Adiabatic Dynamic Logic

Alex G. Dickinson and John S. Denker

Abstract— **With *adiabatic* techniques for capacitor charging, theory suggests that it should be possible to build gates with arbitrarily small energy dissipation. In practice, the complexity of adiabatic approaches has made them impractical. We describe a new CMOS logic family—adiabatic dynamic logic (ADL)—that is the result of combining adiabatic theory with conventional CMOS dynamic logic. ADL gates are simple, general, readily cascadable, and may be fabricated in a standard CMOS process. A chain of 1000 ADL inverters has been constructed in 0.9 μm CMOS and successfully tested at 250 MHz. This result, together with comprehensive circuit simulation, suggest that ADL offers an order of magnitude reduction in power consumption over conventional CMOS circuitry.**

I. Introduction

POWER consumption has become a critical concern in both high-performance and portable applications. Methods for power reduction based on the application of *adiabatic* techniques to CMOS circuits have recently come under renewed investigation. In thermodynamics, an adiabatic energy transfer through a dissipative medium is one in which losses are made arbitrarily small by causing the transfer to occur sufficiently slowly.

If circuits can be made to operate in an adiabatic regime with consequently low energy dissipation, then the energy used to charge the capacitive signal nodes in a circuit may be recovered during discharge and stored for reuse. The efficiency of such a circuit is then limited only by the "adiabaticity" of the energy transfers.

Conventional CMOS circuits are pathologically nonadiabatic. Capacitive signal nodes are rapidly charged and discharged (the energy transfer) through MOS devices (the dissipative medium). At times, the full supply potential appears across the channel of the device, resulting in high device current and energy dissipation.

Younis and Knight [1] and Merkle [2] have proposed adiabatic logic families with near zero dissipation. However, these schemes require reversible logic gates, and in the case of Younis and Knight, each gate requires 16 times the number of devices as conventional logic. Koller and Athas [3], [4] have described a scheme based on transmission gates with limited cascadability. Hinman and Schlecht [5] describe a means for performing adiabatic logic with bipolar devices, or the complex emulation of bipolar devices with NMOS devices. Seitz *et al.* [6] recognized the benefits of adiabatic charging, but used a two-phase clocking strategy that precluded adiabatic operation.

We have developed a new logic family, adiabatic dynamic logic (ADL), based on the application of adiabatic techniques to dynamic CMOS logic structures. ADL is not *completely* adiabatic, but offers an order of magnitude reduction in power consumption using simple, readily cascadable gates that may be fabricated in a standard CMOS process. In this paper, we will describe adiabatic charging, ADL gates, a resonant power supply, and the results of our simulations and test fabrications.

Manuscript received August 11, 1994; revised November 9, 1994.
The authors are with AT&T Bell Laboratories, Holmdel, NJ 07733 USA.
IEEE Log Number 9408736.

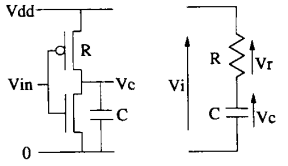

Fig. 1. A static CMOS inverter and its equivalent circuit for the case where the capacitor C is being charged through a device of on-resistance R.

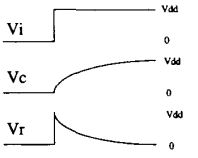

Fig. 2. Voltage waveforms present in the equivalent circuit when charging the capacitor from 0 V to V_{dd} in a conventional manner.

II. Adiabatic Switching

A. Conventional Charging

The dominant factor in the dissipation of a CMOS circuit is the *dynamic* power required to charge capacitive signal nodes within the circuit. Fig. 1 shows a basic CMOS inverter, together with an equivalent circuit of the charging mechanism. Fig. 2 shows the voltage waveforms present when the input of the inverter swings from *high* to *low,* causing the capacitor C to begin charging. At the instant of switching, the full supply potential appears across the on-resistance R of the p-type device—the waveform then decays as the capacitor is charged to V_{dd}. To charge the signal node capacitance C from a supply of potential V_{dd}, a charge $q = CV_{dd}$ is taken from the supply through the p-type device. The total energy $E_T = qV_{dd} = CV_{dd}^2$.

