

Ultralow-Power Operation in Subthreshold Regimes Applying Clockless Logic

New logic circuits promise to provide exceptional throughput and energy savings when the supply voltage is less than the threshold value needed to sustain normal operation.

By RYAN D. JORGENSEN, *Member IEEE*, LIEF SORENSEN, DAN LEET, *Member IEEE*, MICHAEL S. HAGEDORN, *Member IEEE*, DAVID R. LAMB, *Member IEEE*, THOMAS HAL FRIDDELL, *Member IEEE*, AND WARREN P. SNAPP

ABSTRACT | Energy performance requirements are causing designers of next-generation systems to explore approaches to lowest possible power consumption. Subthreshold operation is being examined to stretch low-power circuit designs beyond the normal modes of operation, with the potential for large energy savings. Some of the challenges to be overcome, like 10-100× performance penalties, are being addressed by research into parallelism. However, the uncertainty in timing generated by operating in subthreshold represents a major challenge to overcome. In this paper, first, we will introduce some background information on digital logic subthreshold operation, then provide some background on clockless logic design approaches giving a brief overview of some of the characteristics of the different design styles and focusing on NULL convention logic. Next, we will examine the application of that clockless logic approach to a military system, reviewing the background of the experiment, factors considered in the comparison, and then summarizing the results of the comparisons. Finally, an overview of additional research and development that will be needed to make the technique available to subthreshold designers is presented.

KEYWORDS | Asynchronous logic circuits; circuit modeling; circuit reliability; circuit simulation; circuits; CMOS digital integrated circuits; CMOS integrated circuits; combinational logic circuits; digital circuits; digital integrated circuits; digital systems; electronics; integrated circuit design; integrated circuit reliability; logic circuits; logic design; low-power design; military data processing; sequential logic circuits; ultralarge-scale integration

I. INTRODUCTION

Today's warfighter is becoming reliant upon more and more technological support to achieve strategic and tactical superiority over their enemies. Technological advantages include such items as night-vision equipment, unattended ground sensors, unmanned aerial vehicles, and body worn electronics. All of these systems need to be as small, lightweight, and low power as possible. This leads to the drive for high levels of integration with more and more dense circuits to achieve the minimum size, weight, and power possible.

Since batteries power most devices in the field, reducing power requirements can lead to large savings in the effective footprint of a device, enabling new capabilities that had previously been unattainable.

Two major requirements to deploying these small-size and low-power microsystems are: to exploit energy harvesting from the environment to eliminate or prolong the life of batteries; and to manage the power and performance of the system to fit within the available energy envelope of the power source. These requirements drive the need for novel approaches for dealing with ultralow-power environments, and demand the adoption of new approaches to circuit design.

Manuscript received July 18, 2009; revised September 10, 2009. Current version published January 20, 2010. This work was supported by the United States Air Force and DARPA under Contract FA8750-04-C-0007.

R. D. Jorgenson, L. Sorensen, D. Leet, M. S. Hagedorn, and D. R. Lamb are with Camgian Microsystems, Maitland, FL 32751 USA (e-mail: rjorgenson@camgian.com; lsorensen@camgian.com; dleet@camgian.com; mhagedorn@camgian.com; dlamb@camgian.com).

T. H. Friddell and W. P. Snapp are with The Boeing Company, Seattle, WA 98108 USA (e-mail: thomas.h.fridell@boeing.com; Warren.p.snapp@boeing.com).

Digital Object Identifier: 10.1109/JPROC.2009.2035449

Systems that drive these needs include next-generation unattended ground sensors or wireless sensor systems where the system needs to operate on limited battery or environmentally scavenged power for months or years. These systems can have complex power envelopes that vary widely as the system moves through its cycle of operation. Maximizing the node capabilities within the available energy supply requires a careful analysis of the sensor energy requirements as it moves through its operating cycle, and demands the use of dynamic power management techniques based on the state of the system. While Department of Defense (DoD) applications are the most demanding and challenging, similar needs and issues can be found in many homeland security and structural/environmental monitoring and safety applications.

The intelligence and capability to manage the power must be built into the system itself and includes features such as:

- an intelligent adaptive processing and control architecture that minimizes power consumption requirements by balancing processing capability versus power consumption as the system cycles through its operational states;
- precision power control at the transistor and circuit level that minimizes overall power consumption, while providing on-demand power domains for higher processing throughput when needed and subthreshold operation to minimize power consumption when processor demands are low;
- design techniques and cell libraries that optimize operation across the range from subthreshold to normal mode without major loss of performance;
- multiple on-chip power domains that efficiently support circuit operation across a wide voltage range;
- control of leakage currents in both logic and memory, particularly in the sleep modes where leakage currents can be the dominant energy loss for the system;
- optimized energy processing circuits that account for both the characteristics of the energy harvester and the desired range of voltage and current of the sensor node.

Digital subthreshold logic design, where the operating voltage is below the threshold voltage of the PMOS and NMOS transistors, holds promise for providing ultralow-power operation for these systems, but comes with significant challenges. The most notable example is the widely varying delay characteristics of the transistors across process, voltage, and temperature that force clocked designers to build in increasingly larger design safety margins in their timing, further reducing effective performance.

Clockless circuit design approaches, that replace the fixed timing assumption of the clock paradigm, offer alternative ways to harvest the potential of the subthreshold operation. Though there are many variations on clockless

circuit design approaches, this paper will focus on one in particular, NULL Convention Logic, that was used for the study reported and that is: suitable for application in the subthreshold regime; has been applied to low-power circuit design; has been proven through the fabrication of numerous ASICs of various complexity; and has a design approach geared towards producing circuits that are highly delay insensitive.

In this paper, first, we will introduce some background information on digital logic subthreshold operation, then provide some background on clockless logic design approaches giving a brief overview of some of the characteristics of the different design styles and focusing on NULL Convention Logic design, and presenting results of some of the fabricated chips built using that technology as an aid to illustrate features of the technology as they may apply to subthreshold design. Next, we will examine the application of that clockless logic approach to a military system, reviewing the background of the experiment, factors considered in the comparison, and then summarizing the results of the comparisons. Finally, an overview of additional research and development that will be needed to make the technique available to subthreshold designers is presented.

II. BACKGROUND

A. Digital Logic Subthreshold Operation

Operation of digital CMOS transistors in the subthreshold regime (where the VDD used in operation is below the threshold voltage of the PMOS and NMOS transistors) has proven to be very beneficial for energy constrained systems as it enables minimum energy consumption in logic circuits during active computation and reduces leakage currents in components that must be continuously powered. A good example of this can be found in [1]. This paper reported results on a subthreshold clocked Boolean logic (CBL) system on a chip (SoC) including a microcontroller and a subthreshold 128 kB SRAM powered by a switched capacitor dc–dc converter that delivered variable supply voltages from 0.3 V to 0.6 V. Reported clock speeds included 434 kHz at 500 mV and 1.04 MHz at 600 mV, with a reported minimum operating voltage of 300 mV.

Other examples showing some of the challenges and potential benefits to be had in subthreshold operation include [2]–[5].

In the kinds of sensor applications being addressed by the DoD, where long term operation is required in harsh conditions with limited power available, subthreshold operation is a desirable goal for the lowest energy consumption. There are sacrifices to be made, of course, including the 10–100× or more loss in throughput; but those can be addressed through parallelism, changes in the architecture, or multivoltage level operations in many applications [6].

However, in subthreshold operation the logic design must account for the exponential variation of subthreshold current with V_{th} variation. At 65 nm even a static CMOS logic style does not guarantee functionality in subthreshold. In subthreshold operation, transistor drive current variations increase by $10\times$ or greater, compared to nominal operation [7].

“Sub- V_{th} logic will likely play a key role in many future energy-efficient designs, **but designers must first dedicate all of their efforts to developing variability-resistant designs**” [2].

Process variations can randomly weaken pull-up or pull-down networks thus degrading noise margins of logic gates. In subthreshold registers, inverters with reduced output levels decrease the hold signal noise margin of data rates and affect data retention. Clock buffers with reduced output swing can cause contention thus impeding signal propagation. Gate delay variations can be 300% of nominal, causing major challenges for designers [8], who are left with the choice of designing conservatively and giving up performance gains from smaller geometry processes, or risking timing failures.

There are metal gate processes that can mitigate the variability caused by random dopant fluctuations but line edge roughness is still a factor and these two sources of variability dominate in advanced fabrication processes. As technology progresses to finer and finer geometries, design safety margins will increase.

All of these effects contribute to a very challenging set of problems for circuit designers to overcome: how to exploit subthreshold power advantages, without having to sacrifice all performance in order to guarantee functionality? Is there

a way to dynamically take advantage of the performance inherent in a design, when the exact behavior is not known until after fabrication, and indeed, until operation of the device in the specific environment? There are a variety of devices and logic families that have been proposed to deal with subthreshold operation, with various pros and cons [9]. However, one promising approach to solving this may lie in applying different digital design techniques, possibly borrowing from lessons learned in other low-power design arenas, such as hearing aids [10], [11].

B. Clockless Logic Overview

1) *Clockless Logic Approaches*: An emerging set of technologies capable of addressing the problems experienced in subthreshold design, is asynchronous or clockless logic. While a comprehensive review of clockless logic circuit design is beyond the scope of this paper, some introduction to the key features and characteristics of the most relevant approaches is warranted. More details in depth on the design of clockless logic circuits can be found in [12].

