Understanding Synchronous Dual-Ports

Introduction

This application note explains the basic operations and features of Cypress Synchronous Dual-Port memories. As the demand for faster devices and simpler interfaces increases, Synchronous Dual-Ports became the choice over Asynchronous solutions. For information about Asynchronous Dual-Ports and their operations, refer to application note “Understanding Asynchronous Dual-Port RAMs.”

Asynchronous Dual-Port RAMs (DPRAMs) respond asynchronously to address and control pin changes. This operation places restrictions on input pin timings which tends to limit achievable system performance. It also restricts the DPRAM maximum internal operating speed. Synchronous DPRAMs use external clocking to time read and write operations. The external clocking allows timing specifications that result in reduced DPRAM access and cycle times and hence higher system operating frequency and bandwidth. Synchronous DPRAMs are the choice for higher system operating frequency.

Cypress Synchronous dual-ported RAMs are available in different operating modes depending on the outputs being registered or not, as well as the mechanism used to load the addresses. There are three modes of operation: Pipelined, Flow-through, and Burst. Similar to Asynchronous Dual-Ports, Synchronous devices provide simultaneous access capability to any location in memory. Either port can write and read data into and out of any memory location.

Architecture

Figure 1 shows a block diagram of the general architecture of Sync DPRAMs. 18 data lines (I/O) and 16 Address lines are used in the case of a 64Kx18 (1 Meg) Sync DPRAM (part number CY7C09389). Each port has its own associated control, address and data lines. The Synchronous DPRAM operation is similar to the Asynchronous DPRAM except that operations on each port occur on the LOW to HIGH transitions of the clock signals (CLKL for left port and CLKR for right port). These clock signals are generated externally. The $\overline{FT}$/Pipe signal controls the operating mode for the associated port. For $\overline{FT}$/Pipe Low (connected to GND) the DPRAM is in Flow-through synchronous mode. For $\overline{FT}$/Pipe High (connected to V_{CC}) the DPRAM is in Pipelined synchronous mode.

