# VOLTAGE STACKING FOR SIMPLIFYING POWER MANAGEMENT IN ASYNCHRONOUS CIRCUITS

Andrew Suchanek[1], Zhong Chen[2] and Jia Di[1]

[1]Computer Science and Computer Engineering Department, University of Arkansas, Fayetteville, Arkansas, USA
[2]Electrical Engineering Department, University of Arkansas, Fayetteville, Arkansas, USA

## *ABSTRACT*

*Multiple power domains on a single integrated circuit (IC) are becoming more common due to the increasing complexity of systems on chips (SoCs) as process nodes continue to get smaller. Supplying the correct voltage to each domain requires the use of multiple voltage converters that occupy substantial space either on-chip or off-chip and introduce additional power loss in the conversions. In this paper, an asynchronous paradigm called Multi-Threshold NULL Convention Logic (MTNCL) is used to create a "stacked" architecture that reduces the number of converters needed and thereby mitigating the aforementioned problems. In this architecture, the MTNCL circuits are stacked between a multiple of $V_{DD}$ and GND, where simple control mechanisms alleviate the induced dynamic range fluctuation problem. The GLOBALFOUNDRIES 32nm Silicon-on-Insulator (SOI) CMOS process is used to evaluate and analyze the theoretical effects of parasitic extracted physical implementations in stacking different circuits running different workloads. These results show that the "stacked" architecture introduce negligible overhead compared to the operation of the individual circuits while substantially alleviating the need for voltage converters, which in turn reduces the overall power consumption of the system.*

## *KEYWORDS*

*Voltage stacking; power management; asynchronous; MTNCL*

## 1. INTRODUCTION

Semiconductor technology continues to evolve with newer digital integrated circuits (ICs) designed with smaller process nodes that run at lower voltages. These advances in technology introduce new challenges such as reducing energy consumption while maintaining performance and handling the additional power domains that result from more components being introduced on a single SoC. Therefore, more power supply rails of different voltages are required to power these circuit components [1]. To address this need, the off-chip power supply needs to be down-converted to multiple lower voltages, so each circuit component can reliably operate at their desired voltage level [2]. Voltage converters are traditionally implemented either on-chip or off-chip; but as the number of converters required increases, they become less than ideal because they not only take up substantial on-chip layout area and off-chip space, but also aggregate power loss during voltage conversions.

Prior research [2] takes two topologies of DC-DC converters and compares their performance when the input voltage is 1.5 V and the output voltage is 1.0 V. The first DC-DC converter used is an example of a linear regulator, called the low dropout (LDO) regulator. It is compared to an inductor-based switching regulator called the switched inductor (SI) buck regulator. The data in Table 1 shows clear tradeoffs between both types of DC-DC converters and such tradeoffs vary

DOI: 10.5121/vlsic.2018.9402 

depending on the power supply input and the required output. The efficiency, speed of delivery, area used, and complexity can be major drawbacks when designing the power delivery system and should be considered from time of concept in order to plan ahead for what the various components on-chip will be required in terms of power supply.

Table 1. Comparison table between LDO regulator and SI buck converter at $V_{in}$ = 1.5V and $V_{out}$ = 1.0 V [2].

| Parameter | LDO Regulator | SI buck converter |
|---|---|---|
| **Efficiency at light current load=100μA** | 16% | 13% |
| **Efficiency at heavy current load=5mA** | 60% | 78% |
| **Settling Time** | Less than 2μs | 97μs |
| **Area** | Small | Large |
| **Transient Response** | Fast | Slow |
| **Control Technique** | Simple | Complex |
| **Integration** | On-chip | Both but mostly off-chip |
| **Parameter** | LDO Regulator | SI buck converter |
| **Efficiency at light current load=100μA** | 16% | 13% |
| **Efficiency at heavy current load=5mA** | 60% | 78% |

The data above depicts the problems introduced by implementing on- or off-chip voltage converters. Therefore, instead of using many voltage converters, an alternative method for alleviating such problems was proposed in [3], which created a "stacked" architecture of low-voltage blocks. This methodology not only simplifies the off-chip and on-chip power delivery systems by decreasing the number of voltage converters, but also reduces the chip current draw.

This paper takes the theory proposed in [3] and continues the previous work that implements the voltage stacking methodology using an asynchronous *quasi-delay-insensitive* (QDI) paradigm named Multi-Threshold NULL Convention Logic (MTNCL). Stacking MTNCL circuits allows for a more comprehensive analysis due to MTNCL's robustness, timing flexibility, and minimized electromagnetic interference (EMI). Using MTNCL instead of synchronous logic, different workloads can be run on different types and sizes of circuits implemented in the "stacked" architecture. In addition, MTNCL circuitry has the ability to be put to sleep, which in this "stacked" infrastructure allows the other circuits to continue running while one or more sleeps. Moreover, by incorporating simple control logic, the designer will have the option to choose whether to allow for the non-sleeping circuit to speed up while other circuits sleep or to reduce energy consumption while maintaining acceptable performance.

This paper expands the previous work [4] by developing a physical design methodology in implementing the stacked architecture in the 32nm SOI CMOS process, extracting parasitics from the layouts, simulating the extracted netlists, and analyzing the results, while the work presented in [4] only had schematic designs and simulation results. By solving all challenges in the physical design process and optimizing the physical design steps for circuit stacking, the results and conclusions in this paper are more realistic and comprehensive compared to the previously simulated schematic models.

