

Low Energy Asynchronous Adders

Ilya Obridko and Ran Ginosar

VLSI Systems Research Center

Technion —Israel Institute of Technology

Haifa 32000, Israel

[oilya@tx.technion.ac.il]

Abstract: Asynchronous circuits are often presented as a means to achieve low power operation. We investigate their suitability for low-energy applications, where long battery life and delay tolerance is the principal design goal, and where performance is not a critical requirement. Three adder circuits are studied—two dynamic and one based on pass-transistor logic. All adders combine dual-rail and bundled-data circuits. The circuits are simulated at a wide supply-voltage range, down to their minimal operating point. Leakage energy (at 0.18 μm) is found negligible. Transistor count is found to be an unreliable predictor of energy dissipation. Keepers in dynamic logic are eliminated when possible. A modified version of a two-bit dynamic adder (originally proposed by Chong) is found to dissipate the least amount of energy.

Index terms

Low energy, adder, asynchronous logic.

1 Introduction

Asynchronous logic has been promoted as a means to achieve low power design [1][2][6]. A number of advantages of asynchronous logic that make it appropriate for low power operation have been cited: Asynchronous circuits can stop computing when there is no new input, without the extra complexity of clock-gating logic and without the need to wait for clock restart delays. Power dissipation in large clock distribution trees is eliminated, though partly replaced by local handshake power [10]. When the circuit is speed-independent, supply voltage can be reduced when lower performance can be tolerated without having to retune clock frequencies [5]. More recently, asynchronous low energy (rather than low power) has been addressed [6][7], as this is more appropriate a design goal for extending battery life for mobile and other devices, as well as minimizing the efforts for heat dissipation and cooling expenses. Low power and low energy techniques for asynchronous systems are typically based on minimizing the number of transitions [1]. Other approaches include voltage scaling [5], early-open latch controllers, and data-dependent enabling of the logic [1][3][7].

We focus on simple computing circuits that must dissipate as little energy as possible in applications where performance is non-limiting and the time to complete any computing task is immaterial. A secondary goal is to be able to operate over a very wide range of supply voltage, as is typically the case with some battery-operated devices where voltage regulation is not desirable. The principal implication of a varying supply voltage is a wide range of delays, calling for the speed-independence feature of asynchronous circuits. The most robust speed-independent circuit methodology is based on dual-rail encoding and on quasi-delay-insensitive (qDI) design [1]. Unfortunately, qDI circuits are not necessarily the most energy efficient ones.

Four-phase qDI data signaling is based on alternating valid and null values. Each data bit must toggle from valid to null and back again on every successive data value, even if the data on both sides of

the null have the exact same value. Two-phase qDI protocols help reduce delays but do not improve energy consumption. Bundled data signaling (in both synchronous and asynchronous circuits) eliminates data switching when data values do not change. However, bundled data speed independent logic may not be as tolerant to wide delay variations as qDI circuits, since most bundled data schemes require matched delays and are exposed to the risk of not being long enough, on one hand, while always incurring a worst-case delay, on the other hand.

Another low energy technique prefers large combinational blocks and minimizes the use of pipeline registers. Purely combinational logic could sometimes achieve minimum energy per computation, as long as redundant transitions are avoided.

As a basic test case we consider two-bit adders, which are commonly required for signal processing applications. Although complete CPU or DSP systems may dissipate more energy in other sections, such as their instruction fetch and decode units [7], this is typically due to performance optimization; in low-energy applications where execution rate is not an issue, the data-path is expected to become the energy bottleneck.

We investigate a hybrid bundled data/dual rail approach [1]. The dual-rail part provides completion indication, while the bundled data parts help minimize energy dissipation. As an example, we apply the design methodology to a large adder, and compare it with other published low energy adders [4][9]. The various adders are presented in Section 2. The actual circuits used for our analysis are described in Section 3, and the energy dissipation and simulations results are discussed in Section 4.

2 Low Energy Adder Architectures

In order to achieve high performance in wide adders, carry look-ahead circuits are usually employed. However, such circuits dissipate extra energy. In low-energy applications when performance is not an issue, no look-ahead circuits should be used. Thus, we consider only ripple-carry adders. We also employ those hazard-free asynchronous techniques that block spurious transitions and perform their computations only after all inputs have arrived.

Another energy-related advantage of asynchronous ripple carry adders is their relatively simple completion-detection; in the circuits below, the carry-out of the last stage is considered as the indication of completion, and all sum outputs are assumed to be ready by the time the carry-out becomes valid.

