

M I C R O P R O C E S S O R

www.MPRonline.com

THE INSIDER'S GUIDE TO MICROPROCESSOR HARDWARE

PHILIPS POWERS UP FOR VIDEO

Superpipelined TriMedia Processor Core Gives New Legs to MPEG

By Tom R. Halfhill {11/3/03-01}

With an eye on the growing market for consumer electronics, Philips Semiconductors announced a new TriMedia 32-bit processor core at **Microprocessor Forum 2003**. The swifter core will debut next year in Philips media processors destined for personal video

recorders, wireless networks, high-definition TVs, and other audio/video products.

Unlike some previous TriMedia CPU cores, the new TM5250 won't be offered as licensable intellectual property (IP). Philips shut down its TriMedia Technologies IP-licensing business last year. (See *MPR 5/19/03-03*, "TriMedia Comes Home.") Instead, the TM5250 will spawn a new generation of standard-part Nexperia media processors designed and manufactured by Philips.

The TM5250 is source-code compatible with processors based on TriMedia's DSPCPU32 architecture, which dates to 1994. (Actually, the architecture has an even longer history that stretches back to the Philips LIFE project in

1987; see *MPR 12/5/94-03*, "Philips Hopes to Displace DSPs with VLIW" and *MPR 8/8/90*, "Philips Gives LIFE to VLIW.") Source-code compatibility provides a migration path for current customers of TriMedia 32-bit processors in the TM1000, TM1100, and Nexperia PNX1300 families. The TM5250 is not compatible with the 64-bit DSPCPU64 architecture announced in 1998 (see *MPR 10/26/98-07*, "Philips Advances TriMedia Architecture"), but the incompatibility doesn't matter, because Philips never introduced an implementation of that architecture. Instead, Philips has decided to improve the performance of the 32-bit architecture while maintaining compatibility with existing TriMedia processors.



The Philips TriMedia design team is based in Sunnyvale, California. From left: Tom Hampton, Bill Zhong, Bart Vostermans, Jan-Willem van de Waerd (chief architect), Gerrit Slavenburg (senior fellow), Pedro Rodrigues, Kulbhushan Kalra, Sebastien Mirolo, Jean Paul van Itegem, Carlos Basto, Robert de Gruijl, Sanjeev Das, Chris Yen, Balaji Ramalingam, Dinesh Amirtharaj. Not pictured: Otto Steinbusch, Sharon Zhong, Jan Stuyt, Emiel Peys, Duy Nguyen.

New features in the TM5250 include superpipelining, nine new instructions for video processing, improved L1 caches, an integrated L2 data cache, better prefetching, and a larger array of function units that can keep as many as 29 pending instructions in flight. Table 1 compares the TM5250 to the TriMedia CPU core in the PNX1300.

The TM5250 is a fully synthesizable CPU core based on a standard-cell logic library and standard, single-ported SRAMs for the caches. Philips expects the CPU to run at 500–700MHz and occupy 19.8mm² of silicon when fabricated in a standard TSMC 0.13-micron six-layer-metal process (not the faster, higher-leakage LV process). According to estimates obtained from a Philips gate-level simulation tool and Synopsys PowerTheater, the TM5250 will consume less than 2mW per megahertz, or about 1W at a

conservative target clock speed of 500MHz. A chip based on the new CPU core will be announced early next year.

Superpipeline Compensates for Synthesis

To improve performance with a synthesizable core that lacks the efficiency of a full-custom layout, Philips deepened the TM5250's basic pipeline for simple integer instructions to 11 stages. That compares with five stages in the CPU core of the PNX1300, the most recent TriMedia processor. For some operations, such as floating-point addition, subtraction, and multiplication, the TM5250's pipeline runs as deep as 16 stages. Figure 1 shows the pipelines for various types of operations.

With 29 function units at its disposal, the TM5250 is a very wide machine—appropriate for a descendant of the first VLIW architecture ever implemented as a microprocessor. Most numerous are the four simple ALUs and constant units, which execute the most frequently used integer-arithmetic and Boolean-logic instructions in one cycle. There are three regular ALUs and three DSP ALUs for integer operations that involve rounding and clipping. Additional function units include shifters, DSP multipliers, floating-point ALUs, floating-point multipliers, a floating-point divider, a floating-point compare unit, branch units, and a load/store unit.

