

Scaling Trends for Dual-Rail Logic Styles against Side-Channel Attacks: a Case-Study

Kashif Nawaz, Dinal Kamel, François-Xavier Standaert, Denis Flandre

ICTEAM/ELEN/Crypto Group, Université catholique de Louvain, Belgium.

Abstract. Dual-rail logic styles have been considered as possible alternatives to CMOS for the design of cryptographic circuits (more) secure against side-channel attacks. The state-of-the-art view on this approach is contrasted as they reduce the exploitable side-channel signal while not being sufficient to fully prevent the attacks. Since the limitations of dual-rail logic styles are essentially due to implementation challenges (e.g. the need of well-balanced capacitances), a natural question is to find out how they evolve with technology scaling. In this paper, we discuss this issue based on the relevant case study of an AES S-box implemented in CMOS and a dual-rail logic style, for two (65 nanometer and 28 nanometer) technologies. Our evaluations show that the security vs. performance tradeoff of our dual-rail logic style does not scale well compared to CMOS. It also shows that the scaling trends for CMOS are more positive (i.e. smaller technologies and supply voltages reduce the energy consumption and the side-channel signal). So these results suggest that dual-rail logic style may not be a sustainable approach for side-channel signal reduction as we move towards lower technology nodes.

1 Introduction

Following the first publications on power and electromagnetic analysis against cryptographic implementations, dual-rail (aka dynamic and differential) logic styles appeared as promising candidates to improve security against such attacks. Intuitively, these logic styles aim to solve the issue directly at the circuit level, by trying to reduce the side-channel Signal-to-Noise Ratio (SNR). For this purpose, they typically ensure that the switching activity of the circuits is independent of the manipulated data. However, despite constant switching activity, small data-dependent variations in the current traces can generally be observed, e.g. due to the unbalanced capacitances of the gates differential nodes and their interconnections. Therefore, a large body of work investigated the design of dual-rail logic styles in order to reach the best security vs. performance tradeoff, including but not limited to SABL [36], WDDL [35], DyCML [1], MCML [8] and MDPL [27]. Evaluations based on both simulations and actual measurements then confirmed that getting rid of these data dependencies is extremely challenging [34, 19, 26, 29]. More recent works even showed that filtering effects in concrete measurement setups make these small data dependencies reasonably easy-to-exploit, e.g. thanks to linear regression [16]. To complete the picture,

most of these dual-rail logic styles usually come with significant performance overheads, some of them additionally requiring full custom-design (which allows further control of the hardware, but is making development and deployment significantly more challenging/expensive).

In parallel, recent progresses have shown that mathematical countermeasures against side-channel analysis, and in particular the mainstream shuffling and masking techniques [6, 13, 21], can only lead to significant security improvements if the side-channel SNR has been sufficiently reduced beforehand [33, 37]. This raises the problem of finding effective (and if possible efficient) hardware techniques allowing to fulfil this condition. Intuitively, it can be done by reducing the side-channel signal, which is what dual-rail logic styles achieve, or by increasing the noise.

Eventually, since the evaluation of secure hardware technologies goes together with technology scaling, another problem is to find out which of those approaches has more potential for the future.

In this paper, we therefore tackle this question of the comparative advantage of dual-rail logic styles as a mean of reducing the side-channel signal over standard CMOS in front of technology scaling. More precisely, we investigate how much the security vs. performance tradeoff of these design styles scales, based on the simple yet reflective case-study of CMOS and Dynamic and Differential Swing-Limited Logic (DDSSL) AES S-boxes, implemented in 65 and 28 nanometer technologies. DDSSL is yet another (full-custom) dual-rail logic style which has already been analyzed based on simulations and actual measurements [30]. Our choice of DDSSL arises from the fact that its design using 65 nanometer bulk technology shows $1.5\times$ lower power consumption at the expense of $1.125\times$ increase in area, while increasing the security $10\times$ when compared to CMOS [15]. This compares positively with the typical power/energy and area costs obtained with the previously listed dual-rail logic styles. Therefore, DDSSL can be considered a good candidate to illustrate technology scaling trends (as discussed in conclusions, we expect other dual-rail logic styles to follow similar trends). In this respect, our main conclusions are twofold.

First, and looking at the tradeoff between the side-channel SNR and the implementation performances (here measured with the energy per operation, which is a quite reflective metric to compare cryptographic designs [17]), we see that the comparative advantage of DDSSL over CMOS is vanishing with technology scaling, and we explain this trend by the imbalances in DDSSL gates that gain impact with technology scaling.

Second, and more positively, we also see that technology and supply voltage scaling have a positive impact on the security vs. performance tradeoff of CMOS devices, essentially because such a scaling comes with energy gains and side-channel signal reductions.

Our case study therefore suggests that signal reduction using dual-rail logic styles may not be the best approach w.r.t technological scaling. It also suggests that the design of noisy CMOS implementations (which is a natural consequence of scaling [10]) appears as a promising strategy for ensuring a sufficiently small

side-channel SNR allowing (e.g.) masking and shuffling to be effective in future technologies. Note that by noisy implementations, we do not mean measurement noise or additional external noise but intrinsic noise at the device level (i.e. transistor, interconnect, resistor, ...) as a result of technological scaling.

