WRL
Technical Note TN-45

Circuit and Process Directions for Low-Voltage Swing Submicron BiCMOS

Norman P. Jouppi,
Suresh Menon,
Stefanos Sidiropoulos

digital **Western Research Laboratory** 250 University Avenue Palo Alto, California 94301 USA

The Western Research Laboratory (WRL) is a computer systems research group that was founded by Digital Equipment Corporation in 1982. Our focus is computer science research relevant to the design and application of high performance scientific computers. We test our ideas by designing, building, and using real systems. The systems we build are research prototypes; they are not intended to become products.

There is a second research laboratory located in Palo Alto, the Systems Research Center (SRC). Other Digital research groups are located in Paris (PRL) and in Cambridge, Massachusetts (CRL).

Our research is directed towards mainstream high-performance computer systems. Our prototypes are intended to foreshadow the future computing environments used by many Digital customers. The long-term goal of WRL is to aid and accelerate the development of high-performance uni- and multi-processors. The research projects within WRL will address various aspects of high-performance computing.

We believe that significant advances in computer systems do not come from any single technological advance. Technologies, both hardware and software, do not all advance at the same pace. System design is the art of composing systems which use each level of technology in an appropriate balance. A major advance in overall system performance will require reexamination of all aspects of the system.

We do work in the design, fabrication and packaging of hardware; language processing and scaling issues in system software design; and the exploration of new applications areas that are opening up with the advent of higher performance systems. Researchers at WRL cooperate closely and move freely among the various levels of system design. This allows us to explore a wide range of tradeoffs to meet system goals.

We publish the results of our work in a variety of journals, conferences, research reports, and technical notes. This document is a technical note. We use this form for rapid distribution of technical material. Usually this represents research in progress. Research reports are normally accounts of completed research and may include material from earlier technical notes.

Research reports and technical notes may be ordered from us. You may mail your order to:

Technical Report Distribution
DEC Western Research Laboratory, WRL-2
250 University Avenue
Palo Alto, California 94301 USA

Reports and notes may also be ordered by electronic mail. Use one of the following addresses:

Digital E-net:	`JOVE::WRL-TECHREPORTS`
Internet:	`WRL-Techreports@decwrl.pa.dec.com`
UUCP:	`decpa!wrl-techreports`

To obtain more details on ordering by electronic mail, send a message to one of these addresses with the word ‘‘`help`’’ in the Subject line; you will receive detailed instructions.

Circuit and Process Directions for Low-Voltage Swing Submicron BiCMOS

Norman P. Jouppi, Suresh Menon, and Stefanos Sidiropoulos

March, 1994

Abstract

Low-swing (<600mV) submicron BiCMOS circuits have many advantages over full-swing BiCMOS, CMOS, or small-swing bipolar circuits. We show that the optimal speed fan-in for low-swing BiCMOS logic circuits is generally in the range of 7 to 20, depending on the process characteristics and gate topology. This high fan-in means that the bipolar device parasitic capacitances primarily determine the circuit speed and speed-power products, instead of f_T as in the case of low fan-in mux/demux communication circuits. SiGe HBT BiCMOS circuits are attractive for logic circuits not primarily for their higher f_T, but rather for their increased maximum device currents for a given parasitic capacitance and for their smaller V_{be}, which can lower chip power dissipation. Finally, for small-swing BiCMOS circuits to be competitive with CMOS they must also be built from the same lithography as CMOS circuits, have local interconnect for inter-device intra-gate wiring, and be built with a full-custom design methodology.

digital **Western Research Laboratory** 250 University Avenue Palo Alto, California 94301 USA

Table of Contents

List of Figures

1. Introduction

Low-swing BiCMOS circuits typically have logic swings of less than 600mV and use ECL or CML-based logic structures [1]. These swings are significantly smaller than those used in even 1.5V CMOS. Since the time to charge a wire at the output of a gate is proportional to the logic swing, low-swing BiCMOS circuits have a potential inherent speed advantage over CMOS circuits.

Low-swing BiCMOS circuits can use CMOS RAM cells for memory. This offers a significant density advantage (about 4:1) over pure bipolar RAM cells, while providing about the same access times if BiCMOS peripheral circuits are used. For example, in microprocessors a factor of four increase in RAM density can result in a three-fold reduction in cache miss rates. Because cache misses can severely limit the performance of many applications on modern microprocessors, improved RAM density is very important for their performance. This RAM density advantage can also be very useful when implementing buffer memories of ATM switch chips. This gives low-swing BiCMOS circuits a significant advantage over pure bipolar circuits.

The conventional use of BiCMOS circuits for logic uses the bipolar device simply to aid in driving the capacitive load seen by a CMOS gate and not for performing the logic function itself. Here logic swings equal to the supply voltage are used. As the MOS supply voltages scale down with lithography, the V_{be} drop of the output transistor in a conventional full-swing BiCMOS gate becomes a larger and larger percentage of the logic swing and begins to greatly degrade the performance. The use of full-swing BiCMOS circuits has not shown significant promise below 2V supplies, unless both NPN and PNP bipolar devices are available [4].

In contrast, low-swing BiCMOS circuits use bipolar transistors for computing logic functions as well as for driving wires. ECL logic structures work well with logic swings of only 600mV. The supply voltages for ECL BiCMOS logic circuits are not limited by the supply voltage limits of the MOS devices. As the MOS supplies scale down to 1.5V from 5V, interfacing CMOS circuits and ECL circuits becomes easier due to the smaller differences in swings. Thus low-swing BiCMOS circuits can benefit from MOS supply scaling rather than suffer from it, as full-swing BiCMOS circuits do.

Unlike full-swing BiCMOS circuits, ECL-based low-swing BiCMOS logic circuits dissipate static power. However, the use of MOS memories can save considerable power over the power dissipated by pure-bipolar circuits. Also, small-swing active-pull-down circuits [11, 7, 12] have recently been demonstrated that can reduce the static power of the output of a logic gate by almost an order of magnitude. Thus, although the power of low-swing BiCMOS logic circuits will be larger than that required for CMOS, we do not believe it will be prohibitively large in many applications.

In Section 2 we give circuit examples on how BiCMOS can be useful for logic and memory circuits. Section 3 gives process directions based on these circuits and lists other requirements for the successful use of low-swing BiCMOS circuits. Section 4 summarizes the paper.