# 2. BACKGROUND

## 2.1. SYNCHRONOUS VOLTAGE STACKING

Research proposed in [3] uses voltage stacking for delivering $n$ times the normal operating voltage to $n$ circuits that are stacked upon one another in series. By increasing the voltage level required for delivery and decreasing the chip current draw, their research could potentially achieve several benefits:

1. $I^2R$ power loss is reduced by a factor of $n^2$ due to board and package resistance;
2. IR drop is reduced by a factor of $n^2$ because of the IR drop reduction by a factor of $n$ over *n-stacked* cores; and
3. The remaining voltage regulators and converters benefit from a lower step-down ratio that improves their efficiency and reduces their design complexity.

The work in [3] was carried out in a 150nm Fully Depleted Silicon-on-Insulator (FDSOI) CMOS process and implemented with synchronous logic. Although partially successful, it was only under certain strict constraints that the system would work properly. The voltage fluctuation between the different stacked layers caused a significant inner current mismatch that prevented the system from working properly if the cores were not the same circuitry or if the workload between the stacked cores varied by any significant amount. This was due to the strict setup and hold time requirements in the clocked synchronous paradigm. Therefore, although the theory has potential, using synchronous logic limits the ability to achieve any significant advantages across a wide scope. This is shown to be alleviated by using asynchronous logic and will be explained further below.

## 2.2. ASYNCHRONOUS CIRCUITS

Asynchronous circuits are sequential digital logic circuits that are clockless. There are two main asynchronous design styles: the *bounded-delay* model or the *delay-insensitive* (DI) model. Timing analysis is still required in the *bounded-delay* model because it is designed according to the worst-case propagation delay. The DI model is considered to be correct-by-construction because the delays in both logic elements and wires are assumed to be unbounded; but based upon how a circuit is designed, there can exist arbitrary gate and wire delays that could make the timing model too constrained for some practical circuits using the DI model [5]. Therefore, *quasi-delay-insensitive* (QDI) logic evolved from the DI model in the mid-1980s by separating the wires into critical and non-critical paths. QDI methodology is a commonly used practice in asynchronous circuit design since it is truly correct-by-construction and timing analysis is not required.

NULL Convention Logic (NCL) is a QDI asynchronous circuit design methodology that incorporates multi-rail logic [6-7]. Dual-rail NCL incorporates two rails that provide three valid states and one invalid state. The two rails are mutually exclusive of each other, meaning that both rails cannot be asserted at the same time, i.e., the invalid state. The other three states are the NULL state ($rail^0$ and $rail^1$ are both 0), the DATA0 state ($rail^0$=1, $rail^1$=0, equal to logic 0) and the DATA1 state ($rail^0$=0, $rail^1$=1, equal to logic 1). There are 27 fundamental NCL gates that make up the majority of their BOOLEAN equivalents. NCL architecture uses these gates and a method called hysteresis, where all of the inputs of a logic gate must be de-asserted before the output is able to return to a logic 0. This ensures that all of the gates in a combinational logic block propagate their data completely before returning to 0, which creates a NULL wave that separates the subsequent DATA waves.

The static implementation of an NCL gate is seen in Figure 1(a), where the logic required to output a 1 is in the set logic block. From there, the logic to keep the output high with hysteresis is

in the hold1 logic block. When all inputs have been returned to 0, the reset logic block turns on the NFET of the inverter connected to Z so that the output is pulled down to 0. The hold0 logic block keeps Z low until all inputs have arrived for the next DATA wave. On both ends of the NCL combinational logic blocks are sets of DI registers and completion detection logic that are used to create a single-stage or multi-staged pipelined designs. These sets use handshaking signals to communicate with one another to designate if they are expecting a DATA or NULL wave next.

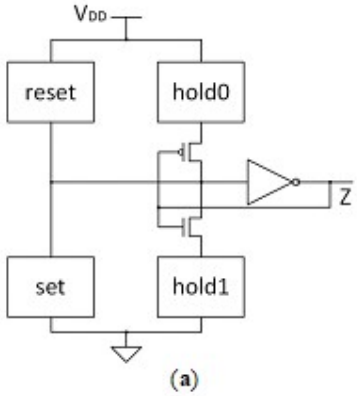


Figure 1. (a) NCL Static Gate Implementation;

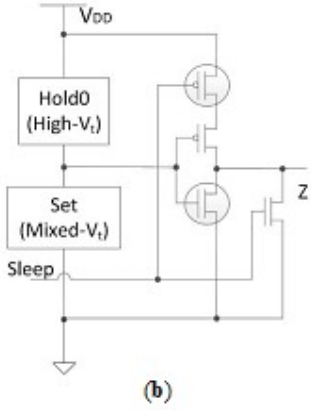


(b) MTNCL Static Gate Implementation.

Traditionally, when transistors are "turned off", the amount of current that continues to flow is more noticeable as the process node gets smaller. This is referred to as leakage current and the resulting leakage power associated with it has become more of an issue when considering overall power dissipation as technology continues to introduce smaller and smaller processes. In synchronous designs, Multi-Threshold CMOS (MTCMOS) is created to reduce the power dissipation by incorporating two different transistors, one with a higher threshold voltage ($V_t$) than the other. The High-$V_t$ transistors allow less leakage current to flow when the transistor is "turned off" but do so at a much slower switching speed. The Low-$V_t$ transistors are used when switching speed is crucial to the circuit's performance, but every critical path from power to ground and the output include a High-$V_t$ transistor to minimize leakage. If switching speed is not crucial, High-$V_t$ transistors are used to minimize the flow of leakage current and any unwanted power loss.

