

First Trimedia Chip Boards PCI Bus

VLIW Multimedia Engine Aimed at PCs, Set-Top Boxes

by Brian Case



Philips' forthcoming TM-1 processor will accelerate a variety of multimedia functions. At last month's Microprocessor Forum, Philips revealed that the first chip in the TriMedia family will contain a powerful VLIW core (see [081603.PDF](#)) and an array of autonomous DMA units, functioning as a complete real-time system.

TM-1 will ship with a multimedia repertoire similar to that of Chromatic's Mpact (see [091404.PDF](#)), including MPEG-1 and -2 (audio, video, and system decoding), 3D graphics (front- and back-end), V.34 data/fax modem, audio synthesis (FM, wavetable), and H.261 and H.320 video conferencing. Philips will offer Windows 95 drivers (see [0915ED.PDF](#)) for its chip but, unlike Chromatic, will not offer register-level Sound Blaster compatibility.

Since it is based on a general-purpose processor, TM-1 can operate as either a standalone CPU or an accelerator in a PC. To simplify system design in both cases, the chip includes glueless interfaces to local high-speed synchronous DRAM (SDRAM) and the standard PCI expansion bus. Though a complete signal-processing system, it is not a mixed-signal chip; where required, analog circuitry must be implemented externally.

The TM-1 chip can acquire multimedia input data streams (audio and video), store and process them in local memory, and output multimedia streams, all in real time. On-chip floating-point hardware speeds 3D graphics operations.

Volume production of TM-1 is planned for 4Q96 in 0.5-micron four-layer-metal CMOS; a shrink to 0.35-

micron CMOS will follow. Philips promises an eventual sub-\$50 chip price. The initial target frequency is 100 MHz; at that speed and 3.3 volts, the chip will dissipate 4 W. Fine-pitch EDQUAD or SuperBGA packages (see [091304.PDF](#)) will be offered.

Coprocessors Handle Basic Tasks

In the general case, a multimedia system must perform several concurrent tasks. For example, in a video phone, the system must input live video and audio, compress it, and send it to a remote system. At the same time, the system must receive compressed video and audio, decompress it, display the video, and play the sound.

The VLIW DSPCPU core is capable of performing all these tasks, but only the compression and decompression require the power and generality of the CPU. A very small amount of simple hardware can perform I/O functions. Furthermore, the high data rates encountered with high-quality MPEG-2 streams would leave a single shared CPU spending most of its time performing conceptually simple tasks and servicing I/O interrupts.

As shown in Figure 1, the DSPCPU has autonomous DMA units for data I/O. In addition, it has a variable-length decoder (VLD) coprocessor and an image coprocessor (ICP) to offload simple but high-bandwidth tasks from the processor core.

The video and audio DMA units relieve the DSPCPU of mundane chores associated with data I/O. In addition to moving data to and from the SDRAM, they format data to make processing efficient. For example, the video-in unit can demultiplex each byte of packed YUV data to three separate arrays in the SDRAM.

The video units deal with CCIR 601/656 YUV 4:2:2 data. To reduce data storage requirements, the video-in DMA can scale horizontally by 2:1, and the video-out DMA can scale horizontally by 1:2. The output rate is programmable from 10 to 38 MHz, with a super-fine resolution of 0.02 Hz. The video-out unit can also overlay graphics with alpha blending.

Audio DMA can deal with either 8- or 16-bit samples, in either mono or stereo. Sampling rates are programmable from 0 to 80 kHz with the same 0.02-Hz resolution available for the video rate. This extremely fine control over data rates aids in audio/video synchronization during playback. The audio or video can be speeded or slowed in imperceptible amounts as needed to maintain synchronization.

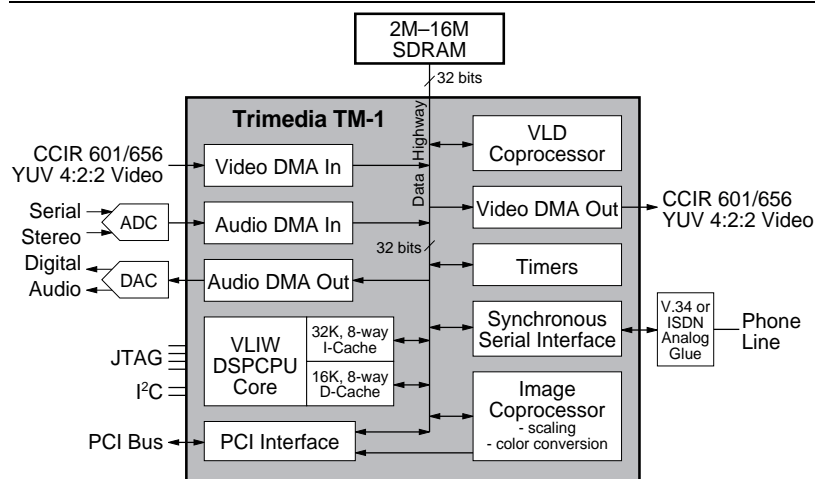


Figure 1. TM-1 combines a powerful VLIW CPU with four autonomous DMA units.