2. Small-swing BiCMOS circuits

One of the advantages of BiCMOS small-swing circuits over bipolar ECL circuits is the availability of MOS current sources. Figure 1 shows an OR/NOR gate using nMOS current sources. In order to behave as a current source, the nMOS transistors must be in their saturated region. Small nMOS devices can provide currents of 100µA and be kept in saturation as long as $V_{ds} > V_{gs} - V_t$ and V_{ds} can be as low as 0.6V. In contrast a bipolar current source would require a $V_{swing} = 0.6V$ drop across the current source resistor for best tracking and an additional drop of 0.8V across the current source transistor to keep it completely out of saturation. The net result is that a traditional -4.5V current switch supply and a -3.3V emitter follower supply can be reduced to -3.7V and -2.5V, respectively. This can easily save 20% or more of the power of a bipolar-only chip.

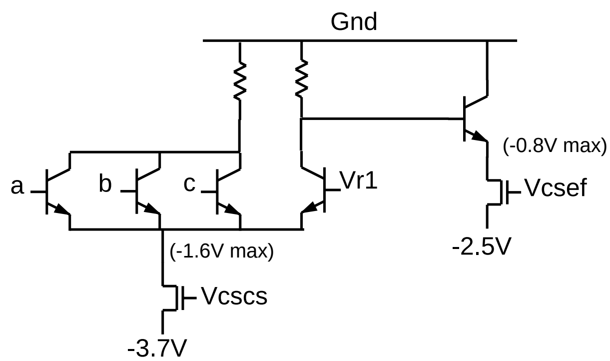

Figure 1: OR/NOR gate using nMOS current sources

The use of a NMOS current source can be limited by either channel punch through or oxide breakdown. Since the current source device usually has at least 2X the minimum channel length, the channel punch through for a 2.5V process should be at least 3.5V. The oxide breakdown is usually significantly higher than the minimum channel width punch through voltage, so it should be at least 3.5V as well. The supply for the gate current switch (V_{ee}) in Figure 1 is -3.7V. This does not present a problem for the use of nMOS current sources since the highest voltage ever seen at the drain of the MOS device is -1.6V, resulting in a V_{ds} of 2.1V. The maximum V_{ds} of the emitter-follower current source is 1.7V. To insure saturation with a V_{ds} of 0.6V and a V_t of 0.6V, V_{cscs} must be 1.2V or less above the negative supply. Thus the nMOS current source operating point is well within the channel punch through and oxide breakdown limits for a 2.5V process, and would likely work even with a 1.5V process.

Another significant advantage of the nMOS current sources is that the nMOS transistors have no gate current corresponding to the base current of a bipolar current source. This makes the distribution of the current source reference voltage much easier since the resistance of the distribution network is not a first-order concern and therefore no IR drops occur in the distribution network.

Large amounts of on-chip memory are crucial for many applications such as microprocessors. Low-swing BiCMOS memory circuits have many advantages over pure bipolar or pure CMOS circuits. Because a CMOS memory core does not dissipate significant power, the memory core can be powered from the larger power supply by using a diode drop on the upper supply and regulator circuit from the bottom supply. This allows bipolar pull-ups and active nMOS pull-down circuits to be used to drive RAM word lines without speed degradation, since the MOS RAM core has its upper supply shifted by a diode drop as well. Other small-swing circuits, such as wired-ORs, can be very useful for building fast decoders. Bipolar cascode circuits enable very fast sensing. The combination of a CMOS core with BiCMOS peripheral circuits can achieve about the same density as pure CMOS but with about 2X higher performance.

2.1. Function delay vs. gate complexity

When implementing very complex logic functions, such as those required by a 64 bit microprocessor, there are many possibilities for restructuring the design's logic equations. Any logic equation can be represented in two levels of logic (e.g., canonical sum-of-products form), however this extreme approach can result in an explosion of the fan-in per logic stage for complex functions. Other structures of the logic equations are possible that use very small fan-ins (e.g., 2 or 3) but have very many stages of logic. For example, 64-bit carry lookahead adders could be constructed from 6 stages of 2 bit groups, 3 stages of 4 bit groups, or 2 stages of 8 bit groups. In this section we discuss the best logic structures for low-swing BiCMOS circuits.

Figure 3 shows the delay versus the fan-in of a low-swing BiCMOS NOR gate implemented in the 0.6μm process of Table 2. The gate delay is measured by simulating a 19-stage ring oscillator. All the devices in the gate are minimum size and both the current switch and the emitter follower are operated at a 350μA current. A 10X increase in the gate fan-in (from an inverter to a 10 input NOR gate) results in only a 2.2X increase in gate delay. The second curve in Figure 3 shows the delay of a NOR gate with the same fan-in when implemented as a two stage network. The delay of the two stage gate network is larger than the delay of a single higher fan-in gate until a fan-in of 14 is reached.

	0.8um [5]	0.6um [13]
device type	single-poly	double-poly
A_E	0.8 X 1.6um	0.5 X 2.0um
β	90	100
f_T	15 GHz	20 GHz
R_b	700 ohms	250 ohms
C_{js}	10.2ff	8.0ff
C_{jc}	2.9ff	3.5ff
C_{je}	3.3ff	6.0ff

Figure 2: Bipolar transistor parameters used in the simulations

If speed-power product is used as the metric, the crossover for splitting a logic function into more than one stage pushes out even further. Figure 4 shows the same comparison in terms of speed-power product. The large steps in the 2-stage curve occur when another gate must be added to the gate tree to handle the increased fan-in, while the small steps occur when the fan-in of a gate in the 2-stage network increases by one. For example, the large step between a fan-in of 7 and 8 occurs when going from two fan-in of 3 gates feeding a gate with fan-in of 3 to three

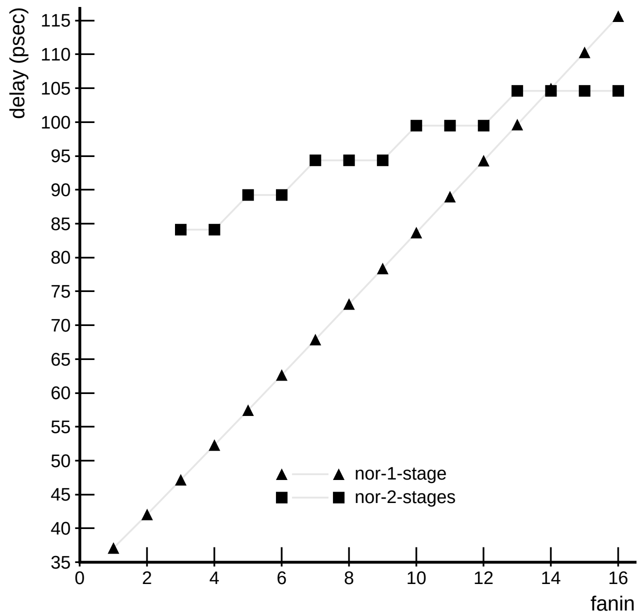

Figure 3: ECL NOR gate delay versus fan-in

gates with fan-in of 3 feeding a gate with fan-in of 3. Looking at the trends of the 1-stage and 2-stage speed-power products, it can be seen that the lines are diverging. Thus it is always optimal from a speed-power standpoint to implement a wide NOR function in a single stage of logic. An implication for circuit noise margins is that it makes sense to allow a very large ΔV_{be} due to current sharing among in OR/NOR structures. By limiting the maximum OR/NOR fan-in to 32, a noise allowance of about 115mV would be sufficient.