Cautionary note. The results in this paper are based on simulations. While we admit that in general, they can lead to shortcomings (e.g. regarding the shape/linearity of the leakage traces), the experiments in [30] showed that they can be used as a good predictor for the amount of information leakage in dual-rail logic styles. Since the goal in this paper is to discuss general scaling trends, we believe simulations can therefore be used as an interesting indication of how the comparative advantage of CMOS over DDSLL scales. Note that anyway, we do not expect the shape/linearity of the leakage traces to be significantly different in 65 and 28 nanometer technologies, nor for CMOS vs. DDSLL, since the main linearization factor is due to filtering effects in the measurement setup and not the internal transistor behavior. So as usual with simulations, they should be interpreted with care (which we try to do in the paper). But as usual with simulations as well, they are a useful tool to get some hints about the best solutions to investigate up to (more expensive) tape outs.

The rest of the paper is structured as follows. Preliminaries are in Section 2. Our target implementation and evaluation settings are in Section 3. The comparative study between CMOS and DDSLL is in Section 4. The positive impact of technology and supply voltage scaling for CMOS implementations is in Section 5. Finally, our conclusions and a discussion of the relevance of this case study are in Section 6.

2 Preliminaries

2.1 Logic styles: CMOS and DDSLL

Traditional CMOS circuits have shown a data-dependency in the power consumption leading to exploitable side-channel information, e.g. thanks to Differential Power Analysis (DPA) [18], Correlation Power Analysis (CPA) [5] and Template attacks [7]. The power consumption of a CMOS circuit is modeled by the following equation:

$$\begin{aligned} P &= P_{dyn} + P_{stat}, \\ &= \frac{1}{2} N_{nodes} \alpha_F C_L V_{DD}^2 f_{clk} + I_{leak} V_{DD}, \end{aligned} \quad (1)$$

where P_{dyn} is the dynamic power consumption, P_{stat} the static power consumption, N_{nodes} is the number of nodes in the circuit, α_F represents the activity factor of the design, C_L is the load capacitance, V_{DD} is the supply voltage, f_{clk} represents the clock frequency and I_{leak} denotes the leakage current. We assume here that one operation is executed per clock cycle, and thus $f_{clk} = f_{op}$, which determines the target throughput of the application. The operation period T_{op}

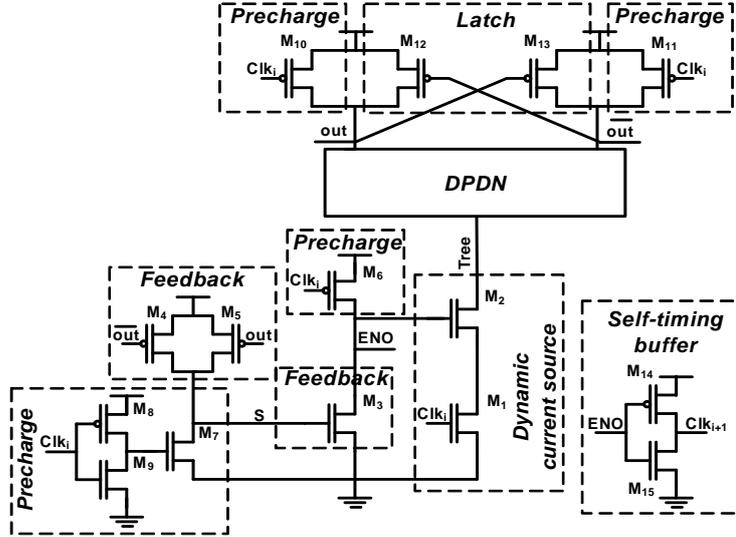


Fig. 1: Schematic of a generic DDSLL gate.

$= \frac{1}{f_{op}}$ should be more than the critical path delay T_{del} to guarantee correct functionality. The data-dependency of the CMOS logic comes from both its dynamic and static power consumptions. In the dynamic part, α_F directly depends on the data being processed. In the static part, the I_{leak} is the data-dependent parameter. Although the former is dominant, static power consumption can also be exploited if its value is sufficiently high and the operating frequency is low enough allowing the reduction of noise via simple averaging techniques [24, 28].¹ Yet, in our following experiments, dynamic power indeed dominates.

In comparison with CMOS, we investigate the Dynamic and Differential Swing-Limited Logic (DDSLL) which aims at low-power implementations and of which the dynamic power consumption is given by:

$$P_{dyn} = \frac{1}{2} N_{nodes} \alpha_F C_L V_{DD} V_{swing} f_{clk}, \quad (2)$$

where α_F equals 1 because all dynamic and differential logic styles ensure one output transition per clock cycle (independent of the data being processed), and V_{swing} is the output voltage swing. DDSLL gates are designed to have limited swing (i.e. $< V_{DD}$) in order to reduce the dynamic power consumption and hence the energy per operation.

¹ Note that advanced technologies usually provide multiple flavors such as low-power and high-performance along with different device choices such as high and low threshold voltages, providing circuit designers with various options to reduce the power consumption – and the leakage power as well – which may modify the respective importance of these source of leakages.

Figure 1 shows the circuit of a generic DDSLL gate. A Differential Pull-Down Network (DPDN) is used to evaluate the required function. It mainly consists of NMOS transistors. The DDSLL gate employs a dynamic current source to significantly reduce the static power consumption similar to what is achieved in the DyCML logic style. The cut-off of this current source is performed via a feedback network which signals the end of an operation, so that the self-timing buffer creates a clock signal Clk_{i+1} declaring the termination of the current evaluation phase. This clock signal is used to feed the following DDSLL gate. The precharge transistors are used to precharge the differential outputs of a DDSLL gate to the supply voltage before an evaluation of the gate’s function takes place, and the latch transistors preserve the evaluated voltage at the differential output nodes.