2.1 The Nielsen Adder

Nielsen [2][3] combines two types of dynamic adder circuits (Figure 1). The least significant half of the n -bit adder employs carry-kill and carry-generate logic to speed up computation. The most significant half of the same adder employs ripple carry adder circuits without any carry acceleration. All adders produce dual-rail carry-out and single-rail sum outputs.




Figure 1 Nielsen Adder

The energy minimization idea behind the Nielsen adder is based on data slicing into the most and least significant parts. The calculation of the lower half is always performed, and the upper part is enabled only when the inputs are large. The lower half is designed to produce the carry out signal as soon as it can be computed, while the upper part is designed to reduce the completion detection tree structure.

2.2 The Chong Adder

Chong *et al.* [4] introduce low-energy adders that also produce dual-rail carry-out and single-rail sum outputs. Two bits are combined and complex gates are employed to further minimize energy. The schematic description is depicted in Figure 2; note that completion detection depends only on the last carry-out.




Figure 2 Chong Adder

2.3 Path Transistor Logic Adder

Single-ended pass transistor logic (SPL) and complementary pass transistor logic (CPL) are advocated as suitable for low energy design especially for arithmetic functions [8][9]. The main reason is that arithmetic functions are based on many XOR gates, and pass-transistor logic enables efficient implementations of XOR gates. CPL and SPL methods (named PTL below) contribute to the energy minimization by the small number of pass transistors, which are usually NMOS, and produce very compact and regular designs. Another advantage of PTL is that VDD-to-GND paths, which may lead to short-circuit energy dissipation, are eliminated.

The main disadvantage of PTL is the delay of the circuit, which is more sensitive to voltage scaling than CMOS logic. Another drawback is the degradation of the voltage swing to one V_{TH} away from the supply. Voltage swing restoration buffers are required, increasing the transistor count and energy dissipation.

3 Low Energy Adder Circuits

Three full-adder (FA) circuits are compared for energy and transistor count: A dynamic FA from Nielsen's adder, a Chong dynamic two-bit FA, and a PTL FA. The dynamic circuits are naturally suited for use in asynchronous systems, while the output of the PTL FA is enabled (and swing-restored) by the Request signal, as in Figure 3.




Figure 3 Pass-gate with NOT at the output of PTL cell.

3.1 Dynamic Full-Adder

The ripple-carry adder, used in the upper half of the Nielsen's adder, is modified by removing some logically redundant transistors that were employed for timing balance. The dynamic FA uses a single-rail sum, dual-rail inputs, and dual-rail carry-out. A complete adder using this dynamic FA is shown in Figure 4. Note that the adder is reset by *REQ*. This allows quick execution of the return-to-zero part of the handshake. Likewise, when *REQ* rises, all the stages in the chain start their calculations simultaneously. *Sum1,2,3* and *cout1,2,3* contain only NMOS pull down logic. The FA circuit is shown in Figure 5. The keepers are marked by dashed-lines; we have found that eliminating them in this circuit does not affect energy dissipation (in contrast with Chong's FA, below).

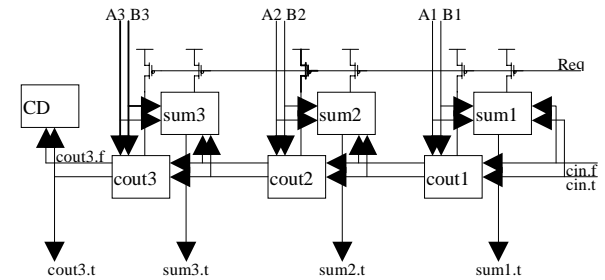


Figure 4 Adder Based on Dynamic FA

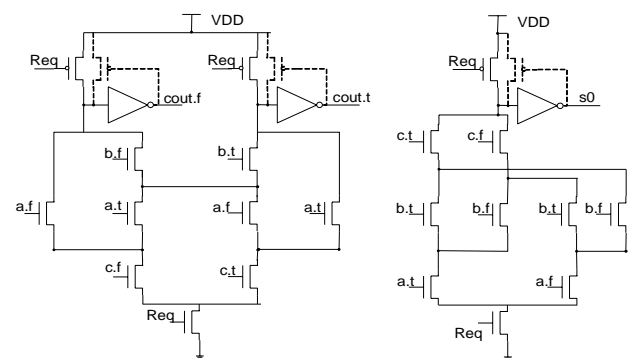


Figure 5 Dynamic FA Circuit

3.2 Two-Bit Chong's Full-Adder

The two-bit FA circuit from Chong's adder [4] was modified by eliminating the keepers (marked by dashed-lines in Figure 6); we have found out that the circuit dissipates less energy and its functionality is unaffected without the keepers.




