A Clocked, Static Circuit Technique for Building Efficient High Frequency Pipelines

Eric Gayles Kevin Acken Robert M. Owens Mary Jane Irwin

Architecture and VLSI CAD Group, Department of Computer Science and Engineering
The Pennsylvania State University, University Park, PA 16802

Abstract

This paper presents a CMOS circuit methodology for designing pipeline stages which are both faster than comparable domino based stages and that also have increased functional capability. The basic gates offer considerably faster switching speeds than domino, while also eliminating the feedback and buffering circuitry required by domino gates for reliable operation. In addition to faster gates, the dual-rail nature of the proposed circuit technique provides greater logic functionality per gate. This results in a reduction of the number of gate delays required for implementing complex functions of high fan-in. Several benchmark circuits were simulated in a 0.5 μm, 3.3 V CMOS process. The results show that the proposed circuit technique provides significant speed improvement over domino.

1. Introduction

Due to the increasing frequency of microprocessors, two trends have evolved. The first is that the use of domino circuits has become more prevalent. Secondly, the functional capability of a single pipeline stage is being dramatically reduced due to higher clock rates. The problem lies with this latter trend. Not only are microprocessors running at faster speeds, but aggressive architectural features such as branch prediction, and out-of-order execution are becoming commonplace. Clearly these features require complex logic which, due to the ever shrinking cycle time, is being spread across several pipeline stages. As a result, the overall performance loss due to branch misprediction penalties and interrupts becomes more significant. Furthermore, though domino gates offer faster switching speeds than conventional CMOS gates, the required buffering after each gate along with the feedback circuitry to handle charge sharing and leakage problems can undermine the performance benefits of domino.

In this paper, we propose a CMOS circuit technique which provides faster gates than both conventional CMOS and domino gates. It will be shown that with the proposed technique, gate delays for high fan-in gates, especially those containing long series connected MOSFET chains (SCMCs), are much less than those of comparable domino gates. In addition, greater logic functionality is available per gate, due to the use of dual-rail logic. Therefore, functions like three-way XORs can be realized as a single gate. Dual-rail logic also eliminates any inverters which would otherwise be required for function realization and not for current drive, reducing the number of gates along critical paths. Section II will describe the basic domino circuit circuit design. In section III, we describe our circuit technique, discussing both the gate topologies and the clocking issues involved. In section IV, we describe a high speed 64 bit adder built using the proposed techniques, and section V concludes the paper.

2. Domino Logic

Domino logic offers speed advantages over conventional static gates because of reduced input and self-loading capacitance. It also allows fast implementations of n-based NAND and p-based NOR gates [1, 2, 3]. Figure 1 illustrates the basic domino circuit structure. Because of the problems domino gates have with charge sharing and leakage current, many domino designs use a V_{dd} keeper transistor as shown in figure 2. The problem with this technique is that the gate's switching speed is reduced. The V_{dd} keeper transistor is always conducting at the end of the precharge phase. The pull down network must then fight against the active pull-up of the keeper transistor in order to drive the gate low.

Consider the circuit shown in figure 1. When realized in domino logic using minimum sized n-transistors, driving a capacitive load of 0.058 pF, the circuit's worst case delay was measured at 2.01 ns. When a minimum sized V_{dd} keeper was added to each gate, the high fan-in NAND