The operation of a DDSLL gate is quite simple. It works in two modes: precharge and evaluation. In the precharge mode, when the clock signal Clk_i is low, the outputs out and \overline{out} are precharged to V_{DD} . There is no current path from V_{DD} to GND because transistor M_1 is switched off. However, transistors M_6 and M_2 are switched on. Next during the evaluation phase, Clk_i goes high turning on the transistor M_1 while M_2 is still on (both forming the dynamic current source) as node ENO was previously charged to V_{DD} in the previous precharge phase, thus creating a path to discharge one of the output nodes to GND . The discharge path through the DPDN network is the one with the lowest impedance depending on the inputs being processed. As one of the outputs falls below the threshold voltage of the feed-back transistors (M_4 , M_5), one of them will turn on which in turn will discharge the node ENO to GND thus starving the current source. The design of a DDSLL circuit comprising of several functions can benefit from resource sharing (the dynamic current source, parts of the precharge circuit, the feedback circuit and the self-timing buffer of functions that evaluate at the same time), therefore reducing the area cost and the overall power consumption.

In all differential design styles, the unbalanced capacitances are considered as a source of information leakage. They mainly come from either routing imbalances or from internal imbalances. In this paper, we only consider the latter ones. (As mentioned in Section 3.3, additionally considering post-layout simulations or variability should amplify the trend we put forward). In this respect, we note that the only way to eliminate internal imbalances is to design circuits with perfectly symmetrical differential gates, which is usually very expensive in terms of area and speed. Hence, our DPDN designs exploit the binary decision diagram technique from [11] in order to improve performances by exploiting more complex gates with minimum imbalance caused by internal capacitances (details can be found in [15]). So essentially, what we show next is that (here internal) imbalances have an increasing impact with technology scaling.

2.2 Evaluation metrics

Evaluating logic styles across technology scaling and supply voltage scaling is a challenging task, since certain metrics may favour one logic style over another. In order to be as fair as possible in our performance and security evaluations,

we therefore selected generic metrics that are generally more reflective of the “global performance level” of an implementation, and can capture any type of information leakage.

More precisely, and as far as performances are concerned, the energy per operation is a quite discriminant metric, as it corresponds to an integral over time, and therefore is not “compressible” (via architectural tweaks) beyond what is allowed by the total combinatorial cost of an implementation [17]. Concretely, the energy consumption of a logic circuit can be calculated by integrating the power consumption over the time required for the target operation (in our case-study, the AES S-box):

$$\begin{aligned}
 E_{op} &= \int_t (P_{dyn} + P_{stat}) dt, \\
 &= \underbrace{\frac{1}{2} N_{sw} C_L V_{DD} V_{swing}}_{\text{Dynamic}} + \underbrace{V_{DD} I_{leak} T_{del}}_{\text{static}}, \tag{3}
 \end{aligned}$$

where $N_{sw} = \alpha_F N_{nodes}$ is the number of switching nodes for the operation and T_{del} is the circuit delay (following the investigations in [4]). In the case of CMOS logic style, V_{swing} is equal to the supply voltage being used.

As for the security metrics, our choice was dictated by various constraints and features. First, the shape of the instantaneous dynamic power (i.e. the side-channel signal) changes significantly depending on the technology and the supply voltage used. Therefore, it is important to consider these changes while evaluating the CMOS and DDSLL logic styles, and to consider a multivariate analysis. Second, our target implementations are not masked, and therefore our simulations exclusively exploit first-order leakages. So while in general, a fair comparison of our logic styles would require to compute a mutual information metric [32], in this particular case we can simplify our evaluations by (i) applying a dimensionality reduction to our traces, namely a Principal Component Analysis (PCA) which will capture the shape of the noise-free simulated traces [2], and (ii) computing Mangard’s SNR on the reduced traces [20], which is equivalent to the mutual information metric in this case [22, 9].

For this purpose, we denote a power trace as \mathbf{l} , its corresponding random variable as \mathbf{L} , and assume that this random variable is a function of the S-box input X , and a noise random variable \mathbf{N} . The multivariate traces are reduced to univariate ones thanks to PCA, which we denote as: $l = \text{PCA}(\mathbf{l})$. Giving the trace l a subscript x corresponding to the input and a superscript i corresponding to an index (since the trace for a plaintext x can be measured multiple times), Mangard’s SNR is defined as:

$$\text{SNR} = \frac{\hat{\text{var}}_x(\hat{\text{E}}_i(L_x^i))}{\hat{\text{E}}_x(\hat{\text{var}}_i(L_x^i))}, \tag{4}$$

where $\hat{\text{E}}$ (resp. $\hat{\text{var}}$) denotes the sample mean (resp. variance) operator. In our following simulations, this SNR will be computed for noise-free traces. This

amounts to maximizing the signal $\hat{\text{var}}_x(\hat{\mathbf{E}}_i(L_x^i))$ (i.e. we ignore the denominator of Equation 4 in this case).

Note that for readability, our following results only report the quantification of our experiments with this security metric. However, various other (heuristic) choices could be considered, e.g. computing the SNR for the most informative samples in the traces, or considering more dimensions after the application of the PCA. The same holds with performance metrics (e.g. the throughput over area ratio could be used as alternative efficiency metric). In our study, none of these variations (that we also browsed) lead to different conclusions regarding the two main trends outlined in introduction.

