

Comparison of High Speed Voltage-Scaled Conventional and Adiabatic Circuits

David J. Frank

IBM T. J. Watson Research Center
P.O. Box 218, Yorktown Heights, NY 10598

Abstract

The power versus frequency performance of a micro-pipelined conventional CMOS logic family is compared with that of three similarly pipelined energy-recovering logic families. Using a circuit simulator, the supplies and operating voltages of each family are optimized for minimum power consumption at each frequency. One of the energy-recovering logic families is shown to be capable of substantially lower dissipation than the conventional case, one is comparable, and one is worse.

Introduction

Adiabatic switching is a recently advocated circuit technique for reducing the power dissipation in digital logic by recovering some of the energy that would be dissipated in conventional logic(1,2). A variety of circuit approaches to creating adiabatic switching logic have been proposed, including both retractile and micro-pipelined techniques. Analyses of the retractile approaches have shown that they can only do better than voltage-scaled CMOS at *very* low frequency(2,3). This work presents a comparison of conventional circuits and adiabatic circuits in the high speed micro-pipelined regime, using voltage-scaling for all circuit families considered.

Circuits and Methodology

The four circuit families that have been compared are: (a)TSPC (True Single Phase Clock, a conventional CMOS logic family)(4), (b) 2N-2N2D (one of the first adiabatic families)(5), (c) 2N-2N2P (a more recently proposed variation on the same idea as 2N-2N2D)(6), and (d) Hot Clock nMOS (as originally proposed by Seitz)(7,8). An example of a 2-input NAND gate is shown for each of these families in Fig. 1(a-d). Note that (b) and (c) are dual rail logic families, which simplifies some of the logic. (a) uses a single clock, while (b) and (d) use dual (non-overlapping) clocks, and (c) uses a set of 4 sine waves. Fig. 2 show the clock

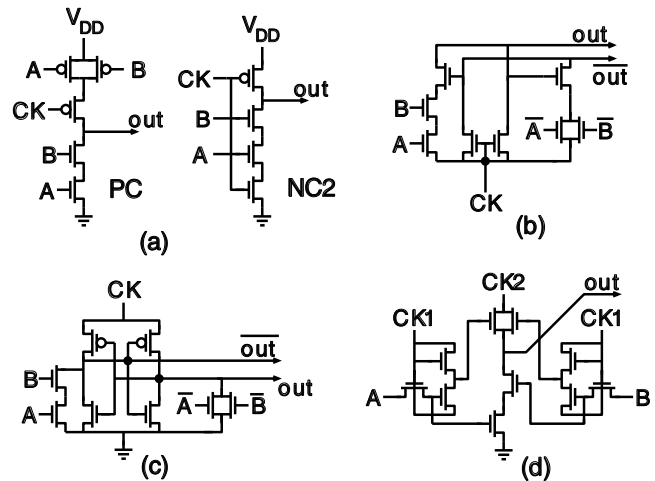


Figure 1. Two-input NAND gates for each logic family: (a)TSPC, (b) 2N-2N2D, (c) 2N-2N2P, and (d) Hot Clock nMOS. For TSPC, two versions are shown, corresponding to PC and NC2 stages.




Figure 2. Clock waveforms used for the different circuit families.

waveforms used.

All four circuit families yield a micro-pipelined type of architecture, and so are well suited to arithmetic functional blocks and other DSP types of circuits. Fig. 3 shows the basic circuit that is simulated for each logic family. It is a slice of a 32 bit adder, including representative carry-lookahead circuits. The circuit is modified as little as necessary to adapt it to each of the 4 logic families. The TSPC version is implemented in a repeated PC2-PC-NC2-NC pipeline (see (4)), with the most complex gates formed in PC2 or NC2 to reduce capacitance. Since it is a little more complex, the TSPC adder circuit is shown in more detail in Fig. 4. The carry-lookahead gates are intended to be more complex than would generally be used in a circuit design. By guaranteeing that these circuits work under worst case conditions, it is expected that circuits with simpler logic gates will function reliably.

To provide a forward-looking view of what technology will be capable of, the simulations reported here assume 0.1 μm conventional bulk MOSFETs throughout. The nominal

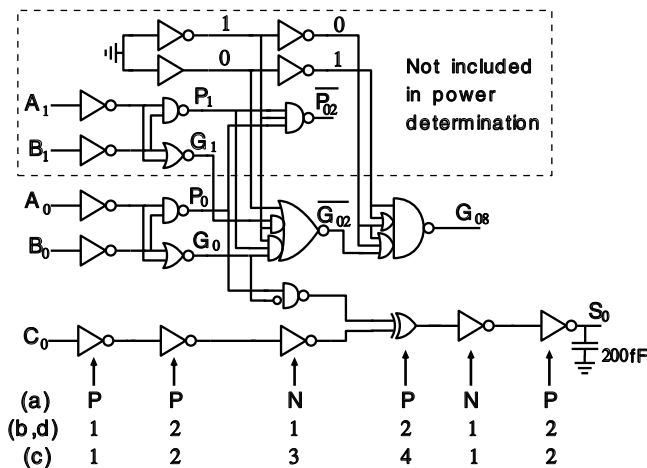


Figure 3. Schematic diagram of a bit slice, as simulated. The rows at the bottom indicate either the circuit type for logic family (a) (TSPC), or the clock phase for the other logic families. The circuits in the dotted box were simulated for margin purposes, but not included in the dissipation.

