Non-Dissipative Rail Drivers for Adiabatic Circuits

Saed G. Younis Thomas F. Knight, Jr

MIT Artificial Intelligence Laboratory
Cambridge, MA 02139

Energy dissipation of CMOS circuits is becoming a major concern in the design of digital systems. Earlier, we presented a new form of CMOS charge recovery logic (SCRL), with an energy dissipation per operation that falls linearly with operating frequency, as opposed to the constant energy required for conventional CMOS circuits [1].

These SCRL circuits, along with most adiabatic circuit techniques proposed to date, require a set of gradually swinging power supply rails that in effect force all charge transfers within the system to occur quasistatically. Proposals to date for generating these swinging rails have relied on a power MOSFET to gate the oscillation of an inductor, forming an RLC circuit. Even under ideal conditions, dissipation in this MOSFET degraded the overall energy savings of SCRL circuits from $1/T$ dependence to $1/\sqrt{T}$. SCRL and other adiabatic circuits thus exhibited inferior overall energy saving performance when compared supply voltage scaling of conventional CMOS circuits.

In this paper, we present a technique for generating the required rail waveforms *without* the series power MOSFET to gate the inductor. This new rail driver circuit relies on adding multiple harmonics of the base frequency to generate a rail waveform of *any* desired shape. Our Harmonic Rail Driver (HRD) can be built using only passive reactive components or by using correctly trimmed transmission line segments. It is non-dissipative to within the achievable Q's of these components. Using HRDs to power and control SCRL circuits, we restore the overall dissipation of SCRL circuits to its attractive $1/T$ dependence. In addition, adding harmonics allows us to construct a rail voltage that very closely approximates a linear voltage ramp instead of the sinusoidal waveform generated by previous techniques, improving the constant factor in the energy expression from $\pi^2/8$ to 1. This change markedly improves the performance of SCRL circuits, making their energy saving performance *comparable* to what is obtained by scaling the supply voltage in conventional CMOS when V_{TH} is held constant, for all operating frequencies.

Alternatively, with HRDs, we can push the energy-speed tradeoff in the direction of higher performance, since it becomes possible to over-drive a CMOS system *beyond* its maximum dc operating frequency for a given supply voltage, so long as the devices do not reach saturation, trading higher energy consumption for increased performance.

The usefulness of HRDs extends beyond driving adiabatic circuits, to other applications such as driving the periodic control lines of heavily capacitive memory chips non-dissipatively.

We also discuss a set of novel techniques aimed at improving the immediate, as well as the long term, viability of SCRL circuits through reducing the cost of reversibility.

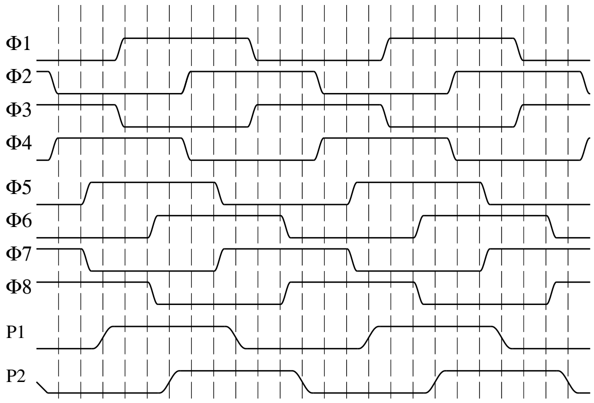

Figure 1. Rail timing for 4 phase SCRL.

1 Introduction

Since our introduction of the CRL and SCRL forms of charge recovery logic, we have focused our attention on reducing the energy dissipation in the computing circuits of digital systems. Theoretically, the energy dissipation of any computer that maintains a constant information entropy is linearly proportional to the speed at which the computation is forced to proceed. Having achieved this goal for digital circuitry with the SCRL logic family, the non-negligible dissipation in the power rail drivers for these circuits became the dominant source of dissipation in the overall system.

