

Write Buffers Enhance 486 Performance

Vendors Tout Cache-Like Performance At Lower Price

By Mark Thorson

MetaDesign Semiconductor (formerly Matra Design Semiconductor), Headland, and Toshiba have all recently announced write buffers that can be used with Intel's 486 microprocessor. A write buffer is a type of FIFO memory that buffers write cycles from the CPU, allowing the CPU to continue execution without waiting for the memory write to complete. The write buffer takes responsibility for propagating the write cycle to memory, and the data and address for the write cycle are held in the write buffer until memory is ready to receive the write.

With the introduction of faster versions of the 486, it has become increasingly difficult to run the CPU without wait states even when an external second-level cache is provided, because there is so little time available for memory access and chip-to-chip crossings. A write buffer has no RAM arrays and does not require propagating signals through another chip, so it is less

difficult to build write buffers that can reach the speeds of the latest processors.

Adding an external write buffer to the 486 is claimed to be a low-cost alternative to using an external second-level cache to increase 486 performance. The on-chip cache of the 486 is a writethrough cache, i.e., it propagates all write cycles to the processor's bus. As a result, the cache delivers a massive reduction in read traffic, but has no effect on write traffic. By adding a write buffer, most bus cycles can be satisfied without wait states, which is said to provide a performance increase similar to that of adding an external cache.

The 486's internal write buffer is four levels deep (i.e., up to four write cycles can be pending). The CPU will only see wait states on a write when it attempts to perform another write when the write buffer is full. External write buffers add additional levels of buffering, increasing performance when a burst of writes occurs that overflows the 486's on-chip write buffer.

MetaDesign's write buffer is a generic device intended for use in systems using either Intel or non-Intel CPUs (such as Motorola's 680x0). Headland and Toshiba have incorporated write buffers within their 486-compatible system logic chip sets.

Price & Availability

MetaDesign's WB 416 and WB 418 are packaged in 52- and 64-pin PQFPs, respectively. In 10K-unit quantities, pricing is \$6.50 at 33 MHz and \$10 at 50 MHz. These are planned to drop to \$3.50 and \$6.50, respectively, late in the fourth quarter.

Available at 33 MHz, Headland's HTK340 chip set costs \$45 in 1K quantity. It is in volume production now.

Toshiba's SLIK+ and SLIK+ Enhanced chip sets are very similar, the difference being that SLIK+ Enhanced supports error-correcting codes for memory and it provides parity generation and checking on the processor and expansion buses. The SLIK+ chip set costs \$200 in 1K quantity at 33 MHz, or \$220 at 50 MHz. The SLIK+ Enhanced chip set costs \$220 in 1K quantity at 33 MHz, or \$240 at 50 MHz. Both chip sets are in volume production.

When asked about the remarkably high prices of these chip sets, Toshiba said they are trying to establish a market price, and that high-volume customers would be likely to receive discounts.

MetaDesign Semiconductor, 2895 Northwestern Parkway, Santa Clara, CA 95051; 408/986-9000; fax 408/748-1038.

Headland Technology, 46221 Landing Parkway, Fremont, CA 94538; 510/623-7857; fax 510/656-0397.

Toshiba, 9775 Toledo Way, Irvine, CA 92718; 714/455-2000; fax 714/859-3963.