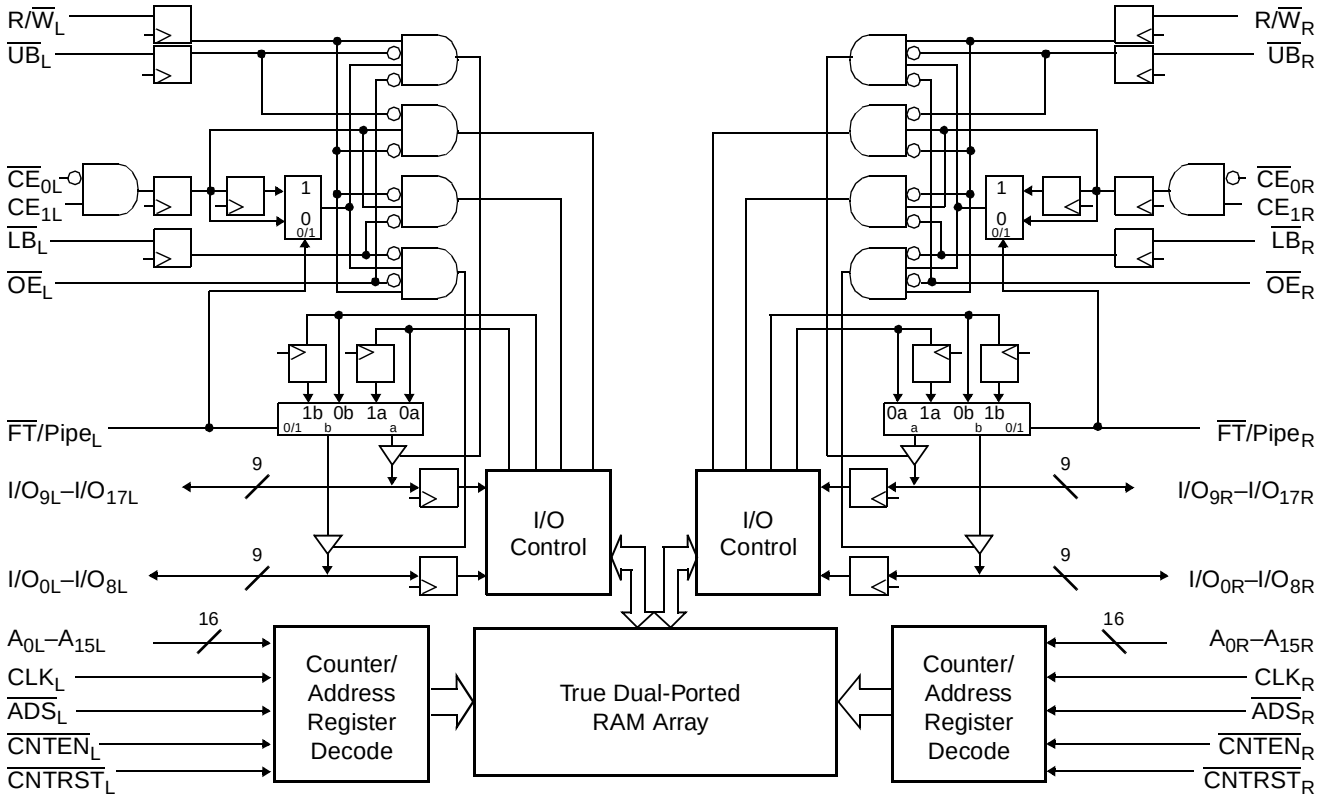

Figure 1. Logic Block Diagram of a 1-Meg Synchronous DPRAM

Dual chip enables are available for the ease of depth expansion of the Sync DPRAMs. The presence of $\overline{CE}_0$ and CE_1 eliminates the need for any external logic.

Each port has a burst counter used to stall the operation of the address input and utilize the internal address generated by the internal counter for fast applications. The internal address counter can increment through the entire address space of the DPRAM and wraps around to zero. The counter is loaded by asserting the $\overline{ADS}$ input LOW during the clock LOW to HIGH transition. With the count enable ($\overline{CNTEN}$) asserted for a port, the address counter will increment on each LOW to HIGH transition of that port's clock signal. The counter is reset to zero on a LOW to HIGH clock transition with count reset ($\overline{CNTRST}$) asserted.

Operating Mode

The $\overline{FT}$/Pipe pin determines the mode of operation for a port. When asserted LOW, the port is said to be in Flow-through sync mode. When asserted HIGH, the port is said to be in Pipelined sync mode. The main difference between these two operating modes is the relationship of data out to clock transition.

In Pipelined Sync mode, output data is stored in registers before being read out on the I/O lines, while in Flow-through sync mode the output is read immediately from the memory array onto the data I/O lines. The address is registered for both modes of operation.

Synchronous dual-ports also feature a burst mode operation. In burst mode, only the first address of a consecutive sequence of addresses must be loaded to the burst counter. Both Flow-through and Pipelined Cypress DPRAMs have the burst feature.

Read Operation

In order to read data from memory, several signals must be in the required state during the setup and hold timing window around the clock LOW to HIGH transition that initiates the read operation. The desired addresses must be applied to the address input pins (A), the chip must be selected by asserting $\overline{CE}_0$ LOW and CE_1 HIGH, and the R/$\overline{W}$ signal must be HIGH.

These signals must be present and valid a specific amount of time (setup time t_{SA}, t_{SC}, and t_{SW}) before clock transitions, and must remain valid for a specific amount time (hold time t_{HA}, t_{HC}, and t_{HW}) after clock transitions. At the rising edge of the clock (CLK) the address is registered and the read cycle begins. The output enable ($\overline{OE}$) is an asynchronous signal that can be asserted at any time. Data appears at the output lines if $\overline{OE}$ is LOW, otherwise the outputs are three-stated, and data will not appear at the outputs.

When the DPRAM is configured in Flow-though Sync mode, the read cycle begins after setting up the required signals and registering the address as described above. Data is valid t_{CD1} (clock to data valid for Flow-through) after the rising edge of the clock as shown in *Figure 2*. t_{CD1} is less than one clock cycle, therefore data appears at the outputs before the next rising edge of CLK.

In Pipelined Sync mode the output data is registered. Data from the read operation is clocked into the output register and transferred to the output lines after the second rising edge of the clock as shown in *Figure 2*. The output data is valid t_{CD2} (clock to data valid for pipelined) after the second clock transition. The use of registers at the output introduces a delay time—latency—before the first data appears at the output lines. Usually t_{CD1} (Flow-through) is higher than t_{CD2} (Pipelined) and therefore the total amount of time it takes to read a complete packet of information is less during Pipelined Sync mode.

Sync burst DPRAMs differ from Async DPRAMs in the way they use the address during the memory request sequence. Only the first address of a consecutive sequence of addresses must be loaded by the sync burst DPRAMs, which accomplish subsequent accesses by using a clock to advance an address counter inside the chip. Pipelined-burst DPRAMs take the idea one step further by using a register at the output stage of the RAM to hold the next piece of data in the burst sequence, which results in an extremely fast clock-to-output time (t_{CD2}). The only delay is the holding register's access time. *Figure 3* shows Burst Read operation using the Address Counter. The external address is constantly loaded into the counter on the rising CLK edge when $\overline{ADS}$ is LOW. When $\overline{ADS}$ is HIGH and $\overline{CNTEN}$ is LOW, the internal address counter is incremented. Data processed with the Pipelined-burst mode DPRAM is already waiting at the output stage once the pipeline has started.

