Electroid-Oriented Adiabatic Switching Circuits

David J. Frank and Paul M. Solomon
IBM T. J. Watson Research Center
P.O. Box 218, Yorktown Heights, NY 10598

Abstract

A dual-rail CMOS adiabatic switching circuit approach is described which follows the electroid model of Hall. These circuits can operate in either the retractile cascade or the reversible pipeline architectures. A novel adiabatic circuit technique for generating retractile cascade clock power signals from multiphase sinusoidal AC inputs is presented, along with experimental verification for a simplified version. Design optimization considerations and experimental results for a switched inductor power supply are also presented. Finally, the operation of a reversible adiabatic 4 bit ripple counter is described. Its operation is verified experimentally and its dissipation is compared with that of a voltage scaled conventional CMOS 4 bit ripple counter.

I. Introduction

The increasing commercial importance of portable battery-powered electronic applications and the increasing power density of high performance chips create a growing need for low power circuit techniques. Since power dissipation in conventional CMOS varies as CV^2, most approaches involve trying to lower the voltage, the switching factor and/or the capacitance. It has been known for many years, however, that computation does not in theory require the dissipation of 1/2 CV^2 for every logic operation. In fact, dissipation is only strictly necessary when data is destroyed,[1] and it is theoretically possible to do computation reversibly so that data need never be destroyed.[2] It has recently been recognized that these reversible computation ideas can be put into practice using conventional CMOS technology, but using a different switching paradigm--that of adiabatic switching.[3-5]

Adiabatic switching operates according to two principles. First, no FET is turned on while there is voltage across it. The two sides of the FET are first brought to the same voltage, and THEN the gate voltage is applied to turn it on. Second, when the voltage on the source of a turned-on FET is varied, the drain must be floating, and the variation must be sufficiently slow that the drain follows fairly closely, with little voltage drop occurring across the FET. Hence the name 'adiabatic switching': the drain adiabatically follows the source. For a linear ramp of time duration T (long compared to the RC time), this type of switching dissipates $CV^2\,(RC/T)$ per switching event,[6] where R is the source-to-drain resistance of the on-state FET (assumed constant) and C is the loading capacitance. For $T > 2RC$ the energy consumption is less than that of conventional CMOS, and asymptotically approaches zero as $T \to \infty$.

A variety of circuit approaches to creating adiabatic switching logic have been proposed. The present work follows the electroid switch concept of Hall,[3] using 2 state logic ('0' and '1') rather than the 3 state logic ('0', '1', and 'quiescent') which is common in other adiabatic switching approaches. Although this concept can be used to create reversible pipeline architectures, the present work focuses mostly on the retractile cascade, which is illustrated in Fig. 1. The basic idea is to achieve reversibility by retaining the original input data until one is finished with the output. This process can be cascaded to many stages, with data and supply signals being progressively applied, and then removed in the reverse order.

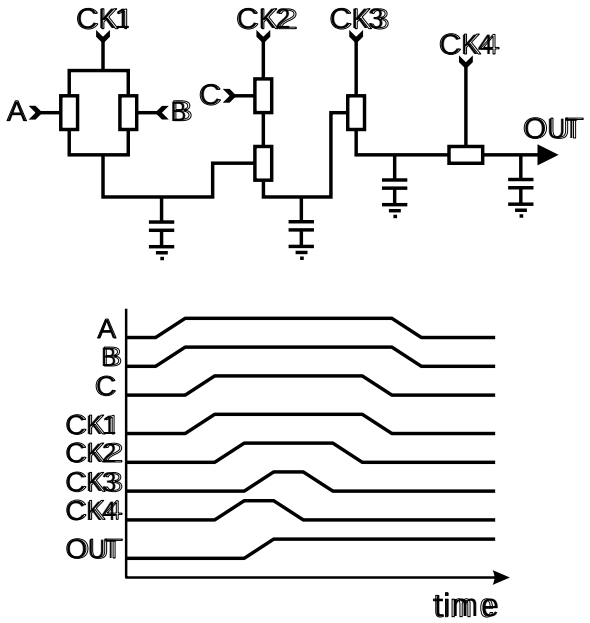

Figure 1. Adiabatic switching using retractile cascade: example circuit and waveforms. The sequence of operations is as follows: (1) apply data to open or close a series/parallel combination of switches, (2) ramp up the supply voltage, adiabatically charging the output if there is a connected path through the switches, (3) keep the data applied and the supply high until the output is no longer needed, (4) ramp the supply voltage back down, retracting the charge, if any, from the output node capacitance, and (5) remove the original data.