The VLD unit decodes Huffman-encoded MPEG streams, relieving the DSPCPU of this task. This unit performs memory-to-memory DMA. At the high bit rates of MPEG-2, too much DSPCPU time would be devoted to detokenization, which would waste the special capabilities of the VLIW core.

The ICP can operate as either a memory-to-memory or memory-to-PCI DMA unit. In memory-to-memory mode, the ICP can perform either horizontal or vertical image filtering and resizing. The filtering is flexible, and the coefficients are programmable.

In memory-to-PCI mode, the ICP can perform horizontal resizing followed by color-space conversion from YUV to RGB. This mode is used, for example, to dump a decompressed video frame into a window on a PC video screen. The ICP transfers the image to the frame buffer of a PCI-bus video card using a full, per-pixel occlusion bit mask to handle arbitrary overlapping windows. The bit mask determines which pixels are actually stored in the frame buffer for display.

Data Highway Has Traffic Cop

The internal data-highway bus consists of separate 32-bit data and address buses and is used for communication and access to the SDRAM by all the autonomous units. The bus implements a burst packet protocol to match the capabilities of the SDRAM. All units depend on gaining access to the bus in a timely manner to perform their tasks and satisfy the real-time constraints inherent in a live multimedia system.

To guarantee satisfactory real-time behavior, the internal bus is mediated by a sophisticated bus arbiter. By writing to registers in the arbiter, software assigns a fraction of available bus (SDRAM) bandwidth to each master. In return, the arbiter guarantees that each bus master will be given no less than this minimum bandwidth and that each master will suffer no more than a maximum associated latency.

Unused bus bandwidth is assigned by a fixed priority scheme. Requests from the DSPCPU are granted and serviced within one cycle; requests from other masters are serviced within a few cycles.

Philips says that the bus arbitration scheme is a feature that makes TM-1 a real-time system instead of just a highly integrated microprocessor. Since the arbiter performs low-level deadline scheduling in hardware, this seems to be a reasonable assertion.

TM-1 System Configurations

In a TM-1 system or subsystem, SDRAM provides local program and data storage, and the PCI bus is used

to interface with a host processor in a PC application or with additional peripherals in an embedded system. For CCIR 601-compliant video sources, no glue logic is required. As a pioneer in audio and video electronics, Philips already offers interface chips for non-CCIR devices. The video source can be programmed through the two-wire serial I²C interface. Analog audio requires separate analog-to-digital (ADC) and digital-to-analog (DAC) converter chips.

For applications like video conferencing, data communications can be performed over a V.34 modem or ISDN. TM-1 performs the signal processing and provides a dedicated serial interface, while external circuitry provides the analog land-line interface.

Figure 2 shows a block diagram for a standalone embedded system. Because it is based on a general-purpose processor and integrated peripherals, TM-1 needs only SDRAM and program ROM to form a processing system. If TM-1 chips and SDRAM eventually become cheap enough, TM-1 could form the core of inexpensive consumer multimedia devices.

In general, such a system might have both audio/video input and output plus extra peripherals, but for some common applications like set-top boxes or game machines, only video output plus peripherals (e.g., joystick, CD-ROM) would be required. To reduce die cost, Philips can build stripped-down versions of TM-1 for specific purposes. As noted in a previous article (*see 081603.PDF*), a chip with an LCD display controller instead of video output is planned.

Figure 3 shows a PC system with a TM-1-based add-in card. The card would be useful simply as a playback engine for high-quality MPEG multimedia, but the full generality of TM-1 could also be exploited, since programs can be downloaded from the host.

A PC-based video-conferencing application would use simultaneous audio/video input and output and a



Gert Slavenburg, chief scientist of Philips' TriMedia group, explains the VLIW design of TM-1.

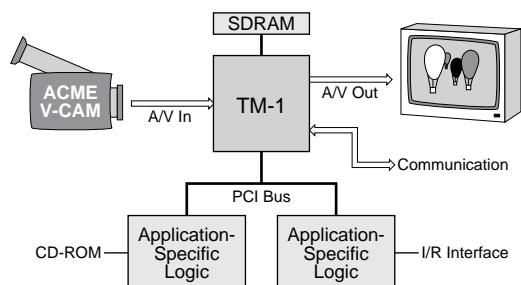


Figure 2. A standalone system, such as a set-top box, requires only a TM-1 chip, SDRAM, analog interfaces for video and audio, and applicable PCI peripheral interfaces.