Write Operation

The write operation is identical for Flow-through and Pipelined Sync DPRAMs since the registers at the output lines in Pipelined mode affect only the read operation. Write operation signal waveforms are shown in *Figure 4*. The Pipelined sync mode write operation latency is zero, whereas the latency is one cycle for read operation. This is because after the addresses, chip enables, and write enable are set and all signals have been latched into the DPRAM, the write operation continues independent of other system operations.

When writing simultaneously to the same location by both ports, the integrity of the data written is not guaranteed. The memory location might contain new data, old data, or some transitional value. The left port clock and the right port clock require a clock-to-clock set-up time of t_{CCS} (refer to datasheet for t_{CCS} values). If the time between the two clocks is less than t_{CCS}, external logic is required to implement arbitration. The same scenario is true when writing to a location from one port and reading out that location from the other port. A t_{CWDD} (Write port clock High to Read Data Delay) is required to assure that the read data is the newest and last piece of data written to that location.

The Address Counter (Burst Mode) can be used to perform a back-to-back Write or Read operation without any latency or idle cycle at the time when a new address is selected. The new address is loaded in the same way the first address is loaded. *Figure 4* shows the process of loading the external addresses A_{n+2} and A_{n+3} while $\overline{ADS}$ is LOW and $\overline{CNTEN}$ is HIGH. For each external address, the internal address is incremented by deasserting $\overline{ADS}$ (HIGH) and asserting $\overline{CNTEN}$ (LOW). Please refer to datasheets of Sync DPRAMs for switching waveforms of read and write operations with Address Counter Advance.