Electroid switch retractile cascades can be implemented using dual-rail CMOS transmission gates (T-gates).[7] T-gates have a low on-state resistance for the full range of voltages, from 0 to V_{DD}. Dual-rail signals are required, since one must have both a signal and its complement to drive the gates. Thus, four FETs are required for each electroid switch.

It has been argued that retractile cascade logic is not useful because it requires a complex set of waveforms and too many logic delays.[8] Section II describes relatively simple circuits by which such waveforms can be created, and section III describes how retractile cascade logic can be used in a system to achieve latency similar to that of conventional CMOS circuits. Experimental results of electroid switch circuitry, including a toggle flip flop circuit, are discussed in section IV.

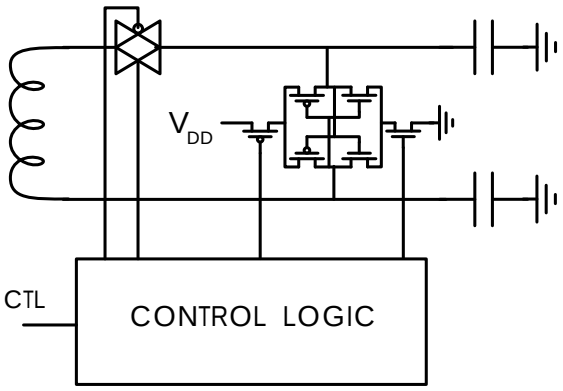

Figure 2. Schematic diagram of an energy-recovering dual-rail switched inductor adiabatic signal source.

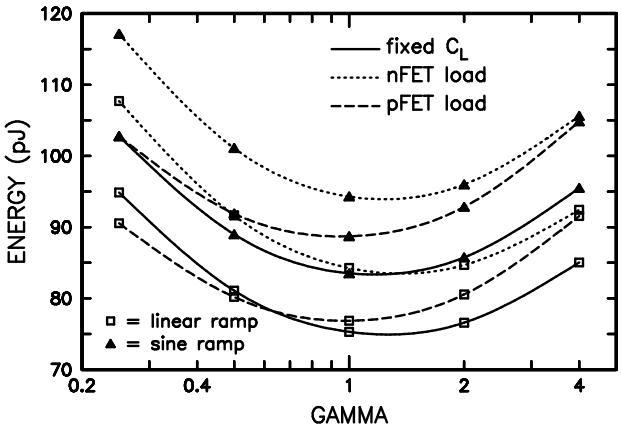

Figure 3. Minimum total switching energy at the optimum transmission gate width versus γ, the pFET to nFET width ratio. This optimization includes the energy used in charging and discharging the gates. 1.2 μm technology is assumed, and the load capacitance is 160 pF for the fixed case, and 100 mm device width for the nFET and pFET loads. The rise time is 1 μs.

II. Energy-Recovering Clock Generators

(a) Switched Inductor Supply

Retractile cascade circuits require energy-conserving variable voltage clock supplies to be applied to each circuit with the proper timing. Although simple conventionally controlled switched-inductor supplies are not expected to be very well suited to this task, due to the number needed, they do recover energy and are a useful vehicle for studying the minimization of energy dissipation in adiabatic circuits.

As Athas, et al.[9] have shown, a simple way to obtain a dual-rail energy conserving power source is to use a "ping pong" circuit, Fig. 2, which shifts energy back and forth between two essentially equal load capacitances. (A single-sided supply could be created by replacing one of the capacitors with a DC supply of $V_{DD}/2$.) Initially, the T-gate is open, one capacitor is charged to V_{DD} and the other is at zero. When the latch is deactivated and the T-gate is closed, voltage appears across the inductor and the LC circuit begins to oscillate. After a half cycle the capacitor voltages have been swapped, the T-gate is opened and the latch is activated, maintaining the outputs at their new, switched voltages. The control portion contains standard (dissipative) CMOS logic gates which turn on and off the CMOS switches in the ('non-dissipative') switch portion. The sequence of logic gates in the control section is chosen to provide the proper relative timing of the control signals to eliminate dissipation due to crowbar currents.

An initial optimization of this type of circuit for minimum dissipation has been given by Athas, et al.[8] The following analysis adds minimization with respect to the T-gate pFET to nFET width ratio γ.

