

INTRODUCTION

In introducing our new Zero Bus Turn-around (ZBT) SRAMs, we are frequently asked how they compare with Late Write SRAMs from other manufacturers. This application brief highlights the differences.

Existing synchronous SRAMs are inefficient when read and write accesses are alternated. The ubiquitous pipelined burst SRAM (PBSRAM) has two idle clock cycles when a read follows a write. Best case bus efficiency is only 50% for alternating reads and writes. ZBT and Late Write both address this inefficiency but only ZBT totally eliminates idle cycles on the bus.

COMPARISON

ZBT comes in two versions: flow-through and pipelined. This terminology is used in exactly the same way as with the ubiquitous synchronous burst SRAM such as the IDT71V432, IDT71V537, etc. The pipelined ZBT SRAM supplies read data in two clocks while the flow-through device requires only a single clock. The difference in read cycles between the devices is the register in the data output path. This extra register means that the pipelined part has greater access latency, but it also permits it to run at higher clock rates than the flow-through device. See Figures 1 and 2.

access timing. With the pipelined device, read data may be sampled on the third clock, while the SRAM samples write data also on the third clock. With the flow-through devices, the data transfer takes place on the second clock rather than the third. The result is that both versions of ZBT can achieve 100% utilization of the data bus for any sequence of reads and writes.

Motorola markets two versions of their Late Write SRAMs: latched (MCM69Lxxx) and registered (MCM69Rxxx). Both devices use a differential clock. This is a good feature at extreme clock rates, but is an unusual and unnecessary requirement when operating at more conventional clock rates.

The first Late Write SRAMs are of the registered type. Their functionality is shown in Figure 3. When reading, they function the same as the pipelined ZBT SRAMs: data is available after the second clock for sampling on the third clock. But unlike the pipelined ZBT, write data is supplied on the second clock instead of the third. Therefore, a read and write (addresses A4 and A5) cannot be initiated on successive clock cycles because both accesses would be attempting to transfer data on the same clock. Conversely, if a read follows a write (addresses A5 and A6), there will be one idle cycle on the data bus between the transfer of the write data (D5) and the read data (D6). For alternating reads and writes, the bus utilization is only 67%.

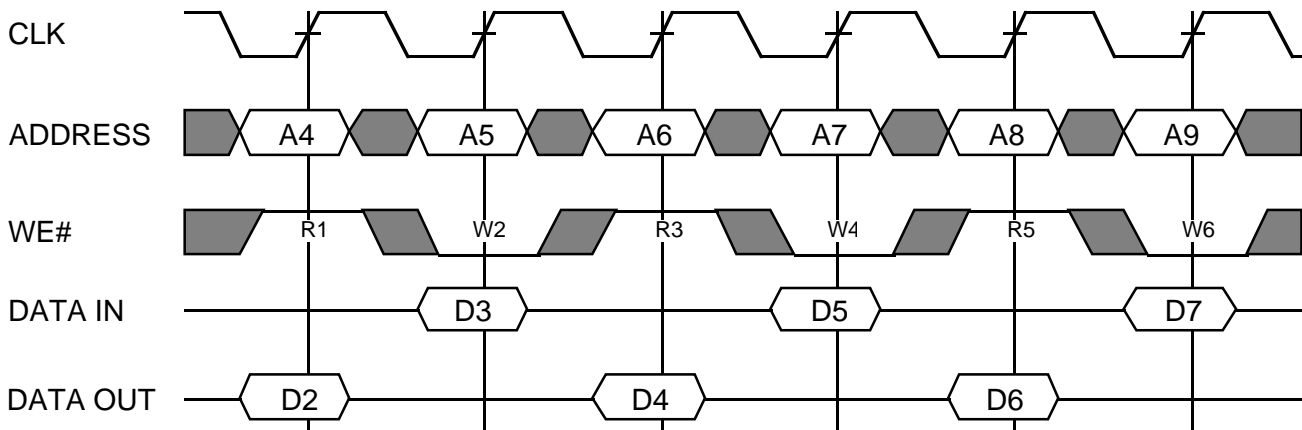


Figure1: ZBT Operation - Pipelined Version

Note that in our discussions of timing, we refer to the rising clock edge on which the address and read/write control signal are supplied to the SRAM as the "first clock." This is the reference point for all other clock cycles.