1066-1395/97 $10.00 © 1997 IEEE

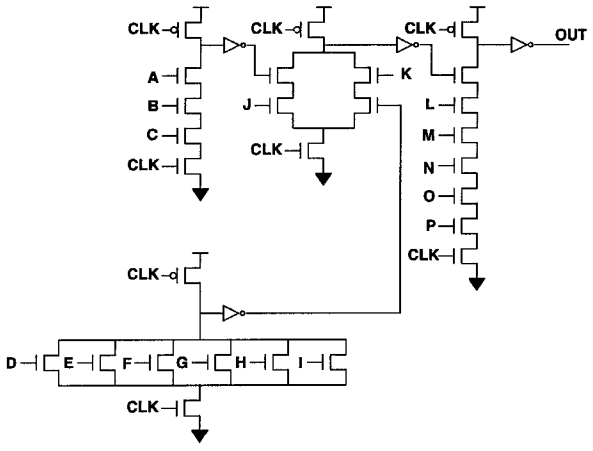

Figure 1. A domino circuit

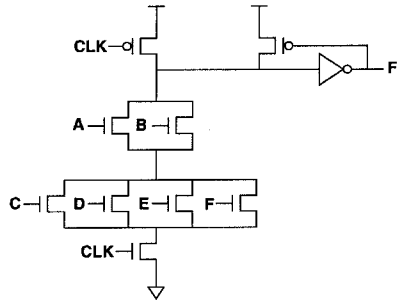

Figure 2. A domino gate with V_{dd} keeper

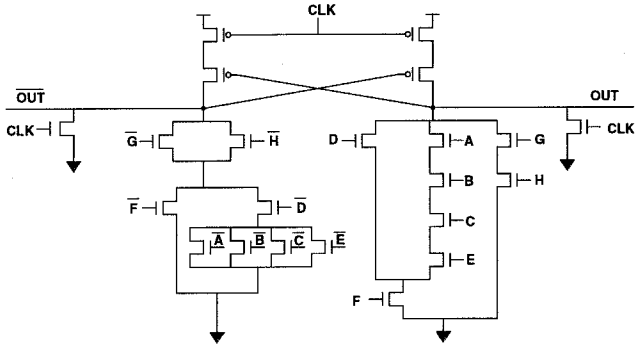

Figure 3. An SPSD gate

gates were no longer capable of generating a zero. By doubling the n-transistor size, a delay of 3.33 ns was attained. While at some point, sizing the n-transistors would allow adequate pull-down drive within the high fan-in gates, the increased input capacitance would increase the delay of the previous stage. The fact that the conventional CMOS implementation of the circuit had a delay of only 2.72 ns shows that domino logic does not always offer dramatic performance improvements. For simple NAND or NOR gates, the pull-up current of a weak feedback transistor is easily overcome. However, for complex boolean functions which contain lengthly SCMCs, domino logic may not provide significant speed improvements.

The potential for charge sharing also limits the complexity of domino gates. Large complex gates potentially contain large parasitic capacitances within their pull-down networks. These capacitances not only affect gate speed, but reliability as well.

3. Sympathetic Precharged Static Domino

3.1. Gate Design and Operation

As an alternative to domino logic for high frequency design, we propose a circuit technique which is more effective at realizing complex boolean functions of high fan-in. The proposed circuit style results in gates with static values at their output nodes, which is a significant advantage over dynamic domino. Even more, the delay of gates implemented using this technique varies only slightly as a function of gate complexity. This is in stark contrast to both conventional and domino logic.

Figure 3 shows the structure of a Sympathetic Precharged Static Domino (SPSD) gate. Just as with domino logic, either all of the inputs should be precharged low, or the inputs should be stable at the beginning of the evaluate phase of the cycle. The gates outputs are precharged low, in sympathy with the evaluation direction. The gate is precharged during the high phase of the clock. When the clock transitions low, the gate is *released* to stabilize based on its inputs. The speed advantage of SPSD over conventional n-based domino derives from the fact that the output nodes are not discharged through an n-network. Therefore, the complexities of the n-networks have only a minor impact on the gate's delay. In SPSD, a closed n-network only serves to maintain a low value in order to create an imbalance in the rates of charge accumulation between the two output nodes. During evaluation, only one of the n-type networks starts conducting. Since the evaluate signal is low and both outputs are charged to 0 V, both sides of the gate will try to pull their output nodes high. However, the side with the closed n-network immediately shorts V_{dd} and Ground. This prevents that side's output node from rising but so far, while