There are three energy dissipation terms that depend on the T-gate dimensions. The adiabatic charging dissipation of the T-gate varies as $r_{eff}(\gamma)/\ Tw_n$, where w_n is the width of the T-gate nFET, and r_{eff} is the γ-dependent effective T-gate resistance. Both the dissipation of the conventional CMOS control circuitry and the dissipation associated with the depletion capacitance of the non-adiabatically charged end of the T-gate vary linearly with $(1+\gamma)\ w_n$. The optimum nFET width is found by minimizing the sum of these energies, yielding

$$E_{switch} = 2C_{ave}V_{DD}^2\sqrt{\frac{(1+\gamma)\ (1.9c_g + c_d)\ r_{eff}(\gamma)}{T}}$$

which displays the expected $T^{-1/2}$ dependence[8]. c_d is the depletion capacitance per unit FET width, c_g is the gate-to-channel capacitance per unit width, 1.9 is a numeric prefactor that takes into account the details of the control circuitry, and C_{ave} is the average capacitance of the load. The γ dependence of this energy can be determined from numerical simulations, and is plotted in Fig. 3 for 6 cases. As can be seen, the minimum switching energy for a sine ramp and fixed capacitive load occurs for $\gamma \approx 1.15$. Note that this is lower than the optimum ratio for conventional CMOS circuits (usually around 2). This appears to be due to the increased importance of capacitance in the adiabatic case, coupled with a decrease in the importance of low pFET resistance because it is in parallel with the nFET.

There is also dissipation in the inductor, which is given by

$$E_L = \frac{\pi^2}{2T} r_l C^2 V_{DD}^2$$

for a single transition (half an RF cycle). C is the capacitance seen by the inductor ($1/2C_L$ here), and r_l is the effective series resistance of the inductor. Parasitic capacitance C_p causes additional dissipation of $1/2C_pV_{DD}^2$ when the primary transmission gate Q1/Q2 is closed and again when the inductor is shorted. This parasitic capacitance includes the internal capacitance C_l of the inductor, and pad and package capacitance for the inductor lead. For small C_L or for long T this parasitic dissipation can dominate, but for large C_L it is usually negligible.

Combining these energies, the total dissipation is given by

$$E_{TOT} = C_L V_{DD}^2 \left(2\sqrt{\frac{2.15\ (1.9c_g + c_d)\ r_{eff}}{T}} + \frac{\pi^2}{8T} r_l C_L + \frac{C_p}{C_L} \right)$$

Note that the $C_LV_{DD}^2$ on the right is just the energy required to charge the two C_L's conventionally. Hence, the terms in parentheses give the fraction of dissipated energy relative to the conventional case, and the first term represents a Q-independent lower bound on dissipation fraction.

Experimental results for such a supply are discussed in section IV.

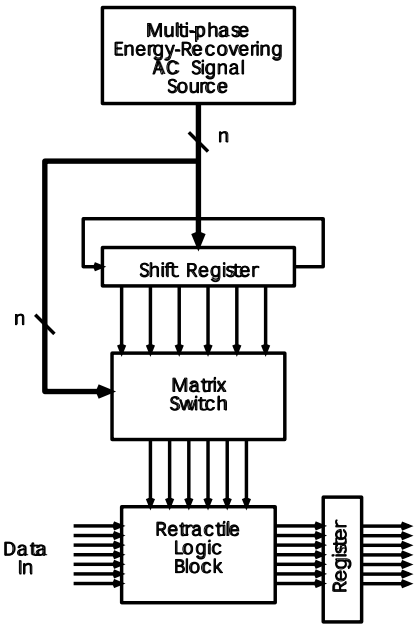

Figure 4. Block schematic of system and clock generator.

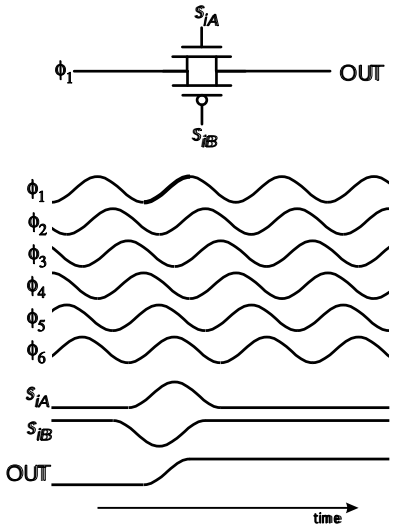

Figure 5. An example of using a T-gate to selectively pass a single rising edge of a sine wave, ϕ_1, to an output signal line. Single cycles of sine waves shifted -60° and +60° are used to control the nFET and pFET, respectively, assuming that $V_T / V_{DD} = 0.25$.

(b) Resonant AC

Since a large system needs a large number of differently timed clocks, it is impractical and inefficient to generate each of them individually by separate switched inductors. Also, switched inductor supplies do not have the highest possible efficiencies. Figure 4 is a schematic block diagram of a method of generating an arbitrary number of differently timed energy conserving clock signals from a small set of RF signals, which can be efficiently generated by resonators.

