Secondary Cache Interface

13

The R4000SC and R4000MC versions of the R4000 processor contain interface signals for an optional external secondary cache. This interface consists of:

- a 128-bit data bus
- a 25-bit tag bus
- an 18-bit address bus
- various static random access memory (SRAM) control signals.

The 128-bit-wide data bus minimizes the primary cache miss penalty, and allows the use of standard low-cost SRAMs in the design of the secondary cache.

The remainder of the System interface signals are described in Chapter 8.

13.1 Data Transfer Rates

The interface to the secondary cache maximizes service of primary cache misses. The Secondary Cache interface, **SCData(127:0)**, supports a data rate that is close to the processor-to-primary-cache bandwidth during normal operation. To ensure that this bandwidth is maintained, each data, tag, and check pin must be connected to a single SRAM device.

The **SCAddr** bus, together with the **SCOE***, **SCDCS***, and **SCTCS*** signals, drives a large number of SRAM devices; because of this, one level of external buffering between the processor and the cache array is used.

13.2 Duplicating Signals

The buffered control signals control the speed of the Secondary Cache interface. Critical control signals are duplicated by design to minimize this limitation: the **SCWR*** signal and **SCAddr(0)** have four versions so that external buffers are not needed to drive them. When an 8-word (256-bit) primary cache line is used, these signals can be controlled quickly, reducing the time of back-to-back transfers.

Each duplicated control signal can drive up to 11 SRAMs; therefore, a total of 44 SRAM packages can be used in the cache array. This allows a cache design using 16-Kbyte-by-64-bit, 64-Kbyte-by-4-bit, or 256-Kbyte-by-4-bit standard SRAM.†

The benefit of duplicating **SCAddr(0)** is greater in systems that use fast sequential static cache RAM and an 8-word primary cache line. If **SCAddr(0)** is attached to the SRAM address bit that affects column decode only, the read cycle time should approximate the output enable time of the RAM. For fast static RAM, this cycle time should be half of the nominal read cycle time.

† Other cache designs within this constraint are also acceptable. For example, a smaller cache design can use 22 8-Kbyte-by-8-bit static RAMs; this design presents less load on the address pins and control signals, and reduces the overall parts count.

13.3 Accessing a Split Secondary Cache

When the secondary cache is split into separate instruction and data portions, assertion of the high-order **SCAddr** bit, **SCAddr(17),** enables the instruction half of the cache.

It is possible to design a cache that supports both joint and split instruction/data configurations of less than the maximum cache size; in doing so, **SCAddr(12:0)** must address the cache in all configurations. **SCAddr(17)** must support the split instruction/data configuration, and any of **SCAddr(16:14)** bits can be omitted, because of the fixed width of the physical tag array.

13.4 SCDChk Bus

The secondary cache data check bus, **SCDChk**, is divided into two fields to cover the upper and lower 64 bits of **SCData**. This form is required by the 64-bit width of internal data paths.

13.5 SCTAG Bus

The secondary cache tag bus, **SCTag**, is divided into three fields, as shown in Figure 13-1. The *CS* field indicates the cache state: invalid, clean exclusive, dirty exclusive, shared, or dirty shared. The *PIdx* field is an index to the virtual address of primary cache lines that can contain data from the secondary cache. Bits 18:0 contain the upper physical address.

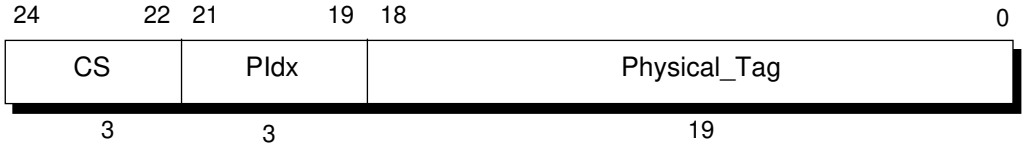

Figure 13-1 SCTag Fields

The **SCDCS*** and **SCTCS*** signals disable reads or writes of either the data array or tag array when the opposite array is being accessed. These signals are useful for saving power on snoop and invalidate requests since access to the data array is not necessary. These signals also write data from the primary data cache to the secondary cache.

13.6 Operation of the Secondary Cache Interface

The secondary cache can be configured for various clock rates and static RAM speeds. All configurable parameters are specified in multiples of **PClock**, which runs at twice the frequency of the external system clock, **MasterClock**.

During boot time, secondary cache timing parameters are programmed through the boot-time mode bits, as described in Chapter 9. Table 13-1 lists the secondary cache timing parameters. The following sections describe secondary cache read and write cycles.

