The Trimedia TM-1 PCI VLIW Media Processor

Gerrit A. Slavenburg

Selliah Rathnam

Henk Dijkstra

1996 (eighth) Hot Chips Symposium

final slides - July 19, 1996

TM-1 Block Diagram

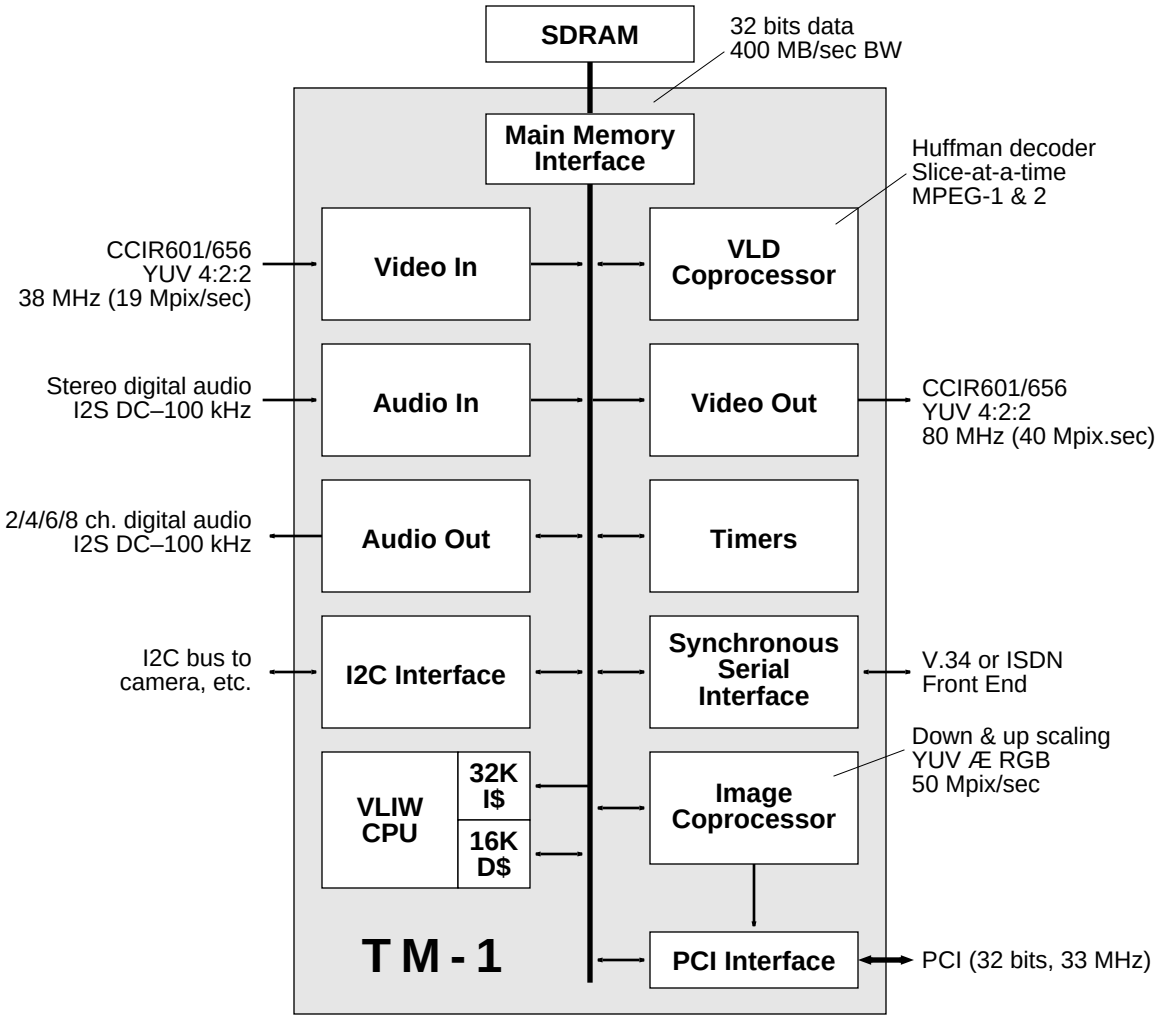

System on a chip

for system details, see Proceedings of the 1995 Micro Processor Forum.

