

ZBT™ Primer

Prepared by: Don Taylor, Benjie Zeller, and Andrea Crocker

The emergence of higher bandwidth networking systems and infrastructures has driven SRAM vendors to search for faster data throughput solutions. One such SRAM architecture standardized by Freescale, IDT, and Micron is called Zero Bus Turnaround™ (ZBT). The ZBT RAM eliminates bus latency by providing a more efficient use of the system bus at a significant cost reduction to the customer. Motorola's new 4M-bit ZBT SRAM is the ideal solution for high-end networking and data communication systems that require back-to-back reads and writes.

The ZBT RAM is a synchronous fast static RAM that eliminates bus latency. It eliminates idle or "unused" bus cycles during system bus read and write cycle transitions to provide simpler control logic, 100% bus efficiency for bus bandwidth, and data throughput without design restrictions. The ZBT RAM eliminates the two idle cycles required by other synchronous SRAMs to transition the bus from a write to a read, resulting in as much as 100% greater bus bandwidth. Consequently, designers need to be aware of the potential hazards that accompany 100% bus utilization, such as bus contention.

This application note serves as a primer to introduce the designer to ZBT potential applications, architecture, and issues associated with frequency, bandwidth, bus contention, and temperature effect. Comparisons between ZBT and other memory devices aid the designer in deciding whether the ZBT RAM is an appropriate design choice.

SUGGESTED ZBT RAM APPLICATIONS

- Switch/Hub Shared Fabric and Router Tables
- Network Interface Cards (NIC)
- LAN and WAN Switches
- Networking Computer
- Gigabit Switching

ZBT RAM FEATURES

- Single 3.3 V ± 5% Power Supply
- 4.0 ns Access Time (Pipelined)
- 10 ns Access Time (Flow-Through)
- Selectable Burst Order (Linear or Interleaved)
- Internally Self-Timed Write
- Two-Cycle Deselect

- Byte Write Control
- ADV Controlled Burst
- 100-Pin TQFP Package

PRODUCT FAMILY CONFIGURATIONS

Configuration	Part No.	Mode
128K x 36	MCM63Z736	Pipelined
256K x 18	MCM63Z818	Pipelined
128K x 36	MCM63Z737	Flow-Through
256K x 18	MCM63Z819	Flow-Through

ZBT RAM versus BurstRAM

The ZBT RAM provides higher bandwidth than standard BurstRAMs in non-PC applications. The ZBT RAM pinout is similar to the pinout of the BurstRAM; therefore, some minimal board redesign may optimize bandwidth. For detailed information, see Freescale Application Note AN1729/D, *BurstRAM to ZBT RAM*.

ZBT RAM ARCHITECTURE

Figure 1 illustrates a simplified external ZBT operation. The address and control logic load the ZBT input registers on the rising edge clock transitions. Address A0 (with control logic) is asserted for a minimum setup time before the falling clock edge of the first clock cycle.

Figure 2 illustrates the defined data timing parameters. The control logic determines whether a read or write command is asserted. The associated data I/O cycle is asserted on DQx two clock cycles later. During the write cycle, D1 must be setup at the inputs before the end of the first clock cycle time t_{DVKH} . Next, D1 must be held for hold time t_{KHDX} after the data is latched on the rising edge of the second cycle. For a read cycle, data out is valid t_{KHQV} after the rising edge of the second clock cycle until t_{KHQX} after the rising edge of the third clock cycle. Input data needs to be valid for a hold time following the low to high clock transition at the end of the third cycle. New address and control signals may be presented to the ZBT RAM during each clock cycle. The relationship between the address and control nth clock cycle and its associated n+2 data clock cycle remains continuous.

ZBT and Zero Bus Turnaround are trademarks of Integrated Device Technology, Inc., and the architecture is supported by Micron Technology, Inc. and Freescale Semiconductor, Inc.

© Freescale Semiconductor, Inc., 2004. All rights reserved.

