HITACHI

SuperH®
RISC engine

The SH-5 Architecture

The fifth-generation, 64-bit SuperH® RISC Engine architecture—co-developed by Hitachi and STMicroelectronics—is an optimized solution for high-performance, low-cost consumer and next-generation embedded applications

Technology/Marketing White Paper
(October 1999)

For more information:

- *Hitachi:*
  Phil Gomes, Weber Group (650-463-8611)
- *STMicroelectronics:*
  J.P. Rossomme, Public Relations, (602-485-2262)

The SH-5 Architecture

I. Introduction

II. Overview of the SH-5 Architecture

III. The First SH-5 product

IV. Appendix

Overviews of products based on the SuperH architecture

The SH-5 Architecture

The fifth-generation, 64-bit SuperH® RISC Engine architecture—co-developed by Hitachi and STMicroelectronics—is an optimized solution for high-performance, low-cost consumer and next-generation embedded applications

I. Introduction

1.1 SuperH generations share a common goal

From the earliest technology discussions that led to the creation of the SuperH® RISC engine architecture nearly a decade ago, to the development efforts now under way or planned, there has been one basic engineering and marketing goal for the product line. The essential parts of that common goal are

- to provide an extended series of upward-compatible microcontroller (MCU) and microprocessor (MPU) devices
- to offer optimized balances of performance, power consumption, integration and die size
- to allow customers to take full advantage of windows of market opportunity
- to deliver economical devices that customers can use to build systems that offer the price/performance levels needed to achieve high sales volumes

The four generations of SuperH Cool Engine™ RISC processors currently in production conform to an aggressive, periodically updated technology roadmap. Enthusiastic customer response worldwide has earned the architecture a leadership position worldwide in the 32-bit embedded RISC market.

To supply customers with advanced processors for the products and systems of the next decade, the SuperH roadmap specifies a fifth-generation architecture (and, beyond that, a sixth). Development of the fifth-generation architecture was guided by the overall SuperH series engineering and marketing goal described previously. To fulfill that goal, given today's evolving, escalating market requirements, the development team had to overcome many design challenges. Specifically, they had to create a microprocessor core that enables next-generation system-on-a-chip (SOC) consumer products, provides enhanced performance for multimedia applications, and reduces customers' time to market. The Hitachi and STMicroelectronics (ST™) design team accomplished this and more.

1.2 Updated SuperH roadmap

The latest revision of the technology roadmap for the SuperH architecture (right) puts the key features of the SH-1, SH-2, SH-3 and SH-4 RISC series into perspective. It also shows the performance targets for the SH-5 RISC engine architecture.

1.3 Cooperative development effort

Hitachi developed four generations of the SuperH architecture and the dozens of MPU/MCU devices in the SH-4, SH-3, SH-2 and SH-1 series. For the fifth-generation architecture, Hitachi formed a strategic alliance with STMicroelectronics (ST) in December 1997—a true technology and marketing partnership. The agreement initiated an in-depth collaboration to develop (using a common design methodology) 64-bit, 700- to 1000-MIPS SuperH MPUs for applications such as interactive set-top boxes (STBs), telecom/datacom networks, digital video products, and automotive multimedia systems.

As part of the agreement, ST licensed from Hitachi the SH-4 core to manufacture and market the ST40-series CPUs. Other current licensees of SuperH technology include Seiko-Epson, NEL and Sony.

1.4 Key aspects of the Hitachi–ST alliance

By combining their engineering talent, Hitachi and ST are significantly accelerating the introduction of higher-performance, low-cost processors based on the fifth-generation SuperH Cool Engine RISC architecture—devices that will populate Hitachi's SH8000 series and ST's ST50 series.

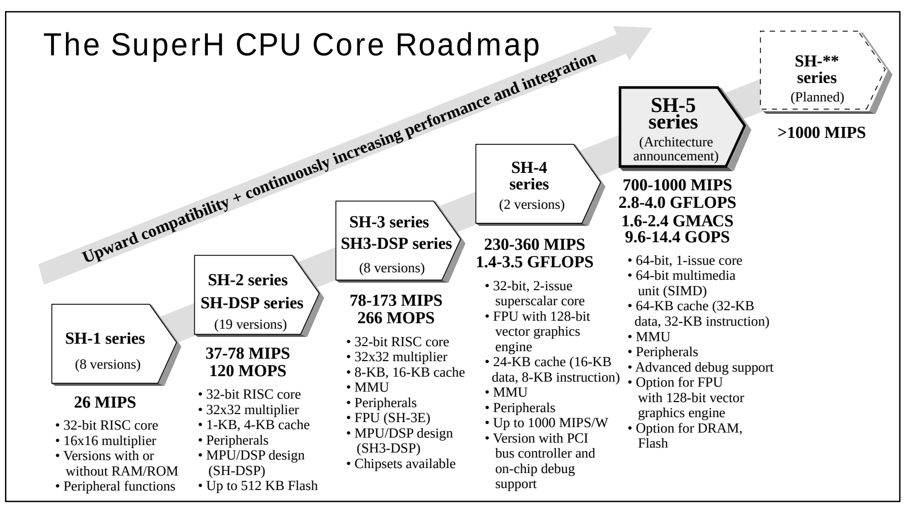

The technology roadmap for the SuperH architecture extends through five generations of products; a sixth generation is now being planned. The upward-compatibility gives system engineers considerable design flexibility. Systems can be upgraded for higher performance and greater functionality, while investments in hardware and software development are preserved.

Both companies have leadership positions in key markets. Hitachi is #1 worldwide in embedded RISC. ST is #1 in digital consumer set-top box CPUs. Both companies expect a strong positions in future embedded computing markets such as HDTV, digital imaging, multimedia, broadband networks, cable systems, VoIP equipment, monitors and displays, and wireless products. Together, the two companies shipped 33 million 32-bit RISC processors in 1998 (Hitachi shipped 26 million; ST shipped 7 million). Total shipments of SuperH devices are expected to exceed 100 million by the end of 1999.

The technology/marketing partnership between Hitachi and ST is creating an architectural standard for embedded systems at the 64-bit level. Design teams are developing the fifth-generation architecture in San Jose, CA, with support provided by worldwide resources of both companies. Other distinguishing features of the partnership include the

- co-development of an advanced 0.15-µm process technology, necessary to meet aggressive chip speed, power and cost objectives for the fifth-generation architecture and future SH-4/ST40 products
- pooling of the companies' intellectual property, both hardware and software
- sharing of development/integration expertise and product support resources
- guarantee of full compatibility between the CPUs produced by both companies.

1.5 Making the architecture system-centric

By applying their combined strengths, Hitachi and ST have fundamentally advanced the SuperH product line, making the jointly-developed fifth-generation architecture the first to allow the implementation of sophisticated SOC products. Older-generation SuperH chips are RISC-core-centric designs that combine a fast CPU with common peripherals and memory. The fifth-generation architecture, by contrast, is system-centric. It enables system-on-a-chip devices that integrate the CPU, an ultra-high-speed on-chip interconnect bus, complex subsystems and common peripherals.

The CPU will be just a small part of typical chip designs built with the fifth-generation architecture. Therefore, an integral part of the architecture is a well-developed SOC methodology that allows the rapid implementation and debugging of cost-effective silicon solutions.

The SOC methodology allows extensive utilization of the valuable reusable libraries of memory arrays, intelligent macros, peripheral functions, on-chip debugging circuits, and other designs that have been amassed by Hitachi and ST. The companies will use their strong integration expertise to integrate selected silicon assets with the fifth-generation 64-bit RISC processor core. This will enable them to offer truly compatible, fully-second-sourced SOC and chipset solutions to their respective customers with quicker turnaround times and guaranteed cross-vendor upward compatibility.

