

SDDS-NCL Design: Analysis of Supply Voltage Scaling

Ricardo A. Guazzelli, Fernando G. Moraes, Ney L. V. Calazans, Matheus T. Moreira
GAPH - FACIN - Pontifícia Universidade Católica do Rio Grande do Sul, Porto Alegre, RS, Brazil
ricardo.guazzelli@acad.pucrs.br,
{ney.calazans,fernando.moraes,matheus.moreira}@pucrs.br

ABSTRACT

Despite their substantial power savings, voltage scaling design increases the concern about sensitivity to manufacturing process and operating conditions variations. These can induce significant delay changes in fabricated circuits. An elegant approach to cope with these issues is to employ quasi delay-insensitive asynchronous design styles, which allow relaxing timing assumptions, enabling simpler timing closure when compared to clocked solutions. This work explores the effects of supply voltage scaling on a specific class of quasi-delay-insensitive circuits called spatially distributed dual spacer null convention logic (SDDS-NCL). It first analyzes basic SDDS-NCL gates from a 65 nm cell library. The analysis explores the effects of supply voltage scaling on isolated cells, encompassing static power, energy and delay trade-offs. Next, it shows the results of a similar analysis applied to a 324-cell case study circuit. Results indicate that the evaluated class of circuits can significantly benefit from sub- and near-threshold operation to trade off energy efficiency and performance.

Categories and Subject Descriptors

B.6.3 [Logic Design]: Design Aids—*Optimization*; B.7.1 [Integrated Circuits]: Types and Design Styles—*Standard Cells, VLSI*

General Terms

Design

Keywords

Voltage Scaling, Asynchronous Circuits, NCL, SDDS-NCL

1. INTRODUCTION & RELATED WORK

Many contemporary applications impose low power as a main design constraint. As examples, distributed sensor networks and wearable devices require low power and/or energy consumption as a fundamental feature. In addition,

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

Copyright SBCCI'15, August 31 - September 04, 2015, Salvador, Brazil ©2015 ACM ISBN 978-1-4503-3763-2/15/08...\$15.00 DOI: <http://dx.doi.org/10.1145/2800986.2800999>.

with decreasing minimum feature size, transistors become increasingly leaky, augmenting static power dissipation of integrated circuits, which imposes new challenges for designers to meet stringent power constraints [1]. This motivates the research on design techniques for minimizing energy and/or power while reaching some minimum performance constraint. These efforts usually focus on high performance strong inversion operation (or super-threshold region) and are implemented at the architectural level, where designers can reduce the computation workload or improve architecture to achieve better power optimization [2]. At the circuit level, a compelling approach to lower power consumption is reducing the voltage supply, also called supply voltage scaling, or simply voltage scaling (VS). As the supply voltage is quadratically related to power dissipation, VS is a very effective low power design technique [3]. Taking this to the extreme, some systems operate in the sub-threshold region of transistors, meaning with supply voltage of a few hundreds mV.

Unfortunately, delay variations increase significantly at low supply voltage levels [4]. As a consequence, problems such as critical path changes at different voltage levels and increased hold and setup times uncertainty on registers can emerge. This makes the design of clocked circuits much more challenging, as the operating frequency needs tighter control, and significant margins may need to be added to the clock. This can compromise performance, complicate clock distribution design and cause area and power overheads. Asynchronous circuits are an alternative to overcome such issues [1]. These circuits do not rely on a discrete notion of time and can be designed using different templates, most of which employ either bundled-data or quasi-delay-insensitive (QDI) templates. Considering VS, QDI appears as a practical option, given that it allows a much more relaxed timing than synchronous or bundled-data designs [5].

There are different ways to design a QDI circuit, and among the styles in contemporary literature, Null Convention Logic (NCL) stands out because it enables semi-custom design and several works demonstrate its efficiency [6–10]. In fact, a recently proposed optimization of NCL logic, called spatially distributed dual spacer NCL (SDDS-NCL) [11] is particularly interesting, because it allows optimizations over previous propositions and enables using conventional electronic design automation (EDA) tools [12]. However, SDDS-NCL originally employed only super-threshold operation at nominal supply voltages, and there is no assessment of if and how it can benefit from VS.

As the sub-threshold regime presents high sensibility to

process, voltage and temperature (PVT) variations, timing uncertainties may compromise circuit functionality. To overcome this, NCL design allows to mitigate several timing issues, as Jorgenson et al. describe [13]. Here, authors employ NCL on the sub-threshold regime and use a military application as case-study. Authors show that voltage scaling and power gating together may achieve 7-20x energy savings. Their implementation provides significant power and performance advantages, compared to a conventional synchronous design. Also, the delay insensitive characteristic of NCL provides a robust design, enabling lower sensibility to PVT variations and longer lifetime. In [9], Parsan and Smith evaluate the basic blocks of an NCL design, comparing several gate topologies of a 4x4 NCL multiplier. According to authors the static topology achieves the lowest supply voltage (200mV), while still maintaining correct functionality. Authors conclude the static topology is the best option when targeting low voltage/energy/power, if speed and area are not major concerns. These works motivate the use of QDI for asynchronous design for VS applications. However, they restrict attention to NCL only, missing the richness of SDDS-NCL, a larger gate family, as Section 2 depicts.

