



Integrated Device Technology, Inc.

A NEW GENERATION OF TAG SRAMS—THE IDT71215 AND IDT71216

APPLICATION
NOTE
AN-136

By Kelly Maas

INTRODUCTION

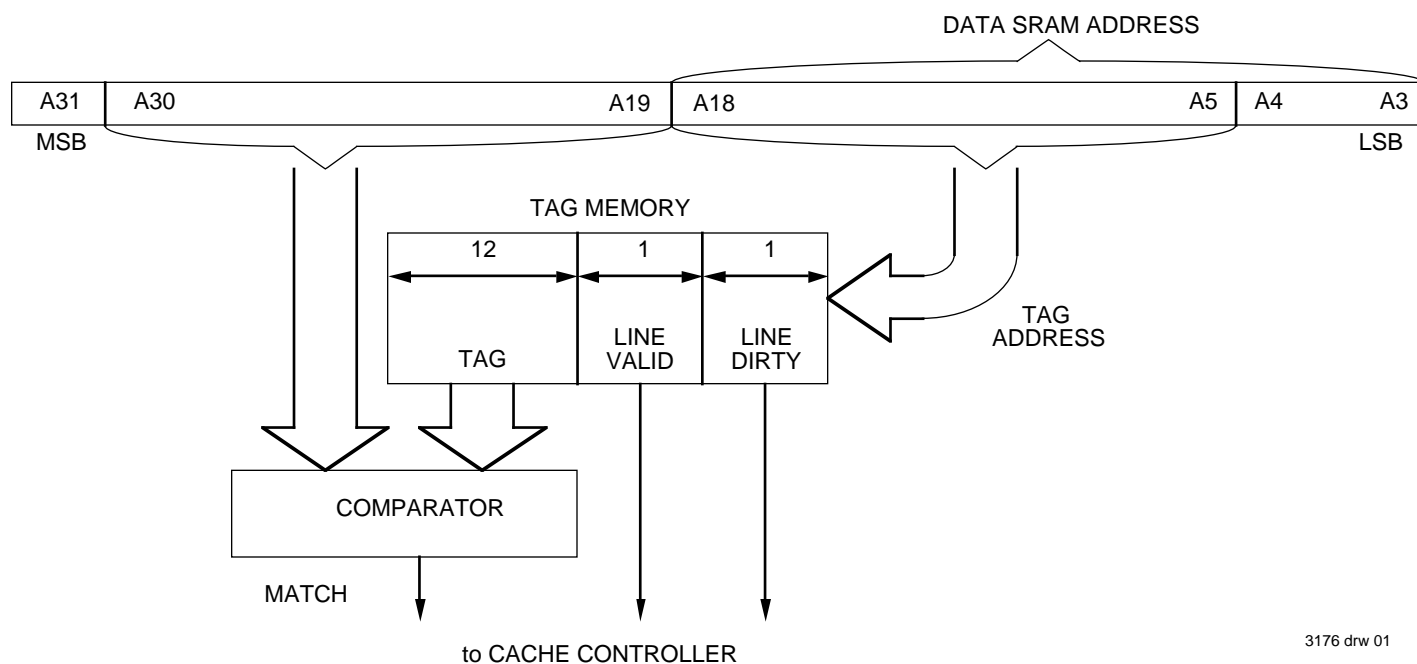
The 71215 and 71216 represent a new generation of integrated Tag SRAMs. Just as earlier Tag SRAMs such as the 71B74 were better suited for tag applications than conventional SRAMs, the 71215/16 go a step further by integrating new features to significantly ease the design of high performance cache subsystems for today's high speed processors. These Tag RAMs are designed for easy interfacing to Intel and PowerPC processors, but are very flexible and can easily be used in other applications as well.

This application note first provides some background information on caches, then describes in detail the architecture and operation of the 71215 and 71216. This is followed by three application examples, then a brief discussion of cache coherency protocol implementation using these Tag RAMs. Since the 71215 and 71216 are very similar, the descriptions and explanations in this application note apply to both unless otherwise noted.

CACHE AND TAG BASICS

For those new to caches, a brief review of cache basics may be worthwhile. A cache is a memory that provides a CPU with high speed access to a subset of the data from main memory. Our discussions are focused on the secondary cache, which is also known as the L2 cache, but it is not much different from the faster primary (L1) cache residing inside most CPUs.

The cache consists of a controller, a data memory and a tag memory. The purpose of the data memory is to store the active data from main memory, and is composed of either synchronous burst or asynchronous SRAMs. The tag memory stores indexes (part of the CPU address field) that indicate which data is stored in the cache. Additionally, most caches also require at least one bit of memory for each cache entry, to indicate the valid or dirty status of that entry. Figure 1 shows how the CPU address field relates to the cache and the tag memory. This example includes valid and dirty status bits, and represents a 512KB cache, 2GB cacheable address space, 32-byte line size, and 8-byte word size.



3176 drw 01

Figure 1. CPU Address Field and the L2 Cache (Showing 512 KB cache size and 2 GB cacheable main memory)

Integrated Tag RAMs operate as ordinary SRAMs, but have an additional access mode in which a word of data (an index) is internally read (but not driven off-chip) and compared with the CPU address driven onto the Tag RAM's data bus. Figure 2 shows the basic architecture of an integrated Tag SRAM. The comparator indicates whether the cache holds the data for the address supplied by the CPU or other bus master. This is a critical timing path since this tag "hit" or "miss" must be determined before the cache memory access can be completed (or even started, in many cases). Note that tag memories connect only to the CPU address bus and never to the CPU data bus.