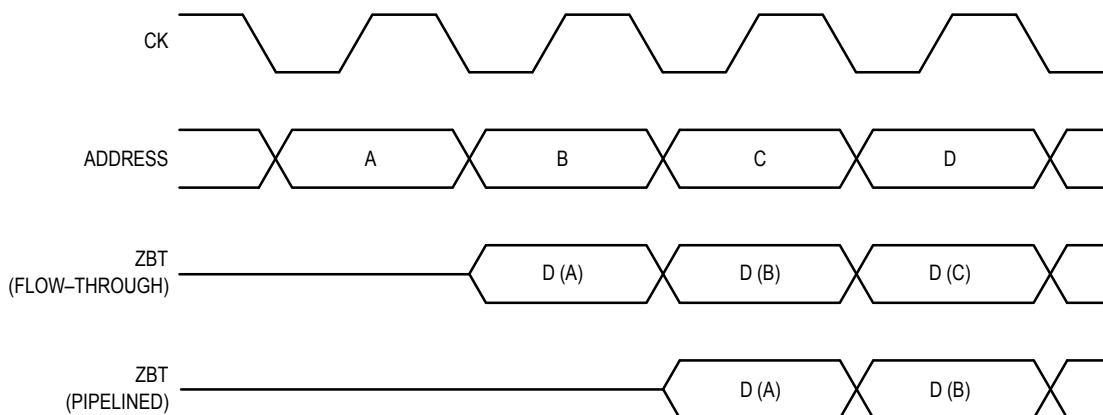


Figure 1. ZBT Operation

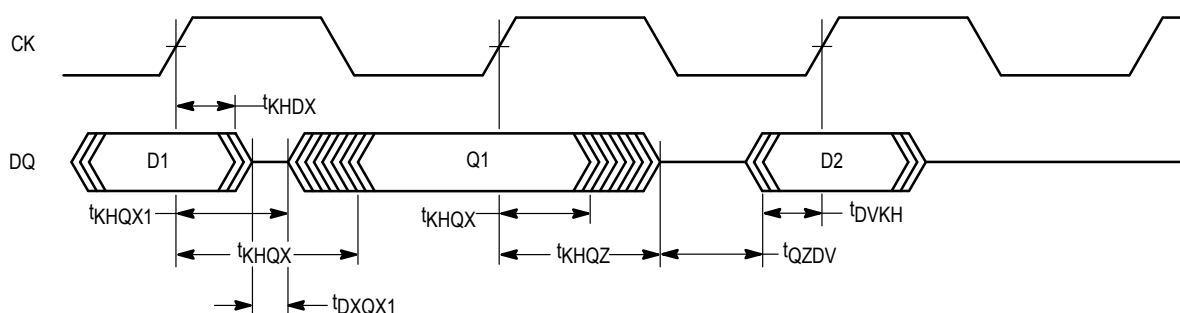


Figure 2. Timing Parameters

Figure 3 shows a simplified ZBT block diagram. The addresses, control signals, and data inputs load positive edge-triggered registers. The output data flows from a positive edge-triggered register, while inputs only need to meet setup and hold times. The ZBT output is controlled by a gate that enables or three-states the output buffer asynchronously with the external \overline{OE} pin. The \overline{OE} pin may either be held low or toggled as needed. Output is internally controlled by logic to turn the buffer on or off with the read/write control signal sampled two clock cycles earlier. The mux ensures coherent operation when a read follows a write to the same address that occurs on the next clock cycle. The control logic comparator circuitry monitors the contents of the address bus to mux the data on the I/O pins into the output register, if the current and previous addresses are identical.

Activating the clock gating control input pin \overline{CKE} instructs the ZBT RAM to ignore the rising edge of the clock cycle, thus leaving the internal registers unchanged. The output remains on when \overline{CKE} asserts during a read cycle. Self-timed write/write cycles complete internally. \overline{SBx} only needs to meet the required clock input setup and hold times, eliminating the need to generate off-chip write pulses. The output

enable pin (\overline{G}) can operate asynchronously or hold low to internally and synchronously control the output buffer.

BANDWIDTH UTILIZATION

Bandwidth is a major consideration in the selection of SRAM. Throughput varies between different SRAMs at the same frequency due to differing wait states during read-to-write transitions. Tables 1 and 2 illustrate bandwidth relationships between the different SRAM families at typical frequencies for both pipelined and flow-through models. The bandwidth of the slowest BurstRAM far exceeds that of a high performance asynchronous SRAM.

Bandwidth may be expressed in several ways. The maximum bandwidth is the best case and most frequently used SRAM metric. Most SRAMs attain maximum bandwidth when performing consecutive reads or writes; however, ZBT RAMs attain maximum bandwidth when performing any combination of reads and writes. The BurstRAM does not suffer significant bandwidth loss in bursting PC L2 cache applications, although it loses 50% of its maximum bandwidth in a read/write mode. The worst-case bandwidth usually occurs when there are continuous read-to-write-to-read transitions.