1.6 Additional benefits of the Hitachi–ST alliance

In addition, by combining their extensive expertise in systems software, and by leveraging their relationships with third-party suppliers, Hitachi and ST will be able to

- provide on-chip debugging capabilities that are powerful, non-intrusive and cost-effective
- give customers access to a comprehensive span of effective, time-saving software development tools
- offer a wide range of software drivers and middleware that customers can use for product differentiation
- support an exceptionally broad range of operating systems and third-party application software packages.

1.7 Current SuperH series devices

There are now over 35 different MPU/MCU products based on the SuperH architecture, the most robust and diverse in the industry, providing customers with unmatched flexibility for meeting their system design targets. The popular chips are now used in over 3200 products.

Despite the diversity, devices within the SuperH family share many characteristics due to the multi-faceted main engineering/marketing goal consistently applied to the evolution of the architecture. Among the characteristics pervasive within the product line are the following:

- configurations optimized for embedded applications and designed to provide efficient, cost-effective total solutions
- upward-compatible instruction sets based on 16-bit fixed length instructions that provide the high code efficiency needed to reduce system memory requirements
- efficient chip designs with small die sizes to facilitate high quantity manufacture and ready availability
- balanced combinations of performance and power, offering the moderate to fast speeds needed for higher throughput and the low power dissipation (Cool Engine operation) that allows the use of low-cost plastic packages, extends battery life in portable products, and minimizes cooling problems
- competitive costs so customers' designs can attain high volumes in price-sensitive high-volume markets
- broad operating system support for customer convenience and maximum design versatility

The fifth-generation architecture is upward compatible with the SH-4 series, which is described in the Appendix (Section 4.1 page 12). Descriptions of devices in the SH-3, SH-2 and SH-1 series are also included in the Appendix.

II. Overview of the SH-5 Architecture

2.1 Major design objectives

The aims of the SH-5 architecture are low-power operation and small chip size, coupled with high clock frequencies and high levels of performance (>700 MIPS). These are key requirements for successful next-generation embedded implementations in price-sensitive markets. The basic feature set of the SH-5 RISC engine was determined by the needs of immediate consumer applications such as set-top boxes, DVD players, HDTV, as well as the requirements of future applications such as handheld PCs and notebook computers, voice-over-IP (VoIP) equipment, in-vehicle computing, voice recognition equipment, and gaming and entertainment products.

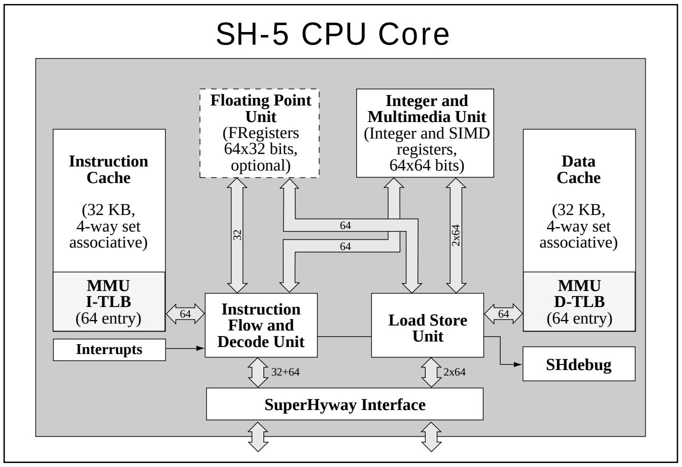

The SH-5 64-bit SuperH RISC core is a scalar, single-issue design that interfaces to a 200-MHz, pipelined, split-transaction on-chip bus. The high-performance core has a 7-stage integer pipeline, multimedia processing unit, and (optionally) a 64-bit floating point unit.

To preserve customer's investments in hardware/software development, upward compatibility is maintained with previous generations of SuperH series devices. To ease software development, the architecture is designed for applications written in C/C++ and Java™, running advanced operating systems that require higher performance processors, including Windows® CE, pSOS™, VxWorks™, Linux™, OS-9™ and JavaOS.

2.2 Key aspects of the SH-5 architecture

The SH-5 is a scalar, single-issue 64-bit RISC design with the following key distinguishing features.

- Two operating modes are supported:
  - The SHmedia mode is a clean-slate definition. It has a complete instruction set that supports 32-bit instruction codes and delivers high multimedia performance for integer, "packed arithmetic/SIMD" and floating point operations. It can perform powerful parallel executions on 8-, 16- and 32-bit objects, and easily mixes scalar and multimedia operations. The SHmedia mode is typically used for time-critical routines.
  - The SHcompact mode is a complete instruction set that supports 16-bit instruction codes for higher code density, including legacy instructions of earlier-generation SuperH RISC devices. This mode provides user-mode instruction compatibility with software written for SH-4 series MPUs. The SHcompact mode is generally used to reduce the storage requirements of code that is not especially time critical (typically, this is most of the code).
- Mode switching occurs dynamically at branch instructions.
- A split-branch approach achieves zero delay on branches most of the time by hiding pipeline "flushes" (clear/refill operations) that otherwise would delay code execution.
- The carefully chosen SIMD core instructions were built into the SHmedia mode from the beginning. They operate on three operands, each of which may have eight 8-bit, four 16-bit, or two 32-bit values. This enables throughput as high as 9.6 GOPS at 400 MHz. SIMD supports signed/unsigned/fraction and saturate/modulo operations.
- A removable IEEE-754 double-precision FPU provides a vector sum of products and matrix row transform for 3D graphics. The FPU performs 4 multiply and 3 addition operations every cycle, achieving 2.8 GFOPS at 400 MHz.
- An ultra-high-speed, on-chip 64-bit interconnect bus—the SuperHyway™ bus—that delivers high levels of interconnectivity. It supports a memory-mapped, packet-based split-transaction protocol and achieves 3.2 GB/s bandwidth when operating at 200 MHz (half clock speed).
- The SHdebug link lets engineers nonintrusively analyze system behavior. For example, they can trace execution flow, watchpoint locations, and on-chip bus traffic. The watchpoint capability is a valuable development aid because it allows continuous operation while data is reported on specified events. Control and related information are obtained by a low-cost, high-bandwidth connection between a JTAG port and an adapter board in the host system.

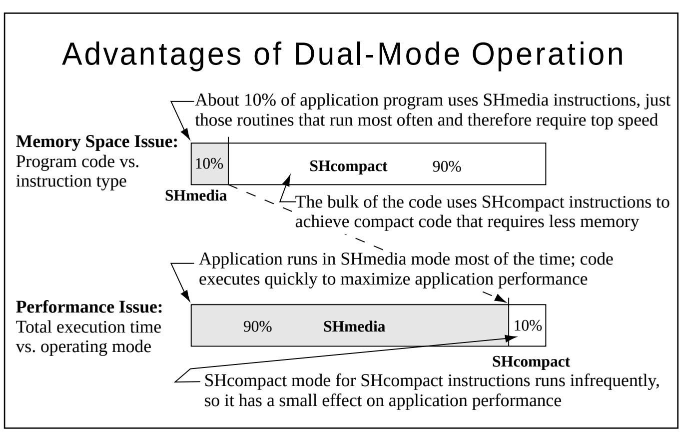

With the SH-5 architecture, performance-critical application code can be optimized for speed in SHmedia mode, and code that isn't time critical (because it doesn't run very often) can be optimized for reduced memory size and reduced bandwidth in SHcompact mode.

