

The Design of High-Performance Microprocessors at Digital

Thomas F. Fox
Semiconductor Engineering Group
Digital Equipment Corporation
Hudson, MA 01749 USA

Abstract--Today's high-performance single-chip CMOS microprocessors are the most complex and challenging chip designs ever implemented. To stay on the leading edge, Digital's microprocessor designers strive for global optimizations across the design space, develop innovative high-speed custom circuit design techniques, and use state-of-the-art manufacturing technology for fabrication. With microprocessor performance increasing at a rate of 50% per year, quick time-to-market is crucial. Logic and circuit bugs that would prevent first-pass silicon from booting the operating system must be eliminated before tape out if unacceptable schedules delays are to be avoided. Quick time-to-market is achieved by: fostering good design practice through the use of a consistent methodology and detailed design reviews; accounting for manufacturing variations and long-term reliability concerns from the start of the design; and subjecting the chip models to rigorous logical and electrical verification prior to prototype fabrication. This paper presents a practical view of Digital's design process.

I. Introduction

The design of high-performance microprocessors poses many complexity management and technical challenges. These challenges are intensified by the need to bring designs to market quickly. This paper describes Digital Equipment Corporation's approach to high-performance microprocessor design as practiced on projects such as the 100 MHz CISC NVAX microprocessor [1] and the 200 MHz RISC Alpha AXP 21064 microprocessor [2]. Summaries of the characteristics of these microprocessors are given in Tables I and II, and a description of the 0.75 micron CMOS-4 technology used to manufacture them is given in Table III. Details of execution vary from project to project but the overall design procedure is similar. Digital's approach to

design is the result of continuous learning through many generations of microprocessor projects spanning a decade and a half.

II. Nature of the Design

The design process described here is that used for high-performance single-chip CMOS microprocessors designed from scratch. The design team consists of people from four main disciplines: computer architects who design the machine's microarchitecture, verification engineers who develop tests to validate the functionality of the various logical design representations of the chip, logic/circuit design engineers who do the design of the chip's logic and circuits, and layout designers who create the artwork for the masks.

Members of the design team work closely with system designers and compiler writers to identify design optimizations beyond the boundary of the chip. The CMOS technology used for chip fabrication is developed in-house and is strongly influenced by the microprocessor designers. Almost all of the CAD tools used for design, verification, and layout are developed in-house in response to project needs.

Table II
The Alpha AXP 21064 Microprocessor

Architecture	RISC
Transistor Count	1.7 million
Die Size	16.8 mm x 13.9 mm
Power Dissipation	30 watts
Clock Frequency	200 MHz

Table III
CMOS-4 Technology Description

Minimum Feature Size	0.75 micron
Minimum Effective Channel	0.5 micron
Gate Oxide Thickness	10.5 nm
NMOS/PMOS Threshold Voltage	0.5 V / -0.5 V
Power Supply	3.3 V
Substrate	P-epi with n-well
Metal 1 (Contacted Pitch)	2.25 micron
Metal 2 (Contacted Pitch)	2.625 micron
Metal 3 (Contacted Pitch)	7.5 micron

Table I
The NVAX Microprocessor

Architecture	CISC
Transistor Count	1.3 million
Die Size	16.2 mm x 14.6 mm
Power Dissipation	16 watts
Clock Frequency	100 MHz

III. Overview of the Design Process

An overview of the design flow is given in Fig. 1 [3]. At the start of the project, the architects write a flexible software model of the chip's proposed microarchitecture. This performance model is used to study the effects of various microarchitectural design decisions on overall CPU performance and to help refine the microarchitecture.

The chip floorplan is devised early in the design process to estimate and track the die size and to provide area-impact data for microarchitectural trade-offs. Once the chip-level floorplan is stable, the section design teams develop more detailed section floorplans and the chip's power dissipation is estimated.

As the chip specification stabilizes, a detailed register-transfer level (RTL) model is written, using a programming language such as C. Some of the ramifications of higher level assumptions become apparent for the first time and many logic/circuit feasibility studies are spawned. The RTL model is used to verify that the proposed microarchitecture executes the machine's instruction set correctly.

The RTL model is also used as a specification for logic/circuit design. Some less speed-critical circuits are synthesized from the RTL model but, due to the demanding timing requirements, the majority are custom designed by engineers.