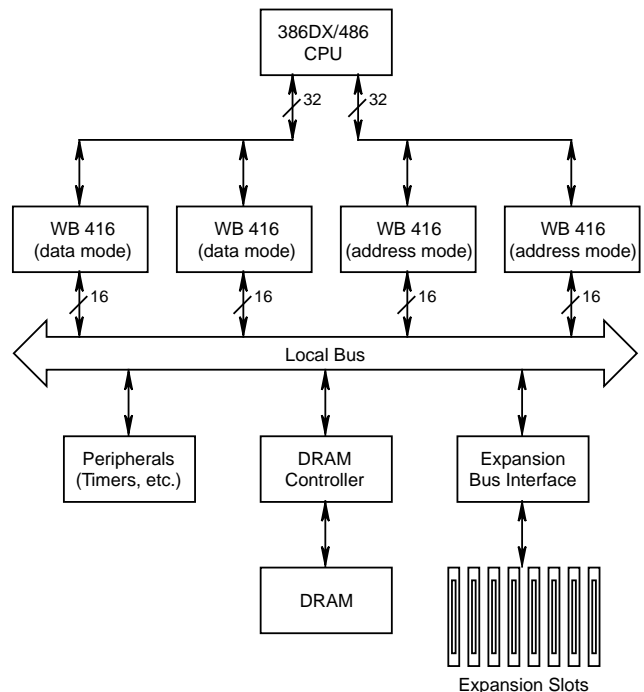


Figure 1. Typical system using MetaDesign's write buffer.

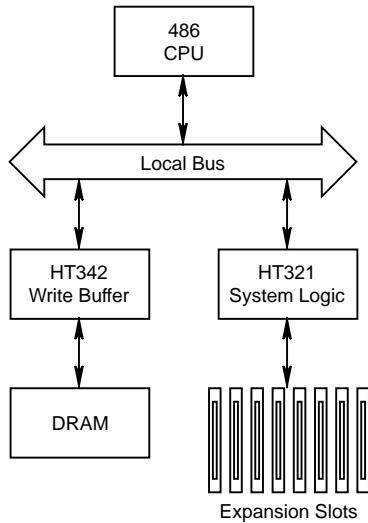


Figure 2. Block diagram of system using Headland's ISA chip set with integral write buffer.

MetaDesign's WB 416 and WB 418

MetaDesign's WB 416 and WB 418 are promoted as being well-suited for upgradeable systems in which the CPU resides on a replaceable module that plugs into the system board. MetaDesign also anticipates applications in modular PCI-bus systems and as an enhancement for some cache-based designs. (E.g., most writethrough caches have a single-level write buffer; MetaDesign can increase that to four levels.)

External logic is required to handle the bus protocol of the specific application. Each chip is either 16- or 18-bits wide, depending on packaging. (A 32-bit version is planned.) Figure 1 is a diagram showing a typical configuration. For a 386DX or 486, there are two chips strapped for address mode and two data-mode chips interposed between the CPU and a 386- or 486-like local bus.

The MetaDesign chips perform snooping at all four levels (i.e., the write buffers compare the address tag of buffered write cycles against the address of CPU read cycles). If a read cycle is a snoop miss on the write buffer, the write buffer can be bypassed and the read cycle is propagated directly to memory. If the read cycle is a snoop hit, the data can be read out of the write buffer chips. Reordering of read and write cycles can be a problem for I/O devices, but I/O devices in PCs are generally mapped into the I/O address space, which is not affected by the write buffers.

If consecutive processor writes are hits in the top level of the pipe, they can be combined. When a write is strobed into the FIFOs, the address FIFO raises a signal if the address matched the previous cycle. The data FIFO signals if the pattern of enabled bytes is compatible with the bytes already loaded for that address (i.e.,

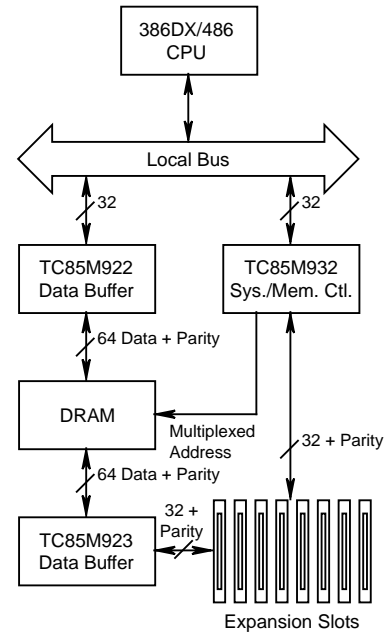


Figure 3. Block diagram of a system using Toshiba's Micro Channel chip set with integral write buffer.

if none of the bytes in the current cycle overwrite bytes updated in the previous cycle). External logic can then tell the FIFOs to combine the two write cycles into one by asserting the EN# pin.