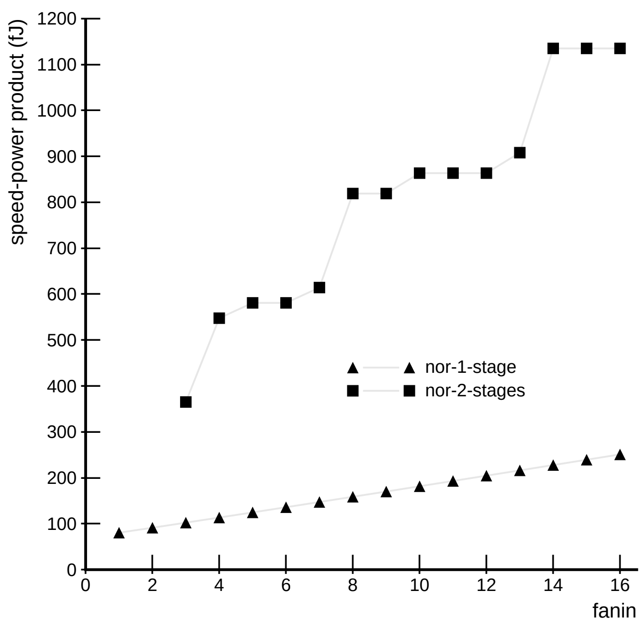

Figure 4: ECL NOR speed-power product versus fan-in

These high optimal fan-ins occur for other ECL processes as well. For example, the one-stage vs. two-stage crossovers for the 0.8um process in Table 2 and a low-stress trench-isolated 0.8μm process [9] are 14 as well. Unfortunately, in the gate array and standard cell design methodologies that have been common with ECL circuits to date, most circuits in the cell libraries have had fairly small per-stage fan-ins. It is not uncommon for the maximum fan-in to be only 8, and the average fan-in to be only 3 or 4. This results in poor circuit delay and power dissipation in comparison to optimal fan-in circuits.

2.2. Optimal fan-ins for CMOS vs. low-swing BiCMOS

One factor which is overlooked in many comparisons of CMOS and ECL circuit technologies is that ECL has better fan-in and fan-out capabilities than CMOS. Figure 5 plots the ratio of a CMOS static NAND over an ECL NOR gate delay versus varying gate fan-in and fan-out. Thus the X-axis in Figure 5 represents the ‘‘logic power’’ of each circuit style. The delays of the gates are from simulation of gates built in two contemporary 0.8 um CMOS [10], and BiCMOS [5] technologies. The ECL gate uses minimum devices and switch and emitter follower currents of 200μA. Figure 5 shows that a single stage gate implemented in static CMOS becomes much slower than a corresponding gate in ECL as the gate complexity and fan-out requirements increase. For fan-in = fan-out = 1, the ECL gate is only 3.3 times faster than the CMOS gate. Thus when comparing CMOS and ECL ring oscillator delays, the ECL gates may not appear to be much faster. However, for logic applications an ECL inverter is largely a useless circuit since most gates can produce true and complement outputs and gates have high overall current gain, so that the taper buffers common in CMOS circuits are not required. As the usefulness of the gate logic function increases, the speed advantage of ECL over CMOS increases. This shows that logic comparisons that compare small "toy" logic equations with fan-ins of only two or three are biased towards CMOS. Real applications, such as 64-bit adders, afford many opportunities for very high fan-in gates.

Of course this comparison is not the whole story. Other circuit techniques are available in both CMOS and ECL for improving the performance of high fan-in gates. For example, dynamic logic families in CMOS avoid the extra capacitance of many large p-channel devices or the high-resistance of many stacked p-channel devices. Differential CMOS logic families also can offer reduced delays, but at the expense of increased power dissipation. These more advanced circuit families are not applicable in all circumstances, but are generally used widely in modern high-performance microprocessors. Similarly, wired-OR circuits in ECL (emitter dotting) offer reduced delay and power over an ECL OR gate. Differential and cascode circuits can provide very high speeds for ECL fan-ins of 50 or more. Unfortunately these circuits are not typically provided or even allowed in gate array or standard cell design systems, which are predominately used for ECL logic design.

2.3. Communication logic circuits vs. computer logic circuits

In communication circuits one of the most important design criteria is the maximum sustainable bandwidth, while in logic applications one of the most important criteria is the minimum latency. This leads communication circuits to typically limit gate fan-ins to a maximum of two, and to use many gates in series to provide the equivalent logic functionality of larger fan-ins.

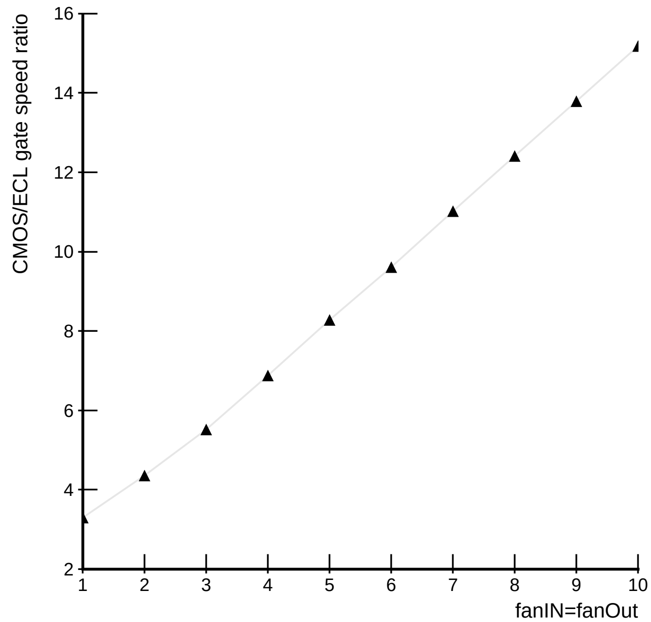

Figure 5: Ratio of 0.8um CMOS static NAND to 0.8um ECL NOR gate delay

While this allows higher bandwidths to be sustained, it increases the overall latency and so is not acceptable in logic circuits. Figure 6 plots the bandwidth vs. latency vs. speed-power product of implementing various multiplexors from 2 through 16 inputs either as a single gate or as a tree of 2-input multiplexor gates. The bandwidth advantages of using only 2-input multiplexor building blocks is clear; even the bandwidth of a 3 or 4-input multiplexor is dramatically less. However the gate delay crossover between one large fan-in multiplexor and a tree of 2-input multiplexors does not occur until a fan-in of 11 is reached. Again, the speed-power products of the two implementations diverge, meaning the single gate always has a better speed-power product. Gates with somewhat larger fan-in than 11 pay only a small delay penalty, but have a large power advantage. This difference in optimal gate fan-ins between communication and logic circuits can have a significant effect on the importance of different bipolar transistor parameters, as we shall see in the next section.