Figure 2. Basic Integrated Tag SRAM Architecture

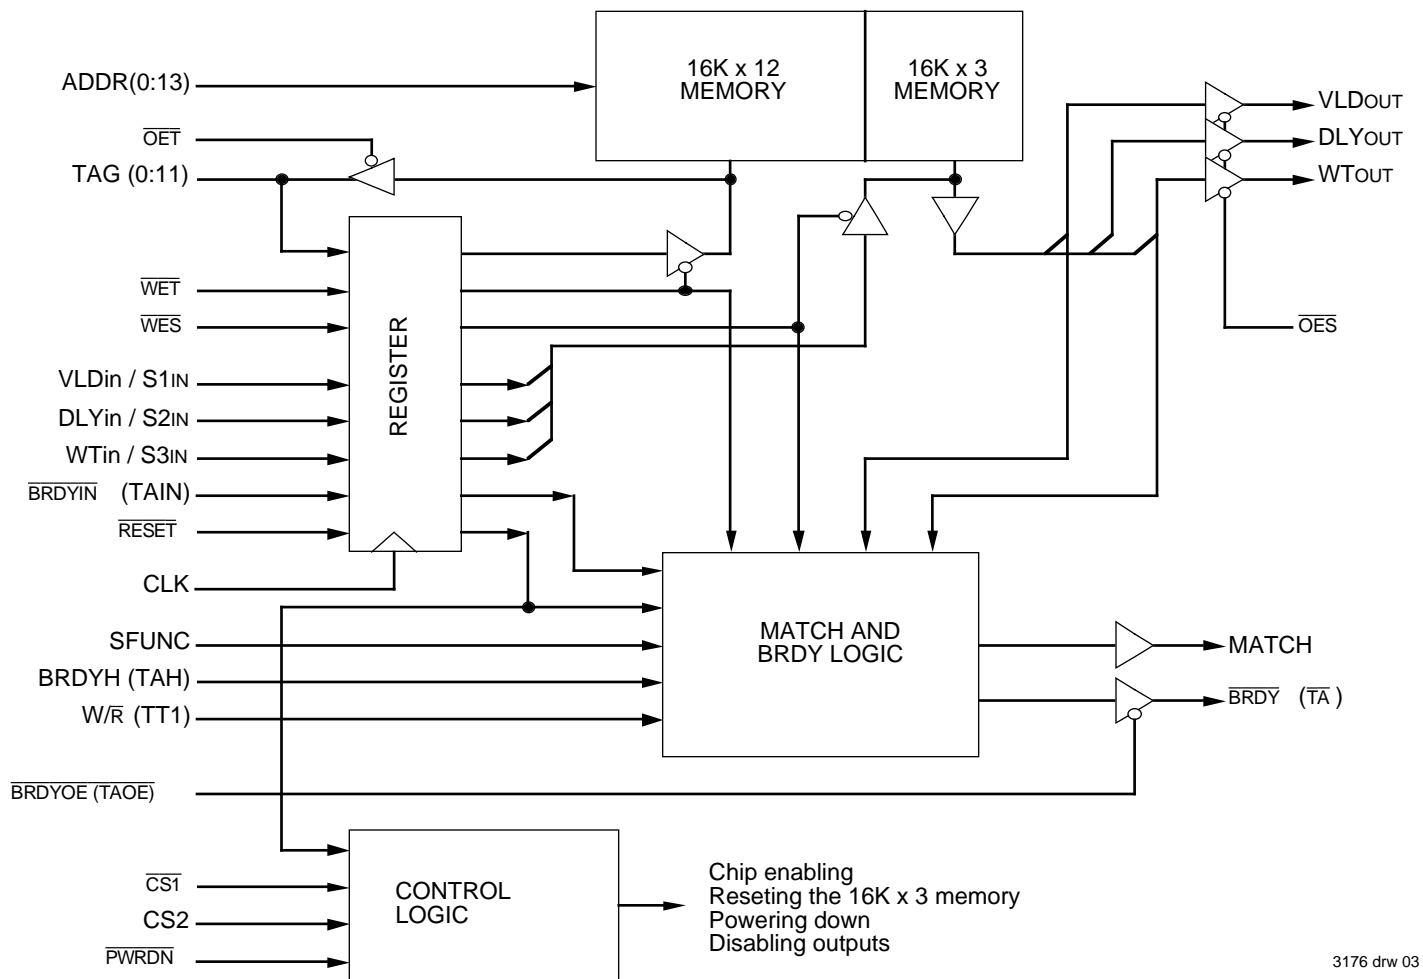
An additional feature of the Tag SRAM is that a portion of the memory is resettable. This permits use of one bit of the data field as a "valid" status bit. Upon system initialization, when the cache contains random data, a quick reset will clear the valid bit for every cache line so that all initial cache accesses will result in a miss. A miss then causes the address to be loaded into the Tag RAM, data from main memory to be loaded into the data RAMs, and the valid bit to be set true. If not included in the Tag RAM, this function requires an additional 1-bit wide SRAM.

The reset feature of earlier Tag RAMs was sufficient for implementation of a valid bit, but nothing more. Today's secondary caches frequently implement four-state write-back protocols such as MESI, with multiprocessor applications requiring five states (e.g. MOESI) or more. Hence, most caches need a two- or three-bit status memory that is accessed separately from the tag memory. It is used in conjunction with the match output to determine the response to a CPU memory access or a snoop. (A snoop is an operation initiated by the system in order to maintain coherency between the cache(s) and main memory.) This has typically been handled with yet another RAM - a conventional separate I/O SRAM organized as either x1 or x4. The 71215/16 includes a three-

bit status memory on chip.

THE 71215 AND 71216

As shown in Figure 3, these 16K x 15 RAMs are configured internally as two memories: 16K x 12 for tag and 16K x 3 for status. These two memories share the address bus but are controlled independently. An important new feature is extra pins and logic for generating $\overline{\text{BRDY}}$ (Intel's Burst Ready) and $\overline{\text{TA}}$ (PowerPC's Transfer Acknowledge). These are CPU input signals which are time critical in zero wait state secondary caches. I/O's are 3.3V compatible and there is a low power standby mode. All writes are synchronous as with burst data SRAMs, while all reads and compares are asynchronous for minimum delay. Two opposite polarity chip select pins are provided for easy depth expansion.



3176 drw 03

Figure 3. Simplified 71215 / 71216 Block Diagram (71216 signal names are in parenthesis)

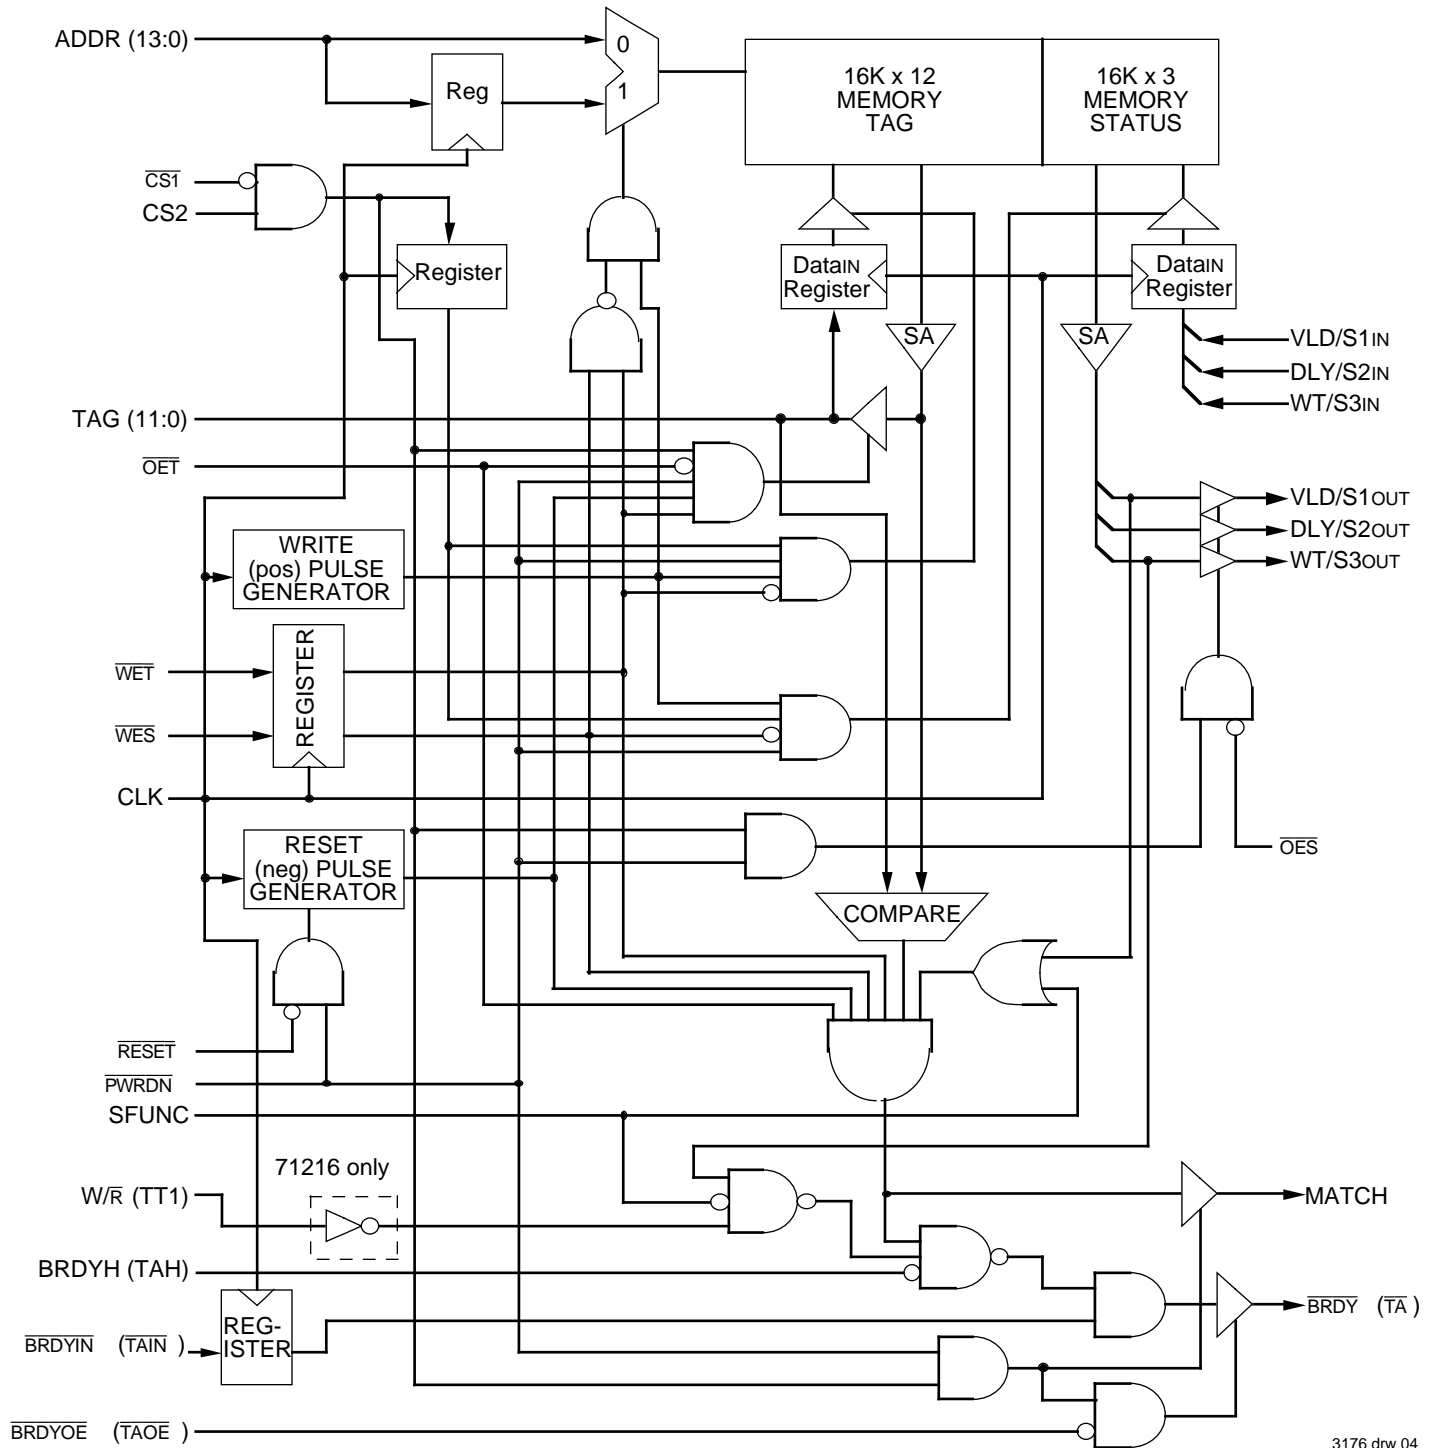
With a 16K x 12 tag memory, the 71215 and 71216 are wider and deeper than most Tag RAMs. For a typical 64-bit CPU with a 32-byte line size, the 16K depth supports a 512KB cache while the 12-bit tag field supports 2GB of cacheable main memory. Thus, only a single component is required for most applications. Table 1 shows the relationships between Tag RAM size, cache size, and cacheable main memory size. The Tag depth is equal to the cache size divided by the line size. The Tag width is equal to the base-2 log of the ratio of main memory size to cache size.