2.3 A true SOC methodology

The design of the SH-5 architecture uses a true system-on-a-chip design methodology. Thus, the SH-5 CPU core is a reusable hard macro. It will be offered as an ASIC module that can be moved easily between manufacturing facilities and used as the basis for a wide range of MPUs. Optimized chips for different applications can be developed using the SH-5 core and required modules selected from libraries of reusable peripheral functions and complex subsystems, such as on-chip Flash, embedded DRAM, MPEG decoders and a FireWire™ interface. Portability of the SH-5 core is enabled by its interface to the powerful, flexible SuperHyway on-chip bus architecture.

Advanced computer-aided engineering (CAE) tools have been used to develop the SH-5 architecture, and these tools facilitate the development of targeted chip variations. The state-of-the-art transparent, non-intrusive debug support integrated into the core will help reduce development time and time-to-market, as will the SH-5 compilers being developed by Hitachi, ST and GNU suppliers. This debug capability will also be useful for end-product system service and support needs.

2.4 Features of the architecture

2.4(a) General

The SH-5's dual-mode instruction set architecture (ISA) gives system engineers the flexibility to achieve a wide span of design objectives. For example, the dynamic mode switching allows a compiler to optimize both code density and performance. SHmedia modes and SHcompact modes can be mixed on boundaries separated by branch instructions.

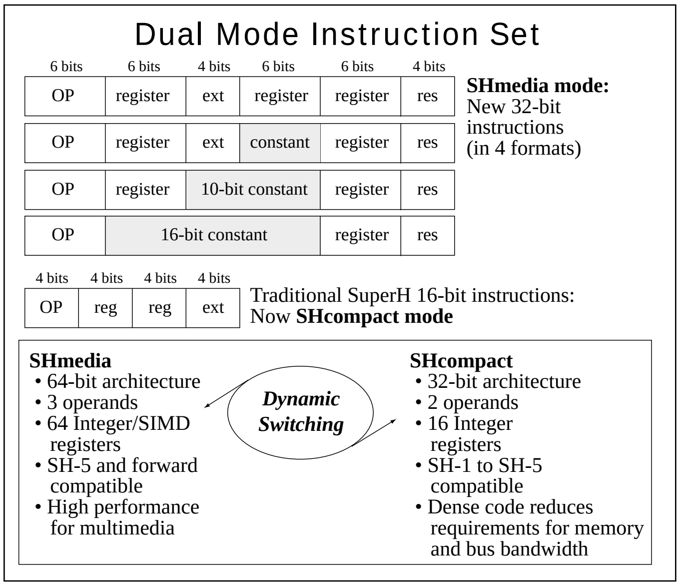

Four of the 32-bits in the SH-5's SHmedia instructions are reserved for future enhancements.

The SHmedia mode includes in its complete set of 32-bit instructions a set of SIMD instructions for multimedia applications, including compare, addition, subtraction and shifts (with and without saturation); fractional multiplication and multiply accumulate; absolute, sum of difference (for motion estimation); and condition, move, data conversions and re-arrangement. The SIMD instructions support signed, unsigned, and fractional data types and saturate and modulo results.

The SIMD instructions are fully integrated into the CPU hardware that supports the SHmedia, and they complement the integer instruction set. SIMD instructions execute on multiple 8-bit (B), 16-bit (W) and 32-bit (L) data elements organized (packed) into sixty-four 64-bit general-purpose registers.

The large number of general-purpose registers in the SH-5 architecture is very useful for multimedia applications. Compute-intensive multimedia inner loops can be supported, as can techniques such as loop unrolling, data prefetching, software pipelining and instruction scheduling. Many algorithms contain both scalar integer and SIMD instructions, and there are enough registers to store both, which eliminates unnecessary data movement. (To further reduce data movement, the SH-5's separate 32/64-bit floating point registers aren't used for SHmedia integer instructions.)

In addition, the wealth of registers in the SH-5 CPU allows the use of simple, efficient software conventions for passing

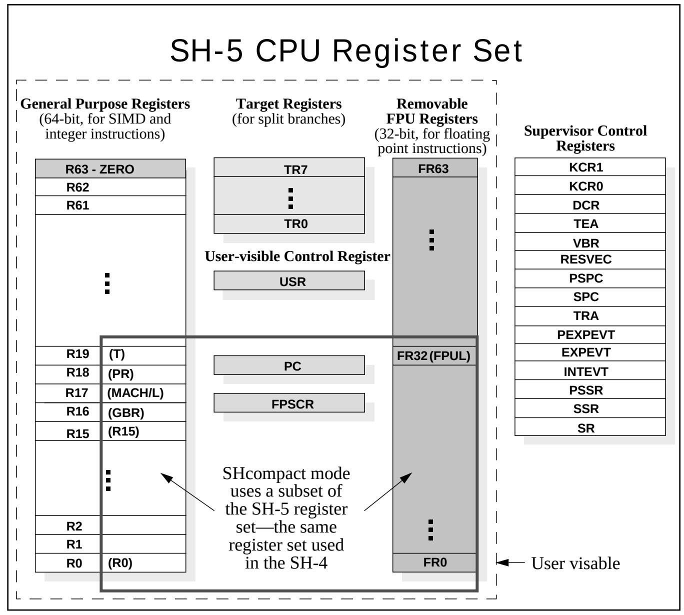

The SH-5 architecture's large general purpose register file (sixty-four 64-bit registers) is used in the SHmedia mode to efficiently execute code for multimedia applications. The SHcompact mode uses a set of 16-bit registers identical to the set in the SH-4 architecture. The FPU registers can be removed from the core if they aren't needed, thus saving die area.

multimedia parameters via on-chip registers. This is aids the design of efficient, high-performance optimizing compilers.

The multimedia registers are separated from the FPU registers. The FPU registers can be removed to build lower power, more economical CPU implementations for applications that do not require floating point instructions.

2.4(b) SHmedia mode

The SHmedia mode has a set of 203 32-bit, fixed-length instructions. The SHmedia instructions

- can execute on three operands
- address all sixty-four 64-bit general purpose registers
- support integer, floating point and SIMD arithmetic
- provide simple decode for fast implementations
- reserve 4 bits for future architectural enhancements

The 32-bit instructions in the SHmedia mode have 3-operand encoding for 64-bit registers (Rm + Rn → Rd). They efficiently support 32-bit and 64-bit (address width) software.

2.4(c) SHcompact mode

The SHcompact mode has a set of 201 16-bit fixed-length instructions. The SHcompact instructions

- can execute on two operands
- address sixteen 32-bit general purpose registers, the same set that is used in the SH-4 architecture
- support integer and floating point arithmetic
- produce dense code that reduces storage requirements
- offer user-mode compatibility for software written for SH-4/ST40 series processors

Preliminary data for the SH-5 instruction set is presented in Section 4.5 of the Appendix (page 13).

2.4(d) Split-branch architecture

The SH-5 architecture has a unique branch method in the SHmedia mode that can achieve zero branch penalty most of the time. This eliminates the wasted clock cycles that would otherwise be needed to refill the pipeline.

Branches are split into two parts: prepare-to-branch (prepare target address) and branch (branch to target). In the prepare-to-branch part, the target address is loaded into one of the eight target branch registers. The hardware automatically begins to fetch two instructions at that address if the “target-likely-to-be-used” bit is set. If both the prepare target address and the branch predictions are true, there is no exception delay.