3 AES S-box implementations

3.1 AES S-box

For the sake of simplicity, and in order to demonstrate the technology trends of CMOS and DDSLL, we chose an 8-bit AES S-box as the benchmark circuit. More specifically, we considered a combinatorial implementation of the S-box from [14], based on the architecture proposed in [31, 23]. Thanks to mapping the elements of the original field $\text{GF}(2^8)$ to the composite field $\text{GF}(((2^2)^2)^2)$, the gate complexity and the power consumption can be greatly reduced. The adopted S-box consists of 3 stages: a transformation stage to map the elements to the $\text{GF}(((2^2)^2)^2)$ field, an inversion stage and an inverse transformation stage to map the elements back to $\text{GF}(2^8)$, grouped with the affine transformation.

3.2 Target designs

The CMOS and DDSLL AES S-boxes were implemented in a full-custom fashion, using the CADENCE Virtuoso tools, in a low-power 65 nanometer bulk technology and 28 nanometer FDSOI technology. The CMOS implementation of the S-box is made only of 2-input AND/NAND and 2-input XOR/XNOR gates. The total number of transistors is 1,530, with a logic depth of 22 gates. On the other hand, the DDSLL S-box accounts for 1275 transistors, with a logic depth of 13 gates.

The gate design of CMOS and DDSLL S-boxes in both 65 nanometer bulk and 28 nanometer FDSOI technologies is kept identical, i.e. the gates used and the number of transistors remain unchanged. However, we respected the minimum feature size of each technology (to decrease the switching capacitance, hence the energy per operation) and resized the transistors' widths adequately to guarantee functionality. In 65 nanometer bulk technology, standard threshold voltage (SVT) transistors are used to reduce the static current while maintaining good performances with respect to the circuit delay. For benchmarking purposes, we also implemented the CMOS and DDSLL S-boxes using the low threshold voltage (LVT) transistors available from the same technology. In 28 nanometer FDSOI technology, both SVT and LVT transistors were again used to maintain a fair

comparison with the 65 nanometer technology. Yet, for readability, our following results only report the quantification of our experiments with SVT transistors (trends are again identical for LVT transistors). Eventually, the widths of the DDSLL S-box transistors were chosen such that the voltage swing is sufficient for the circuit to operate correctly at the lower limit of the supply voltages of each technology.

In our experiments, all the S-boxes are fed with buffered inputs to maintain equal fan-ins and have realistic inputs, yet the DDSLL S-box additionally has a buffered clock. Also, all the outputs of the S-boxes are loaded with equal fan-out buffers. Each S-box is provided with a separate supply voltage than that of the input/output buffers so that the buffers' energy consumptions are not taken into account in our evaluations.

3.3 Simulation settings

Simulations for the above designs are done at the schematic level (without any extracted post-layout capacitances) using Eldo simulator based on SPICE models provided by the industrial foundries at room temperature of 25°C. In this respect, we note that any imbalance in the parasitic elements would affect the difference between the delay of the differential routes of the DDSLL S-box, which would impact its power consumption, leading to a higher (exploitable) signal being observed by practical attacks [30]. And this is expected to get only worse with technology scaling and variability, since balancing the capacitances in an implementation naturally becomes more challenging with smaller circuits and smaller routing capacitances. Therefore, and as previously mentioned, taking the parasitic routing capacitances into consideration could only amplify our observation (which is that the comparative advantage of DDSLL over CMOS is vanishing with technology scaling). A similar statement holds for other effects that we did not take into consideration in this work such as crosstalk and the influence of process, voltage and temperature (PVT) variations.

Consequently, we assume our results correspond to an ideal scenario and the inclusion of more physical default(s) should only deteriorate the performance of dual-rail logic styles compared to CMOS.

The frequency of operation is chosen to be 10MHz, which is in accordance with the usual operating frequencies for cryptographic applications (see, e.g.[3]). The supply voltage is swept across a range of 500mV, in steps of 100mV starting from the nominal voltage of each technology, namely, 1.2V and 1V for the 65 nanometer bulk and the 28 nanometer FDSOI technologies, respectively. The lower limit of the supply voltages is imposed by the correct functionality of the circuit at the target frequency for a given implementation.

Note that operating at high supply voltages and using the 28nm technology node allows the circuit implementation to run at higher frequencies. But the frequency choice has no impact on our results since we have chosen to use the energy per operation as the evaluation metric and not the total power consumption.

To calculate the energy per operation, we considered 1000 random input signals. As for the security analysis, the S-box input signal has 256 possible values

whose transitions are chosen between 0 and an arbitrary input. Restricting the inputs to a subset from the 256^2 possible inputs was mainly motivated by practical simulation constraints (memory and simulation time) and is not expected to strongly impact the comparison between the logic styles.

4 Comparative scaling trends: CMOS vs. DDSLL

In this section, we aim to compare the CMOS and DDSLL logic styles and study how their security versus performance tradeoff evolves with technology and supply voltage scaling. To be able to do that in a comprehensible manner, we plot the ratios between our (security and performance) metrics computed for both CMOS and DDSLL, one in function of the other. More precisely, Figure 2 shows the PCA signal ratio between CMOS and DDSLL S-boxes versus the energy per operation ratio between these logic styles for the 65 nanometer bulk and the 28 nanometer FDSOI technologies (designed using SVT devices). The different points represent the supply voltages we used that span over a range of $0.5V$ starting from the nominal supply of each technology ($1.2V$ and $1V$ for the 65 and 28 nanometer technologies, respectively). This study allows us to make the following observations:

- 1 By reducing the supply voltage, the energy per operation ratio of CMOS with respect to DDSLL is decreasing for the 65 and 28 nanometer technologies. This reduction can be explained by the fact that the E_{op} value of the DDSLL style decreases almost linearly with the supply voltage, because it maintains nearly the same voltage swing while E_{op} of CMOS decreases quadratically (see Equation 3).
- 2 As for the PCA signal ratio between CMOS and DDSLL (for both technologies), it also decreases with the supply voltage scaling. Again, this is due to the fact that the voltage swing of the DDSLL S-box is kept almost unchanged leading to a slow reduction rate of the transient power consumption. Hence the PCA signal with V_{DD} scaling is less compared to that of CMOS.
- 3 Most importantly, Figure 2 illustrates that at 28 nanometer technology, the PCA signal ratio between CMOS and DDSLL is less than that of the 65 nanometer technology for similar E_{op} ratios. This figure neatly puts forward that the security vs. performance tradeoff between these logic styles does not scale positively for DDSLL, even though we do not consider any routing parasitics or PVT variations in our simulations.