Table 13-1 Secondary Cache Timing Parameters

Symbol	Number of Cycles
t_{Rd1Cyc}	4-15 PCycles
t_{Rd2Cyc}	2-15 PCycles
t_{Dis}	2-7 PCycles
t_{Wr1Dly}	1-3 PCycles
t_{Wr2Dly}	1-3 PCycles
t_{WrRC}	0-1 PCycles
t_{WrSUp}	3-15 PCycles

Read Cycles

There are two basic read cycles: 4-word read and 8-word read.

Each secondary cache read cycle begins by driving an address out on the address pins. The output enable signal **SCOE*** is asserted at the same time.

This section describes both 4-word and 8-word read cycles, including timing diagrams.

4-Word Read Cycle

The 4-word read cycle has two user-accessible timing parameters:

t_{Rd1Cyc}	read sequence cycle time, which specifies the time from the assertion of the **SCAddr** bus to the sampling of the **SCData** bus
t_{Dis}	cache output disable time, which specifies the time from the end of a read cycle to the start of the next write cycle

Figure 13-2 illustrates the 4-word read cycle, including the two user-accessible timing parameters.

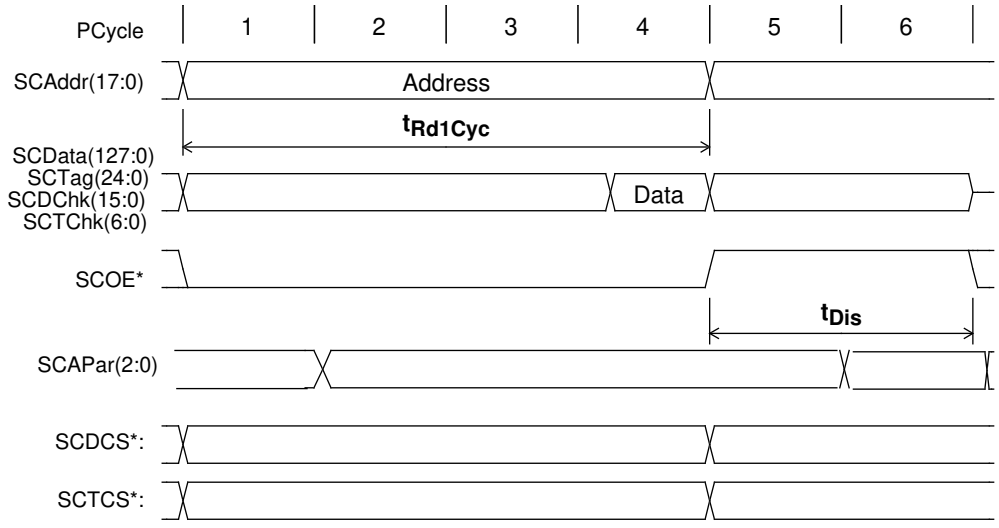

Figure 13-2 Timing Diagram of a 4-Word Read Cycle

8-Word Read Cycle

The 8-word read cycle has an additional user-accessible parameter beyond that of the 4-word read cycle described above: t_{Rd2Cyc}, the time from the first sample point to the second sample point.

In an 8-word read cycle, the low-order address bit, **SCAddr(0)**, changes at the same time as the first read sample point.

Figure 13-3 illustrates the 8-word read cycle, including the three user-accessible timing parameters.