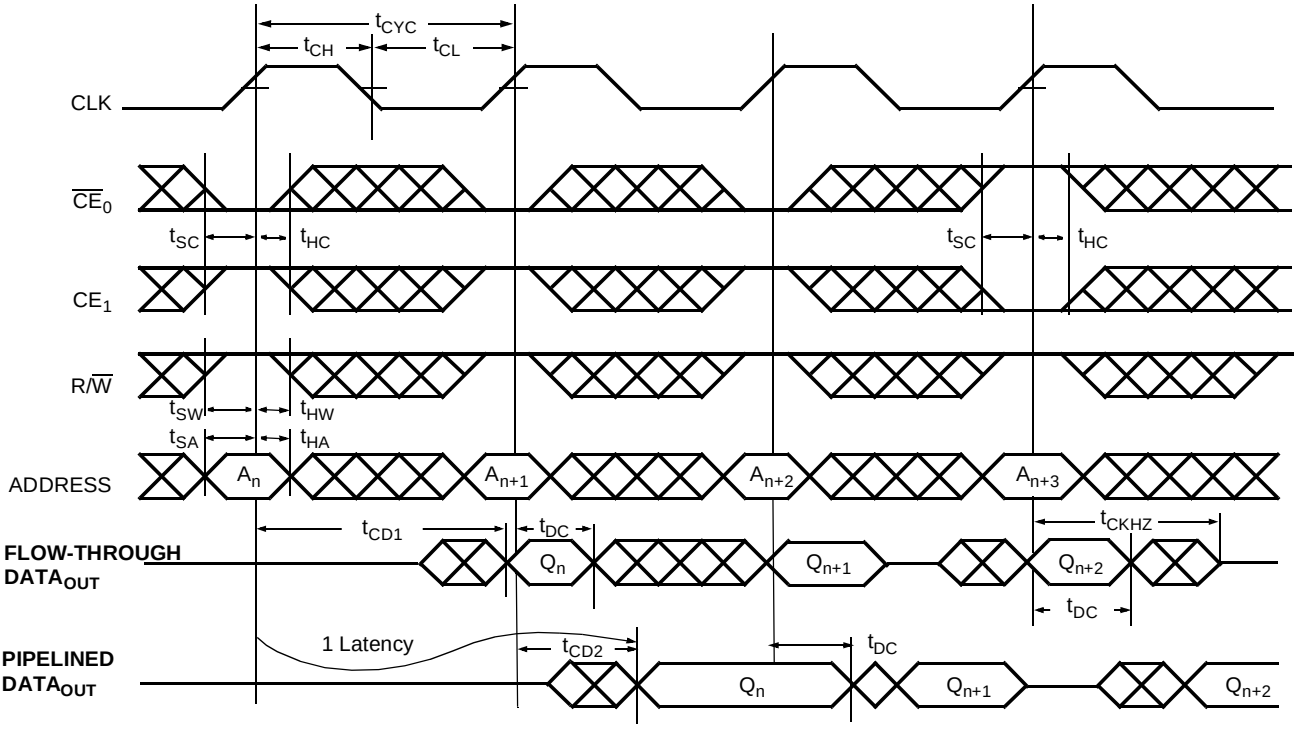

Figure 2. Read Operation of Flow-through and Pipelined

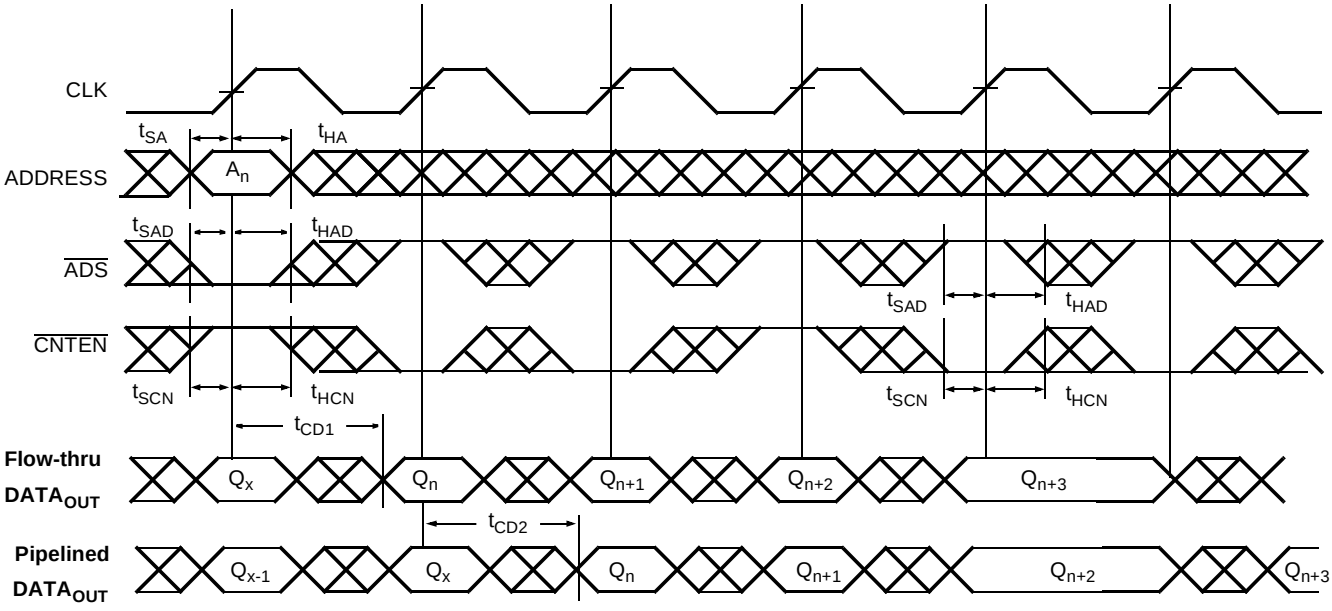