This work presents a first step to fill this gap, exploring the effects of VS on SDDS-NCL gates and on a 324-gate case study circuit. It employs the 65nm ASCEnD cell library [14], [10]. This analysis discusses the trade-offs of static power, energy efficiency and performance as voltage is scaled. At the cell level we evaluate several gates with different strengths, number of inputs and functionality. At the circuit level, we present the analysis of a case-study 8-bit Kogge-Stone adder. The obtained results show that SDDS-NCL designs are good candidates for VS applications, as we observed energy improvements of up to 14x when scaling the supply voltage. More importantly, due to their QDI nature, no extra care is required with timing constraints when submitting an SDDS-NCL design to VS. In fact, during the experiments we observe the circuit gracefully scales its performance, power and energy efficiency figures under VS.

2. SDDS-NCL DESIGN

Asynchronous design techniques generically classify in either bundled-data and quasi delay-insensitive (QDI) template families. Despite the fact that bundled-data designs can benefit to some extent from the use of conventional design tools, due to its similarity to synchronous circuits, these require extra care with the definition and verification of timing constraints between data and control signals. Accordingly, such circuits can be very sensitive to delay variations [5], even in the super-threshold regime, making them often inadequate to VS applications. An alternative to avoid such issues is to encode control signals within data communication channels, which partly defines QDI design. Martin and Nyström report QDI as the most practical asynchronous template, due to its relaxed timing constraints and robustness to delay variations [5], which also make the template suited for VS applications.

In the 90s, Theseus Logic proposed the NCL logic family [6] to implement QDI asynchronous logic. Since then, NCL has been applied to deal with power problems [15, 16], to design high speed circuits [17, 18] and fault tolerant applications [19], among other uses. An evolution of NCL design is NCL+, proposed in [10]. A fundamental difference between NCL and NCL+ is that the latter relies on a different

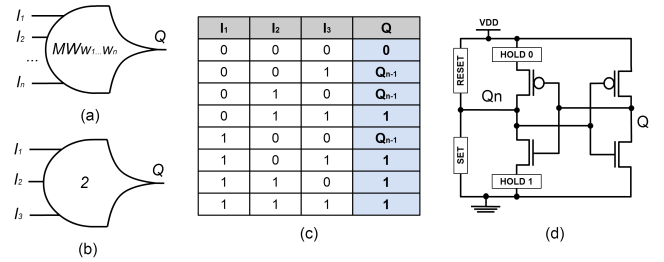


Figure 1: NCL basics: (a) generic NCL gate symbol, (b) NCL2-of-3 gate, (c) truth table for the NCL2-of-3 gate and (d) generic NCL static topology.

handshake protocol called return-to-one (RTO) [20]. A recent work [11] showed that with a basic set of constraints for technology mapping [12] it is possible to combine NCL and NCL+ in a single design. This coupling constitutes a new design style called SDDS-NCL, which can provide substantial improvements w.r.t. pure NCL or NCL+ alone [11, 12]. The design of SDDS-NCL circuits requires NCL and NCL+ gates. These couple a threshold function with positive integer weights assigned to inputs to the use of a hysteresis mechanism. Most NCL gates are thus sequential circuits. Figure 1(a) shows a generic symbol for an NCL gate noted $MW_{w_1..w_n}$ -of- N , where M is the threshold function, N is the number of inputs and the w_i s are input weights. Depending on the NCL gate function, each input can have a specific weight. Weight 1 is assumed whenever these are omitted. For example, Figure 1(b) shows the symbol of an NCL gate with $M = 2$, $N = 3$ and $w_1 = w_2 = w_3 = 1$. Given its specific function and naming style, we call this NCL2W111-of-3 or just NCL2-of-3.

The output of an NCL gate switches according to the following premises: (1) a high-to-low transition can only occur when all inputs are at logic 0; (2) a low-to-high transition occurs when the sum of weights for inputs at logic 1 is bigger than or equal to the threshold M . In case the input values and their weights do not combine to reach M , the output holds its previous state. For example, Figure 1(c) shows the NCL2-of-3 gate truth table. The output of this gate will only switch to 1 when 2 or more inputs are at 1, and it will only switch to 0 when all inputs are 0. For all other cases, the output keeps its previous state. Regarding transistor topologies, NCL gates can be realized with several distinct approaches [9, 10] As previous work showed the static NCL topology is the best option for low voltage operation, this work focus only on these. Figure 1(d) illustrates a generic NCL gate static topology formed by a memory element and four logic blocks: SET, RESET, HOLD1 and HOLD0. SET is responsible for producing logic 1 at the output, according to the functionality of the gate and RESET acts to force 0 on the output. HOLD1 and HOLD0 control the feedback loop, that can retain the output value when neither RESET nor SET blocks are active. In fact, regarding transistors arrangement, HOLD1 acts as a complement of RESET and HOLD0 as a complement of SET.