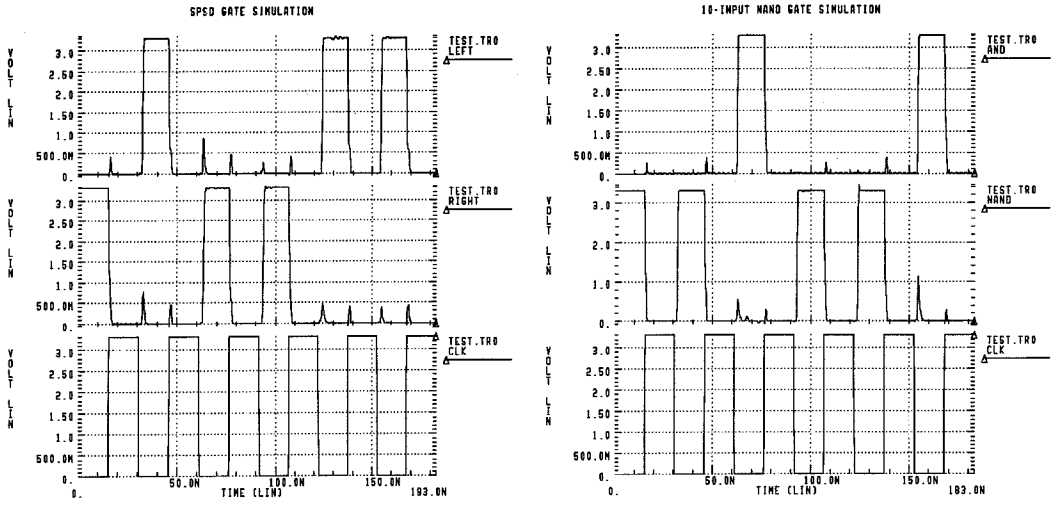

Figure 4. SPSD Gate Output profiles

the p-transistor it drives continues to conduct current. As a result of this imbalance, the output on the other side of the gate quickly rises to V_{dd}. As it rises, the short circuit within the gate is broken and the gate stabilizes with 0 V on one output and V_{dd} on the other.

3.2. SPSD Gate Speed

To estimate the speed advantage of SPSD logic over conventional static and dynamic domino, we compared the delay of NAND gates having 2 to 10 inputs in the various logic families. A NAND gate was used as the basis for comparison because its delay increases with fan-in. A comparison of NAND gates delays tells us how an arbitrary complex gate's delay is effected by the maximum series transistor path within the circuit. What we found was that for series transistor networks greater than 3, an SPSD gate design is faster.

To compare the logic families, each gate was driven from a set of edge-triggered registers and fed an edge-triggered register. All of the n-transistors were minimum size while the p-transistors where twice as wide. Hspice [4] was used to simulate the worst case condition for the gates. The SPSD NAND side of the tested gates face their worst case condition when all but the bottom transistor is on. This is because the capacitances in the pull-down network have an effect on the pull-up delay. However, the linear resistances of the transistors in the pull-down network isolate these capacitances to the extent that they have a marginal effect on the delay. The complete results are graphed in figure 5.

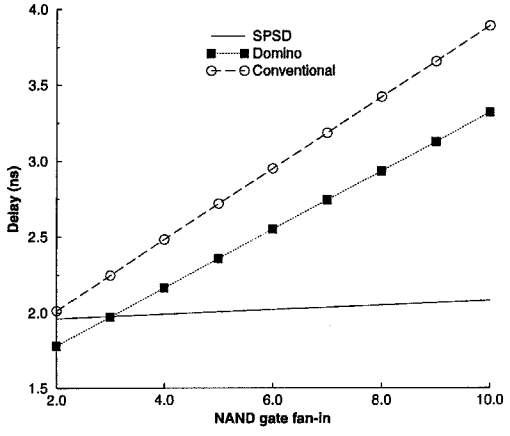

Figure 5. Comparative delay of conventional static, domino, and SPSD

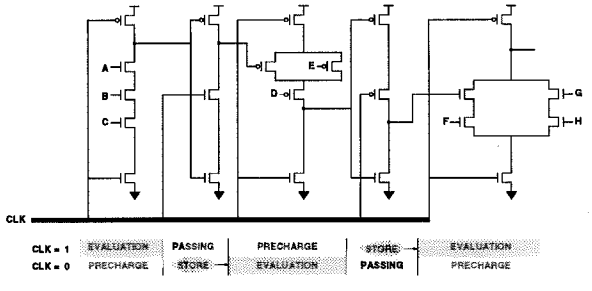

Figure 6. A True-Single-Phase domino pipeline

From the graph, it is clear that the delay for the SPSD gates changes marginally as a function of fan-in. This result indicates that extremely fast gates of high fan-in can be realized.

For small gates SPSD offers little advantage over domino. However, for gates requiring lengthly series transistor networks, SPSD offers very high performance. Figure 4(a) shows the output profile for several cycles of a 16 input SPSD gate implementing $f = abcd + efgh + ijkl + mnop$. Figure 4(b) shows the output curves for a 10-input NAND gate. Its characteristic is nearly identical with that of the previous example, except that the NAND side of the gate tends to spike noticeably higher than the AND side. This is caused by the imbalance in the two output capacitances. The AND side has a self loading output capacitance of 13 drains, whereas the NAND side's self loading capacitance is 4 drains. Yet, even with this imbalance, the gate's performance is not significantly effected. If there is a significant inequality in the output capacitances, the widths of the four p-type transistors can be sized to compensate for the imbalance. This is why two separate clocked p-transistors are recommended in the pull-up structure.