Figure 6 Chong's Dual-bit implementation.

3.3 PTL Full-Adder

The PTL FA of [9] has been appended with the Request-enabled output inverter and adapted to produce dual-rail carry-out (Figure 7).




Figure 7 PTL FA Circuit

4 Simulation Results

For fair comparison with Chong's two-bit FA, all designs were simulated as two-bit circuits. All three FA circuits were designed (at the schematic transistor level) for TSMC 0.18 μ m technology and simulated with Cadence Spectre. The simulated circuits included completion detection. All outputs were loaded by 10fF capacitors. Since voltage scaling serves as the principal means for energy reduction, all simulations were conducted by V_{DD} sweeping over 0.7—1.5V (where 1.8V is the nominal V_{DD} for the technology). All 32 input combinations (of two 2-bit numbers plus carry-in) were simulated in each case, and energy dissipation was averaged over all 32 cases.

Ten cycles of valid-then-empty inputs were simulated, with a long idle period in the middle (Figure 8). Thus, measurements results

are more reliable, and by varying the idle period we were able to determine that leakage accounts for less than 1% of the total energy consumed by the adder.

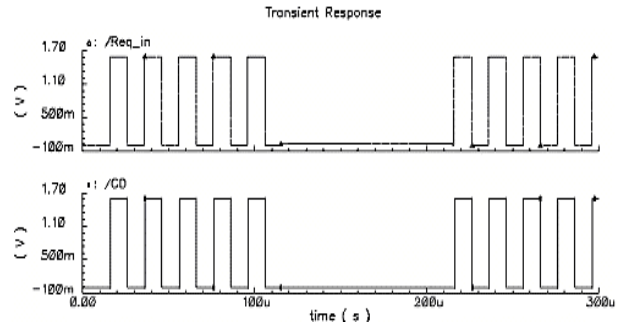


Figure 8 Input conception for fair adders' simulation. (CD – Completion Detection signal)

Figure 9 shows the transistor counts for three 2-bit FA circuits.

Figure 10 presents the energy dissipation of the three circuits versus V_{DD} , averaged over 32 runs of ten additions each of the two-bit adders, including the idle times. Other circuits have also been simulated, but their energy consumption far exceeded that of these three circuits.

We can learn from the simulations that Chong's adder dissipates the least amount of energy. PTL dissipates a bit more, but less than Nielsen's FA. All three circuits demonstrate robustness to a wide variation of voltage levels. Chong's FA produces the best result thanks to its dual-bit structure, reducing the logic size, eliminating redundant wiring and consequentially reducing the number of transitions. These observations provide a strong incentive to design larger blocks of logic in order to gain maximal energy reduction.

We checked the transistor count of the adders in order to investigate their impact on energy. The conclusion was that mere transistor count is not a sufficient predictor of energy dissipation. PTL FA requires the largest number of transistors (40% of them were employed in the Request-enabled output buffer that was required to make it "asynchronous"). Still, the PTL FA dissipates on average 14% less energy than the dynamic FA. Also, despite the fact that PTL FA requires 17% more transistors than Chong's FA, it dissipates only about 10% more energy. Chong's FA contains 8.5% fewer transistors but consumes 20% less energy than the (single-bit) dynamic FA, thanks to producing only one carry-out signal. The dynamic FA calculates a carry-out signal per every bit, thus dissipating more energy.




Figure 9 TransistorCount Comparison

5 Conclusion

We have investigated some novel adder circuits and have been able to identify the low-energy ones. Delay was ignored in this analysis so as to emphasize low energy over all other parameters. Our next research goal is to investigate the E_t and E_t^2 metrics [6][9][10]. In addition, we plan to consider 2-bit and 3-bit circuits for further energy reduction.

Acknowledgment

We are grateful to Dana Amburg, Michael Moreinis, Yevgeny Perelman and Akadiy Morgenshtein, who have helped with ideas and CAD tools.