Applications

- standalone / PC
- video confer-encing
- settop decoder
- DVD decoder
- 3-D graphics
- audio en/de/ transcode (AC-3, MPEG)
- audio synthesis

TM-1 VLIW engine

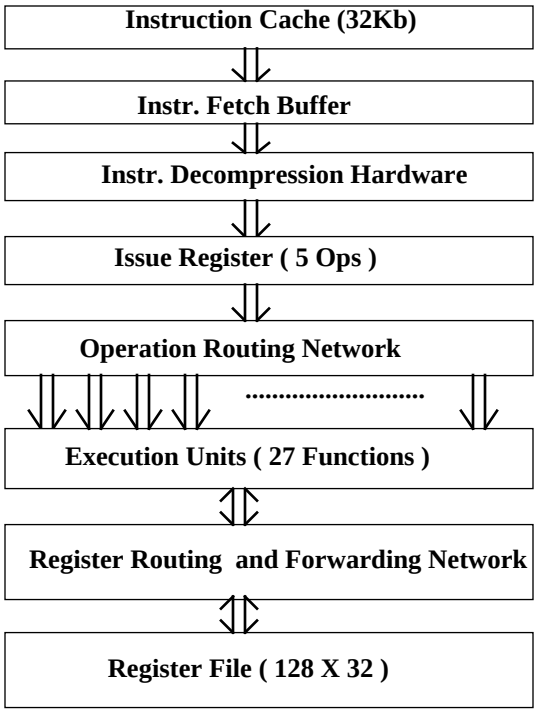

- 5 RISC operations/cycle
- 32 kByte Icache (compressed)
- dual port 16 kByte Dcache
- conditional (guarded) operations
  $R_g : R_{dest} = imul\ R_{src1}, R_{src2}$
- multimedia operation set

Func Type	Qty	Latency	Recovery
Const	5	1	1
ALU	5	1	1
Memory	2	3	1
Shift	2	1	1
DspAlu	2	2	1
DspMul	2	3	1
Branch	3	3	1
Falu	2	3	1
Ifmul	2	3	1
Fcomp	1	1	1
Fdiv/Fsqrt	1	17	16

TM-1 cache architecture

- I-cache contains compressed VLIW instructions
- 32 kByte I-cache, 8 way set associative, block size 64 bytes

- 16 kByte D-cache, 8 way set associative, block size 64 bytes
- advanced D-cache features:
  - copyback, allocate on write, non-blocking cache
  - 11 cycle read-miss, 3 cycle write-miss penalty
  - (pseudo) dual ported
  - streaming
  - dual entry copyback buffer
  - programmer controlled prefetching
  - programmer controlled alloc
  - optional cache locking
  - performance evaluation support

TM-1 MPEG-2 decoder stats

DVD-batman bitstream, variable rate, 4 - 9 Mbit/sec

NO programmer prefetch/alloc

100 MHz TM-1 CPU load for video decoding = 64 %

averaged 3.95 useful RISC operations/VLIW operation

CPI contributor	amt
instruction issue	1.00
I-cache misses	0.07
D-cache misses	0.27
D-cache bank conflict stalls	0.03
total CPI (clock cycles/VLIW instruction)	1.37

TOTAL CPU + memory system performance 3.95/1.37 = 2.9 ops/clock

(the majority of these operations are SIMD multi-media operations used in the DCT and Motion Compensation code)

DVD decoder resource usage overview (100 MHz)

	CPU load	SDRAM load (= highway load)	effective ops/ cycle[a]
program stream decoder	< 5 %	< 2 %	*no recent data*
video decoder[b]	64 %	18 %	2.9
audio decoder[c]	4.6 %	0.95 %	2.4
subpicture decode/insert VBI decode PCI decode	*no data available yet, code under development*		

a. speedup over 1 issue/cycle machine divided by CPI due to cache misses. Effective ops include multimedia ops.
b. DVD-Batman stream, 4-9 Mbit/sec., first 500 images, no programmer prefetch
c. 44.1 ksample/sec stereo 16 bit MPEG L2 audio

Problem: DVD video excursions. 3 runs, 44 out of 500 images exceed 85% CPU load. Frame 2-9 (91 %), 109-118 (95 %), 315-340 (88 %).