One VLIW instruction bundle may contain up to five operations, including as many as three branches. To avoid the deeper-pipeline penalty of taking a branch—at least nine clock cycles—the TM5250 is the first TriMedia processor that has dynamic branch prediction. It predicts which cache way holds the branch target code, saving two clock cycles for each correctly predicted taken branch. Like previous TriMedia processors, the TM5250 also avoids branch

Feature	TriMedia TM5250	Nexperia PNX1300
Architecture	DSPCPU32	DSPCPU32
Architecture Width	32-bit	32-bit
Core Freq (Worst)	500MHz	200MHz
Memory Bus Freq	200MHz	200MHz
Effective Bus Freq*	400MHz	200MHz
Pipeline Depth	11–16 stages	5–7 stages
Function Units	29 total	27 total
Constant	4	5
Shifter	3	2
Simple ALU	4	—
ALU	3	5
Branch	3	3
DSP ALU	3	2
DSP Multiplier	2	2
Floating-Point ALU	2	2
Floating-Point Mul	2	2
Floating-Point Div	1	1
Floating-Point CMP	1	1
Load/Store	1	2
IEEE-754 Floating Point	Yes	Yes
SIMD Capabilities	1 x 32-bit 2 x 16-bit 4 x 8-bit	1 x 32-bit 2 x 16-bit 4 x 8-bit
L1 Cache (I/D)	64K/16K	32K/16K
L2 Cache	128K	—
L2 Cache Freq	250MHz	—
Process	0.13µm 6LM	0.25µm 4LM
Die Size	19.8mm ²	35mm ²
Core Voltage	1.2V	2.5V
Power (Typical)	1W**	2.5W
MediaStone Score	425	200
EEMBC ConsumerMark [†]	51.3 OTB	23.3 OTB
EEMBC ConsumerMark [†]	284.6 OPT	110 OPT
Availability	Early 2004	Now

Table 1. The new TriMedia TM5250 has significant enhancements over the TriMedia core in the Nexperia PNX1300 (formerly known as the TriMedia TM1300). *The TM5250 uses DDR-SDRAM memory on a 200MHz bus, so the effective bus frequency is 400MHz. **Philips estimate. †Certified EEMBC ConsumerMark scores are based on out-of-the-box (OTB) and optimized (OPT) tests. MPR has extrapolated the simulator-based TM5250 raw scores to the 500MHz target clock frequency.




Figure 1. An unusual feature of the TM5250 pipeline is the two stages required for way prediction when fetching instructions from the eight-way set-associative instruction cache. Following those stages are four stages for actually fetching and decompressing the VLIW instructions; three stages to decode the instructions, access the register file, and perform operand bypassing; one to six execution stages; and a final writeback stage. Note that a nonpipelined function unit carries out floating-point divides and square roots.

penalties by permitting any instruction to be guarded (conditional). This allows the compiler to statically eliminate some branches at design time.

The single-ported SRAMs in this design—optimized for area and power consumption, not for speed—can't run at the full clock frequency of the processor, so Philips had to redesign the cache subsystem. Consequently, the instruction cache runs at only half the core frequency, requiring two clock cycles to access. However, the fetch bandwidth is the same as in the PNX1300, because the TM5250 fetches twice as many bits per clock cycle: 512 bits vs. 256. In addition, Philips doubled the size of the TM5250's instruction cache over that of the PNX1300: 64KB vs. 32KB. The cache is eight-way set-associative and is organized into 128-byte lines. To preserve deterministic behavior in real-time applications, programmers can lock as much as half the cache, line by line.

Allocating two cycles for cache access should make it easier to port the synthesizable model to a future 90-nanometer (nm) process. Standard-cell logic generally scales better than compiled memory at smaller geometries, so as the performance gap between the logic and memory grows, the two-cycle cache will be less susceptible to timing problems.

Another aspect of the redesigned instruction cache is the effect on the pipeline: two extra stages are devoted to way prediction and cache-tag lookup. Two clock cycles are required to check the cache tags and another two cycles to look up the compressed instruction information. (The VLIW instructions are compressed to save space in the cache.) As long as the TM5250 keeps fetching instructions sequentially, the extra stages will probably have little effect on performance. Branches will levy a penalty, of course, which is the reason the TM5250 has dynamic branch prediction and conditional execution. All things considered, Philips believes the redesigned instruction cache will sustain the processor's higher performance while reducing the cache's power consumption.