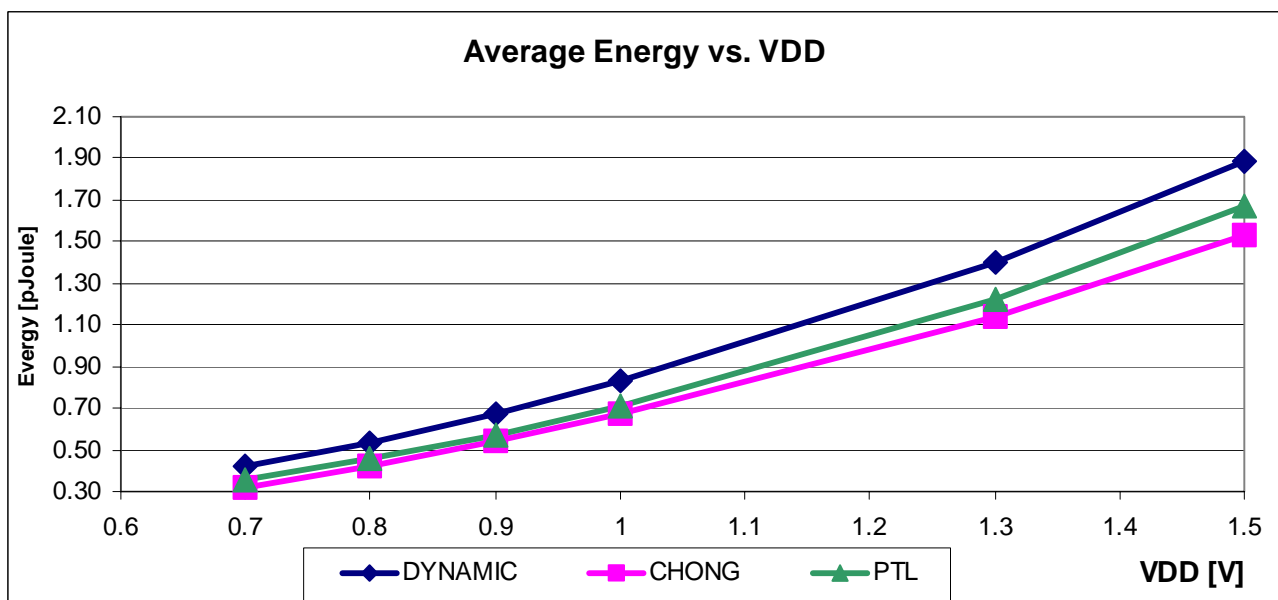


Figure 10 Simulation results

References

- [1] J. Sparso and S. Furber, *Principles of Asynchronous Circuit Design: A Systems Perspective*: Kluwer Academic Publishers, 2001.
- [2] L.S. Nielsen and J. Sparsø, "A Low-power Asynchronous Data-path for a FIR Filter Bank," *Int. Symp. Adv. Res. Async. Circuits and Systems (ASYNC '96)*, pp. 18 - 21, 1996
- [3] L.S. Nielsen, "Low-power Asynchronous VLSI Design," Ph.D. Thesis, Department of Information Technology, Technical University of Denmark, 1997.
- [4] K.S. Chong, B.H. Gwee and J.S. Chang, "Low-voltage Asynchronous Adders for Low Power and High Speed Applications," *Int. Symp. Circuits and Systems (ISCAS)*, 2002.
- [5] L. S. Nielsen, C. Niessen, J. Sparso, "Low-power operation using self-timed and adaptive scaling of the supply voltage," *IEEE Trans. VLSI Systems*, 2:391-397, 1994.
- [6] A.J. Martin, "Remarks on low-power advantages of asynchronous circuits," *Europ. Solid-State Circuits Conf. (ESSCIRC)*, 1998.
- [7] A.J. Martin, M. Nyström *et al.*, "The Lutonium: A Sub-Nanojoule Asynchronous 8051 Microcontroller," *IEEE Int. Symp. Async. Systems and Circuits*, May 2003.
- [8] M. Munteanu, I. Bogdan *et al.*, "Single-Ended Pass Transistor Logic for Low-Power Design," *IEEE Asilomar Conf. Signals Systems and Computing*, pp. 364-368, 1999.
- [9] L. Bisdounis, D. Gouvetas and O. Koufopavlou, "Circuit techniques for reducing power consumption in adders and multipliers," in D. Soudris, C. Piguet and C. Goutis, "Designing CMOS Circuits for Low-Power," Kluwer Academic Publishers, 2002.
- [10] A.J. Martin, "An asynchronous approach to energy-efficient computing and communication," *SSGRR 2000*, August 2000.