Figure 3. Flow-through Burst Vs. Pipelined Burst Read Operation

.

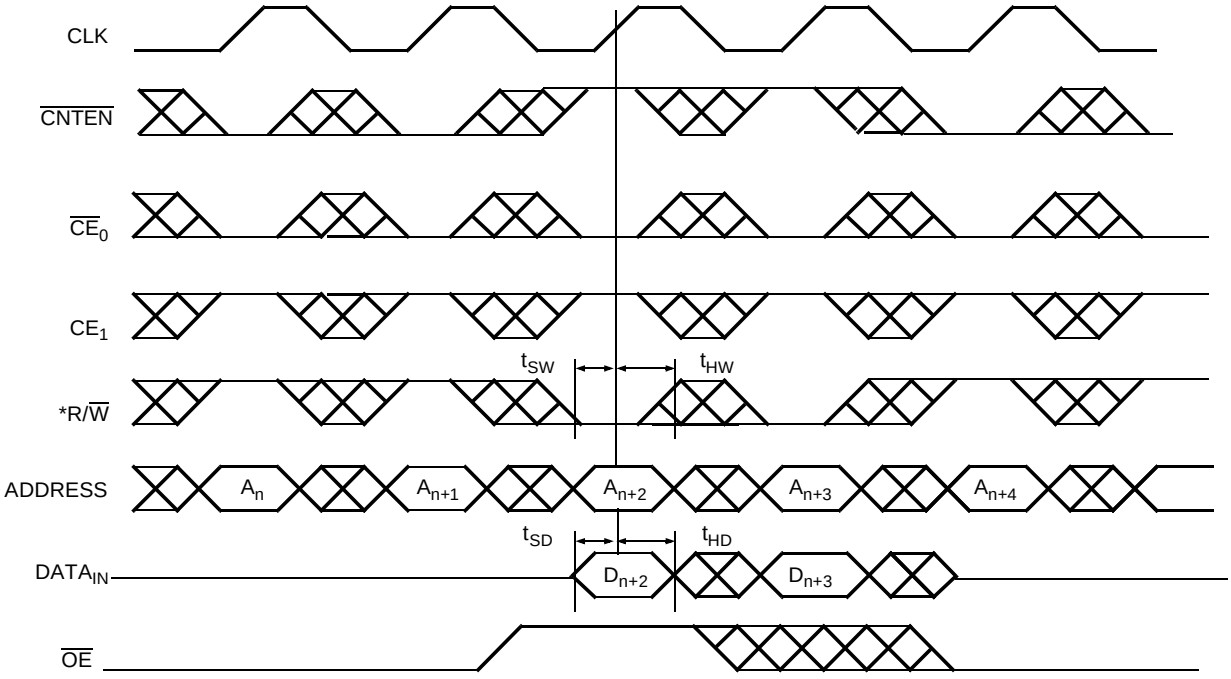

Figure 4. Write Operation for Synchronous DPRAM (all modes)