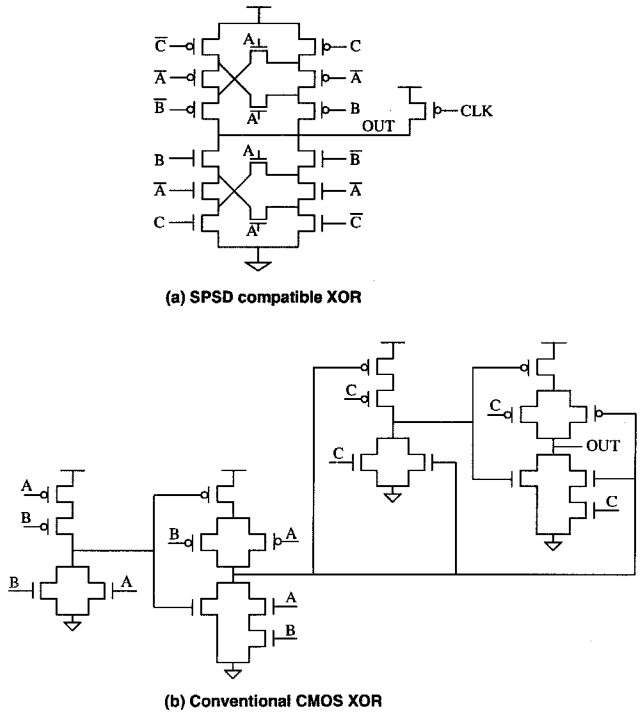

Figure 7. Comparison of two-input XOR gates

3.3. Functional Efficiency

For high frequency pipelines, fast complex gates provide a means for improved efficiency. For example, consider the True-Single-Phase-domino [5] circuit shown in figure 6. The single SPSD gate shown in figure 3 allows the same operating frequency, without three separate pipeline stages. In addition, SPSD represents a complete logic family, unlike domino or conventional CMOS, which allows for more efficient functional mapping onto gates. This tends to reduce the number of gate along critical paths, since for example, gates inserted to generate an OR from a NOR gate or to complete an XOR computation disappear. In particular for XOR functions, figure 7 shows how the precharged complementary gate can be used efficiently. XORs are not easily realized in conventional CMOS or in domino. Therefore being able to provide high speed XOR functionality *at any point* in the pipeline stage is one significant advantage of SPSD logic.

3.4. Clocking

Aside form providing fast gates, a circuit technique needs to provide an efficient methodology for connecting gates together. One reliable method for constructing SPSD based pipeline stages is shown in figure 8 The first gate of the stage is a SPSD gate. It feeds a precharged complementary CMOS gate. The precharged complementary CMOS gate serves to guard the rest of the pipeline from the false spike on one of the SPSD gate's outputs. The p-transistors in the complementary gate are sized to provide sufficient current from V_{dd} to the output to eliminate the effect of small voltage spikes on the inputs. The complementary gate requires a locally inverted clock. As with the SPSD gate, its inputs must be precharged low, however, they must settle prior to the evaluation of the SPSD gate. If this cannot be guaranteed, then inverters are used instead to squash the spikes on the SPSD gate's outputs. The third gate in the pipeline stage is a p-transistor based DCVS [6] gate. Its inputs must either be stable at the time evaluation begins, or must be precharged high. This ensures that the cross-coupled transistors cannot latch into an incorrect state. The forth gate is a clock delayed SPSD gate. After the fourth gate, we can continue with a precharged complementary CMOS gate, and so on.

The advantages of this circuit methodology become apparent when you consider that even for extremely complex logic, each gate will have a delay equivalent to a simple complementary CMOS gate. Moreover, each gate does useful work in the computation, unlike the required inverters following each gate in conventional domino. What may be even more important from the standpoint of reliability, is that all of the signals within the pipeline stage are static. This is also an advantage in low power designs where significant savings can be achieved by disabling nonactive units.