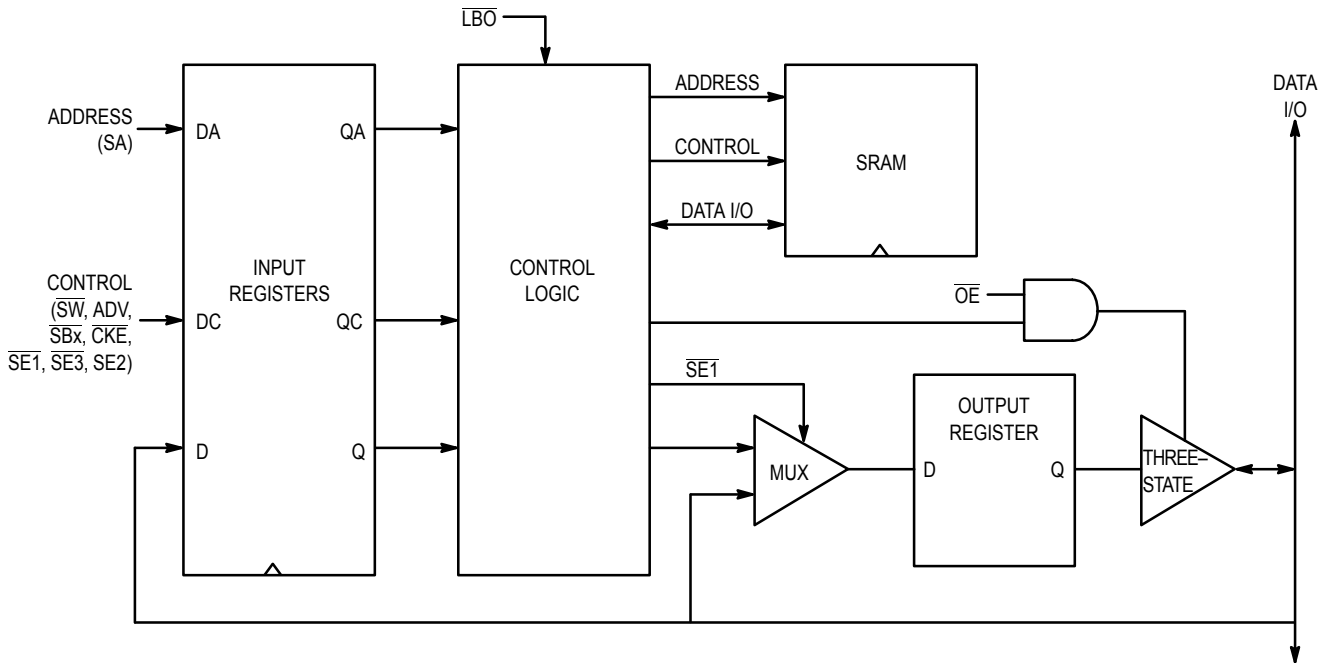


Figure 3. Simplified Block Diagram

Table 1. Bandwidth for Flow-Through Devices

Device	(MHz)	(MB/sec)	Bandwidth	SW Burst of Four Bandwidth
BurstRAM	83	332	221	295
ZBT	100	400	400	400
Late Write	166	664	664	664

Table 2. Bandwidth for Pipelined Devices

Device	(MHz)	(MB/sec)	Bandwidth	SW Burst of Four Bandwidth
Asynchronous	100 (- 10)	75	50	50
BurstRAM	133	533	267	462
ZBT	133	533	533	533
Late Write	200	800	533	711

PIPELINED BurstRAM, LATE WRITE, AND ZBT RAM TIMING COMPARISON

Figure 4 compares timing between the following synchronous SRAMs: MCM69P737 (Pipelined BurstRAM), MCM69R737 (4M Late Write SRAM), and MCM63Z736 (Pipelined ZBT RAM). The timing sequence shows back-to-back read and write operations in a read-read-write-write-read sequence.

The BurstRAM performs back-to-back reads or writes. It has two idle clock cycles on the data bus and one deselect cycle on the address bus during the write-to-read transition. The BurstRAM performs a read-write transition in one clock cycle. The read-read-write-write-read sequence takes ten cycles to complete.

The Late Write RAM performs back-to-back reads or writes, but has one idle cycle during the transition from write to read. There are no idle cycles on the data bus or any deselect cycles on the address bus during the turnaround from write to read. The read-read-write-write-read sequence takes nine cycles to complete.

The ZBT RAM performs back-to-back read or write operations on each clock cycle without any clock cycle penalties during read-to-write or write-to-read transitions. There are no idle cycles on the data bus or any deselect cycles on the address bus during the write-to-read turnaround. The read-read-write-write-read sequence takes eight cycles to complete.