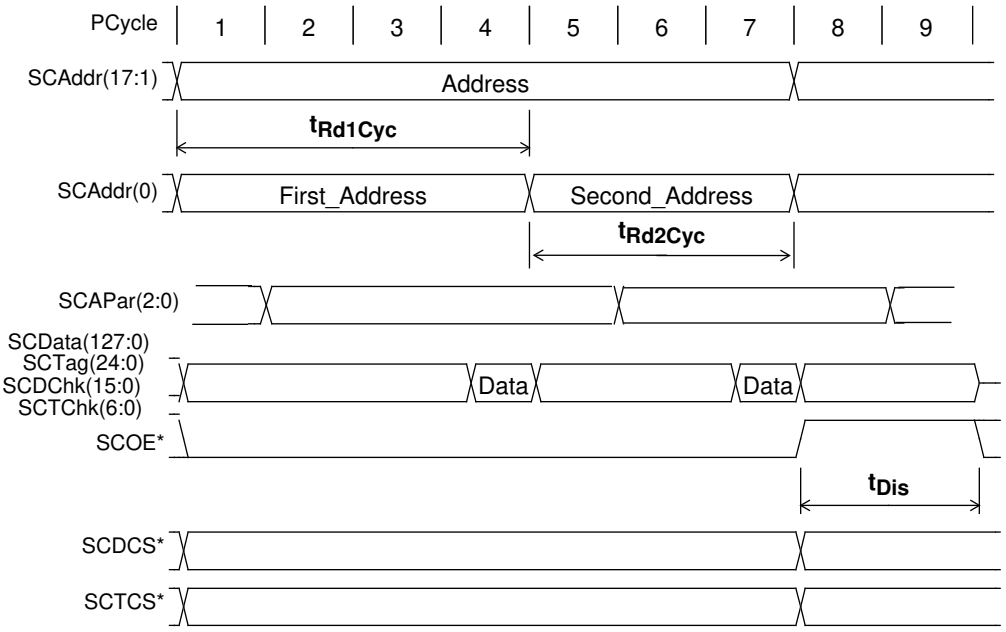

Figure 13-3 Timing Diagram of an 8-Word Read Cycle

Notes on a Secondary Cache Read Cycle

All read cycles can be aborted by changing the address; a new cycle begins with the edge on which the address is changed. Additionally, the period t_{Dis} after a read cycle can be interrupted any time by the start of a new read cycle. If a read cycle is aborted by a write cycle, **SCOE*** must be deasserted for the t_{Dis} period before the write cycle can begin.

Read cycles can also be extended indefinitely. There is no requirement to change the address at the end of a read cycle.

Write Cycles

There are two basic write cycles: a 4-word write cycle and an 8-word write cycle. The secondary cache write cycle begins with the assertion of an address onto the address pins.

This section describes both 4-word and 8-word write cycles, including timing diagrams.

4-Word Write Cycle

A 4-word write cycle has three timing parameters:

t_{Wr1Dly}	delay from the assertion of the address to the assertion of **SCWR***
t_{WrSUp}	delay from assertion of the second data doubleword to the deassertion of **SCWR***
t_{WrRc}	delay from the deassertion of **SCWR*** to the beginning of the next cycle

The timing parameter t_{WrRc} is 0 for most cache designs. Note that the upper data doubleword and the lower data doubleword are normally driven one cycle apart; this reduces the peak current consumption in the output drivers.

Figure 13-4 illustrates the 4-word write cycle. Either the upper or lower data doubleword can be driven first.