When creating schematics, circuit designers use technology-specific rules of thumb to estimate the interconnect capacitance on signal lines, etc. SPICE is used for critical-path timing analysis. Layout for the schematics is then generated, accurate capacitances are extracted from the layout, and the circuits are resimulated to ensure that they still meet the specifications. A library of common circuit and layout structures, such as latches, is used to reduce the design and verification efforts.

Layout design is created on a sub-section basis. A typical sub-section consists of approximately ten related schematics. Each schematic and then each sub-section is checked for connectivity and correctness as the layout is completed. Sub-sections within a section are then integrated and verified (sections contain from five to fifteen sub-sections) before the complete chip composite of about 500 schematics is assembled.

Schematic verification and layout design are performed concurrently during much of the project. Although this overlap leads to design changes that increase the amount of layout rework, the chip development schedule would be extended significantly if schematic verification and layout design were done serially.

Once the layout design is completed, the final stages of layout verification ensure that the assembled chip layout meets the reliability and electrical integrity requirements for

global nodes such as power, clocks, and signals that cross section boundaries.

Although the simplified design flow shown in Fig. 1 is top-down, in reality the whole design process is highly iterative, often spanning many levels of abstraction simultaneously.

IV. Designing for High Performance

The performance of a microprocessor is inversely proportional to the product of clock cycles per instruction (CPI) and the machine cycle time. So the design team strives to maximize microprocessor performance by minimizing the *product* of CPI and cycle time rather than by minimizing each one independently.

A. Global Design Optimizations

Many opportunities to boost the performance of the design will be lost unless there is excellent communication among designers who work on different sections of the design as well as among designers working at different levels of abstraction. This communication is difficult to achieve because of the sheer complexity of the design and the large size and inter-disciplinary nature of the design team. In order to foster this communication, the offices of all the designers are located together. Additionally some senior designers from each discipline are responsible for identifying and resolving chip-wide technical issues.

In order to achieve a well balanced design, for example, tradeoffs must be made between the chip's microarchitecture

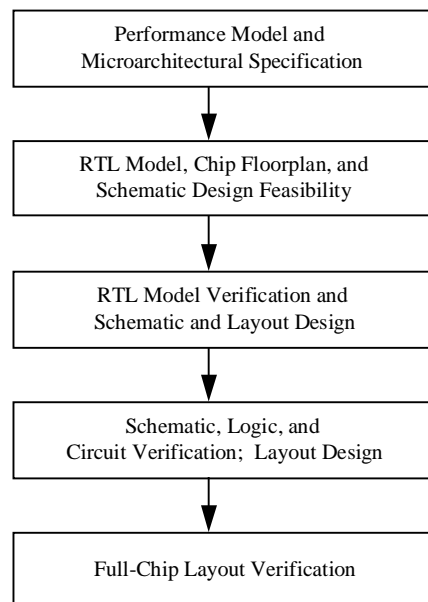


Fig. 1 Design Flow

and circuit implementation. For these tradeoffs to result in maximum performance, the architects must convey to the logic/circuit designers the *intent* of the design as well as the options for change at the microarchitectural level; and the logic/circuit designers must convey to the architects the freedoms and limitations of custom CMOS design along with the relevant characteristics (speed, power, area, etc.) of a variety of microarchitectural building blocks. Thus, through close cooperation, architects and logic/circuit designers can often eliminate microarchitectural bottlenecks and realize superior implementations.

B. High-Speed Custom Circuit Design

The schematic design style used is well suited to circuit design at the transistor level. Schematics are drawn with a restricted hierarchy: apart from the usual symbols for transistors and logic gates, the only commonly used global symbols are the standard latches. In addition, *in situ* local symbol definition is used extensively to facilitate efficient reuse of circuits on the defining schematic itself. Thus the structure of the circuit at the transistor-level can be readily discerned from the schematic. This helps designers to more easily visualize subtle interactions among transistors and facilitates optimization of the circuits. Moreover, the schematics can be used to convey to the layout designer the desired physical placement of transistors and wires.

A well established four-phase clocking scheme was used on the NVAX chip [4]. The Alpha AXP 21064 chip, on the other hand, used a two-phase single-wire clocking scheme to reduce the machine cycle time penalties that result from clock skew [5], [6]. With this scheme, clock rise and fall times must be tightly controlled across the chip to prevent latch race through. New CAD tools were developed to analyze the clock grid and check for clock methodology violations. Many new high-speed custom circuit techniques were developed to work with the new clocking scheme.