We have developed a simple and intuitive solution that succeeds in eliminating all of the non-linear dissipation in these rail drivers, thus making the overall energy dissipation of SCRL circuits track the theoretical minimum of $1/T$ where T corresponds to half a clock period.

We begin by revisiting the rail driver circuit that is commonly employed in adiabatic circuits today. We then present our new rail driver, the Harmonic Rail Driver (HRD). We show how to generate approximations to linear ramps using only discrete passive components. We then discuss techniques for constructing HRDs from one or more correctly trimmed transmission line segments instead of discrete reactive elements, such as capacitors and inductors.

Next, we examine the overall dissipation of SCRL circuits that are driven by HRDs and find that the overall circuit dissipation of a SCRL implementation is lower than that of conventional CMOS with supply voltage scaling for all operating frequencies.

Finally we revisit reversible computing and suggest ways in which non-reversible SCRL circuits can be operated with vanishingly small dissipation per bit of information discarded.

2 The Rail Driver Revisited

For proper operation, an SCRL circuit is controlled by gradually clocking a number of control/power lines we call rails. This clocking can proceed in several ways. One correct clocking pattern in shown in Figure 1. Adiabatic sequencing requires generating a waveform that gradually alternates between two voltage levels. Upon reaching one level, the waveform remains constant while other activity proceeds in the sequence. Typical methods for generating these these waveforms use the circuit shown in Figure 2.

A rail is electrically modelled as a capacitor in series with a resistor. The capacitor is the

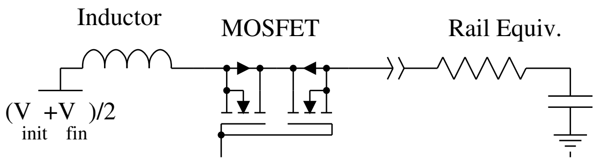

Figure 2. CRL rail driver circuit.

sum of the capacitances that the rail is driving and the resistor is the equivalent resistance of the devices through which the capacitances are driven. Properly designed SCRL circuits will have identical RC time constants for each gate branch. Suppose that the initial voltage on the rail was V_{init} and we wish to swing the rail to V_{fin}. To start the swing, we connect the rail through an inductor to a DC power supply at $(V_{init} + V_{fin})/2$. Current starts to build up in the inductor and the rail starts the swing towards V_{fin}. At the moment that the current drops back again to zero we disconnect the inductor. The rail should now be at V_{fin}. The action of connecting and disconnecting the rail is performed by the series power MOSFET and provides the required "dead zone" between transitions.

There are two difficulties in fabricating circuits with high energy recovery factors using the above circuit. The first results from the swings having a sinusoidal form. Ideally, charge recovery circuits require a linear voltage ramp during swings [3]. Using a sinusoid reduces the energy saving performance of these circuits by a factor of $\pi^2/8$ at all operating frequencies.

The second and most serious difficulty is that of the dissipation associated with controlling the action of the power MOSFET. Athas, Koller and Svensson [2] have studied this problem and have shown that dissipation in the power switch forces the energy dissipation of the overall system to follow $1/\sqrt{T}$ instead of the theoretical minimum that drops linearly with $1/T$.

The problem comes from the tradeoff between lowering the energy dissipated due to the ON resistance of this MOSFET and between the energy needed to control it. Making the MOSFET wider decreases its dissipation due to on resistance. However it also increases the gate capacitance of this power MOSFET thus increasing the energy needed to control it. Assuming that this power MOSFET is controlled non-adiabatically, the optimal point in this tradeoff forces the overall dissipation of the switch to follow $1/\sqrt{T}$.

In [2] the authors analyze driving the switch with yet another sinusoidal ramp circuit which itself is driven sinusoidally and so on. From their analysis we see that in theory, we could recover $1/T$ behavior provided we carry out this nesting on these rail driver at infinitum. However, they report that practical considerations limit the nesting to 2 or 3 levels only.