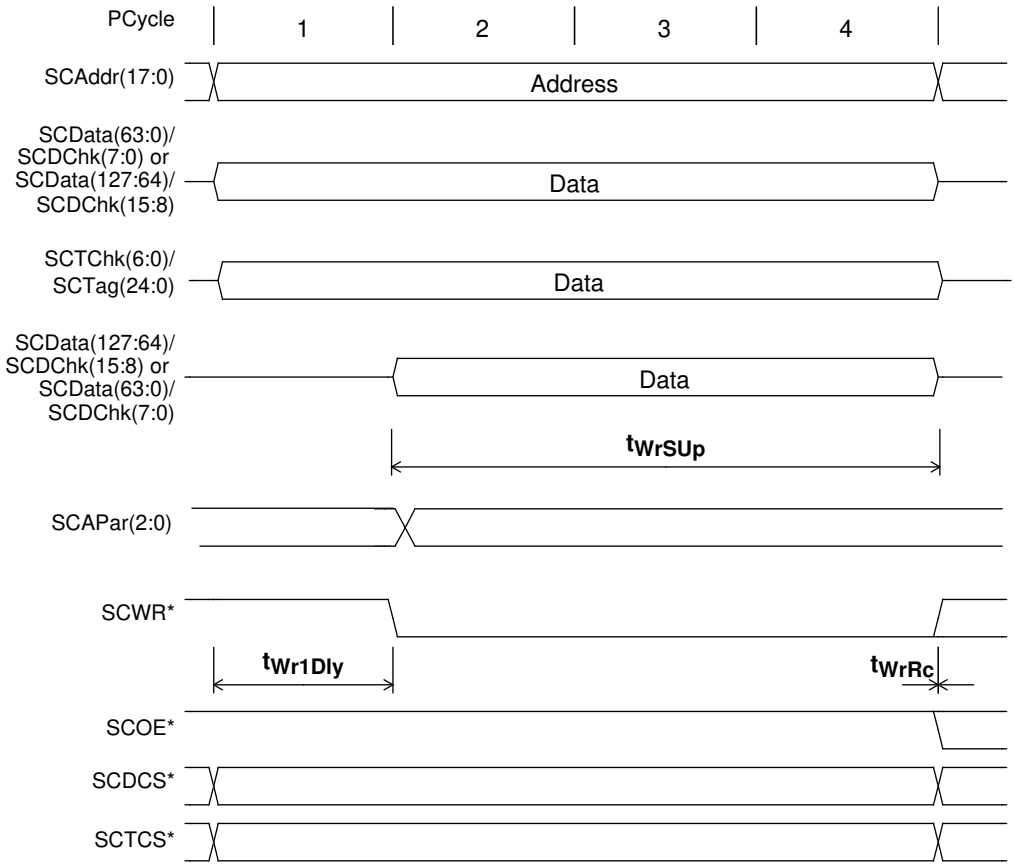

Figure 13-4 Timing Diagram of a 4-Word Write Cycle

8-Word Write Cycle

An 8-word write cycle has one additional parameter beyond those used by the 4-word write cycle: t_{Wr2Dly}. This is the time period that begins when the low-order address bit **SCAddr(0)** changes and ends when **SCWR*** is asserted for the second time. The lower half of **SCData** is driven on the same edge as the change in **SCAddr(0)**.

Figure 13-5 illustrates the 8-word write cycle.

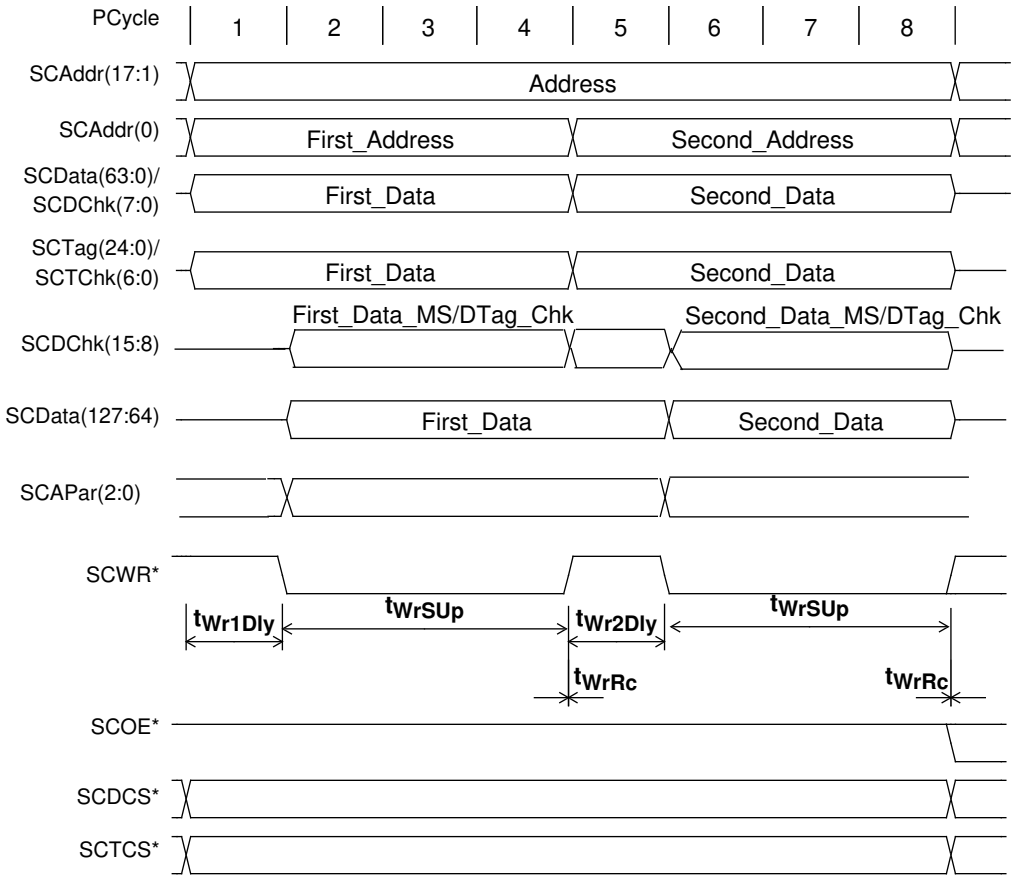

Figure 13-5 Timing Diagram of an 8-Word Write Cycle

Notes on a Secondary Cache Write Cycle

When receiving data from the System interface, the first data doubleword can arrive several cycles before the second data doubleword. In this case, the cache state machine enters a wait-state that extends **SCWR*** until t_{WrSUp} period after the second data item is transmitted.