3. Requirements for competitive small-swing submicron BiCMOS

In the previous section we discussed the bipolar device characteristics which would be most favorable for low-swing BiCMOS circuits. This section presents the resulting process features and CAD/design methodology requirements for competitive low-swing BiCMOS circuits.

3.1. Bipolar device parameter delay sensitivities

Figure 7 shows how the delay of an 8-input multiplexor varies as f_T, C_{jc}, C_{js}, and C_{je} are increased or decreased by up to a factor of two for the 0.6μm process parameters given in Table 2. One of the first things to notice is that a factor of two reduction in f_T (from 20Ghz to 10GHz) results in less than 10% speed degradation of the multiplexor. Instead, the device capacitances C_{jc} and C_{js} are by far the most important device properties for large fan-in multiplexors. Figure

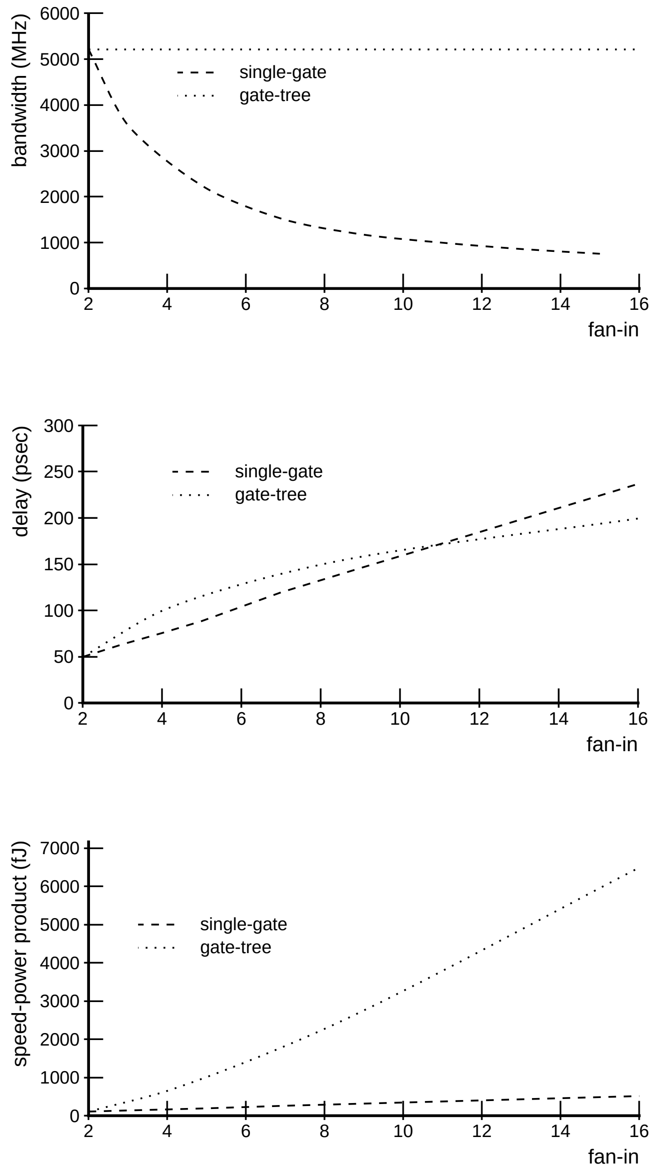

Figure 6: Multiplexor bandwidth, latency, and speed-power product

8 shows how the delay of an 8-input NOR gate varies as various transistors parameters are varied. Again C_{jc} and C_{js} are the dominant terms, although f_T is relatively more important than for the multiplexor. This device parameter sensitivity is in sharp contrast to the sensitivity of small 2-fan-in differential communication circuits, where the $\Delta V \times C$ delay terms are much smaller due to the smaller fan-ins and smaller differential swings. Here f_T alone is a good predictor of circuit bandwidth [14].

The most common technology benchmarks for ECL logic circuits are single-ended swing ring oscillators. These circuits have similar device parameter sensitivities as 2-fan-in differential circuits. Figure 9 shows a sensitivity analysis for a ring oscillator with five stages of buffers and five of inverters. For logic applications, however, an ECL inverter or buffer is largely a useless circuit. If logic applications are at all being considered as a target of process development, much better benchmarks would be fan-in = fan-out = 8 multiplexors and NOR gates.