NCL+ gates have a similar functionality. However, the assumption of the RTO protocol mandates the switching function of an NCL+ gate to be the reverse of its NCL counterpart: the output will only switch to 1 when all inputs are at 1 and to 0 when threshold M is reached by the inputs at 0 [10]. For other input combinations, the output keeps its previous value. The symbol to represent NCL+ gates is similar to

the NCL symbol of Figure 1(a), except for a “+” symbol on its top right corner. Also, the naming style of NCL+ gates is similar to that of NCL, but the NCL label is replaced by NCL+. To design these gates, the same topologies used for NCL can be employed, as the only difference is how to implement blocks SET/HOLD0 and RESET/HOLD1 [10]. Again, this work addresses only static NCL+ gates. Classically, no negative function was supported in either NCL or NCL+. However, this limitation was overcome in SDDS-NCL, as [11, 12] discusses. In this way, to every classic NCL and NCL+ gate there is a corresponding negative unate version gate. This can be useful for circuit synthesis optimizations, because internal inverters already present in these cells can be reused. Throughout this paper, negative unate gates will have the “I” prefix in their label to indicate their distinct functionality. For instance, the negative unate version of gate NCL2W111-of-3 is INCL2W111-of-3.

3. GATE-LEVEL ANALYSIS

This Section explores the effects of VS over a basic set of NCL and NCL+ gates selected from ASCEnD [14], a cell library of gates supporting asynchronous design. We first detail the experimental environment developed to evaluate cells, and next present the obtained results. All experiments consider the use of a 65nm bulk CMOS technology with typical transistors operating at 25°C. The nominal supply of the target technology is 1V.

3.1 Gate Subset Selection

Because ASCEnD counts with several hundreds of cells, effective gate level analysis can benefit from selecting a representative cell subset. Of course, the subset cannot compromise the analysis generality. Our choice is to include gates with four distinctive features related to VS effects: logic family (NCL or NCL+), unateness, number of inputs and driving strength. Comparing families indicates how NCL and NCL+ families behave as voltage is scaled. Different unateness implies different transistor arrangements, being thus relevant to VS. As supply voltage reaches the near-/sub-threshold region, gates with large number of inputs may not be able to provide correct operation due to transistor stacking effects [4]. Thus, the number of inputs range must be assessed to define an appropriate constraint. This work considers 2-,3- and 4-input gates. Evaluating gates with widely distinct driving strengths allows assessing behavior changes under VS. This work thus evaluates ASCEnD cells with three driving strengths: X2, X4 and X13. The selected cell subset in our experiments counts with 31 different cells. However, we discuss here only 9 of these, which we found to be representative of the overall results and that enable to respect this article’s space limitations. Table 1 shows the selected 9 gates and their characteristics.

3.2 Experimental Environment

The experimental setup comprises four steps for cell characterization, simulation and results generation. Figure 2 shows the setup, highlighting the employed tools and their ordering. The first step relies on LiChEn [21], a tool for electrical characterization of asynchronous cells. LiChEn generates simulation files considering all possible input-to-output transition arcs and static states of target cells, which would be too laborious if manually generated. The output is a SPICE deck with the stimuli for exercising all cell arcs and

Table 1: Selected gates with respective unateness, topology type, number of inputs and drive strength.

Gate Name	Unateness	Family	In	Drive
INCL1W11-of-2X4	Negative	NCL	2	X4
INCL2W11-of-2X4	Negative	NCL	2	X4
NCL2W11-of-2X2	Positive	NCL	2	X2
NCL2W11-of-2X4	Positive	NCL	2	X4
NCL2W11-of-2X13	Positive	NCL	2	X13
NCL3W11-of-3X4	Positive	NCL	3	X4
NCL5W2211-of-4X4	Positive	NCL	4	X4
INCL+1W11-of-2X4	Negative	NCL+	2	X4
NCL+5W2211-of-4X4	Positive	NCL+	4	X4




Figure 2: VS cell evaluation: experimental setup.

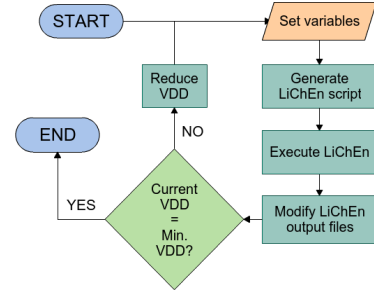


Figure 3: Flowchart for VS characterization. LiChEn is invoked several times, each time generating all characterization files for a given V_{dd} .

states, which reduces manual analysis time and error. As multiple V_{dd} levels are needed for a VS analysis and LiChEn only supports one V_{dd} per characterization, it is necessary to iterate LiChEn calls, using a different V_{dd} in each call. Thus, each cell undergoes the flow in Figure 3 that details the first two steps of the experimental setup (LiChEn and Voltage Scaling Script). This flowchart shows that LiChEn is called multiple times and V_{dd} decreases in steps of 50 mV until it reaches a pre-defined, minimum V_{dd} . We determined this minimum experimentally here to be $V_{dd} = 0.1V$, which ensures correct operation for all gates in the selected subset. All experiments consider a typical process corner and 25°C.

