



Tech Info Library

Workgroup Server 9150: External Cache (4/94)

Revised: 4/26/94
Security: Everyone

Workgroup Server 9150: External Cache (4/94)

Article Created: 26 April 1994

TOPIC -----

This article describes the external cache of the Workgroup Server 9150 (WS 9150).

DISCUSSION -----

The optional external cache in the WS 9150 is a direct mapped, write through, 2nd level (L2) cache capable of 2-1-1-1 operation. It is selectable to either unified or instruction only operation, and can be a variety of sizes by the choice of components installed in the cache SIMM. Two pins on the SIMM socket are read by the High speed Memory Controller (HMC) at reset, and are intended to indicate the cache size to the Operating System software so meaningful diagnostics can be performed. If the SIMM is not present, pullup resistors on these pins disable all L2 cache related operations. If the SIMM is installed, the cache can be enabled or disabled from the HMC control register, although updates will still be performed to insure coherency.

Tag Organization

The L2 cache tags all 32 bits of the PowerPC 601 processor's physical address. Address bits A(0:14) comprise the tag, A(15:26) index the set location, and A(27:31) form the line offset as shown in figure 4. In addition to A(0:14) the tag also compares a valid bit. The valid bit is cleared by a system reset and set by a cache line allocate operation.




Figure 4: Address Tag Organization

Cache Coherency

The L2 cache maintains a single valid bit for each 32 byte line. The valid bits are cleared by a system reset and individually set when a line is allocated in the cache. When the L2 cache is operating in the unified mode, a cache line is never invalidated other than system reset.

The PowerPC 601 processor broadcasts primary cache manipulation instructions so that other caches in the system can mirror the functionality of the instruction.

Since the HMC makes the assumption that the L2 cache is bound to main memory rather than the processor, it ignores these broadcasts. Executing these invalidate memory cycles potentially discards useful data. In order to keep the L2 cache coherent without invalidating, all PowerPC 601 processor's write transactions, force a tag compare. Cache-inhibited write data is tag compared due to the possibility that the data was formerly cache-enabled and allocated.

Configuring the L2 cache as instruction-only alters the line allocation policy but does not affect the write-update procedure. This prevents stale instruction data that could be caused by relocating program code.

Whenever the tag matches on a write cycle, data is updated in the cache simultaneously with the write-buffer being filled. Bus cycles that write less than the full 64 bit data bus width, utilize the cache data RAM synchronous write-enables to select the appropriate byte lanes. To permit writes to execute as quickly as possible, the cache data RAM write-enables are qualified with the state of the tag match. As a result, the first beat of a two-clock write that misses the cache will be aborted asynchronously.

Since all 32 address bits are significant to the cache, but not to the physical memory decode (there are physical aliases of both DRAM and ROM), access of physical memory aliases is not recommended as it can result in stale data in the L2 cache. This feature allows cache testing.

Line Allocation

Cache lines are allocated only on PowerPC 601 processor's burst read memory cycles. Read data is forwarded to the CPU in conjunction with the line fill. The L2 cache is write-through so write transactions are never allocated. Setting the HMC control register bit, L2_INST, forces the L2 to cache only instruction fetches. The PowerPC 601 processor's TC(0) output distinguishes whether the bus read data is instruction or data.

The L2 cache attempts to allocate cache lines similarly to the PowerPC 601 processor primary cache. The PowerPC 601 processor bus transaction modifiers that affect line allocation are cache-inhibit and transfer-burst. Cache-inhibiting an address space prevents the L2 cache from allocating lines in that address space. Thus memory areas such as DMA buffer space that might be marked cache-inhibited will not thrash the L2 cache. The HMC also watches the transfer-burst signal for single-beat transactions. Because the granularity of the L2 cache valid data is 32 bytes, single-beat transactions are not allocated.

Additionally, single beat reads are never looked up in the cache even if the cache-inhibit line is de-asserted.

The cacheable address spaces are RAM at \$00000000-\$3FFFFFFF, Macintosh ROM at \$40000000-\$4FFFFFFF, and PowerPC ROM at \$FF000000-\$FFFFFFF. Other address locations, such as the PDS (processor direct slot), are not cached in the L2, but may be cached in the PowerPC 601 processor's primary cache.

Testing

Since all of the address bits are tagged, aliases of physical RAM space will appear distinct to the cached, but not to uncached data. This allows testing to determine proper cache operation. Normally the L2 cache ignores ROM writes. However, by setting the L2ROMW bit, ROM writes will update the cache. For normal ROMs, the cache will now be inconsistent with memory, allowing ROM caching to be tested.

Support Information Services

Copyright 1994, Apple Computer, Inc.

Tech Info Library Article Number:15215