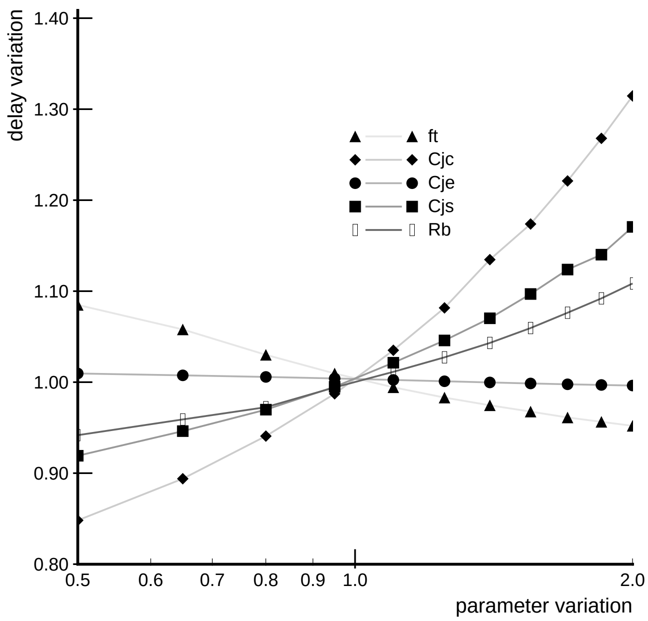

Figure 7: 8-input multiplexor delay sensitivity analysis

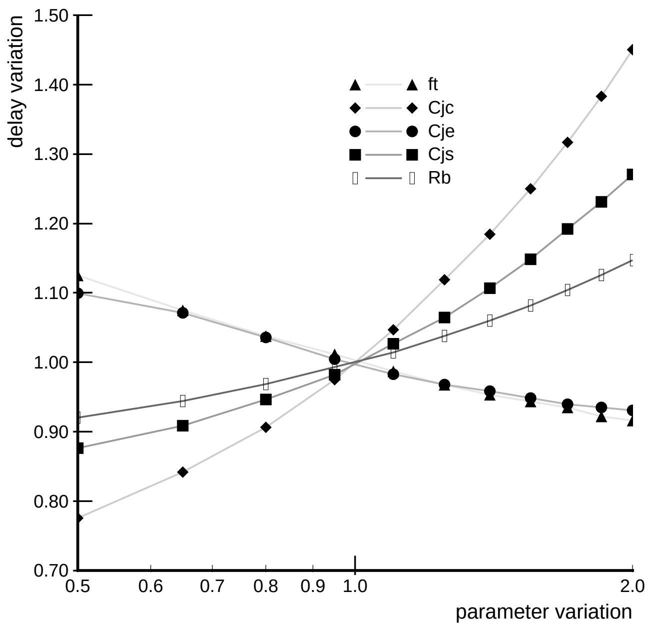

Figure 8: 8-input NOR gate delay sensitivity analysis

We can define a logic speed figure of merit for a bipolar device which is the reflects the average sensitivity of the dominant delay terms for the multiplexor and NOR gates. The average sensitivity to C_{jc} in Figures 7 and 8 is 38% while the average sensitivity to C_{js} is 22%. C_{js} is less important than C_{jc} because it is reversed-biased and there is no Miller effect. Thus our simple figure of merit is:

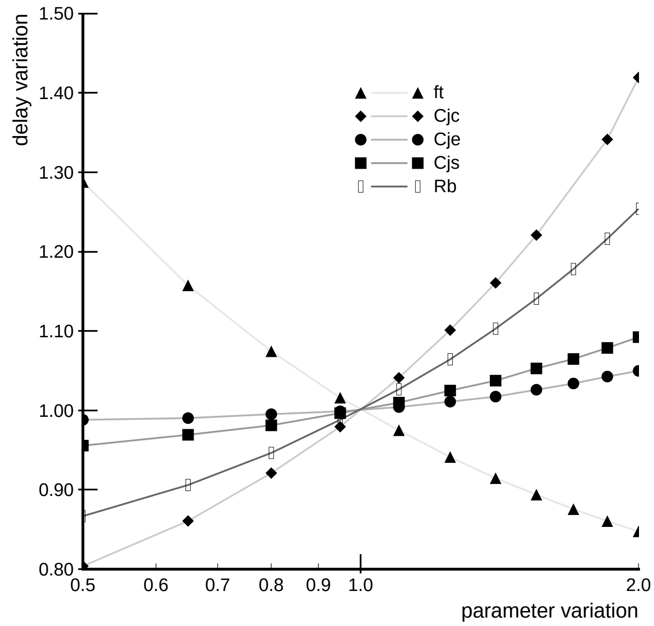

Figure 9: High-speed ring oscillator delay sensitivity analysis

$$Logic\_speed_{FOM} = \frac{I_{max}}{0.38 \times C_{jc} + 0.22 \times C_{js}}$$

However, this does not take into account power, which is not an unlimited resource on a VLSI chip. Dividing by the current to get a speed-power product figure of merit (the power supply voltage remains constant so it can be omitted):

$$Speed\_power\_product_{FOM} = \frac{1}{0.38 \times C_{jc} + 0.22 \times C_{js}}$$

Finally, circuit density is also a measure of computational power [3]. Combining the two figures of merit above and dividing by the device area, we get a systems figure of merit:

$$System\_performance_{FOM} = \frac{I_{max}}{(0.38 \times C_{jc} + 0.22 \times C_{js})^2 \times A_{device}}$$

This figure of merit is quite different than traditional bipolar transistor optimization criteria.

3.2. Lithography

One of the biggest limitations of gate array and standard cell ECL circuits in comparison to full-custom CMOS circuits has been their poorer circuit density and integration. This has often been compounded by the availability of coarser lithography in contemporary VLSI bipolar processes in comparison to CMOS processes. Circuit density is one of the most important parameters in determining overall system performance [3]. For example, with a lithographic feature size better by 1.4X, twice the number of components are available on-chip. This can directly translate to 2X better system performance in microprocessors by allowing multipliers to retire twice as many bits per cycle, processors to issue twice as many instructions per cycle, etc. A factor of three advantage in circuit performance can all to easily be thrown away with coarser lithography. Thus it is essential that the bipolar devices in a BiCMOS process be jointly developed with the CMOS devices in the same time frame as pure CMOS processes with the same lithography.

Simultaneous development of CMOS and bipolar devices is easiest if the bipolar device shares as many steps with the CMOS process as possible. Simultaneous development is also aided by having a bipolar device which is scalable with lithography. In these respects single-poly devices have many advantages over double-poly or more exotic bipolar device structures.

3.3. Interconnect

Just as lithography is crucial for density, so is adequate interconnect. Full-custom CMOS circuits significantly improve their density through the use of silicided local diffusion and polysilicon wiring. In many double-poly processes, the use of silicide for local interconnect between device terminals is not allowed. Thus typically in ECL gate arrays only metal is used for device connections. The recent design of a full-custom ECL microprocessor has shown that if local interconnect is available, the majority of intra-gate wiring connections can be made without the use of metal [6]. This combined with the wire planning which is done in custom designs allows the devices to be packed at minimum spacing across an entire die, and significantly improves system density and performance.

3.4. Impact of heterostructures

One very promising process development for low-swing BiCMOS logic circuits is SiGe HBT BiCMOS processes [2]. SiGe HBT BiCMOS is promising for two primary reasons: increased current densities and a reduced V_{be}. As we saw with our logic speed figure-of-merit, the logic speed depends primarily on the maximum device current divided by device capacitances. Since SiGe HBTs can be developed with similar device parasitics for the same device structure, but allow much higher current densities, they should give much higher logic speeds. Also, because the V_{be} of the SiGe HBT can be about 0.2V less than a Si BJT, the power supply of the chip can be lowered almost proportionally. With modern active pull-down circuits and full-custom design, the vast majority of the power would be dissipated in the gate current switches themselves. Thus a reduction in the gate current switch power supply voltage would result in a commensurate power dissipation reduction.

3.5. CAD/Design methodology requirements

Although the use of small-swing BiCMOS circuits can give a performance advantage over CMOS circuits, it is important not to throw this potential performance advantage away by using an inappropriate design style. ECL logic circuits have historically been used in multichip gate-array processors with low density and performance in comparison to full-custom CMOS microprocessors. This has led many people to the erroneous conclusion that CMOS circuits have become faster than ECL circuits. We believe a more accurate conclusion is that ECL design techniques have remained mired in a design technique over the past decade which throws away much of their performance (e.g., gate arrays or standard cells), while CMOS full-custom design techniques have continued to improve, negating most of the inherent speed advantage of ECL.