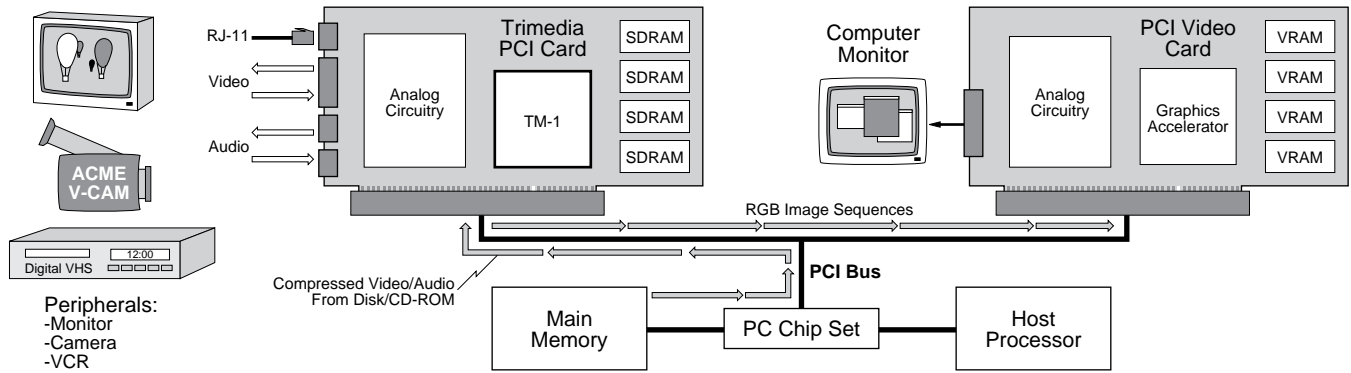


Figure 3. System diagram of a TM-1 application implemented on a PCI card. The TM-1 card could be designed to accept and process multimedia streams from several sources, including a modem port, video-camera input, and compressed video from the host.

land-line interface. Since TM-1 performs all compression/decompression and data communication, the host PC remains responsive to user interaction and could be used to support collaborative editing of a shared document, for example. Also, since mass storage is inexpensive, a PC conferencing application could save the entire video conference to disk for later review.

The PC add-in card could also perform video editing functions in a business or home setting. With its capability to mix audio/video decompression and 3D graphics, the card could be a platform for advanced realistic games and industrial simulations.

At the Forum, Philips' Gert Slavenburg showed a mock-up of a multimedia accelerator on a short PCI card. The card had audio/video input and output and an RJ-11 telephone line interface. Philips claims the card could be built and sold for less than \$300 retail in 1996, but the cost of SDRAM is a critical issue. Assuming it could replace a \$100 V.34 modem and provide high-quality multimedia, this might make a compelling product.

More VLIW DSPCPU Core Details Revealed

Slavenburg revealed that the VLIW processor core has 27 execution units, not 25 as reported previously; the characteristics of the units are detailed in Table 1. The

Function Unit	# Of Units	Latency	Throughput
Constant	5	1	1
Integer ALU	5	1	1
Load/Store	2	3	1
DSP ALU	2	2	1
DSP Mult	2	3	1
Shifter	2	1	1
Branch	3	3	1
Integer/FP Mult	2	3	1
FP ALU	2	3	1
FP Compare	1	1	1
FP Sqrt/Div	1	17	16

Table 1. TM-1 execution-unit characteristics. All unit latencies and issue rates are exposed to software, so the compiler must arrange operations appropriately to guarantee correct program operation.

number of execution units was tuned to match the computing requirements of actual multimedia code. Philips says it has spent several years coding real audio/video applications, such as MPEG and H.320 codecs.

As shown in Table 1, the simple integer execution units have the expected single-cycle latency, but most other units take two or three cycles to compute a result. In particular, branches and memory operations have three delay slots.

In the DSPCPU architecture, all latencies are exposed to software; there is no hardware interlocking. Consequently, the compiler is responsible for scheduling operations so that results are not used before they are available. The compiler's job can be quite complex, considering that it must attempt to maximize performance by filling five operation slots in each instruction while considering various operation latencies.

Figure 4 shows an example of how the guarding feature of the architecture helps the compiler pack operations in branch-delay slots. For ease of illustration, the VLIW instructions in this example have been simplified from five to three operations, but the instructions shown still constitute valid TM-1 code.

The branch condition is computed into r11 by the ILEQ (integer less-or-equal) operation in slot 1 of cycle 102. The CJMPT operation causes a branch if the LSB of r11 is set. The following three instructions in cycles 104 through 106 are the branch-delay slots. Per the definition of delayed branches, these operations are always executed, regardless of the branch outcome.

One of the delay-slot operations, BITINV (bit invert), inverts the branch condition. The other eight delay-slot operations are guarded by either the branch condition or its inverse, depending on whether the operation is part of the fall-through or branch-taken path. Thus, guarding lets the compiler fill branch-delay slots with relative ease.