Only half of the energy is applied to storing the signal on the capacitor—the other $\frac{1}{2}CV_{dd}^2$ is dissipated as heat, primarily in the device on-resistance R. Note that the dissipation is independent of this resistance: it is a result of the capacitor charge being obtained from a constant voltage source V_{dd}.

To drive the inverter output low, the n-type device is used to discharge the $\frac{1}{2}CV_{dd}^2$ energy stored in capacitor C by short circuiting the capacitor and dissipating energy as heat.

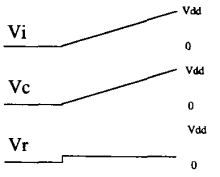

Fig. 3. Voltage waveforms present in the equivalent circuit when charging the capacitor from 0 V to V_{dd} in an adiabatic manner.

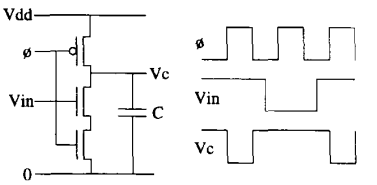

Fig. 4. A conventional dynamic CMOS inverter and associated clock, input, and output waveforms for a 1–0–1 input and resulting 0–1–0 output.

Hence, the total charge/discharge cycle has required an energy CV_{dd}^2—half being dissipated in charging, and half being used for information storage before it too is dissipated during discharge.

B. Adiabatic Charging

Adiabatic switching can be achieved by ensuring that the potential across the switching devices is kept arbitrarily small. In Fig. 2, it can be seen that the potential V_r across the switch resistance is high in the conventional case because of the abrupt application of V_{dd} to the RC circuit.

Adiabatic charging may be achieved by charging the capacitor from a time-varying source, as shown in Fig. 3. This source has an initial value of $V_i = 0$ V—the ramp increases towards V_{dd} at a slow rate that ensures $V_r = V_i - V_c$ is kept arbitrarily small. This rate is set by ensuring that the period of the ramp $T \gg RC$.

In fact, the energy dissipated is (From Athas [4])

$$E_{\text{diss}} = I^2RT = \left(\frac{CV_{dd}}{T}\right)^2 RT = \left(\frac{RC}{T}\right)CV_{dd}^2. \quad (1)$$

A linear increase in T causes a linear decrease in power dissipation. Adiabatic discharge can be arranged in a similar manner with a descending ramp.

Now, if T is sufficiently larger than RC, energy dissipation during charging $E_{diss} \rightarrow 0$, and so the total energy removed from the supply is $\frac{1}{2}CV_{dd}^2$—the minimum required to charge the capacitor and hence hold the logic state. This energy may be removed from the capacitor and returned to the power supply adiabatically by ramping V_i back down from V_{dd} to 0 V. As a result, given a suitable supply, it should be possible then to charge and discharge signal node capacitances with only marginal net losses.

Note that the RC time constant of a typical CMOS process is about 100 ps. If we set T to ten time constants, the resulting delay through an adiabatic gate would be 1 ns.

III. An Adiabatic Dynamic Logic Gate

A. Conventional Dynamic Logic

A conventional dynamic CMOS inverter is shown in Fig. 4. When the clock Φ is *low*, the output is always pulled to V_{dd}; when the clock goes *high*, the output remains *high* if the input is *low*, or goes *low* if the input is *high*. Note that if the output remains *high*, it does so not as a "driven" value, but as charge stored on the output node capacitance C.

