

Digital's DECchip 21066: The First Cost-focused Alpha AXP Chip

by

Dina L. McKinney, Masooma Bhaiwala, Kwong-Tak A. Chui,
Christopher L. Houghton, James R. Mullens, Daniel L. Leibholz,
Sanjay J. Patel, Delvan A. Ramey, and Mark B. Rosenbluth

ABSTRACT

The DECchip 21066 microprocessor is the first Alpha AXP microprocessor to target cost-focused system applications and the second in a family of chips to implement the Alpha AXP architecture. The chip is a 0.675-micrometer (μm), CMOS-based, superscalar, superpipelined processor that uses dual instruction issue. It incorporates a high level of system integration to provide best-in-class system performance for low-cost system applications. The DECchip 21066 microprocessor integrates on-chip, fully pipelined, integer and floating-point processors, a high-bandwidth memory controller, an industry-standard PCI I/O controller, graphics-assisting hardware, internal instruction and data caches, and an external cache controller. Cost-saving packaging techniques and an on-chip, analog phase-locked loop enable the chip to meet the cost demands of personal computers and desktop systems. This paper discusses the trade-offs and results of the design, verification, and implementation of the DECchip 21066 microprocessor.

INTRODUCTION

The DECchip 21066 microprocessor is the first cost-focused implementation of the Alpha AXP architecture.[1] The chip integrates system functions that are normally found in a microprocessor chip set with a high-performance, superscalar microprocessor core to deliver high-end personal computer (PC) performance and low overall system cost. The DECchip 21066 device is also the first microprocessor to integrate a Peripheral Component Interconnect (PCI) local bus controller.[2] This open industry-standard I/O bus allows direct connection to high-performance peripheral components supplied by many vendors. The bus also allows direct connection to commodity Industry Standard Architecture (ISA) or Extended ISA (EISA)-based PC peripherals through a simple bridge chip. The DECchip 21066 microprocessor integrates a high-bandwidth memory controller that directly sequences an external secondary cache and main system memory. In addition, an on-chip, phase-locked loop (PLL) multiplies a low-frequency reference clock to a high-frequency CPU core clock, thus eliminating the need for a high-frequency board oscillator. The combination of cost-saving functional

integration and features that ease system design reduces the overall CPU subsystem cost and shortens the time-to-market of high-volume Alpha AXP PCs.

CMOS TECHNOLOGY

The first chip to fully implement the Alpha AXP architecture, the DECchip 21064 microprocessor was designed in a 0.75-micrometer (μm) (drawn) complementary metal-oxide semiconductor (CMOS) process.[3,4] The physical implementation of the DECchip 21066 microprocessor was achieved through the use of a 10 percent CMOS process shrink, which reduced the minimum feature size from 0.75 μm (drawn) to 0.675 μm . This reduction in the feature size enabled the DECchip 21064 integer unit, floating-point unit, and caches to be combined with a new memory controller, PCI I/O controller, and PLL on a die with approximately the same area as the original DECchip 21064 device. The maximum core clock speed of the DECchip 21066 microprocessor is specified as 166 megahertz (MHz). This speed allows the internal PLL to multiply a 33-MHz reference clock (the same reference clock as on the PCI bus) by five to generate a clock for the CPU core. This cycle time was set to be slightly less aggressive than the 200-MHz maximum core clock speed for the DECchip 21064 microprocessor to allow a greater number of yielding parts and thus lower part cost.

TARGET MARKET

One target market for the DECchip 21066 microprocessor is Alpha AXP PCs running the Microsoft Windows NT operating system. The bulk of the application base for this operating system is provided by vendors who target the Microsoft Windows operating system running on PCs based on Intelx86 microprocessors. Software vendors have substantial expertise with this architecture and, until recently, the advanced software tool base for Windows development has been targeted exclusively at the Intelx86 architecture. To provide compelling motivation for application developers to supply portable Windows NT applications and for customers to adopt a new architecture, Alpha AXP PCs must be price competitive with Intelx86 PCs and must deliver substantially higher performance.

The availability of the Windows NT operating system and high-performance PCs based on Intel's Pentium microprocessor dictated that a set of competitive products be available in early 1994.[5] Furthermore, the performance of Pentium-based PCs set the lower bound of acceptable performance for PCs based on the Alpha AXP technology. These schedule and performance goals were met by implementing the DECchip 21066 microprocessor as a high-integration design variant of the DECchip 21064 microprocessor. This strategy permitted a design cycle of only 10 months.

TARGET PERFORMANCE

The DECchip 21066 microprocessor provides a CPU subsystem that is comparable in cost to a 66-MHz Intel486 DX2-based PC with Pentium-class performance.[6] Because of its excellent price/performance ratio, the DECchip 21066 microprocessor is attractive to low-end workstation products running the DEC OSF/1 AXP and OpenVMS AXP operating systems, in addition to Windows NT platforms.

A performance-limiting factor for many applications is the bandwidth to the second-level cache. The DECchip 21064 microprocessor maximizes bandwidth with a 128-bit data bus, which allows two secondary cache read sequences to fill a line in the primary cache. DECchip 21066 engineers chose to implement a 64-bit data bus. This implementation resulted in a less-expensive package, a smaller die, and lower system cost, while reducing by only 20 percent the performance that would have been achievable with a 128-bit bus. The fast core clock and close proximity of the secondary cache to the core still enable the DECchip 21066 microprocessor to deliver high performance.