To illustrate this point, consider the recent remapping of a Unisys mainframe from many ECL gate arrays into many CMOS gate arrays [8]. The Unisys 2200/900 uses a 1.5µm bipolar ECL technology and has a performance of 40 MIPS. The Unisys 2200/500 uses 0.8µm CMOS gate arrays and has a performance of 10 MIPS. In this case when the same design styles are used, even though the lithography used in the ECL machine is worse by a factor of two, the ECL machine still has four times the performance of the CMOS implementation. Does this mean that a full-custom small-swing BiCMOS microprocessor should be expected to have four times the performance of a similar full-custom CMOS microprocessor? Our experience with a full-custom 1µm ECL microprocessor [6] has lead us to believe that a significant performance advantage can be obtained with full-custom small-swing circuits.

4. Conclusions

Low-swing (<600mV) submicron BiCMOS circuits have many advantages over full-swing BiCMOS, CMOS, or small-swing bipolar circuits. Low-swing BiCMOS circuits offer a significant speed advantage over CMOS circuits while offering better density and lower power dissipation than small-swing bipolar circuits. The static power dissipation of low-swing BiCMOS circuits does remain higher than than of pure CMOS circuits. However, unlike conventional full-swing BiCMOS circuits, which lose their advantages over pure CMOS circuits at reduced supply voltages, small-swing bipolar circuits become more attractive with MOS supply voltage scaling.

The optimal speed fan-in for low-swing BiCMOS logic circuits is generally in the range of 7 to 20, depending on the process characteristics and gate topology. When speed-power is considered, the optimum is to always use a single stage of logic where possible. These degrees of fan-in are much larger than have been historically provided in ECL gate array or standard cell libraries.

The best process characteristics for implementing low-swing BiCMOS logic and memory circuits are quite different from the best process characteristics for communication circuits. Logic circuits have high fan-ins and fan-outs in comparison to communication circuits, and have larger single-ended swings in comparison to the smaller differential swings of communication circuits. Because of this the most important bipolar device characteristics are just the maximum bipolar device current over the device capacitances. The importance of f_T can be lower by almost an order of magnitude for logic circuits in comparison to communication circuits. Although the higher f_T of SiGe would be important for communication circuits, it is primarily the higher device current densities supported by the SiGe devices along with their lower V_{be} that are attractive for logic circuits.

Whatever process is used for implementing small-swing BiCMOS circuits, for them to be competitive with CMOS they must be built from the same lithography as CMOS circuits, have local interconnect for inter-device intra-gate wiring, and be built with a full-custom design methodology. Otherwise the circuit speed afforded by the small-swing BiCMOS will be squandered away.

References

[1] Chih-Liang Chen. 2.5V Bipolar/CMOS Circuits for 0.25um BiCMOS Technology. *IEEE Journal of Solid-State Circuits* 27(4):485-491, April, 1992.

[2] D. L. Harame, et. al. A High-Performance SiGe-Base ECL BiCMOS Technology. *Tech. Dig. of the IEDM* :19-22, December, 1992.

[3] John L. Hennessy and Norman P. Jouppi. Computer Technology and Architecture: An Evolving Interaction. *Computer* 24(9):18-29, September, 1991.

[4] Mitsuru Hiraki, et. al. A 1.5V Full-Swing BiCMOS Logic Circuit. *IEEE Journal of Solid-State Circuits* 27(11):1568-1574, November, 1992.

[5] A. Iranmanesh, et. al. A 0.6um Single-Poly Advanced BiCMOS Technology for ASIC Applications. *VLSI Technology Symposium* :87-88, June, 1990.

[6] Norman P. Jouppi, et. al. A 300Mhz 115W 32b Bipolar ECL Microprocessor. *IEEE Journal of Solid-State Circuits* 28(11):1152-1166, November, 1993.

[7] Norman P. Jouppi. A Fully-Compensated APD Circuit with 10:1 Ratio Between Active and Inactive Current. *Bipolar and BiCMOS Circuits and Technology Meeting* :111-114, October, 1994.

[8] Reinhardt Krause. Unisys Improves CMOS with Motorola Devices. *Electronic News* :1, September 20, 1993.

[9] S. Matsuda, et. al. A Low-Stress Trench Isolation Structure and its Electrical Characteristics of Sub 20 ps High-Speed ECL. *VLSI Technology Symposium* :73-74, May, 1993.

[10] MOSIS 0.8μm BiCMOS process parameters. Fabricated through Hewlett-Packard.

[11] Hyun J. Shin. Self Biased Feedback-Controlled Pull-Down Emitter Follower for High Speed Low-Power Bipolar Logic Circuits. *VLSI Circuits Symposium* :27-28, May, 1993.

[12] Stefanos Sidiropoulos, Norman P. Jouppi, and Suresh Menon. A Speed, Power, and Supply Noise Evaluation of ECL Driver Circuits. *Bipolar and BiCMOS Circuits and Technology Meeting* :119-122, October, 1994.

[13] N. Tamba, et. al. A 1.5ns BiCMOS SRAM with 11K 60ps Logic Gates. *ISSCC Dig. Tech. Papers* , February, 1993.

[14] P. K. Tien. Propagation Delay in High-Speed Silicon Bipolar and GaAs HBT Digital Circuits. *International Journal of High-Speed Electronics* 1(1):101-124, March, 1990.

WRL Research Reports

‘‘Titan System Manual.’’
Michael J. K. Nielsen.
WRL Research Report 86/1, September 1986.

‘‘Global Register Allocation at Link Time.’’
David W. Wall.
WRL Research Report 86/3, October 1986.

‘‘Optimal Finned Heat Sinks.’’
William R. Hamburgen.
WRL Research Report 86/4, October 1986.

‘‘The Mahler Experience: Using an Intermediate Language as the Machine Description.’’
David W. Wall and Michael L. Powell.
WRL Research Report 87/1, August 1987.

‘‘The Packet Filter: An Efficient Mechanism for User-level Network Code.’’
Jeffrey C. Mogul, Richard F. Rashid, Michael J. Accetta.
WRL Research Report 87/2, November 1987.

‘‘Fragmentation Considered Harmful.’’
Christopher A. Kent, Jeffrey C. Mogul.
WRL Research Report 87/3, December 1987.

‘‘Cache Coherence in Distributed Systems.’’
Christopher A. Kent.
WRL Research Report 87/4, December 1987.

‘‘Register Windows vs. Register Allocation.’’
David W. Wall.
WRL Research Report 87/5, December 1987.

‘‘Editing Graphical Objects Using Procedural Representations.’’
Paul J. Asente.
WRL Research Report 87/6, November 1987.