Clockless logic circuit designs are conceptually similar to CBL designs, in the sense that both circuits have *registers* for storing the inputs and results of a calculation and *computational elements* for transforming the data flowing in a circuit. In a CBL design, the sequencing of the data from register to register is controlled by a (usually) global signal, the clock. In clockless logic circuit design, the sequencing of the data from register to register through the computation elements is controlled by some other means, an *asynchronous control* (Fig. 1).

There are numerous approaches to designing without clocks, with various pros and cons depending on the design

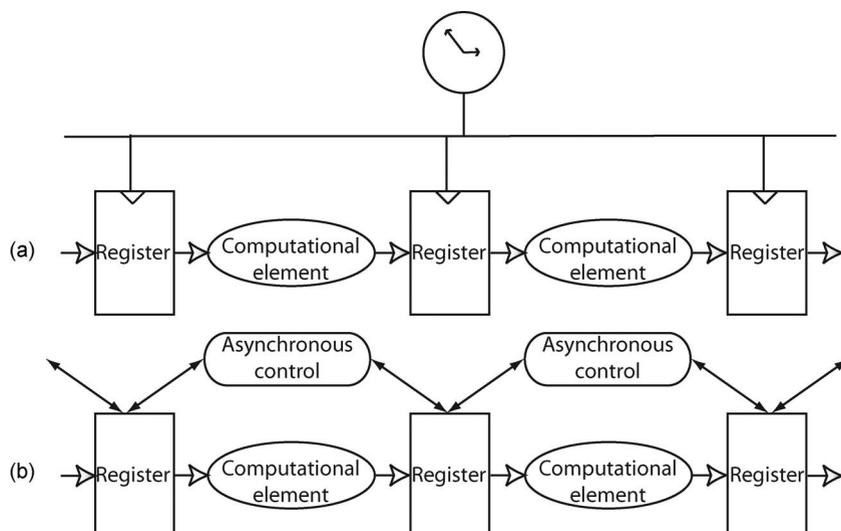


Fig. 1. (a) Comparison of CBL design structure, with (b) a general clockless logic circuit design structure.

style. Some of the major potential benefits include the following.

- Robust operation across PVT variations due to the elimination of the clock.
- Logically determined circuit design [13]. Circuits are designed to function independent of the timing assumptions normally inherent in CBL design approaches. This results in designs that will function over the entire operating range of the transistors.
- Power management with very low latency—no clock gating or restabilization issues on power-up. Circuits automatically go quiescent when they are waiting for data, and immediately resume operation upon receipt of that data. CBL designs need to carefully manage the gating of the clock(s) in order to achieve similar results, and have a latency issue upon receipt of new data.
- Low EMI and crosstalk [14]. Circuits designed using clockless techniques generally have much lower EMI crosstalk signatures than equivalent CBL designs. The lack of a single unifying control signal (the clock) enables the system to distribute the switching of the transistors in time, resulting in lowered EMI and lowered substrate noise signatures. For building systems, clockless logic circuit design's low EMI and low crosstalk characteristics will also be an advantage. Integrating low voltage or sensitive analog/RF components on the same substrate with large amounts of digital logic can be challenging due to the coherent noise generated by the simultaneous clocking of large numbers of circuit elements. Applying a clockless approach to this removes that global synchronization, reducing the noise and enabling much higher performance analog/RF circuitry on chip. This represents a major advantage for sensor systems that must reliably detect weak signals in a noisy and energy starved environment.
- Modular composition and delay insensitive interfacing. The ability for individual blocks to automatically self-synchronize their data rates permits the designer to concentrate on the logical structure of the data flow, and not crafting complicated timing relationships to ensure clock synchronization between regions operating at different voltage levels or speeds. Using the appropriate clockless logic circuit design technology, a designer can essentially design with as many voltage regimes as make sense, without having to derive multiple clock signals to accommodate all of the speed variations implied by multiple voltage domains. Indeed, this capability can be exploited to self-synchronize different circuit regions operating at different rates no matter what the cause; this can be used in what is sometimes called globally asynchronous locally synchronous (GALS) designs to

simplify complicated clock domain interfacing, even at normal operating voltage levels. GALS is becoming commonplace for complex ASIC designs where multiple clock domains make for easier design. In subthreshold, where multiple voltage domains (effectively clock domains) may be essential, the GALS concept is highly desirable.

Some of the significant potential drawbacks of clockless logic circuit design include the following.

- Complicated design approaches unfamiliar to CBL designers—many approaches develop new languages and design styles that add additional barriers to entry for designers not steeped in clockless logic circuit design practices.
- Lack of mainstream EDA tool support—existing EDA tools from mainstream suppliers are becoming more and more refined at synthesizing, inserting test support, simulating, and laying out clocked designs. Adapting these to clockless logic circuit design approaches is no small feat, but more work is being done in this area demonstrating the use of commercial EDA tools for clockless logic circuit designs [15].
- Area/performance penalties—although there are many examples of clockless logic circuit designs that are faster than their clocked equivalents [16]–[20], in many cases, clockless logic circuit designs incur a cost in terms of area or performance relative to the equivalent clocked functionality.

Clockless logic circuit designs can be grouped into four broad classes [21].

- Self-timed circuits rely upon bounded delay assumptions for gates and wires.
- Speed independent (SI) circuits assume that all wire delays are zero, but that gates can have arbitrary delays. This technique is not relevant to deep sub-micron circuit design where the gates operating at nominal are fast (zero delay) and the wires are slow.
- Quasi-delay insensitive (QDI) circuits relax the strict delay insensitive assumptions by making assumptions that signal branching in wires in some limited cases are assumed to have equal (or relatively matched) delays. These are termed *isochronic forks*.
- Delay insensitive (DI) circuits work correctly under arbitrary assumptions of delays of gates and wires. In practice, this results in uninteresting designs that are essentially static in nature with no decisions.

Of the four classes, only the self-timed and the QDI circuits will be described further; DI is in essence an academic curiosity, and SI circuits impose an assumption about zero wire delay that is increasingly invalidated in smaller geometry processes.

2) *Self-Timed Circuits*: Self-timed circuit designs are the closest conceptually to a CBL design. In this case, each

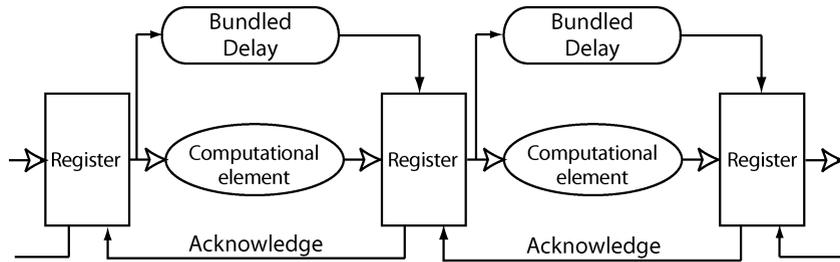


Fig. 2. Bundled data design.

datapath through a computational element has a matched timing signal to trigger the latching of the data. This is called “bundled delay,” because a data word is bundled with a matched timing signal, and the assumption is that the arrival of the timing signal is sufficient to guarantee that the entire bundle of data has been calculated and received. The arrival of the bundle signal is equivalent to the arrival of the clock edge. The receiving register sends an acknowledgement back to the originating register in order to close the control loop and prevent data overruns. This localized timing/control approach can be conceptually considered akin to very fine-grained clocking. A simplified block diagram of a bundled data approach is shown in Fig. 2.

The most commercially successful example of this design approach is from Handshake Solutions, who have produced a variety of commercial designs including an

8051 based microprocessor, and in collaboration with ARM, have developed a clockless ARM9 core. The ARM996HS processor automatically adapts to variations in environmental conditions (such as supply voltage, current, and temperature).

Fig. 3 [22] shows the results of simulations in a 0.13 μm process comparing a CBL (ARM968E-S) and clockless (ARM996HS) ARM core, with essentially equivalent gate count. As the figure shows, the clockless core has significantly lowered current peaks, as well as much lower electromagnetic emissions as a result of the distributed switching of the logic gates in time, inherent in all clockless logic circuit designs. The increased power efficiency of the clockless ARM9 (0.045 mW/MHz versus the CBL version’s 0.13 mW/MHz) is also a benefit.