Figure 10 shows logic for a complex pipeline. The n-transistors were all minimum size, and the p-transistors were 6λ. The output transitions for each gate are shown in fig-

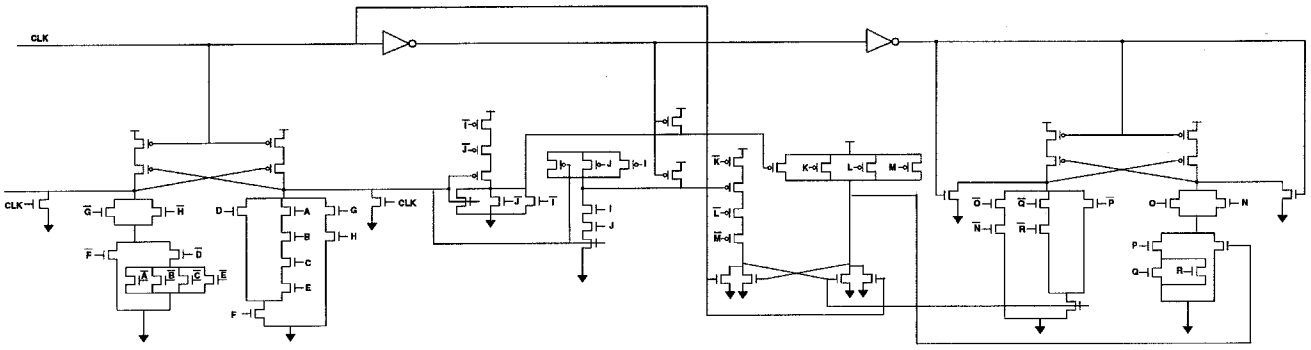

Figure 8. A 4 gate SPSD based pipeline stage

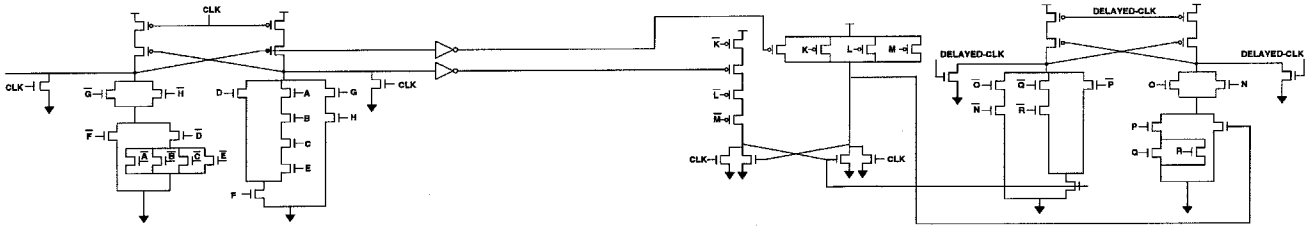

Figure 9. A 3 gate SPSD based pipeline stage

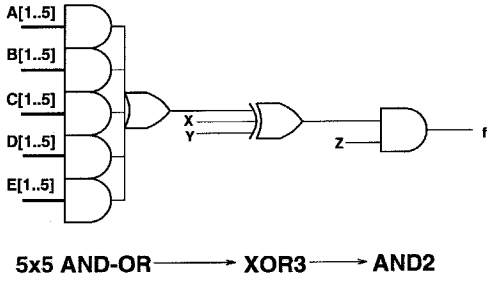

Figure 10. Pipeline logic

Adder	Precision	Circuit Style	Delay	
RC	32	Conventional	25.32	
RC	32	Trans. Gate	19.20	
CS	32	Conventional	11.55	
CLA	32	Conventional	8.58	
CLA	32	Domino	6.82	(eval)
CLA	64	SPSD	4.6	(eval)

Table 1. Adder comparisons for ripple-carry(RC), carry-select(CS), and carry-lookahead(CLA).

ure 11. Due to the gate complexity, it would be impossible to realize the same logic in three gate delays using regular domino logic. SPSD gates provide a means of implementing complex logic within a evaluation cycle.

From figure 8, it is apparent that if the logic for the precharged complementary static gate cannot be chosen efficiently, it contributes the least to speed improvements. For these cases we recommend the alternative strategy is shown in figure 9 where a buffering stage replaces the precharged complementary gate with a buffering stage. In this arrangement, the buffers directly drive the p-based DCVS gate.