Although LiChEn significantly automates the characterization process, there are issues in near-/sub-threshold operation it disregards. First, it uses an artificially generated input ramp for dynamic and internal arcs, which is acceptable in super-threshold operation. But in near-/sub-threshold operation this can lead to unrealistic behavior due to increased susceptibility to variations. In addition, LiChEn does not consider the possibility of noise in the input, which in near/sub-threshold operation cannot be overlooked, as cells can be unable to represent correct logic levels under noise. Thus, the VS Script modifies the characterization files to provide a more realistic scenario. For input slope generation, we add two inverters in series to each cell input.

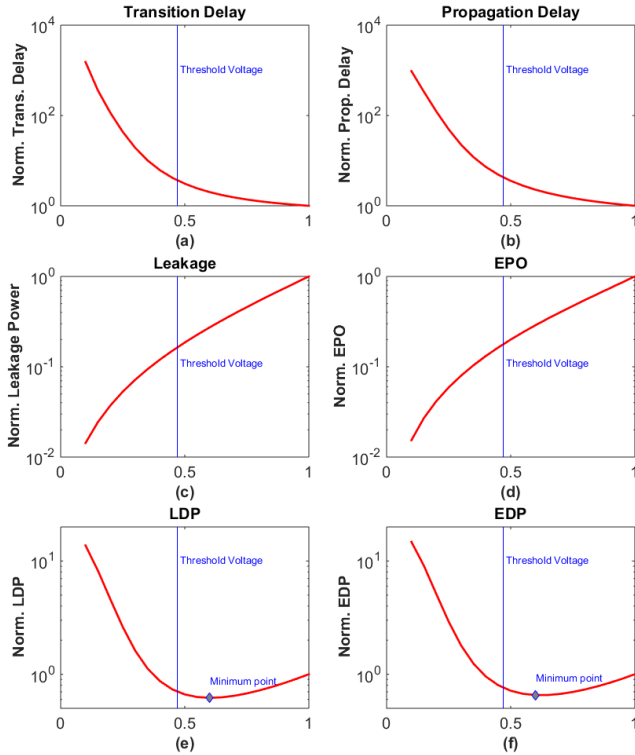


Figure 4: Delay and power analysis results for an NCL2W11-of-2X2 under supply voltage scaling: (a) Transition delay; (b) Propagation delay; (c) Leakage power; (d) EPO; (e) LDP and (f) EDP. Note that x-axes represent the supply voltage and all y-axis values are normalized to the nominal supply voltage values.

These have the same drive strength as the NCL gate. To introduce input noise margin considerations, we modify the waveform of each input such that the high logic level is 90% of V_{dd} and the low logic level is 10% of V_{dd} .

3.3 Simulation Results

Figure 4 details the analysis for a single cell, an NCL2W11-of-2 with X2 driving strength, named NCL2W11-of-2X2. Six aspects are considered: transition delay; propagation delay; leakage power; energy per operation (EPO); leakage-delay product (LDP) and energy-delay product (EDP). Transition delay was measured as the average of the slope in the output for all delay arcs, measured between 10% and 90% of V_{dd} . Propagation delay, also called cell delay, was measured as the average time it takes for a transition in an input propagate through the cell and cause a transition in the output for all delay arcs. For measuring this parameter, we considered 50% of V_{dd} as the switching threshold for both rising and falling transitions. Leakage power was measured as the average power in the power source for all static states of the cell multiplied by V_{dd} . EPO was measured as the average energy consumed for all delay arcs of the cell. Energy was measured as the V_{dd} multiplied by the integral of the current in the power source from the starting point of a transition in the input of the cell until the output switches. Note that to simplify the discussion all charts show values normalized with regard to the nominal supply voltage behavior.

Observing Figures 4 (a) and (b), it is possible to see an

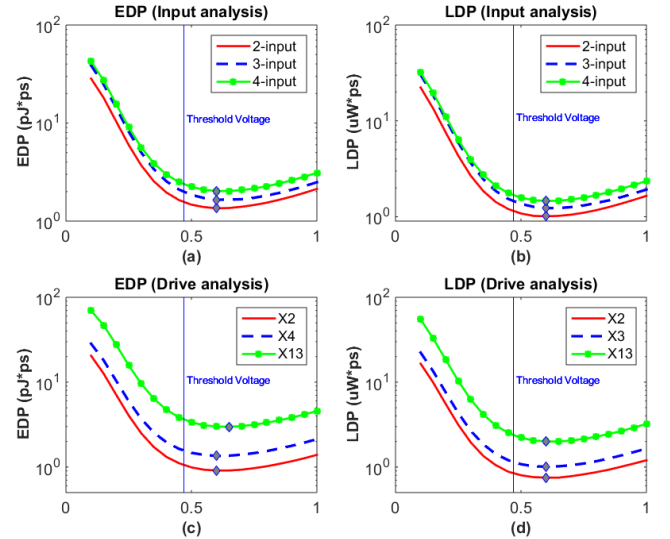


Figure 5: Analysis regarding number of inputs and drive strength for an NCL gate under supply voltage scaling: (a) EDP and (b) LDP for 2-,3- and 4-input NCL gates; (c) EDP and (d) LDP for NCL gates with driving strengths X2, X4 and X13.

increase in transition and propagation delays as supply voltage is reduced. The increase gets more pronounced (see the slope of the charts) at the sub-threshold region, where a 100x-1000x increase can be observed. It is an intuitive notion that such increase in propagation delay has a direct impact on performance of a system. A metric that is not as intuitive but also has a big impact at system level is the increase in transition delay. This is because, at this level, as the slew in the output of a cell increases, the slew in the input of the next cell also increases, directly affecting its delay. It is obvious, sub-threshold operation should be avoided for high performance applications.