It is worth emphasizing that similar observations were made by comparing CMOS to DDSLL S-boxes using LVT devices in both technologies. Therefore, changing the device type leads only to either better performance or more power savings, but the technology and supply voltage scaling trends remain the same. Also, and for the sake of completeness, we conducted the same experiments using the maximum SNR (before PCA) as a security metric and the technology and supply voltage scaling trends remained unchanged.

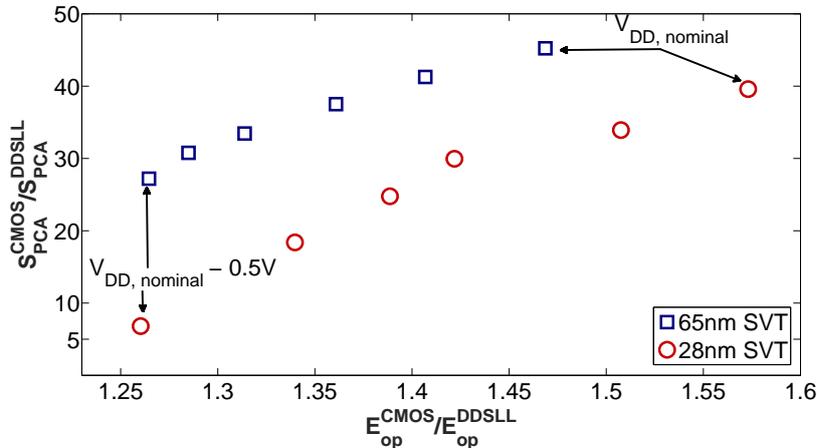


Fig. 2: Evolution of the tradeoff between the PCA signal ratio and the E_{op} ratio for CMOS vs. DDSLL using 65 & 28 nanometer technology nodes.

In addition, we note that changing the frequency of operation does not impact our conclusions as long as the dynamic power consumption dominates. We simulated the S-boxes down to 100kHz and the technology and voltage scaling trends were again the same.

Eventually, we note that technology scaling is advancing at a fast pace and secure implementations will soon follow (given the fact that up until now applications such as smart cards tend to lag by one or two technologies). Also, circuit designers generally aim at scaling the supply voltage in order to further reduce the energy consumption of the digital circuits (sometimes operating below the transistor subthreshold voltage leading to minimum energy per operation). Both trends lead us to conclude that the observations in this section may rapidly have concrete relevance.

5 Technology and V_{DD} scaling trends for CMOS

Since the comparative advantage of DDSLL over CMOS vanishes in our case study, one natural complementary question is whether scaling trends for CMOS circuits lead to a more positive conclusion. In this section, we answer this question by focusing on the impact of technology and V_{DD} scaling on CMOS circuits only. For this purpose, Figure 3(a) reports the energy per operation of the CMOS S-box for different supply voltages, using both 65 and 28 nanometer technologies. We recall that the minimum V_{DD} was chosen such that the CMOS S-box operates correctly at the 10MHz target frequency for each technology. As expected, the energy per operation of CMOS decreases almost quadratically with the reduction of the supply voltage (see Equation 3). The figure also shows clearly that technology scaling from 65 to 28 nanometer reduces the energy per operation

(of the CMOS S-box) by a factor of $2.2\times$. This reduction is compliant with the expected technology scaling trend as explained in [12].

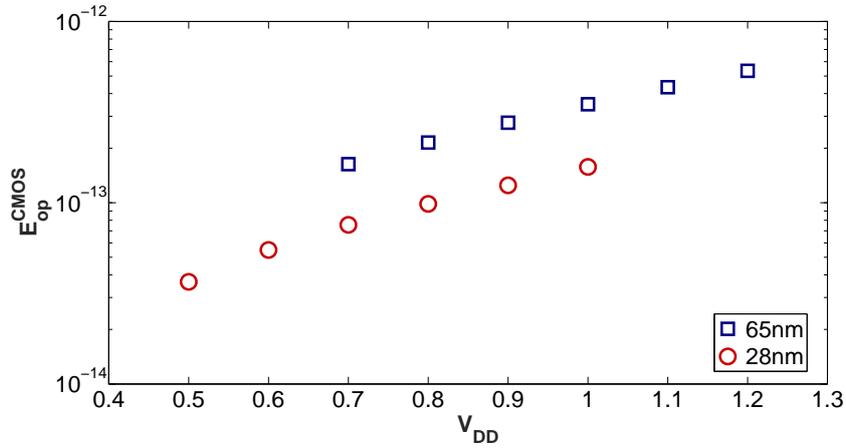
Similarly, Figure 3(b) shows the signal after PCA of the CMOS S-box for different V_{DD} values, using both 65 and 28 nanometer technologies. The PCA signal of CMOS decreases as the supply voltage scales down. In the 65 nanometer technology, the S_{PCA} reduction is more than one order of magnitude across the whole V_{DD} range (1.2V to 0.7V) and it reaches even more than two orders of magnitude in the 28 nanometer technology, by reducing the supply voltage from 1V down to 0.5V. As for the technology scaling from 65 to 28 nanometer, it is also clear that the signal after PCA of the CMOS S-box decreases by a factor of $2.3\times$ at comparable supply voltages. So the scaling trends for PCA signal is also positive for CMOS. Here, we note that an analytical explanation of these observations is more challenging, since the small data-dependent current variations that lead to exploitable side-channel signal are much harder to capture theoretically. (Yet, we can suppose that the aforementioned signal reductions essentially originate from the interactions between reductions of the current and variations of the load capacitance).