First, note that the desired waveforms (see Fig. 1) can be obtained by using switches to select segments from simple repetitive waveforms at the desired times. These segment selecting switches can be implemented as CMOS transmission gates, while state-holding switches can be single nFETs or pFETs, as appropriate. If the gate control signals for these switches were generated by conventional CMOS, one would be back to the situation with the switched-inductor supply in II(a) where dissipation varies as $T^{-1/2}$. It is possible, however, to generate these transmission gate control signals adiabatically from single cycles of multiphase sinusoidal input clocks. An example of this is shown in Fig. 5 where a single rising edge of ϕ_1 is selected and passed to the output. For 6 phase sinusoidal inputs, the phase that precedes ϕ_1, namely ϕ_6, will have reached 1/4 of V_{DD}, the peak-to-peak voltage swing, at the instant when ϕ_1 is at its minimum. Thus, to switch the T-gate on at this instant, one can set V_T to $0.25V_{DD}$, and then use a single cycle of ϕ_6 to control the T-gate nFET. This will turn the nFET on during the first 2/3 of the transition. Similarly, a single cycle of ϕ_2 can turn the pFET on during the latter 2/3 of the transition. (They are both on during the center 1/3.) For purposes of explanation, switching is assumed to occur at exactly V_T. Conceptually, other numbers of phases may be used, with suitable adjustment of V_T.

To generate a complex set of output waveforms, one must have single cycles available for switching the output at almost every possible rising or falling edge of the simple periodic input clocks. The best way to provide this is to create a switching pulse for *every* possible rising or falling edge. This can be done using a chain of adiabatic switching logic, in which the stages work together so as to progressively switch each stage in the chain, just once. Adiabatic switching circuitry is good at this, and there are many ways

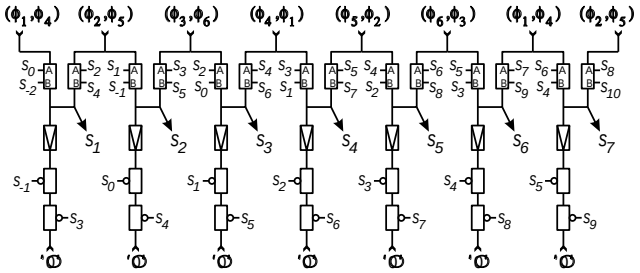

Figure 6. Example of an adiabatic shift register chain. Note that this is a dual rail circuit; each line represents a pair of signal wires. The inputs are 6 phase sine waves as in Fig. 5, and $V_T / V_{DD} = 0.25$. The lower portion of the circuit serves to enable static operation.

in which it can be done. Fig. 6 shows one example, using the symbols defined in Fig. 7. This circuit is in effect a generalization of the shift register described by Athas, et al.[9] Each of the wires shown in the schematic represents two wires (for the two rails of the logic), and the electroid switches connected to fixed logic level '0' only require 2 FETs (1 nFET and 1 pFET) since the other two could never be turned on. This fixed logic level represents DC supply connections, 0 V on the first rail and V_{DD} on the second rail. Inversion (indicated by the small circles) is always available for free in this dual-rail circuit, simply by switching the rails. Using 6 phase sine waves with peak-to-peak amplitude V_{DD} and $V_T / V_{DD} =$ 0.25, the output pulses of this chain are similar to the s_i control signals shown in Fig.5 except that each pulse has a flat top for 1/6th of a cycle.

The operation of this chain can be seen by considering s_4, the output pulse of the fourth stage of the chain. The 'true' rail of s_4 rises with one of the rising edges of ϕ_4, and then falls with the next falling edge of ϕ_5. The electroid switch connecting ϕ_4 to s_4 is controlled by the preceding pulse signals s_3 and s_1, which lead the rising edge by 1/6th and 3/6ths of a cycle, respectively, as required for these flat-topped signals. The electroid switch connecting ϕ_5 to s_4 is controlled by the following pulse signals s_5 and s_7, which lead the falling edge by 3/6ths and 1/6th of a cycle, respectively, again as required. After the pulses have passed by, s_1, s_3, s_5 and s_7 are all

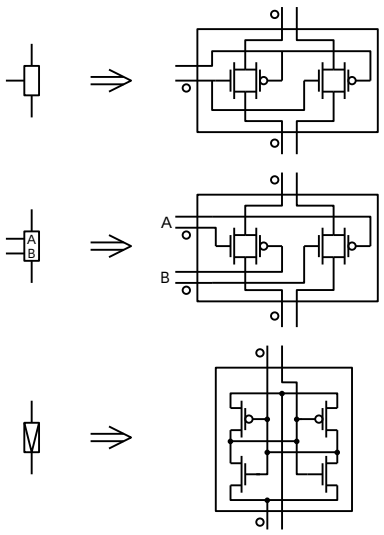

Figure 7. Definitions of symbols in Fig. 6.