* the same waveform for $\overline{ADS}$

There are two latencies that could occur when using the Synchronous DPRAMs. The first is for a Pipelined read operation, as shown in *Figure 2*. The second type of latency happens when the R/$\overline{W}$ enable is used to do the I/O bus turnaround and it occurs on the first Write as shown in *Figure 5a*. This "dead" cycle is labeled as NO OPERATION, and occurs for both Flow-through and pipelined operating modes. During "NO OPERATION," data in memory at the selected address may be corrupted and should be re-written to ensure data integrity. *Figure 5b* shows that by controlling the $\overline{OE}$ and using it to asynchronously disable the I/O drivers for bus turnaround, the "dead" cycle no longer exists.

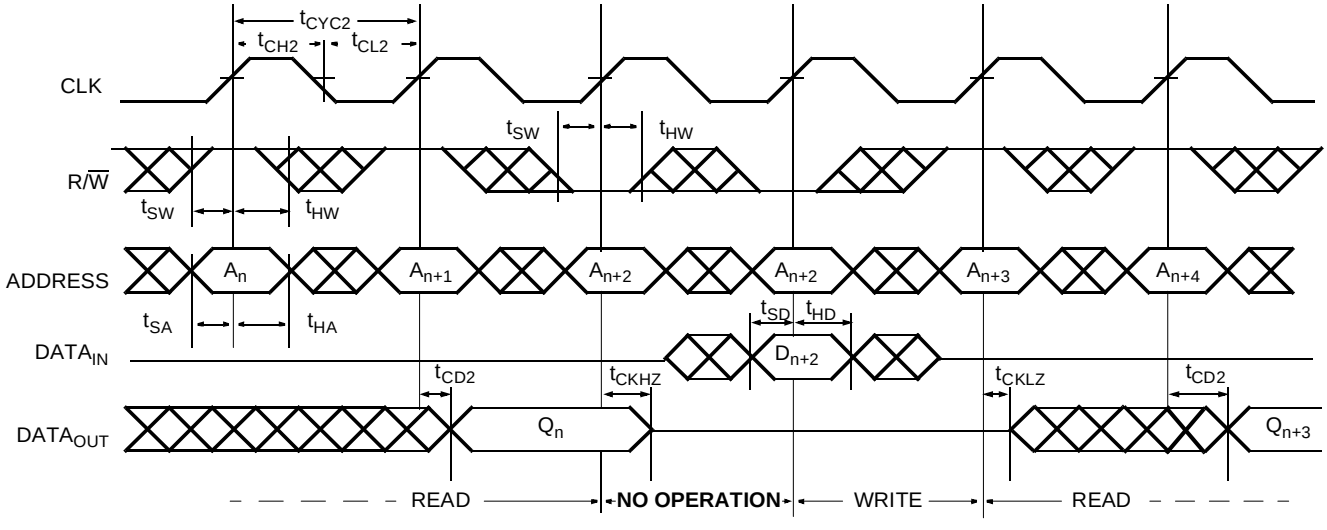

Figure 5a. Pipelined Read-to-Write-to-Read with $\overline{OE}$ = V_{IL}

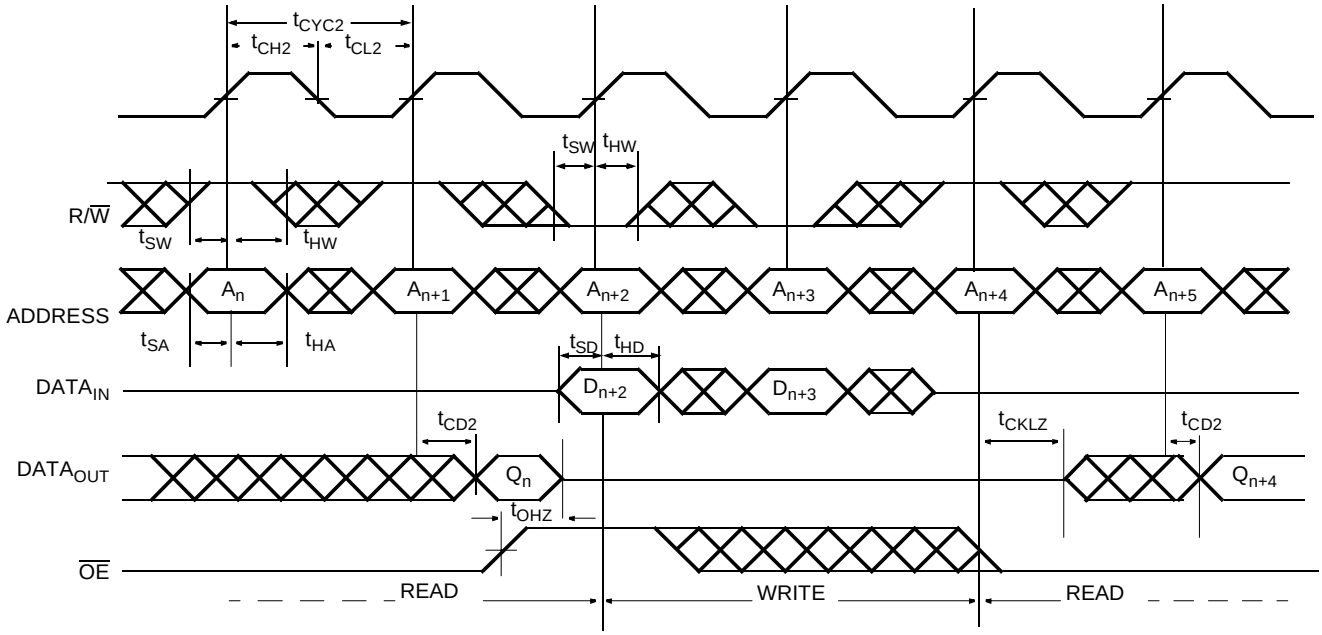

Figure 5b. Pipelined Read-to-Write-to-Read with $\overline{OE}$ Controlled

Width and Depth Expansion

Synchronous Dual-Ports come with two chip enables, $\overline{CE}_0$ and CE_1, for the purpose of simplifying the process of cascading these devices in both width and depth expansion without the need for external logic. *Figure 6* shows how to connect two 64Kx18 Synchronous DPRAMs in Depth Expansion to achieve a 128Kx18 organization. The same two devices can also be expanded in width. *Figure 7* shows the use of the x18 1-Meg DPRAM to achieve a x36 interface.