TABLE 1: REQUIRED TAG RAM SIZE AS A FUNCTION OF CACHE SIZE AND MAIN MEMORY SIZE (For 32-byte line size and direct mapped cache architecture.)

Cache Size	Cacheable Main Memory Size				
	64MB	256MB	1GB	2GB	4GB
128KB	4K x 9	4K x 11	4K x 13	4K x 14	4K x 15
256KB	8K x 8	8K x 10	8K x 12	8K x 13	8K x 14
512KB	16K x 7	16K x 9	16K x 11	16K x 12	16K x 13
1MB	32K x 6	32K x 8	32K x 10	32K x 11	32K x 12

3176 tbl 01

For a 1MB cache and 4GB of cacheable main memory, two of the devices may be cascaded in depth without any timing penalty apart from increased capacitive loading. This is accomplished with the two Chip Select pins. A low order address signal may be connected to CS1 on one chip and to CS2 on the other so that at any given time, one is selected and the other is deselected. The deselected chip ignores all control inputs (except RESET and PWRDN) and tri-states its outputs so that the two chips can be conveniently bussed together. As expected, worst case timing delays from the Chip Select inputs are the same as for the Address inputs. When only a single 71215 or 71216 is used in an application, CS1 is tied to Vss and CS2 is tied to Vcc.



3176 drw 04

Figure 4. Detailed 71215 / 71216 Block Diagram (71216 pin names are in parenthesis)

The 71215/16 is shown in more detail in Figure 4. The tag memory is controlled by the Write Enable Tag (\overline{WET}) and Output Enable Tag (\overline{OET}) pins. During writes, \overline{WET} is synchronous to CLK, as are the input data (TAG0 - TAG11) and address (A0 - A13). Note that \overline{WET} has no effect on the TAG output buffers, so \overline{OET} must be high to disable the outputs during writes. Reads are performed by deasserting \overline{WET} and asynchronously asserting \overline{OET} . For cache architectures in which the tag is never read (e.g. write-through caches), \overline{OET} may be tied to VCC. When both \overline{WET} and \overline{OET} are high, the 71215/16 is in the match mode, where the TAG0 - TAG11 inputs are compared with the stored data and are used to generate the MATCH and \overline{BRDY} / \overline{TA} outputs. In both read and

match modes, the address path is flow-through for the fastest possible response to a new address.

The three status bits of the 71215/16 are labeled VLD/S1, DTY/S2, and WT/S3. The reason for the dual names is that their functions vary, dependent on the state of the static Status Function (SFUNC) input signal. When SFUNC is low, the status bits are said to be in a “dedicated” mode and are referred to as Valid, Dirty and Write-Through. See Figure 5. When SFUNC is high, the status bits play no special role within the 71215/16 and are simply referred to as Status 1, Status 2 and Status 3. See Figure 6. The functionality of VLD and WT in the dedicated mode is described later. DTY/S2 does not have any special functionality within the 71215/16.