INTERNAL COMPONENTS

The block diagram in Figure 1 shows the major DECchip 21066 microprocessor components. Fetched instructions from an on-chip 8-kilobyte, direct-mapped instruction cache are dual-issued to the integer, floating-point, and load-and-store (addressing box) units. The on-chip, 8-kilobyte, direct-mapped data cache initially services memory loads. Stores are written through the data cache and absorbed into a write buffer. The memory controller or PCI I/O controller handles all references that pass through the first level of caches.

[Figure 1 (Major Components of the DECchip 21066 Microprocessor) is not available in ASCII format.]

The memory controller handles memory-bound traffic. The controller first probes a direct-mapped, write-back, write-allocate cache and then sequences main memory to fill the caches, if necessary. The PCI I/O controller handles I/O-bound traffic and performs programmed I/O read and write operations on behalf of the CPU. Direct memory access (DMA) traffic from the PCI is handled by the PCI controller in concert with the memory controller. DMA read and write operations are not allocated in the secondary cache. The memory and PCI interfaces were designed specifically for uniprocessor systems and do not support multiprocessor implementations.

SYSTEM APPLICATION

Figure 2 shows a sample system block diagram using the DECchip 21066 microprocessor. In this configuration, the memory controller sequences both the static random-access memory (SRAM) secondary cache and the dynamic random-access memory (DRAM) main memory. The secondary cache comprises tag and data SRAMs with identical read and write access times. System software selects specific random-access memory (RAM) timing values in increments of the number of core clock cycles.

[Figure 2 (Sample System Block Diagram) is not available in ASCII format.]

The design supports four separate banks of DRAM, each of which can be timed independently. This feature adds flexibility in memory organization and upgrading. One of the four banks can be configured with video RAMs (VRAMs) for low-cost graphics applications. The memory controller directly sequences the VRAMs and supports simple graphics operations, such as stippling and bit-focused writes.

Even though the DECchip 21066 microprocessor operates at 3.3 volts (V), all pads directly drive transistor-transistor logic (TTL) levels and can be driven by either 3.3-V or 5-V components. This eliminates the need for either special voltage conversion buffers or 3.3-V memory parts, which can add to system cost or affect performance.

The PCI bus is a high-bandwidth bus with a number of attractive features. In addition to its ability to handle DMA and programmed I/O, the PCI bus allows for special configuration cycles, extendibility to 64 bits, 3.3-V or 5-V components, and faster timing. The base implementation of the PCI bus supports 32 bits of multiplexed address and data with a bus clock of 33 MHz, yielding a maximum burst bandwidth of 132 megabytes per second (MB/s).

The PCI bus is directly driven by the microprocessor. In Figure 2, some high-speed peripheral devices, such as graphics and small computer systems interface (SCSI) adapters, connect directly to the PCI bus. An ISA bridge chip allows access to lower-cost, slower peripheral devices, such as modems and floppy disk drives.

MEMORY CONTROLLER FEATURES

A primary goal of the DECchip 21066 project was to simplify the system design to enable customers to build products with minimal engineering investment. Consequently, the memory controller directly connects to industry-standard DRAMs and single in-line memory modules (SIMMs), with buffers required only for electrical drive. To span the wide range of system applications from embedded controllers to midrange workstations, the timing of address, data, and control signals is highly programmable. This feature supports varying speeds and physical organization of the

DRAMs and clock speeds of the DECchip 21066 microprocessor itself.

The memory controller supports from one to four banks of memory, which allows main memory to range from 2M bytes to 512M bytes. The system may include a secondary cache of 64K bytes to 2M bytes. This optional cache also connects directly to the microprocessor, and the timing of control signals is programmable in increments of chip cycles.

Several simple, graphics-assisting features allow systems that need to save the cost of dedicated graphics control hardware to control the system frame buffer in software. The decision to include graphics features was based on the ability to provide significant acceleration of key functions for minimal hardware cost in the memory controller.

DYNAMIC MEMORY INTERFACE

The DRAM data interface consists of a 64-bit-wide bidirectional bus that is shared with the secondary cache. The interface provides two control signals for an optional memory data bus transceiver that may be required because of bus loading. All main memory read operations return a full 64 bits of data. Main memory write operations normally consist of a full quadword (64-bit) write. Write requests for less than a quadword of data automatically result in a read-modify-write sequence, unless the software has enabled masked writes on the memory bank. During masked writes, the error correction code (ECC) pins provide a byte mask that is externally combined with a column address strobe (CAS) signal to form a byte mask strobe for the array. ECC checking is not supported on memory banks that have masked writes enabled.

An optional 8-bit bidirectional bus provides quadword ECC checking on the main memory and the secondary cache. Each bank of main memory may be protected by ECC at the cost of including 8 extra DRAM bits per bank. The ECC used by the DECchip 21066 microprocessor allows correction of single-bit errors and detects double-bit and 4-bit nibble errors. The memory controller automatically corrects the single-bit memory read errors before data is passed to the read requester. When the ECC identifies a double-bit or a 4-bit nibble error, the memory controller does not attempt to correct it. When any ECC error is detected, the memory controller stores the error condition along with the address of the error. System software must scrub the error from physical memory.