In effect, this design approach hides pipeline flushes that would otherwise delay code execution. A flush does occur whenever there is a branch, but the flush is hidden by the fact that two instructions are already in the target buffer.

Compilers can often schedule a prepare-to-branch early, thus allowing the processor to prefetch instructions at the branch target so that those instructions are ready when needed by the branch instruction. This software approach eliminates the need for complex branch prediction hardware in the RISC chip. Also, the methodology is extendible to the longer pipelines needed to achieve higher clock speeds.

Note that the prepare target address (likely) instruction is executed only once—prior entering the loop. Because PTA isn’t repeated inside the loop, it imposes only a one-instruction execution delay. Also, SH-5 branch instructions combine conditional branching with the compare instruction. This eliminates a separate compare instruction prior to a conditional branch.

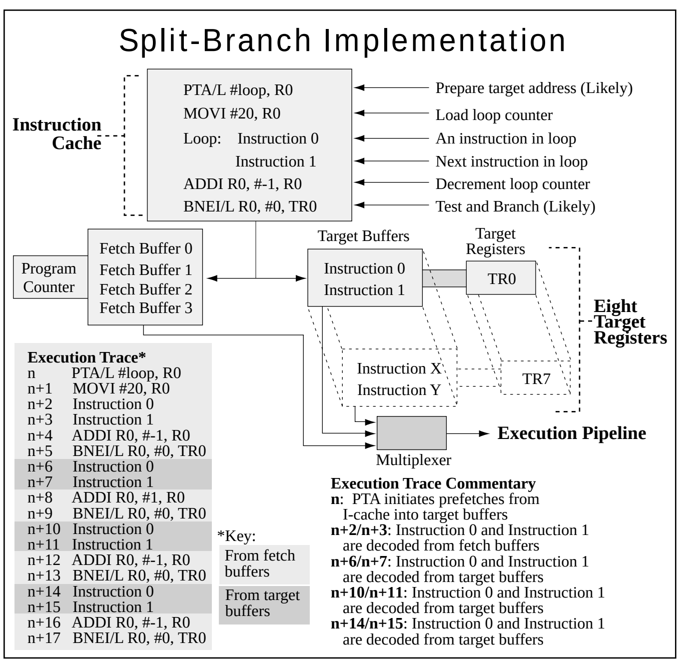

Branch penalties can be eliminated because the SH-5 CPU can prefetch target addresses into one of eight branch registers prior to the execution of a branch instruction.

2.5 Hardware implementation

2.5(a) General

The 64-bit SH-5 single-issue integer CPU core has a 7-stage pipeline: Fetch-1 (F1), Fetch-Decode (FD), Decode (D), Execute-1 (E1), Execute-2 (E2), Execute-3 (E3), and Writeback (W). The pipeline uses a decoupled pipe-file to store results before writing results back to the register set in the writeback stage. It allows zero-penalty branching and has full forwarding capability. Support is provided for pipelined back-to-back MAC instructions and pipelined stores.

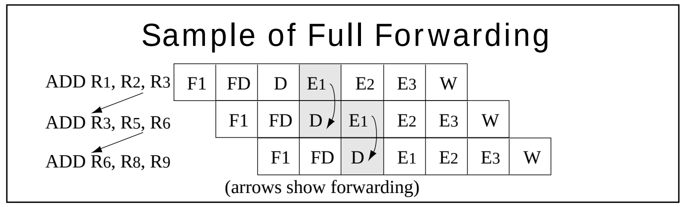

Full forwarding of results from execute stages of the 7-stage pipeline eliminates pipeline stalls. Simple instructions, like the ADD shown here, produce results at the end of E1. Those results are immediately available for subsequent instructions. The SH-5’s advanced branch architecture hides instruction fetch pipeline stages to allow zero-overhead branching.

The CPU core executes many instructions with one cycle pitch, and full data forwarding is used to ensure minimum data stalls and maximum throughput. Two decoders in the core module allow high clock rates while supporting the SHmedia and SHcompact modes (only one mode is used at a time). A mode-switching branch allows operation of the CPU to shift dynamically between SHmedia code and SHcompact code.

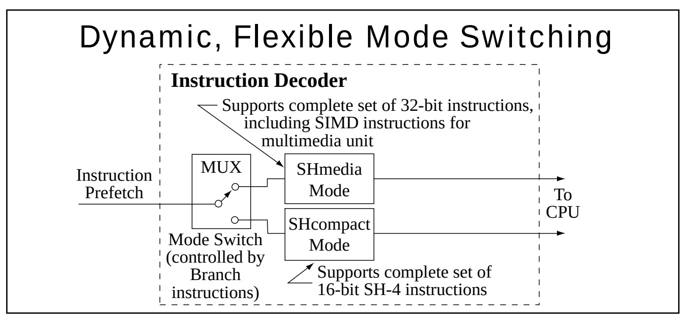

Operation of the SH-5 architecture can switch from SHmedia mode to the SHcompact mode at boundaries separated by branch instructions.

2.5(b) Virtual caches and memory management

The CPU core includes separate 32-KB virtual instruction and data caches that are four-way set associative (32-byte cache line) and optimized for high speed and low power. Fully-associative translation look-aside buffers (TLBs) with 64 entries are provided for each instruction and data cache for memory management, including memory protection and translation.

Virtual cache provides several advantages over a physical cache approach. It allows the CPU core to access the cache without turning on the TLB (the TLB is accessed only when there is a cache miss). This decreases power dissipation and increases data throughput. The virtual cache also decreases the chances for TLB misses.

Instructions and data can be locked to implement privilege and user modes to ensure that time-critical data isn't removed from the cache.

2.5(c) Multimedia unit

The integer and multimedia units share sixty-four 64-bit general purpose registers. The SIMD (single-instruction, multiple data) instructions can be performed on 8 pieces of 8-bit data, 4 pieces of 16-bit data, and 2 pieces of 32-bit data.

The SIMD instructions are part of the SH-5's 32-bit instruction set. They are highly efficient when large amounts of parallelism exist on multiple pieces of data. The data must be organized (packed), and the 64-bit registers can be configured to handle the required number of bits; for example, eight 8-bit data words. Once the packed data is loaded, SIMD instructions perform multiple operations of the same type simultaneously.

SIMD performance is outstanding. The multimedia unit performs 4 MACs every cycle; that is, it can multiply-accumulate (4x16-bit multiply, 3x32-bit add, and 1x64-bit accumulate) on two 64-bit registers, each of which contains four 16-bit packed integers. This is 8 arithmetic operations per clock cycle. Thus, at the 400 MHz clock speed, the SH-5 architecture performs 3.2 billion operations per second (3.2 GOPS). Also, because SIMD allows 4 accumulates, four 16-bit MACs per clock cycle, the architecture performs 1.6 billion MACs per second, earning it a 1.6 GMAC rating.

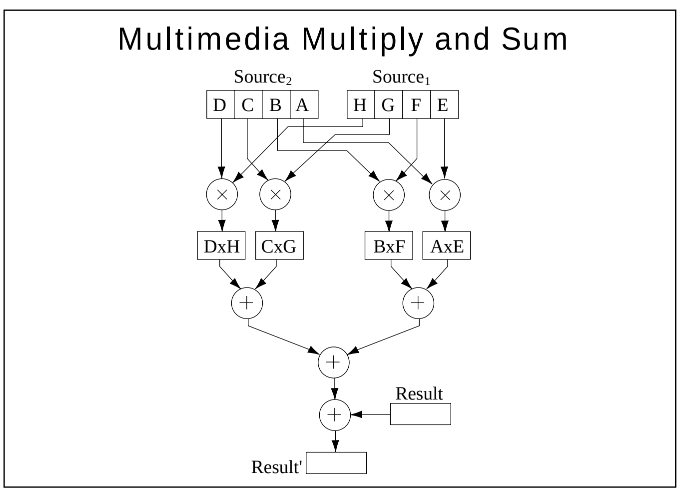

In this sum-of-products operation, the SH-5 performs—in one cycle—four 16x16–>32 multiplications, then reduces the results to a single 64-bit accumulation. This is 1.6 GMACS, the architecture's peak performance. (For traditional MACs, which use two multiplications and provide separate 32-bit results for each MAC, the SH-5 delivers 800M MACS speed.)