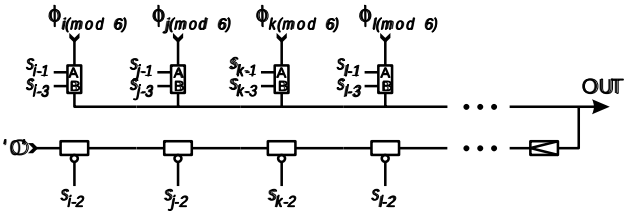

Figure 8. Schematic of matrix switch for a given output line including latching. The ϕ input waveforms are switched to the output line at times that are determined by the *s* pulses. The output rises at s_i, falls at s_j, rises again at s_k, falls again at s_l, and so on.

low, and s_4 will undergo no further state changes until the chain is triggered again. The flat-topped character of the pulses allows the single signal s_2 to deactivate the latching behavior of the floating latch (Fig. 7c) during the rising edge of s_4, and the single signal s_6 to continue the deactivation during the falling edge. Used in this manner, the floating latch behaves reversibly and adiabatically, but does dissipate more than an electroid switch because it passes into and back out of a high impedance state. The same process outlined above happens progressively at each stage of the chain. The chain can be triggered by conventional logic, or by other adiabatic logic, or one can connect the end back to the beginning, so as to create a ring which repetitively generates the desired output waveforms.

Given the marching series of single pulses, one can obtain an output waveform with a transition at any desired time simply by using the appropriate single pulse to activate a transmission gate between the sinusoidal input and the output at the desired time as shown in Fig. 8. This leaves the matter of latching. There are three possibilities. (1) No latching is used. The outputs could simply be left floating between transitions. This solution may not be satisfactory, however, if other signals are capacitively coupled to the output in question while it is floating, since such signals will cause the output to drift. (2) One can use many switches to DC, and as many of the single cycle waveforms as necessary to cover the time between switching transitions. (3) The simplest approach, shown in Fig.8, uses the floating latch. The flat-topped signal pulses generated by the shift register in Fig. 6 are used, and the second preceding signal is used to deactivate the latching for each transition.

The preceding explanations assume that the FETs turn on quite sharply when the gate-to-source voltage reaches V_T. In reality, of course, the FETs make a gradual transition from high resistance to low resistance as they turn on, at the same time as the supply signal on the source is passing through its minimum or maximum. The overall dissipation in this circuit depends on the relative rates at which the resistance drops and the current (proportional to *dv/dt* of the supply signal) rises. This should be evaluated using numerical simulations to obtain more precisely the optimum V_T / V_{DD} ratio. For example, simulations show that the optimum supply voltage is around 4.5 V for the 1.2 μm technology used in Sec. V, where the threshold voltages are 1.4 V. Under these conditions, the 6 phase shift register chain dissipates about 1.8x more energy than it would if the same T-gates were fully on for the entire duration of every transition.

III System Approach

The schematic block diagram in Fig. 4 shows not only the clock generator concept, but also how a system might be configured to use this type of reversible logic.[7] Retractile combinatorial logic operates upon the contents of registers or latches, and produces new values that are stored in additional registers or latches. These results become the operands for the next stage of computation, while the first stage retracts. The clock supply signals are generated adiabatically as described above.

After the logic operation has been retracted, the input signals must be removed. They could be removed by inverse functions as in a reversible pipeline architecture, but this may incur excessive circuit overhead[8] and not actually save any energy. Pragmatically, they can simply be erased before the next data is written into them. In this case the logic design and architecture correspond quite closely to conventional design practice, making it straightforward to implement. The latency involved in this approach is essentially the same number of logic delays as would be involved in the conventional design, except that here the logic delays are the transition times, which are deliberately lengthened to achieve energy recovery. The throughput is reduced a factor of two because of the retractile process, but can be increased by parallelism at a circuit cost that may compare favorably with the reversible pipeline architecture.

The energy dissipation expected for such adiabatic logic has been discussed in Ref. 7 in comparison to conventional logic. The adiabatic clock signal generation is expected to dissipate energy at a rate similar to that of the adiabatic logic in a well balance system, resulting in a doubling of the total system dissipation relative to that of the logic alone.