Multi-Threshold NULL Convention Logic (MTNCL) [7-15] is created by implementing MTCMOS power-gating in NCL. MTNCL uses both Low-$V_t$ and High-$V_t$ transistors and introduces a sleep function that replaces the functionality of hysteresis. The static gate implementation in Figure 1(b) shows that the Hold0 logic block of transistors do not need to be fast switching because the sleep signal actually forces the circuit to 0 when it is asserted, so they are all High-$V_t$ transistors. When the sleep signal is de-asserted the Set logic block of transistors uses mostly Low-$V_t$ transistors for faster switching speeds to assert a valid output. After a gate is asserted, it is de-asserted when the sleep signal is enabled because it serves as a power-gating mechanism. The sleep signal is provided by the completion logic from the stage one step ahead in the pipeline and just like in the NCL pipeline architecture, MTNCL has sets of DI registers and completion detection logic on both ends of any MTNCL combinational logic that use handshaking with the previous and future stages to request either a NULL or DATA wave.

## 3. MTNCL VOLTAGE STACKING

### 3.1. BASIC MTNCL VOLTAGE STACKING

The immediate benefit of voltage stacking is the reduction of the voltage converters on- or off-chip, the space/effort needed to implement them, and the power loss they introduce when operating. Simple MTNCL double-stacked and triple-stacked implementations are shown in Figure 2(a) and Figure 2(b), respectively, where the circuits are placed in series with one another and in parallel to a capacitor for each level of the stack.

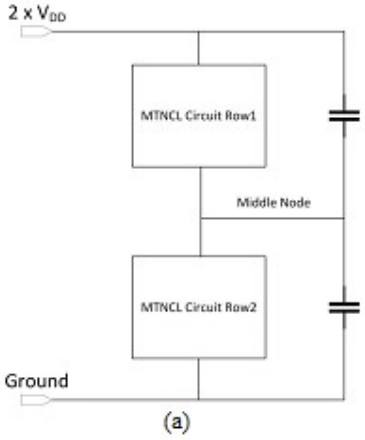


Figure 2. (a) Simple MTNCL Double Stacked Implementation;

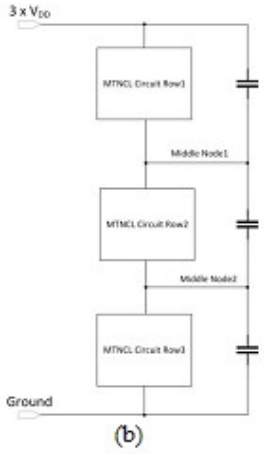

(b) Simple MTNCL Triple Stacked Implementation.

The capacitors function as bypass capacitors to ensure that the middle nodes remain oscillating at the correct multiple of VDD. The process used is the GLOBALFOUNDRIES 32nm SOI CMOS, where the designated transistor voltage is 0.9 V, so VDD will represent that voltage for this paper.

Previous work in [4] shows that when stacking the same circuits and running different workloads, the middle nodes oscillate in the desired range. It is by introducing the stacking of different circuits, mainly in size, that a fluctuation of the middle node's voltage is seen. Figure 3 shows the result of Double-Stacking an 11-bit by 7-bit Dadda Multiplier on top of an 11-bit Ripple Carry Adder (RCA). Since the multiplier (651 gates) is roughly four times larger than the RCA (141 gates), the middle node fluctuates towards the larger of the two at around 1 V to 1.1 V. However, due to the robustness of MTNCL, both circuits function correctly.

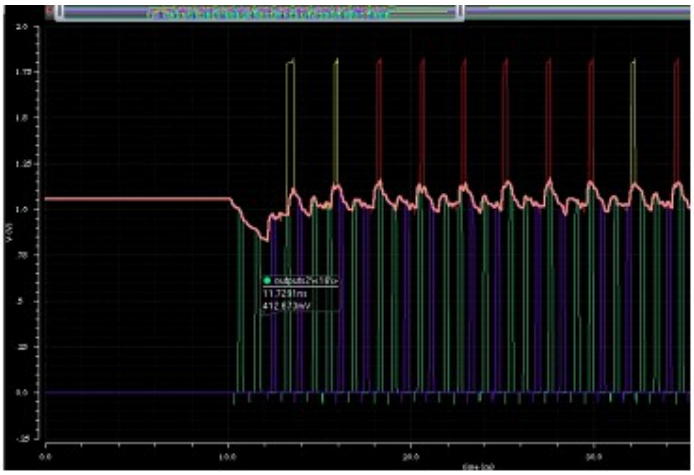

Figure 3. MTNCL Double-Stacked Simulation with Multiplier on Row 1 and RCA on Row

## 3.2. Advanced MTNCL Voltage Stacking

The effect of middle node shifting is much more severe when one circuit is put to sleep while the other one is still running. While the active MTNCL circuit operates correctly, its performance is substantially slowed down due to the reduced dynamic range (i.e., the range of the circuit's supply voltage, in between its VDD pin and GND pin). The Double-Stacked implementation will be used to illustrate how to mitigate this problem through implementing the advanced stacking structure, which introduces the Bypass and Awake transistors as shown in Figure 4.