In multimedia mode, the SH-5 architecture can be used to rapidly execute the sum of absolute-differences operation needed for MPEG encoding. For example, in one instruction cycle—just 2.5 ns—it performs a packed sum of absolute-difference-accumulate operation that accumulates 8 pieces of 8-bit data. This is 9.6 GOPS performance.

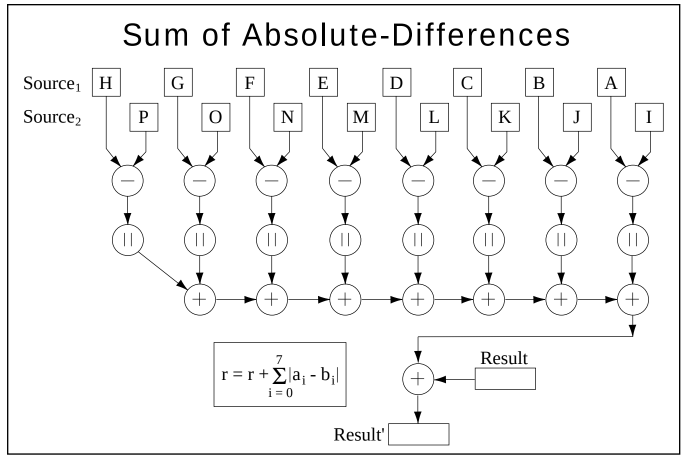

In MPEG encoding, information is transmitted only when a pixel changes from one frame to the next. The sum of absolute-differences operation that enables the data reduction can be performed by the SH-5's multimedia unit every 2.5 ns. This achieves a throughput of 8 subtracts, 8 absolutes and 8 adds, resulting in a 9.6 GOPS performance rating.

2.5(d) Removable floating point unit

The removable floating point unit (FPU) supports IEEE-754-compatible single-precision and double-precision operations, as well as a set of special-purpose operations for 3D graphics. The FPU, which has a 9-stage pipeline structure, performs the vector inner product and matrix transformation operations used for processing 3D graphics. Its flexible register set can function as sixty-four 32-bit registers for single-precision operations, thirty-two 64-bit registers for double-precision operations, sixteen 128-bit vectors for four single-precision operations, or any combination that takes advantage of sixty-four 64-bit registers.

Each of the FPU's four floating point multipliers (fmuls) can receive two 32-bit values and produce a multiplied result that is passed to a four-input floating point adder. The hardware reads two 128-bit vectors (two sets of four 32-bit values) out of register files, multiplies the four 32-bit pairs at the same time, adds the four products together, and puts the 32-bit result back into the register file. This provides the equivalent of 288-bit data crunching (2 x 128 + 32 = 288). (Operation of the SH-5 FPU is comparable to that of the FPU used in SH-4 series MPUs.)

The capabilities of the FPU are illustrated by the following computation, a 1x4 matrix multiplied by a 4x4 matrix. This computation is performed in seven clock cycles.

$$FR_0*FR_{16} + FR_1*FR_{20} + FR_2*FR_{24} + FR_3*FR_{28} \rightarrow FR_0$$
$$FR_0\ FR_{17} + FR_1*FR_{21} + FR_2*FR_{25} + FR_3*FR_{29} \rightarrow FR_1$$
$$FR_0*FR_{18} + FR_1*FR_{22} + FR_2*FR_{26} + FR_3*FR_{30} \rightarrow FR_2$$
$$FR_0*FR_{19} + FR_1*FR_{23} + FR_2*FR_{27} + FR_3*FR_{31} \rightarrow FR_3$$

The fully pipelined SH-5 architecture can repeat these 16 fmuls and 12 fadds (28 operations) every four clock cycles, for an average of seven floating point operations per cycle. The CPU and double-precision fmov allow registers to be loaded from, and stored to, cache during these four cycles, so the operations are sustainable. (The flow diagram for this matrix computation is shown on page 9.)

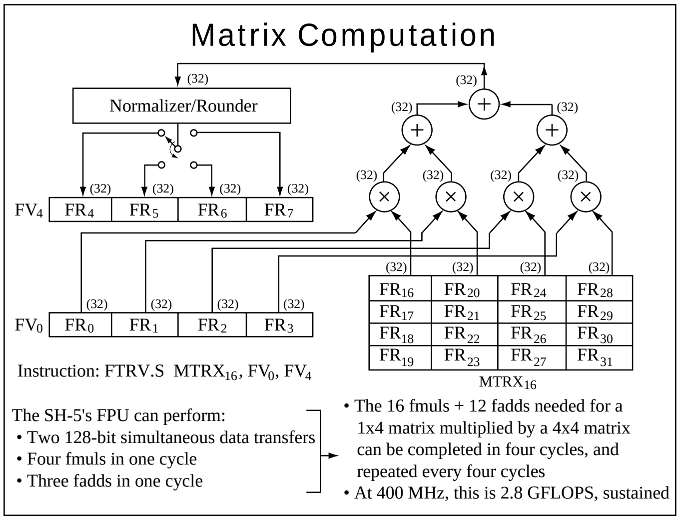

The optional FPU in the SH-5 can perform matrix operations at a sustained rate of 2.8 GFLOPS.

At its 400 MHz clock speed, the SH-5 architecture with the FPU achieves 2.8 GFLOPS performance, sustained.

Products with and without the FPU are planned by both Hitachi and ST.

2.5(e) SuperHyway bus

The bus interface unit in the SH-5 CPU core connects to the 64-bit, very high bandwidth on-chip SuperHyway bus—an essential feature for successful SOC integration. The 200-MHz (half clock speed) pipelined bus is compatible with the Virtual Socket Interface (VSI) protocols. It seamlessly connects the libraries of VSI virtual components offered by Hitachi and ST.

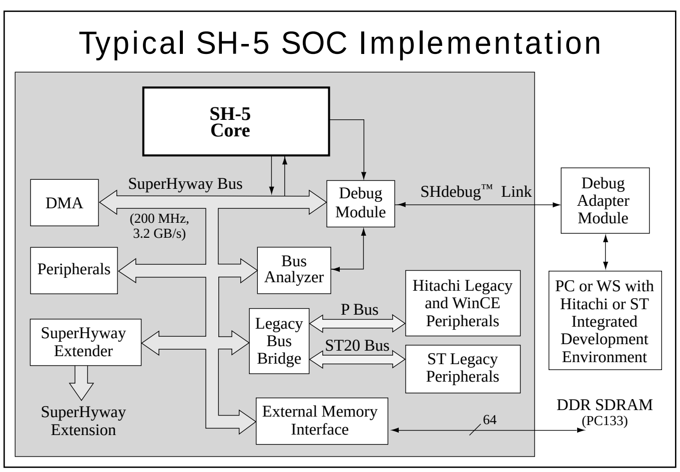

The SH-5 core can be embedded in a wide range of SOC designs, with the SuperHyway bus linking the required VSI virtual components.