A similar behavior occurs regarding power consumption. Figures 4 (c) and (d) respectively show the leakage and EPO reductions due to supply voltage scaling. Again, when the supply voltage reaches the sub-threshold region, a higher decay on these figures is observed (10x-100x decrease), which demonstrates the potential gains of going sub-threshold when low power is a requirement. Analyzing delay and power metrics alone provides some intuition about their trade-offs. However, to better visualize them, LDP and EDP each combine delay and power in a single metric. The former is the product of leakage power and propagation delay, which provides a deeper insight in how the energy efficiency of the cell is impacted by VS in idle states. The latter is the product of EPO and propagation delay. Figures 4 (e) and (f) respectively show the LDP and EDP curves for the NCL2W11-of-2X2 gate. When $0.6V \leq V_{dd} \leq 1V$, power reduction is overcomes the delay increase. This translates in a decrease of LDP and EDP in this region. However, when $V_{dd} < 0.6V$, delay increases faster than power reduces, significantly increasing EDP and LDP. As a consequence, the minimum point for both curves is around $V_{dd} = 0.6V$. These results indicate that the best trade-off between performance and energy efficiency for this cell is in the near-threshold region.

The whole set of selected gates underwent the same evaluation. To assess the obtained results we divide the analysis in four main parts: (i) how the number of inputs influ-

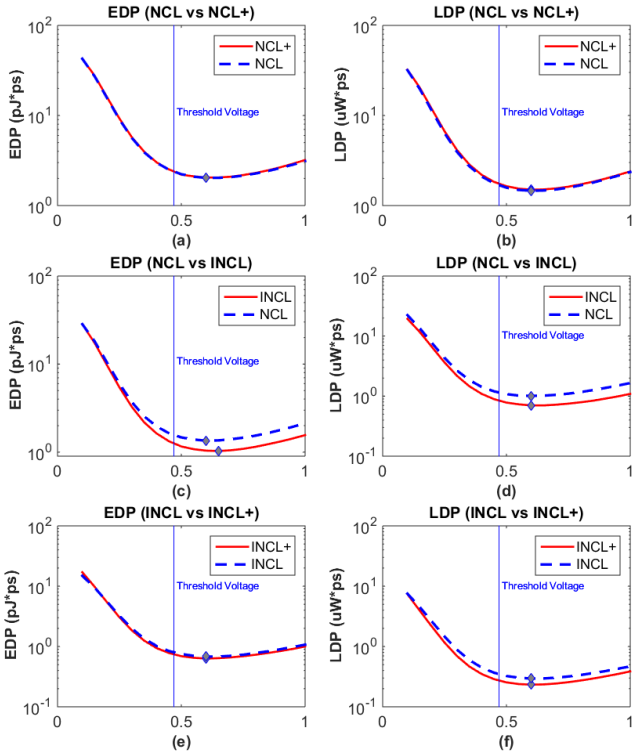


Figure 6: NCL function and families analysis. (a-c-e) EDP and (b-d-f) LDP of NCL and NCL+ gates (positive and negative unate).

ences the characteristics of the cells; (ii) how different driving strengths impact cells; (iii) how NCL+ cells behave as voltage is scaled and how they compare to NCL cells; and (iv) how negative unate cells are affected by VS and how they compare to positive unate cells. Figures 5 (a) and (b) represent part (i) of the analysis. There, we evaluate cells NCL2W11-of-2X4, NCL3W111-of-3X4 and NCL5W2211-of-4X4. As the charts show, increasing the number of inputs of the cells also compromises EDP and LDP figures. This is because as more transistors are stacked, signal integrity is compromised and balancing dynamic and leakage currents in low voltages gets tricky. The charts presented in Figures 5 (c) and (d) show part of (ii) the analysis respectively depicting EDP and LDP curves. This comparison is based on the measured parameters for cells NCL2W11-of-2X2, NCL2W11-of-2X4 and NCL2W11-of-2X13. As the charts show, increasing driving strength worsens EDP and LDP. This is because as transistors are enlarged, they get more leaky, which compromises energy efficiency of cells. Hence, VS applications should ideally employ small driving strength cells.