In [3] we revisited their analysis and suggested a number of techniques that can partially recover the linear relationship between energy dissipation and operating frequency. In what follows, we describe how to recover the linear $1/T$ dependence fully.

3 Harmonic Rail Driver

As described above, present rail drivers undesirably require a dissipative switch and produce a sinusoidally shaped waveform instead of the desired linear ramp. The harmonic rail driver (HRD) succeeds in eliminating both of these drawbacks in a simple and intuitive manner. Any periodic waveform can be constructed by the linear addition of appropriately scaled and phased cosine waves using a Fourier series expansion. In particular, for a square

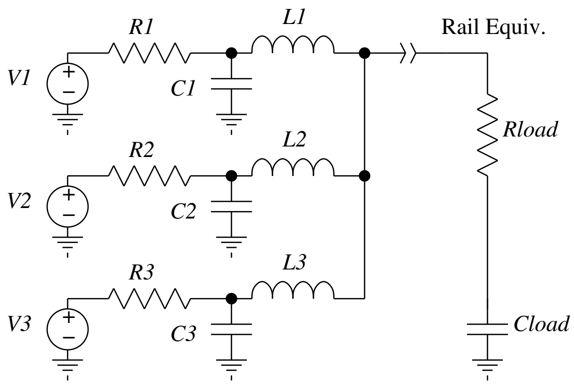

Figure 3. HRD circuit containing 3 terms.

wave, $x(t)$, the expansion follows

$$x(t) = \frac{1}{2} + \frac{2}{\pi} cos \frac{2\pi t}{T} - \frac{2}{3\pi} cos \frac{6\pi t}{T} + \frac{2}{5\pi} cos \frac{10\pi t}{T} - \ldots \quad (1)$$

where T is the period of the square wave. Fortunately, truncating the series to a small number of terms produces an imperfect square wave with *sloping sides*. This is an approximation to the ideal linear ramp that we desire. Reducing the magnitude of the added harmonics both reduces the slopes of the square wave sides and reduces ripple, both highly desirable for our application.

A harmonic rail driver containing only the first three harmonic terms is shown in Figure 3. Throughout our discussion we will assume that R_{Load} is small enough that the circuit is under damped. Initially we will also assume that $C_{Load} \gg C_1 \simeq C_2 \simeq C_3$. This condition is not necessary but makes understanding of the circuit operation simpler. We will relax it later on.

At $t = 0$, all the currents in the network are at zero, and the voltage across the load, v_{Load}, is also at zero. Additionally, we assume that $v_{C1}(t_0) = V_1$, $v_{C2}(t_0) = V_2$ and $v_{C3}(t_0) = V_3$, where V_1, V_2, V_3 are equal to the alternating coefficients in our series. Starting from this condition, currents start to build up in each inductor. The sum of these developing currents are fed into C_{Load}. With C_{Load} assumed large, the frequency of oscillation is each inductor is governed by its respective capacitor and is independent of C_{Load}. Hence the circuit reduces to three independent LC oscillators with $\omega_n = 1/\sqrt{L_n C_n}$ for $n = 1, 2, 3$. Letting $\omega_1 = 2\pi/T$, $\omega_2 = 3\omega_1$ and $\omega_5 = 5\omega_1$, the LC branches will independently oscillate at $f_0, 3f_0, 5f_0$, which are the required harmonics in our expansion. Letting C_{Load} be much larger than the other capacitances simplifies picking up the correct component values but has the side effect of amplifying the voltage swings in C_1, C_2 and C_3. This is not necessarily bad since these modes do not go through any dissipative element. In addition, having a large C_{Load} makes f_0 nearly independent of variations in C_{Load} that are usually present in non-complementary SCRL circuits. It is important to note here that variations in C_{load} must still be limited to small percentages since wild fluctuations in C_{load} still lead to proportional fluctuations in V_{load}. If it is still desired to increase C_1, C_2 and C_3 to the point where they are comparable to C_{Load} we can easily do this in small numerical steps. Given C_{Load}, we start with $C_{Load} \gg C_1 \simeq C_2 \simeq C_3$ and pick the inductors to get $f_0, 3f_0$ and $5f_0$. We then incrementally increase C_1, C_2 and C_3 while adjusting L_1, L_2 and L_3 to maintain $f_0, 3f_0$ and $5f_0$ as the natural frequencies of the system. We converged on the solution incrementally since finding the general solution to the poles of the system equation with seven reactances was somewhat intractable.