However, there is a cost trade-off evident in this design case: the clockless core as reported has operating

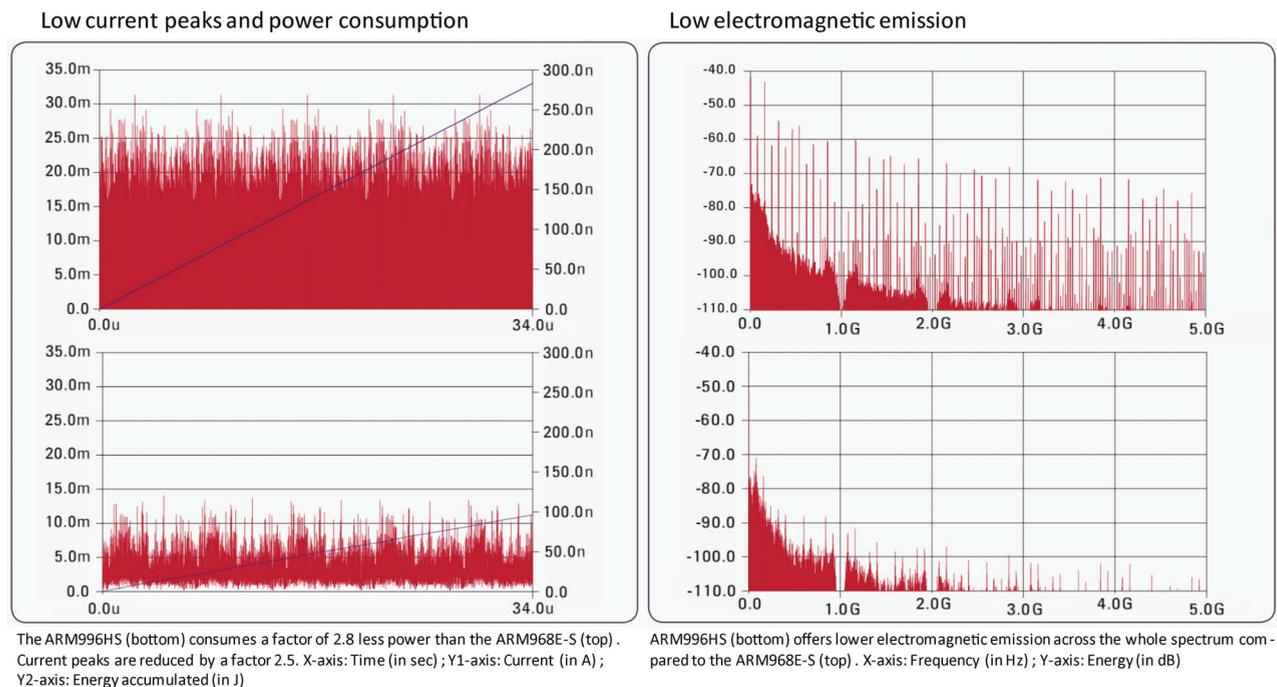


Fig. 3. Simulation comparison of a clockless and CBL ARM9 processor in 0.13 μm TSMC process.

performance equivalent to a CBL design of 77 MHz under nominal conditions, and equivalent to 50 MHz operation under worst case conditions. This compares to the CBL design's performance of 100 MHz across the entire operating range.

The assumption of the bundled delay becomes a weakness when looking at applying self-timed designs to subthreshold applications. In fact, in subthreshold operation, where process, voltage, and temperature (PVT) variations can cause exponential variations in subthreshold current, self-timed designs are at more of a disadvantage than a CBL approach. The bundled delays need to be precalculated and realized in silicon, forcing a very conservative timing approach in order to guarantee functionality. If the timing is not conservative enough, the circuit breaks, with no way to fix it without a respin of the IC. At least in a CBL design, the clock rate can usually be turned down after manufacturing to enable the circuit to function if the timing estimates were wrong.

3) *Quasi Delay Insensitive Circuits*: Quasi delay insensitive (QDI) approaches try to maintain as much delay insensitivity as possible, making them ideal choices for operating in unknown or wildly varying delay regimes, such as in subthreshold. Martin *et al.* [23] did some of the earliest work on QDI microprocessors, including the development of a clockless version of the MIPS R3000 processor [24]. This work and applications of the techniques have mainly focused on high performance circuit applications, where the goal was to deliver high-speed circuits, rather than lowest energy designs.

A broad overview and review of QDI in general was provided by Martin *et al.* in [25], and will not be repeated here. Instead, one version of QDI, NULL convention logic, will be considered.

4) *NULL Convention Logic*: NULL Convention Logic (NCL) [26], an example of a QDI approach, was selected for consideration based on a number of factors, including the ready availability of design tools based on familiar digital design approaches to develop the circuits (e.g., VHDL data entry capability), the expertise of the researchers in working with the technology, the proven ability of the technology to generate functional chips and the known robustness of the design approach to timing variations. The technology has been used to produce a number of clockless logic designs, ranging from small scale demonstrations to microcontrollers, where the emphasis has been on designing for power efficiency, at moderate performance levels, with low noise/EMI characteristics suitable for mixed-signal designs and desirable characteristics for an ultralow-power military electronics system.

a) *NCL operation*: NCL replaces the synchronizing function of the clock with two primary features.

- A delay insensitive data encoding that provides an unambiguous signal when DATA is available by

introducing the concept of the NULL state (no valid data present), resulting in a DATA-NULL cycle of operation. At its simplest, the encoding scheme can encode a bit as dual-rail, where one rail is energized to represent a DATA value of 1, and the other rail is energized to represent a DATA value of 0. Both rails energized is a logical error, and neither rail energized represents a NULL state. Alternative encoding schemes are possible (e.g., a rail is energized for the NULL state, and the nonenergized state corresponds to the DATA meaning).

- Completion detection circuitry, to watch the propagation of the DATA and NULL and to synchronize the communications of adjacent registration stages. In the case of a simple dual-rail encoded calculation, the completion circuitry verifies that one signal has been received for each bit value in the calculated DATA.

Fig. 4 shows the basic threshold gate symbology used for NCL, as well as a simple "combinational" example implementing a half adder function. NCL gates are threshold gates that "count" the number of inputs that are either DATA or NULL. In a two-valued system, such as is typical of CMOS design, one electrical level is assigned to represent each state (e.g., asserting an output at VDD is equivalent to the DATA state, and deasserting the output to VSS is the NULL state). An M-of-N threshold gate switches to DATA when M inputs are DATA and switches back to NULL when all N inputs are NULL. In the NCL half adder example all of the wires begin in a NULL state, and then DATA values are applied to the appropriate wires to represent the logical values for A and B. The thicker lines show the individual wires that would be energized for the case where variable A is a 1, or A is a 0 value; only one of the two wires would be set to the DATA state (typically VDD). As the gates reach their DATA threshold the system then propagates the values to the outputs, setting the appropriate wires for C(arry) and S(um). The NCL circuit will output the same truth table as the CBL circuit, but note that a time reference is required (depicted as a clock) for the CBL circuit to manage the input and output of data.

NCL integrates the data transformation and control functions into a single logic expression and produces circuits that are clockless, data driven, and effectively delay insensitive. An illustration of the integrated data/control flow is shown in Fig. 5. This figure shows an NCL circuit fragment in its conceptual parts. Initially, the circuit is in a NULL state (one of the two available electrical levels), except for the two "ACK" lines feeding the input and output registers, which are DATA (opening the registers). DATA (the opposite electrical level to the NULL state) then comes into the input register, propagates through the combinational circuitry and through the output register (shown by the bold arrows). This is then passed through the completion detection circuitry, which determines when a calculation has finished. The completion detection

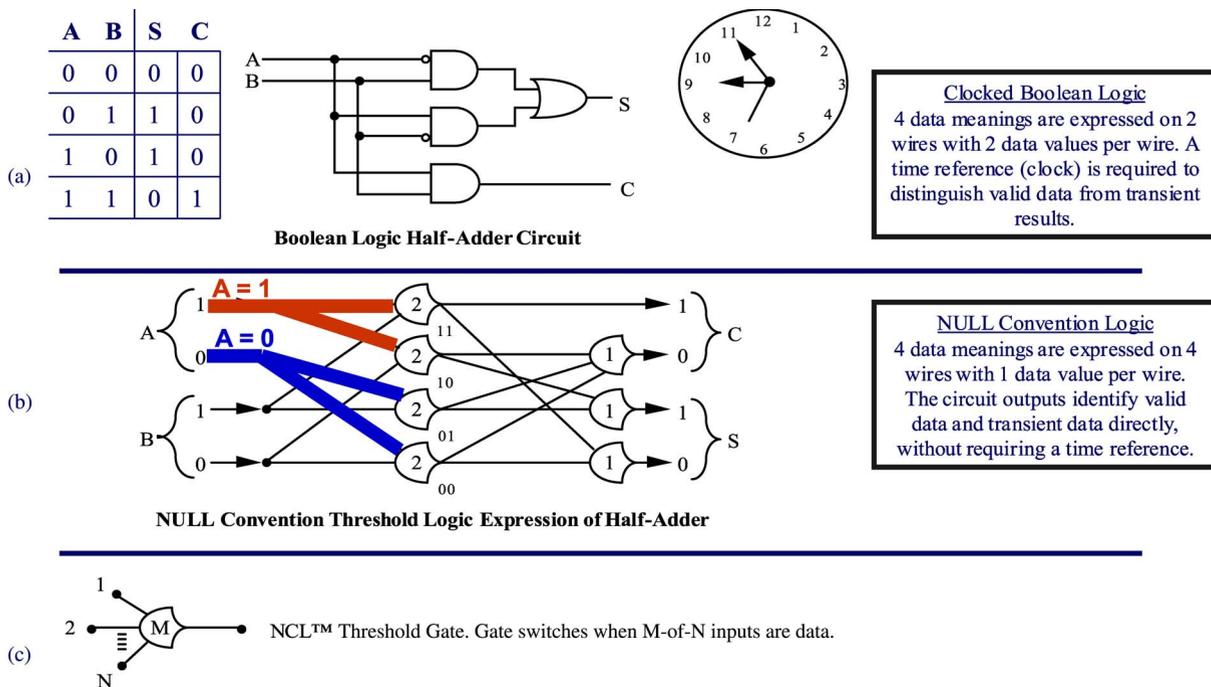


Fig. 4. NULL Convention Logic signaling protocol compared to CBL. (a) Truth table for a half adder function, and a combinational logic implementation with a time reference. (b) A NULL Convention Logic implementation of the same truth table. (c) NCL threshold gate symbology.

block then ACKnowledges the input register, and requests a NULL. The NULL wavefront then follows the path that DATA took, and “resets” the circuit to its initial (NULL) state.