Physically, the SuperHyway bus consists of dual 64-bit read/write buses. The 2x64-bit implementation is used in order to support full duplex operation—simultaneous 64-bit request/response transfers.

The SuperHyway bus uses a memory mapped, split-transaction protocol that is physically addressed and cache coherent. Transactions can contain up to 32 bytes of data, and traffic is directed by a DMA controller. At the 200 MHz peak bus speed, 128 bits (16 bytes) can be transmitted every cycle, so peak bandwidth is 3.2 GB/s (200 MHz x 64 bits x 2 buses).

Because the SuperHyway bus can pipeline requests, it can tolerate high-latency modules. For PCI peripherals, it supports a cache snoop protocol of the physical address space.

The SH-5 uses a 400-MHz internal clock source to drive the core. Timing signals for the on-chip function modules and the SuperHyway bus are derived from that master clock by circuits that maintain proper phase relationships between signals in the edge-triggered, static chip design.

2.5(f) SHdebug capabilities

For the SOC designs enabled by the SH-5 architecture, the on-chip SuperHyway bus cannot be accessed by external logic analyzers and other traditional debug tools, so powerful on-chip system debug support is essential. Advanced, non-intrusive debug capabilities support a complex debug feature set that is transparent to the target application software. The built-in debugging capabilities promote the rapid, efficient and effective development of reliable real-time systems using inexpensive external debugging tools, while also supporting the conventional on-chip debug tools required for some environments.

The CPU includes a watchpoint controller (WPC) and 12 programmable watchpoint channels. Chain latches support combinations of watchpoint conditions for filtering and conditional tracing. During continuous system operation, engineers can observe instruction execution, operand access, branch tracing, breakpoints, and single stepping. They can also gain insight into TLB misses, cache aliases, interrupts, operand accesses, pipeline freeze cycles, and more.

A programmable set of preconditions is available for all of the debug facilities. A programmable set of actions is also available. Possible actions include raise a debug exception, generate a trace message, alter performance counters, and alter the event counters. These features can be used for complex on-chip filtering of debug events. For example, branch tracing, bus tracing or performance monitoring can be initiated only within specific instruction address ranges, or when specific bus states occur. This type of filtering can be used to ensure that debugging is non-intrusive, and to make efficient use of trace bandwidth.

The on-chip debug module contains the debug links and a trace FIFO. The external debugger connects to either the JTAG port or the SHdebug link. The JTAG port conforms to the industry standard. It provides a very low cost control/observation port. The SHdebug link is a Hitachi/ST-specific innovation that has of a 9-pin interface. It provides real-time high-speed control and trace of the SH-5 system.

The JTAG port and SHdebug-link provide the same feature set. Therefore, debugger tools can scale across pin-limited connections (JTAG ports) or high-speed, wider interfaces (SHdebug links). Both of these interfaces are usable even when power-saving modes are activated.

The debug link gives an external debugging tool full access to the SuperHyway bus. Thus, debuggers can read or write to all addressable locations. This capability also allows the SH-5 to access external memory via the debug link. Therefore, debug code and data for the SH-5 can be held separately from the system's normal memory interfaces.

An on-chip trace message FIFO works in conjunction with the debug links to implement deep tracing of on-chip CPU and SuperHyway states. Various trace modes are available (wrap, trace hold and RAM FIFO), as is time stamping (optional).

A bus analyzer, another ST/Hitachi innovation, allows the observation of specific transactions on the SuperHyway bus. The analyzer provides the same features as the CPU debug facilities (preconditions and actions such as raise a debug exception, generate a trace message, alter performance counters, etc.). This system debug feature makes possible an entirely new class of system-on-chip debug functionality. On-chip states can be observed/detected for functional and timing related debugging.

2.5(g) Power-saving modes

The architecture has four operating modes: normal mode and three power-saving modes (2 standby modes and a CPU sleep mode). Key circuits such as the cache and the clock distribution system are specially designed for power efficiency. The CPU core, including the cache, TLBs, and SuperHyway bus,

Power Dissipation vs. Mode

		Status of On-Chip Functions			
	Mode	**CPG**	**CPU**	**On-Chip Memory**	**Peripheral Modules**
	Normal	operating	operating	operating	operating
Power down	**Sleep**	operating	halt	hold	operating (DMA: halt)
Power down	**Standby**	halt	halt	hold	specified module halt
Power down	**Module Standby**	operating	operating	operating	specified module halt

To maximize battery life in portable products, the SH-5 architecture has three power-down modes.

dissipates <800 mW. With the FPU, dissipation is <1000 mW.

2.5(h) Process technology

Hitachi and ST will build SH-5 chips using the latest production process: a jointly-developed 0.15-µm, 6-layer copper metal CMOS technology. The process accommodates the SH-5 core, plus extensive libraries of Hitachi and ST legacy peripherals. Thus, a broad range of standard and application-specific products can be fabricated.

2.6 System features

The SH-5 architecture is designed for efficient execution of applications written in C/C++ and Java. It has the features that are needed to work with the latest embedded operating system kernels, including the Windows CE, JavaOS, pSOS, VxWorks, Linux and OS-9 products.

The architecture includes a memory management unit (MMU) and has both user and privilege modes. There are three programmable vector base registers for reset, interrupt handling and trap functions. A separate debug vector enables the non-intrusive debug capability.

Chip Characteristics

	CPU Core	CPU Core, Cache/TLB and SuperHyway bus	CPU Core Cache/TLB, SuperHyway bus and FPU
Power Dissipation (mW)	<400	<800	<1000
Chip Area (mm^2)	3	11	14
Process Technology (0.15-µm CMOS)	• Copper metal • Power supply voltage = 1.5 V • Frequency = 400 MHz		

The SH-5 architecture is implemented in an advanced process jointly developed by Hitachi and ST to produce small-size chips that offer high performance and low to moderate power dissipation.

To implement sophisticated control systems, the CPU supports 16 levels of interrupt priority and provides a nonmaskable interrupt (NMI). For improved performance, the SH-5 architecture uses separate offsets for interrupts and TLB misses.

Various CPU mechanisms are provided to improve the performance of exception handling, interrupt handling and context switching:

- Two 64-bit control registers are provided for the exclusive use of the operating system. Typically they used to improve the performance of entry and exit code sequences for exception and interrupt handlers. Additionally, software conventions may be used to reserve general-purpose registers for use by the kernel.
- The SH-5's Applications Binary Interface (ABI) provides one 64-bit control register that can be used by the kernel to hold a temporary value.
- The floating point unit can be disabled. This allows a kernel to optimize context switches for threads of execution that do not require floating point operations. In particular, if either zero threads or exactly one thread uses floating point operation, then no context saving is needed for the floating point state.
- The CPU maintains "dirty" bits for the general-purpose and floating-point register sets. One dirty bit is used for each group of 8 consecutive registers, so there are 8 dirty bits for the general-purpose registers and another 8 dirty bits for the floating point registers. A dirty bit is set when there is a write to one of the registers in its group. An operating system can use this information to optimize context switches. For example, if a thread hasn't been written to a register group since that group was last context switched in, then those registers need not be saved the next time it is context switched out.
- The amount of register state required to execute an SHcompact thread of execution is a small subset of the full SHmedia state. If a program uses the SHcompact mode exclusively, then only the SHcompact-visible register state has to be saved.

These optimizations can drastically reduce the number of instructions and amount of memory bandwidth required for a context switch.