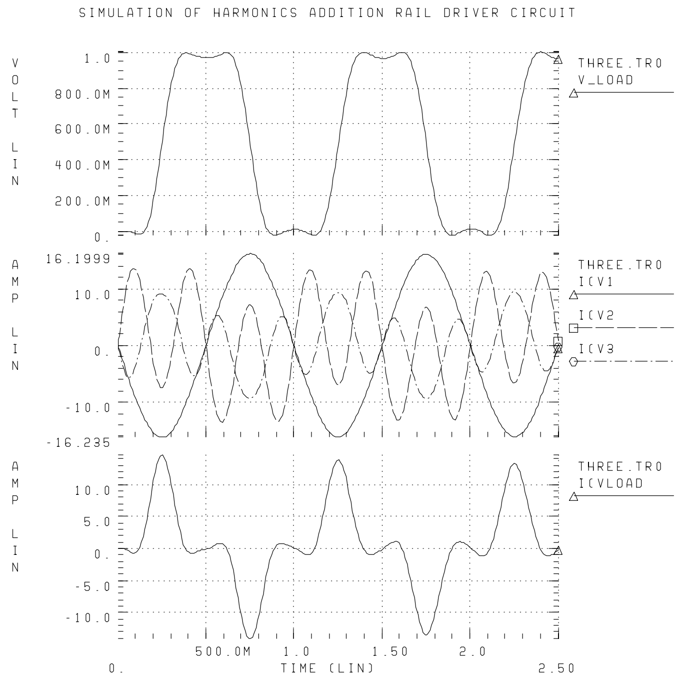

Figure 4. Spice simulation of a HRD circuit containing 3 terms.

As an example, Figure 4 shows the waveforms from a SPICE simulation of the above circuit with $C_{Load} = 2$ F, and $C_1 = C_2 = C_3 = 1$ F, and $L_1 = 0.0314$ H, $L_2 = 0.00346$ and $L_3 = 0.00163$ H. The initial conditions were $V_1 = 3.02$ V, $V_2 = -0.807$ and $V_3 = 0.231$ V. Obviously, the above values can be scaled to yield any desired frequency and amplitude. The first trace in the Figure shows the waveform appearing across the load, C_{Load} in series with R_{Load}, and demonstrates the ability of the HRD circuit to generate linear voltage ramps with flat "dead zones". The second trace shows the three current components that sum up to the load current which is shown in the third trace.

Finally, a single HRD circuit can be non-dissipatively multiplexed among a number of rail in order to reduce the number of required inductors in the the system. The multiplexor comutator is placed in series with the load and is used to share one HRD circuit with multiple rails. We note here that the switch of the multiplexor is a non-dissipative adiabatically controlled switch unlike the power MOSFET in the old rail drivers. This is true because of the existence of a finite and long "dead zone" in the switched waveform during which the MUX switch can be toggled *gradually* and without dissipation. Integrating everything but the inductors and capacitors inside the silicon chip results in a chip with SCRL behavior with very few external components.