NCL circuits run at the data rate and not under the control of a global clock. If no data is being presented to the circuits, they inherently go into a sleep mode and wait for the next data set to arrive. The recovery from the sleep mode occurs with very low latency because there are no phase locked loops and clocks to restabilize.

The robustness of the NCL circuits is a result of their average case behavior—a slow logic gate in an NCL circuit does not slow down all the rest of the logic gates (Fig. 6). Thus a slow gate may not slow the overall circuit behavior and does not necessarily result in a failure to meet throughput requirements. Even if the circuit does slow

down it will continue to propagate correct data and will provide a signal (the acknowledge) that can be monitored to determine that a performance error has occurred and corrective action should be taken. In comparison, a slow gate in a CBL circuit will result in the propagation of bad data to the output with no indication of a problem.

Little research has been done in examining the application of NCL to high throughput systems; the focus has been primarily on low-power medium throughput designs. The cost of the NULL/DATA cycle operating paradigm in terms of speed varies widely depending on the nature of the design being considered. The NCL circuits (for an equivalent pipelining depth) require propagation of 2× the number of wavefronts for each calculation. This is balanced against the setup, hold, and jitter margins needed by the CBL design. In simple designs, with long pipeline

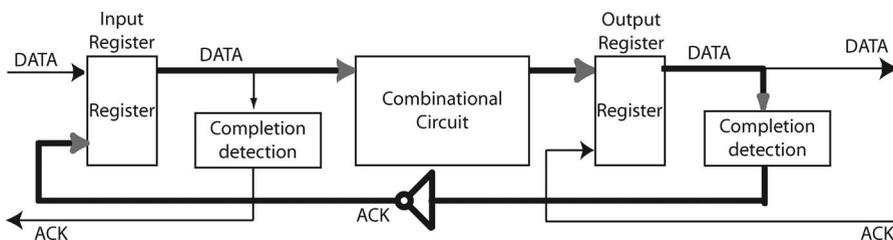


Fig. 5. NULL Convention Logic operation.

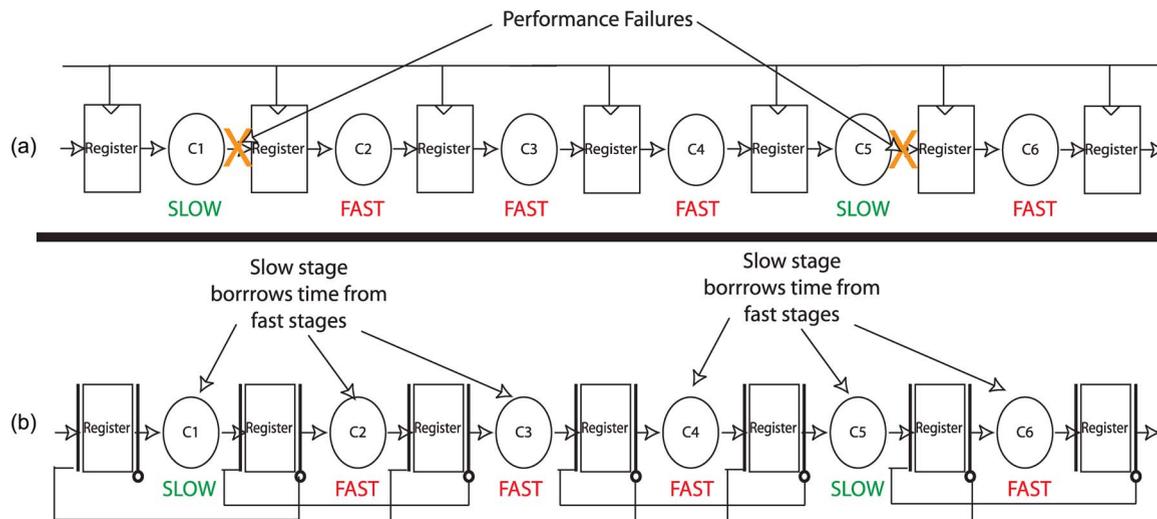


Fig. 6. (a) CBL circuits produce erroneous results if any one element fails to meet timing. (b) NCL circuits adjust to timing variations, either borrowing slack from faster stages to maintain the desired rate, or automatically adjusting the throughput to the maximum functional rate.

stages and simple clocking requirements, the penalty can be 2x. In more complex designs, timing margins for the clock grow, and CBL designs suffer more degradation over PVT variations than does the corresponding NCL design, thus the penalty decreases. In practice, the speed comparison is even more uncertain, since the clockless paradigm offers alternate architectures for solving the problem that are not available to CBL designs, potentially increasing the effective speed of the clockless designs; even when forced to use the same architecture, most times the speed issue is addressed by adding additional pipelining to the clockless design to compensate for the extra cycles. In any case, while a possible significant penalty at nominal operating voltage, the premise of ultralow-power (ULP) designs is that speed will be traded for lower power.

QDI circuit design approaches like NCL can offer intrinsic benefits of circuits that are data driven and essentially delay insensitive with no timing issues to break the functionality. This provides robust operation across a wide range of PVT variation.

b) *NCL EDA support*: NCL designs can be captured using VHDL, a language familiar to CBL designers, helping

to reduce the barrier to adoption of clockless logic. Place and route, and dynamic simulation can all be accomplished using commercially available tools.

NCL circuit designs have approaches to circuit test developed, including automated scan insertion and built-in self test. However, the EDA support for these approaches is not as comprehensive as that available for CBL designers. Basic approaches have been shown to operate both in industry and in academia, and the principles demonstrated on simple systems, using techniques to adapt CBL tools to the job, or based on custom tools developed specifically for the purpose [27], [28].

c) *NCL circuit characteristics*: NULL convention logic designs typically see a penalty in terms of area when compared to CBL designs. Table 1 summarizes a comparison of an NCL version of a small microcontroller versus two flavors of CBL designs (two different synthesis runs of the CBL core trading off area versus energy efficiency), fabricated in 0.25 μm CMOS. As can be seen, the NCL design occupies approximately 1.6–1.7x the area of the CBL designs. The transistor count difference is in the range of approximately 1.14–1.21x that of the CBL designs.

Table 1 Comparison of NCL and Two CBL Microcontroller Designs

Design	Area (mm ²)		# of Cell Types	Total Cells	Total Transistors	Energy (mW/ MHz)	Processor Speed
	Cells	Routed					
NCL8051	249,679	295,935	162	5,698	58,753	47	36.5 MHz (effective)
DW8051A	153,867	179,772	72	3,410	51,482	157.4	36.5 MHz (actual)
DW8051B	145,699	169,740	123	3,305	48,482	182.4	36.5 MHz (actual)

(Energy and performance results were simulated with PowerMill™)

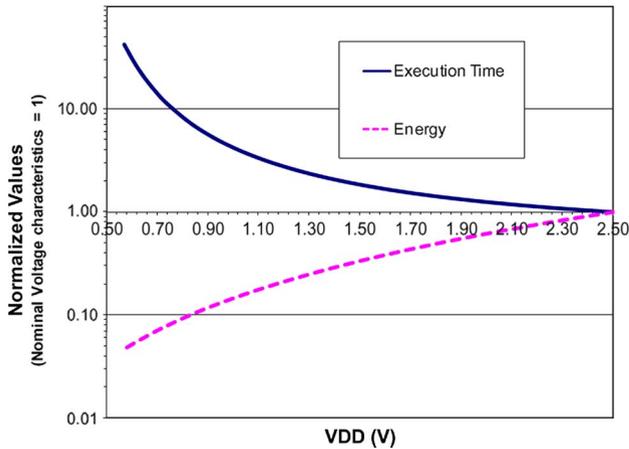


Fig. 7. Clockless microcontroller operation as a function of power supply voltage.

While area penalties are an obvious disadvantage for cost sensitive commercial applications, it is not likely to be a significant disadvantage for demanding military/aerospace (or even commercial) applications where power considerations are critical. As the table shows, the clockless design provides more than a 3× savings in power/MHz (or energy) versus the clocked designs.

Fig. 7 shows the operation of a clockless microcontroller as a function of power supply voltage. The graph shows the plot of the energy consumption and execution

times of the microcontroller normalized to 1 at the nominal VDD of 2.5 V. It can be seen that the microcontroller operates continuously over the entire measured voltage range, even as the system performance varies over a 40× range. In practice, the core functioned up to at least 3.3 V, but that was beyond the recommended voltage levels for the process (risking destruction of the transistors). Operation over such wide voltage ranges is particularly attractive for battery powered unattended systems.

