

High-Performance Microprocessor Design

Ujjwal Kumar, Abhishek Kumar, Viplove Kumar

Dronacharya College of Engineering, M. D. University, Rohtak, Haryana, India

Abstract—Three generations of Alpha microprocessors have been designed using a proven custom design methodology. The performance of these microprocessors was optimized by focusing on high-frequency design. The Alpha instruction set architecture facilitates high clock speed, and the chip organization for each generation was carefully chosen to meet critical paths. Digital has developed six generations of CMOS technology optimized for high-frequency design. Complex circuit styles were used extensively to meet aggressive cycle time goals. CAD tools were developed internally to support these designs. This paper discusses some of the technologies that have enabled Alpha microprocessors to achieve high performance.

Index Terms Alpha, CMOS digital integrated circuits, computer architecture, flip-flops, integrated circuit design, logic design, microprocessors.

I. INTRODUCTION

DIGITAL introduced the Alpha 21064 in 1992, the highest performance microprocessor in the industry at that time. Digital has delivered three generations of high performance

Alpha microprocessors through process advancements, architectural improvements, and aggressive circuit design techniques the last five years, the clock frequency of the Alpha microprocessor has increased from 150 to 600 MHz. The 21264, the third-generation Alpha, has been designed to operate at 600 MHz with improved performance over the 21164. It was designed to operate at 200 MHz in a 0.75- μ m n-well CMOS process, allowing for roughly 16 gate delays per cycle including latching. Power dissipation is 30 W from a 3.3-V power supply at 200 MHz. The die measures 2.3 cm, and contains 1.68 million transistors, half of which are dedicated to non-cache logic. The second-generation Alpha microprocessor, the 21164 is fabricated in a 0.5- μ m n-well CMOS process. It was designed to operate at 300 MHz using a 3.3-V supply, and it dissipates 50 W. The number of gate delays per cycle was reduced from 16 to 14 on this design to provide an additional 10% reduction in cycle time beyond process scaling. The die is roughly 3.0 cm and contains 9.3 million total transistors. The no cache transistor count is tripled from the previous

generation design to 2.5 million to operate at 300 MHz; migration of this design to a 0.35- μ m process has increased the operating frequency to 600 MHz. The 21264 the third-generation Alpha microprocessor.

It is designed in a 0.35- μ m n-well CMOS process, and is targeted to operate at 600 MHz. The number of gate delays per cycle has been further reduced to 12, again providing an additional 10% reduction in cycle time relative to the previous design. A nominal supply voltage of 2.2 V is used to limit power dissipation to an estimated 72 W, but the design and process can operate reliably up to 2.5 V. The die is 3.1 cm and contains 15.2 million transistors. The non-cache transistor count is more than double that of the 21164.

To achieve high performance without impacting time-to market a careful balance among micro architectural features, process complexity, and circuit design style was required on each of these microprocessors. The use of high-performance circuit design techniques required the development of many custom CAD tools, and added to the complexity of the circuit verification task.

(I) ARCHITECTURE

The Alpha instruction set architecture is a true 64-bit Load/store RISC architecture designed with emphasis on high clock speed and multiple instruction issue. Fixed-length instructions, minimal instruction ordering constraints, and 64-bit data manipulation allow for straightforward instruction decode and clean micro architectural design. The architecture does not contain condition codes, branch delay slots, adaptations from existing 32-bit architectures, and other bits of architectural history that can add complexity. The chip organization for each generation was carefully chosen to gain the most advantage from micro architectural features while maintaining the ability to meet critical circuit paths.

The 21064 is a fully pipelined in-order execution machine capable of issuing two instructions per clock cycle. It contains one pipelined integer execution unit and one pipelined floating-point execution unit. Integer instruction latency is one or two cycles, except for multiplies which are not

pipelined. Floating-point instruction latency is six cycles for all instructions except for divides. The chip includes an 8-kB instruction cache and an 8-kB data cache. The emphasis of this design was to gain performance through clock rate while keeping the architecture relatively simple. Subsequent designs rely more heavily on aggressive architectural enhancements to further increase performance. The quad-issue, in order execution implementation of the

21164 was more complex than the 21064, but simpler than an out-of-order execution implementation. It contains two pipelined integer execution units and two pipelined floating-point execution units. The first-level cache was changed to nonblocking. A second-level 96-kB unified and cache was added on-chip to improve memory latency without adding excessive complexity. Integer latency was reduced to one cycle for all instructions, and was roughly halved for all MUL instructions. The floating-point unit contains separate add and multiply pipelines, each with a four-cycle latency. Floating-point divide latency is reduced by 50%. The trend of increased architectural complexity continues with Digital's latest Alpha microprocessor. The 21264 gains significant performance from six-way-issue and out-of-order execution. It contains four integer execution units and two floating-point execution units. The size of the instruction and data caches was increased from 8 to 64 kB, eliminating the need for an on-chip cache. Integer multiply latency was reduced and full pipelining improved throughout. The floating-point latency remained at four cycles, but the divide latency was reduced by another 50%. In addition, the ISA was extended to include square root and to support multimedia instructions.