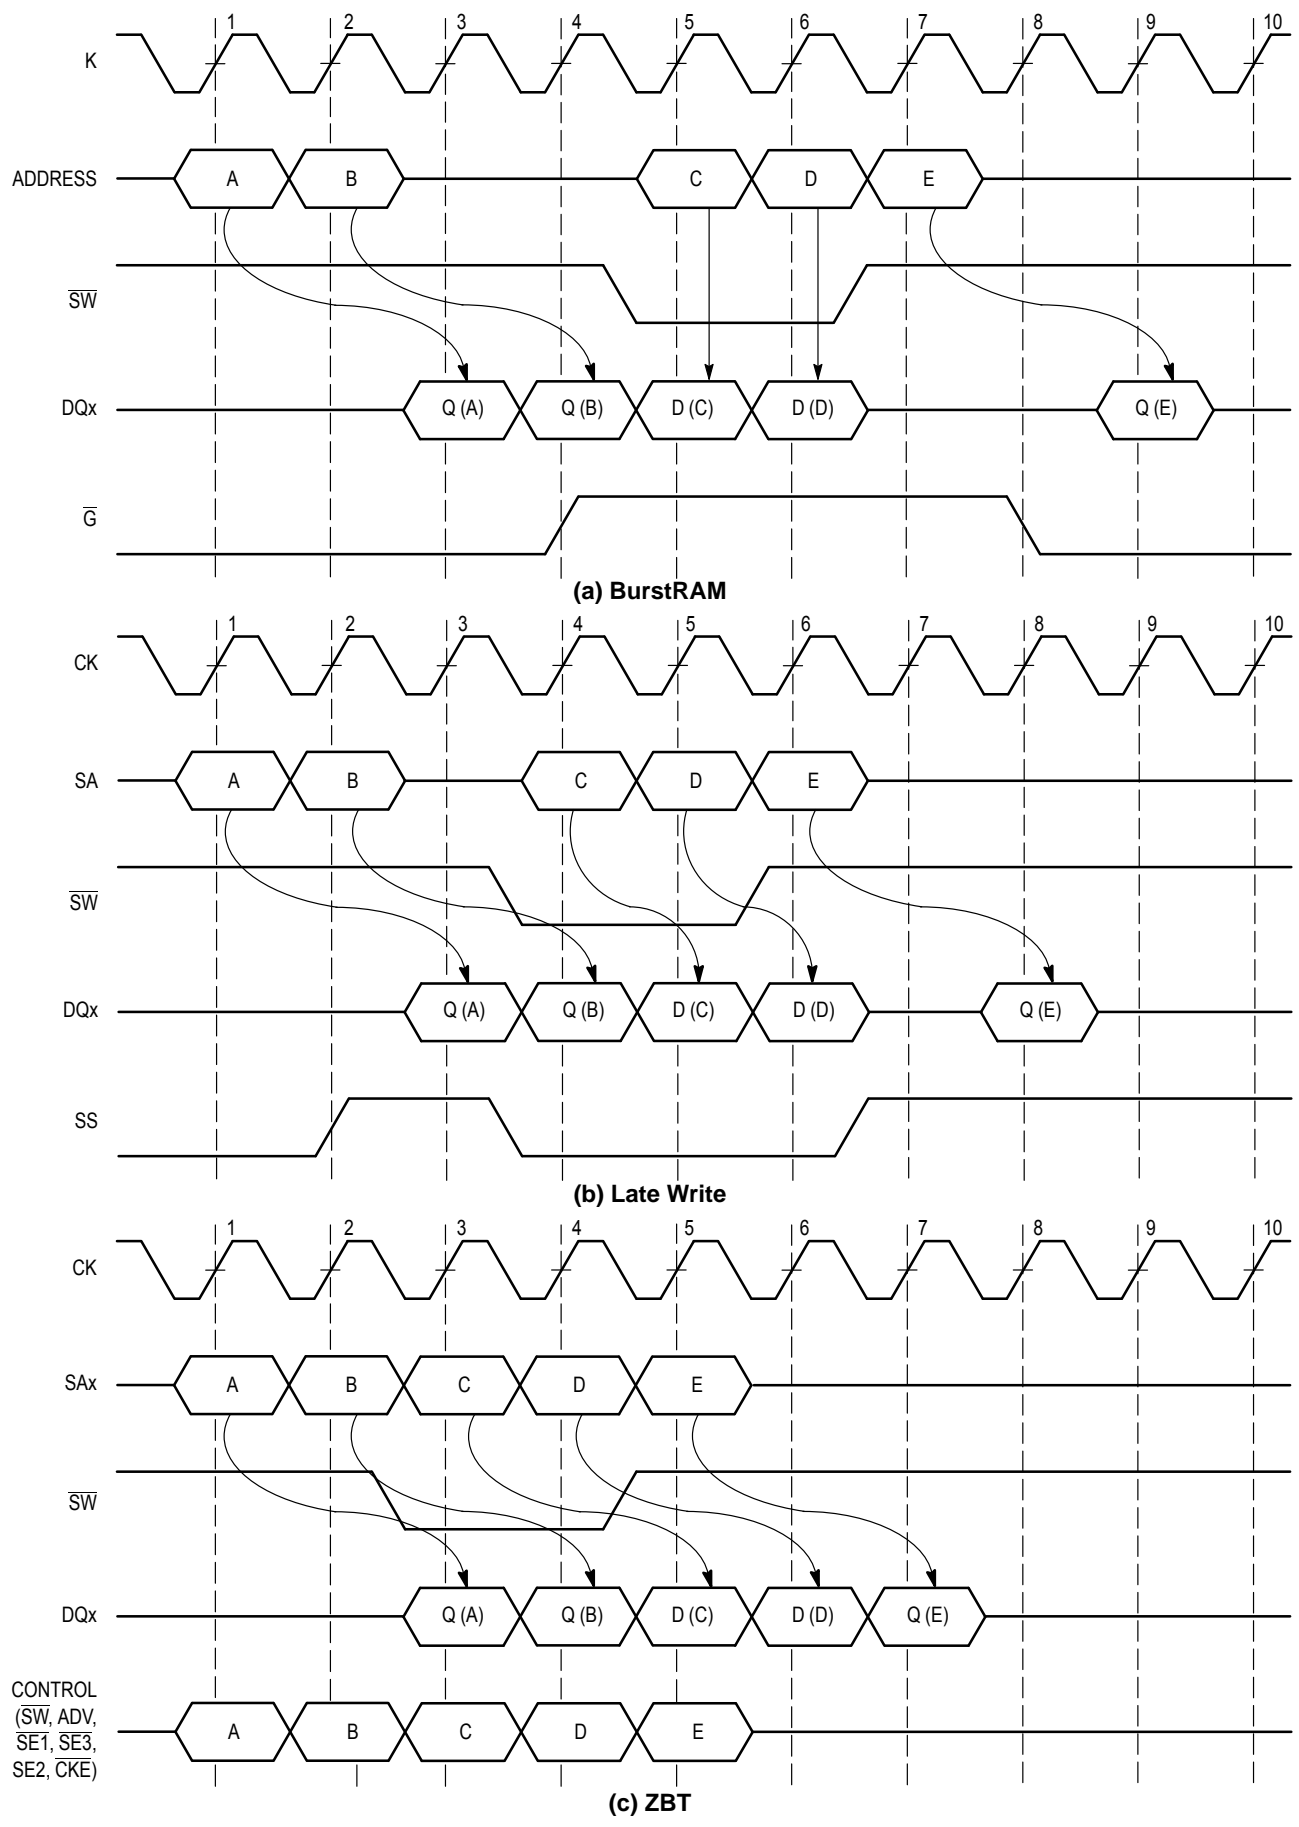


Figure 4. Comparison of BurstRAM, Late Write, and ZBT Timing

Consider the case of continuous read–write–read–write cycles. The effective bus frequencies for BurstRAM, Late Write, and ZBT RAMs are 50%, 60%, and 100%; respectively, of the clock frequency after a two–clock latency. The Freescale ZBT RAM provides up to 100% performance improvement over the standard pipelined BurstRAMs and up to 50% performance improvement over Late Write RAMs in high read and write data bus switching applications. The control (Figure 4c) represents the state of the read/write signal (\overline{SW}), synchronous load/advance signal (ADV), the synchronous chip enable signals (\overline{SE} , $\overline{SE3}$, SE2), and the clock enable signal (CKE). The control signals correspond simultaneously with the addresses they affect and the associated output always occurs at clock cycle $n+2$. It is then entered into the SRAM on the rising clock edge of cycle $n+3$. This allows the maximum system bus time to drive the address and control lines to the SRAM.

For example, the address and control signal A of clock cycle 1 must be setup by the end of cycle 1, so the information can be latched onto the SRAM at the rising edge of the cycle 2 transition. The DQx cycle Q(A) occurs during cycle 3. This pattern continues for every D or Q data cycle, such that the address and control cycles precede it by two clock cycles.

CONTROL LOGIC

The ZBT RAM is an excellent replacement for BurstRAMs in non-PC applications that use simplified control logic and fewer control signals. Specifically, two ZBT control signals can replace five BurstRAM control signals. The standard BurstRAM uses \overline{ADSP} , \overline{ADSC} , and ADV to load commands and perform burst operations. The ZBT RAM uses the advance/load pin ($\overline{ADV/LD}$) to perform the same commands as

the BurstRAM's ADV, \overline{ADSP} , and \overline{ADSC} . The SRAM receives a read or write command when the ADV is low due to the state of pin \overline{SW} . A deselect command is asserted when one of the three chip enables is inactive. The optional burst command is performed when ADV is high. The ZBT RAM uses a read/write pin (\overline{SW}) with byte write enables (\overline{SBx}) for global writes. This allows byte writes and global writes to be performed minus one control signal. The ZBT global writes are performed by asserting all four byte writes.