The mechanism used for interrupts and exceptions—an evolution of the mechanism used in the SH-4 series—facilitates system design. It provides 16 levels of interrupt priority and a single nonmaskable interrupt. All maskable interrupts can be ignored, and all exceptions caused by instruction execution are precise. A "panic" mode saves the processor state for debugging, and all traps and interrupts are vectored to one of seven locations. Debug exceptions and interrupts are separated to allow ICE-type debug support.

The SH-5 architecture implements a flat address space with a 32-bit virtual address range and simple address modes (indexed and scaled). Support is provided for signed/unsigned loads of 8/16 bit words; signed load of 16-bit long words and 64-bit quad words; store of byte, word, long and quad words; instructions for unaligned memory access for long and quad words; and little and big endian operation.

2.7 Performance summary

The SH-5 SuperH architecture raises the industry-leading SuperH processor family to new levels of performance. The SH-5 CPU core, running at 400 MHz at 1.5 V, delivers excellent general purpose, multimedia and floating point performance. The performance ratings of the architecture are summarized in the table below.

Performance Achievements at 400 MHz

Dhrystone Benchmark	• v1.1 • v2.1	714 MIPS 604 MIPS
Multimedia	• Multiply, add, accumulate • Math operations	1.6 GMACS (16-bit integer) 3.2 GOPS (16-bit MulAdd) 9.6 GOPS (8-bit SumAbsDif)
Floating Point	• Single precision	2.8 GFLOPS ((1x4)*(4x4) matrix)
SuperHyway Bus Bandwidth		3.2 GB/s
Power Efficiency*		1000 MIPS/W

* CPU core only, 400 MHz, 1.5V

Performance of the SH-5 is outstanding when considered in combination with the cost-effectiveness and power efficiency it offers.

2.8 Software development

Software development tools with common interfaces will be available from Hitachi, ST and third-party suppliers. A common ABI and ELF/DWARF format ensure that the binary files produced by any of the compilers can be linked without modification. Compiler optimizations include hoisting prepare target instructions (for split instructions) and pipeline-optimized scheduling, including SIMD.

A library of SIMD functions allows software engineers to avoid assembly language programming and write SIMD code in C without having to analyze register allocations. Instead, a statement like the following can be used:

```
v5 = msubs_w(mmulfxrp_w(v1,v2), mmulfxrp_w(v3,v4))
```

Here, the C variables v1–v5 each represent four 16-bit values.

To help customers model their SH-5 based SOC designs and achieve "right first time" implementations, a complete toolchain is available from Hitachi, ST and third-party vendors. For example, SuperHyway models are available from Hitachi and ST for the emulators offered by Synopsis and MetaSystems.

III. The First SH-5 Product

3.1 Overview

The first product with the SH-5 CPU core co-developed by Hitachi and STMicroelectronics will be a device designed for use in development systems, application reference platforms for third-party developers and customers, and customer products, including handheld products. This device will be included in evaluation platforms that will allow benchmark tests and will facilitate system design.

Both Hitachi and ST plan to use the first SH-5 architecture product to develop derivative chips and application-specific products. It will form the basis of Hitachi's SH8000 series and ST's ST50 series.

First Product with the SH-5 Architecture

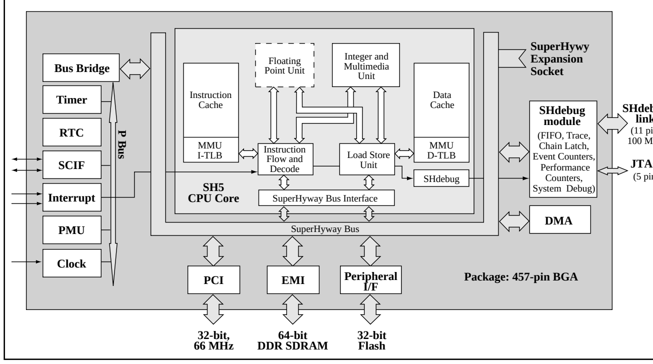

The first product to use the SH-5 architecture will be a chip designed primarily to aid the development activities of customers and third-party support suppliers.

3.2 Hardware details

The first SH-5 product will give system engineers an extensive array of hardware functions, including

- a 400-MHz SH-5 CPU core chip that includes a standard debug module with special extensions for system debug
- a 3.2 GB/s SuperHyway bus
- a standard set of peripherals:
  - a UART serial interface with DMA, 16-bit FIFO, and software-configured baud clock generator
  - an interrupt controller with programmable priorities that allows up to 16 interrupts
  - three timers: a watchdog timer and two 32-bit timers with auto-reload, configurable inputs, and interrupts
  - a low-power real-time clock that includes calendar, alarm and IRQ functions, controlled by a controller with a software-programmable PLL
  - a power management controller
  - a 16-bit programmable I/O that responds to level-sensitive or edge-sensitive interrupts

- a PCI v2.1, 32-bit 33/66-MHz interface that allows bus mastering to main memory and supports 4 external bus masters; this interface allows users to connect many standard peripherals to the SH-5 core before ASICs or application-specific companion chips are developed
- an interface to external SDRAM and DDR DRAM devices, capable of handling 16/32/64-bit data at up to 133 MHz speeds; this I/F supports 4 open banks and SDRAM self-refresh for standby mode
- a DMA controller with 4 programmable channels
- an interface to up to 64 MB of Flash memory or ROM; this I/F handles 8/16/32-bit data and a 26-bit address

In addition, the first SH-5 product will offer an extensive range of debug capabilities:

- breakpoint (BRK instruction) channels
- single-step (BRK-STEP) channel
- four instruction-address range (IA) channels
- two operand-address range (OA) channels
- two instruction-value (IV) channels
- a branch (BR) channel,
- two CPU-performance (WPC_PERF) channels
- two bus-analyzer (SHWYBA) channels
- fast printf (for application/RTOS instrumentation)
- resources for CPU-mode and bus-state preconditions:
  - four 48-bit performance counters
  - four 16-bit event counters (for counter preconditions)
  - seven chain latches (so debug facilities can be combined for complex on-chip filtering).
  - 1K debug message FIFO

3.3 Chip production

Samples and development platforms of the first SH-5 based products are expected to be available in the fourth quarter of 2000. Volume production is expected to begin mid-year 2001.

3.4 Summary

The Hitachi/STMicroelectronics partnership has extended the SuperH architecture to 64 bits. The fifth-generation RISC engine provides the outstanding balance of performance, features, power dissipation, and cost-effectiveness needed to enable new generations of products in price-sensitive markets. The co-developed CPU core, produced with a new 0.15-µm production process, establishes a new standard for embedded systems design. The SOC methodology inherent in the SH-5 architecture, plus the software interoperability and libraries of peripheral IP modules shared between Hitachi and ST, will be used to populate the SH8000 and ST50 product lines with optimized silicon solutions for a wide range of applications.

IV. Appendix

Overviews of products that use the SuperH architecture

4.1 Hitachi's SH-4 series and ST's ST40 series

The two devices in Hitachi's SH-4 series and ST's ST40 series are based on a 2-issue superscalar, 32-bit SuperH RISC core. They offer a 16 KB data cache, 8 KB instruction cache, 64/32/16/8-bit external bus and 2 GB address space. The 200/167-MHz chips include a powerful FPU with a 128-bit vector graphics engine optimized for 3D graphics that can process up to 7 million polygons per second. They also have a 32-bit MAC, MMU, SDRAM I/F, and user break controller (UBC), plus an extensive on-chip debugging capability.

The SH-4 and ST40 series devices achieve 360 MIPS for the Dhrystone 2.1 benchmark. One version has a PCI interface for easy connectivity to standard peripheral products.

