✕ XILINX ®

Synthesizable 200 MHz ZBT SRAM Interface

XAPP136 (v2.0) January 10, 2000   Author: Shekhar Bapat

Summary

The Virtex™ series and the Spartan™-II family of FPGAs provide access to a variety of on-chip and off-chip RAM resources. In addition to the on-chip SelectRAM™ and Block SelectRAM+™, a Virtex or Spartan-II design can interface to Megabytes of external high-speed SRAM and DRAM. The combination of the high-speed SelectI/O resource and on-chip CLKDLL (Clock Delay-Locked Loop) enables the interface to operate at maximum RAM speeds. A Virtex or Spartan-II interface to ZBT (Zero Bus Turnaround) SRAM provides interleaved Read/Write without wasteful turnaround cycles.

Introduction

Beyond the obvious benefits of speed and size, the Virtex series and Spartan-II family offer several memory-related advantages.

- SelectRAM
- Block SelectRAM+
- High-speed Interfaces

SelectRAM

Virtex and Spartan-II FPGAs provide SelectRAM or LUTs (Look-Up Tables) configured as small bits of RAM. The SelectRAM can be configured as a single-port 32 x 1 RAM or a dual-port 16 x 1 RAM (one read and one write port), or a single-port 16 x 2 RAM. It provides shallow RAM distributed throughout the chip and is well suited for certain applications.

Block SelectRAM+

Block SelectRAM+, available in 4K blocks, is a fully-synchronous true dual-port memory. Each port allows reads and writes on independent clocks and can access the memory in 4K x 1, 2K x 2, 1K x 4, 512 x 8, or 256 x 16 configurations. These independently-configurable ports enable the RAM blocks to be used as buffers for high-speed data streams and to funnel data to different width/speed combinations. The blocks can be combined to create wider or deeper memory. The dual-port Block SelectRAM+ has been used to create FIFOs with independent clocks running at 170 MHz. Block SelectRAM+ offers advantages in many networking and telecommunication applications that require memory updates without delaying read access.

High-speed Interfaces

The third part of the memory hierarchy comprises pre-engineered interfaces to off-the-shelf memory parts from leading memory vendors. The combination of the high-speed SelectI/O resource and on-chip CLKDLL enable the interface to operate at the maximum RAM speeds. The Virtex and Spartan-II off-chip memory solutions currently include interfaces to ZBT SRAM (pipelined and flowthrough), SDRAM, SGRAM and DDR SDRAM. These pre-engineered solutions incorporate a synthesizable Verilog net list, timing constraints in the Xilinx UCF format, and special switches for software (where applicable). The focus of this document is on interfacing Virtex devices with a particular static RAM, namely the ZBT SRAM.

A Primer on ZBT SRAM

The term "Zero Bus Turnaround" was coined by IDT in October 1996. IDT, Micron, and Motorola developed the ZBT technologies independently, and all three companies are offering compatible ZBT products based upon the same architecture.

ZBT SRAM devices are synchronous SRAMs that provide maximum system throughput by utilizing every bus cycle. As the name implies, there are no turnaround cycles required when switching between read and write cycles. Consequently, there are no wasted cycles and the system actually delivers the stated bandwidth. This feature is particularly beneficial in applications with many random, intermixed read and write operations on the data bus as opposed to long bursts of reads or writes. Examples of such applications include network interface cards, LAN and WAN switching, gigabit switching, ATM switching, and switch or hub-shared fabric and router tables. Figure 1 shows alternate writes and reads and the resulting bus utilization for ZBT SRAM and contrasts it with the regular SRAM. Traditional pipelined SRAMs have two turnaround cycles and consequently suffer a larger degradation in performance.
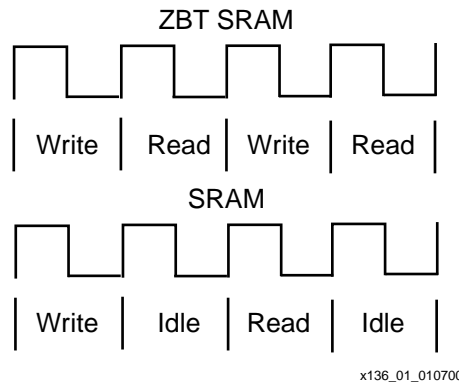
ZBT SRAM

| Write | Read | Write | Read |

SRAM

| Write | Idle | Read | Idle |

x136_01_010700

Figure 1: ZBT SRAM vs. Regular SRAM

ZBT SRAM devices come in pipelined and flowthrough varieties. Pipelined devices can operate at higher frequencies than flowthrough devices, but flowthrough devices have shorter latency. Control signals on both types are identical, however, and are simpler than for a Sync Burst SRAM. Many designers opt to use only the read/write control signal, further simplifying the controller and minimizing FPGA I/O usage. The control signals "Chip Enable" and "Output Enable" can be used to create a bank of ZBT SRAMs.