BUS CONTENTION

All devices mentioned above will control the data bus from time to time. Bus contention is an important element in the overall design of a major system. The SRAM controller must separate the write and read with a short high-Z state before the SRAM output driver can be activated. Bus contention occurs when two or more devices drive the bus to opposite logic levels. The amount of bus contention for system design is determined by the amount of time two contending outputs are driving the bus until one enters high impedance.

Figure 5 illustrates device X's attempt to drive the bus to logic 1, while device Y attempts to drive the bus to logic 0, resulting in a large current path between the contending devices from V_{CC} through the "on" transistors of both devices to ground. This current path is prevented, if either device drives the bus to the same logic value or one of the devices is driven to high-Z. Bus contention problems are difficult to identify and eliminate, since they usually last less than 5 ns. Bus contention is observed as noise on power supply and data lines connecting the contending devices. The system design must be robust enough to eliminate the possibility of bus contention to reduce potential long-term system reliability failures.

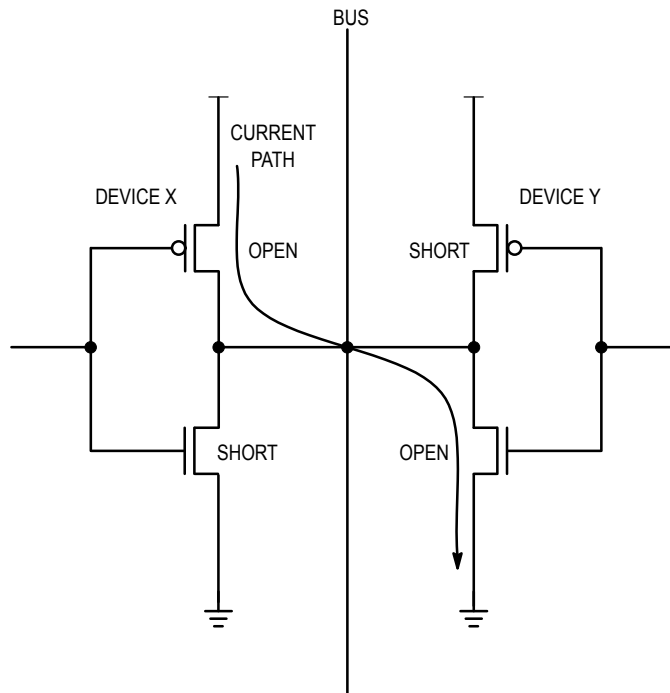


Figure 5. Bus Contention Current Path

The output of a synchronous SRAM turns on and off relative to the rising clock edge t_{KHQX1} and the falling clock edge of the data-in hold time for t_{KHDX} of a contending device. The timing diagram shown in Figure 6 illustrates both timing parameters. The goal is to design the system so $t_{DXQX1} > 0$ and $t_{QZDV} > 0$, where

$$\text{Bus Margin (DQ)} = t_{DXQX1} = t_{KHQX1} - t_{KHDX}$$

and

$$\text{Bus Margin (QD)} = t_{QZDV} = t_{KHKH} - t_{KHQZ} - t_{DVKH}$$

The bus margin is the amount of time the device currently writing to the bus has to transition to high-Z. For example, during a write-to-read transition, if the clock operates at 133 MHz and $t_{KHQX1} = 1.5$ ns and $t_{KHDX} = 0.5$ ns, then

$$\text{Bus Margin (D} \geq \text{Q)} = 1.5 \text{ ns} - 0.5 \text{ ns} = 1.0 \text{ ns}$$

1.0 ns is the amount of time the SRAM controller has to drive the bus to high-Z before the ZBT SRAM begins its transition to driving the bus.

Similarly, for a read-to-write transition, if $t_{KHQZ} = 3.5$ ns and $t_{DVKH} = 1.7$ ns, then

$$\text{Bus Margin (Q} \geq \text{D)} = 7.5 \text{ ns} - 3.5 \text{ ns} - 1.7 \text{ ns} = 2.3 \text{ ns}$$

2.3 ns is the amount of time the ZBT RAM has to drive the bus to high-Z before the SRAM controller begins its transition to driving the bus. It is highly recommended to a 0.5 ns timing margin of error to compensate for clock skew, supply voltage variation, and temperature. The length of time the SRAM controller drives the bus is largely a function of the components' process technology. Note: It may appear that the bus contention can not be avoided if a system designer uses the data sheet parameters for t_{KHDX} and t_{KHQX1} , because data sheets specify the worst case conditions for t_{KHDX} and t_{KHQX1} . The worst case conditions are t_{KHDX} (maximum) and t_{KHQX1} (minimum), which will not occur simultaneously under one set of conditions.