The HRD circuit discussed so far imposes no restrictions on the relative phases of the separate oscillators. In time these will drift, thus altering the shape of the generated wave. In addition, the circuit has no way of replenishing the energy that is being lost every cycle due to R_{Load} and the finite Q's of the inductors. We compensate for this lost energy by connecting a voltage source through a source resistance R_{source} directly to the rail. The waveform generated by this source is the waveform that we desire across the load, C_{Load}. In this connection, the voltage source continuously injects the correct amount of energy needed to maintain the desired waveform across the load. As always, the source resistance of the power supply should be made as small as practical in order to minimize the energy

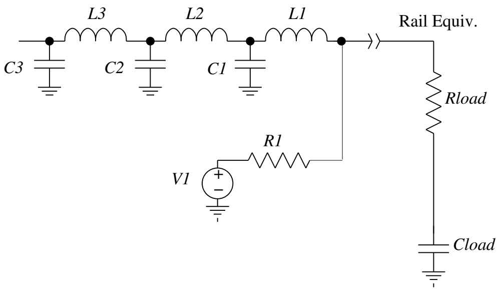

Figure 5. Another example of an HRD implementation.

dissipated in it as a percentage of the energy reaching the system through it.

Finally, we stress that an HRD using a finite number of branches can losslessly sustain waveforms that only containing the harmonics that are supported by the branches. Any component of the waveform that is at a frequency not supported by a resonant branch in the HRD is dissipated in R_1, R_2, and R_3.

Please note the these HRDs are just as capable in generating non sinusoidal waveforms non-dissipatively across inductive loads as well.

4 Generalization of HRDs

The circuit in Figure 3 serves only as one example out of the countless number of ways in which HRDs can be built. In essence, an HRD is a linear reactive circuit that when combined with the slightly dissipative reactive load, will yield a system that has more than one underdamped natural frequency of oscillation. By design, these frequencies will be the frequencies that are present, with non-negligible power, in the spectrum of the non-sinusoidal waveform that we hope to develop across the load. Figure 5 shows a circuit that conforms to the above definition and would therefore be able to function as an HRD given a proper selection of component values.

5 Transmission Line HRDs

Figure 6 shows a way to build HRDs without using discrete reactive circuit elements, such as inductors and capacitors. Like the previous reactive implementation, the purpose of this circuit is to non-dissipatively generate a periodic linear ramping signal across the capacitive load C_{Load}. Just as the case with the reactive implementation, this circuit achieves its goal by adding a number of odd harmonics, $3f_0$, $5f_0 \ldots$ to the base operating frequency of the signal f_0. As illustrated in the figure, the circuit is composed of an unterminated transmission line of impedance Z_0, that is trimmed to a length of $\lambda/2$, where $\lambda = c/f_0$ and c is the speed of propagation of electromagnetic waves in those transmission lines. Alternatively, the line could be trimmed to $\lambda/4$ and short-circuited at the far end.

To explain its operation, we start by assuming that the impedance of the load capacitor, $1/j\omega C_{load}$, is much larger than Z_0. With this assumption, the transmission line appears to be unloaded. Because an unloaded transmission line of length $\lambda/2$ can resonantly support all the odd harmonics, we find that in a short time the periodic waveform of V_1 appears across C_{Load} with no modification provided that the waveform has only odd harmonic components. More importantly, once in resonance, there is no voltage across the source resistor. Hence, the system is able to non-dissipatively generate any periodic waveform,

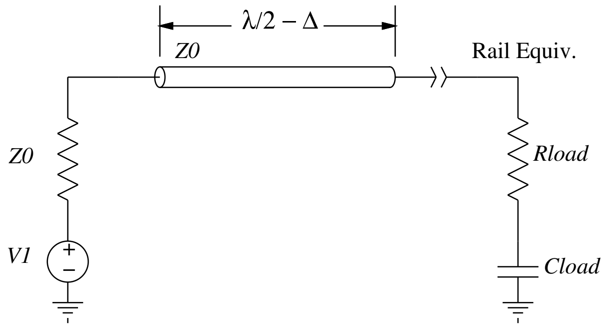

Figure 6. HRD circuit constructed out of trimmed transmission lines.

including the desired linear ramp, across C_{Load}. Since we know that a single transmission line is capable of supporting all odd harmonics, we need only one transmission line for the circuit to work properly given that C_{Load} is kept small.