SECONDARY CACHE

The DECchip 21066 microprocessor supports a secondary cache designed with industry-standard asynchronous SRAMs. The cache is

direct mapped with 8-byte blocks and uses a write-back policy. Designers chose a secondary cache block size that is smaller than that of the 32-byte primary internal caches to simplify the allocation of an external cache block during write operations. With the smaller block size, it is not necessary to fetch the missing quadwords of the block from DRAM when less than a full internal cache block is written, as would be the case with a larger block size. In addition, the smaller block size provides one dirty bit per quadword, so that the resolution of dirty (not yet written back) data is one quadword rather than four, thus reducing the number of write-back operations to the DRAM.

The DECchip 21066 microprocessor writes the cache tag into the SRAMs on cache allocation operations and receives and compares the tag internally on cache lookup operations. Read-fill operations of the internal caches are pipelined as the microprocessor drives the next cache index to the SRAMs prior to determining the hit or miss result of the current lookup. A cache hit, which is the more common occurrence, causes SRAM access to achieve maximum speed. In the case of a miss, the microprocessor re-drives the miss address and reads the DRAMs.

The chip divides the memory space into cacheable and noncacheable regions based on address. To save time, accesses to the noncacheable region skip the cache-probe step and immediately access the DRAM. This feature may be used, for example, to optimize accesses to a frame buffer, which typically is not cached.

The cache tag field, including the dirty bit, may optionally be protected by a single parity bit. In systems with write-through caches, tag parity errors are not necessarily fatal since the correct data can be fetched from DRAM. With a write-back external cache, such as the one used by the DECchip 21066 microprocessor, the bad parity may be on a dirty location. In this case, bad parity is a fatal error, since the copy in DRAM is no longer current.

GRAPHICS-ASSISTING FUNCTIONS

The graphics-assisting logic of the DECchip 21066 microprocessor provides some basic hardware enhancements to improve frame buffer performance over a standard, simple frame buffer system. The graphics-data-path--assist logic is targeted at reducing the number of instructions executed during inner loop graphics operations. By reducing the number of inner loop instructions, the microprocessor offers an improved graphics performance while keeping the overall graphics subsystem cost in line by using a simple frame buffer design rather than a more expensive graphics accelerator. The graphics-assist logic provides the opportunity to design a low-cost, entry-level graphics option and may be used in conjunction with a high-performance graphics accelerator.

The hardware graphics-assist logic consists of a direct interface to VRAM parts and the ability to perform some simple graphics-oriented data manipulations. To facilitate the use of VRAMs, the memory controller supports both full- and split-shift register load operations. External video monitor control logic generates monitor timing and signals the DECchip 21066 memory controller when VRAM shift register loads are required. An internal linear address generator keeps track of the video display's refresh address.

The features of the graphics data path provide the ability to perform simple frame buffer--style data operations, e.g., transparent stipple writes, plane masked writes, and byte writes (with minimal external logic). In transparent stipple mode, the memory controller can conditionally substitute the foreground data for the frame buffer data on memory writes.

Pixel depth may range from 1 bit to 32 bits. A special feature of the graphics-assist logic is the ability to perform graphics data manipulation operations on both VRAM and standard DRAM memory banks. The DECchip microprocessor emulates graphics operations targeted at memory banks without VRAMs by means of a read-modify-write sequence. This method allows the same graphics firmware to operate on either VRAM or DRAM memory banks and allows DRAM to be used as additional off-screen memory.

A low-cost video option board may be built around the microprocessor by adding a bank of VRAMs and a video controller. Typical video controller logic consists of a video timing generator, a RAMDAC (video digital-to-analog [D/A] converter with color mapping), a hardware cursor, and a few other simple video data path components.

PCI BUS INTERFACE

The DECchip 21066 microprocessor is the industry's first microprocessor to implement an interface that connects directly to the PCI bus. This PCI bus interface, called the I/O control (IOC), runs asynchronously with the rest of the microprocessor's core logic. Asynchronous design was chosen to enable optimal system performance by setting the chip's core clock to its maximum frequency without being limited to an integer multiple of the PCI clock frequency.

The IOC can be viewed as two separate controllers: one for DMA and the other for core requests. The DMA controller handles all peripheral-initiated DMA operations to the system memory; the core controller handles the loads and stores to either the PCI devices or the IOC internal registers. Since the PCI bus uses a 32-bit address and the DECchip 21066 microprocessor uses a 34-bit address, a PCI address to which the IOC responds must be translated to an equivalent address in the microprocessor's address space. The IOC provides two types of address-translation

mechanisms, direct and indirect (also called scatter-gather).

During a direct-mapped address translation, the lower bits of the PCI address are concatenated with a page address that is stored in an IOC register to form a 34-bit translated address. In an indirect-mapped address translation, certain bits of the PCI address are concatenated with a translated base address that is stored in an IOC register to form a 34-bit address. This address indexes into a scatter-gather map table in system memory where the page address resides. The page address is then concatenated with the lower bits of the PCI address to form the translated address. The IOC contains two programmable windows, each of which can be programmed to respond with direct or indirect address translation. To facilitate fast translation for the indirect address translation, the IOC contains an eight-entry, fully associative, translation look-aside buffer. To simplify the design, DMA burst length is limited to occur within an 8K-byte page boundary. Bursts that extend beyond a page boundary are broken into separate transfers.