6 Conclusion

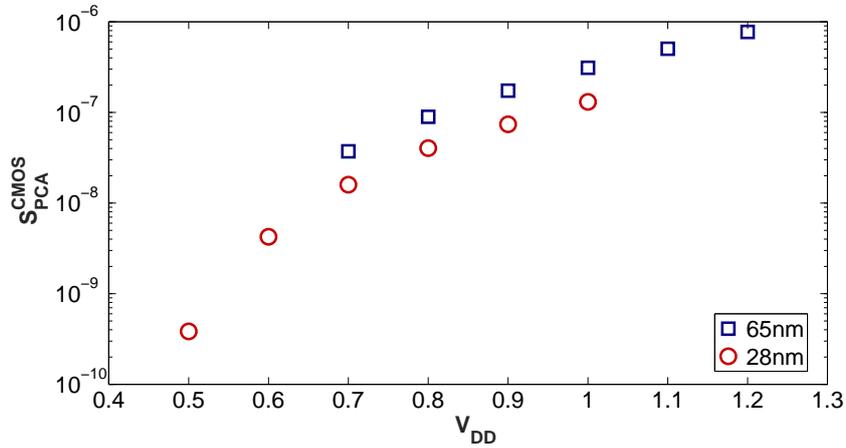
In this case study, we analyzed for the first time the comparative scaling trends of CMOS and dual-rail logic styles. In short, our evaluations suggest that the interest of DDSLL over CMOS, expressed in terms of a security vs. performance tradeoff, vanishes as circuit sizes shrink. To a good extent, we believe a similar conclusion should be obtained for other dual-rail logic styles. In particular, from the security point-of-view, they all suffer from capacitance imbalances to some extent, and this phenomenon can only be magnified in smaller technologies. While our case study was based on an AES S-box, we believe similar trends should also be obtained for full AES implementations. Indeed, similar energy trends will be integrated over more clock cycles, and the signal variations exploited in a DPA are anyway focused on the first cipher rounds. These results therefore suggest that reducing the SNR in advanced technologies may be better achieved by exploiting the naturally increasing intrinsic noise level than by reducing the signal using dual-rail logic styles. It also suggests the design of noisy and efficient CMOS implementations as an interesting scope for further research.

We finally note that while dual-rail logic styles may not be a sustainable solution for signal reduction purposes, it remains possible that they are helpful ingredients of physically secure implementations for other reasons (e.g. in order to facilitate the independence condition that is required for masking to deliver its security promises [39]), which is an interesting scope for further research.

Acknowledgments. 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.



(a)



(b)

Fig. 3: Scaling trends of E_{op} and the PCA signal for the CMOS S-box across a supply voltage range of 500mV for 65 and 28 nanometer tech.

References

1. Mohamed Allam and Mohamed Elmasry. Dynamic current mode logic (DyCML): A new low-power high-performance logic style. *IEEE Journal of Solid-State Circuits*, 36(3):550–558, 2001.
2. Cédric Archambeau, Eric Peeters, François-Xavier Standaert, and Jean-Jacques Quisquater. Template attacks in principal subspaces. In Louis Goubin and Mitsuru Matsui, editors, *Cryptographic Hardware and Embedded Systems - CHES 2006, 8th International Workshop, Yokohama, Japan, October 10-13, 2006, Proceedings*, volume 4249 of *Lecture Notes in Computer Science*, pages 1–14. Springer, 2006.