As experience was gained with the new circuit design style, the design team built up a body of knowledge about its appropriate use. New CAD tools were needed to recognize these types of circuits and check that they were being used appropriately on the chip. Design teams are less likely to make widespread use of new design styles on “must succeed” microprocessor projects without the support of a strong, responsive CAD group.

C. High-Performance Manufacturing Technology

Digital has been developing advanced CMOS technologies for more than a decade. The microprocessor design teams occupy the same building as the technology developers. These teams collaborate closely to ensure that the technology is well suited to high-performance microprocessor design. As a result, the 0.75 micron CMOS-4 technology has high-current transistors with thin gate oxides, a low resistance

third metal layer used mainly for clock and power distribution, and employs local interconnect to support a dense RAM cell for on-chip caches [7].

An overview of the interactions between the microprocessor design and the CMOS technology development is given in Fig. 2. This interaction affords an important opportunity for optimization across disciplines. The microprocessor design team helps the technology developers identify and focus on areas of high leverage. Early circuit design investigations aid in establishing the goals for transistor performance (saturation currents and threshold voltages) and the electrical properties of interconnect (resistivity, capacitance, and electromigration). Knowledge about die size, yield, RAM cell size, and number

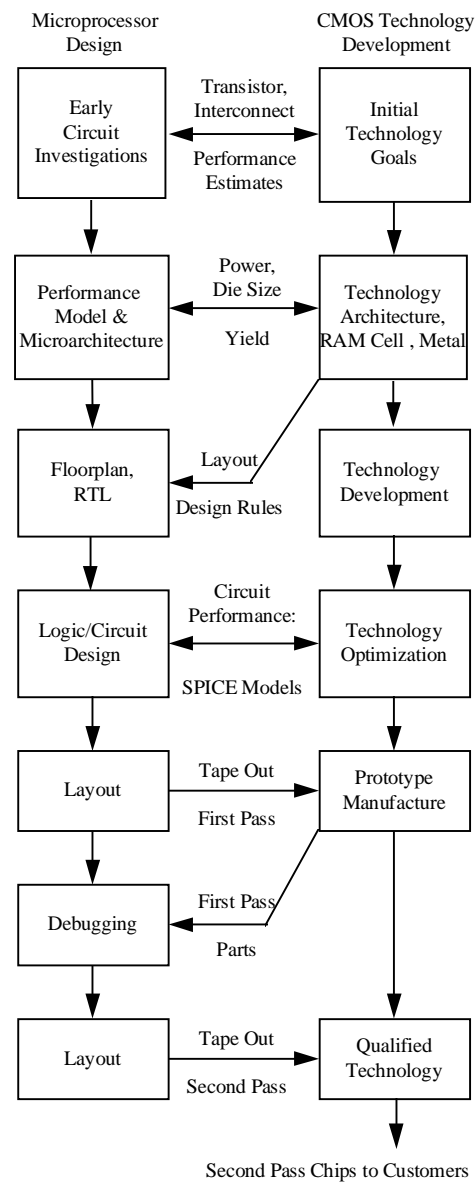


Fig. 2 Interactions between Microprocessor Design and CMOS Technology Development

of metal layers helps with microarchitectural design and floorplanning; and sets the stage for RTL modeling. Information from detailed circuit design is used to optimize the technology (transistor currents versus leakage, etc.). Factors that determine the chip reliability, such as electromigration and hot carrier stresses, are accounted for throughout. When the chip is taped out, the technology is ready to manufacture prototypes. While chip debug is in progress, the CMOS technology is being readied for volume manufacturing. Second-pass parts are subjected to rigorous functional verification and thorough reliability checks before being shipped to customers.

This concurrent engineering of the technology and the microprocessor has both advantages and challenges. The main advantages are that the microprocessor designers work with the CMOS technology developers to optimize the technology for high-performance microprocessors and that a state-of-the-art technology is ready when the chip reaches tape out. But because the technology is still under development while the chip is being designed, the characteristics of the technology may change: the SPICE models are subject to revision, predicting manufacturing variations is difficult, the reliability criteria for the technology are not known with precision, etc. All these uncertainties must be managed with little margin for error.

V. Minimizing Bugs and Schedule Delays

With microprocessor performance increasing at a rate of 50% per year, it is vital that the chip boot the operating system on first-pass silicon -- preferably at full-speed -- so that chip and system debugging can proceed rapidly and products can be brought to market quickly. After booting the operating system, the microprocessor can run more instructions in a few minutes than all the billions of instructions simulated on the chip models prior to tape out! The likelihood that disabling logic or circuit bugs elude detection prior to tape out is greatly reduced by fostering good design practice, accounting for manufacturing variations from the start, and subjecting the chip models to rigorous logical and electrical verification.