In another study a test chip was designed and fabricated with both analog and digital circuits so that the impact of the digitally introduced substrate noise on the analog circuits could be directly measured. NCL and CBL pseudo-random number generators were designed and implemented with an interdigitated layout. The circuits were then placed into operation, and the substrate induced crosstalk generated by the operation of these circuits on the analog circuits was measured and compared. As can be seen (Fig. 8), the NCL implementation demonstrated 25 dBm lower coupled noise compared with the equivalent CBL design.

The NCL approach, when compared to equivalent CBL designs, provides lower power, lower EMI, lower substrate noise (enabling easier integration of sensitive analog or RF circuits), and reliable, maximum throughput delivered automatically based on the environmental conditions (temperature, voltage, process), rather than throughput being artificially constrained to worst-case conditions as it is in a CBL system. This latter characteristic, the reliability of the functionality derived from logically determined

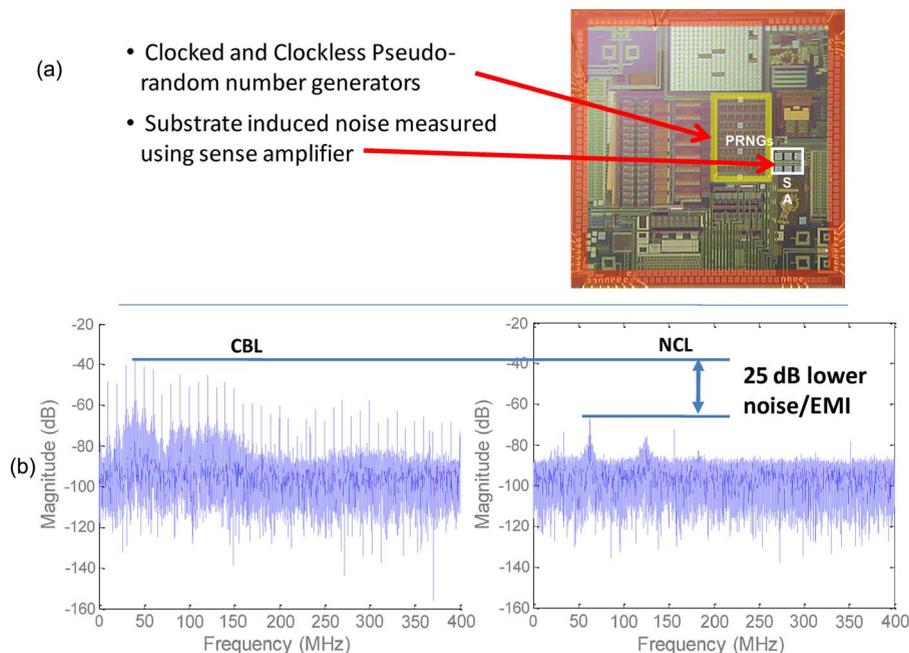


Fig. 8. (a) Low noise characteristics of NCL showing the photomicrograph of the test circuit. (b) Measured substrate induced noise for the CBL and NCL circuits.

system design, is what enables NCL systems to adapt and achieve maximal performance even in subthreshold operation.

In the next section, we perform some analysis and simulation, projecting the application of this technology to DoD systems operating in subthreshold, using a test case based on a 65 nm CMOS process.

III. APPLYING CLOCKLESS LOGIC TO REAL SYSTEMS

A. Setting the Stage

To investigate the potential value of a clockless logic approach to subthreshold digital design in the context of military ASICs, the authors chose to apply both analytical and simulation approaches. A set of standard cells were created for both the NCL and CBL approaches in a 65 nm bulk CMOS process, and the effects of subthreshold PVT variations were modeled. A set of 5-tap FIR filters were then designed and simulated across PVT variations as well. These results were then used to scale up to projections of the effects on larger circuits and to estimate the performance of each design approach under circumstances consistent with military electronics.

B. Factors in the Comparison

In performing the comparisons of the NCL design and the CBL design, the comparison was between subthreshold NCL and subthreshold CBL without the ability to adapt to temperature, with both circuits designed for a 65 nm process. For NCL it was assumed that either power shutoff (PSO) or adaptive voltage scaling (AVS) making use of NCL's integral completion signals could be used to save energy when NCL completes the computations ahead of time. AVS is preferable to PSO because it addresses both dynamic and static energy whereas PSO only eliminates leakage energy and PSO requires state save/re-initialization.

CBL clock rate adaptation over temperature was not considered in the comparison. Rough adaptation to temperature is sometimes done at the chip-level in circuits, (e.g., as in high end microprocessors) but this is done globally. It is rough because there are always measurement errors due to the difference in location between the temperature sensor and logic gates under control and control errors due to errors in modeling speed vs. temperature, variability sensitivity vs. temperature, power supply response time vs. logic speed, speed vs. VDD, and the need to maintain timing closure as the clock and/or power supply voltage are changed. These errors cannot easily be removed so although compensation is sometimes done it is never an exact match to the critical path performance as it is with NCL. Other techniques, e.g., Razor [29] (a form of AVS), can compensate over temperature as does NCL, but Razor has drawbacks including its own control system complexity and the large margin ($2\times$) needed to guarantee

fixed computational rate in the presence of more than 1/2 clock slack failures. Nevertheless, the circuit does adapt and can save energy, but it was excluded in the comparison.

The following factors were considered in the comparison and analysis, in an effort to make a projection of the qualitative benefits that might be realized from an NCL design versus a CBL design approach. These factors overlap so they must be considered together as a whole rather than independently. The percentages associated with each of these factors are not independent and are not multiplicative.

1) Operational temperature history and mission duration:

With a simulated chip at room temperature NCL and CBL operated at nearly the same Energy per Operation (E/Op). There are no savings when operating at the low temperature. In fact in the implementation we simulated NCL was $5.5\times$ higher energy per operation E/Op at -40°C , but very little time is spent at this temperature even in a cold environment due to self-heating.

As temperature rises, NCL energy savings accumulates over time. The operational history and its profile across the clock tree is therefore important. Since systems usually start out cold and warm up over time, most ICs spend the majority of their time at higher temperature rather than the minimum that they must operate at. As the temperature stabilizes, the total energy savings of NCL versus CBL will continue to grow, converging to a constant performance advantage versus the CBL design.

Some missions are shorter but many missions last much longer than the 2.7 hr time curve considered. Unattended Aerial Vehicle (UAV) missions can last between 1 and 20 days for long endurance UAVs. Since the energy savings accumulates over time, the longer the mission the more energy is saved. From the viewpoint of mission duration, *the comparison is pessimistic and the lower value of the energy savings range was assigned a factor of $7\times$* . If the temperature profile had started at room temperature, the factor would have been more like $7.4\times$. If the mission had been twice as long the savings would have been $8.7\times$.

2) Packaging and total power dissipation:

A 20°C temperature rise was factored into the mission E/Op savings projection based on an assumption of a typical package and total power dissipation. Since circuits operate faster in this process as temperature increases, there is a speed advantage to operating the circuit fast and the designer would select packaging and design to operate at higher temperatures. *High performance designs operating near thermal limits achieve a $32\times$ E/Op*

advantage with NCL over CBL and the upper energy savings was assigned the value of $30 \times$ E/Op savings.

- 3) **On-chip temperature profiles:** Most chips have a nonuniform temperature profile across the span of the clock tree. The temperature difference is typically between 0°C and 30°C . The NCL simulation we did showed a $3 \times$ lower E/Op between 55°C and 25°C . This places a minimum value on the ratio of NCL E/Op to CBL E/Op because CBL cannot be compensated within the clock tree, whereas, NCL has local completion signals that could be used for local PSO.
- 4) **Compatibility with paralleling and pipelining:** Uncorrelated device-to-device variations cannot be compensated with clock-based logic that emulates the computational path. A measurement of the computational path itself must be made and NCL does this. The impact this has on E/Op for designs implemented in nanometer processes is that CBL must maintain margin, and as the register-to-register delay decreases with added paralleling and pipelining, the throughput quickly becomes limited by the minimum timing margin. There is also a significant global clock tree E/Op cost associated with both paralleling and pipelining, whereas NCL does not drive long clock lines, therefore, a CBL circuit is limited in its depth of pipelining and paralleling well before there is a single gate between registers. NCL can be more heavily pipelined and paralleled. In a process with $2 \times$ the variability of the 65 nm process, with heavy pipelining, NCL can achieve $3 \times$ the speed of CBL and this combined with power shutoff would result in $3 \times$ lower E/Op than CBL, hence the upper range of $7 \times$ (from item 1) times $3 \times = 21 \times$ and this was rounded down and listed as the upper savings for most military projects.
- 5) **Compatibility with power shutoff (PSO) or adaptive voltage scaling (AVS):** The power savings with NCL is predicated upon the use of either PSO or AVS. Without the ability to either significantly reduce the leakage (PSO), or to take advantage of the improved energy efficiency of matching the execution rate to the data rate (AVS), the NCL system becomes constrained by its environment and would not save energy over the CBL system in this analysis. The estimated savings was based on PSO to be conservative and this assumes some form of reinitialization capability or state savings [30] for the NCL pipeline. Reinitialization and reset requires a small amount of additional time and not all designs have enough slack to accommodate it. However, since subthreshold throughput changes orders of magnitude over temperature, it is highly likely that timing margins would allow PSO implementation for nearly all designs that operate in subthreshold.