3. D. Bellizia, S. Bongiovanni, P. Monsurro, G. Scotti, and A. Trifiletti. Univariate power analysis attacks exploiting static dissipation of nanometer CMOS VLSI circuits for cryptographic applications. *IEEE Transactions on Emerging Topics in Computing*, PP(99):1–1, 2016.
4. David Bol, Dina Kamel, Denis Flandre, and Jean-Didier Legat. Nanometer MOS-FET effects on the minimum-energy point of 45nm subthreshold logic. In *Proceedings of the 2009 International Symposium on Low Power Electronics and Design, 2009, San Francisco, CA, USA, August 19-21, 2009*, pages 3–8, 2009.
5. Eric Brier, Christophe Clavier, and Francis Olivier. Correlation power analysis with a leakage model. In *Cryptographic Hardware and Embedded Systems - CHES 2004: 6th International Workshop Cambridge, MA, USA, August 11-13, 2004. Proceedings*, pages 16–29, 2004.
6. Suresh Chari, Charanjit S. Jutla, Josyula R. Rao, and Pankaj Rohatgi. Towards sound approaches to counteract power-analysis attacks. In Wiener [38], pages 398–412.
7. Suresh Chari, Josyula R. Rao, and Pankaj Rohatgi. Template attacks. In *Cryptographic Hardware and Embedded Systems - CHES 2002, 4th International Workshop, Redwood Shores, CA, USA, August 13-15, 2002, Revised Papers*, pages 13–28, 2002.
8. Zeynep Toprak Deniz and Yusuf Leblebici. Low-power current mode logic for improved DPA-resistance in embedded systems. In *International Symposium on Circuits and Systems (ISCAS 2005), 23-26 May 2005, Kobe, Japan*, pages 1059–1062. IEEE, 2005.
9. Alexandre Duc, Sebastian Faust, and François-Xavier Standaert. Making masking security proofs concrete - or how to evaluate the security of any leaking device. In Elisabeth Oswald and Marc Fischlin, editors, *Advances in Cryptology - EUROCRYPT 2015 - 34th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Sofia, Bulgaria, April 26-30, 2015, Proceedings, Part I*, volume 9056 of *Lecture Notes in Computer Science*, pages 401–429. Springer, 2015.
10. Swaroop Ghosh and Kaushik Roy. Parameter variation tolerance and error resiliency: New design paradigm for the nanoscale era. *Proceedings of the IEEE*, 98(10):1718–1751, 2010.
11. L. Giancane, P. Marietti, M. Olivieri, G. Scotti, and A. Trifiletti. A new dynamic differential logic style as a countermeasure to power analysis attacks. In *Electronics, Circuits and Systems, 2008. ICECS 2008. 15th IEEE International Conference on*, pages 364–367, Aug 2008.
12. Wilfried Haensch, Edward J. Nowak, Robert H. Dennard, Paul M. Solomon, Andres Bryant, Omer H. Dokumaci, Arvind Kumar, Xinlin Wang, Jeffrey B. Johnson, and Massimo V. Fischetti. Silicon CMOS devices beyond scaling. *IBM Journal of Research and Development*, 50(4-5):339–362, 2006.
13. Christoph Herbst, Elisabeth Oswald, and Stefan Mangard. An AES smart card implementation resistant to power analysis attacks. In Jianying Zhou, Moti Yung, and Feng Bao, editors, *Applied Cryptography and Network Security, 4th International Conference, ACNS 2006, Singapore, June 6-9, 2006, Proceedings*, volume 3989 of *Lecture Notes in Computer Science*, pages 239–252, 2006.
14. D. Kamel, F. X. Standaert, and D. Flandre. Scaling trends of the AES S-box low power consumption in 130 and 65 nm CMOS technology nodes. In *2009 IEEE International Symposium on Circuits and Systems*, pages 1385–1388, May 2009.

15. Dina Kamel, Mathieu Renaud, David Bol, Francois-Xavier Standaert, and Denis Flandre. Analysis of dynamic differential swing limited logic for low-power secure applications. *Journal of Low Power Electronics and Applications*, 2(1):98, 2012.
16. Dina Kamel, Mathieu Renaud, Denis Flandre, and François-Xavier Standaert. Understanding the limitations and improving the relevance of SPICE simulations in side-channel security evaluations. *J. Cryptographic Engineering*, 4(3):187–195, 2014.
17. Stéphanie Kerckhof, François Durvaux, Cédric Hocquet, David Bol, and François-Xavier Standaert. Towards green cryptography: A comparison of lightweight ciphers from the energy viewpoint. In Emmanuel Prouff and Patrick Schaumont, editors, *Cryptographic Hardware and Embedded Systems - CHES 2012 - 14th International Workshop, Leuven, Belgium, September 9-12, 2012. Proceedings*, volume 7428 of *Lecture Notes in Computer Science*, pages 390–407. Springer, 2012.
18. Paul C. Kocher, Joshua Jaffe, and Benjamin Jun. Differential power analysis. In Wiener [38], pages 388–397.
19. François Macé, François-Xavier Standaert, and Jean-Jacques Quisquater. Information theoretic evaluation of side-channel resistant logic styles. In Paillier and Verbauwhede [25], pages 427–442.
20. Stefan Mangard. Hardware countermeasures against DPA ? A statistical analysis of their effectiveness. In *Topics in Cryptology - CT-RSA 2004, The Cryptographers' Track at the RSA Conference 2004, San Francisco, CA, USA, February 23-27, 2004, Proceedings*, pages 222–235, 2004.
21. Stefan Mangard, Elisabeth Oswald, and Thomas Popp. *Power analysis attacks - revealing the secrets of smart cards*. Springer, 2007.
22. Stefan Mangard, Elisabeth Oswald, and François-Xavier Standaert. One for all - all for one: unifying standard differential power analysis attacks. *IET Information Security*, 5(2):100–110, 2011.
23. Nele Mentens, Lejla Batina, Bart Preneel, and Ingrid Verbauwhede. A systematic evaluation of compact hardware implementations for the Rijndael S-Box. In *Topics in Cryptology - CT-RSA 2005, The Cryptographers' Track at the RSA Conference 2005, San Francisco, CA, USA, February 14-18, 2005, Proceedings*, pages 323–333, 2005.
24. Amir Moradi. Side-channel leakage through static power - should we care about in practice? In Lejla Batina and Matthew Robshaw, editors, *Cryptographic Hardware and Embedded Systems - CHES 2014 - 16th International Workshop, Busan, South Korea, September 23-26, 2014. Proceedings*, volume 8731 of *Lecture Notes in Computer Science*, pages 562–579. Springer, 2014.
25. Pascal Paillier and Ingrid Verbauwhede, editors. *Cryptographic Hardware and Embedded Systems - CHES 2007, 9th International Workshop, Vienna, Austria, September 10-13, 2007, Proceedings*, volume 4727 of *Lecture Notes in Computer Science*. Springer, 2007.
26. Thomas Popp, Mario Kirschbaum, Thomas Zefferer, and Stefan Mangard. Evaluation of the masked logic style MDPL on a prototype chip. In Paillier and Verbauwhede [25], pages 81–94.
27. Thomas Popp and Stefan Mangard. Masked dual-rail pre-charge logic: DPA-resistance without routing constraints. In Josyula R. Rao and Berk Sunar, editors, *Cryptographic Hardware and Embedded Systems - CHES 2005, 7th International Workshop, Edinburgh, UK, August 29 - September 1, 2005, Proceedings*, volume 3659 of *Lecture Notes in Computer Science*, pages 172–186. Springer, 2005.