A. Good Design Practice

Methodology documents reinforce a necessary degree of design consistency across the team. These documents generally contain information on node-naming conventions, circuit sizing guidelines, global clock waveforms, latching and clocking rules, standard latches, guidelines on good circuit design practice, guidelines for layout designers, design for reliability, suggestions for CAD tool usage, etc. Although design consistency helps to reduce the risk of

circuit bugs, it should not be allowed to hamper creativity and innovation.

Design reviews are held to check circuit schematics for deviations from good circuit design practice. These reviews focus on circuit integrity rather than circuit speed or logical correctness.

B. Manufacturing Variations and Long-Term Reliability

The performance of a CMOS technology depends on the current-driving capability (I_{ds}) of the PMOS and NMOS transistors, as well as the capacitances of the interconnect, gates, sources, and drains that the transistors must charge and discharge. All of these parameters vary in manufacturing. In order to ensure that chips will function correctly and at the planned speed, chip designers use SPICE models (fast, slow, typical, and mixed) to account for these manufacturing variations when the chip is being designed.

The fast SPICE (FF) models are used to predict the speed of the fastest chip, the maximum power dissipation, the transient current demands on the power supply, the maximum voltage drops in the on-chip power and ground routing, the worst-case current density (checked to ensure that the electromigration limits are not violated) in the power supply and signal lines, and the maximum rate at which the chip's signal pins will transition. Power supply current transients and signal-pin transition rates have important implications for the electrical design for chip packages. Both place limits on how much inductance can be tolerated in the leads. The maximum power dissipation has obvious thermal implications for package and heat sink design.

The slow SPICE (SS) models are used to predict the speed at which the slowest chips will run.

The typical SPICE (TT) models are used to check that the circuits on the chip will run at the desired speed when the manufacturing line is running at the center of its range. The bulk of the circuit design work is done using these typical SPICE models.

The correct operation of some CMOS circuits is particularly sensitive to the ratio of the PMOS to NMOS currents so mixed SPICE (FS and SF) models are used to determine the noise margin for circuits (DC circuit analysis) rather than to predict the speed of which circuits will run (AC analysis). Circuit designers use these DC simulations to determine the safe bounds for the ratios of transistor sizes in a variety of common circuit structures. These bounds are incorporated into the design methodology for the project, and CAD tools are used to search the circuit schematic database for structures that violate the methodology.

By using the SPICE models in the manner just outlined, designers verify that the circuits will function correctly in spite of variations in the manufacturing line.

Environmental effects, such as operating temperature and power supply variations, are also taken into account. In addition, reliability issues, such as electromigration and hot carrier wear out, are factored into the design methodology and compliance is verified by CAD tools.

C. Logical Verification

The RTL model is verified by running a combination of pseudorandom tests, standard architecture tests, and handcrafted implementation-specific tests. Many billions of clock cycles are simulated prior to first-pass tape out. When bug fixes or design improvements are made, flaws can inadvertently creep into the design. So pseudorandom and regression tests are run on the updated model before time-consuming schematic and layout changes are made. Design changes and issues are posted to an electronic bulletin board along with the implications for other parts of the design. This tracking procedure keeps the team members apprised of changes [8].

When all the schematics have been designed for a section of the chip, the transistor wirelists for the circuit schematics are automatically translated to a logically equivalent model. The logical integrity of this schematic-based model is then verified on a two-state, gate-level logic simulator. Before tape out, the schematic-based logical model of the whole chip is subjected to an extensive battery of tests totaling about 100 million CPU cycles. In some cases, the operating system is booted on this model [9].

The full-chip transistor wirelists are also compiled and simulated on an event-driven three-state switch-level simulator to ensure that the chip initializes correctly when the chip reset signal is asserted.

D. Electrical Verification

If the chip has a slow critical-path, it will run at less than the target cycle time. While this is not desirable on first-pass silicon, system debugging can proceed albeit at a lower clock frequency, and the slow path can be corrected in the second pass. However, if a race (e.g., path A must evaluate before path B to prevent incorrect logical operation) is lost, the chip will not function correctly at *any* frequency and debug is severely impeded. For this reason, great pains are taken to avoid races wherever possible. Because subtle, insidious races can be difficult for designers to find, CAD tools that can detect their presence are extremely useful.