Headland's HTK340

Headland's HTK340 chip set for the 486 is an adaptation of their HTK320 chip set for the 386DX. Figure 2 is a block diagram of an HTK340-based system. It replaces the HT322 cache/DRAM controller used in the HTK320 with the HT342 write buffer/DRAM controller. The four-level write buffer is full-featured, supporting both read snooping at all levels and assembly of byte- and word-length cycles into a smaller number of wider cycles. Write buffers for both the address and data paths to DRAM are included in the HT342.

Toshiba's SLIK+ and SLIK+ Enhanced

Toshiba's SLIK+ and SLIK+ Enhanced chip sets for 386DX- and 486-based Micro Channel systems include a four-level write buffer split between two chips, the memory data buffer and the system/memory controller. The former holds the data path, and the latter holds the address path. Figure 3 is a block diagram of a SLIK+ Enhanced system.

Unlike the Headland and MetaDesign chips, the Toshiba implementation does no snooping. To avoid the possibility of a read cycle referencing data held in the write buffer, all read cycles must be stalled until the write buffer is flushed. The Toshiba chip set also does no assembly of byte- or word-length cycles into wider cy-

cles. The design was developed by Micral.

Benchmark Performance

Table 1 shows benchmark scores provided by MetaDesign, and Table 2 shows scores from Headland. The data from MetaDesign was collected on a system that used a 386DX chip set (VLSI's TOPCAT), with proprietary logic in a PLD for retrofitting the 486DX or 486DX2 CPU to a 386DX-like local bus. The data from Headland compares an HTK340 prototype against both cache-based and cacheless systems purchased on the open market.

(Note that Headland's data includes scores for non-Headland chip sets. These scores may vary from scores reported by the vendors for those chip sets, because the system board designer may have chosen an implementation which does not run the chip set at its maximum performance. Headland disclaims responsibility for accurately characterizing the performance of its competitors.)

MetaDesign's scores consistently show a large performance improvement, on the order of 7%–40%. Only the floating-point benchmarks show no gain, since the write buffer does not affect the FPU interface. In a float-

Benchmark	With Write Buffer	Without Write Buffer	Improvement
Byte CPU Index (AT)	7.66	6.29	21.8%
Byte CPU Index (386)	2.57	2.2	16.8%
Byte FPU Index (AT)	49.11	49.11	0.0%
Byte FPU Index (386)	6.87	6.87	0.0%
PowerMeter v1.5 MIPS	20.731	14.933	38.8%
PowerMeter v1.5 Dhrystone	26.6	19.1	39.3%
PowerMeter v1.5 Whetstone	5022	4670	7.5%
PCLab Instr. Mix (8088)	0.92	1.08	17.4%
PCLab Instr. Mix (386)	0.87	1.02	17.2%
PCLab String Sort/Move	0.23	0.29	26.1%
PCLab Prime Sieve	0.13	0.14	7.7%
PCLab Conventional Memory Read	0.2	0.32	60.0%
PCLab Conventional Memory Write	0.15	0.27	80.0%

Table 1. Benchmark statistics provided by MetaDesign. The Byte and PowerMeter benchmarks indicate increased performance with increasing score, while the PCLab benchmarks indicate increased performance with decreasing score. All scores were measured using a 25-MHz Topcat/386 system retrofitted with a module containing a 486DX2 running at 50 MHz (internal clock) and the MetaDesign write buffer.

ing-point application that performed many floating-point stores, the write buffers would presumably show a benefit.