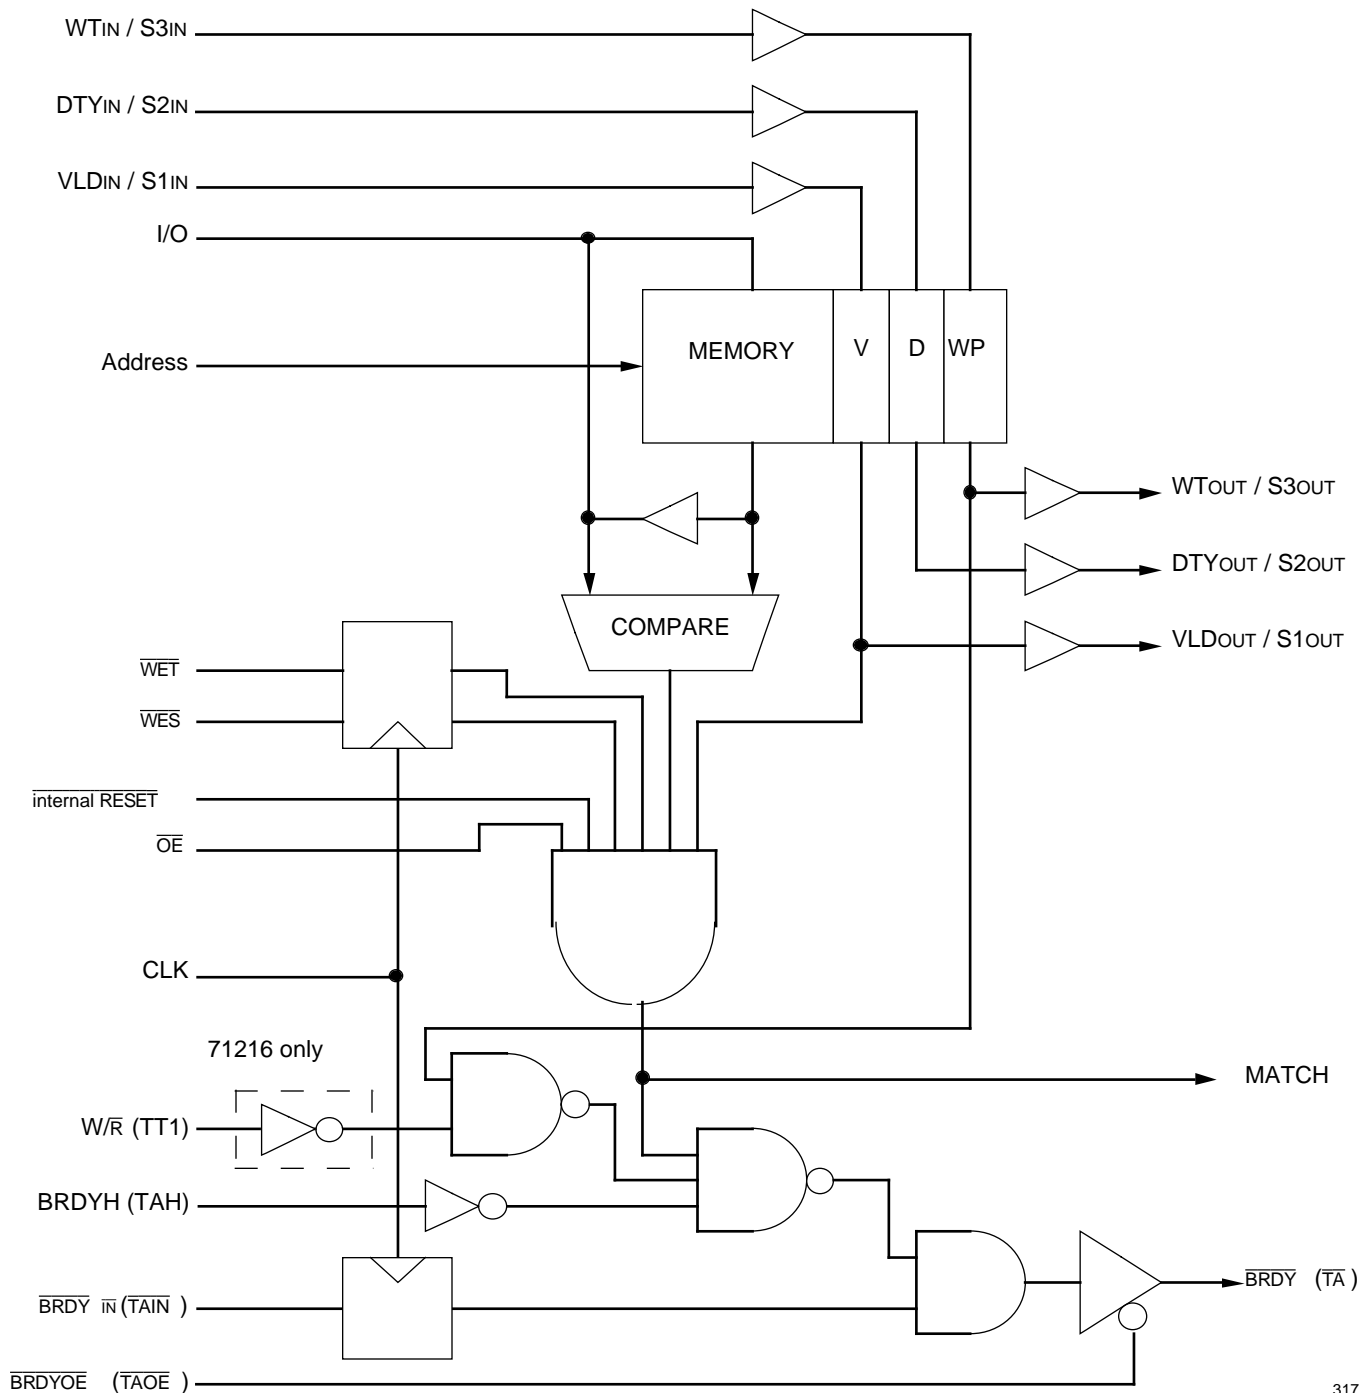
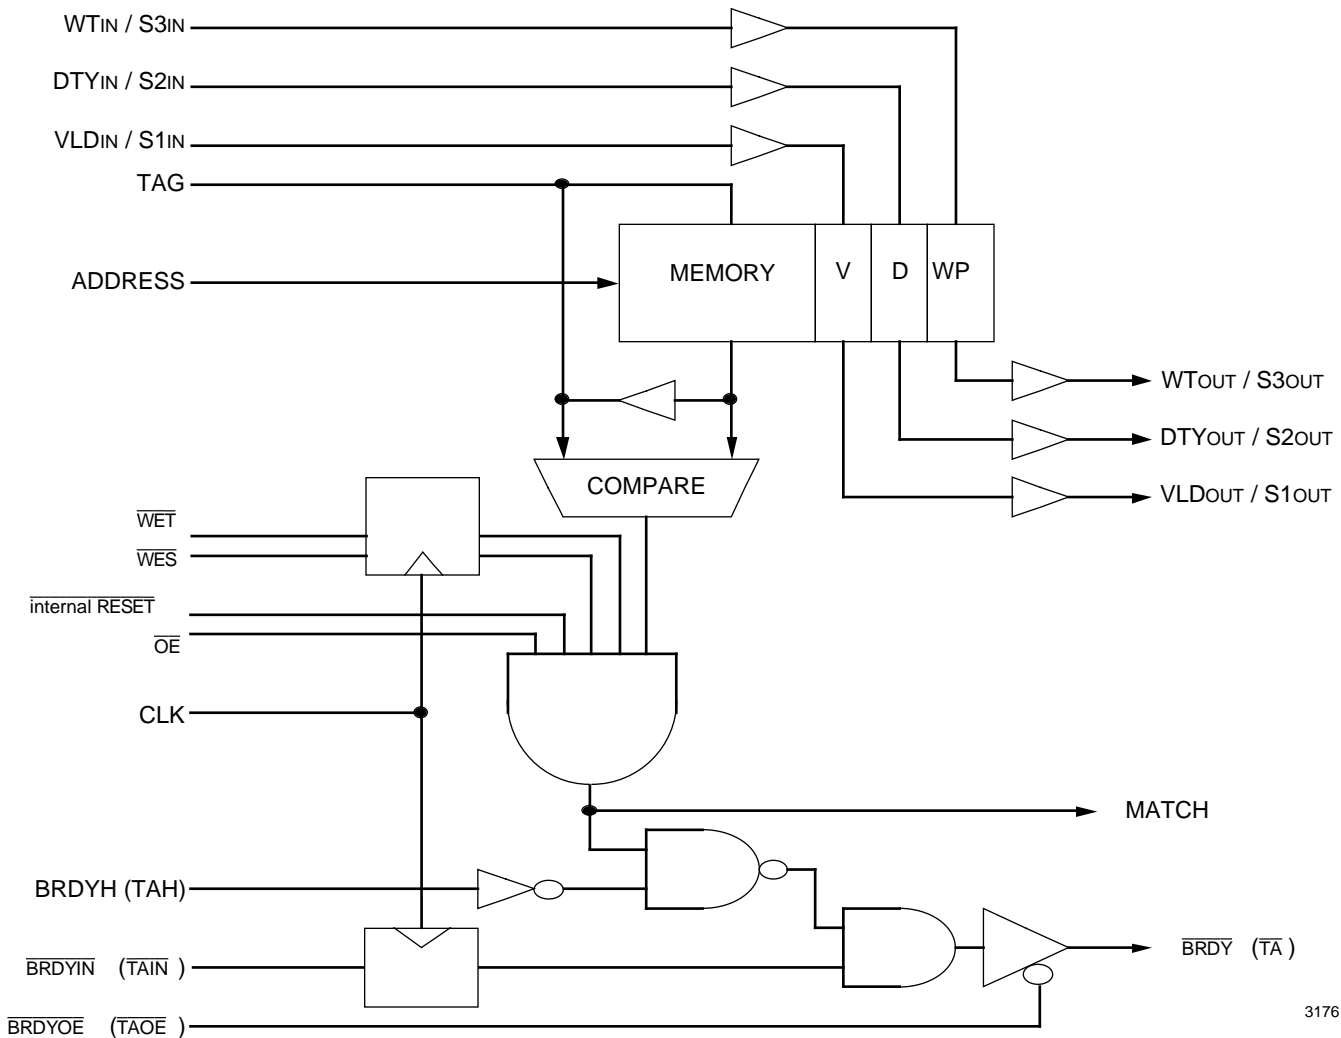