Note also that the CJMPT branch could itself be guarded if necessary. This technique allows some compound branch conditions to be implemented without executing a separate explicit operation.

Also, guarding, like conditional execution in the ARM and other architectures, can sometimes eliminate branches entirely. With a branch delay of three cycles, removing branches is important in TM-1 programs.

TM-1 takes an interesting approach to interrupt servicing. Interrupts are allowed only when the processor executes an interruptible conditional jump. Thus, it is incumbent on the compiler to insert interruptible jumps with sufficient frequency to prevent excessive interrupt latency. The advantage of this approach is minimal interrupt-service overhead. The compiler ensures that a portion of the register file sufficient for interrupt service is not in use when an interruptible jump executes. Thus, interrupt handlers need not save context before beginning interrupt service.

Multimedia a Popular Market

TM-1 appears to fit its intended markets well. The computational requirements of simultaneous, high-quality audio/video compression and decompression are beyond the capabilities of a standard microprocessor. Dedicated solutions, such as MPEG-only chips, can perform one task well, but for a little more die size and cost, devices like TM-1 can implement a reprogrammable multifunction consumer device or PC enhancement card with the same number of chips.

When Philips announced its plans nearly a year ago for a VLIW-based single-chip multimedia processor, it seemed the company had a unique approach. At last month's Microprocessor Forum, however, Chromatic announced its Mpack chip with a similar implementation, similar goals, and similar performance.

Assuming that both Philips and Chromatic achieve their goals, there will be at least two good solutions for accelerating high-quality multimedia in PCs. Mpack will sample and be in production earlier, but TM-1 has technical and marketing advantages. First, TM-1 has floating-point hardware, which is important for 3D

Price & Availability

Samples of TM-1 are planned for 2Q96 with production shortly thereafter. Software development tools are available now. For more information, contact the Trimedia Group of Philips Semiconductor (Sunnyvale, Calif.) at 408.991.3838; fax 408.991.3300.

graphics. For standard multimedia algorithms, however, Mpack is adequate and might reap a die-size advantage over TM-1.

Second, TM-1 is supported by an optimizing C compiler and development system that Philips is making available to customers. Chromatic has stated bluntly that it will control the development of Mpack code. Since both Philips and Chromatic will supply canned software for standard functions, this difference may not mean much to some customers. Where proprietary or additional functions are important, however, Philips has an edge. OEMs can differentiate their TM-1 products by supporting additional standards. Embedded applications, such as multimedia kiosks, need no host processor to supplement TM-1.

Third, Philips says TM-1 has a lower load on the host system for most functions, especially MPEG, than the Chromatic chip. Fourth, Philips is a multinational conglomerate with plans for its own consumer-electronics applications of Trimedia chips, which virtually guarantees high-volume production.

Philips hopes to establish Trimedia as an industry standard for multimedia applications. TM-1 has the technical attributes and Philips has the production muscle to achieve this goal, but being a little late to market can nullify these advantages. To have a chance against Mpack and low-cost special-purpose MPEG chips, Philips must execute its plans flawlessly, produce inexpensive chips, and hope the market is still receptive when TM-1 ships. ♦

Cycle	Issue Slot 1	Issue Slot 2	Issue Slot 3
⋮	⋮	⋮	⋮
101	#13 ♦ r11	ld32 r12(4) ♦ r13	me8 r101,r100 ♦ r23
102	ileq r11,r15 ♦ r11	iaddi r12,#220 ♦ r17	#LABEL ♦ r18
103	r11: st32 r12(8) r14	ld32x r17,r19 ♦ r16	cjmnt r11,r18
104	bitinv r11 ♦ r21	r11: st32 r12(12) r13	r11: me8 r99,r98 ♦ r42
105	r21: st32 r12(12) r17	r21: fir8 r101,r102 ♦ r23	r11: me8 r97,r96 ♦ r43
106	r21: #1234 ♦ r24	r21: fir8 r103,r16 ♦ r25	r11: me8 r95,r94 ♦ r44
(Not taken) 107	fir8 r104,r105 ♦ r26	fir8 r106,r107 ♦ r27	iadd r42,r43 ♦ r45
⋮	⋮	⋮	⋮
(Taken) 107	me8 r93,r92 ♦ r45	me8 r91,r90 ♦ r46	#0xffff000 ♦ r18
⋮	⋮	⋮	⋮

Figure 4. Code fragment illustrating how guarding helps fill branch-delay slots. For ease of illustration, only three-slot VLIW instructions are used. Branch on 'true' is in issue slot 3 in cycle 103. Some operations in the delay slot are guarded with r11 (thus executed if branch is taken), while others are guarded with r21 (thus executed if branch not taken).