CPU addresses are translated to an equivalent address in the PCI address space through one of two types of address translation, sparse or dense, depending on the target region of the address. For sparse-space access, the lower 32 bits of the CPU address are shifted right by 5 bits to generate a 27-bit address. This address, concatenated with a 5-bit field from a register in the IOC, maps the 32-bit PCI address. Transfers of up to 8 bytes can be completed in this address space. For dense-space addressing, the lower 32 bits of the CPU address are directly mapped to the PCI address. Only unmasked operations are allowed, and up to 32 bytes of write data and 8 bytes of read data can be transferred in this address space. The IOC improves write bandwidth by buffering two 32-byte writes from the DECchip 21066 core.

The priorities of the IOC design were to support peak PCI bandwidth and at the same time to meet the tight project schedule and limited die area constraint. A 32-byte data queue buffers DMA data. A larger queue would require more die area; a smaller queue might stall the DMA bursts due to the synchronization delay and system memory access latency. The IOC and memory controller are connected by a 64-bit data bus. Since the IOC PCI bus supports only 32 bits, two 32-bit data words are transferred at a time to minimize read-modify-write operations that would otherwise be issued by the memory controller. The IOC improves the DMA read bandwidth by prefetching up to 32 bytes of data.

A hardware semaphore provides the asynchronous handshaking between the core clock and PCI clock domains. Approximately two months were dedicated to the logic and physical implementation of the hardware semaphore and the manual verification of the asynchronous communication and signal timing. This effort was required because the existing computer-aided design (CAD) tools did not have a formal verification methodology that would sufficiently validate the asynchronous aspects of the IOC

architectural design and the physical implementation.

Developing the test strategy was a challenge for two reasons: (1) the presence of asynchronous sections and (2) the need to guarantee that the semaphore would yield predictable results on a production tester operating at the maximum design frequency (166 MHz for the core and 33 MHz for the PCI bus). Even when the two clocks are running synchronously, they must be controlled and aligned such that the hardware semaphore guarantees reliable and predictable results. Each core clock phase is only 3 nanoseconds at 166 MHz. Taking into account the logic delay and setup time required by the hardware semaphore, this clock skew requirement is essentially impossible to achieve. The inaccuracy of the tester and the test hardware further complicates the problem. The solution adopted was to incorporate extra logic to reduce the sample rate of the hardware semaphore that runs on the core clock without actually slowing down the clock itself. With the extra logic, we were able to demonstrate the ability to test the IOC and the DECchip 21066 microprocessor at the targeted clock frequencies.

LOGIC DESIGN VERIFICATION PROCESS

A pseudorandom design exerciser is the main verification tool for the DECchip 21066 microprocessor. The exerciser takes input from Alpha AXP code streams and PCI bus commands. A template-based text manipulator generates test patterns from the templates written by verification engineers to target specific sections of the microprocessor design. The text manipulator provides the primary source of randomness. It selects templates based on probabilities to create the test pattern. Each test pattern runs on a design model. We use two different design models, a register transfer level (RTL) model written in the C programming language and a structural model generated from the circuit schematics. Finally, the same test pattern runs on the instruction set processor (ISP) level reference model to determine the correctness of the design model.[7] Figure 3 illustrates the logic design exerciser process flow.

[Figure 3 (Logic Design Exerciser Process) is not available in ASCII format.]

To determine if the design model operates correctly, the exerciser performs checks at three points: (1) during the execution of the design model, (2) during the execution of the reference model, and (3) after both models complete. Verification software in the design model performs additional checks for illegal conditions, such as multiple bus drivers, protocol violations, and invalid states. This code is kept to a minimum and covers design aspects not easily checked by comparing state information.

Previous verification projects have used similar random test

methods; however, they did not have to ensure correct operation of asynchronous I/O transactions.[7,8] To deal with this problem, the DECchip 21066 verification team enhanced the reference model to evaluate CPU and PCI transactions. The reference model receives the same stimuli as the design model plus additional I/O bus and event information. While executing the test pattern, the model verifies that the I/O information is consistent.

By periodically comparing functional states in a test pattern, we circumvented adding detailed timing information to the reference model. The model needs only valid state for comparison at synchronized capture points rather than at a clock phase boundary. In essence, we traded off being able to determine the timing correctness of our design for a simpler, faster reference model. We augmented the design model with verification code to check critical timings and performance features. Detailed static and circuit timing analysis was accomplished by other CAD tools.

SHARED MEMORY

Several features of the DECchip 21066 microprocessor require the exercisers to employ a shared memory structure to properly test the chip.

- o The microprocessor invalidates an internal data cache line on a DMA transfer to that cache line address.
- o An exclusive PCI access unconditionally clears the processor's lock flag. A nonexclusive DMA write also clears the lock flag if the DMA address is within the locked 32-byte range.
- o The DECchip 21066 secondary cache does not allocate on a DMA operation. Hence, the only way to test DMA sequences that hit in the secondary cache is to first access those addresses from the CPU.

The prime challenge to implementing a shared memory is that simultaneous accesses to the same address location by the CPU and DMA could result in unpredictable behavior. To circumvent this problem, we developed a model for software access to shared memory and added synchronization hooks into the reference model and the I/O trace file. The shared memory access model and synchronization hooks extended the exerciser to fully verify the DECchip 21066 shared memory capability.