Despite the added architectural complexity, clock frequencies have continued to improve due to circuit design enhancements and advances in process technology.

B. TECHNOLOGY

Digital Semiconductor has developed six generations of

CMOS process technology, with a new technology for each major microprocessor design. The microprocessor design occurs in parallel with the development of the manufacturing process. Therefore, close cooperation is required between the process development and microprocessor

design teams to perform this concurrent design and ensure optimum chip performance. The processes were optimized for high-frequency microprocessor design. In particular, emphasis is placed on low τ 's and very short τ 's which increase drive current at the cost of higher leakage. Close interaction between the circuit design team and the process development team also results in the following benefits.

- 1) Early process information and timely updates of technology parameters are provided to the design teams, allowing circuit design to start before the process is fully defined.
- 2) Early design work provides valuable feedback to the process team to ensure that target process performance is met.
- 3) Major process features such as number of interconnect layers, interconnect pitch, and device characteristics are managed in the context of the overall chip design.
- 4) The design of critical structures such as RAM arrays and data paths can be optimized through process and circuit design.
- 5) Scaling issues for future process shrinks may be uncovered.

A. Definition of Design Rules

One of the key areas where close collaboration is required between design and process development teams is the definition of layout design rules. Aggressive design rules can result in increased circuit density, and can potentially improve overall chip performance. However, design rules that are too aggressive will complicate manufacturing, and may impact yield. On the other hand, slack design rules may result in increased die size, resulting in increased distances between critical structures. This increased distance results in higher capacitance, larger routing delays, and lower chip performance. Often, the process team can be more aggressive if limits are placed not only on the minimum widths and spaces of structures, but also on the maximum widths and spaces. The 21264 implements metal fillers to limit the maximum spacing between adjacent lines. The fill metal is automatically placed in the design and tied to or . For large areas of fill metal, stress relief holes are automatically placed in the fill pattern. Metal fill may increase the capacitance of nearby signal lines, but they also result in improved interlayer dielectric uniformity. The improved uniformity allows the process to be targeted more aggressively.

3) CAD TOOLS AND VERIFICATION

Custom circuit techniques allow designers to build very high-speed circuits. However, the use of these custom circuits requires design expertise and detailed postlayout electrical verification. Commercially available EDA design systems and point CAD tools have very limited support for these techniques. Therefore, Digital developed an extensive suite of in-house CAD tools to facilitate the design of custom VLSI microprocessors. Internally developed CAD tools are used extensively in all phases of microprocessor design, from initial performance evaluation, through circuit implementation and final design verification. The internally developed CAD suite includes tools for schematic and layout entry, two-state and three-state logic simulation, RTL versus schematic equivalence checking, static timing analysis and race verification, parameter and netlist extraction, and electrical analysis and verification.

A. Electrical Verification

Electrical verification covers all circuit issues that are not related to logic functionality such as timing behaviour, electrical hazards, and reliability. Electrical hazards result from noise sources interfering with the logical functions of the chip, and include charge sharing, interconnect capacitive and inductive coupling, power supply IR noise, and noise-induced minority carrier charge injection. Reliability checks include metal and via electro migration, transistor hot-carrier damage, ESD, and latch-up failure. The primary goal of the electrical verification tools is to verify that all circuits conform to the project design methodology.

The design methodology defines an acceptable set of circuit styles and sizing rules that, when followed, ensures functionality with minimal analysis. The methodology also forces a consistent design style to be used project-wide, which has the added benefit of simplifying CAD design. However, occasionally, there is a need to design circuits outside the methodology to meet performance or area goals. In these instances, additional manual verification is required to ensure functionality. Detailed analysis requires complex models with many process and circuit variables. Many checks are complex and hard to define as procedures that can be completely automated.