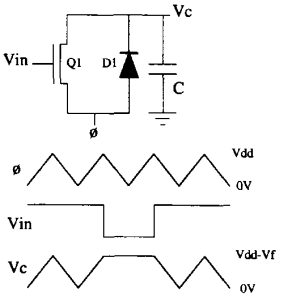

Fig. 5. An adiabatic dynamic logic CMOS inverter and associated clock, input, and output waveforms for a 1–0–1 input and resulting 0–1–0 output.

Clearly, there is abrupt (nonadiabatic) charging and discharging of the output node in this circuit. Also (as with static CMOS circuits), there is an additional loss mechanism resulting from direct current flowing from the supply to ground during the switching operation, when both n-type and p-type devices are conducting simultaneously.

In working to develop an adiabatic logic family, our attention was drawn to dynamic circuits because of their time-sequential nature. All nodes are charged *(precharged)* during one half cycle of the clock. The same nodes then remain charged or are discharged *(evaluated)* during the next half cycle. Conceptually, at least, if one half cycle of the clock were a rising ramp, and the other a falling ramp, one might be able to produce adiabatic operation.

B. The ADL Inverter

The ADL inverter is shown in Fig. 5. It is comprised of a diode $D1$ with forward drop V_f, a transistor $Q1$ of threshold V_t, and an implicit capacitance C that represents the load of subsequent gates. Note the simplicity of the circuit, and the lack of dc paths to ground. All power is supplied by the time-varying "clock supply" Φ.

As with a conventional dynamic CMOS gate, there are two basic stages of operation. In the *precharge* stage, the clock-supply ramps from 0 to V_{dd}, precharging the output through the forward-biased diode $D1$ to $V_{out} = V_{dd} - V_f$. In the *evaluate* stage, the clock supply ramps down from V_{dd} to 0. If the input is *high*, $Q1$ conducts, and so the output is driven *low*, V_{out}=0. If the input is *low*, the output remains *high* at $V_{out} = V_{dd} - V_f$.

Clearly, an inversion has occurred between input and output. In addition, all transitions of the output node are driven by a ramped voltage supply, and are hence potentially adiabatic.

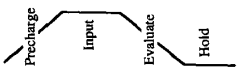

Fig. 6. The four-phase clock waveform required for interconnecting ADL gates. The vertical arrows indicate the flow of data from the *hold* time of one gate to the *evaluate* time of a subsequent gate.

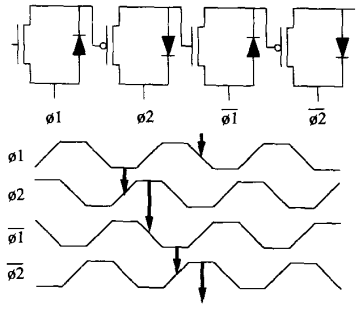

Fig. 7. A "cycle" of four interconnected ADL inverters and their associated four-phase clocks.

IV. Systems of ADL Gates

A. Cascaded Gates

Adiabatic circuit operation requires that only small potentials appear across the channel of a device while it is conducting. This must be taken into consideration when interconnecting gates, as it implies that devices should never "turn on" while there is a nonzero voltage present across the device channel.

Between *precharge* and *evaluate,* there are only small potential differences present in the ADL inverter circuit: both clock-supply and output nodes are near V_{dd}. We make use of this observation by adding a stage of constant clock voltage between these two stages. During this *input* stage, the input signal may safely transition without causing a nonadiabatic transition within the circuit. Similarly, we add a *latch* stage between *evaluate* and *precharge* to ensure that the output is latched for a finite time. The resulting clock waveform is shown in Fig. 6.

To connect two gates, it is necessary to synchronize their respective clock supplies such that the *latch* stage of the first gate coincides with the *evaluate* stage of the second. However, as the first gate begins precharging, it will cause the input of the second gate to undergo a *low* to *high* transition, resulting in a nonadiabatic transition in the second gate. The solution to this problem is shown in Fig. 7: the second stage is constructed as an inverted version of the first, precharging *low* and evaluating *high.* The precharge of the first gate can then only cause the input device of the second gate to *stop* conducting: an allowable adiabatic transition.