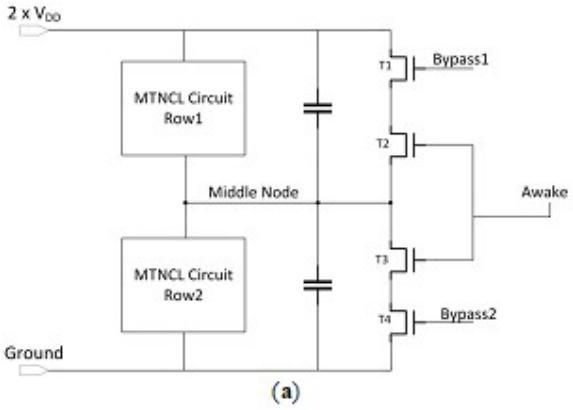


(a)

Figure 4. MTNCL Double-Stacked Circuits with Bypass and Awake Transistors.

Adding the Bypass and Awake transistors allows the designer to manipulate the middle node voltage and current flow. The additional logic is implemented in parallel to the MTNCL circuits in order to provide a separate path from one supply rail to the next that circumvents the idle circuit. This logic is controlled by separate signals generated at the system controller level indicating when the circuit is being put to sleep for an extended period of time. When either circuit is running, the *Awake* signal stays high, turning both the innermost transistors (T2 and T3) on. If either circuit is put to sleep for an extended period, their respective *Bypass* signal will also be enabled turning the transistor (T1 for row 1 or T4 for row 2) in the same row on and shorting either 2×VDD or GND to the middle node. Creating this path allows the middle node, which normally shifts drastically towards the circuit that is still running, to be pulled in the opposite direction, thereby increasing the dynamic range and speed for the working circuit. The *Awake* signal is set low (turning off transistors T2 and T3) when both circuits are put to sleep for an extended period, thereby blocking the direct path from 2×VDD to GND because the two *Bypass* signals are enabled. Table 2 shows the status of each transistor based upon what mode each circuit is currently in.

Table 2. Transistors T1-T4 status based upon circuits' modes

| Circuits Mode | T1 Status | T2 Status | T3 Status | T4 Status |
|---|---|---|---|---|
| **Both Circuits Active** | Off | On | On | Off |
| **Only Row 1 Circuit Active** | Off | On | On | On |
| **Only Row 2 Circuit Active** | On | On | On | Off |
| **Both Circuits Sleeping** | On | Off | Off | On |

Additionally, transistors T1-T4 can be sized differently to manipulate exactly how much the dynamic range is increased, which directly ties to the speed and energy consumption of the circuit still running. That is, if the circuit on row 1 is sleeping, sizing T1 and T2 will have a direct effect on the Energy Delay Product (EDP) of the circuit on row 2. The same is true for sizing T3 and T4 when the circuit on row 2 is sleeping. Manipulating the sizes of these transistors can produce faster speeds that consume more energy or slower speeds that conserve active energy, but there exists an ideal EDP for every combination. The following results use the multiplier and RCA stacked on top of one another and compare various transistor sizes. Figure 5 shows the RCA circuit on row 1 is active while the multiplier on row 2 is sleeping. Figure 6 shows the multiplier circuit on row 1 is active while the RCA is sleeping on row 2. Figures 7 and 8 show the results from the circuits being switched to the opposite rows. That is, both setups have the multiplier and the RCA active on row 2 while the opposite circuit is sleeping on row 1.

The plots of each figure show that there is an optimal EDP for each situation that will result in the lowest energy and delay combination. Thus, by adjusting the size of the transistors, T1-T4, the designer can actually manipulate the design and choose to increase the speed, conserve energy or simply find the best EDP. This allows for a much broader scope of choices when implementing this architecture based upon the circuit's desired functionality and performance.

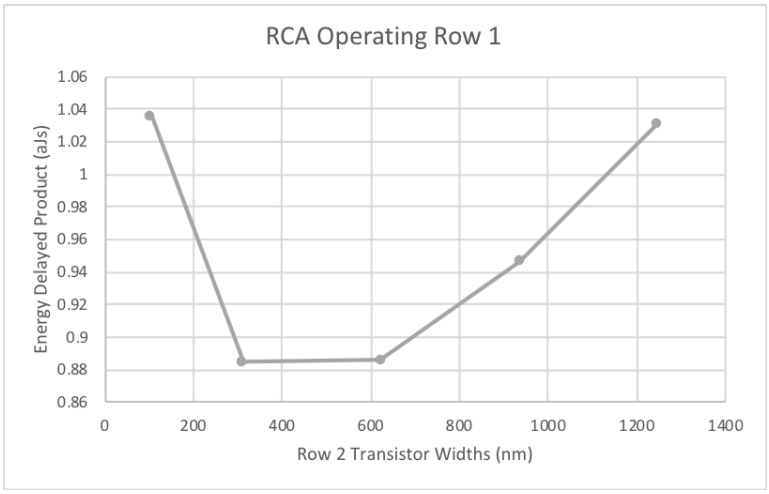


Figure 5. EDP Curve for the RCA Operating on Row 1 while the Multiplier is Sleeping on Row 2.