Solution: quality degradation (will disappear with prefetch & faster TM-1's)

TM-1 performance sample : 3-D setup (1)

input:

- description of the 3 vertices of a triangle in 3-D
- per vertex : screen coordinates, 4 color values, Z, 1/w, u/w and v/w (totalling 10 float values/vertex, or 30 float values total)

output:

- 44 fixed point values per triangle to off-chip raster engine

computation:

- sort vertices by Y coordinate, compute all derivatives
- 4 float divisions
- 83 float multiplications
- 14 float comparisons
- 93 other float operations (fadd, fsub, fixpoint convert)
- 24 integer multiply operations
- total of 410 operations, incl. loads/stores

TM-1 performance sample: 3-D setup (2)

Table 1: 3D triangle setup duration in CPU clock cycles/triangle

	TM-1 (cycles)	P5-133 (cycles)[a]	P6-200 (cycles)	TM-1 over P5	TM-1 over P6
CPU (no multimedia ops)	210	2526	2440	12x	12x
CPU (with multimedia ops)	96[b]	2526[c]	2440	26x	25x
CPU (with multimedia ops) + memory system	170[d]	3591	3480	21x	20x

a. Visual C++ 4.0 ‘maxspeed’ optimized, L2 cache 256 kBytes, 32 MByte DRAM
b. this code achieves 4.68 useful operations/cycle
c. MMX does not speedup triangle setup computation
d. no programmer prefetch

Architectural statistics

(min/avg/max)

application category	#1 dyn % of guarded ops[a]	#2 fine grain par. (speedup)[b]	#3 memory system (CPI)[c]	#4 effective ops per cycle #2/#3[d]	#5 SDRAM bandwidth util.%[e]
general purpose (10 large programs including some spec92 benchmarks)	23%/26%/30%	1.41/2.46/3.61	1.03/1.37/2.39	1.13/1.87/2.59	3%/21%/66%
MPEG video decode	14.3%	3.95	1.37	2.88	18%
MPEG audio decode	5.2%	4.08	1.70	2.40	0.95%
3-D workload (3 programs)	17%/18%/19%	3.67/4.10/4.68	1.05/1.41/1.77	2.64/2.98/3.49	2%/24%/46%

a. operations with dynamically computed guard as a percentage of total operations executed - shows frequent use of conditional execution
b. execution cycles on RISC (1 op/cycle) with same instructionset/execution cycles on TM-1 VLIW CPU - without cache effects
c. Cycles Per Instruction = total # clock cycles / total number of VLIW instructions executed - shows quality of cache system
d. dividing column 2 by column 3 results in actual RISC instruction equivalents per clock cycle (CPU + memory system)
e. memory utilization. 100% corresponds to 400 MByte/sec SDRAM bandwidth utilization

TM-1 silicon status

CPU Test Chip:

- DSPCPU, caches, SDRAM i/f, timers, Vin, PCI i/f, 80 MHz issue rate
- CTC 1.0 silicon (0.5 u 4LM CMOS) selective sampling May 96
- CTC 1.1 silicon (0.5 u 4LM CMOS) sampled to EAP customers Jul 96
- 3.3 Volt, 2.5 Watt (operating)

TM-1 chip:

- complete TM-1 system
- 0.50 u 4LM CMOS 100 MHz issue rate samples Oct 96
- 0.50 u 4LM CMOS 100 MHz RFS (volume avaiability) Q1 97
- 0.35 u 4LM CMOS 132 Mhz issue rate samples Q1 97
- 3.3 Volt, 4 Watt