The shared memory implementation proved useful in discovering two bugs in the shared memory functionality. Both bugs were tricky to generate and would have been extremely difficult to find without the shared memory capabilities of a random exerciser.

The random verification methodology enabled us to meet the aggressive 10-month schedule. By separating functional

correctness and timing correctness, we were able to quickly implement and stabilize the exerciser environment. The exerciser can easily be ported to future generations of the DECchip 21066 microprocessor; we have already used the tool to verify the second-pass design.

PLL DESIGN ISSUES

To reduce the module-level cost of a system based on the DECchip 21066 microprocessor, the chip includes a frequency synthesis system based on a phase-locked loop design.[9-18] The module frequency reference is a standard crystal oscillator, rather than an expensive, surface acoustic wave (SAW) oscillator. This design reduces radio frequency emissions, making qualification of the design easier and reducing system enclosure cost. Different speed binnings of the chip can run at maximum performance levels in the same basic module design. This result is possible because of the asynchronous, fixed clock rate of the PCI bus, the on-chip caches, and the programmable timing of the memory interface.

Incorporating a PLL subsystem in a digital CPU chip design requires that several problems be addressed. For example, increased noise sensitivity, which is a concern for digital designs, is even more critical for the analog circuits in a PLL. Often, simple guidelines for digital designs can ensure reasonable noise immunity; however, analog circuits frequently require detailed SPICE simulations that incorporate package- and chip-switching noise models.[19] Designers must consider making changes to otherwise fully functional circuits to ensure noise immunity. These changes can range from appropriate decoupling of critical nodes with high-quality, on-chip capacitors to consideration of major circuit changes or additions. The overall solution must be some combination of the following three alternatives: (1) careful circuit design based on noise environment simulations, (2) decoupling that is sufficient to reduce noise to tolerable levels, and (3) a better package that reduces loop inductances to reduce switching noise on the internal power supplies.

Each of the three choices comes with an associated cost or risk. The cost benefits of having a PLL at the system level must be sufficient to justify the extra design time, chip area, and risk of excess jitter and low yields, which increase chip cost. On the first Alpha AXP microprocessor project, these factors were not critical because high-end systems can afford the extra module and system cost to reduce the risk associated with delivering the product to market. With the lower-cost targets of the DECchip 21066 project, eliminating expensive SAW oscillators in favor of less-expensive, standard crystal devices is critical and justifies the chip-level costs and risks associated with introducing a new analog subsystem into the chip design.

Figure 4 illustrates how the PLL is a closed-loop feedback

system. An output is compared with an input to determine the degree of error. The error signal is then used to make adjustments within the elements of the system to reduce the error. This process occurs continually; under normal operating conditions, the error approaches zero. The major components of the PLL are

- o A phase frequency detector and a charge pump. The phase frequency detector inputs the reference and feedback signals and measures the phase error. The charge pump outputs a current proportional to this error, using digital pulse-width modulation to achieve an analog signal.
- o A filter. A simple, low-pass, resistor-capacitor (RC) filter averages the error signal, removing the digital noise of the pulse-width modulation used by the charge pump and setting the response rate of the PLL while ensuring loop stability. The filter is composed of an on-chip, 100-ohm resistor and the external capacitor.
- o A voltage-controlled oscillator (VCO). The VCO consists of a voltage-to-current buffer and a current-controlled oscillator. The output frequency of the VCO is proportional to the input voltage.
- o A feedback frequency divider (K_NDIV). In the linear control theory model of the PLL, the feedback signal is the oscillator phase attenuated by the gain $1/N$, where N equals the number of clock pulses input for every output pulse. This produces an oscillator phase output that is N times the input, which is the key to synthesizing a higher-frequency clock from a lower-frequency reference. A digital divider reduces the frequency of the VCO, thus reducing its phase as well. In this manner, a digital divide by N performs an analog attenuate-by- N function.

Although not strictly part of the PLL itself, the following two design elements are major components of the frequency synthesis system:

- o Power regulators. To reduce clock jitter, the PLL on-chip power is regulated using the noisy module 5-V power input. The use of power regulators yields a power source with less AC noise than the global 3.3-V power used by the digital logic and I/O pads.
- o A simple, programmable frequency divider (K_QDIV). This divider reduces the VCO frequency while ensuring a 50 percent duty cycle. The K_QDIV output is then buffered adequately to drive the global clock node.

[Figure 4 (Overview of the DECchip 21066 Phased-locked Loop Frequency Synthesis System) is not available in ASCII format.]

The RC filter and VCO create a textbook second-order feedback loop. This model uses linear control theory methods to analyze loop stability. The on-chip resistor is a key element in loop stability because it creates the transfer function zero that ultimately ensures stability. The off-chip capacitor, however, subjects the PLL to chip-switching noise coupled between the signal etches of the package. To minimize the effects of this AC noise on clock jitter, the design includes additional decoupling capacitors, such as the one shown in Figure 5. The capacitors add additional high-frequency poles. These poles are placed well above the PLL bandwidth, and are thus largely ignored when analyzing stability, but well below the package resonance frequencies that would cause significant jitter.