Conventionally it would cost $1/2CV_{DD}^2$ energy per bit to erase the data in the registers. It is possible, however, to reduce the cost of erasure down to $1/2CV_T^2$ by using the data itself to indirectly control the switch through which the node is adiabatically discharged.[4] The switch opens, however, when the voltage gets down to the threshold voltage, and a shorting switch must be used to remove the last of the energy. The energy loss may be minimized by buffering the latch from the external load using adiabatic buffers, so that $C_{latch} << C_L$, where C_{latch} is the capacitance of the latch and C_L is the average capacitive load of a logic circuit.[10] N_{stg}, the number of stages one can use before this dissipation becomes dominant is determined from

$$\frac{N_{stg} r (C_L V_{DD})^2}{T} = \frac{C_{latch} V_T^2}{2}$$

Using reasonable values as an example, if $V_{DD} = 4V_T$, $C_{latch}/C_L = 0.2$, and $rC_L/T = 0.001$, then $N_{stg} = 6$, indicating that erasure may be acceptable after blocks containing more than 6 stages of logic.

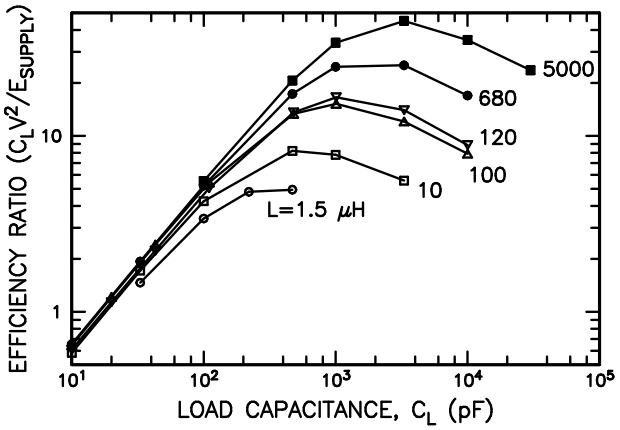

Figure 9. Experimental results for switched inductor supply: energy reduction ratio versus capacitive load.

Note that for $rC_L/T > 0.006$, all stages may be latched, so that this is a viable option (as in the approach of Denker, et al.[11]) for realizing a high throughput at modest energy savings.

IV. Experiments

Using 1.2 μm ground rules, a small test chip was fabricated at a Si foundry to test some of the adiabatic switching concepts. Among other things, the chip included a switched-inductor power supply, four bit counters, and a simplified clock generator circuit.

The switched-inductor supply used a value of w_n = 1 mm, which is large enough to drive a reasonable load yet does not occupy too much area. Since the dependence of dissipation on γ is weak, γ = 1 was chosen for the experimental design to simplify the layout. ($r_{eff} \approx 17\Omega$mm for γ = 1.) Measured results for this supply are shown in Fig. 9 in the form of energy dissipation reduction factor versus capacitive load, for several different inductors. The control logic dissipation is nearly constant at 175 pJ per output transition, and the internal, package, and test fixture capacitances contribute about 13 pF to the resonant frequency. Note that the peak efficiency for the 5 mH inductor corresponds to an energy reduction factor of 45x compared to the energy that would be dissipated conventionally to charge or discharge the two C_L loads, representing recovery of 97.8% of the energy. The low capacitance asymptote corresponds to the C_L-independent energy dissipation associated with the control circuitry and the inductor and pad parasitics.

A reversible dual-rail toggle flip-flop was constructed from CMOS bi-directional switches using the circuit shown in Fig. 10. As in Fig. 6, each of the wires shown in the schematic represents two wires. The fixed logic level ‘1’ represents DC supply connections, V_{DD} on the first rail and 0 V on the second rail. Note that this circuit uses only 2-state electroid switch logic, contrary to some speculation that such a circuit function necessarily requires 3-state logic (‘0’, ‘1’, and quiescent).

The operation of the flip-flop is illustrated in Fig. 11, where the signals on the inputs, the internal nodes and the outputs are shown versus time. For clarity, only the signal on the first of the two rails of each node is shown. Thus, each of the switches is closed when the input is in the ‘high’ or logical ‘1’ state, and open when the input is ‘low’ or at logical ‘0’. Switches S1-S4 allow the connection of CK1 or $\overline{CK1}$ to Q and QB. Only one of these switches is closed during any one transition of CK1, as is determined by the signals A, B, C, and D. The states of signal nodes A-D are determined by switches S5-S8 which connect them to CK2 at the appropriate time. Half switches H1-H8 allow static operation by holding