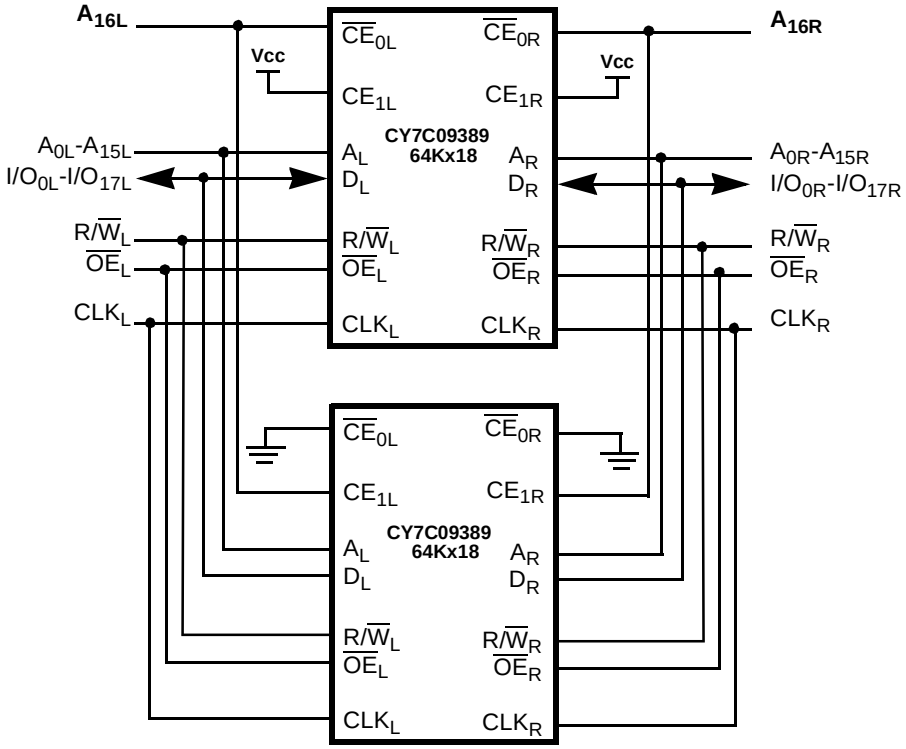

Figure 6. Depth Expansion of Synchronous Dual-Port RAMs

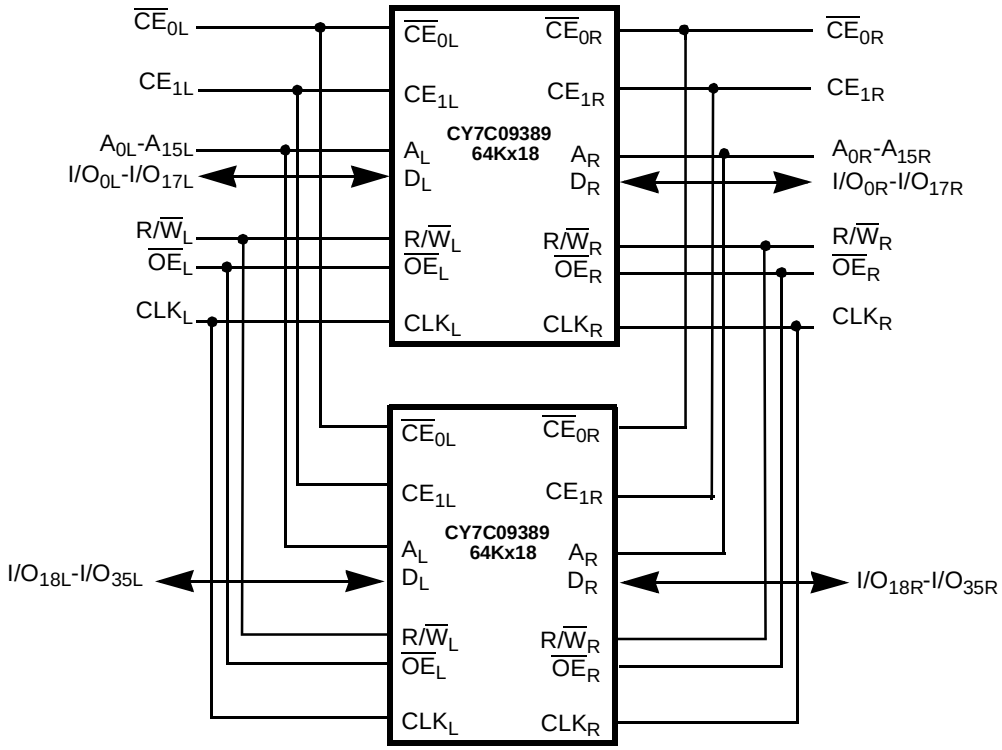

Figure 7. Width Expansion of Synchronous Dual-Port RAMs

Applications

Dual-Ports in general are widely used in communications, which includes the exchange of data between different processors. DPRAMs could be used in any system that requires data to be moved from one device to another. These systems include datacommunications networks (ethernet, LAN, switches, etc.) and multiprocessor systems where the system bus and the processor bus operate at different speeds. Synchronous DPRAMs are the ideal devices for high bandwidth applications where the outputs are registered (Pipelined sync mode) with fast clock to data time (t_{CD}) and short cycle time. Burst mode can be used for applications that require fast cycle time.

In a backplane switch (*Figure 8*), Synchronous DPRAMs can be used to provide high speed buffering for data packets. This minimizes backplane congestion by isolating critical controlling functions. It also provides independent random access for both the CPU and the Ethernet ASIC (controller & transceiver). Synchronous DPRAM also provides independent port clocks for different data rates. The DPRAMs are used to contain buffer indentifiers or pointers to shared buffer memory space allocated to packets from specific ports. The synchronous DPRAM interface gives more time for the controller to generate input signals.

The primary goal of any high-performance memory system is to eliminate processor wait time. To achieve this goal, the RAM system must be designed and optimized to provide for minimal latency when requests are made by the processor. Two ways of offering a minimal latency period are through increased memory size and increased memory speed. An even more important factor in the elimination of wait states is the reduction in access and transfer time to and from the RAM.

Figure 9 shows a high-speed memory system implemented using Cypress Sync DPRAMs as an in-line cache to interface a processor local bus to a system bus. This solution eliminates the problems that occur when a shared bus architecture is implemented. The use of the Synchronous DPRAM provides a direct interface between the microprocessor bus and the system bus through the presence of two independent ports. The left port address lines are directly connected to the processor bus while the right port address lines are not connected to the system bus to avoid overloading the bus. The presence of the dual addressing scheme enables simultaneous operations of read-write, write-write, or read-read in the same clock cycle. This means that an 83-MHz DPRAM will perform at a speed equivalent to a 166-MHz single address RAM. The use of the Synchronous DPRAM thus provides a cost effective method of designing an inline cache.

Conclusion

With synchronous interfaces becoming the desired and preferred choice, Cypress is becoming the market leader when it comes to synchronous Specialty Memory products (FIFOs and DPRAMs). Cypress DPRAMs are the deepest and fastest devices available world wide. They are available in 5V and 3.3V Async. and Sync. products. They come in x8,x16 as well as in x9 and x18 for parity use. For a complete listing of Cypress Dual-Ports (as deep as 1 Meg), please visit our web site at http://www.cypress.com.