Data Caches Revamped, Too

The TM5250's data pipe received even more attention from the designers than the instruction cache, because media processing is typically a data-intensive task. The L1 data cache is only 16KB—one-fourth the size of the instruction cache—so it can run at the full clock frequency of the core when implemented with the same SRAM memories as the instruction cache. It is four-way set-associative, organized into 64-byte lines, and can transfer enough data to refill a cache line in one clock cycle.

In addition, the TM5250 is the first TriMedia processor to have an L2 cache. The 128KB cache runs at half the CPU frequency, is eight-way set-associative, and is organized into 128-byte lines. The bus interface to main memory is 64 bits wide, but there are 512-bit interfaces between the L1 and L2 caches and the refill, copy-back, and prefetch

units. An L1 cache miss that hits the L2 cache incurs a penalty of four or five clock cycles.

Some unusual features improve the performance of the data caches. For one thing, the L1 cache uses an allocate-on-write-miss policy instead of the more common fetch-on-write-miss policy. It keeps track of modified data at byte resolution, not line resolution, so the TM5250 can avoid unnecessary fetches from memory after a write operation that misses the cache. This cache design should be especially suitable for processing streaming data with audio/video codecs, because the processor rarely has to modify the data stream later.

Another new feature is that the TM5250 suffers no read/write penalty for using unaligned data. The load/store unit can read and write aligned or unaligned data at the same speed—a useful feature for a SIMD media processor that handles many different data types. A 32-bit load instruction can fetch four consecutive elements of eight-bit video data on any byte alignment in memory. Other CPU architectures can fetch unaligned data, of course, but they often require multiple instructions to fetch multiple elements. Prime examples are the patented LWL (load word left) and LWR (load word right) instructions in the MIPS architecture.

A related feature allows software programmers to define “memory prefetch regions” for automatically loading noncontiguous data in memory. Programmers can define up to four prefetch memory regions and stride lengths. When a load instruction misses the L1 data cache, and the target memory address falls within a defined region, the prefetch hardware automatically starts loading data into the L2 cache from the target address and stride offset.

For example, if the memory data represents a video image with 720 pixels per line, the prefetch hardware can begin loading the first data elements for the second line of the image while the program is operating on data elements for the first line of the image. This mechanism—in combination with the ability to read and write data on any byte alignment without penalty—allows the TM5250 to optimize memory accesses with streaming media codecs, no matter how the data is configured in memory.

To boost throughput, the TM5250 can use double-data-rate (DDR) SDRAM on its 200MHz memory bus, yielding an effective bus frequency of 400MHz. The existing PNX1300 uses regular SDRAM on a 200MHz memory bus, which provides only half as much bandwidth. However, the memory latency with both types of memory is essentially the same, and latency affects CPU performance more than bandwidth does. The real-world improvement is therefore somewhat less dramatic than the difference in bus frequencies suggests.

Analyzing the EEMBC Benchmarks

To measure the performance of the new CPU core and compare it with other processors, Philips used a Quickturn

machine to run the EEMBC consumer benchmark suite on a cycle-accurate RTL simulation of the TM5250. Philips already has EEMBC consumer benchmark results for another TriMedia processor, the 166MHz TM1300. All EEMBC scores have been certified by EEMBC Certification Labs (ECL).

In addition, Philips ran its own MediaStone benchmark suite on the TM5250 simulation and on a 200MHz TriMedia PNX1300 processor. The Philips MediaStone suite includes three tests from the EEMBC consumer suite plus additional consumer-oriented algorithms, such as AC3 audio decoding, MPEG2 audio/video decoding, motion-compensated de-interlacing, motion estimation, and film detection.

For this early round of benchmarking, Philips didn't optimize either the EEMBC or MediaStone source code for the TM5250. (EEMBC rules allow vendors to rewrite the source code and use other techniques when measuring the optimized "full-fury" benchmarks.) Instead, Philips used the EEMBC and MediaStone source code that was previously optimized for the PNX1300. That means the early benchmark tests don't take advantage of the new instructions and features in the TM5250, such as unaligned load/store instructions, new MPEG-4/H.264 instructions, and the improved prefetching scheme. (Philips won't disclose all the new instructions in detail until next year.) Therefore, the EEMBC and MediaStone scores reported here probably represent the lower range of the TM5250's actual performance.

Under EEMBC rules, vendors must report benchmark scores for simulated processors as the number of loop iterations performed at 1MHz. Vendors must also specify a realistic target clock frequency for their simulated processor.