[Figure 5 (Use of a Decoupling Capacitor to Reduce Noise in an Analog Current Mirror) is not available in ASCII format.]

The implementation of the oscillator can take many different forms. The PLL design described in this paper uses a ring oscillator based on simple differential buffers. The reasons for this design decision are as follows:

- o Power supply current is nominally constant because the summation of current into each buffer in the ring smoothes the total.
- o Differential buffers used in the ring offer potential power supply rejection ration (PSRR) benefits.
- o Multiple phase outputs are possible. Although the DECchip 21066 microprocessor does not use multiple phase outputs, other applications of this basic oscillator design have exploited the feature.
- o Designers' experience with this oscillator implementation minimizes the risk, as compared to possible multivibrator designs.

The current design provides only for setting the clock frequency synthesis ratio at power up. Future designs may include the ability to control the synthesis ratio through software. Such control can provide a means of lowering chip power levels as well as more flexibility over using different clock rates in the same basic module design.

The off-chip capacitor in the low-pass loop filter does have an associated cost. Package coupling introduces noise on the control voltage, so placement of this capacitor is important. Future designs will likely use on-chip filters to reduce this constraint.

The power regulators are effective, but using 5 V will not be possible as CMOS technologies move to shorter channel lengths.

The trade-offs between on-chip regulation and different oscillator buffer designs that have better jitter characteristics for noisy power levels are yet to be investigated.

PACKAGE DEVELOPMENT

The DECchip 21066 microprocessor is the first CMOS microprocessor with a cost-focused package. Designers followed tight cost constraints to allow the package to be used on follow-on chip variants for the embedded microprocessor market. With a goal of 200-MHz operation, significant electrical analysis was necessary to ensure adequate package performance.

We learned from previous designs that the main concern regarding package design is ensuring the integrity of the on-chip supply voltage. For a given chip design and performance goal (which determines power dissipation, specifically, the chip's supply current characteristics), primarily one package characteristic and one chip characteristic influence on-chip supply integrity. These characteristics are package supply loop inductance (L_l), which is the intrinsic parasitic inductance of the supply path through the chip package, and on-chip supply decoupling capacitance (C_d). The on-chip supply integrity is generally proportional to the square root of C_d/L_l .

The supply loop inductance affects on-chip supply integrity because variations in the chip's supply current generate inductive noise voltage on the on-chip supply voltage. The magnitude of supply current variations is dependent on chip clock frequency. At 200 MHz, variations in supply current can be as much 3 amperes (A). If not reduced by some method, the inductive noise generated across the supply loop inductance would be excessive.

Three methods are typically used in microprocessor chip and package designs to decrease supply loop inductance: (1) increase the number of package supply pins, (2) increase the number of internal package layers (sometimes called planes), and (3) include on-package decoupling capacitors to shunt the supply loop inductance component of the package pins.

Given DECchip 21066 package cost constraints, designers ruled out several package features used in previous designs, notably the DECchip 21064 design. For example, on-package decoupling capacitors were considered too expensive, since they would have added approximately 50 percent to the package cost. To limit the package size, the number of available package pins was constrained to 287. Specifically, the number of package supply pins was reduced to 66 percent of the number implemented on the DECchip 21064 microprocessor. The number of internal package supply planes was reduced from four to two. Looser tolerances on internal package geometry allow the DECchip 21066 packages to be manufactured in the package vendors' less-costly process lines;

however, less-than-optimal supply routing results.

Given the constraints and the lack of complete package simulation models, we developed a method that combined simulation and empirical analysis to ensure adequate on-chip supply integrity. The expected package supply loop inductance was determined through a two-step process:

1. We modified the packages of a set of DECchip 21064 parts (by removing the on-package decoupling capacitor and reducing the number of supply pins) to closely match the proposed package configuration for the target chip. We determined the L1 of the degraded DECchip 21064 packages from the following relationship derived from the classical formula for the resonance frequency of an inductive and capacitive tank circuit:[20]

$$L1 = \frac{1}{4 \cdot \pi^2 \cdot Fr^2 \cdot Cd},$$

where Fr is the measured supply resonance frequency of the DECchip 21064 supply network, and Cd is the known on-chip supply decoupling capacitance for the DECchip 21064 microprocessor.

2. We used simulation to determine the effect on the supply loop inductance of eliminating expensive internal package features.

The resulting expected supply loop inductance for the DECchip 21066 packages was factored into the following relationship, which must be satisfied for the on-chip supply integrity (on-chip supply noise voltage) of the DECchip 21066 microprocessor to be equivalent to that of the DECchip 21064 device:

$$\frac{\text{SQRT}(Cd_{21064}/L1_{21064})}{S21064} = \frac{\text{SQRT}(Cd_{21066}/L1_{21066})}{S21066}$$

= Q21064
= quality factor of the DECchip 21064 on-chip supply,

where S is the maximum difference in chip supply current (Is). For example, the DECchip 21064 microprocessor operates with a minimum Is of 6 A and a maximum Is of 9 A (both at 200 MHz). Therefore, S21064 is equal to 3 A.

The amount of on-chip decoupling capacitance specified by this relationship was impractical to implement in the DECchip 21066 design. Because achieving the on-chip supply integrity of the DECchip 21064 design was not possible, we investigated the extent

of supply integrity degradation that could accompany a reliably functioning chip. We learned from system tests that DECchip 21064 chips with a quality factor as low as one-third the quality factor of original DECchip 21064 chips function correctly up to a chip frequency of 135 MHz.