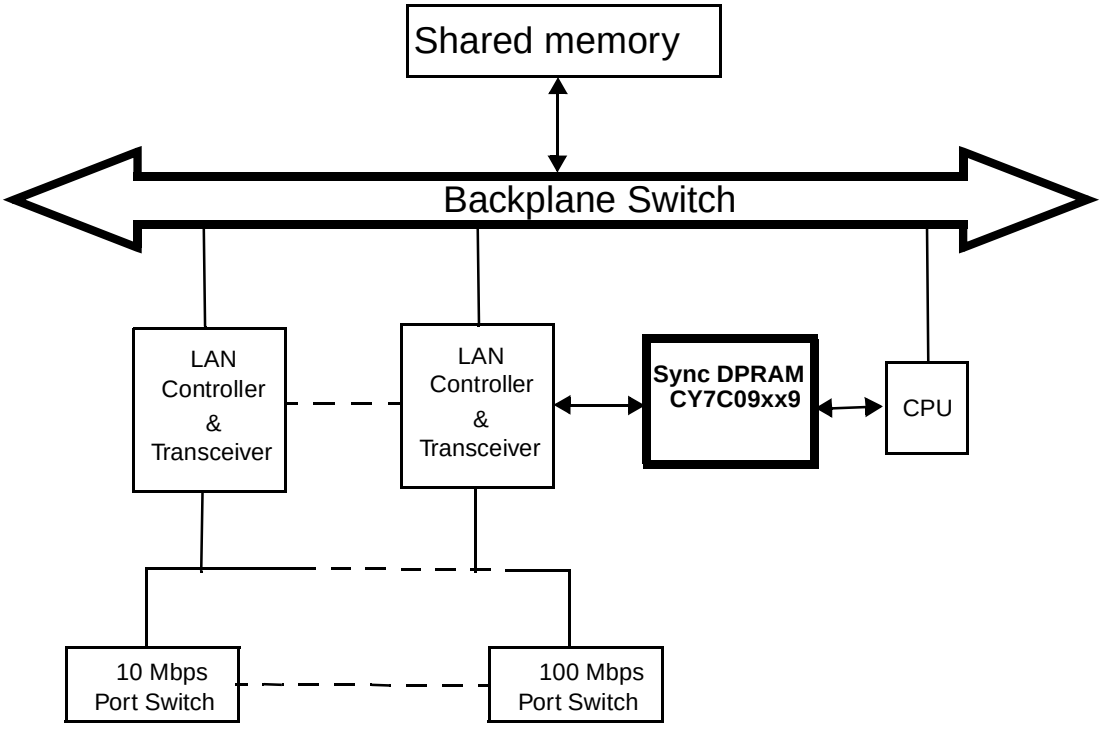

Figure 8. Backplane Switch using the Synchronous DPRAMs

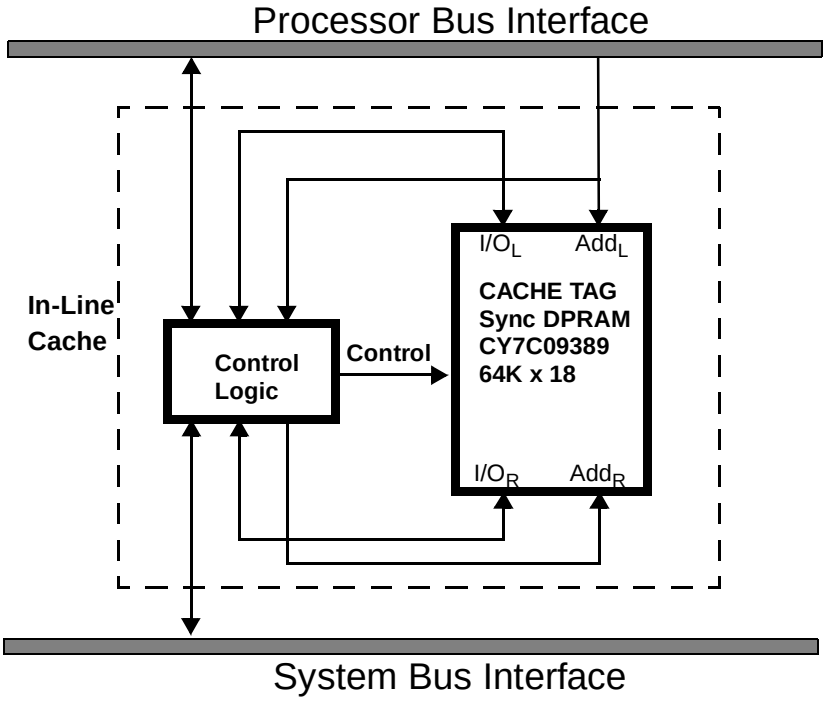

Figure 9. In-line Cache using Synchronous DPRAMs