‘‘The USENET Cookbook: an Experiment in Electronic Publication.’’
Brian K. Reid.
WRL Research Report 87/7, December 1987.

‘‘MultiTitan: Four Architecture Papers.’’
Norman P. Jouppi, Jeremy Dion, David Boggs, Michael J. K. Nielsen.
WRL Research Report 87/8, April 1988.

‘‘Fast Printed Circuit Board Routing.’’
Jeremy Dion.
WRL Research Report 88/1, March 1988.

‘‘Compacting Garbage Collection with Ambiguous Roots.’’
Joel F. Bartlett.
WRL Research Report 88/2, February 1988.

‘‘The Experimental Literature of The Internet: An Annotated Bibliography.’’
Jeffrey C. Mogul.
WRL Research Report 88/3, August 1988.

‘‘Measured Capacity of an Ethernet: Myths and Reality.’’
David R. Boggs, Jeffrey C. Mogul, Christopher A. Kent.
WRL Research Report 88/4, September 1988.

‘‘Visa Protocols for Controlling Inter-Organizational Datagram Flow: Extended Description.’’
Deborah Estrin, Jeffrey C. Mogul, Gene Tsudik, Kamaljit Anand.
WRL Research Report 88/5, December 1988.

‘‘SCHEME->C A Portable Scheme-to-C Compiler.’’
Joel F. Bartlett.
WRL Research Report 89/1, January 1989.

‘‘Optimal Group Distribution in Carry-Skip Adders.’’
Silvio Turrini.
WRL Research Report 89/2, February 1989.

‘‘Precise Robotic Paste Dot Dispensing.’’
William R. Hamburgen.
WRL Research Report 89/3, February 1989.

‘‘Simple and Flexible Datagram Access Controls for Unix-based Gateways.’’
Jeffrey C. Mogul.
WRL Research Report 89/4, March 1989.

‘‘Spritely NFS: Implementation and Performance of Cache-Consistency Protocols.’’
V. Srinivasan and Jeffrey C. Mogul.
WRL Research Report 89/5, May 1989.

‘‘Available Instruction-Level Parallelism for Superscalar and Superpipelined Machines.’’
Norman P. Jouppi and David W. Wall.
WRL Research Report 89/7, July 1989.

‘‘A Unified Vector/Scalar Floating-Point Architecture.’’
Norman P. Jouppi, Jonathan Bertoni, and David W. Wall.
WRL Research Report 89/8, July 1989.

‘‘Architectural and Organizational Tradeoffs in the Design of the MultiTitan CPU.’’
Norman P. Jouppi.
WRL Research Report 89/9, July 1989.

‘‘Integration and Packaging Plateaus of Processor Performance.’’
Norman P. Jouppi.
WRL Research Report 89/10, July 1989.

‘‘A 20-MIPS Sustained 32-bit CMOS Microprocessor with High Ratio of Sustained to Peak Performance.’’
Norman P. Jouppi and Jeffrey Y. F. Tang.
WRL Research Report 89/11, July 1989.

‘‘The Distribution of Instruction-Level and Machine Parallelism and Its Effect on Performance.’’
Norman P. Jouppi.
WRL Research Report 89/13, July 1989.

‘‘Long Address Traces from RISC Machines: Generation and Analysis.’’
Anita Borg, R.E.Kessler, Georgia Lazana, and David W. Wall.
WRL Research Report 89/14, September 1989.

‘‘Link-Time Code Modification.’’
David W. Wall.
WRL Research Report 89/17, September 1989.

‘‘Noise Issues in the ECL Circuit Family.’’
Jeffrey Y.F. Tang and J. Leon Yang.
WRL Research Report 90/1, January 1990.

‘‘Efficient Generation of Test Patterns Using Boolean Satisfiablilty.’’
Tracy Larrabee.
WRL Research Report 90/2, February 1990.

‘‘Two Papers on Test Pattern Generation.’’
Tracy Larrabee.
WRL Research Report 90/3, March 1990.

‘‘Virtual Memory vs. The File System.’’
Michael N. Nelson.
WRL Research Report 90/4, March 1990.

‘‘Efficient Use of Workstations for Passive Monitoring of Local Area Networks.’’
Jeffrey C. Mogul.
WRL Research Report 90/5, July 1990.

‘‘A One-Dimensional Thermal Model for the VAX 9000 Multi Chip Units.’’
John S. Fitch.
WRL Research Report 90/6, July 1990.

‘‘1990 DECWRL/Livermore Magic Release.’’
Robert N. Mayo, Michael H. Arnold, Walter S. Scott, Don Stark, Gordon T. Hamachi.
WRL Research Report 90/7, September 1990.

‘‘Pool Boiling Enhancement Techniques for Water at Low Pressure.’’
Wade R. McGillis, John S. Fitch, William R. Hamburgen, Van P. Carey.
WRL Research Report 90/9, December 1990.

‘‘Writing Fast X Servers for Dumb Color Frame Buffers.’’
Joel McCormack.
WRL Research Report 91/1, February 1991.

‘‘A Simulation Based Study of TLB Performance.’’
J. Bradley Chen, Anita Borg, Norman P. Jouppi.
WRL Research Report 91/2, November 1991.

‘‘Analysis of Power Supply Networks in VLSI Circuits.’’
Don Stark.
WRL Research Report 91/3, April 1991.

‘‘TurboChannel T1 Adapter.’’
David Boggs.
WRL Research Report 91/4, April 1991.

‘‘Procedure Merging with Instruction Caches.’’
Scott McFarling.
WRL Research Report 91/5, March 1991.

‘‘Don’t Fidget with Widgets, Draw!.’’
Joel Bartlett.
WRL Research Report 91/6, May 1991.

‘‘Pool Boiling on Small Heat Dissipating Elements in Water at Subatmospheric Pressure.’’
Wade R. McGillis, John S. Fitch, William R. Hamburgen, Van P. Carey.
WRL Research Report 91/7, June 1991.

‘‘Incremental, Generational Mostly-Copying Garbage Collection in Uncooperative Environments.’’
G. May Yip.
WRL Research Report 91/8, June 1991.

‘‘Interleaved Fin Thermal Connectors for Multichip Modules.’’
William R. Hamburgen.
WRL Research Report 91/9, August 1991.

‘‘Experience with a Software-defined Machine Architecture.’’
David W. Wall.
WRL Research Report 91/10, August 1991.

‘‘Network Locality at the Scale of Processes.’’
Jeffrey C. Mogul.
WRL Research Report 91/11, November 1991.

‘‘Cache Write Policies and Performance.’’
Norman P. Jouppi.
WRL Research Report 91/12, December 1991.