A transmission line with a 50Ω impedance and operating at a base frequency of $f_0 = 25$ MHz requires that

$$C_{Load} \ll \frac{1}{2\pi f_0 Z_0} \quad (2)$$

$$\ll \; = 127\text{pF} \quad (3)$$

in order to appear unloaded. Therefore, at that frequency, a *small* C_{Load}, must not exceed $\simeq$ 20pF. Any higher value and the behavior of the circuit becomes very different.

A problem arises from the non-negligible capacitive load at the end of a transmission line, which has the effect of increasing the perceived electrical length of that transmission line and altering its supported resonant frequencies away from f_0, 3 f_0, $5f_0 \ldots$ The problem can be slightly improved by shortening the physical length of this transmission line to obtain the desired odd harmonic resonances. Unfortunately, the impedance of a capacitor is frequency dependent and hence the correct length of the transmission line for resonance at the desired harmonics is different for every harmonic.

Thus, it is desirable to have transmission lines with as low Z_0 as possible. Fortunately, we can achieve this in a variety of ways. We know that ideally, the impedance of a transmission line, Z_0, is equal to $\sqrt{L/C}$. Therefore, increasing the unit capacitance of a transmission line will simultaneously lower its impedance. We can increase the capacitance of a line by periodically shunting the line with discrete capacitors that are spaced closely enough to approximate an increased distributed line capacitance. At minimum, the number of added shunt capacitors should equal the number of needed harmonic frequencies. This is because each added capacitor effectively lowers the impedance of the line for a particular frequency depending on its position. Alternatively, the distributed capacitance of the line can be increased by altering the geometry of the line, by bringing the conductors closer together, or by using an insulator with a high dielectric constant if the resources to do so are available. Another way to obtain a low impedance transmission line is to parallel a number of higher impedance transmission lines.

Finally, we can use a broadband transmission line transformer to "transform" the impedance of our resonating line to a lower value. A good source for empirical data and design of a variety of transmission line transformers is presented by Sevick in [6].

Just as in previous implementations of HRDs, this implementation can non-dissipative support any waveform that is made entirely out of the harmonics that are supported by the transmission line resonantor. Any other frequencies will be dissipated in the source resistor and will arrive altered at the load.

Note that transmission line HRDs, like discrete HRD implementations, can be used to drive inductive loads as well. In his case the lines should be lengthened to accommodate a non-negligible inductive load rather than shortened as in the capacitive load case.

6 SCRL Operation vs. Supply-Voltage Scaling

In [4] Indermaur and Horowitz sought to evaluate charge recovery circuits for low power CMOS applications. In their analysis they used SCRL as a representative of adiabatic switching techniques. Following the analysis in [5] we see that energy dissipation for SCRL, E_{SCRL} follows

$$E_{SCRL} = \frac{1}{2}C(\triangle V)^2(1 - e^{(\frac{-\pi\alpha}{\omega_d})}) \tag{4}$$

Using Equation 4 as the basis for comparison, Indermaur and Horowitz report that charge recovery circuits at 5.0 Volts consume more energy than conventional CMOS at 5.0 Volts with both operating at the maximum frequency. In comparing the lowering clock speed of SCRL circuits with the linear scaling the supply voltage of conventional CMOS they report that conventional CMOS continues to outperform SCRL until the scaled supply voltage of conventional CMOS reaches 1.0V. In addition, they assert that once the dissipation of the power MOSFET in the old rail drivers is factored in, the $1/\sqrt{T}$ energy dissipation dependence of the overall SCRL system becomes much worst that conventional CMOS for all interesting regions of operation.