Because the system testing was not exhaustive and because this testing could not be performed above 135 MHz, we wanted to further increase our quality factor. We added on-chip decoupling capacitance in areas not otherwise utilized and implemented package features that would minimize the supply loop inductance. Consequently, the DECchip 21066 microprocessor ultimately achieved a quality factor equal to 65 percent that of the DECchip 21064 microprocessor.

As a result of system testing the DECchip 21064 at reduced quality factors, Q21066 equal to 65 percent of Q21064 was deemed adequate. The recent opportunity to test and measure DECchip 21066 parts validated the decision to reduce the quality factor to this level.

SUMMARY

This paper describes the DECchip 21066 microprocessor, the first cost-focused chip to implement the Alpha AXP architecture. The project goal to achieve best-in-class system performance for cost-focused applications was realized by integrating a high-performance CPU, a high-bandwidth memory controller, a phase-locked loop, and a low-cost package, and by being the first microprocessor in the industry that connects directly to the PCI bus. A pseudorandom verification strategy that leveraged an architectural reference with an I/O stream was a key decision that helped achieve first-pass silicon success on a compressed 10-month schedule.

ACKNOWLEDGMENTS

The authors would like to thank the entire DECchip 21066 team. The efforts of all those involved in the microprocessor architecture, the design, the implementation, the logic verification, and the layout made it possible to bring this project to fruition.

REFERENCES AND NOTE

1. R. Sites, ed., Alpha Architecture Reference Manual (Burlington, MA: Digital Press, 1992).
2. Peripheral Component Interconnect (PCI) (Hillsboro, OR: PCI Special Interest Group, rev. 2.0, April 30, 1993).

3. D. Dobberpuhl et al., "A 200-MHz 64-bit Dual-issue CMOS Microprocessor," *Digital Technical Journal*, vol. 4, no. 4 (Special Issue 1992): 35-50.
4. B. Zetterlund, J. Farrell, and T. Fox, "Microprocessor Performance and Process Complexity in CMOS Technologies," *Digital Technical Journal*, vol. 4, no. 2 (Spring 1992): 12-24.
5. Intel Pentium Processor Data Book, vol. 1 (Santa Clara, CA: Intel Corporation, 1993).
6. Intel486 DX Microprocessor Data Book (Santa Clara, CA: Intel Corporation, 1991).
7. G. Darcy III et al., "Using Simulation to Develop and Port Software," *Digital Technical Journal*, vol. 4, no. 4 (Special Issue 1992): 181-192.
8. W. Anderson, "Logical Verification of the NVAX Chip Design," *Digital Technical Journal*, vol. 4, no. 3 (Summer 1992): 38-46.
9. D. Su et al., "Experimental Results and Modeling Techniques for Substrate Noise in Mixed-Signal Integrated Circuits," *IEEE Journal of Solid State Circuits*, vol. 28, no. 4 (April 1993): 420-430.
10. F. Gardner, *Phaselock Techniques*, 2d ed. (New York: John Wiley & Sons, 1979).
11. R. Best, *Phase-Locked Loops* (New York: McGraw-Hill, 1984).
12. P. Gray and R. Meyer, *Analysis and Design of Analog Integrated Circuits* (New York: John Wiley & Sons, 1984).
13. M. Wakayama and A. Abidi, "A 30-MHz Low-Jitter High Linearity CMOS Voltage-Controlled Oscillator," *IEEE Journal of Solid State Circuits*, vol. SC-22, no. 6 (December 1987): 1074-1081.
14. K. Kurita et al., "PLL-Based BiCMOS On-Chip Clock Generation for Very High-Speed Microprocessors," *IEEE Journal of Solid State Circuits*, vol. 26, no. 4 (April 1991): 585-589.
15. S. Miyazawa et al., "A BiCMOS PLL-Based Data Separator Circuit with High Stability and Accuracy," *IEEE Journal of Solid State Circuits*, vol. 26, no. 2 (February 1991): 116-121.
16. I. Young, J. Greason, and K. Wong, "A PLL Clock Generator with 5 to 110 MHz of Lock Range for Microprocessors," *IEEE Journal of Solid State Circuits*, vol. 27, no. 11 (November 1992): 1599-1607.
17. N. Nguyen and R. Meyer, "Start-up and Frequency Stability in High-Frequency Oscillators," *IEEE Journal of Solid State Circuits*, vol. 27, no. 5 (May 1992): 810-820.

18. J. Melsa and D. Schultz, *Linear Control Systems* (New York: McGraw-Hill, 1969).
19. SPICE is a general-purpose circuit simulator program developed by Lawrence Nagel and Ellis Cohen of the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley.
20. W. Hayt et al., *Engineering Circuit Analysis* (New York: McGraw-Hill, 1978).

TRADEMARKS

The following are trademarks of Digital Equipment Corporation:
Alpha AXP, AXP, DEC, DECchip, and OpenVMS.

Intel, Intel486, and Pentium are trademarks of Intel Corporation.

Microsoft and Windows NT are registered trademarks of Microsoft Corporation.

OSF/1 is a registered trademark of Open Software Foundation, Inc.