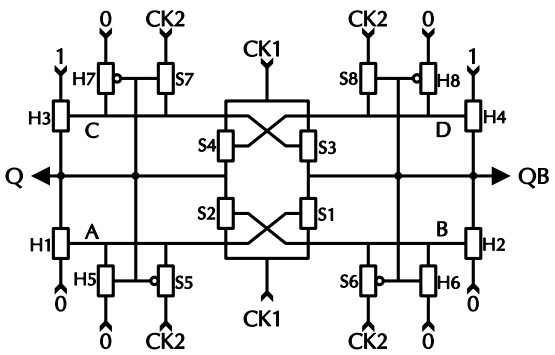

Figure 10. Schematic diagram for a reversible adiabatic toggle flip-flop using 2-state dual-rail electroid switches.

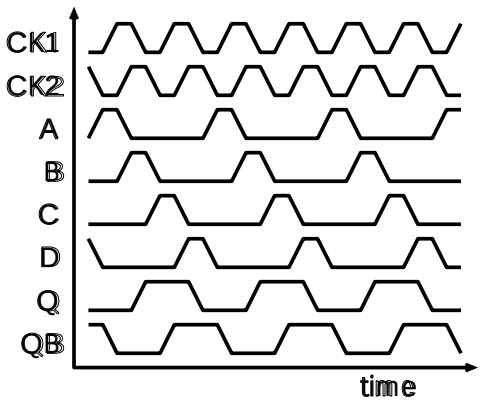

Figure 11. Waveforms for the toggle flip-flop of Fig. 10.

their respective nodes at the proper potentials during times when they would otherwise be floating. The half switches might be unnecessary under some conditions in flip-flops that were regularly clocked, thus resulting in dynamic operation of the flip-flop. As for reversibility, the behavior of this circuit versus time can be exactly reversed by reversing the relative phase of the two clocks, CK1 and CK2.

Four of these toggle flip-flops were built on the experimental chip, connected together to form a 4 bit ripple counter. Functional testing revealed the expected binary counter sequences, but several attempts to measure the dissipated power by use of the switched inductor supplies were unsuccessful due to the very small capacitive load presented by the counter. The simulated dissipation versus input frequency is shown in Fig. 12, along with experimentally measured dissipation for a conventional CMOS ripple counter fabricated on the same chip, with comparable FET widths. The conventional counter is powered in a voltage scaled manner, the voltage at each frequency being just high enough to make it operational at that frequency. As can be seen, the voltage scaled CMOS counter has lower dissipation at all frequencies above about 1 MHz, but the adiabatic counter dissipates less below 1 MHz. This is not a particularly favorable comparison for most applications, but is probably a fair example of the trade-offs between these two types of circuits.

The simplified clock generator circuit was designed to use four phase trapezoidal input clocks rather than sine waves. The shift register had 12 stages and was configured as a ring, with some triggering circuitry to allow it to be started. Four different output waveforms were generated, with patterns ‘0000/1\00000’, ‘000/111\0000’, ‘\0/1\0/11111’, and ‘\00/11\00/11’, where ‘/’ and ‘\’

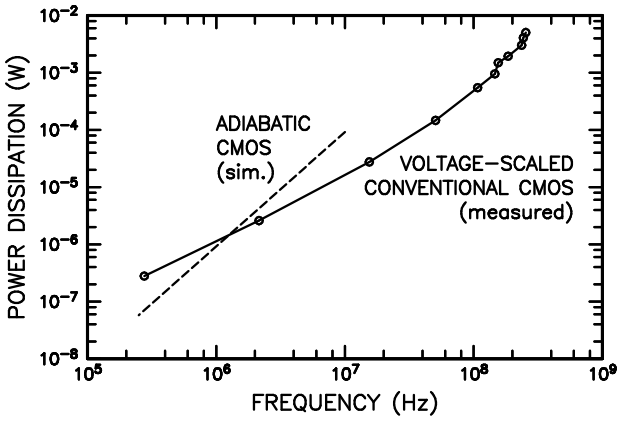

Figure 12. Dissipation versus frequency for an adiabatic switching 4 bit ripple counter (simulated) and for a conventional CMOS 4 bit ripple counter (measured).