TM-1c

- 0.35u 4LM 150 MHz compacted, pin-compatible, enhanced version samples Q3 97

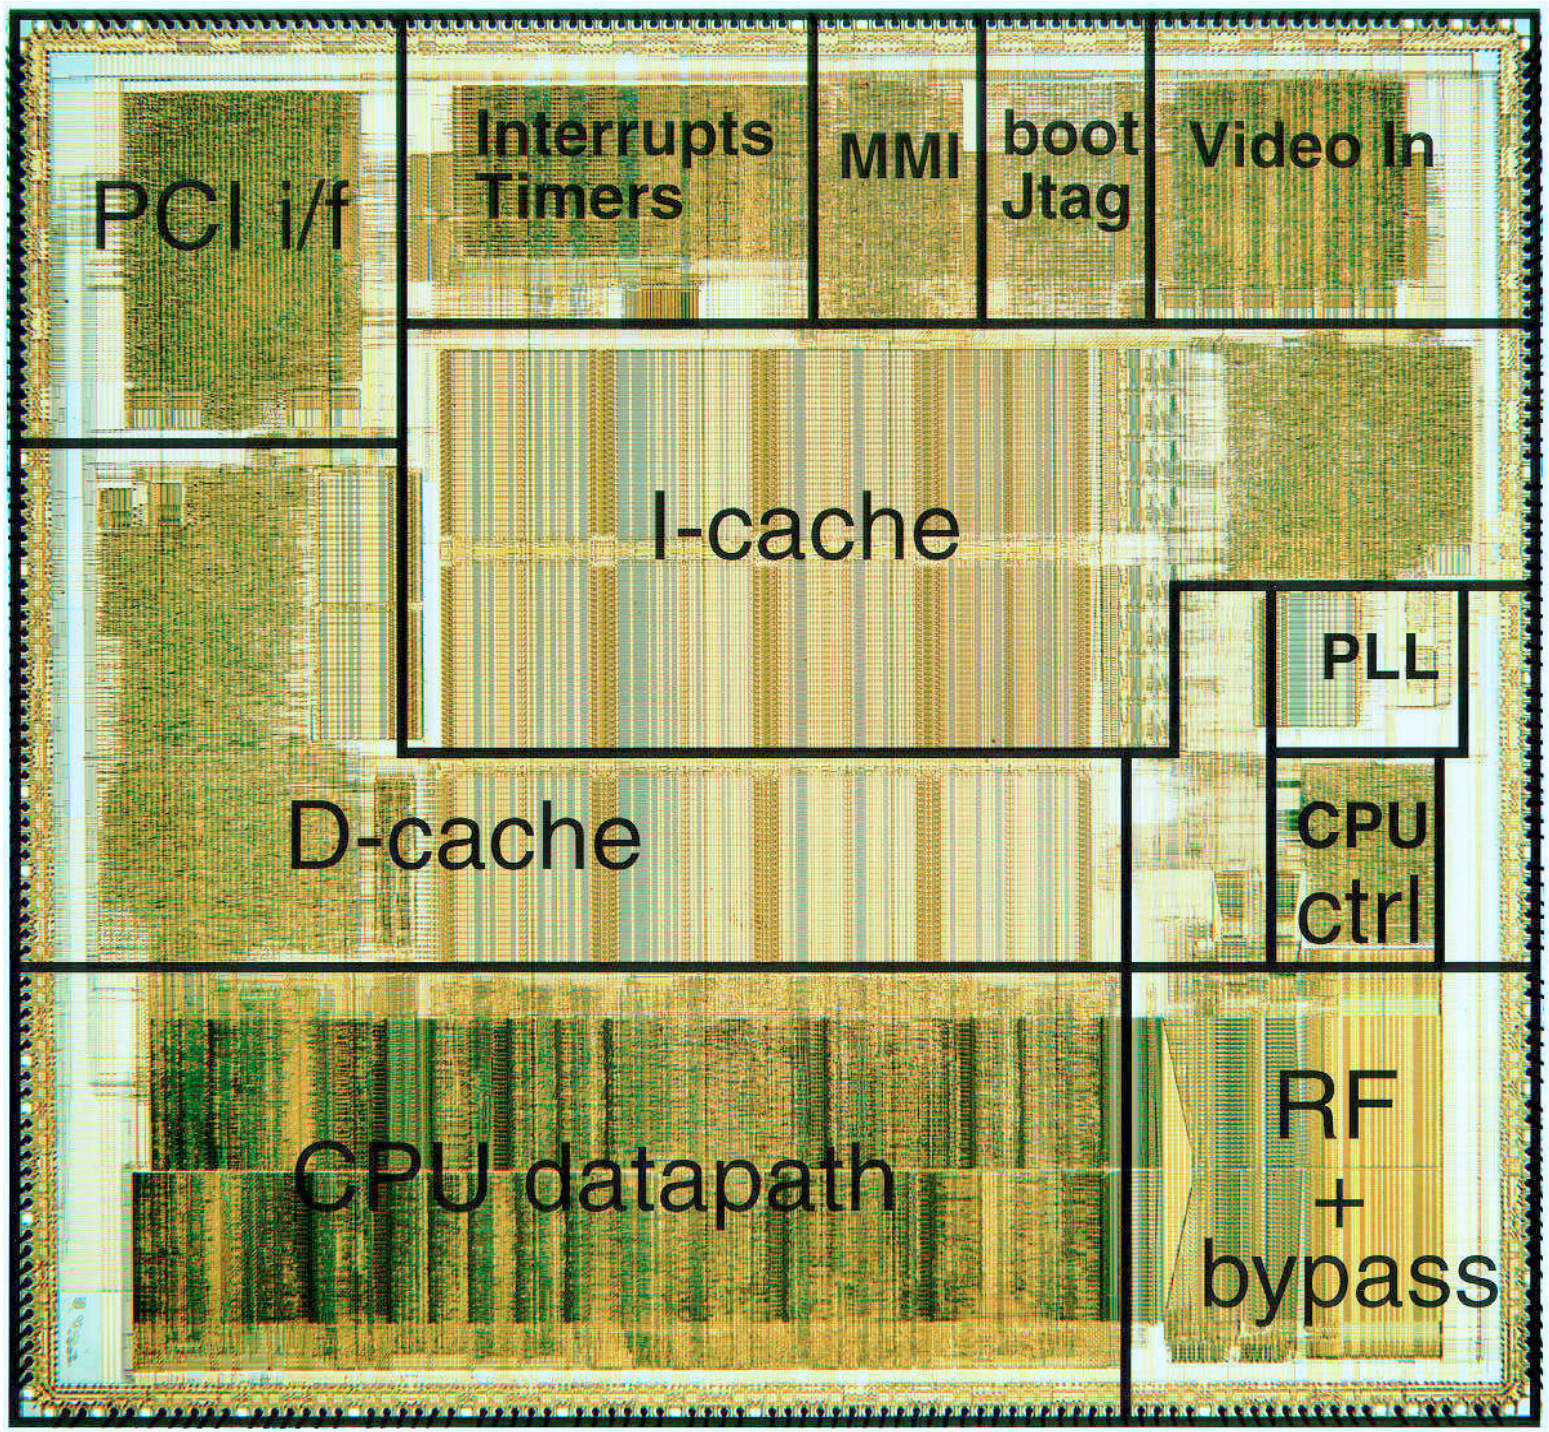

PHILIPS

CTC chip photograph

PHILIPS

TM-1 programming

- Open architecture, 3’rd party application development encouraged
- pSOS+M real-time kernel
- State-of-the-art fine-grain parallelizing C and C++ compilers available to EAP customers today with CTC boards
- Multi-media libraries with well-tuned code available with TM-1 RFS
- complete application software available with TM-1 RFS
  - MPEG-1 decoder
  - MPEG-2 program stream decoder
  - 3-D graphics pipeline
  - PC audio synthesis (FM, wavetable)
  - V.34 modem
- available in 1997
  - H.324 (PSTN) PC video conferencing Q1 97
  - H.320 (ISDN) PC video conferencing Q4 97

Lessons learned

- 21 known bugs in CTC1.0, *none* of which in the VLIW CPU and memory system
  - VLIW's are datapath intensive, very low control area/complexity
  - directed tests + random generation of tests works well for CPU/ memory system verification
  - no equivalent methodology known for verifying peripherals/co-processors and PCI
- the ratio between average CPU load and peak CPU load for software real-time video decoding is greater than expected
  (hence, software video decoders require overpowered CPU's and then mix well with 'best effort' compute intensive tasks such as 3-D)
- VLIW's work extremely well for signal processing and 3-D tasks, and quite good for general purpose programming
- programmer use of prefetch/alloc is very hard