SPICE is a trademark of the University of California at Berkeley.

BIOGRAPHIES

Masooma Bhaiwala A hardware engineer in the Models, Tools, and Verification Group of Semiconductor Engineering, Masooma Bhaiwala has contributed to several projects since joining Digital in 1991. She worked on the I/O controller verification of the DECchip 21066 microprocessor and wrote the PCI transactor that mimics Intel's PCI bus protocol. Recently, she led the DECchip 21066 second-pass verification efforts. Masooma received a B.S.C.E. from N.E.D. University of Engineering and Technology, Pakistan, and an M.S.C.E. from the University of Massachusetts at Lowell.

Kwong-Tak A. Chui Kwong Chui, a principal engineer in the Semiconductor Engineering Group, is currently leading the definition of a core chip set for a RISC processor. He was the co-architect of the PCI interface section of the DECchip 21066. Since joining Digital in 1985, Kwong has worked on memory and I/O bus controller chips for the MicroVAX 3000 and VAX 4000 series systems, and a cache address fanout buffer and multiplexer chip for the DEC 4000 AXP system. Kwong holds a B.S. in computer engineering from the University of Illinois at Urbana-Champaign and an M.Eng.E.E. from Cornell University.

Christopher L. Houghton A principal hardware engineer in the Semiconductor Engineering Group, Chris Houghton is responsible for the signal integrity, packaging, and I/O circuit design of the DECchip 21066 Alpha AXP microprocessor. His previous work includes signal integrity, packaging, and circuit design for the GIGAswitch crossbar chips, digital and analog circuit design for a semicustom cell library, and advanced development of processor-memory interconnect chips. Chris came to Digital in 1985 after receiving a B.S.E.E. from the University of Vermont. He holds one patent.

Daniel L. Leibholz Daniel Leibholz, a senior engineer in the Semiconductor Engineering Group, has been designing integrated Alpha AXP CPUs at Digital for three years. He is an architect of the DECchip 21066 Alpha AXP microprocessor and is currently involved in the design of a high-performance microprocessor. While at Digital he has also worked on the development of massively parallel processors and was awarded a patent for a processor allocation algorithm. Daniel joined Digital in 1988 after earning B.S. and M.S. degrees in electrical engineering from Brown University in Providence, RI.

Dina L. McKinney Principal engineer Dina McKinney supervises the

DECchip 21066 low-cost Alpha AXP verification team. Since joining Digital in 1981, Dina has held design and verification positions on various projects, including PC modules, semicustom chips, and full-custom chips for communication, graphics, and processor products. Previously, Dina was a member of the U.S. Air Force. She holds two year technical degrees from Community College of the Air Force and Gulf Coast College, a B.S.E.E. (1982, summa cum laude) from Central New England College, and an M.S.E.E. (1990) from Worcester Polytechnic Institute.

James R. Mullens James Mullens is a principal hardware engineer in the Semiconductor Engineering Group. He is working on CPU advanced development and is responsible for the architecture and design of the DECchip 21066 dynamic memory controller. In earlier work, he performed verification and test of the system-on-a-chip, designed four ASIC devices used in workstations and DECwindows terminals, led the Corporate Mouse and Tablet projects, and was a design and project engineer for the Mini-Exchange. Jim joined Digital in 1982 after receiving a B.S.E.E. from Northeastern University.

Sanjay J. Patel As a verification engineer for the DECchip 21066 low-cost Alpha AXP project, Sanjay Patel was a key implementor of the DECchip 21066 verification strategy. He is currently working on specifying power management capabilities for a future chip. Sanjay joined Digital in 1992. In his earlier experience as a co-op student at IBM in Kingston, NY, he helped develop a vectorized subroutine library. Sanjay received B.S.E. (1990) and M.S.E. (1992) degrees in computer engineering from the University of Michigan.

Delvan A. Ramey A principal hardware engineer, Del Ramey specializes in analog CMOS development for CPU, communications, and graphics chips. His recent work involves phase-locked loop (PLL) design. Del architected the hybrid analog-to-digital PLL clock and data recovery system in the Ethernet protocol DECchip 21040 product, and designed the frequency synthesis PLL for the DECchip 21066 Alpha AXP microprocessor. He has also developed semicustom chips and libraries. Del joined Digital in 1980, after completing his Ph.D. at the University of Cincinnati. He is a member of IEEE, Eta Kappa Nu, and Phi Kappa Phi.

Mark B. Rosenbluth Consulting engineer Mark Rosenbluth contributes to the architecture, modeling, and logic design of the DECchip 21066 microprocessor. In earlier work, he was involved in the modeling, logic design, and circuit design of the DC222 system-on-a-chip (SOC) VAX microprocessor, the DC358 MicroVAX DMA controller chip, and the J-11 PDP-11 microprocessor. He also helped create semicustom design system and cell libraries. Mark came to Digital in 1977 from RCA Corporation. He

holds a B.S.E.E. from Rutgers University College of Engineering.

=====
Copyright 1994 Digital Equipment Corporation. Forwarding and copying of this
article is permitted for personal and educational purposes without fee
provided that Digital Equipment Corporation's copyright is retained with the
article and that the content is not modified. This article is not to be
distributed for commercial advantage. Abstracting with credit of Digital
Equipment Corporation's authorship is permitted. All rights reserved.
=====