4.2 Hitachi's SH-3 series

Devices in the SH-3 series (8 versions, total, including one in the SH3-DSP series) have a 16 KB cache, 32/16/8-bit external bus and 448 MB address space. The 133/66-MHz chips have up to 16 KB of cache, a MMU, 32-bit MAC, barrel shifter, real-time clock and UBC. They provide performance levels up to 168 MIPS. Special features of various SH-3 versions include serial, PCMCIA, SmartCard, and IrDA interfaces, PLL and a JTAG serial debug interface (SDI).

The 133/66-MHz SH3-DSP device combines an SH-3 series 32-bit RISC CPU and a full-featured, 16-bit integer DSP unit into a powerful, multitasking core. The device has a four-bus structure, 16-KB cache and 16-KB X/Y RAM. It executes all software from one instruction stream and can perform a 16-bit multiply in a single cycle. The chip can shift its operation from a 168-MIPS RISC device to a DSP that can sustain 266 MOPS (532 MOPS, peak), or to any combination in between.

4.3 Hitachi's SH-2 series

Devices in the SH-2 series (19 versions, total, including 3 in the SH-DSP series) have a 4 KB cache, 32/16/8-bit external bus and 128 MB address space. The 66/33-MHz chips, which have up to 4 KB of cache and a 32-bit MAC, provide greater functionality and higher performance (up to 78 MIPS) than SH-1 series chips. Special features of various SH-2 versions include up to 512 KB of Flash, timers suitable for motor control, a bus state controller and CAN 2.0B ports.

The SH-DSP devices are ideal for systems that previously required both an embedded processor and a DSP chip. They combine an SH-2 series 32-bit RISC CPU and a full-featured, 16-bit integer DSP unit into a powerful, multitasking core that uses three buses, a 4 KB cache and 16 KB X/Y RAM to achieve high throughput. Versions have up to 256 KB of on-chip Flash.

4.4 Hitachi's SH-1 series

MCUs in the SH-1 series (8 versions, total) achieve up to 20 MIPS performance. The 20-MHz devices have a 16/8-bit external bus and 32 MB address space. Versions offer up to 8 KB RAM/64 KB ROM, a glueless interface to SRAM and DRAM, a 16-bit MAC, and many peripherals, including special timers for motor control, serial channels, DMA circuits, UBC and more.

4.5 SH-5 instruction set (preliminary data)

Flow control instructions	*Summary*
PTA,PTB	prepare target immediate, target is SHmedia or SHcompact
PTABS, PTREL	prepare target absolute/relative register
B[EQ|EQI|GE|GEU|GT|GTU|NE|NEI]	conditional branches
BLINK	unconditional branch (optional link and mode switch)
GETTR	move from target register
Integer instructions	
MOVI,SHORI	move immediate, shift then or immediate
ADD[.L],ADDI[.L],ADDZ.L	add register/immediate 32/64-bit, with zero-extend 32-bit
SUB[.L]	subtract 32/64-bit
MULU.L	multiply full 32-bit x 32-bit to 64-bit unsigned
CMP[EQ|GT|GTU],CMV[EQ|NE]	compares, conditional moves
AND,ANDI,ANDC,OR,ORI,XOR,XORI	and, and-complement, or, xor (and immediate forms)
SHARD[.L],SHARI	shift arithmetic right dynamic 32/64-bit, immediate 64-bit
SHLLD[.L],SHLLI[.L],SHLRD[.L],SHLRI[.L]	shift logical left/right dynamic/ immediate 32/64-bit
BYTEREV,NSB,NOP	byte reversal, count number of sign bits, no operation
Memory instructions	
LD.[BLWQ],LDX.[BLWQ]	load displacement/indexed 8/16/32/64-bit signed
LD.U[BW],LDX.U[BW]	load displacement/indexed 8/16-bit unsigned
LDHI.[LQ],LDLO.[LQ]	load misaligned high/low part 32/64-bit
ST.[BLWQ],STX.[BLWQ]	store displacement/indexed 8/16/32/64-bit
STHI.[LQ],STLO.[LQ]	store misaligned high/low part 32/64-bit
SWAP.Q	atomic swap in memory 64-bit
ICBI,PREFI	instruction cache block invalidate/prefetch
ALLOCO,OCBI,OCBP,OCBWB	operand cache block allocate/ invalidate/purge/write-back
SYNC[IO]	synchronize instructions or operand data
Multimedia instructions	
MABS.[WL]	absolute signed 16/32-bit with saturation
MADD.[WL],MADDS.[UB|W|L]	add 16/32-bit, add 8/16/32-bit with saturation
MCMPEQ.[BWL],MCMPGT.[UB|W|L]	compare equal/greater-than 8/16/32-bit
MCMV	bitwise conditional move
MCNVS.[WB|WUB|LW]	convert word-to-byte/word-to-ubyte/long-to-word
MEXTR[1234567]	extract 64 bits from 128 bits from byte number
MMACFX.WL,MMACNFX.WL	fractional multiply and accumulate/subtract signed 16-bit
MMUL.[WL]	multiply 16/32-bit
MMULFX.[WL],MMULFXRP.W	fractional multiply 16/32-bit, with round nearest +ve
MMULHI.WL,MMULLO.WL	full multiply signed 16-bit high/low parts
MMULSUM.WQ	multiply and sum signed 16-bit
MPERM.W	permute 16-bits
MSAD.UBQ	sum of absolute differences of unsigned 8-bit
MSHALDS.[WL]	shift arithmetic saturating-left/right 16/32-bit
MSHARDS.Q	shift arithmetic right, saturation to signed 16-bit
MSHLLD.[WL],MSHLRD.[WL]	shift logical left/right 16/32-bit
MSHFHI.[BWL],MSHFLO.[BWL]	shuffle upper/lower-half 8/16/32-bit
MSUB.[WL],MSUBS.[UB|W|L]t	subtract 16/32-bit, 8/16/32-bit with saturation

Floating point instructions	
FABS.[SD],FNEG.[SD],FSQRT.[SD]	absolute/negate/square-root of single/double
FADD.[SD],FSUB.[SD]	add/subtract two single/double
FCMP[EQ|GE|GT|UN].[SD]	compare single/double for =, >=, >, unordered
FCNV.[DS|SD]	double-to-single/single-to-double conversion
FGETSCR,FPUTSCR	move from/to floating-point status/control register
FLOAT.[LD|LS|QD|QS]	long-to-double/long-to-single/quad-to-double/ quad-to-single convert
FMOV.[S|D]	single-to-single/double-to-double move
FMOV.[SL|DQ|LS|QD]	single-to-long/double-to-quad/long-to-single/ quad-to-double move
FMUL.[SD],FDIV.[SD]	multiply/divide two single/ double
FTRC.[DL,SL,DQ,SQ]	double-to-long/single-to-long/double-to-quad/single-to-quad convert
Special-purpose floating point instructions	
FMAC.S	fused multiply accumulate
FIPR.S/FTRV.S	vector dot-product, transform vector by matrix
Floating point memory instructions	
FLD.[SPD],FLDX.[SPD]	load displacement/indexed 32/2x32/64-bit value
FST.[SPD],FSTX.[SPD]	store displacement/indexed 32/2x32/64-bit value
System control/configuration instructions	
BRK,TRAPA,RTE, SLEEP	cause a debug-exception/trap, return from exception, sleep
GETCFG,PUTCFG,GETCON,PUTCON	move from/to configuration/ control register

Hitachi document: PMH15IW001D1
ST document: SH-5/ST50-WP