MPR has extrapolated the raw 1MHz scores to those target frequencies—500MHz in the case of the TM5250.

Results: the TM5250's EEMBC ConsumerMark scores are exceptional. The unoptimized or "out-of-the-box" ConsumerMark (51.3) beats the published scores of every other simulated or silicon processor in this suite. The optimized full-fury ConsumerMark (284.6) is surpassed only by the Tensilica Xtensa V, a configurable processor whose instruction set was customized for the EEMBC tests. Two other configurable processors—Tensilica's older Xtensa III and ARC International's ARCTangent-A4—were the only other simulated CPU cores that came close to the TM5250's optimized ConsumerMark.

The fastest processor in production silicon with a published ConsumerMark score is Motorola's MPC7447, a Power PC chip. It posted an optimized ConsumerMark of 172.2, well below the TM5250's score, despite the advantage of running at a much higher clock frequency (1.3GHz vs. 500MHz). Table 2 shows the detailed results of these benchmarks.

However, there are some interesting anomalies in the EEMBC results. The TM5250's strong ConsumerMark scores rely largely on its superior ability to compress and decompress JPEG files. The optimized scores for the three configurable processors—Xtensa III, Xtensa V, and ARCTangent-A4—don't come close to matching the TM5250 in those two tests. A big reason for their shortfall is that, in two cases, they weren't really trying.

ARC made no special effort to optimize its processor for the JPEG tests, although the optimized JPEG scores did improve a trifle over the out-of-the-box scores, owing to side effects from other optimizations. Tensilica made no attempt at all to optimize Xtensa III for the JPEG tests, which is why the "optimized" JPEG scores are identical to the out-of-the-box

	TriMedia TM5250	TriMedia TM1300	Nexperia PNX1300	Tensilica Xtensa V	Tensilica Xtensa III	Motorola MPC7447	ARC ARCTangent-A4	SuperH SH-4
Clock Speed*	500MHz	166MHz	200MHz	260MHz	200MHz	1.3GHz	150MHz	202MHz
Implementation	Simulation	Chip	Chip	Simulation	Simulation	Chip	Simulation	Simulation
EEMBC Consumer Suite (Out of the Box)**								
Compress JPEG	20.0	11.5	—	12.5	7.9	—	6.5	10.1
Decompress JPEG	40.0	14.3	—	17.8	9.9	—	7.7	12.1
Gray-Scale Filter	275.0	115.9	—	117.4	46.2	—	27.3	70.7
RGB to CMYK	255.0	125.2	—	130.7	76.6	—	63.3	117.2
RGB to YIQ	320.0	146.7	—	88.0	37.2	—	24.5	34.3
ConsumerMark	51.3	23.3	—	22.6	11.5	—	8.4	14.7
EEMBC Consumer Suite (Optimized)**								
Compress JPEG	125.5	65.2	—	57.2	7.9	85.3	6.8	34.3
Decompress JPEG	137.0	76.9	—	83.2	9.9	103.9	8.0	40.4
Gray-Scale Filter	2,125.0	759.4	—	7,014.8	4,884.0	907.3	4,768.1	179.8
RGB to CMYK	1,765.0	482.3	—	6,762.6	5,202.0	992.0	3,904.1	181.8
RGB to YIQ	1,465.0	444.2	—	9,014.2	6,936.0	960.1	5,198.7	135.3
ConsumerMark	284.6	110.0	—	525.9	193.6	172.2	159.3	41.4
Philips MediaStone Benchmark								
Composite	425	—	200	—	—	—	—	—

Table 2. This table shows the highest achievers in EEMBC's consumer benchmark suite. The TriMedia TM5250's ConsumerMark scores beat every other processor except Tensilica's customized Xtensa V. *Clock speeds for processors running in simulation are the vendors' target clock frequencies. **MPR extrapolated the EEMBC scores by multiplying the vendors' target clock frequency by the raw 1MHz simulator scores certified by ECL.

scores. The relatively low JPEG scores for the ARCTangent-A4 and Xtensa III pulled down their composite Consumer-Marks. (ECL derives the composite score from a geometric mean of the individual scores, adjusted by a normalization factor.)

Tensilica did optimize Xtensa V for JPEG, obtaining significant gains in performance. In fact, the raw (iterations per megahertz) JPEG scores for Xtensa V approximately match or exceed the raw JPEG scores for the TM5250. However, Xtensa V loses the JPEG race because the TM5250's clock rate is nearly twice as fast: 500MHz vs. 260MHz. That accounts for the large differences between our extrapolated JPEG scores in Table 1.