Both versions of ZBT permit switching between write and read cycles without any idle data bus cycles. This is achieved by making the write access timing symmetrical with the read

At first glance, the latched version of Late Write does better. It employs the same write timing as the registered Late Write, but with a lower latency read. This is similar to the flow-through version of ZBT, but there are some important differences as shown in Figure 4. The Late Write part uses both edges of the clock, requiring an extra half clock to supply valid read data to the bus (timing tCD). This means that the clock duty cycle becomes a factor in the interface timing, and that the data is




Figure 2: ZBT Operation - Flow-through Version

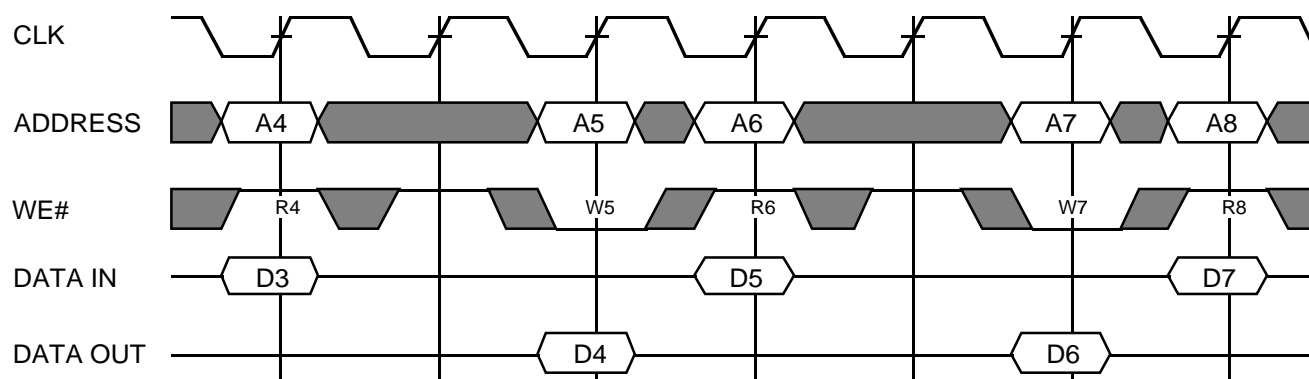


Figure 3: Late Write Operation - Registered Version

available relatively late in the clock cycle. This leaves very little time to meet the setup time into the receiving device and makes it nearly impossible to operate the latched version Late Write SRAMs at their rated clock speed.

Additionally, the latched version Late Write SRAMs use the falling clock edge to enable and disable their data bus output buffers (timing parameter tHZ). This makes it almost impossible to perform a write immediately after a read, since the SRAM outputs do not tri-state quickly enough to permit the writing device to supply data to the SRAM without contention on the bus.

In contrast, all ZBT timing references the rising edge of the clock, with the exception of the asynchronous output enable signal. This includes enabling and disabling of the data bus output buffers. This makes the timing very straightforward.

SUMMARY

As the name implies, ZBT SRAMs from IDT permit users to combine reads and writes in any sequence without any idle cycles on the data bus. This 100% data bus utilization is not just theoretical; it's achievable in practice. This applies both to the low latency flow-through version and the high speed pipelined version.

While similar in some ways to ZBT SRAMs, both versions of the Late Write SRAMs fall short. The registered version has an idle cycle whenever a read follows a write. The latched version can eliminate idle cycles between alternating reads and writes, but only theoretically. Its poor timing essentially guarantees data bus contention if a user attempts to perform a write immediately following a read. And while this problem is reduced a bit by operating the part more slowly, the user must still deal with timing referenced to the falling edge of the clock.