The following example illustrates the effect of bus contention on power dissipation junction temperature (T_J). The junction to ambient temperature gradient is

$$T_{JA} = T_J - T_A = 110^\circ\text{C} - 70^\circ\text{C} = 40^\circ\text{C}$$

and the junction to ambient thermal resistance is $R_{\theta JA} = 40^\circ\text{C/W}$ for a single-layer board. Then the junction-to-ambient power dissipated is

$$P_{JA} = \frac{T_{JA}}{R_{\theta JA}} = \frac{40^\circ\text{C}}{28^\circ\text{C/W}} = 1.43 \text{ W}$$

The operating power supplied under maximum voltage is

$$P_{S,max} = V_{DD} \times I_{DDA} = 3.6 \text{ V} \times 350 \text{ mA} = 1.26 \text{ W}$$

The maximum power to be added by bus contention is then

$$P_{ALLOWED} = P_{JA} - P_{S,max} = 1.43 \text{ W} - 1.26 \text{ W} = 0.17 \text{ W}$$

The ZBT RAM operates at 133 MHz for a 7.5 ns clock cycle. Bus contention occurs only during a write-to-read cycle transition every two cycles. The worst-case bus contention occurs during 6.7% of the total cycle, or 1 ns out of every 15 ns (two duty cycles). The maximum power consumed by the ZBT RAM without bus contention is

$$P_{S,max} = V_{DD} \times I_{DDA} = 3.6 \text{ V} \times 350 \text{ mA} = 1.26 \text{ W}$$

Since the data bits are either 1 or 0, half of the DQ pins are in contention with the SRAM when bus contention occurs. Therefore, 18 pins of a x36 SRAM are in contention, on average. Assume the SRAM driver and controller resistance are both 32 ohms during contention, and a 64 Ω -path exists between power and ground for each data bus bit in contention. Note that neither the SRAM driver nor controller driver resistance are constant, as they alternately turn on and off. Therefore, 56.25 mA is sourced between power and ground for worst case 3.6 V V_{DD} and the current drain of the 18 bits in contention is

$$I_{CONTENTION} = \frac{V_{DD}}{Z_T} \times \text{No. of Pins} = \frac{3.6 \text{ V}}{64 \Omega} \times 18 \text{ Pins} = 1.0125 \text{ A}$$

or

$$P_{CONTENTION} = \frac{V_{DD} \times I_{DDA}}{2} = \frac{3.6 \text{ V} \times 1.0125 \text{ A}}{2} = 1.8225 \text{ W}$$

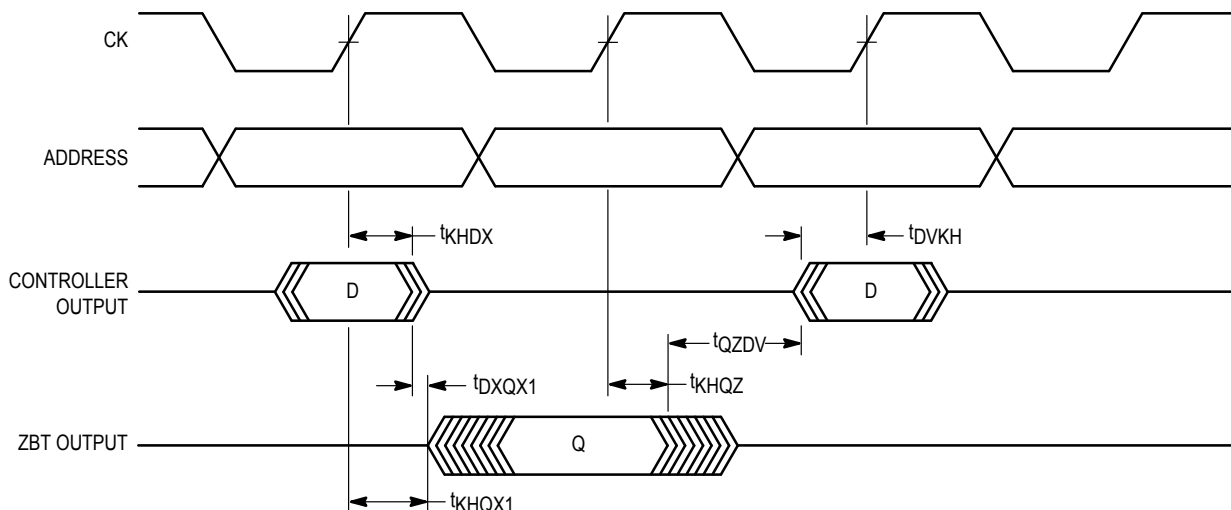


Figure 6. Bus Contention Timing

This may be a minimal power loss since bus contention occurs only 6.7% of the time. The added amount of contention is $1.8225 \text{ W} \times 6.7\% = 0.122 \text{ W}$ for a total of 1.38 W . The time allowed for bus contention is then

$$t_{\text{ALLOWED}} = \frac{P_{\text{ALLOWED}}}{P_{\text{CONTENTION}}} \times 2 t_{\text{KHKH}} = \frac{0.17 \text{ W}}{1.8225 \text{ W}} \times 2 \times 7.5 \text{ ns} = 1.4 \text{ ns}$$