Because there are four stages in a single clock cycle *(precharge, input, evaluate, latch),* it is necessary to place four gates in series before the last gate may feed the first. The clock-supply waveforms consists of two trapezoidal waveforms separated by a quarter period and their complements. A four-gate "cycle" and its associated clock waveforms are shown in Fig. 7.

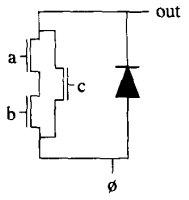

Fig. 8. A complex ADL gate implementing $\overline{a \cdot b + c}$.

B. Clock Voltage Considerations

A requirement of the particular system of gate interconnection that we have described is that when a gate is in its *precharge* state, the voltage present at its output must be such that it not turn on the input device of the subsequent gate, i.e., $V_f < V_{th}$. In the case of a precharge *high* gate, for example, if the precharge voltage $(V_{dd} - V_f)$ is less than the turn-on point of the subsequent p-type device $(V_{dd} - V_{th})$, it will turn on the input p-device of the subsequent gate during precharge. We have found two means of the avoiding this problem. The first is to ensure that the forward drop of the diode V_f is less than the threshold voltage V_{th} of the p-type or n-type device to which it is connected. The second technique is to offset the clock voltages: for example, the first and third gates in Fig. 7 would have a positive offset compared to the clocks for the second and fourth gates to ensure correct operation.

C. Complex ADL Gates

Extension of the ADL inverter to more complex NAND/NOR gates is shown in Fig. 8. The complexity of an ADL gate is low: n FET's + 1 diode for an n input gate compared to $2n$ FET's and $(n + 2)$ FET's for conventional static and dynamic CMOS, respectively. In terms of routing, an ADL system requires four clock lines and two tub bias lines—comparable to some conventional CMOS systems based on two clocks their complements, and two power lines.

In addition to the synchronous interconnection of ADL gates, it is possible under certain conditions to combine gates asynchronously, as shown in Fig. 9, if they have the same precharge type. This arrangement has the advantage of adding an optional inversion function at the output or input of a gate, and potentially doubling the number of gates available for operations in a four-phase clock cycle (eight gates rather than four). The arrangement is, however, sensitive to the rise and fall time of the clock edges. If, for example, the right-hand gate is being left *high* during evaluation, the output of the left-hand gate must be evaluating *low.* If the output of the left-hand gate lags the clock as it falls towards 0, the input device of the right-hand gate may turn on, causing the output of that gate to go *low* erroneously. However, given that ADL in general is based on using relatively slow clock edges to ensure adiabatic operation, this has not been found to be a major limitation.

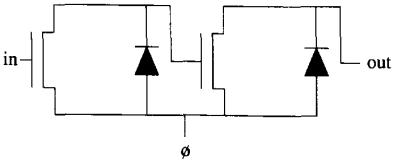

Fig. 9. A pair of gates with identical precharge types (positive) connected to a single clock phase.

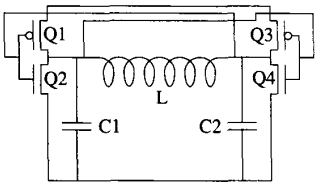

Fig. 10. A resonant oscillator "clock supply" for an ADL system. Note that there are no active devices in the main current path between the two capacitive loads.

V. Energetics

Further insight into an ADL circuit can be gained by thinking of its *evaluate-precharge* sequence as being one of a "write" being followed by an "erase." During *evaluate*, the output node is either left alone or discharged (written). During *precharge*, the diode $D1$ observes the state of the output, and in effect makes a decision as to whether or not the output needs to be erased. It is the use of the diode in this decision-making process that dissipates most of the energy in the gate.

The object of using reversible logic (Younis and Knight [1], for example) is to eliminate the cost of erasing the information held in latches. In ADL, this erasure does not come for free, but it is cheap: the diode makes a purely *local* erasure decision.