AVS, since it addresses both the dynamic energy and the static energy consumption, produces even higher savings in a dynamic energy dominated E/Op budget. Most systems, even at optimum VDD in subthreshold operate with dynamic energy dominating. A factor of $6 \times$ E/Op savings can be achieved over the subthreshold VDD range using AVS based on a comparison of low subthreshold VDD E/Op and near subthreshold VDD E/Op. AVS requires power islands, sensing, and high efficiency power conversion all of which increase the complexity of the design. For small designs or designs that would not switch or regulate a high percentage of the total power budget in this way, the benefit may not justify the added complexity.

- 6) **Expected percentage implemented at each technology node (e.g., 65 nm, 45 nm, ...):** Our simulation was done with a 65 nm process. Many new designs are moving on to 45 nm because of the high density it affords and because of the extremely high transistor bandwidths. There is also a trend to compensate analog circuits with digital calculations allowing analog circuit size and power reductions. With this trend, it is expected that new designs will migrate to the most modern processes available and affordable that support any mixed-signal requirements and the benefits of NCL will increase. Cost does not justify the latest processes in some designs and applications requiring high voltage or high precision analog components will likely remain at 90 nm and larger.
- 7) **Average activity factor:** Activity factor determines the percentage of energy that is dynamic vs. static. Designs having a high activity factor require more dynamic energy. Designs with a low activity factor, especially at 65 nm and below have a much higher percentage of static power. We implemented a datapath architecture in this analysis in order to exercise NCL in a situation that was known to be stressing for NCL. Although the overall activity factor was not estimated for this design, the fact that it was a datapath architecture means that it must compute an output every clock cycle and this means that it has a higher proportion of dynamic energy. Since PSO saves only static E/Op, designs using PSO for energy savings would save the most energy when the circuit has a low activity factor. If AVS were used, both static and dynamic energy could be saved.

With those factors in mind, the results of the comparisons follow.

C. Comparison of NCL to CBL in DoD Systems

Fig. 9 shows the comparison of subthreshold CBL and NCL performance of a 5-tap FIR filter implemented in a

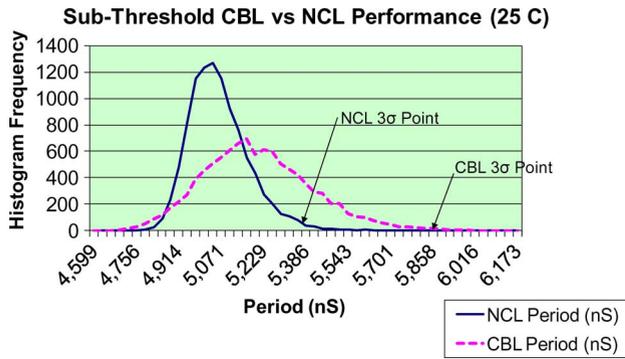


Fig. 9. Subthreshold CBL versus NCL performance.

65 nm bulk CMOS process. The variation in the circuit’s performance under subthreshold operating conditions was simulated using Monte Carlo analysis techniques, varying the transistor characteristics in accordance with the process’ variability. The circuit speeds were then measured by calculating the minimum period of a clocked data source that could be used to drive the circuit. The results were then binned to produce the plot shown. The Histogram Frequency shows the number of circuits from the Monte-Carlo runs that fell into each bin.

As can be seen in Fig. 9 from the distance of the 3σ point from the center of the respective curves, the NCL circuits have significantly less variation in performance than the CBL circuits. This translates to increased yield and reduced die cost for the NCL design versus the CBL design.

The figure shows that, in this example, NCL logic can implement fixed-throughput systems that are 8% to 18% faster than CBL for the same yield due to the necessity to use lognormal statistics in subthreshold operation. Also because of their delay insensitivity NCL circuits with negative timing slack for the selected yield point will still function (but may not meet required performance goals), whereas CBL circuits with negative slack time fail.

In the subthreshold case shown in Fig. 10, an analysis compared a CBL and NCL circuit operating over the military specified temperature range, both designed in a 65 nm process. For the CBL circuit the clock speed required to meet the worst case corner of the envelope is 9.15 kHz. In contrast to this an equivalent clockless NCL circuit responds to the operating conditions. It runs at 11.4 kHz at the worst case corner of the envelope but runs at 198 kHz at high temperatures.

This ability for the NCL circuit to operate at the actual circuit conditions means that the NCL circuit runs approximately 20× the speed of the CBL system at nominal temperature, while the NCL system consumes approximately the same energy per calculation as it did at the worst case temperature. At 125 °C, the NCL circuit runs 250× as fast as the CBL circuit, and consumes nearly 32× less energy per calculation. The energy savings is realized either with power shutoff at the end of the calculation or preferably by voltage scaling which would save more energy in this dynamic energy per operation (E/Op) dominated example.

The graphs on the right of the figure show how dynamic and static (leakage) power contribute to the overall

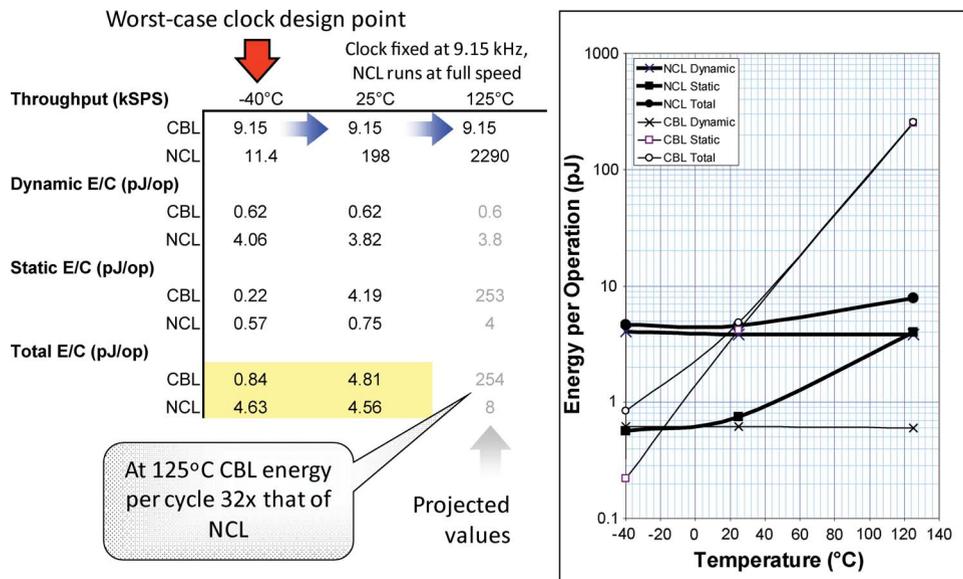


Fig. 10. Comparison of CBL and NCL circuits operating in subthreshold.

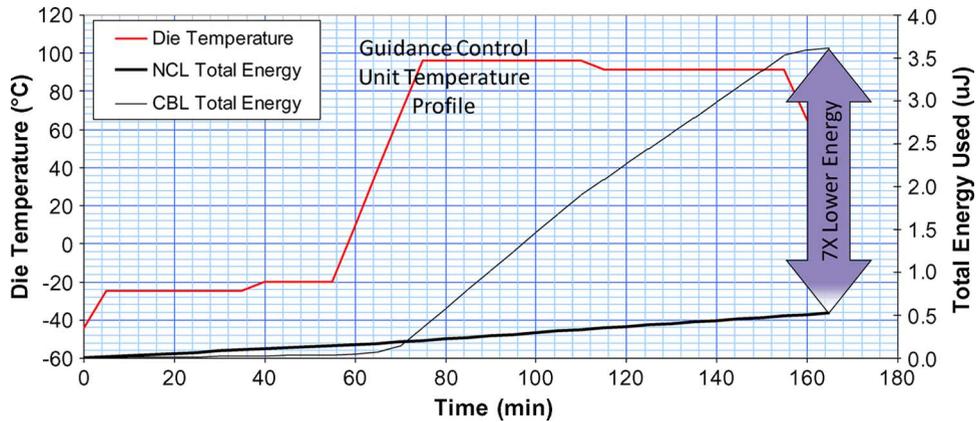


Fig. 11. NCL impact on mission energy requirement.

energy required per cycle (E/C). Note that the CBL systems' static energy consumption increases linearly, because the circuit has to wait for the worst case timing bounded by the clock period.

The impact of this behavior on a military system is shown in Fig. 11 where it can be seen that over the temperature profile shown, NCL circuits combined with power shutoff would use $7\times$ less energy than the CBL circuits not having the capability to adapt to temperature changes.

Military systems operate over a broad temperature range, often with high performance or mixed-signal content producing temperature gradients within the physical extent of the clock tree. When comparing NCL to CBL, based on our long experience with designing and using military ASICs, our perception is that by combining NCL with power shutoff or adaptive voltage scaling at least 80% of military projects could achieve energy savings between that shown in the figure ($7\times$) and $20\times$ lower energy. The remaining systems might achieve even greater energy reduction; up to $30\times$ or more over the life of a mission.

This best reflects our impression of the percentages of military systems that fall within these categories. Due to the complexities, intellectual property issues, and time and expense involved in doing a thorough quantification of the percentages of systems that could achieve these savings, we feel that a more accurate quantification than those we provided would be impractical.