While their conclusions are accurate when using single inductor rail drivers, the situation is completely reversed with the introduction of the harmonic rails driver. HRD circuits are non-dissipative and therefore the overall system continues to have an energy performance that tracks $1/T$. More importantly, equation 4 was developed for SCRL circuits that are driven by sinusoidal profiles which degrade the performance by a constant factor of $\pi^2/8$. Using HRD circuits that approximate ideal voltage ramps, E_{SCRL} now becomes

$$E_{SCRL} = \frac{1}{2}C(\triangle V)^2(\frac{RC}{T}) \tag{5}$$

With the elimination of the $\pi^2/8$ factor, E_{SCRL} now equals CMOS dissipation at the highest operating frequency and supply voltage. Furthermore, using HRDs, the energy dissipation scales linearly with frequency starting from this highest frequency of operation, in contrast to the older sinusoidal drive. In contrast, the energy dissipation of conventional circuits does not scale linearly with supply voltage if the threshold voltage is kept constant, because as V_{DD} approaches the threshold voltage, the reduction in energy dissipation slows due to the poor conductance of weakly turned on CMOS devices.

In summary, the introduction of HRD generators succeeds in tipping the relative energy advantage of SCRL operation versus scaling the supply voltage of same process CMOS in favor of SCRL operation. This, coupled with the fact that in many situations it is more desirable to slow down the clock rather than to alter the process, supply voltage, and interfacing standards, should make SCRL a competitive alternative for low power design even today.

In its pure form, SCRL requires a reverse gate for each forward gate. This in effect doubles the gate count and hence doubles the dissipation of SCRL when compared to conventional CMOS. Later, we sketch how this factor of two can be mitigated.

7 Overdriving with HRDs

With the ability to generate any periodic waveform with relative accuracy, HRD circuits can be used to over drive conventional CMOS circuits and thus to increase energy dissipation in favor of increased clock speed. With previous rail drivers we could not do that because the capacitive load of the circuit set the frequency at which the system becomes critically damped. This is no longer the case when HRDs are used. With HRDs, it becomes possible to position the critically damped point of the system higher than the maximum operating frequency of conventional CMOS circuits.

If the rail waveform have significant overshoot, then for most of the rise and fall times, the nodes are charged at an accelerated rate since the rail voltage exceeds V_{DD} during the transition. Timing the overshoot such that the rail voltage goes back to V_{DD} just as the driven nodes reach V_{DD}, we can limit the swings of the internal nodes to a maximum of V_{DD}. This in effect allows us to operate the circuit at a speed attainable only with a higher V_{DD} but without letting internal signals swing more than V_{DD}.

In essence, just as SCRL and HRDs allow CMOS circuits to consume less energy at lower frequencies, they allow for a higher operating frequency than conventional CMOS at the price of higher energy dissipation. Admittedly, creating an overshoot with just the right timing and the right shape is not a trivial matter. In addition, we must be careful to keep the devices from operating with velocity saturation causing no improvement in performance while consuming more energy. Just as the normal use of SCRL reduces or eliminates the reliability impact of hot electron behavior and electromigration, the operation of of SCRL overdriven in this way accentuates these problems.

8 Cost of Reversibility

Pure SCRL requires that every functional block in the forward direction be complemented by a functional block in the reverse direction performing the inverse of the forward operation. This has the disadvantage of doubling the circuit loads and resistances in certain current paths of the system. In addition, being accustomed to throwing information away, pure reversibility continues to be the major barrier against adopting SCRL techniques today. While it is the opinion of the authors that this will cease to be an issue for the long term, through a rethinking of computer architecture and programming environments as we know them to include reversibility requirements from the outset, we offer here a number of hints on how to use SCRL circuits non-reversibly in order to make them more attractive in the near future.