Figure 5. Dedicated Mode Logic (71216 pin names are in parenthesis)

3176 drw 05

BRDY and Match logic (Generic Status bit mode)
TA names added for ap note



3176 drw 06

Figure 6. Generic Mode Logic (71216 pin names are in parenthesis)

The status bits are accessed through separate input pins and output pins. This avoids the need for fast turn around on this bus as in the following example: a single word write hit to a write back line results in the need to set the state to dirty (also called “modified”). The status memory must go from reading to writing then back to reading in as little as two cycles. If common I/O is preferred, the user may tie the respective input and output pins together. The status memory control signals (\overline{WES} and \overline{OES}) are equivalent to \overline{WET} and \overline{OET} for the tag memory. Also, because the status field is separate I/O, \overline{OES} is normally tied to Vss to permanently enable the status outputs.

The tag and status memories are controlled independently since normal operation of the 71215/16 finds the tag memory in match mode and the status memory in read mode. Often, however, \overline{WET} and \overline{WES} are tied together in a design because the write function tends to be common between them. For those times when only the status bits need to be updated, \overline{WET} , \overline{WES} and \overline{OET} can be asserted together without having

to externally drive the TAG bus. This causes the data read from the tag field to be written back to the same address, resulting in no change to the tag data.

Note that there is only one address register that is used by both memory segments. The address is registered when either \overline{WET} or \overline{WES} is sampled low, and is flow-through when both \overline{WET} and \overline{WES} are sampled high.

The entire status memory is cleared to zeros when \overline{RESET} is sampled low on at least one rising edge of CLK. This can be used to put the cache into a known state after power up, or after a cache flush. Since reset is a type of write, \overline{WET} and \overline{WES} are required to be high during reset. \overline{PWRDN} must also be high, but the state of the chip select inputs does not matter. During reset, \overline{BRDY} / \overline{TA} is driven high, and MATCH is driven low.

MATCH, \overline{BRDY} AND \overline{TA}

As mentioned earlier, the 71215/16 is in match mode when \overline{OET} is high and \overline{WET} is sampled high. This allows the TAG0

- TAG11 inputs (high order bits of the CPU/system address bus) to be compared with internally stored data. When SFUNC is low, the stored VLD bit is combined with this comparator output to generate a MATCH output that is true only when both the tag comparison is true and the VLD bit is high. Thus, an invalid tag entry does not generate a hit. Note that \overline{OES} and \overline{WES} do not affect internal access of the status bits. When SFUNC is high, the status bits are generic and MATCH is simply the output of the comparator. MATCH is driven low when the 71215/16 is not in match mode. When the chip is deselected, MATCH becomes high impedance.

The cache/memory controller has traditionally generated the $\overline{BRDY} / \overline{TA}$ signal to the CPU, using MATCH and other inputs. This is a critical timing path. During a zero-wait-state lead-off, there are only two clock cycles for the CPU to drive the address and other bus signals, and for \overline{BRDY} or \overline{TA} to be returned to the CPU by the cache controller. See Figure 7. Typically there is not enough time to have two chips (Tag RAM and controller) in this timing path. The 71215/16 address this difficulty by incorporating logic for generating $\overline{BRDY} / \overline{TA}$, thereby removing the cache controller from this path. This is shown in Figure 8.

While the cache controller is removed from the primary $\overline{BRDY} / \overline{TA}$ timing path, it must still play a part in generating $\overline{BRDY} / \overline{TA}$. The controller has address and other bus cycle information that is needed to qualify the generation of $\overline{BRDY} / \overline{TA}$. This qualification logic is best placed in parallel with the tag lookup, rather than in series with it. Also, there are cases where the the generation of $\overline{BRDY} / \overline{TA}$ by the 71215/16 must be blocked so that the cache/memory controller can generate it instead.

Logic has been included in the 71215/16 that enables it to qualify $\overline{BRDY} / \overline{TA}$ for one particular case. While a write hit to a write-back line can be handled by the cache alone, a write-through line requires that the write also proceed to main memory. In the former case the cache can respond without wait states and $\overline{BRDY} / \overline{TA}$ is driven low immediately as the result of a tag match and set VLD bit. In the latter case, main memory writes normally require wait states. If a line (or the whole cache) is write-through, the 71215/16 should not drive $\overline{BRDY} / \overline{TA}$ low, so that the cache/memory controller may do so later when the main memory (or write buffer) write is complete. When the Tag RAM is in dedicated status mode (SFUNC low), the stored WT bit determines whether the line is write-through (high) or write-back (low). Note that it may also be used to denote a write protected line. Another pin - W/\overline{R} and TT1 on the 71215 and 71216 respectively - connects directly to the CPU for distinguishing between processor reads and processor writes. These two bits of information are used to block internal generation of $\overline{BRDY} / \overline{TA}$ during a processor write to a write-through line. Without this feature, the cache controller might not have enough time to generate a blocking signal (as described below) based on the WT output from the Tag RAM. If a user wants to gate the VLD bit with MATCH but not use the WT bit in combination with W/\overline{R} or TT1, he should select the dedicated mode (SFUNC low) and tie W/\overline{R} low or TT1 high. Note that the one functional difference between the 71215 and 71216 is the polarity of the W/\overline{R} and TT1 signals.

The cache controller may have additional information and may wish to delay the assertion of $\overline{BRDY} / \overline{TA}$. Thus, the 71215 and 71216 have input pins - BRDYH and TAH respectively - so that the cache controller may force $\overline{BRDY} / \overline{TA}$ high, regardless of the result of the tag comparison inside the 71215/16.

In the case of a cache miss or write through, the system memory controller (usually combined with the cache controller) becomes completely responsible for generating \overline{BRDY} or \overline{TA} for that bus cycle. For flexibility, the 71215/16 incorporates two options for merging its own $\overline{BRDY} / \overline{TA}$ output with that generated by the system memory controller.