IV. ADDITIONAL DEVELOPMENT TO MAKE CLOCKLESS LOGIC AVAILABLE TO SUBTHRESHOLD/ULTRALOW-POWER DESIGNERS

Although there are clearly potential advantages to applying NCL approaches to subthreshold systems as shown, NCL design methodologies are less mature than their

equivalent CBL design methodologies (we are working with a tools/methodology company to address this). While this represents a significant handicap in conventional superthreshold designs, the mature CBL design approaches are not well suited to ULP operation in a subthreshold regime.

Commercially available CBL subthreshold design approaches are not presently available that produce optimal designs. Statistical static timing analysis is required to avoid excessive timing margins, however, process design kit support needed to do the analysis with commercial tools at subthreshold voltages having lognormal statistical analysis capability is not available for many leading edge manufacturing processes. Lack of a robust methodology for NCL based subthreshold design is obviously a problem, but in this case it is not as severe a penalty as it would be in the nominal voltage regime.

Some of the challenges unique to subthreshold design require additional development in order to provide a design environment to ULP designers that is comparable to existing design environments for CBL superthreshold designs.

Operating at the lowest voltage levels requires the careful design and modeling of cell library structures, where the sizing and topology of the cells need to be validated for use in the subthreshold environment. Sizing rules that may work well in the normal operating regime may leave cells unable to function across extreme low voltages. Cell libraries need to be developed and characterized specifically for operation in subthreshold conditions in addition to normal mode operation [31]. Many cells that may function adequately at normal voltage levels may be incapable of operating at very low subthreshold levels. With this in mind, designers need to understand the trade-offs and may need to synthesize their lowest energy designs into a smaller subset of cells than are available for sections of the design that are destined for normal voltage level operation. This applies equally to CBL and NCL approaches.

To truly make subthreshold design routine for designers, approaches for giving the designer feedback on the trade-off of energy versus performance during the design phase need to be developed. Designers interested in the lowest possible energy designs need to be able to understand the trade-offs they are making in regards to throughput, energy consumption, and area. Such tools need to be capable of providing designers with information on the performance of designs across all valid conditions. The kinds of performance information required would include the worst case speeds, best case speeds, energy consumption, but would also need to provide some estimate of the expected nominal throughput: unlike the CBL design which is limited to the worst case performance by the clock, the NCL design will tend to operate much closer to the nominal throughput figures, somewhere between worst case and best case timings.

To truly take advantage of the lowest energy calculations, multiple voltage level schemes need to be developed, analyzed, and demonstrated. Trade-offs between the overhead and complexity of the more complicated voltage schemes versus the granularity of control over power consumption must be analyzed, and design guidelines created. Considerations on boundaries for level shifting need to be explored, to determine the appropriate granularity of subthreshold voltage domains.

Additionally, interfaces to, and capabilities for operating additional components required for creating usable systems in subthreshold need to be investigated. This would include approaches to interfacing efficiently with various classes of memories as well as analog or RF circuitry.

Further, operation with the widely varying characteristics expected in subthreshold will open up whole new categories of challenges for IC test. New approaches for efficiently validating both functionality and performance, across extremely wide operating characteristics need to be examined. Are there new classes of tests, new ways to compress testing that can be developed to better facilitate the validation of circuits operating in subthreshold? What can be done to accelerate testing of devices, so that the cost of testing the chips does not become prohibitively expensive due to long test times on very expensive testers? These are key areas to be researched.

To address one of the key issues with operation in the subthreshold mode, the loss of performance, new approaches will need to be considered. In a typical microsystem architecture there are periods of low performance in which subthreshold operation can be used to reduce the energy usage, but there are periods in which much higher performance is required, for example to perform complex signal and image processing algorithms. A way to deal with this situation is the use of power domains: the use of processing blocks in which the supply voltage is raised or lowered to meet the processing demand.

In a CBL design this requires very careful matching of the power supply voltage with the clock rate and the

availability of a range of clock speeds. An NCL circuit inherently responds to the power supply voltage. This allows the possibility of using the data rate being introduced into a circuit block to dynamically control the power supply voltage and turn the voltage up or down to match the data rate.

Currently we are working with EDA partners to address this mode of operation. Taking advantage of high level knowledge of the NCL circuit functionality, it will be possible for the designer to optimize the power supply voltage and dynamically balance data rate and supply voltage. This ongoing work is intended to automate the process of designing the structures needed to automatically adjust the performance of small scale NCL circuits to provide the most energy efficient operation. Turning the voltage down when there is little or no activity in that subcircuit, or turning it up as the data rates increase, moving through the entire operating range from subthreshold to superthreshold levels. The aspects being explored that make this innovative are the granularity with which this can be achieved, and the ability to move the operating voltage continuously from sub- to superthreshold.

In principle this use of data driven control of the power supply voltage can provide a continuous dynamic control that will allow the circuits to “slide” up and down the power versus performance envelope from subthreshold operation to normal mode operation. Some performance penalty will occur because currently transistor structures cannot be optimized simultaneously for subthreshold operation and normal mode operation, however as results from the DARPA Subthreshold-slope Transistors for Electronics with Extremely Low Power (STEEP) program [32] emerge these can easily be incorporated into an NCL cell library.

V. CONCLUSION

This paper has looked at the challenges associated with emerging battery-powered microsystems and their widely varying power-performance envelopes, and some of the challenges these systems can present. Digital logic subthreshold operation was introduced briefly as a means to potentially achieve very high energy savings, and the major issue of timing uncertainty was presented. Clockless logic was presented as an emerging class of digital logic design techniques as a means to enable subthreshold/ultralow-power operation by mitigating the extreme timing issues associated with significant PVT variation that occurs in this voltage regime. Some representative design results were introduced, along with a brief discussion of operation of some of the different styles of clockless logic.

As proof of concept of the potential of clockless logic's, and particularly NCL's, ability to enhance the value of subthreshold design, results from the consideration of actual systems were presented. The results showed the

significant power-performance advantages that can be had by using a quasi delay insensitive clockless logic style to design circuits for subthreshold operation, resulting in a conclusion that significant energy savings of 7–20× could be had in a large percentage of designs by using a clockless approach, and that even larger savings of up to 30× could be had in some applications. Further, the ability of the NCL approach to adapt to operating conditions automatically across PVT variations showed promise of better throughput than subthreshold CBL designs. Some of the pros and cons of the clockless logic style considered in this paper relevant to subthreshold operation were presented. Finally, this paper outlined some ongoing and future work required in order to enable designers to take advantage of

clockless subthreshold design, including looking at applying dynamic power control to the operation of clockless circuits moving between subthreshold and superthreshold regimes. ■

Acknowledgment

The views, opinions, and/or findings contained in this article/presentation are those of the author/presenter and should not be interpreted as representing the official views or policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the Department of Defense. Approved for Public Release, Distribution Unlimited.

