.: Soft analytical side-channel attacks. In: Sarkar, P., Iwata, T. (eds.) ASIACRYPT 2014. LNCS, vol. 8873, pp. 282–296. Springer, Heidelberg (2014). doi:10.1007/978-3-662-45611-8_15
24. Veyrat-Charvillon, N., Medwed, M., Kerckhof, S., Standaert, F.-X.: Shuffling against Side-Channel Attacks: a comprehensive study with cautionary note. In: Wang, X., Sako, K. (eds.) ASIACRYPT 2012. LNCS, vol. 7658, pp. 740–757. Springer, Heidelberg (2012). doi:10.1007/978-3-642-34961-4_44
25. Xu, T., Wendt, J. B., Potkonjak, M.: Security of IoT systems: design challenges and opportunities. In: ICCAD, pp. 417–423. IEEE/ACM (2014)
26. Weize, Y., Uzun, O.A., Köse, S.: Leveraging on-chip voltage regulators as a countermeasure against side-channel attacks. In: DAC, pp. 115:1–115:6. ACM/EDAC/IEEE (2015)