One approach is to bus the two signals together. This is the preferred approach when the cache (including the 71215/16) is optional, as on a module, since addition or removal of the cache does not affect the way in which the cache controller generates $\overline{BRDY} / \overline{TA}$. Figure 9 shows this approach for the 71215 used with the Pentium. It applies equally to the 71216 and PowerPC. This requires that both $\overline{BRDY} / \overline{TA}$ sources be tri-statable. The \overline{BRDYOE} and \overline{TAOE} input pins of the 71215 and 71216 are driven by the cache/memory controller, and are used to enable or disable the 71215/16 $\overline{BRDY} / \overline{TA}$ output as necessary. To be prepared for a possible hit, each new bus cycle begins with $\overline{BRDYOE} / \overline{TAOE}$ low. In the event of a cache miss, the controller deasserts $\overline{BRDYOE} / \overline{TAOE}$, then takes over responsibility for driving $\overline{BRDY} / \overline{TA}$. This is also the procedure for writes to write-through lines, where even cache hits are responded to by the controller. Also, the controller usually takes over control of $\overline{BRDY} / \overline{TA}$ for the second, third and fourth words of a burst transfer. This is required if either the CPU address is not guaranteed to remain valid throughout the entire bus cycle (a change to the 71215/16 address bus

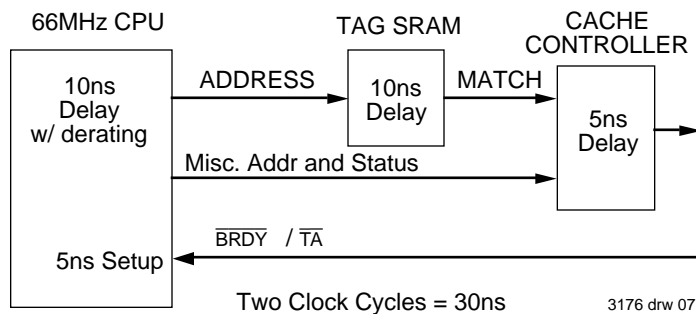


Figure 7. Conventional Tag RAM Usage - Chip Set in BRDY/TA Critical Path

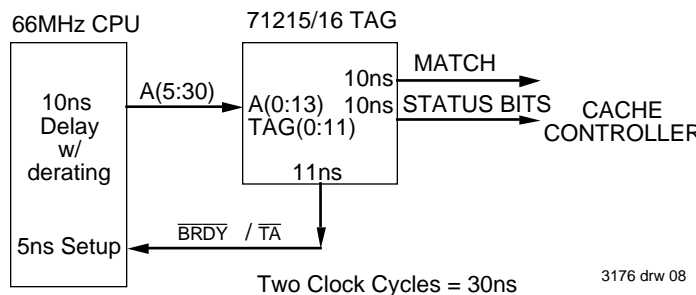
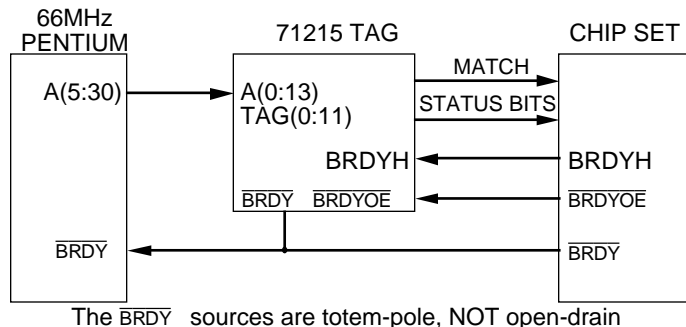


Figure 8. 71215/16 BRDY/TA Timing - Chip Set Removed from the Critical Path


will ripple through to $\overline{BRDY} / \overline{TA}$), or if data SRAM accesses require wait states (it's necessary to toggle $\overline{BRDY} / \overline{TA}$ from cycle to cycle). For those times when neither device is driving $\overline{BRDY} / \overline{TA}$, a pull-up resistor is used to keep the signal high. In this case, it's suggested that the controller drive $\overline{BRDY} / \overline{TA}$ high before putting it in a high impedance state. Thus, the resistor is never used to generate a low to high transition and therefore can be weak (3 K Ω to 20 K Ω). Also, both the 71215/16 and controller can remain off the $\overline{BRDY} / \overline{TA}$ bus for extended periods of time if so desired. With this approach, \overline{BRDYIN} or \overline{TAIN} (Burst Ready Input, Transfer Acknowledge Input) is tied high.



3176 drw 09

Figure 9. Combining BRDY/TA : Bussing Option

The second approach is to have the cache/memory controller drive it's $\overline{BRDY} / \overline{TA}$ output into the $\overline{BRDYIN} / \overline{TAIN}$ input on the 71215/16 at all times. Inside the 71215/16, $\overline{BRDYIN} / \overline{TAIN}$ is registered by the clock then ANDed (negative logic ORed) with the internally generated $\overline{BRDY} / \overline{TA}$. For this approach, $\overline{BRDYOE} / \overline{TAOE}$ is tied permanently low. The controller no longer generates $\overline{BRDYOE} / \overline{TAOE}$, but instead must generate $\overline{BRDY} / \overline{TA}$ one cycle earlier because it is delayed by one cycle in reaching the CPU. Note that $\overline{BRDYH} / \overline{TAH}$ only enables or disables the $\overline{BRDY} / \overline{TA}$ generated inside the 71215/16, and does not affect the propagation of $\overline{BRDYIN} / \overline{TAIN}$ through to the $\overline{BRDY} / \overline{TA}$ output. Figure 10 shows this approach for the 71215 and Pentium.



3176 drw 10

Figure 10. Combining BRDY/TA : Pass-Through Option

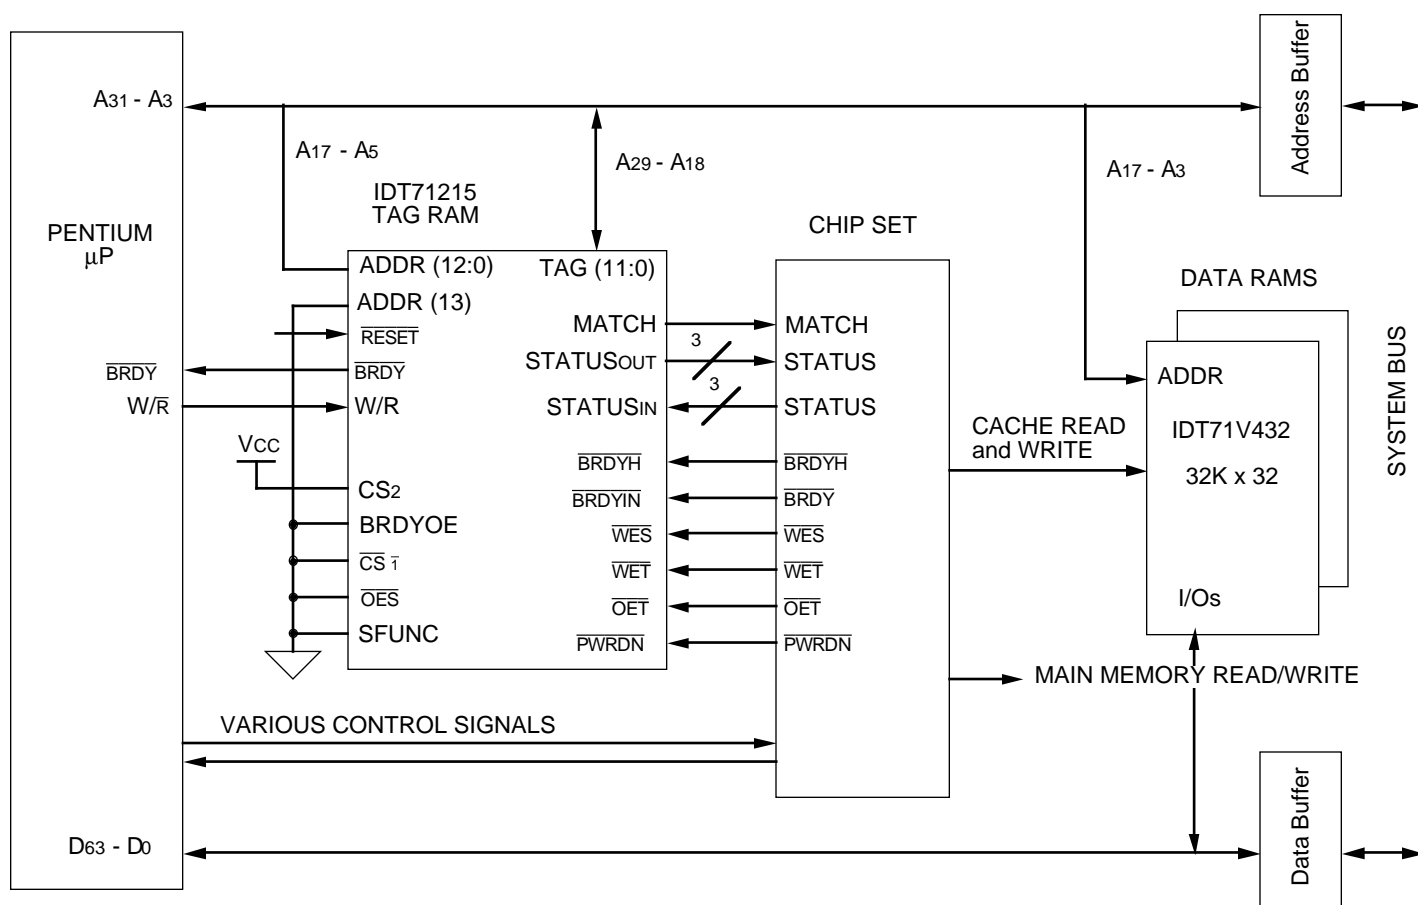
$\overline{BRDY} / \overline{TA}$ functions similar to MATCH (but opposite in polarity) when the 71215/16 is not in match mode. It is high impedance when the chip is deselected (or $\overline{BRDYOE} / \overline{TAOE}$ is high), and otherwise is driven high when out of match mode.