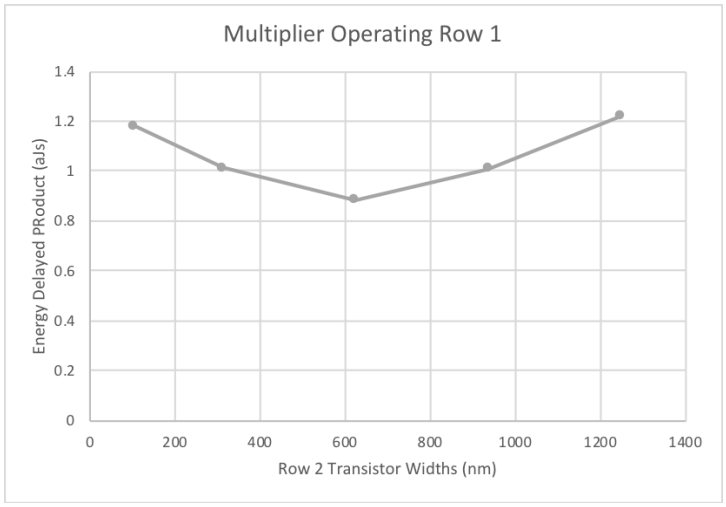


Figure 6. EDP Curve for the Multiplier Operating on Row 1 while the RCA is Sleeping on Row 2.

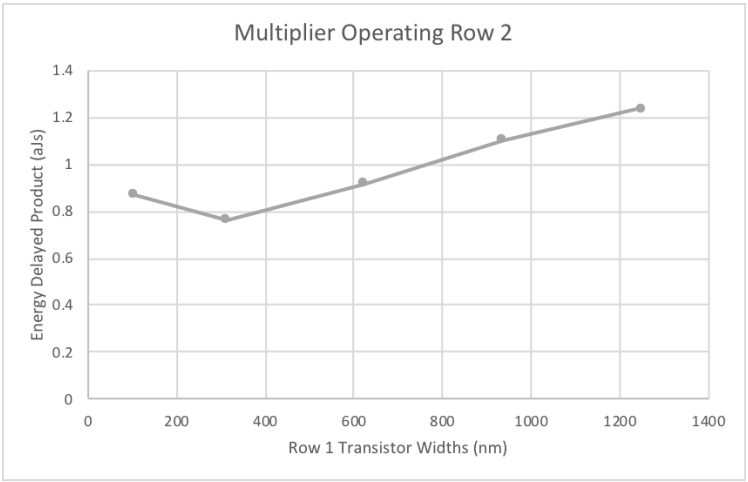


Figure 7. EDP Curve for the Multiplier Operating on Row 2 while the RCA is Sleeping on Row 2.

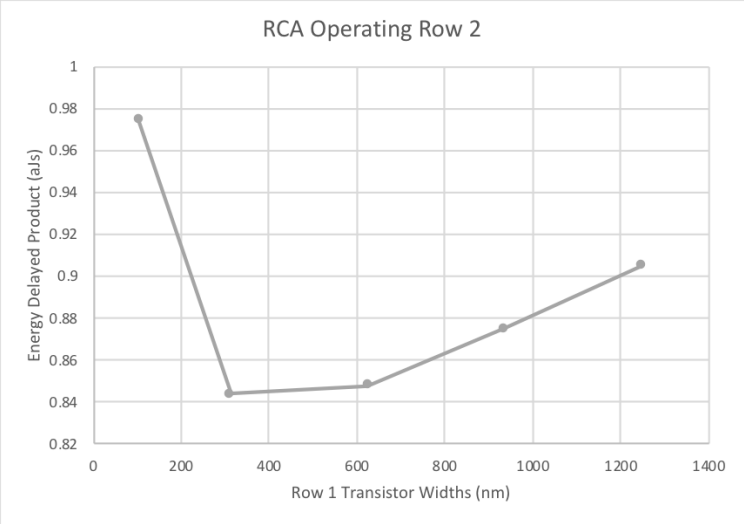


Figure 8. EDP Curve for the RCA Operating on Row 2 while the Multiplier is Sleeping on Row 1.

Another important thing to note are the execution times, energy consumed and resulting EDP of the circuits being stacked compared to their individual performances. Table 3 shows the results from stacking two identical multipliers on top of one another running different workloads and then comparing the results to how each multiplier runs their same respective workload on their own, that is, unstacked. Each data is measured after the multipliers complete processing 12 data inputs. Table 3 clearly shows that the overhead of the stacked architecture is negligible being only about 0.3%. Therefore, by stacking the asynchronous circuits on one another you lose nothing other than the need for voltage converters to shift the voltage down to the required levels. This not only alleviates power loss they incur from their lack of efficiency but removes their area from the IC and the time required to design and implement them.

Table 3. Energy Delay Product Results from Stacked and Unstacked Multipliers.

| Circuit Setup | Execution Time (ns) | Energy Consumption (pJ) | Energy Delayed Product (aJ*s) |
|---|---|---|---|
| **Individual Multiplier 1** | 33.65 | 20.64 | 0.694 |
| **Individual Multiplier 2** | 33.66 | 20.63 | 0.694 |
| **Total Unstacked** | N/A | N/A | 1.388 |
| **Both Running Stacked** | 33.66 | 41.33 | 1.391 |

## 4. PHYSICAL IMPLEMENTATION AND RESULTS

To demonstrate the feasibility of circuit stacking for IC designs and obtain more accurate results, the circuits being used, the multiplier and RCA, are laid out as physical implementations. Cadence Virtuoso is used in order to design the 32nm gate layouts, which are then place and routed in the Cadence Innovus tool. The circuits are place and routed individually and then reimported into Virtuoso, where they are stacked and wired together with the capacitors and additional logic (Bypass and Awake transistors). As a clear illustration, Figure 9(a) shows the cross-section view of two inverters in the 32nm SOI Process and Figure 9(b) shows the final clean implementation in Virtuoso of the Double-Stacked implementation with the RCA on row 1 and multiplier on row 2. Note that the capacitor is placed next to the logic for the clear view. For a real chip, certain types of capacitors (e.g., metal-insulator-metal or MIM) use higher level metals and can be placed on top of the logic circuits to save area.

