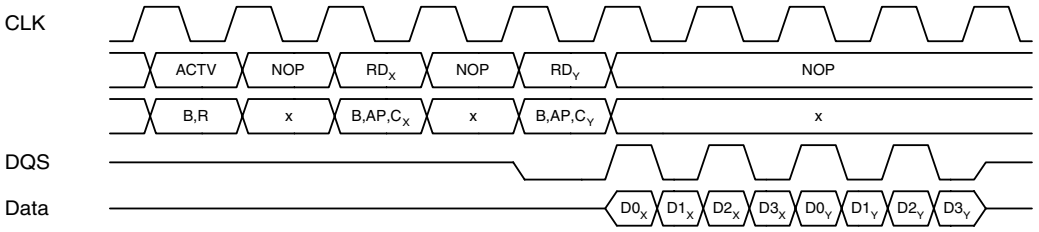

FIGURE 8.9 Consecutive DDR SDRAM reads (CL = 2, BL = 4).

returns to tri-state. Reads may be executed consecutively to achieve high bus utilization by hiding row activation and CAS latency delays, as with an SDR SDRAM.

Unlike an SDR SDRAM, writes are subject to a brief latency between assertion of the write command and delivery of write data. The first write data is presented on the first rising edge of DQS following the write command. DQS is not driven to the SDRAM until just after the write command is presented. This restriction prevents a collision between the SDRAM and the memory controller when a write follows a read by giving time for the SDRAM to turn off its DQS driver. Following the write, DQS can remain driven until a read command is asserted, at which time the SDRAM will need to drive the strobe. Write timing is shown in Fig. 8.10. Writes may also be executed consecutively to more effectively utilize the device interface.

When transitioning between reading and writing, minimum delays are introduced in a situation unlike that of a conventional SDRAM. Because write data lags the write command by a clock cycle, a cycle is lost when following a read with a write, because the write command cannot be issued until the read burst is complete (as with an SDR SDRAM). Going the other way, an explicit single-cycle delay is imposed on issuing a read command following a write burst, thereby incurring a data bus idle time equal to the selected CAS latency plus the single cycle write/read delay.

DDR SDRAM has taken the place of conventional SDRAM in many PC applications. Like SDR SDRAM, DDR devices are commonly available in densities ranging from 64 to 512 Mb in 4-, 8-, and 16-bit wide data buses. Thirty-two-bit devices are also available, although they are not the sweet spot for the industry as a whole.

8.3 SYNCHRONOUS SRAM

Like DRAM, high-performance SRAM transitioned to a synchronous interface to gain performance improvements. Several basic types of *synchronous SRAM* (SSRAM) devices appeared and became

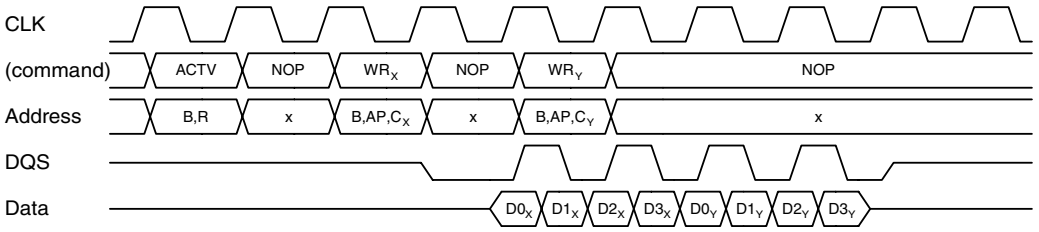

FIGURE 8.10 Consecutive DDR SDRAM writes (BL = 4).

standard offerings of numerous semiconductor vendors. SSRAMs are well suited for applications that require rapid access to random memory locations, as compared to SDRAMs that are well suited to long bursts from sequential memory locations. Many SSRAM devices can be sourced from multiple vendors with identical pinouts and functionality. An SRAM is made synchronous by registering its interface. Two basic types of SSRAMs are *flow-through* and *pipelined*. Flow-through devices register only the input signals and are therefore able to return read data immediately after the requested address is registered on a rising clock edge. Pipelined devices register both the input and output signals, incurring an added cycle of latency when returning read data. These differences are illustrated in Fig. 8.11.

As with SDRAM, there is a trade-off between access latency and clock speed. Pipelined devices can run at substantially faster clock frequencies than flow-through devices, because the SSRAM has a full clock cycle to decode the presented address once it is registered. In applications where clock speeds are under 100 MHz, flow-through SSRAMs may be preferable because of their lower latency. However, a flow-through device exhibits relatively high clock-to-data-valid timing, because the outputs are not registered. This large t_{CO} directly impacts the overall memory system design by placing tighter constraints on the interconnection delays and input register performance of the device that is reading from the SSRAM. For example, a Micron Technology MT55L512L18F-10 8-Mb flow-through SSRAM runs up to 100 MHz and exhibits a 7.5 ns access delay and a 3.0 ns data hold time after the next clock edge.* At a 10-ns clock period, there are 2.5 ns of setup budget to the next clock edge for an input register that is sampling the returned data. This 2.5-ns budget must account for interconnect delay, clock skew, and the setup time of the input flops. Alternatively, the 3 ns of hold time can help increase this timing budget, but special considerations must then be made to shift the data valid window of the input flops more in favor of hold time and less in favor of setup time. This is not always practical. In contrast, Micron's MT55L512L18P-10 8-Mb pipelined SSRAM is rated for the same 100-MHz clock but exhibits a 5.0-ns clock-to-valid delay and a 1.5-ns hold time.† For the added cycle of latency, the setup budget increases to a much more comfortable 5 ns with the same 10-ns clock period. Pipelining also allows the SSRAM to run at a much faster clock frequency: 166 MHz versus 100 MHz for the 8-Mb flow-through SSRAM. By using a pipelined SSRAM, you can choose between more favorable timing margins or increased memory bandwidth over flow-through technology.

An application in which SSRAM devices are used is a cache, which typically performs burst transactions. Caches burst data a line at a time to improve main memory bandwidth. Standard SSRAM devices support four-word bursts by means of a loadable two-bit internal counter to assist

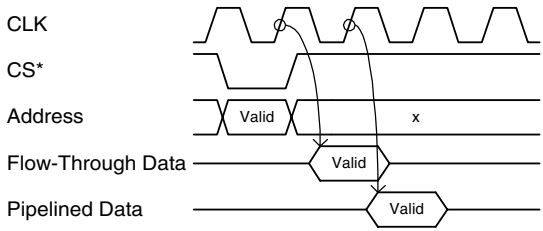

FIGURE 8.11 Flow-through vs. pipelined SSRAM reads.

* MT55L512L18F_2.p65-Rev. 6/01, Micron Technologies, 2001, p. 25.
† MT55L512L18P_2.p65-Rev. 6/01, Micron Technologies, 2001, p. 25.