4. A High Speed 64-bit Adder Using Clock Delayed SPSD

An alternative clocking scheme which employs SPSD gates is shown in figure 12. The clock is slightly delayed at each stage to allow the SPSD gate to settle to a steady state before allowing the next gate to begin evaluation. This is required, otherwise the SPSD gate being driven could temporarily see both X and $\overline{X}$ as high. Delaying the clock to the following gate prevents false evaluations. To demonstrate the idea, a 64-bit carry-lookahead adder was designed and simulated. The designed adder consists of four stages of SPSD gates based on a blocking factor of four: carry generate/propagate, group generate/propagate, carry evaluate, and 4 bit full adder stage. Each SPSD gate uses minimum sized n-transistors and 6λ wide p-transistors in all. A worst

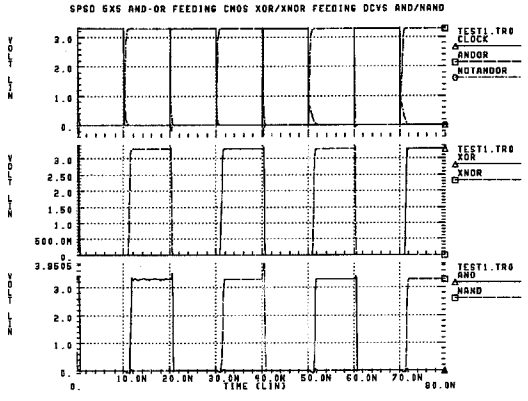

Figure 11. Pipeline logic

case evaluation delay of 4.6 ns was achieved with the blocking factor of 4. The precharge time was set at 1 ns. The CLA adder architecture is shown in figure 13.

The benefit of using SPSD gates to build adders is that the operands can be used directly by complex carry-generator gates, without deriving intermediate propagate and generate signals. Table 1 gives a list of comparative figures to put the SPSD scheme into context. All of the adders where designed and simulated using the same process parameters and having the same transistor sizing. This ensured that the drive strengths of the gate in each of the adders was not biased by transistor sizing.

5. Conclusion

This paper has presented a circuit technique for building high speed CMOS pipelines which are functionally more efficient than conventional domino based pipelines. Along with the eliminating the feedback and buffering circuitry required by domino gates for reliable operation, the SPDP gates offer considerably faster switching speeds than domino, especially with high fan-in gates, and therefore can efficiently be used to implement complex logic, such as branch prediction logic, in a single pipeline stage.

References

[1] R.H. Krambeck, C.M. Lee, and H.S. Law. High-Speed Compact Circuits with CMOS. *IEEE Journal of Solid State Circuits,"* sc-17(3):614-619,June 1982.

[2] V. Friedman and S. Liu. Dynamic Logic CMOS Circuits. *IEEE J. Solid-State Circuits,* sc-19(2):263-266,April 1984.

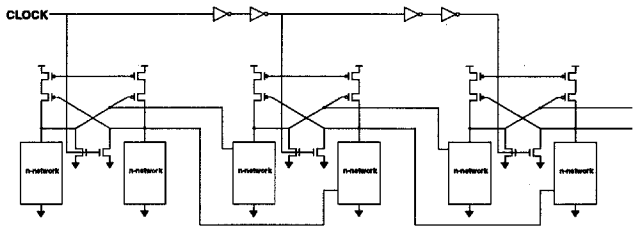

Figure 12. Clock Delayed SPSD domino

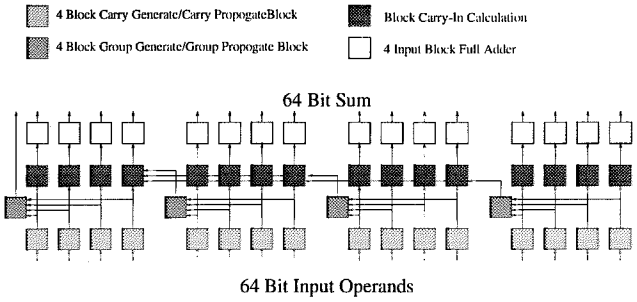

Figure 13. 64 Bit adder architecture

[3] C.M. Lee and E.W. Szeto. Zipper CMOS. *IEEE Circuits and Devices Magazine,* sc-17(3):10-16,May 1986.

[4] Meta-Software. *HSPICE User's Manual H9001.* Meta-Software Inc,Campbell,CA,1990.

[5] Y. Ji-Ren, I. Karlsson, and C. Svensson. A True Single-Phase-Clock Dynamic CMOS Circuit Technique. *IEEE Journal of Solid State Circuits,* sc-22(5):899-901,October 1987.

[6] L.G. Heller W.R. Griffin, J.W. Davis, and N.G. Thoma. Cascode Voltage Switch Logic: A Differential CMOS Logic Family. *Proceedings of International Solid State Circuits Conference,* pp.16-17, 1984.