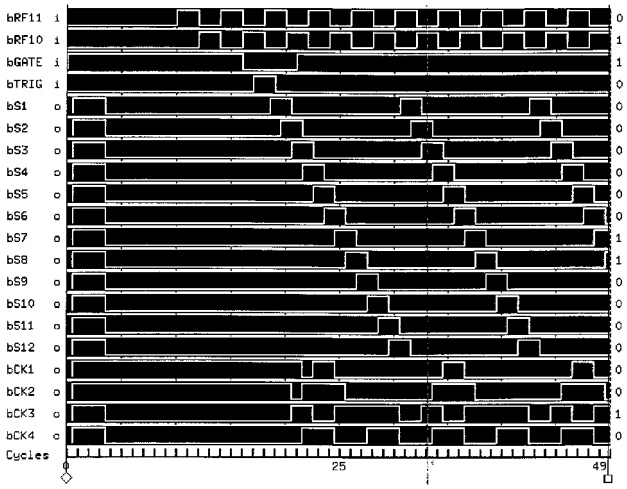

Figure 13. Timing diagram for a simplified clock generator circuit fabricated on test chip. This data was taken on a digital tester to verify correct functioning of the circuit. The time per division is 1 μs.

indicate up and down transitions respectively. Functional testing revealed that the circuit behaved as expected. Fig. 13 shows the output of a digital tester used to verify the proper functioning of the circuit. The first two signals are two of the four input clocks, the next two signals initiate the ring, and the next 12 signals are the progressive pulses being created by the shift register. The last 4 signals are the generated output signals, which show the intended sequences of rising and falling edges.

V. Conclusions

Using the electroid concepts of Hall, implemented using dual-rail CMOS transmission gates, it is possible to design almost any type of digital electronic system. As an example, a 4 bit reversible adiabatic ripple counter has been described. Such systems can use the reversible pipeline approach, or the retractile cascade approach, in which case it is usually most efficient to erase register data after it has been used. This retractile cascade adiabatic switching logic is straightforward to build and use, since it is quite analogous to conventional CMOS design, but previously it has lacked a suitable supply scheme. The clock generator circuitry presented here rectifies this situation, and encourages the further exploration of this type of circuitry.

Looking to the future, it appears that scaling considerations may favor the development of energy-recovery logic. As long as supply voltage can continue to be reduced, conventional CMOS will continue to offer the lowest energy approach to VLSI for most applications, but when V_{DD} can no longer be decreased due to lower limits on threshold voltages because of leakage constraints, then energy-recovery circuits can gain an advantage. This follows from the observation that in this regime CMOS circuits only decrease their energy cost per node linearly with technology scaling, whereas adiabatic switching circuits can decrease their energy cost per node by the cube of the technology scaling. It may be, however, that wire resistance effects will weaken or cancel this advantage, so it remains to be seen how this ‘competition’ will develop.[12]

References

1. R. Landauer, “Irreversibility and Heat Generation in the Computing Process,” *IBM J. Res. Devel.* **5**, 183 (1961).

2. C. H. Bennett, “Notes on the History of Reversible Computation,” *IBM J. Res. Devel.* **32**, 16 (1988).

3. J. S. Hall, “An Electroid Switching Model for Reversible Computer Architectures,” *Proc. Workshop on Physics of Computation, PhysComp’92*, (Dallas, Texas, Oct. 1992), p. 237.

4. J. G. Koller and W. C. Athas, “Adiabatic Switching, Low Energy Computing, and Physics of Storing and Erasing Information,” *Proc. Workshop on Physics of Computation, PhysComp’92*, (Dallas, Texas, Oct. 1992), p. 267.

5. R. C. Merkle, “Reversible Electronic Logic Using Switches,” *Nanotechnology* **4**, 21 (1993).

6. 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 *Proceedings of the 1985 Chapel Hill Conference on VLSI* (Computer Science Press, 1985), pp.1-17.

7. 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.

8. 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.* **2**, 398 (1994).

9. W. C. Athas, L. J. Svensson, J. G. Koller, N. Tzartzanis, and E. Chou, “A Framework for Practical Low-Power Digital CMOS Systems Using Adiabatic-Switching Principles,” *Proc. of 1994 Int’l Workshop on Low Power Design* (Napa Valley, CA), pp. 189-194.

10. S. G. Younis and T. F. Knight, Jr., “Asymptotically Zero Energy Split-Level Charge Recovery Logic,” *Proc. of 1994 Int’l Workshop on Low Power Design* (Napa Valley, CA), pp. 177-182.

11. J. S. Denker, S. C. Avery, A. G. Dickinson, A. Kramer, and T. R. Wik, “Adiabatic Computing with the 2N-2N2D Logic Family,” *Proc. of 1994 Int’l Workshop on Low Power Design* (Napa Valley, CA), pp. 183-187.

12. D. J. Frank and P. Solomon, “Energy Recovery In Future VLSI Logic,” pub. in Japanese in *Low Power LSI* by Nikkei Microdevices (Tokyo, 1994).