‘‘Packaging a 150 W Bipolar ECL Microprocessor.’’
William R. Hamburgen, John S. Fitch.
WRL Research Report 92/1, March 1992.

‘‘Observing TCP Dynamics in Real Networks.’’
Jeffrey C. Mogul.
WRL Research Report 92/2, April 1992.

‘‘Systems for Late Code Modification.’’
David W. Wall.
WRL Research Report 92/3, May 1992.

‘‘Piecewise Linear Models for Switch-Level Simulation.’’
Russell Kao.
WRL Research Report 92/5, September 1992.

‘‘A Practical System for Intermodule Code Optimization at Link-Time.’’
Amitabh Srivastava and David W. Wall.
WRL Research Report 92/6, December 1992.

‘‘A Smart Frame Buffer.’’
Joel McCormack & Bob McNamara.
WRL Research Report 93/1, January 1993.

‘‘Recovery in Spritely NFS.’’
Jeffrey C. Mogul.
WRL Research Report 93/2, June 1993.

‘‘Tradeoffs in Two-Level On-Chip Caching.’’
Norman P. Jouppi & Steven J.E. Wilton.
WRL Research Report 93/3, October 1993.

‘‘Unreachable Procedures in Object-oriented Programing.’’
Amitabh Srivastava.
WRL Research Report 93/4, August 1993.

‘‘An Enhanced Access and Cycle Time Model for On-Chip Caches.’’
Steven J.E. Wilton and Norman P. Jouppi.
WRL Research Report 93/5, July 1994.

‘‘Limits of Instruction-Level Parallelism.’’
David W. Wall.
WRL Research Report 93/6, November 1993.

‘‘Fluoroelastomer Pressure Pad Design for Microelectronic Applications.’’
Alberto Makino, William R. Hamburgen, John S. Fitch.
WRL Research Report 93/7, November 1993.

‘‘A 300MHz 115W 32b Bipolar ECL Microprocessor.’’
Norman P. Jouppi, Patrick Boyle, Jeremy Dion, Mary Jo Doherty, Alan Eustace, Ramsey Haddad, Robert Mayo, Suresh Menon, Louis Monier, Don Stark, Silvio Turrini, Leon Yang, John Fitch, William Hamburgen, Russell Kao, and Richard Swan.
WRL Research Report 93/8, December 1993.

‘‘Link-Time Optimization of Address Calculation on a 64-bit Architecture.’’
Amitabh Srivastava, David W. Wall.
WRL Research Report 94/1, February 1994.

‘‘ATOM: A System for Building Customized Program Analysis Tools.’’
Amitabh Srivastava, Alan Eustace.
WRL Research Report 94/2, March 1994.

‘‘Complexity/Performance Tradeoffs with Non-Blocking Loads.’’
Keith I. Farkas, Norman P. Jouppi.
WRL Research Report 94/3, March 1994.

‘‘A Better Update Policy.’’
Jeffrey C. Mogul.
WRL Research Report 94/4, April 1994.

‘‘Boolean Matching for Full-Custom ECL Gates.’’
Robert N. Mayo, Herve Touati.
WRL Research Report 94/5, April 1994.

‘‘Software Methods for System Address Tracing: Implementation and Validation.’’
J. Bradley Chen, David W. Wall, and Anita Borg.
WRL Research Report 94/6, September 1994.

‘‘Performance Implications of Multiple Pointer Sizes.’’
Jeffrey C. Mogul, Joel F. Bartlett, Robert N. Mayo, and Amitabh Srivastava.
WRL Research Report 94/7, December 1994.

WRL Technical Notes

‘‘TCP/IP PrintServer: Print Server Protocol.’’
Brian K. Reid and Christopher A. Kent.
WRL Technical Note TN-4, September 1988.

‘‘TCP/IP PrintServer: Server Architecture and Implementation.’’
Christopher A. Kent.
WRL Technical Note TN-7, November 1988.

‘‘Smart Code, Stupid Memory: A Fast X Server for a Dumb Color Frame Buffer.’’
Joel McCormack.
WRL Technical Note TN-9, September 1989.

‘‘Why Aren’t Operating Systems Getting Faster As Fast As Hardware?’’
John Ousterhout.
WRL Technical Note TN-11, October 1989.

‘‘Mostly-Copying Garbage Collection Picks Up Generations and C++.’’
Joel F. Bartlett.
WRL Technical Note TN-12, October 1989.

‘‘The Effect of Context Switches on Cache Performance.’’
Jeffrey C. Mogul and Anita Borg.
WRL Technical Note TN-16, December 1990.

‘‘MTOOL: A Method For Detecting Memory Bottlenecks.’’
Aaron Goldberg and John Hennessy.
WRL Technical Note TN-17, December 1990.

‘‘Predicting Program Behavior Using Real or Estimated Profiles.’’
David W. Wall.
WRL Technical Note TN-18, December 1990.

‘‘Cache Replacement with Dynamic Exclusion’’
Scott McFarling.
WRL Technical Note TN-22, November 1991.

‘‘Boiling Binary Mixtures at Subatmospheric Pressures’’
Wade R. McGillis, John S. Fitch, William R. Hamburgen, Van P. Carey.
WRL Technical Note TN-23, January 1992.

‘‘A Comparison of Acoustic and Infrared Inspection Techniques for Die Attach’’
John S. Fitch.
WRL Technical Note TN-24, January 1992.

‘‘TurboChannel Versatec Adapter’’
David Boggs.
WRL Technical Note TN-26, January 1992.

‘‘A Recovery Protocol For Spritely NFS’’
Jeffrey C. Mogul.
WRL Technical Note TN-27, April 1992.

‘‘Electrical Evaluation Of The BIPS-0 Package’’
Patrick D. Boyle.
WRL Technical Note TN-29, July 1992.

‘‘Transparent Controls for Interactive Graphics’’
Joel F. Bartlett.
WRL Technical Note TN-30, July 1992.

‘‘Design Tools for BIPS-0’’
Jeremy Dion & Louis Monier.
WRL Technical Note TN-32, December 1992.

‘‘Link-Time Optimization of Address Calculation on a 64-Bit Architecture’’
Amitabh Srivastava and David W. Wall.
WRL Technical Note TN-35, June 1993.

‘‘Combining Branch Predictors’’
Scott McFarling.
WRL Technical Note TN-36, June 1993.

‘‘Boolean Matching for Full-Custom ECL Gates’’
Robert N. Mayo and Herve Touati.
WRL Technical Note TN-37, June 1993.

‘‘Circuit and Process Directions for Low-Voltage Swing Submicron BiCMOS Circuits’’
Norman P. Jouppi, Suresh Menon, and Stefanos Sidiropoulos.
WRL Technical Note TN-45, March 1994.