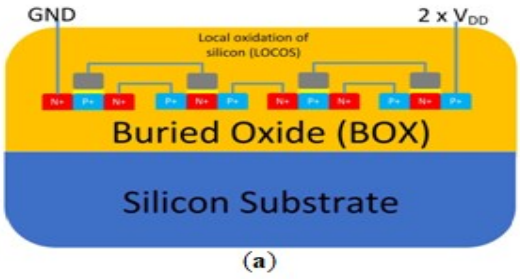


Figure 9. (a) Cross-Section View of Two Inverters Stacked;

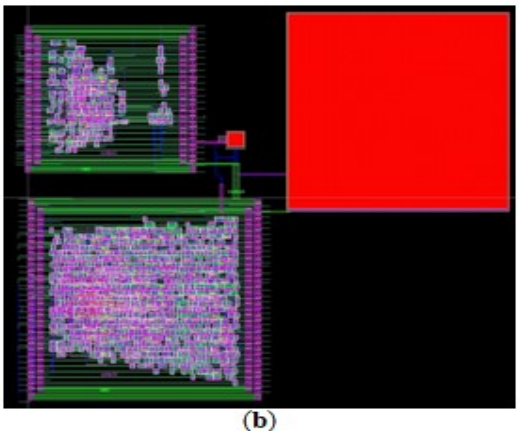


(b) RCA Row 1 and Multiplier Row 2 Design with Capacitors and Additional Logic in Cadence Virtuoso.

Like the design from Figure 9(b), the Double-Stacked implementation with the multiplier on Row 1 and RCA on Row 2 is created in Virtuoso so that both designs can be run through Mentor Graphics Calibre, a Parasitic Extraction (PEX) tool. PEX introduces all of the parasitic information, e.g., wire resistance and capacitance, coupling capacitance, etc. The resulting Spectre netlists are used to simulate both designs with these physical attributes in the Cadence Analog Design Environment (ADE). The first simulation implemented is comparing the EDP of the two stacked designs with both circuits running with the individual PEXed circuits running the same workloads. Table 4 shows that the EDP from the stacked architecture is slightly higher than the combination of the two individual circuit's EDP, but still very small (i.e., 4.8% and 7.6%, respectively). Note that each data in Table 4 is measured after the multiplier completes processing 10 data inputs and the RCA completes 24 data inputs.

Table 4. Energy Delay Product Results from Stacked and Unstacked Multipliers.

| Circuit Setup | Execution Time (ns) | Energy Consumption (pJ) | Energy Delayed Product (aJ*s) |
|---|---|---|---|
| **Individual Multiplier** | 27.3 | 31.03 | 0.8472 |
| **Individual RCA** | 44.1 | 15.4 | 0.6792 |
| **Total Unstacked** | N/A | N/A | 1.5264 |
| **Both Running Stacked Multiplier Row 1-RCA Row 2** | 33.73 | 47.41 | 1.599 |
| **Both Running Stacked RCA Row 1-Multiplier Row 2** | 34.6 | 47.49 | 1.643 |

When analyzed and compared to the pre-PEX simulation results, the numbers are larger due to the added parasitics, but the trend of the EDPs for the RCA and multiplier are similar. Since the multiplier is roughly four times the size of the RCA, accumulative current flow through the multiplier is larger. Thus, the Bypass and Awake transistors parallel to the RCA need to be sized larger to allow more current to flow through, which corresponds to an increase in the dynamic range of the multiplier while it is active. Conversely, the transistors parallel to the multiplier can be sized smaller to achieve the same increase in the dynamic range of the RCA while it is active. Figures 10 and 11 are comprised of simulations from the two post-PEX designs where the circuit on row 1 is operating and the circuit on row 2 is sleeping.


Figure 10. Post-PEX EDP Curve for the RCA Operating on Row 1 while the Multiplier is Sleeping on Row 2.

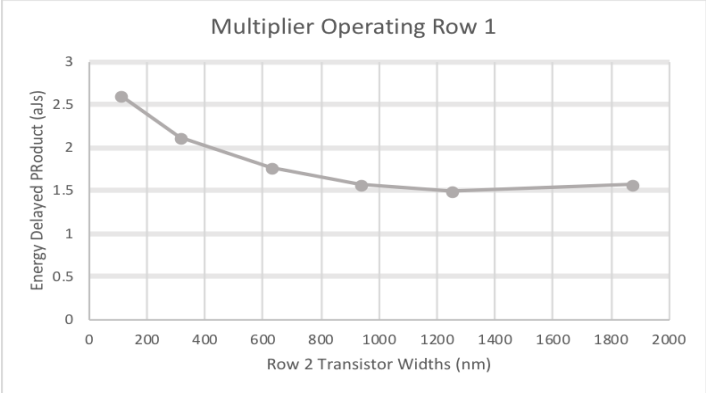


Figure 11. Post-PEX EDP Curve for the Multiplier Operating on Row 1 while the RCA is Sleeping on Row 2.

Similarly, Figures 12 and 13 are comprised of simulations from the two post-PEX designs where the circuit on row 1 is sleeping and the circuit on row 2 is operating.