REDUCED POWER

For the increasing number of applications that require a low power standby mode, the 71215/16 includes an asynchronous power down pin (\overline{PWRDN}). When it is driven low, both the tag and status memories are shut down to save considerably on power consumption. For optimum power savings, all input and bidirectional signals should also be held at CMOS voltage levels (near VCC or VSS). During power down, all outputs are placed in a high impedance state and all data is retained. All writes should be allowed to complete before \overline{PWRDN} is asserted. There is no minimum time that it must be low. When exiting the power down state, there is only a very short delay after the rising edge of \overline{PWRDN} before normal activity can be resumed.

SYSTEM USAGE

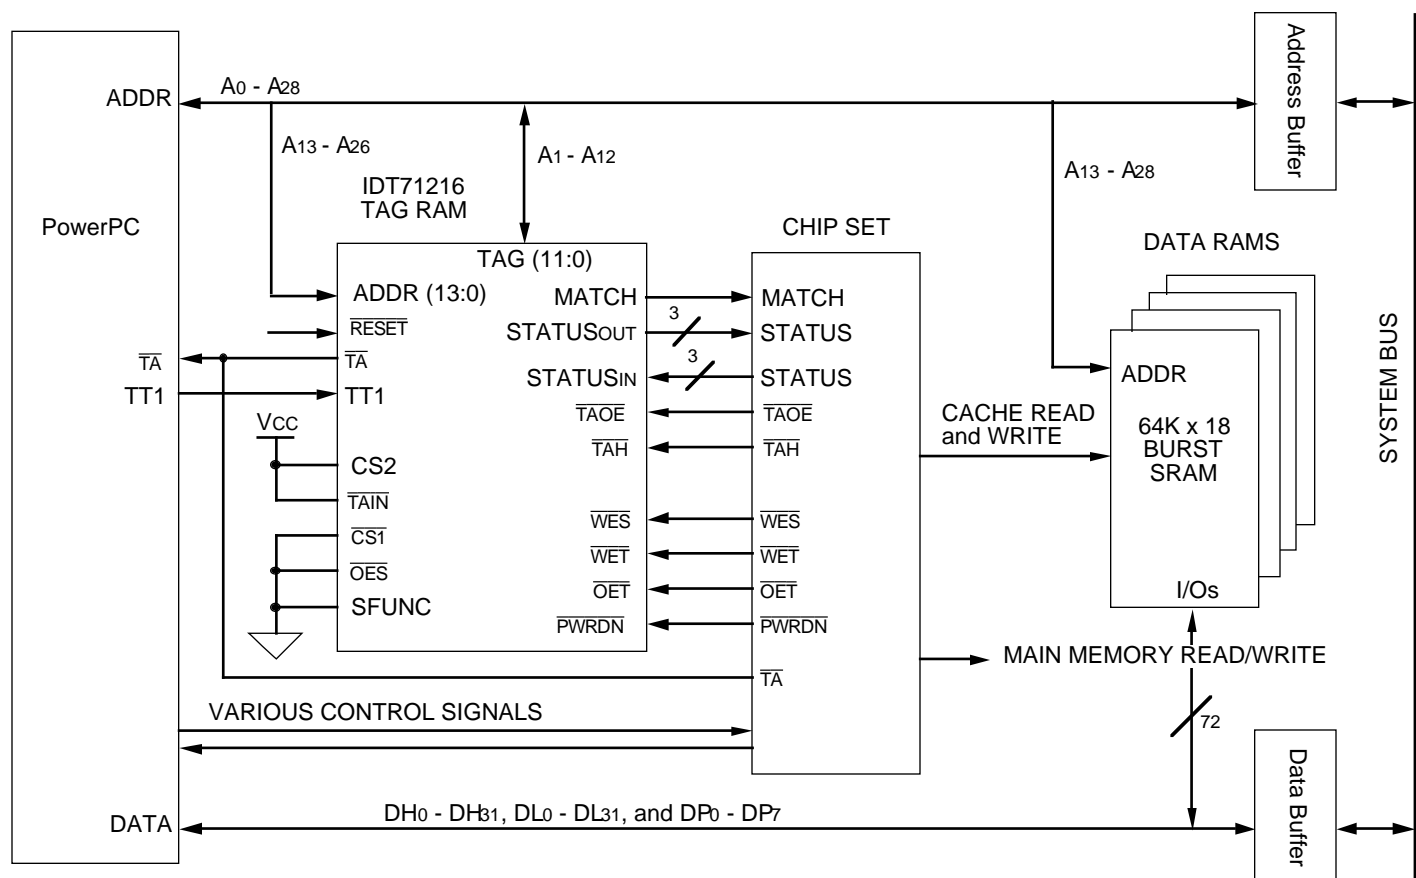
For applications not using the entire 12-bit tag field, the unused TAG I/O pins should be pulled either high or low through 1 K Ω to 5 K Ω resistors. For applications not using the entire 3-bit status field, the unused inputs may be tied directly to VCC or VSS, and the unused outputs are left unconnected. All other unused inputs should be tied either to VCC or VSS as appropriate for their function. This includes unused address signals if only part of the depth of the 71215/16 is used.