Chip Set	Headland HTK340	Symphony 382/461		OPTi 486WB		UMC 82C480		ACC 2046		OPTi DXBB		
		Cache	None	64K	64K	64K	64K	None	None			
486DX2	Byte CPU Index (AT)	13.8	12.9	7.0%	8.6	60.5%	8.6	60.5%	9.1	51.6%	11.1	24.3%
	Byte CPU Index (386)	5.13	5.04	1.8%	3.06	67.6%	3.08	66.6%	3.12	64.4%	4	28.3%
	Byte FPU Index (AT)	65.9	65.3	0.9%	65.9	0.0%	65.9	0.0%	65.9	0.0%	65.3	0.9%
	Byte FPU Index (386)	9.23	9.13	1.1%	9.22	0.1%	9.23	0.0%	9.22	0.1%	9.13	1.1%
	PowerMeter v1.5 MIPS	29.1	26.1	11.5%	26.5	9.8%	27.1	7.4%	27.1	7.4%	25.6	13.7%
	PowerMeter v1.5 Dhrystone	37.3	33.5	11.3%	33.1	12.7%	33.3	12.0%	33.3	12.0%	32.9	13.4%
	PowerMeter v1.5 Whetstone	6531	6447	1.3%	6531	0.0%	6574	-0.7%	6574	-0.7%	6531	0.0%
	PCLab Instr. Mix (8088)	0.6	0.67	11.7%	0.67	11.7%	0.65	8.3%	0.73	21.7%	0.66	10.0%
	PCLab Instr. Mix (386)	0.55	0.61	10.9%	0.61	10.9%	0.6	9.1%	0.67	21.8%	0.62	12.7%
	PCLab String Sort/Move	0.18	0.18	0.0%	0.18	0.0%	0.18	0.0%	0.19	5.6%	0.17	-5.6%
	PCLab Prime Sieve	0.06	0.04	-33.3%	0.04	-33.3%	0.04	-33.3%	0.06	0.0%	0.05	-16.7%
	PCLab Conventional Memory Read	0.11	0.16	45.5%	0.32	190.9%	0.33	200.0%	0.22	100.0%	0.17	54.5%
PCLab Conventional Memory Write	0.05	0.11	120.0%	0.27	440.0%	0.28	460.0%	0.16	220.0%	0.16	220.0%	
486DX	Byte CPU Index (AT)	8.32	8.87	-6.2%	6.17	34.8%	6.33	31.4%	6.72	23.8%	7.84	6.1%
	Byte CPU Index (386)	3.25	3.59	-9.5%	2.36	37.7%	2.43	33.7%	2.49	30.5%	3.02	7.6%
	Byte FPU Index (AT)	33	33	0.0%	33	0.0%	33	0.0%	33	0.0%	32.6	1.2%
	Byte FPU Index (386)	4.61	4.61	0.0%	4.61	0.0%	4.61	0.0%	4.61	0.0%	4.57	0.9%
	PowerMeter v1.5 MIPS	14.7	14.7	0.0%	14.7	0.0%	14.7	0.0%	14.4	2.1%	14.7	0.0%
	PowerMeter v1.5 Dhrystone	18.9	18.9	0.0%	18.9	0.0%	18.6	1.6%	18.5	2.2%	18.9	0.0%
	PowerMeter v1.5 Whetstone	3266	3266	0.0%	3276	-0.3%	3266	0.0%	3276	-0.3%	3276	-0.3%
	PCLab Instr. Mix (8088)	1.18	1.17	-0.8%	1.18	0.0%	1.18	0.0%	1.2	1.7%	1.18	0.0%
	PCLab Instr. Mix (386)	1.13	1.13	0.0%	1.13	0.0%	1.12	-0.9%	1.15	1.8%	1.13	0.0%
	PCLab String Sort/Move	0.32	0.33	3.1%	0.34	6.3%	0.33	3.1%	0.33	3.1%	0.34	6.3%
	PCLab Prime Sieve	0.09	0.1	11.1%	0.1	11.1%	0.09	0.0%	0.09	0.0%	0.09	0.0%
	PCLab Conventional Memory Read	0.16	0.16	0.0%	0.38	137.5%	0.39	143.8%	0.22	37.5%	0.17	6.3%
PCLab Conventional Memory Write	0.16	0.16	0.0%	0.28	75.0%	0.28	75.0%	0.17	6.3%	0.16	0.0%	

Table 2. Benchmark statistics provided by Headland comparing HTK340 chip set with write buffer against commercially available systems with and without cache. The Byte and PowerMeter benchmarks indicate increased performance with increasing score, while the PCLab benchmarks indicate increased performance with decreasing score. All systems have 33-MHz processor bus speed and have 80-ns DRAMs. DX2 systems have a 66-MHz internal clock. Percentage columns show improvement of Headland chip set over each competitor.