As charge q is drawn from the supply V_{dd} to recharge C_{ADL}, the diode drop V_f causes an energy loss:

$$E_{\text{pre}} = qV_f = C_{\text{ADL}}V_{dd}V_f. \tag{2}$$

The energy dissipated during evaluation (discharge) is (1)

$$E_{\text{eval}} = \left(\frac{RC_{\text{ADL}}}{T}\right)C_{\text{ADL}}V_{dd}^2 \tag{3}$$

which will be much smaller than $C_{ADL}V_{dd}V_f$ for $T \gg RC_{ADL}$. The ratio of conventional to ADL energy dissipation is then

$$\frac{E_{\text{conv}}}{E_{\text{ADL}}} = \frac{C_{\text{conv}}V_{dd}^2}{C_{\text{ADL}}V_{dd}V_f} = \left(\frac{V_{dd}}{V_f}\right)\left(\frac{C_{\text{conv}}}{C_{\text{ADL}}}\right). \tag{4}$$

For typical values of $V_{dd} = 5$ V, $V_f = 0.5$ V, and for our measured capacitance ratios of between 1 and 2, we find that the potential power savings cover the range of 10–20 over conventional CMOS logic, before allowing for clock generation losses.

VI. Adiabatic Clock Generation

By itself, the use of adiabatic charging will only improve power consumption by a factor of two: the $\frac{1}{2}CV_{dd}^2$ of resistive heat dissipation is removed, but the $\frac{1}{2}CV_{dd}^2$ stored in the capacitor is still lost during discharge. If we are to realize the full potential of the technique, it is necessary to return that energy to the supply for use in a subsequent cycle.

A resonant circuit will serve to drive the largely capacitive load presented by the ADL system if a sinusoid may be substituted for the trapezoidal supply waveform. We have, in fact, verified that ADL gates will operate from sinewave sources in both simulations and circuit testing.

Fig. 10 shows a simple resonant oscillator for driving two complementary clock-supply nodes of an ADL system, represented by the capacitors C_1 and C_2. If, at time $t_0, V_1 = 5$ V and $V_2 = 0$ V, then the inductor L will serve to pump charge from C_1 to C_2 and back again. The waveform presented across each capacitor will be a sinewave with an amplitude that decreases with time due to the resistive losses in the inductor and ADL circuits. To maintain the amplitude of the supply waveform, switches to the dc supply and ground are provided as shown $(Q1Q4)$, and are activated at appropriate moments to "top-up" the system. The gates on these devices are driven by the resonant circuit itself.

Note that in this circuit, there is only an inductor present in the high-current supply path between the two capacitors. The switches only carry the relatively low top-up current, minimizing the potential resistive losses in the circuit.

The four phases required for an ADL system may be generated by two of these oscillators running with a quarter-period phase difference. Each oscillator may be built on-chip with the circuit, with only the small inductor (about 50 nH) connected between two package pins off-chip.

A difficulty with this resonant oscillator approach is that the load capacitance of an ADL circuit is data-dependent. As these loads change, the oscillator frequency must be actively stabilized to maintain a constant clock rate. One means of avoiding this problem is to adopt fully differential signaling in which, for every high-going node, there is a low-going node resulting in a constant capacitive load to the oscillator.

VII. Simulation and Implementation

Our early detailed simulations of ADL inverter chains indicated that operation was possible over a frequency range from less than 1 MHz to greater than 1 GHz. An average factor of 15 improvement in power dissipation over conventional CMOS was found in the 1–100 MHz frequency range. For example, at 100 MHz, an ADL inverter consumed 1.7 μW (assuming a 100% efficient resonant power supply) compared to 26 μW for a conventional inverter. These results are for $V_{dd} = 5$ V; however, the inverters operated down to $V_{dd} = 2$ V in simulation (with consequently less power).