The negative effects of bus contention are alleviated somewhat by circuit inductance and capacitance due to electrical properties that conditions are not able to change instantaneously. This buys some time before a large current path occurs. The power supply and lead inductance resist instantaneous changes in current and thus, decrease the effects of bus contention by delaying its onset. Capacitive output loading also reduces the bus contention effects since the voltage across a capacitor can not change instantaneously. The capacitor, not the devices in contention, supplies the initial current. The RC constant discharge time will increase proportionately with the load capacitance. Obviously, there are limits to the amount of load capacitance that may be added to reduce bus contention, since the power consumption is proportional to load capacitance. High bus capacitance leads to other chronic problems, such as high power consumption, excessive heat, device breakdown, and shifting threshold voltages.

The primary focus is to ensure the junction temperature remains within the acceptable levels. Assume a worst case bus contention occurs once every two cycles and the driver resistance is constant. The controller is driven low at the same time the SRAM is driven high (worst-case). The ZBT RAM

junction temperature dissipating 1.26 W is 105.3°C at 70°C ambient and θ_{JA} of 28°C/W . The junction temperature increases 3.4°C to 108.6°C with bus contention, which is within specifications. Therefore, a manageable amount of bus contention is acceptable, but must be understood by the designer. The I/O pins and power supply are isolated to prevent driver or core logic damage due to bus contention.

Based on characterization data, the current increases by about 1 mA per I/O during bus contention. The Freescale ZBT RAM can tolerate t_{KHDX} up to 2.6 ns at 143 MHz . Therefore, bus contention up to 1.1 ns is acceptable with $t_{\text{KHQX}} = 1.5 \text{ ns}$. Likewise, 2.1 ns of margin is available for $t_{\text{KHDX}} - 500 \text{ ps}$. The ZBT RAM passes functionality at 133 MHz with $t_{\text{KHDX}} = 3.1 \text{ ns}$, and 100 MHz with $t_{\text{KHDX}} = 4.4 \text{ ns}$. Figure 7 shows the bus contention curve during a back-to-back write-read. Figure 8 illustrates the bus contention curves for 100 MHz , 133 MHz , and 143 MHz .

CONCLUSION

The ZBT RAM is the new synchronous SRAM standard for high performance communication designs and represents a dramatic improvement over other SRAMs for systems requiring back-to-back reads and writes. The ZBT RAM maximizes bus bandwidth and data throughput by eliminating idle bus cycles during turnaround from read-to-write or write-to-read. The control logic for the ZBT RAM provides simplified control of the SRAM with minimal control pins and no external logic. While special design considerations must be taken because of bus contention, the ZBT RAMs offer the best system solution for the networking and communications markets.




Figure 7. Bus Contention Curve




Figure 8. Bus Contention Current Per I/O

Information in this document is provided solely to enable system and software implementers to use Freescale Semiconductor products. There are no express or implied copyright licenses granted hereunder to design or fabricate any integrated circuits or integrated circuits based on the information in this document. Freescale Semiconductor reserves the right to make changes without further notice to any products herein. Freescale Semiconductor makes no warranty, representation or guarantee regarding the suitability of its products for any particular purpose, nor does Freescale Semiconductor assume any liability arising out of the application or use of any product or circuit, and specifically disclaims any and all liability, including without limitation consequential or incidental damages. "Typical" parameters which may be provided in Freescale Semiconductor data sheets and/or specifications can and do vary in different applications and actual performance may vary over time. All operating parameters, including "Typicals" must be validated for each customer application by customer's technical experts. Freescale Semiconductor does not convey any license under its patent rights nor the rights of others. Freescale Semiconductor products are not designed, intended, or authorized for use as components in systems intended for surgical implant into the body, or other applications intended to support or sustain life, or for any other application in which the failure of the Freescale Semiconductor product could create a situation where personal injury or death may occur. Should Buyer purchase or use Freescale Semiconductor products for any such unintended or unauthorized application, Buyer shall indemnify and hold Freescale Semiconductor and its officers, employees, subsidiaries, affiliates, and distributors harmless against all claims, costs, damages, and expenses, and reasonable attorney fees arising out of, directly or indirectly, any claim of personal injury or death associated with such unintended or unauthorized use, even if such claim alleges that Freescale Semiconductor was negligent regarding the design or manufacture of the part.