Headland's scores include both 486DX and 486DX2 comparisons. The 486DX scores show little or no performance differences among systems, whether or not they have a cache or a write buffer. If obviously unreliable scores are rejected, the data for the 486DX2 suggest a performance improvement of about 2%–12% over cache-based systems or 10%–25% over non-cached systems. (For purposes of comparison, the best of the competitor's scores should be used because they are closest to being an ideal design. Among cacheless systems the OPTi system is fastest, while the best cached system is Symphony.)

The scores for the 486DX2 show much greater improvement because the faster CPU generates twice as many internal bus cycles, hence there are more opportunities for the internal write buffer to overflow into the external write buffer.

To interpret these scores, it is necessary to understand the systems on which they were collected. The slow DRAM performance of MetaDesign's upgraded 386DX system (6-5-5-5 wait states on four-cycle bursts at 33 MHz) exaggerates the effect of the write buffer because it imposes a heavy penalty on any cycle to DRAM, which explains the high percentage-wise increase in their benchmark scores when write buffering is enabled.

The data from Headland was collected on real systems with DRAM performance typical of well-tuned system designs. As a result, their benchmark scores—when adjusted for clock speed—are much higher than that for the MetaDesign prototype. However, the performance differences among systems are much smaller, even when comparing cacheless against cache-based systems.

A large amount of the performance improvement reported for systems with write buffers seems to be the result of demented behavior by the benchmarks. Most PC benchmarks are peculiar little programs which exercise obscure aspects of system performance. Their behavior is not at all typical of real application code.

For example, the PCLabs Instruction Mix benchmarks show virtually no differences among well-implemented cached and cacheless systems, yet they show about an 11% performance improvement for the Headland chip set when the 486DX2 processor is used. Despite being blind to the effect of adding a cache, the benchmark somehow senses the existence of the write buffer. It probably performs a large number of writes—enough to overflow the 486's internal write buffer—or it might be performing many byte or word writes to consecutive addresses.

Although nearly all the benchmarks show Headland beating the competition when the 486DX2 is used, the PCLabs Prime Sieve benchmark reports an anomalously low score for the chip set. This is probably due to

behavior in which the benchmark is attempting to read data which has been placed in the pipeline (i.e., a snoop hit). Apparently, there are wait states involved with servicing a snoop hit, and this particular benchmark heavily exercises that condition.

Conclusion

It should be kept in mind that 486-family CPUs already have an on-chip four-level write buffer. Much of the ability of a write buffer to average out peak demands for write access to the local bus has already been obtained. However, the architectural balance of the original 486 design seems to have been changed by the introduction of CPUs with internal clock-doublers. Because of their higher bus utilization, these processors benefit from an external snooping write buffer.

A feature offered by the MetaDesign and Headland chips—but not the on-chip write buffer of the 486—is the ability to assemble byte- and word-length write cycles into a smaller number of wider cycles. Both companies claim that a large proportion of the bus traffic generated by typical MS-DOS applications is byte- and word-length, and that their chips deliver cache-like performance improvement by optimizing access to 16- and 32-bit memory.

The benchmark data doesn't prove very much. One reason write buffers are claimed to improve 486 performance is that its on-chip cache catches most of the processor read cycles, resulting in a disproportionate number of writes in the traffic propagated to the external bus. However, all of the popular PC benchmarks are small programs that mostly run out of the on-chip cache, so they do not show much performance delta among 486 systems, whether or not they have external cache or write buffers.

Intel's 486DX2 data book reports a 3%–9% performance increase when adding an external cache to the 486DX and a 20%–30% increase when adding a cache to a 486DX2. If these benchmarks were a reliable indicator of system performance, they would easily distinguish between cacheless and cache-based systems. (A few, such as the Byte CPU Index benchmarks, do better in this regard than most others.)

Except under certain specific conditions, such as retrofitting a clock-doubler CPU onto an older system or buffering writes to an expansion bus video controller, the case for external write buffers in 486 systems remains unproven. It seems clear that some amount of performance improvement is being delivered, but the size of this improvement is difficult to judge. To be convincing, a well-tuned system equipped with a write buffer will need to show a significant and consistent performance superiority on real-world software, such as the SPEC or BAPCo benchmark suites (see *μPR* 5/27/92, p. 5). ♦