3176 drw 11

Figure 11: Pentium / 71215 Example of 256 KB Cache

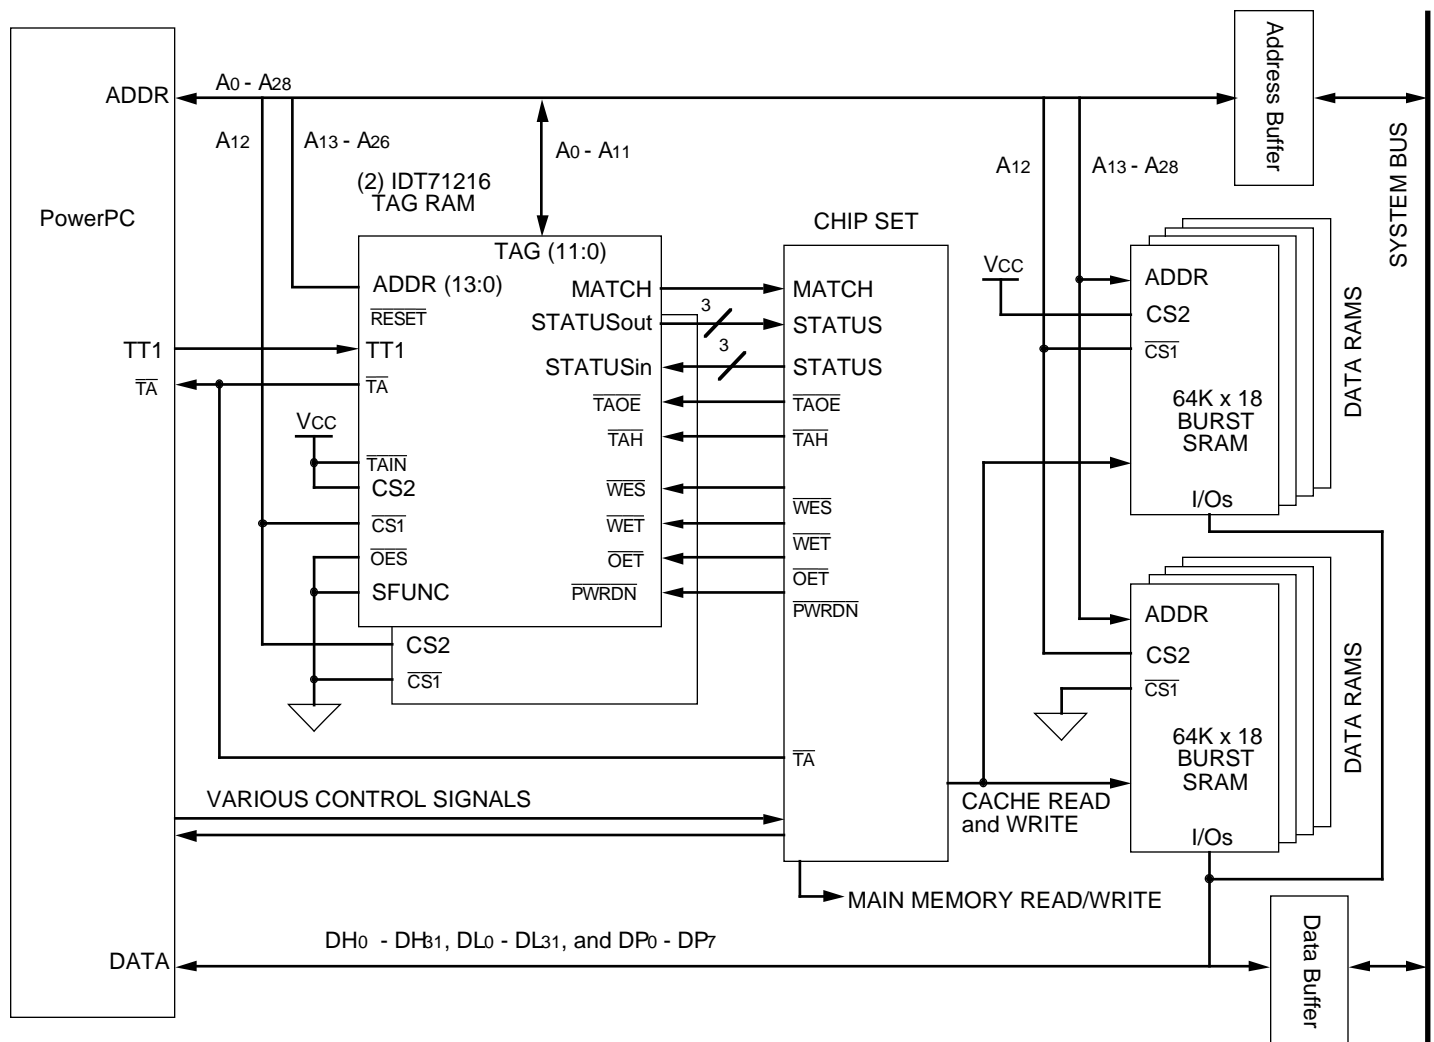
When changing the depth of the cache (and Tag RAM), the TAG field shifts accordingly, so that it remains contiguous with the address field. The example in Figure 11 uses only half of a 71215, and can be compared with Figures 1 and 12. It shows a 256KB cache for the Pentium, uses the $\overline{\text{BRDY}}$ pass-through option from Figure 10, and maps 1GB of main memory into the cache. Two 32Kx32 burst SRAMs are shown as the cache data RAM.



3176 dr w 12

Figure 12: PowerPC / 71216 Example of 512 KB Cache

Figure 12 shows a 512KB cache implementation for the PowerPC, using the full address range of the 71216. This example uses a bussed \overline{TA} implementation shown in Figure 9. The tag size is sufficient to support 2GB of cacheable main memory.



3176 drw 13

Figure 13: PowerPC / 71216 Example of 1 MB Cache

Figure 13 shows a 1MB cache for the PowerPC using the 71216. The implementation is essentially the same as for 512KB, but with two 71216 Tags and two banks of data SRAMs. Except for $\overline{CS1}$ and CS2, all the same signals that were connected to the first Tag RAM should be connected to the same pins of the second Tag RAM. The least significant tag bit of the 512KB cache is used to select between the two Tag RAMs of the 1MB cache. The same is true for the two banks of data SRAMs. The tag field then shifts one bit in the direction of the more significant address bits. Please note that the PowerPC and Intel processors do not have the same address sequence. A0 is the MSB for the PowerPC while A31 is the MSB for Intel's processors.

It is also possible to double the size of the cache and cached address space without doubling up the Tag RAMs. This can be done by doubling the line size of the cache - from 32 bytes to 64 bytes, for example. It is not necessary to have the same line size for both the primary and secondary caches, though it does simplify the cache controller. A more detailed discussion of this topic is beyond the scope of this application note.

The CLK pin should be driven by the same clock that drives the CPU. Although there is no standard for clock skew tolerances between devices, a recommended target is $\pm 1\text{nS}$.

MESI PROTOCOL IMPLEMENTATION

MESI is a cache coherency protocol, implemented in the primary cache of both the PowerPC 601 and the Pentium Processor. With the 71215/16, it is now practical to also implement MESI for the L2 cache. The acronym stands for Modified (write-back data that is dirty), Exclusive (clean write-back data that can later transition to Modified), Shared (write-through data which cannot become Modified) and Invalid. In short, it allows for cache lines to be individually marked as

either write-through or write-back. While the cache controller is responsible for implementing the protocol and controlling the state transitions, the 71215/16's features can be helpful in the implementation.

The following state assignments for MESI are intended to take advantage of the features of the 71215/16 when it is in the dedicated status mode (SFUNC low). Some variations are possible:

TABLE 2: SUGGESTED MESI STATE ASSIGNMENTS

	VLD/S1	DTY/S2	WT/S3
Invalid	0	X	X
Shared	1	0	1
Exclusive	1	0	0
Modified	1	1	0

3176 tbl 02

As described earlier, $\overline{\text{BRDY}} / \overline{\text{TA}}$ generation is blocked when VLD/S1 is low, and during write hits to write-through lines. The cache controller is responsible for all state transitions, including Exclusive to Modified.

SUMMARY

The 71215 and 71216 represent a major step forward in Tag RAMs. They are sized appropriately for the majority of today's cache and main memory requirements, and offer new features that help remove many of the barriers to the implementation of zero wait-state caches.

As this application note is being written, the fastest speed grade of the 71215 and 71216 is 9nS (address to match time), with faster speeds expected in the future. Please contact your local IDT sales office or representative for information on the latest speed availability.