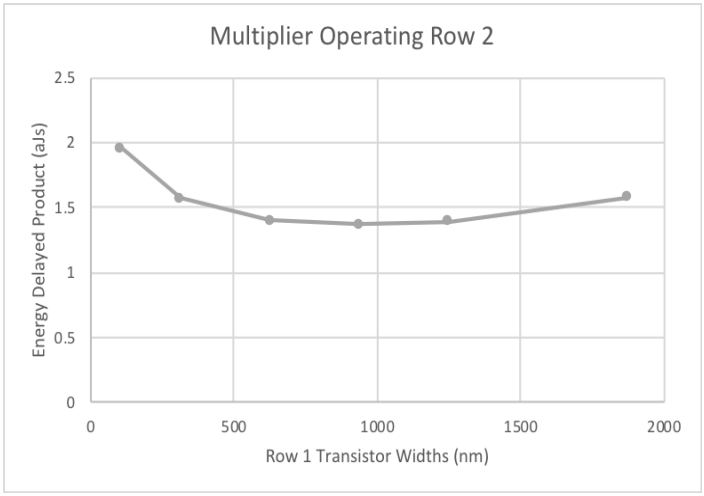


Figure 12. Post-PEX EDP Curve for the Multiplier Operating on Row 2 while the RCA is Sleeping on Row 1.

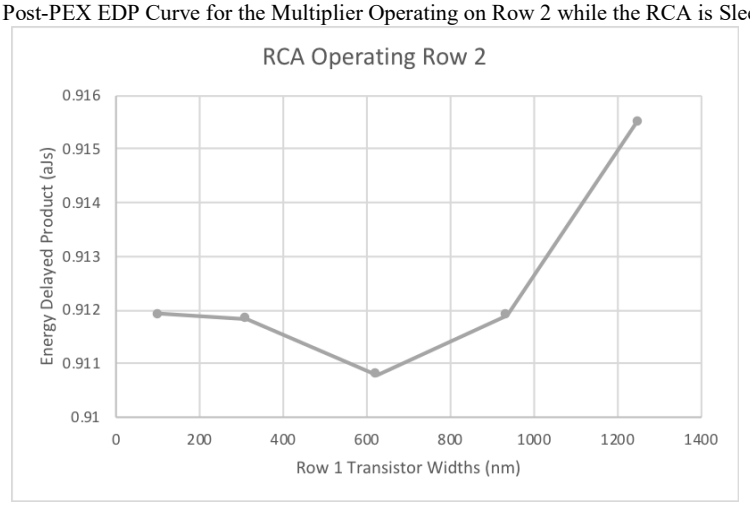


Figure 13. Post-PEX EDP Curve for the RCA Operating on Row 2 while the Multiplier is Sleeping on Row 1.

To show the impacts of the transistors' sizes on the dynamic range and speed of a circuit, the simulation waveforms of the multiplier operating on row 2 while the RCA on row 1 are shown in Figure 14 with 312nm channel width, and in Figure 15 with 1,872nm channel width. The latter is much more stable than the former.

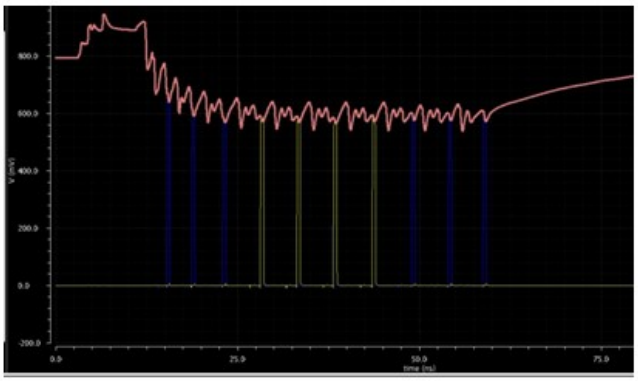

Figure 14. Post-PEX Multiplier Operating on Row 2 with Bypass and Awake Transistors' Widths Sized on Row 1 at 312 nm.

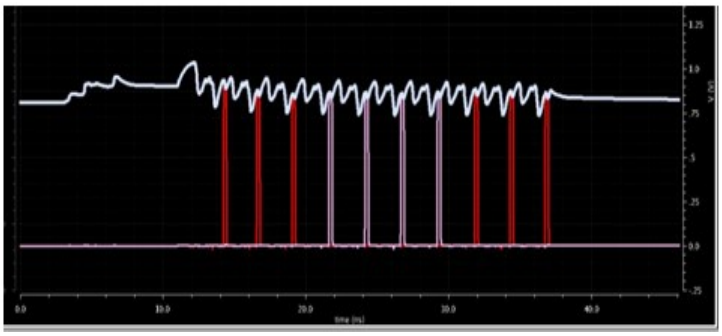

Figure 15. Post-PEX Multiplier Operating on Row 2 with Bypass and Awake Transistors' Widths Sized on Row 1 at 1,872 nm.

The overall analysis shows that simulating the MTNCL Double-Stacked architecture after implementing PEX follows the same trends that the schematic simulation models do. That is, when running either the multiplier stacked on top of the RCA or vice versa, the middle node will fluctuate towards the larger circuit. In addition, when putting either of the circuits to sleep for an extended period of time, the middle node can be manipulated using the additional logic. This ensures that the dynamic range of the active circuit remains at an acceptable level based upon the designer's energy and performance requirements of the overall system.