In our original proposal for CRL circuits, we presented a technique by which we limited the energy wasted by an irreversible bit of information to CV_T^2. The same trick can be employed for SCRL with minor modification of the SCRL gate. If in this portion of the circuit we use devices with very low V_T, in effect we are able to get away with irreversibility while dissipating a vanishingly small amount of energy. The beauty of this approach is that using low threshold devices in this part of the circuit does not increase overall dissipation due to increased sub threshold conduction of these low V_T devices because, whenthese devices are turned off, the voltage difference across them is asymptotically 0.

9 Other Applications of HRDs

HRDs usefulness is not limited to controlling the swinging rails of adiabatic circuits. They are just as useful in non-dissipatively driving any capacitive or inductive loads in a circuit. Please note that HRDs are capable of non-dissipatively producing *any* shape waveform as long as it is periodic and as long as enough resonant modes are included in the HRD to accurately construct the waveform.

One example application is to use HRDs to non-dissipatively drive the control lines (RAS,CAS) in a memory array.

10 Conclusion

In this paper we have demonstrated a simple technique for generating linear ramp voltage swings non-dissipatively. The technique relies on the addition of a number of harmonic waveforms to construct periodic waveforms that are ideally suited to driving SCRL rails. We have shown how the introduction of these harmonic rail drivers restores the equivalence between SCRL operation and supply voltage scaling if V_{Th} is also scaled and to tip the balance in favor of SCRL at low frequencies when V_{Th} is not proportionally scaled. Further, we have shown how HRDs can be built using discrete reactive components or by using trimmed transmission lines. Being non-dissipative, HRDs succeed in restoring the overall energy performance of SCRL to its theoretical minimum of $1/T$. In addition, the ability of HRDs to generate near linear voltage ramps eliminated the $\pi^2/8$ degradation in energy saving that was caused by the sinusoidal waveform shape of previous rail drivers. With this degradation eliminated, SCRL dissipation becomes equal to conventional CMOS dissipation when both are operated at their maximum frequency. Furthermore, at all frequencies lower than the maximum, SCRL circuits driven by HRDs exhibit better energy dissipation characteristics than that of conventional CMOS when only the supply voltage is scaled, making SCRL a more viable option to low power design today. We have also discussed how HRDs can be used to achieve operation at higher frequencies than are achievable with conventional CMOS, albeit at higher energy cost. While HRDs are valuable in driving adiabatic circuits in particular, HRDs also have valuable applications in any application that requires driving a periodic waveform of general shape across a reactive load non-dissipatively.

Finally we sketched how the cost of irreversibility can be lowered by using in low threshold devices in bit-discarding circuits, with little penalty for transistor leakage.

Acknowledgment

This research is supported in part by the Advanced Research Projects Agency under contract N00014-91-J-1698. We acknowledge useful discussions with William Athas, Robert Fano, William Siebert and Tom Simon. We thank L. Svensson for his helpful comments.

References

[1] Younis, S., and Knight, T., "Asymptotically Zero Energy Split-Level Charge Recovery Logic," *Proc. Workshop on Low Power Design*, April 1994.

[2] Athas, W. C., Koller, J. G., and Svensson, L. , "An Energy-Efficient CMOS Line Driver Using Adiabatic Switching," *Proc. 1994 IEEE Great Lakes Symposium on VLSI,* March, 1994.

[3] Younis, S., "Asymptotically Zero Energy Computing with Split-Level Charge Recovery Logic," *MIT AI Lab Tech. Report #AITR-1500,* June 1994.

[4] Indermaur, T., Horowitz, M., "Evaluation of Charge Recovery Circuits and Adiabatic Switching for Low Power CMOS Design," *Proc. 1994 Symposium on Low Power Electronics,* San Diego, October 1994.

[5] Younis, S., Knight, T., "Practical Implementation of Charge Recovering Asymptotically Zero Power CMOS," *Proceedings of the 1993 Symposium on Integrated Systems,* MIT Press, 1993, pp. 234-250.

[6] Sevick, J., "Transmission Line Transformers," *American Radio Relay League,* 2nd edition, 1991.