We have been fabricating a series of ADL test circuits in 0.9 μm CMOS. Because suitable diodes are not readily available in this process, we have been using FET's connected as diodes (gate tied to one end of the channel), with the penalty that the effective forward drop across these devices is quite high—typically about 1 V.

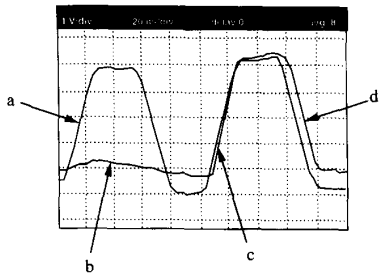

Fig. 11. Electron beam probed output of a 0.9 μm precharge low CMOS ADL inverter with a trapezoidal clock supply **a**. Note that the output of the gate precharges to within V_f of the clock low point **b**. During an evaluate, the output closely tracks the clock **c**, indicating that there is very little potential across the channel of the input device. During precharge **d**, the output again tracks the clock with a V_f drop.

In our first fabrication run, we built a chain of 64 inverters, which was successfully tested to 250 MHz the limit of our test equipment. Our second fabrication run included several chains of 1000 inverters, and established that ADL circuits could maintain good logic levels through many levels of logic. Clock skew tolerance was found to be in excess of 15% between adjacent phases.

We have used electron beam probing to obtain waveforms from the functioning circuits. Fig. 11 shows the operation of an ADL inverter at 10 MHz with alternating input. The output tracks the supply ramp very accurately during evaluation, and trails it by about 1 V during precharge, as expected. Because of the difficulty of accurately measuring the net power (the small difference between large input and output powers) drawn from a high-frequency supply, we have not yet been able to obtain accurate figures for power consumption. We are examining calorimetric approaches to this problem.

We have simulated and built a clock supply circuit (described earlier) from discrete components, and have found it to operate with 85–90% efficiency. Other more sophisticated and easily synchronized circuits have been simulated with similar results, and are presently being fabricated on-chip. We intend to use these to directly drive ADL circuits and accurately measure power consumption by monitoring the dc supply to the oscillators.

VIII. Conclusions

Our experience with ADL circuits indicates a factor of 15 reduction in dissipation in 0.9 μm CMOS. When combined with our 85–90% efficient resonant clock-supply circuits, the logic family offers an order of magnitude reduction in power consumption over conventional CMOS.

ADL circuit design will benefit from the addition of low-threshold diodes and transistors to the CMOS process. At the same time, the low field strengths present in devices used in an adiabatic regime may allow simplification of FET device design.

Acknowledgment

The authors wish to think B. Ackland, S. Avery, D. Inglis, L. Jackel, A. Kramer, and T. Wik for their many valuable contributions to this work.

References

[1] S. G. Younis and T. F. Knight, "Practical implementation of charge recovering asymptotically zero power CMOS," in *Research on Integrated Systems: Proc. 1993 Symp.* Cambridge, MA: M.I.T. Press, 1993.

[2] R. C. Merkle, "Reversible electronic logic using switches," *Nanotechnol.*, vol. 4, pp. 21–40, 1993.

[3] J. G. Koller and W. C. Athas, "Adiabatic switching, low energy computing, and the physics of storing and erasing information," in *Proc. Phys. Computation Workshop*, Dallas, TX, Oct. 1992.

[4] W. C. Athas, J. G. Koller, and L. J. Svensson, "An energy-efficient CMOS line driver using adiabatic switching," Inform. Sci. Inst., CA, Tech. Rep. ACMOS-TR-2, July 1993.

[5] R. T. Hinman and M. F. Schlect, "Recovered energy logic—A highly efficient alternative to today's logic circuits," in *Proc. IEEE Power Electron. Specialists Conf.*, 1993, pp. 17–26..

[6] C. L. Seitz *et al.*, "Hot-clock nMOS," in *Proc. 1985 Chapel Hill Conf. VLSI.* Rockville, MD: Computer Science, 1985.