## 5. CONCLUSIONS

This paper presents a stacked architecture for the asynchronous MTNCL circuits, in order to reduce the number of voltage converters and eliminating the associated problems like additional area and power loss. The delay insensitivity of MTNCL allows the stacked circuits to tolerate dynamic range fluctuations due to the shifting of middle node's voltage, in spite of their different sizes and workloads. Simple logic (4 transistors) can be added and sized to provide the designer with the capability for balancing between energy and speed. The physical implementation and post-PEX simulations demonstrate the minimal overhead of the stacked architecture. Moreover, users are provided with options in balancing between faster speed and lower EDP, through manipulating the sizes of the transistor switches. The MTNCL stacked architecture has the potential to be incorporated into mixed-signal systems to raise the supply voltage of digital components to match or nearly match the supply voltages of analog/RF components. This would simplify the overall power management system design and save energy from the power sources, e.g., batteries.

## REFERENCES


[1] Jeff Falin. "Powering today's multi-rail FPGAs and DSPs, Part 1," *Analog Applications Journal* (1Q 2006).

[2] D. Kilani, B. Mohammad, H. Saleh and M. Ismail, "LDO regulator versus switched inductor DC-DC converter," *2014 21st IEEE International Conference on Electronics, Circuits and Systems (ICECS)*, Marseille, 2014, pp. 638-641.

[3] Sae Kyu Lee, David Brooks, and Gu-Yeon Wei. 2012. Evaluation of voltage stacking for near-threshold multicore computing. In *Proceedings of the 2012 ACM/IEEE international symposium on Low power electronics and design* (ISLPED '12). ACM, New York, NY, USA, 373-378.

[4] A. Suchanek, Z. Chen, and J. Di, "Asynchronous Circuit Stacking for Simplified Power Management," IEEE SoutheastCon 2018.

[5] Martin, Alain J. "The limitations to delay-insensitivity in asynchronous circuits." Springer New York, 1990.

[6] Karl M. Fant and Scott A. Brantd. "NULL Convention Logic: a complete and consistent logic for asynchronous digital circuit synthesis," in ASAP 96, Chicago, IL, Aug. 1996.
[7] S. C. Smith and J. Di, Designing Asynchronous Circuits using NULL Convention Logic (NCL), Morgan & Claypool Publishers, 2009.
[8] B. Bell, W. Bouillon, S. Li, E. Logal, and J. Di, "Application of Multi-Threshold NULL Convention Logic to Adaptive Beamforming Circuits for Ultra-Low Power," 2016 Government Microcircuit Applications & Critical Technology Conference (GOMACTech).
[9] L. Zhou, R. Parameswaran, F. A. Parsan, S. C. Smith, and J. Di, "Multi-Threshold NULL Convention Logic (MTNCL): An Ultra-Low Power Asynchronous Circuit Design Methodology," Journal of Low Power Electronics and Applications, vol. 5, issue 2, pp. 81-100, May 2015.
[10] L. Men and J. Di, "Asynchronous Parallel Platforms with Balanced Performance and Energy," Journal of Low Power Electronics, Vol. 10, No. 4, pp. 566-579, 2014.
[11] L. Men and J. Di, "An Asynchronous Finite Impulse Response Filter Design for Digital Signal Processing Unit," IEEE Midwest Symposium on Circuits and Systems, Aug. 2014.
[12] M. Hinds, B. Sparkman, J. Di, and S. C. Smith, "An Asynchronous Advanced Encryption Standard Core Design for Energy Efficiency," Journal of Low Power Electronics, Vol. 9, No. 2, pp. 175-188, 2013.
[13] L. Zhou, S. Smith, and J. Di, "Bit-Wise MTNCL: an Ultra-Low Power Bit-Wise Pipelined Asynchronous Circuit Design Methodology," 2010 IEEE Midwest Symposium on Circuits and Systems, pp. 217-220, 2010.
[14] A. Alzahrani, A. Bailey, G. Fu, and J. Di, "Glitch-Free Design for Multi-Threshold CMOS NCL Circuits," 2009 Great Lake Symposium on VLSI, pp. 215-220, 2009.
[15] A. Bailey, A. Al Zahrani, G. Fu, J. Di, and S. Smith, "Multi-Threshold Asynchronous Circuit Design for Ultra-Low Power," Journal of Low Power Electronics, Vol. 4, No. 3, pp. 337-348, 2008.

## AUTHORS

Andrew Suchanek received two B.S. degrees, one in Applied Mathematics and the other in Computer Engineering from the University of Arkansas in 2013, and his M.S. degree and Ph.D. in Computer Engineering from the University of Arkansas, in 2017 and 2018, respectively. His research area includes asynchronous integrated circuit design for low power, hardware security, and radiation hardening. He is a member of Tau Beta Pi and IEEE.

Zhong Chen holds a doctorate in electrical engineering from North Carolina State University. He worked for 7 years as ESD specialist at Texas Instruments prior to joining the University of Arkansas in 2015 as an Assistant Professor in Electrical Engineering. His research interests include nanofabrication, high-k materials, thin film deposition, silicon memory device, solid-state molecular memory device, etc.

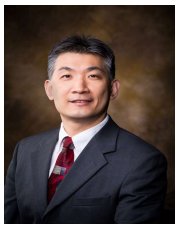

Jia Di received B.S. and M.S. degrees from Tsinghua University, China, in 1997 and 2000, respectively. He completed his Ph.D. in Electrical Engineering at the University of Central Florida in 2004. He joined the Computer Science and Computer Engineering Department of the University of Arkansas as an Assistant Professor in Fall 2004, where he is now a Professor and 21st Century Research Leadership Chair. His research area is asynchronous integrated circuit design and hardware security. Dr. Di has published one book and over 100 papers on technical journals and conferences. He also has 5 U.S. patents. Dr. Di is a senior member of IEEE and an elected member of the National Academy of Inventors.