In the other three EEMBC consumer tests—applying a high-pass gray-scale filter to an image and converting between various color spaces—all three optimized configurable processors handily beat the TM5250, even when running at much lower clock speeds (only 150MHz for the ARCTangent-A4). This demonstrates the advantage of a configurable processor, but it doesn't necessarily take anything away from the TM5250, which wasn't optimized for these tests. Unfortunately, the EEMBC consumer suite doesn't include any MPEG tests, which would be a more relevant benchmark for the audio/video-oriented TriMedia CPU.

One factor in the TM5250's favor is that Philips was conservative when running the cycle-accurate RTL simulation for the EEMBC tests, timing the memory system at a 5:2 core-to-bus ratio. Most of the other synthesizable processors simulate the memory system running at the full core frequency.

Performance Scales With Clock Speed

To provide another view of the TM5250's performance—albeit without the imprimatur of EEMBC certification—Philips ran its own MediaStone benchmarks on the same RTL simulation for comparison with a 200MHz PNX1300. (As mentioned above, MediaStone includes three EEMBC consumer-suite tests plus some additional video tests.)

Result: the TM5250's MediaStone score is 425 at 500MHz, compared with 200 for the PNX1300 at 200MHz. That is 2.125 times the performance of the existing processor—at 2.5 times the CPU clock frequency. In other words, the MediaStone performance of the new processor core falls just short of keeping pace with its faster clock speed. This result would seem to question the value of the TM5250's architectural enhancements.

Bear in mind a few factors, however. Philips had no time to rewrite the MediaStone benchmarks to use the new instructions, prefetching features, and other enhancements in the TM5250. Huge gains are possible by optimizing a media codec for a specific microarchitecture. But that kind of fine-tuning is a highly skilled task, and it will be a while before Philips rewrites its extensive codec library.

Price & Availability

Although the TriMedia TM5250 is a fully synthesizable processor core, Philips no longer offers TriMedia cores for broad IP licensing. Instead, the TM5250 will be available in a standard-part media processor that Philips plans to announce in early 2004. For more information, visit www.semiconductors.philips.com/products/nexperia/.

Another factor is that direct clock-frequency comparisons between the TM5250 and PNX1300 are purely hypothetical. The older processor simply cannot run at 500MHz, even if fabricated in the same process as the TM5250, because its shorter pipeline is a bottleneck. And even if the PNX1300 could run at 500MHz, its slower memory interface, lack of an L2 cache, and less-sophisticated prefetching would give it handicaps the higher core frequency couldn't overcome.

Then, too, 500MHz is a conservative target for the TM5250. Philips thinks the core could reach 700MHz in a standard TSMC 0.13-micron process. That clock rate would push the composite MediaStone score to 565—again, without using the core's new features. In a higher-performance 0.13-micron or 90nm process, Philips thinks the TM5250 could attain 900MHz, a very respectable clock rate for a synthesizable media-processor core. As Figure 2 shows, that clock rate would boost the MediaStone score to 691.

More to the point, the TM5250 achieves Philips's design goals: it can simultaneously encode and decode MPEG-2 or MPEG-4 streams for personal video recorders; decode 720p Windows Media 9 or standard-definition H.264; and, with the assistance of some acceleration hardware, decode a high-definition MPEG-2 stream while simultaneously encoding standard-definition MPEG-2/MPEG-4 video or a small picture-in-picture frame of high-definition MPEG-2 video.

Nevertheless, it says something that, despite numerous enhancements at the architectural and microarchitectural




Figure 2. Estimated MediaStone scores. RTL simulations indicate the TM5250's performance will scale fairly well with increases in clock speed. Philips believes clock frequencies up to 700MHz are possible in a standard 0.13-micron process; faster frequencies will require a 0.13-micron LV or a 90nm process.

levels, the TM5250 strains to improve performance beyond what might be expected from the climb in clock frequency alone. It says that previous TriMedia designs are already well optimized for media processing, leaving relatively little low-hanging fruit to harvest, and that media processors scale less

easily than desktop processors. By introducing several new features that help maintain data throughput, Philips is doing well to keep the almost 10-year-old TriMedia architecture competitive in a market that continues to attract new and more-exotic architectures. ♦

To subscribe to Microprocessor Report, phone 480.609.4551 or visit www.MDRonline.com