The charts presented in Figures 6 (a) and (b) show part (iii) of the analysis, using cells NCL5W2211-of-4X4 and NCL+5W2211-of-4X4. Accordingly, the minimum point of EDP and LDP of both NCL and NCL+ cells are in the near-threshold region (0.6V~0.65V) and all curves indicate a significant increase when supply voltage hits the sub-threshold region. Furthermore, NCL and NCL+ presented similar efficiency, which demonstrates that NCL+ is also suited for VS applications. Figures 6 (c)-(f) show the results for part (iv) of the analysis. To do so we rely on cells NCL2W11-of-2X4, INCL2W11-of-2X4, INCL1W11-of-2X4 and INCL+1W11-of-2X4. As the figure shows, negative unate functions present

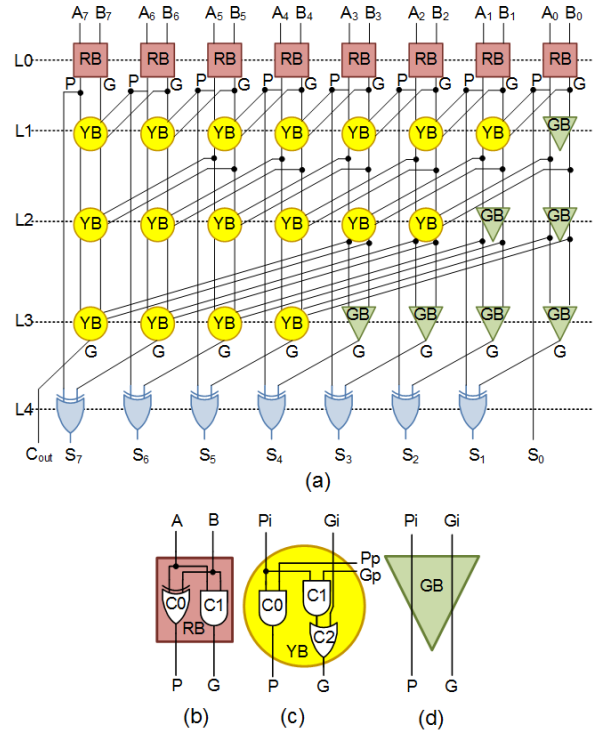


Figure 7: The 8-bit Kogge Stone adder and its basic blocks: (a) block diagram; (b) red box; (c) yellow box; (d) green box.

slightly worse EDP and LDP when compared to their positive unate counterparts. However, the increase is not substantial and negative unate cells presented variations in these figures similar to those observed for positive unate cells. These results motivate the use of SDDS-NCL for VS applications, as it relies on both NCL and NCL+ families employing either positive or negative unate functions.

4. AN ADDER CASE-STUDY

Another set of experiments allowed assessing the effects of voltage scaling on a case-study SDDS-NCL design. To do so, we designed an 8-bit Kogge-Stone adder and mapped it to the cells of the ASCEnD library. The adder is similar to the one described in [11] and its block diagram appears in Figure 7. Note that this block diagram represents its classic single-rail implementation, for the sake of simplicity. For the synthesis process this was translated to dual-rail and mapped using the approach proposed in [12].

The choice for this adder as a case-study was due to the fact that it is employed in a variety of real life applications. The choice for a small version (8 bits) target the reduction of synthesis and simulation time. Nevertheless, this does not compromise the generality of the results. In fact, in such an environment, all cells other than those fed by the primary inputs are driven by other NCL gates and scaling the voltage of the whole circuit impacts all cells at the same time. This is different from the controlled environment presented in the previous Section, where each cell was separately evaluated.

Synthesis relied on the method proposed in [12] and on the Cadence Framework. After synthesis, the tool automatically exported the mapped netlist to a Verilog description, which was then exported to a Spice description using Mentor Calibre. This description was employed in an analog

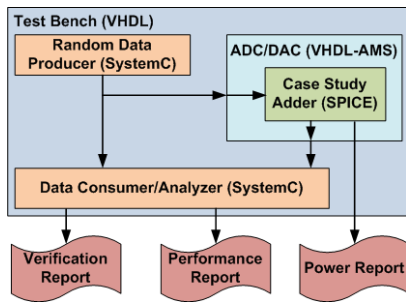


Figure 8: Environment for validation and performance and power analysis of the case study circuit.

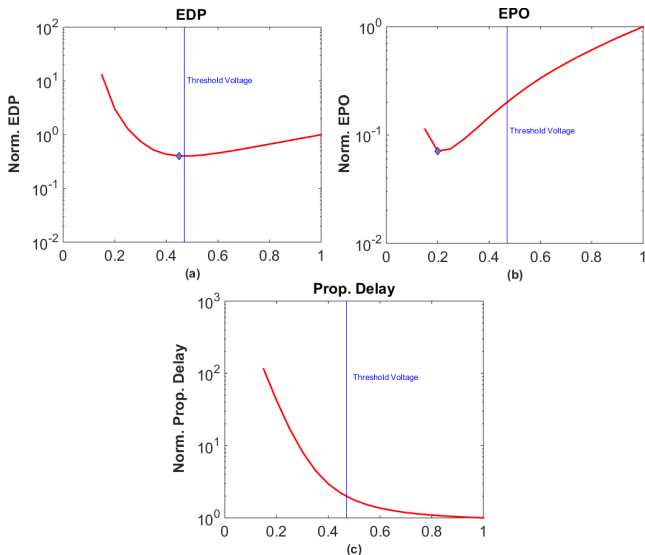


Figure 9: Results for the 8 bit Kogge-Stone case study: (a) energy delay product, (b) energy per operation and (c) propagation delay.

mixed signal environment for validation and analysis. As Figure 8 shows, this environment had a random data producer designed in SystemC to provide stimuli to the case study adder.