Exact chip-wide analysis is impractical; instead, the tools perform design filtering. The tools filter out all circuits that can easily be validated while identifying the small number of circuits that may

have problems and require additional analysis. This approach focuses design attention on potential problem areas, and therefore helps improve overall design efficiency. The CAD tools perform over 100 unique electrical checks. Some of the major areas of focus are circuit topology violations, dynamic node checks, including charge sharing, IR noise, injection, and leaker usage, interconnect coupling, noise margin checks, writeability checks, latch checks, beta ratio checks, gate fan-in and fan-out restrictions, transistor size and stack height limitations, max/min edge rates and delays, and power consumption. Some of the checks are applied to all circuit styles, while other checks are required for specific circuit types. The CAD tools require a large amount of design information to perform these checks, including electrical parameter extraction from layout, device electrical characteristics, relative circuit locations, timing information, and transistor connectivity.

B. Functional and Logical Verification

The functional complexity and time-to-market pressures of microprocessor design necessitated the development of an extensive functional verification strategy covering all phases of the design process from functional definition through manufacturing tests. During the micro architectural design phase of chip development, a two-state RTL behavioural model is the primary verification vehicle. This model provides a balance between design detail and simulation speed. The model is combined with abstract behavioural models of the other system components to verify correct operation of the processor in the system environment. Once logic design is complete, a two state gate-level simulation model is extracted from the circuit schematics. This model is used to ensure that the schematics match the RTL model. Finally, to verify correct initialization of the circuits at power-up a three-state switch-level model is extracted from the circuit schematics. A wide variety of simulation stimuli is used to verify the design, including hand-coded test patterns and randomly generated test patterns. Coverage analysis guides the verification process. Many of the manufacturing test patterns are derived from the simulation stimuli. Fault simulation is used to direct test enhancement.

CONCLUSION

This paper has reviewed Digital's approach to high-performance microprocessor design. Three generations of Alpha microprocessors have been

designed and optimized for performance by focusing on high-speed design using fully custom circuit design techniques, incorporating state-of-the-art architectural features, and utilizing high-performance CMOS technology for fabrication. An extensive suite of in-house

CAD tools are used to analyse complex electrical and logical behaviour to ensure functionality. All three microprocessors booted multiple operating systems using first-pass silicon, validating the design and verification methodologies.

REFERENCES

- [1] D. Dobberpuhl *et al.*, "A 200 MHz 64 b dual-issue CMOS microprocessor," *IEEE J. Solid-State Circuits*, vol. 27, pp. 106–107, Nov. 1992.
- [2] W. Bowhill *et al.*, "A 300 MHz 64 b quad-issue CMOS microprocessor," in *ISSCC Dig. Tech. Papers*, Feb. 1995, pp. 182–183.
- [3] B. Gieseke *et al.*, "A 600 MHz superscalar RISC microprocessor with out-of-order execution," in *ISSCC Dig. Tech. Papers*, Feb. 1997, pp. 176–177.
- [4] R. Sites and R. Witek, *Alpha AXP Architecture Reference Manual*, 2nd ed. Boston, MA: Digital, 1995.
- [5] J. Edmondson *et al.*, "Superscalar instruction execution in the 21164 Alpha microprocessor," *IEEE Micro*, vol. 15, pp. 33–43, Apr. 1995.
- [6] J. Kowaleski *et al.*, "A dual-execution pipelined floating-point CMOS processor," in *ISSCC Dig. Tech. Papers*, Feb. 1996, pp. 358–359.
- [7] B. Benschneider *et al.*, "A 300-MHz 64-b quad-issue CMOS RISC microprocessor," *IEEE J. Solid-State Circuits*, vol. 30, pp. 1203–1214, Nov. 1995.
- [8] L. Heller and W. Griffin, "Cascode voltage switch logic: A differential CMOS logic family," in *ISSCC Dig. Tech. Papers*, Feb. 1984, pp. 16–17.
- [9] D. Priore, "Inductance on silicon for sub-micron CMOS VLSI," in *1993 Symp. VLSI Circuits, Dig. Tech. Papers*, May 1993, pp. 17–18.
- [10] P. Gronowski and W. Bowhill, "Dynamic logic and latches: Practical implementation methods and circuit examples used on the ALPHA 21164," in *1996 Symp. VLSI Circuits—Proc. VLSI circuits Workshop*.
- [11] D. Dobberpuhl *et al.*, "A 200-MHz 64-bit dual-issue CMOS microprocessor," *Digital Tech. J.*, vol. 4, no. 4, pp. 35–50, 1992.