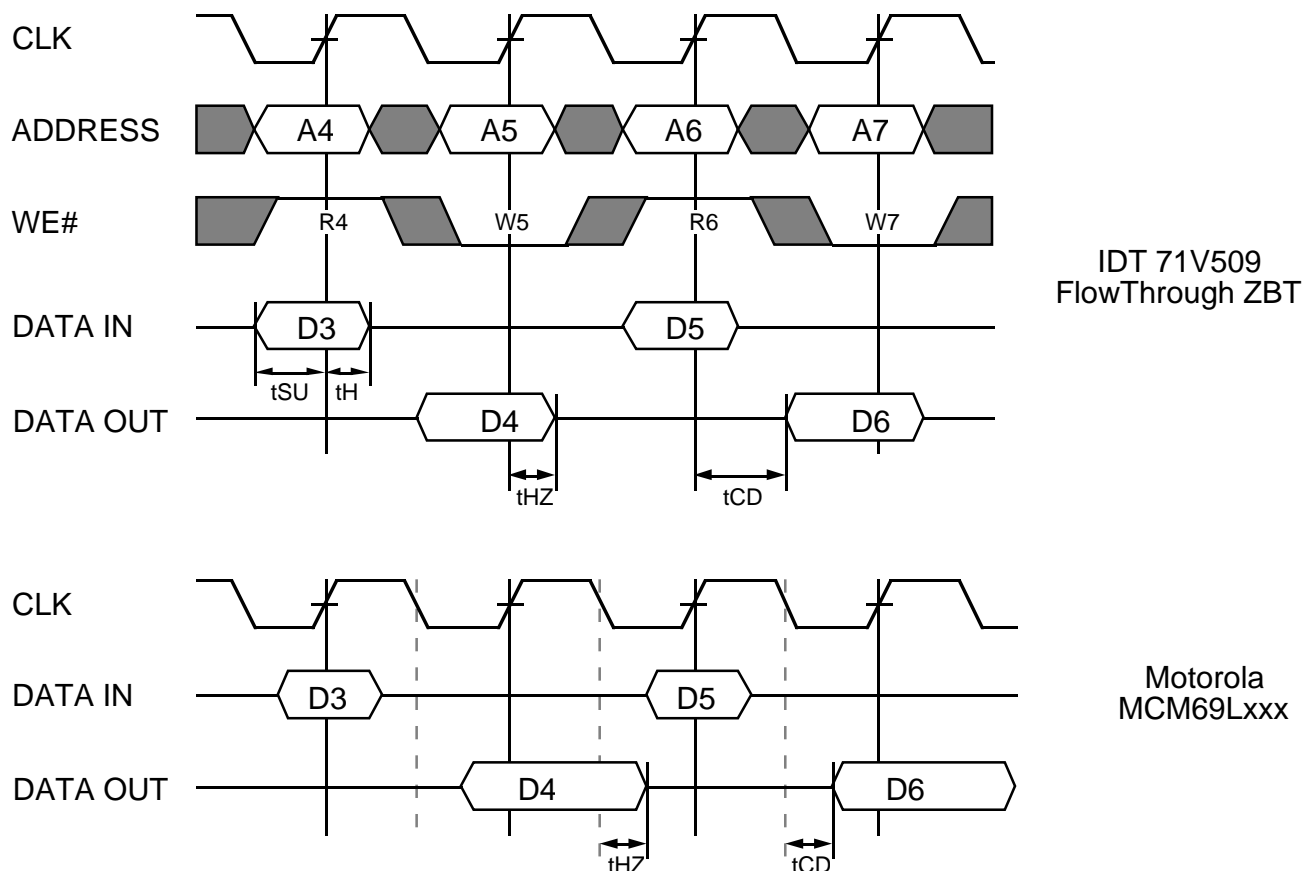


Figure 4: ZBT Flow-through Version Compared with Late Write Latched Version

Integrated Device Technology, Inc. reserves the right to make changes to the specifications in this data sheet in order to improve design or performance and to supply the best possible product.

Integrated Device Technology, Inc.

2975 Stender Way, Santa Clara, CA 95054-3090

Telephone: (800) 544-SRAM

Fax: (408) 754-4547