The generated outputs were computed by a data consumer and analyzer, also designed in SystemC. The interface between the Spice netlist of the case-study and the digital testbench was realized using analog to digital and digital to analog converters (ADC and DAC) described in VHDL-AMS. The environment relied on Cadence analog and digital simulators, Spectre and NCSim, respectively.

Figure 9 summarizes the obtained results. The charts display three characteristics of the circuit as a function of its supply voltage level: (a) EDP, (b) EPO and (c) latency. Note again, all results are normalized. The minimum operating voltage for the case study was 0.15V. Below this point, the circuit did not operate correctly, due to delay variations on the cells. Also, we observed during simulation that for voltage levels between 0.15V and 1V, performance and power scaled with voltage and correct functionality of the circuit was not compromised. As Figure 9(b) shows, EPO substantially reduces as the voltage is scaled until 0.2V, where it is 14x lower than at nominal voltage. However, at the minimum operating voltage (0.15V) there is a modification in the trend of the chart and it consumes more EPO

than at 0.2V. This is because at this point leakage power is highly significant and the cells have longer short circuit periods while switching, until the active PUN/PDN network overpowers the cut off PDN/PUN network. For ultra-low-power applications the 14x reduction on the Energy Per Operation (EPO) may seem attractive. However, as Figure 9 (c) shows, propagation delay increases dramatically for voltages below 0.5V. For example, at 0.2V propagation delay is increased by roughly 42x, which can be prohibitively constraining for some applications. In this context, the EDP metric allows an analysis of the sweet spot for defining an operating voltage taking into account both propagation delay and EPO. As Figure 9 (a) shows, this point is between 0.45V and 0.5V, in the near-threshold region.

At this point, the slope in propagation delay just started to increase substantially and the EPO is still reducing. Taking for example 0.45V supply voltage, the EDP is roughly 2.5x smaller than the EDP at 1V. This translates to a 5.4x reduction on EPO for a cost of 2.2x in propagation delay. Albeit the reduction in energy is not as large as that at minimum operating voltage, the costs in propagation delay are much smaller. These results corroborate the analyses presented in previous sections, validating the usage of SDDS-NCL for VS applications, and confirm that the most efficient delay and energy voltage for NCL and NCL+ cells (positive and negative unate) is in the near-threshold region. Therefore, if energy efficiency is the target, SDDS-NCL designers should focus on near-threshold design. However, if performance can be sacrificed for ultra low power operation, designers should target minimum voltage and the robustness of SDDS-NCL will accommodate delay variations gracefully, without compromising the functionality.

5. CONCLUSIONS AND FUTURE WORK

This article presented the first results obtained for designs using SDDS-NCL under VS. The authors believe the presented analysis is of big interest for the VLSI research community due to the following reasons: (i) low power design is a requirement that is increasingly common in IC design; (ii) VS was shown to be a very efficient approach to reduce the power of a system, but clocked designs impose huge challenges for VS applications; (iii) QDI templates are a promising asynchronous solution to ease and optimize VS application design; and (iv) SDDS-NCL provides a semi-custom and efficient design style for QDI circuits.

In this context, we highlight the following set of contributions and conclusions of this work: *a)* The evaluated cells were designed targeting a nominal voltage. Our results showed that without any optimizations, they could correctly operate at voltage levels as low as 15% of the nominal supply. This fact alone confirms the robustness of the cells and motivates their use for VS applications. Also, these results enabled a first validation of the use of the ASCeND library for VS applications; *b)* SDDS-NCL designs require both NCL and NCL+ families. Our experiments verified that NCL+ is as efficient as NCL for VS applications, indicating that both families can be used for that purpose. This was an important step, because these families employ different transistor arrangements and there was no such analysis available in the literature. Because these families are similarly efficient, they can be either used separately, in pure NCL or pure NCL+ designs, or jointly, in SDDS-NCL designs, without compromising power and performance trade offs; *c)* Another set of

experiments explored how negative unate cells behave under VS. We observed that they are not as efficient as positive unate ones. However, because they allow optimizations that can be explored in SDDS-NCL, their overhead for VS is not as substantial as the improvement they allow in the synthesis of SDDS-NCL circuits. Therefore, the results motivate the use of the full families of cells for SDDS-NCL design; *d*) The obtained results from the case-study points that the sub-threshold region brings energy reduction around 14x, but with a substantial delay increase, of 42x. This implies that, for SDDS-NCL designs, the sub-threshold region should be considered for applications that requires very limited power dissipation and with constraints that can accept large delays, such as sensor networks. However, if a more balanced trade-off between delay and energy is desirable, operation in near-threshold supply voltages should be considered. Accordingly, the case-study shows a 5.4x energy reduction and 2.2x delay increase in this supply region. Hence, the trade offs observed for near-threshold could be attractive to systems that implement a standby mode, which can reduce the supply voltage when the system is not in use, or when it is executing low priority tasks only. This allows the system to reduce energy consumption without sacrificing performance.