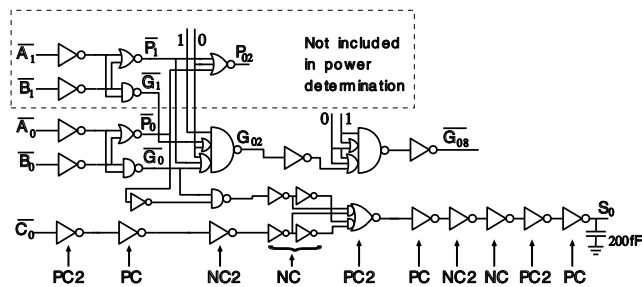


Figure 4. Schematic diagram of the TSPC bit slice, as simulated, showing in more detail the arrangement of the different blocks. In two cases there are double stages within a NC or PC block.

nFET width is taken to be 2 μm , and the nominal pFET width is 4 μm . For series stacked devices, the width is in all cases increased with the number of devices in series, and some tapering is used in the TSPC family. For 2N-2N2D the nominal nFETs are 2 μm , while in 2N-2N2P the nominal nFET and pFET devices are 1 μm and 4 μm , respectively, because there are always two parallel nFET paths for pull-down. In Hot Clock nMOS, the nominal bootstrap isolation devices are taken to be 0.6 μm , as a minimum dimension, while the nominal driven pull-up devices are 2.4 μm , and the nominal state holding pull-down devices are 1 μm .

The capacitance on the final output node, S_0 , is chosen to be large enough to represent a 1 mm line being driven to carry the signal to some other part of the chip. The inverters leading up to it are scaled up appropriately. Also, extra capacitance is added to some of the nodes to account for logic loads that would exist in the full 32 bit adder. By estimating the size of the entire adder (see Table 1), it is estimated that the typical internal nodes of the adder have short wires (5-20 μm), with capacitances of order 2 fF. Since this capacitance is much smaller than the capacitance due to the logic gates, the capacitance of the internal wires is neglected.

The necessity of using the fully driven form of Hot Clock nMOS shown in Fig. 1(d) was created by this lack of wiring capacitance. The simpler form of this logic in which the output floats when low suffered from severe charge-sharing problems. A different approach to eliminating these problems is to use CPL (Complementary Path Logic) -type circuits in conjunction with the bootstrapped output drivers, as has been analyzed by Athas and Tzartzanis (8). This type of logic requires complementary signals throughout, and also eliminates floating outputs.

The small wiring capacitance in this circuit macro is detrimental to many proposed adiabatic circuits, such as REL, 2N-2N2D, CAMOS, and some variations of hot-clock nMOS. These circuits have floating nodes during some part of the clock cycle, and only work well if there is substantial capacitance to ground. In the presence of such capacitance, they can work quite efficiently, but in its absence, they easily fail to have adequate operating margins because of unfavorable capacitive coupling between the floating nodes and the clocks.

The circuit simulations were constructed so as to locate the minimum dissipation conditions for both supply voltage, V_{DD} , and threshold voltage, V_T , while still maintaining operating margins. These margins are implemented by requiring that the circuit operate correctly not only at nominal conditions, but also under worst case conditions. Worst case conditions consisted of the supply voltage (AC or DC) being up to 10% high or 10% low, and the FETs simultaneously varying up to gate length +30% and $\Delta V_T = 20\text{mV} + 0.05 \cdot V_{DD}$, or down to gate length -30% and $\Delta V_T = -(40\text{mV} + 0.05 \cdot V_{DD})$. The V_T shifts proportional to V_{DD} are intended to account for voltage drops

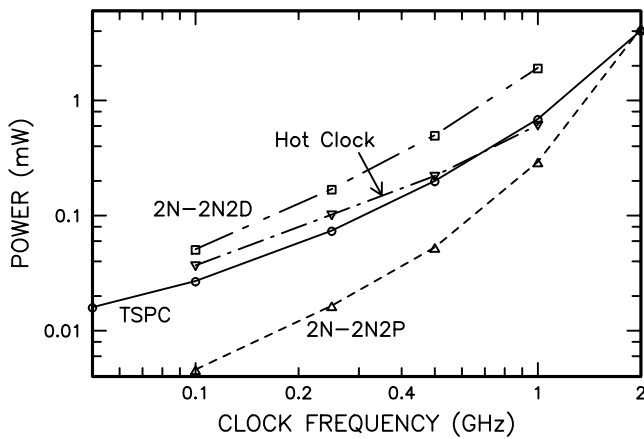


Figure 5. Worst case power versus clock frequency, for full voltage scaling of each circuit family. This is the power used by a one bit slice of a 32 bit adder.

and inductive effects in the on-chip wiring. For TSPC the power includes both the DC supply and the clock driver dissipation, while for the adiabatic families only the logic dissipation is computed, since the supply efficiencies are unknown (but are expected to be high). Correct circuit operation at each set of conditions was verified by simulating a 10 bit input sequence designed to test a variety of state change conditions.