A static timing verification tool is used to identify critical-paths that would otherwise remain undetected due to the complexity of the design. This tool works on the transistor wirelist and accepts capacitances and resistances extracted from layout. Because the timing verifier does not predict circuit delays accurately, all questionable paths are simulated

using SPICE. The slow paths are then redesigned to meet the target cycle time.

Design reviews are supplemented by an in-house CAD tool that identifies violations of a wide set of circuit design rules: noise margin, capacitive coupling, charge sharing, beta ratios, etc.

VI. The Role of Custom CAD Tools

High-performance microprocessor design projects place new demands on existing CAD tools and always generate a long wish-list of new tools. Support from an in-house CAD group that develops the specialized design and verification tools required is invaluable. Tools are sometimes developed or customized on very short notice. On chip design projects that are pioneering new design techniques, it is just not possible to anticipate all the CAD requirements in advance.

As already mentioned, design teams are less likely to make widespread use of new design styles on complex “must succeed” projects without the support of a strong, responsive CAD group.

VII. Trends

Very high-performance microprocessors will be designed to provide the vast amounts of computational power that will be needed to run human interface applications such as video compression and decompression, cursive handwriting recognition, speaker-independent continuous speech recognition, etc. As algorithms mature and standards emerge in these areas, instruction sets will be expanded and special-purpose hardware will be added to microprocessors so that these important applications can be run efficiently.

The problems that tomorrow's microprocessor designers must solve will make today's issues seem easy. As design teams grow in size, the work of more designers must be coordinated. Technology feature sizes will drop to 0.25 microns and below and clock frequencies will approach 0.5 GHz, forcing designers to deal with a broader range of electrical issues: circuits that rely on limited swing differential signaling, on-chip inductive effects, signals with RC delays that approach the clock period, multiple clock domains, severe electromigration restrictions, hot carrier robustness, etc. Out-of-order instruction execution will pose new logical verification challenges. Power consumption must be held in check as tens of millions of transistors are used to implement increasingly complex architectures.

VIII. Conclusions

Digital has produced microprocessors with industry-leading performance by optimizing across the design space, applying innovative high-speed custom circuit design, and

using high-performance CMOS technology for fabrication. By fostering quality design practices, accounting for manufacturing variations and long-term reliability concerns, and subjecting the chip models to rigorous logical and electrical verification prior to tape out, these microprocessors have booted the operating system at full-speed on first-pass silicon.

REFERENCES

- [1] R. Badeau et al., "A 100 MHz macropipelined CISC CMOS microprocessor," *1992 IEEE ISSCC Digest*, vol. 35, pp. 104-105, 255.
- [2] D. Dobberpuhl et al., "A 200 MHz 64b dual-issue CMOS microprocessor," *1992 IEEE ISSCC Digest*, vol. 35, pp. 106-107, 256.
- [3] D. R. Donchin et al., "The NVAX CPU chip: design challenges, methods, and CAD tools," *Digital Technical Journal*, vol. 4, no. 3, Summer 1992, pp. 24-37.
- [4] R. Badeau et al., "A 100-Mhz macropipelined VAX microprocessor," *IEEE J. Solid-State Circuits*, vol. 27, no. 11, Nov. 1992, pp. 1585-1598.
- [5] D. W. Dobberpuhl et al., "A 200 MHz 64b dual-issue CMOS microprocessor," *IEEE J. Solid-State Circuits*, vol. 27, no. 11, Nov. 1992, pp. 1555-1567.
- [6] J. Yuan and C. Svensson, "High-speed CMOS circuit techniques," *IEEE J. Solid-State Circuits*, vol. 24, no. 1, Feb. 1989, pp. 62-70.
- [7] B. Zetterlund, J.A. Farrell, and T. F. Fox, "Microprocessor performance and process complexity in CMOS technologies," *Digital Technical Journal*, vol. 4, no. 2, Spring 1992, pp. 12-24.
- [8] W. Anderson, "Logical verification of the NVAX CPU chip design," *IEEE ICCD Proceedings*, Oct. 1992, pp. 306-309.
- [9] V. Peng, D.R. Donchin, Y.T. Yen, "Design methodology and CAD tools for the NVAX microprocessor," *IEEE ICCD Proceedings*, Oct. 1992, pp. 310-313.