Another important point is that these first results enable us to look deeply into the effects of VS on SDDS-NCL and potential applications. Accordingly, as future work we will analyze physical effects, like delays on wires, and effects that can be significant on modern technologies, like noise, crosstalk and charge sharing. Moreover, we can now optimize cells for VS applications at the layout design phase. Hence, it is also part of future work the design of a library optimized for VS applications. We will also explore how to perform static timing analysis for different voltage levels. To do so, we plan to use the LiChEn tool and recharacterize the libraries to generate models that encompass VS. Finally, after such evaluation we consider validating a VS application design on a test chip.

6. ACKNOWLEDGEMENTS

This work was partially supported by CNPq, FAPERGS and CAPES. Authors acknowledge CNPq through processes no. 472126/2013-0, 401839/2013-3, 312556/2014-4 and 2001-47/2014-5), FAPERGS through process no. 302625/2012-7 and CAPES through process no. BEX 2129/14-0.

7. REFERENCES

- [1] P. Beerel *et al.*, *A Designer's Guide to Asynchronous VLSI*. Cambridge University Press, 2010.
- [2] C.-I. Jeong *et al.*, "Standard cell library design with voltage scaling and transistor sizing for ultra-low-power biomedical applications," in *EDSSC*, 2013, pp. 1–2.
- [3] I. J. Chang *et al.*, "Exploring asynchronous design techniques for process-tolerant and energy-efficient subthreshold operation," *IEEE JSSC*, vol. 45, no. 2, pp. 401–410, Feb 2010.
- [4] J. Kwong and A. Chandrakasan, "Variation-driven device sizing for minimum energy sub-threshold circuits," in *ISLPED*, 2006, pp. 8–13.
- [5] A. Martin and M. Nystrom, "Asynchronous techniques for system-on-chip design," *Proc. of the IEEE*, vol. 94, no. 6, pp. 1089–1120, 2006.
- [6] K. Fant and S. Brandt, "NULL Convention LogicTM: a complete and consistent logic for asynchronous digital circuit synthesis," in *ASAP*, 1996, pp. 261–273.
- [7] F. Parsan *et al.*, "Gate Mapping Automation for Asynchronous NULL Convention Logic Circuits," *IEEE TVLSI*, vol. 22, no. 1, pp. 99–112, 2013.
- [8] J. Cheoljoo and S. Nowick, "Technology mapping and cell merger for asynchronous threshold networks," *IEEE TCAD*, vol. 27, no. 4, pp. 659–672, 2008.
- [9] F. Parsan and S. Smith, "CMOS implementation comparison of NCL gates," in *MWSCAS*, 2012, pp. 394–397.
- [10] M. Moreira *et al.*, "NCL+: Return-to-One Null Convention Logic," in *MWSCAS*, 2013, pp. 836–839.
- [11] —, "Spatially Distributed Dual-Spacer Null Convention Logic Design," *JoLPE*, vol. 10, no. 3, pp. 313–320, 2014.
- [12] —, "Semi-custom NCL Design with Commercial EDA Frameworks: Is it Possible?" in *ASYNC*, May 2014, pp. 53–60.
- [13] R. Jorgenson *et al.*, "Ultralow-power operation in subthreshold regimes applying clockless logic," *Proc. of the IEEE*, vol. 98, no. 2, pp. 299–314, Feb 2010.
- [14] M. Moreira *et al.*, "A 65nm standard cell set and flow dedicated to automated asynchronous circuits design," in *SoCC*, 2011, pp. 99–103.
- [15] L. Zhou *et al.*, "Bit-Wise MTNCL: An ultra-low power bit-wise pipelined asynchronous circuit design methodology," in *MWSCAS*, Aug 2010, pp. 217–220.
- [16] X. Guan *et al.*, "Performance Analysis of Low Power Null Convention Logic Units with Power Cutoff," in *APWCS*, April 2010, pp. 55–58.
- [17] J. Wu and M. Choi, "Latency & Area Measurement and Optimization of Asynchronous Nanowire Crossbar System," in *I2MTC*, May 2010, pp. 1596–1601.
- [18] Y. Yang *et al.*, "A high-speed asynchronous array multiplier based on multi-threshold semi-static NULL convention logic pipeline," in *ASICON*, 2011, pp. 633–636.
- [19] F. Lodhi *et al.*, "Modified null convention logic pipeline to detect soft errors in both null and data phases," in *MWSCAS*, 2012, pp. 402–405.
- [20] M. Moreira *et al.*, "Return-to-One Protocol for Reducing Static Power in QDI Circuits Employing m-of-n Codes," in *SBCCI*, 2012.
- [21] —, "Lichen: Automated electrical characterization of asynchronous standard cell libraries," in *DSD*, 2013, pp. 933–940.