28. Santos Merino Del Pozo, François-Xavier Standaert, Dina Kamel, and Amir Moradi. Side-channel attacks from static power: when should we care? In *Proceedings of the 2015 Design, Automation & Test in Europe Conference & Exhibition, DATE 2015, Grenoble, France, March 9-13, 2015*, pages 145–150, 2015.
29. Francesco Regazzoni, Thomas Eisenbarth, Axel Poschmann, Johann Großschädl, Frank K. Gürkaynak, Marco Macchetti, Zeynep Toprak Deniz, Laura Pozzi, Christof Paar, Yusuf Leblebici, and Paolo Ienne. Evaluating resistance of MCML technology to power analysis attacks using a simulation-based methodology. *Trans. Computational Science*, 4:230–243, 2009.
30. Mathieu Renaud, Dina Kamel, François-Xavier Standaert, and Denis Flandre. Information theoretic and security analysis of a 65-nanometer DDSLL AES S-Box. In Bart Preneel and Tsuyoshi Takagi, editors, *Cryptographic Hardware and Embedded Systems - CHES 2011 - 13th International Workshop, Nara, Japan, September 28 - October 1, 2011. Proceedings*, volume 6917 of *Lecture Notes in Computer Science*, pages 223–239. Springer, 2011.
31. Akashi Satoh, Sumio Morioka, Kohji Takano, and Seiji Munetoh. A compact Rijndael hardware architecture with S-Box optimization. In *Advances in Cryptology - ASIACRYPT 2001, 7th International Conference on the Theory and Application of Cryptology and Information Security, Gold Coast, Australia, December 9-13, 2001, Proceedings*, pages 239–254, 2001.
32. François-Xavier Standaert, Tal Malkin, and Moti Yung. A unified framework for the analysis of side-channel key recovery attacks. In Antoine Joux, editor, *Advances in Cryptology - EUROCRYPT 2009, 28th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Cologne, Germany, April 26-30, 2009. Proceedings*, volume 5479 of *Lecture Notes in Computer Science*, pages 443–461. Springer, 2009.
33. François-Xavier Standaert, Nicolas Veyrat-Charvillon, Elisabeth Oswald, Benedikt Gierlichs, Marcel Medwed, Markus Kasper, and Stefan Mangard. The world is not enough: Another look on second-order DPA. In Masayuki Abe, editor, *Advances in Cryptology - ASIACRYPT 2010 - 16th International Conference on the Theory and Application of Cryptology and Information Security, Singapore, December 5-9, 2010. Proceedings*, volume 6477 of *Lecture Notes in Computer Science*, pages 112–129. Springer, 2010.
34. Kris Tiri and Ingrid Verbauwhede. Securing encryption algorithms against DPA at the logic level: Next generation smart card technology. In Colin D. Walter, Çetin Kaya Koç, and Christof Paar, editors, *Cryptographic Hardware and Embedded Systems - CHES 2003, 5th International Workshop, Cologne, Germany, September 8-10, 2003, Proceedings*, volume 2779 of *Lecture Notes in Computer Science*, pages 125–136. Springer, 2003.
35. Kris Tiri and Ingrid Verbauwhede. A logic level design methodology for a secure DPA resistant ASIC or FPGA implementation. In *2004 Design, Automation and Test in Europe Conference and Exposition (DATE 2004), 16-20 February 2004, Paris, France*, pages 246–251. IEEE Computer Society, 2004.
36. Kris Tiri and Moonmoon Akmal Ingrid Verbauwhede. A dynamic and differential cmos logic with signal independent power consumption to withstand differential power analysis on smart cards. In *Solid-State Circuits Conference, 2002. ESSCIRC 2002. Proceedings of the 28th European*, pages 403–406. IEEE, 2002.
37. Nicolas Veyrat-Charvillon, Marcel Medwed, Stéphanie Kerckhof, and François-Xavier Standaert. Shuffling against side-channel attacks: A comprehensive study with cautionary note. In Xiaoyun Wang and Kazue Sako, editors, *Advances in*

- Cryptology - ASIACRYPT 2012 - 18th International Conference on the Theory and Application of Cryptology and Information Security, Beijing, China, December 2-6, 2012. Proceedings*, volume 7658 of *Lecture Notes in Computer Science*, pages 740–757. Springer, 2012.
38. Michael J. Wiener, editor. *Advances in Cryptology - CRYPTO '99, 19th Annual International Cryptology Conference, Santa Barbara, California, USA, August 15-19, 1999, Proceedings*, volume 1666 of *Lecture Notes in Computer Science*. Springer, 1999.
 39. Alexander Wild, Amir Moradi, and Tim Güneysu. GliFreD: Glitch-free duplication - towards power-equalized circuits on FPGAs. *IACR Cryptology ePrint Archive*, 2015:124, 2015.