Results

The 0.1 μm circuits simulated here all show excellent performance compared to present technology, with operating speeds up to more than 2 GHz for the best circuits. A plot of the dissipation versus frequency for all 4 families is shown in Fig. 5, and the optimized nominal voltage conditions are indicated in Fig. 6. Note that the supply voltage scales down with frequency for all of these circuits, resulting in at least quadratic power reduction with frequency over most of the range. The threshold voltages rise somewhat with decreasing frequency, as required to reduce leakage current contributions to dissipation. Table 1 shows the values of the various parameters at 500 MHz.

The worst case dissipation plotted in Fig. 5 is found by perturbing about the nominal operating conditions using the margin conditions described above, to find that set of conditions that results in the highest power. (It is this worst case power that is minimized in the procedure described above.) As can be seen, the conventional circuit performance and that of the Hot Clock nMOS are substantially comparable. The Hot Clock nMOS is probably significantly degraded by the need to dissipatively create dynamic inverses at each input, but this is necessary to achieve robust operation under the assumed circuit conditions. The CPL form of Hot Clock nMOS would avoid the dissipation associated with dynamic inverters, and might well give lower power performance.

Although the 2N-2N2D and 2N-2N2P logic families are in many ways quite similar, their performance is dramati-

Table 1: Voltages, Power, and Area, for each circuit family, optimized at 500 MHz

Logic family	V_T	V_{DD}	Power (mW)	Est. Area (μm^2)
TSPC	0.23	0.45	0.20	13000
2N-2N2D	0.285	1.45	0.50	7400
2N-2N2P	0.25	0.60	0.053	5900
Hot Clock nMOS	0.26	1.18	0.22	5300

cally different. The reason for this difference lies in the floating nodes of the 2N-2N2D. High supply voltages are required to maintain operating margins on these nodes against unwanted capacitive coupling to the following stages' clock signals. 2N-2N2P fares much better because its latching character always clamps the output nodes to the desired state. Indeed, the supply voltage for 2N-2N2P can be scaled down to 0.5 V at 100 MHz, while still maintaining margins for the complex gates in this adder.

Finally, the efficacy of the adiabatic switching technique is seen in the comparison between TSPC and 2N-2N2P.




Figure 6. (a) Optimum nominal threshold voltage versus frequency, and (b) optimum nominal supply voltage versus frequency.

Both logic families can run at very low voltages as the speed is decreased, but the energy-recovering logic achieves up to 6X less dissipation, because the energy is recovered into the power supplies instead of being converted to heat. Thus, even if the power supplies do not recover energy perfectly, it should still be possible to do better than the conventional CMOS circuit.

In conclusion, it has been shown that there is at least one adiabatic logic family that may be able to do better than conventional CMOS in the high speed micro-pipelined arena. The challenge for this family, 2N-2N2P, is the distribution and phase locking of 4 highly efficient sinusoidal power sources. None of the other adiabatic approaches evaluated here appear very promising in the low wiring capacitance regime considered.

References

1. W. C. Athas, L. J. Svensson, J. G. Koller, N. Tzartzanis, and E. Chou, "Low-power digital systems based on adiabatic-switching principles," *IEEE Trans. VLSI Sys.*, vol. 2, p. 398, 1994.
2. P. Solomon and D. J. Frank, "The case for reversible computation," *Proc. of 1994 Int'l Workshop on Low Power Design* (Napa Valley, CA), pp. 93-98.
3. T. Indermaur and M. Horowitz, "Evaluation of charge recovery circuits and adiabatic switching for low power CMOS design," *1994 Symp. Low Power Electronics* (San Diego, CA), pp. 102-103.
4. C. Svensson and J. Yuan, "Ultra high speed CMOS design," *VLSI 93*, T. Yanagawa and P. Ivey, Eds., Elsevier Science B. V., 1994, pp. 273-282.
5. A. Kramer, J. S. Denker, S. C. Avery, A. G. Dickinson, and T. R. Wik, "Adiabatic computing with the 2N-2N2D logic family," *1994 Symp. VLSI Circuits., Digest*, pp. 25-26.
6. A. Kramer, J. S. Denker, B. Flower, and J. Moroney, "2nd order adiabatic computation with 2N-2P and 2N-2N2P logic circuits," *Proc. 1995 Int. Symp. Low Power Design* (Dana Point, CA), pp. 191-196.
7. C. L. Seitz, A. H. Frey, S. Mattisson, S. D. Rabin, D. A. Speck, and J. L. A. van de Snepscheut, "Hot-Clock nMOS," in *Proc. 1985 Chapel Hill Conf. on VLSI* (Computer Science Press, 1985), pp.1-17.
8. W. C. Athas and N. Tzartzanis, "Energy Recovery for Low-Power CMOS," *Proc. 1995 Chapel Hill Conf. on VLSI*, pp. 415-429.