REFERENCES

- [1] J. Kwong, Y. Ramadass, N. Verma, M. Koesler, K. Huber, H. Moormann, and A. Chandrakasan, "A 65 nm Sub-Vt microcontroller with integrated SRAM and switched-capacitor DC-DC converter," in *Proc. IEEE ISSCC*, Feb. 2008, pp. 318–319.
- [2] H. Soeleman, K. Roy, and B. C. Paul, "Robust subthreshold logic for ultra-low power operation," *IEEE Trans. VLSI Systems*, vol. 9, no. 1, pp. 90–99, Feb. 2001.
- [3] C. H.-I. Kim, H. Soeleman, and K. Roy, "Ultra-low-power DLMS adaptive filter for hearing aid applications," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 11, no. 6, pp. 1058–1067, Dec. 2003.
- [4] A. Wang and A. Chandrakasan, "A 180-mV subthreshold FFT processor using a minimum energy design methodology," *IEEE J. Solid-State Circuits*, vol. 40, no. 1, pp. 310–319, Jan. 2005.
- [5] L. Nazhandali, B. Zhai, J. Olson, A. Reeves, M. Minuth, R. Helfand, S. Pant, T. Austin, and D. Blaauw, "Energy optimization of subthreshold-voltage sensor network processors," in *Proc. 32nd Int. Symp. Comput. Archit. (ISCA'05)*, Jun. 2005.
- [6] V. Sze, R. Blazquez, M. Bhardwaj, and A. P. Chandrakasan, "An energy efficient subthreshold baseband processor architecture for pulsed ultra-wideband communications," in *IEEE Int. Conf. Acoust., Speech Signal Process.*, May 2006.
- [7] S. Hanson, B. Zhai, K. Bernstein, D. Blaauw, A. Bryant, L. Chang, K. K. Das, W. Haensch, E. J. Nowak, and D. M. Sylvester, "Ultralow-voltage, minimum-energy CMOS," *IBM J. Res. Develop.*, vol. 50, no. 4/5, pp. 469–490, Jul./Sep. 2006.
- [8] B. Zhai, S. Hanson, D. Blaauw, and D. Sylvester, "Analysis and mitigation of variability in subthreshold design," in *Proc. Int. Symp. Low Power Electron. Design*, Aug. 2005, pp. 20–25.
- [9] R. Vaddi, S. Dasgupta, and R. P. Agarwal, "Device and circuit design challenges in the digital subthreshold region for ultralow-power applications," *VLSI Design*, vol. 2009, Article ID 283702, 14 pp. 2009, DOI: 10.1155/2009/283702.
- [10] L. S. Nielsen and J. Sparso, "Designing asynchronous circuits for low power: An IFIR filterbank for a digital hearing aid," *Proc. IEEE*, vol. 87, no. 2, pp. 268–280, Feb. 1999.
- [11] K. S. Chong, B. H. Gwee, and J. S. Chang, "Energy-efficient synchronous-logic and asynchronous-logic FFT/IFFT processors," *IEEE J. Solid-State Circuits*, pp. 2034–2045, Sep. 2007.
- [12] *The "Asynchronous" Bibliography*, 2004. [Online]. Available: <http://www.win.tue.nl/async-bib/>
- [13] K. M. Fant, *Logically determined Design: Clockless System Design With NULL Convention Logic Hoboken*. Hoboken, NJ: Wiley Interscience, 2005.
- [14] J. Le, C. Hanken, M. Held, M. Hagedorn, K. Mayaram, and T. S. Fiez, "Comparison and impact of substrate noise generated by clocked and clockless digital circuitry," in *Proc. CICC 2006*, Sep. 2006, pp. 105–108.
- [15] M. Ferretti and P. A. Beerel, "High performance asynchronous design using single-track full-buffer standard cells," *IEEE J. Solid-State Circuits*, pp. 1444–1454, Jun. 2006.
- [16] J. Teifel and R. Manohar, "A high-speed clockless serial link transceiver," in *Proc. Int. Symp. Adv. Res. Asynchronous Circuits Syst.*, May 2003, pp. 151–161.
- [17] [Online]. Available: <http://research.sun.com/features/async>
- [18] M. Singh and S. M. Nowick, "MOUSETRAP: Ultra-high-speed transition-signaling asynchronous pipelines," in *Proc. Int. Conf. Comput. Design (ICCD)*, Nov. 2001, pp. 9–17.
- [19] B. D. Winters and M. R. Greenstreet, "Surfing: A robust form of wave pipelining using self-timed circuit techniques," *Microprocess. Microsyst.*, vol. 27, no. 9, pp. 409–419, Oct. 2003.
- [20] P. Golani, G. D. Dimou, M. Prakash, and P. A. Beerel, "Design of a high-speed asynchronous turbo decoder," in *13th IEEE Int. Symp. Asynchronous Circuits Syst. (ASYNC'07)*, 2007, pp. 49–59, ASYNC.
- [21] A. Davis and S. M. Nowick, "An Introduction to Asynchronous Circuit Design," Dept. Comput. Sci., Univ. Utah, Tech. Rep. UUCS-97-013. [Online]. Available: <http://www1.cs.columbia.edu/async/publications/davis-nowick-intro-tr.pdf>
- [22] Handshake Solutions, *ARM996HS Processor Leaflet*, Feb. 2006. [Online]. Available: http://www.handshakesolutions.com/assets/downloadablefile/ARM996HS_leaflet_feb06-13004.pdf
- [23] A. J. Martin, S. M. Burns, T. K. Lee, D. Borkovic, and P. J. Hazewindus, "The design of an asynchronous microprocessor," in *1989 Caltach Conf. Very Large Scale Integration*.
- [24] A. J. Martin, A. Lines, R. Manohar, M. Nystroem, P. Penzes, R. Southworth, U. Cummings, and T. K. Lee, "The design of an asynchronous MIPS R3000 processor," in *Proc. 17th Conf. Adv. Res. VLSI*, R. B. Brown and A. T. Ishii, Eds., 1997, pp. 164–181.
- [25] A. J. Martin and M. Nystrom, "Asynchronous techniques for system-on-chip design," *Proc. IEEE*, vol. 94, no. 6, pp. 1089–1120, Jun. 2006.
- [26] K. Fant and S. Brandt, "NULL convention logic: A complete and consistent logic for asynchronous digital circuit synthesis," in *Proc. 1996 Int. Conf. Appl. Specific Syst., Archit., Process. (ASAP 96)*, Aug. 1996, pp. 261–273.
- [27] S. Masteller and L. Sorenson, "Cycle decomposition in NCL," *IEEE Des. Test Comput.*, vol. 20, no. 6, pp. 38–43, Nov./Dec. 2003.
- [28] S. Kakarla and W. K. Al-Assadi, "Testing of asynchronous NULL Conventional Logic (NCL) circuits," in *Proc. IEEE Region 5 Conf.*, Apr. 2008, pp. 1–6.
- [29] D. Ernst, S. Das, S. Lee, D. Blaauw, T. Austin, T. Mudge, N. S. Kim, and K. Flautner, "Razor: Circuit-level correction of timing errors for low-power operation," *IEEE Micro*, vol. 24, no. 6, pp. 10–20, Nov. 2004.
- [30] A. Bailey, A. Al Zahrani, G. Fu, J. Di, and S. C. Smith, "Multi-threshold asynchronous circuit design for ultra-low power," *J. Low Power Electron.*, vol. 4/3, pp. 337–348, Dec. 2008.
- [31] B. H. Calhoun and A. P. Chandrakasan, "Characterizing and modeling minimum energy operation for subthreshold circuits," in *Proc. Int. Symp. Low Power Electron. Design (ISLPED)*, Aug. 2004.
- [32] GOMAC Tech 2009, Session 24 STEEP-Subthreshold Low power electronics.

ABOUT THE AUTHORS

Ryan D. Jorgenson (Member, IEEE) received the B.S. degree in electrical engineering with a minor in Computer Engineering from the University of Calgary, Canada, in 1988.

He has more than 20 years of experience in the development of system, board, and ASIC level products and clockless design methodologies. He has managed and worked on a wide variety of ASIC and product development projects for consumer, industrial, security, defense, and scientific applications, including most recently programs focused on applying low-power technology to unattended ground sensors and their related systems. Previously at Theseus Logic, Inc., he helped to develop and implement clockless logic designs and design tools based on NULL Convention Logic. He is currently Vice President of Advanced Systems, Camgian Microsystems, Maitland, FL.



Lief Sorensen received the B.S. degree in chemistry from Andrews University and the M.S. degree in electrical and computer engineering from the University of California.

At Hewlett Packard, he investigated the ITT digital TV chip set and prototyped a project to integrate real-time video with HP's graphics displays, as well as being responsible for the lab-wide design and implementation of a new test architecture to streamline effort needed to verify ASICs and graphics hardware subsystems. Recent efforts include CyclePath, an NCL Circuit Analysis Tool, development of a Shipping Container Security System, and design of low-power remote wireless networks for global monitoring. He is currently on Senior Technical Staff, Camgian Microsystems, Maitland, FL.



Dan Leet (Member, IEEE) received B.S. and M.S. degrees in computer engineering from the University of Central Florida in 2006 and 2007 respectively, with a focus on digital systems.

He is currently a Design Engineer, Camgian Microsystems, Maitland, FL. Since joining Camgian in 2006, he has been involved in the development of several integrated sensor system programs, including environmental sensing, infrared imagery, wireless communications, and radar. His research interests include mixed signal and RF integration, asynchronous IC design, and low-power system architectures.



Michael S. Hagedorn (Member, IEEE) received the Ph.D. degree in electrical engineering from the University of Minnesota.

He is currently on Senior Technical Staff, Camgian Microsystems, Maitland, FL. He has over 20 years of industry experience as an architect and designer of mainframe and microcomputers. He has been the principal investigator on many programs and has designed NCL systems for both low-power and high throughput applications. He is working on the design of the key RF components for a Radar-on-a-chip under a DARPA MTO SBIR, as well as working on the design of a compact, low-power Unattended Ground Sensor architecture. He has published numerous papers on the subject of NCL design and has patents pending in the area.



David R. Lamb (Member, IEEE) received the B.A. and M.A. degrees from Oxford University and the M.Sc. and Ph.D. degrees from London University.

He is currently Chief Operating Officer, Camgian Microsystems, Maitland, FL. Prior to joining Camgian, David was CEO of Theseus Logic Inc. He has over 30 years experience in the semiconductor field and has held senior industry and university positions in both the United States and the U.K. He has extensive experience in developing advanced technologies and was Associate Director at the Honeywell Technology Center, where he managed a broad portfolio of technology areas including data acquisition systems, wireless communications, photonics, information processing systems, and flat panel/head mounted displays.

Dr. Lamb received the Lund Award, Honeywell's highest award for managerial excellence.



Thomas Hal Friddell (Member, IEEE) received the B.S.E.E. degree from WSU in 1976 and the M.S.E.E. degree from UW in 1988.

He is an Associate Technical Fellow with over 33 years of engineering experience at Boeing, Seattle, WA. His design experience includes analog, digital, and mixed signal hardware/software design and modeling. His recent ASIC design work includes design of an integrated circuit for pressure sensor processing, IR detector interface electronics design, radiation hardened mixed-signal structured ASIC generators, a 40 GHz DDS DAC, asynchronous digital processing implemented at 45 nm, and he led an X-band Radar-on-a-Chip design project. Recent DARPA projects include ultralow-power studies and the Navigation-Grade Integrated Micro Gyro (NGIMG).



Warren P. Snapp, photograph and biography not available at the time of publication.