Pipelined ZBT SRAM

Pipelined ZBT SRAMs are late-late-read RAMs in the sense that data appears two clock cycles after address (Figure 2). In order to allow interleaved reads and writes, it is also a late-late-write RAM. During a read operation, valid data appears two clock cycles after the address and control signals. Similarly, during a write operation, the data is input two clock cycles after the address and control signals are registered.
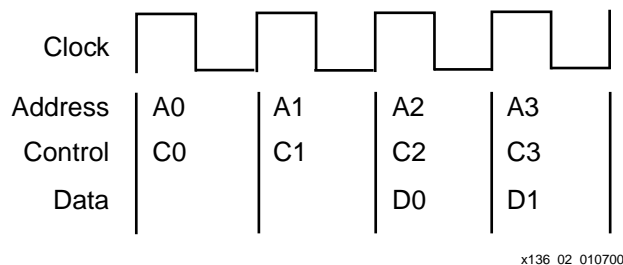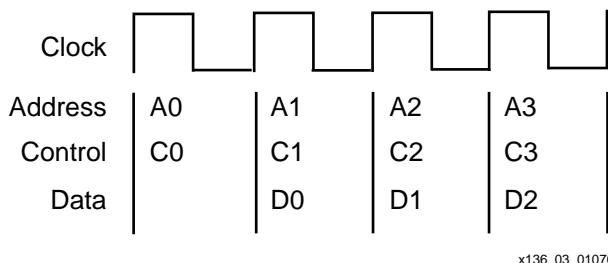
Clock				
Address	A0	A1	A2	A3
Control	C0	C1	C2	C3
Data			D0	D1

x136_02_010700

Figure 2: Bus Operations for Pipelined ZBT SRAM

Flowthrough ZBT SRAM

Flowthrough ZBT SRAMs are late-read RAMs in that data appears one clock cycle after address (Figure 3). They are also late-write RAMs.


x136_03_010700

Figure 3: Bus Operations for Flowthrough ZBT SRAM

Device Interface with ZBT SRAM

Figure 4 shows a simple interface between a Xilinx device and a ZBT SRAM. In addition to the mandatory address and data signals, the interface only needs a read/write (RW) control signal. To ensure high performance, the design uses two CLKDLLs: one to de-skew and generate a 2x controller clock and another to de-skew and generate a board-level 2x clock. The result is a high-speed, de-skewed clock driving the controller and the ZBT SRAM. The ZBT SRAMs are available with LVTTL I/Os. The Virtex SelectI/O resource supports LVTTL signaling with a range of driver strengths from 2 mA to 24 mA and a choice of fast or slow slew rates.
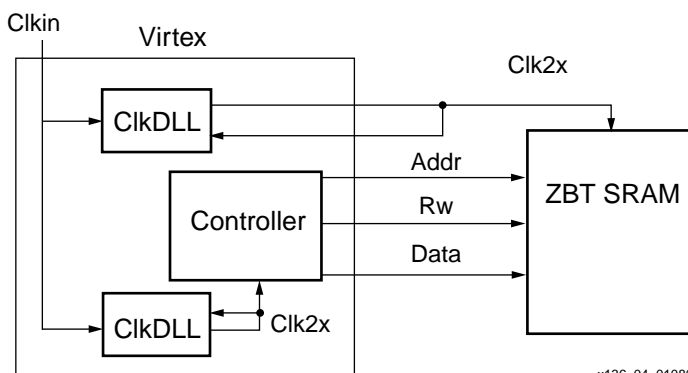

x136_04_010800

Figure 4: Virtex device and ZBT SRAM

A ZBT SRAM controller design in synthesizable Verilog for a 64K x 36 pipelined device, as well as for a 64K x 36 flowthrough device, is provided. Controllers for other sized ZBT SRAMs are easily derived from these two designs. In these examples, RW is the only control signal and is used to control the 3-state drivers for the 36 bits of data. In order to run the Virtex Pipelined ZBT SRAM interface at 200 MHz (100 MHz for flowthrough), CLKDLLs are employed to de-skew and double the clock frequency.

During a write operation to the pipelined ZBT SRAM, data must be provided to the RAM two cycles after the address and control signals. Assuming that data arrives in lock-step with address and control, the controller drives the data to the RAM through two pipelined stages. While the control signals to the RAM are sent directly, the control signals inside the controller are delayed through an equal number of pipelined stages to match the data latency. The pipelined example uses a flip-flop tree to minimize the fan-out for the RW signal, which in turn facilitates high-speed data access.

Xilinx's SelectI/O resource provides a wide variety of standards and, for LVTTL, offers a wide range of driver strengths. The example design employs a fast slew driver with 16 mA drive strength.

x136_06_010800

Figure 6: Contention Current

A Verilog design (Figure 5) consists of three top-level modules: a controller, a tester, and a DLL. The DLL module provides de-skewed and 2x multiplied clocks for Virtex and board-level designs. The tester module receives a reset signal from outside the Virtex device and generates the address, data, and read/write signals for the controller module. The tester module also compares the data values that are read back from the RAM to the expected values and generates an error signal if the values do not match.
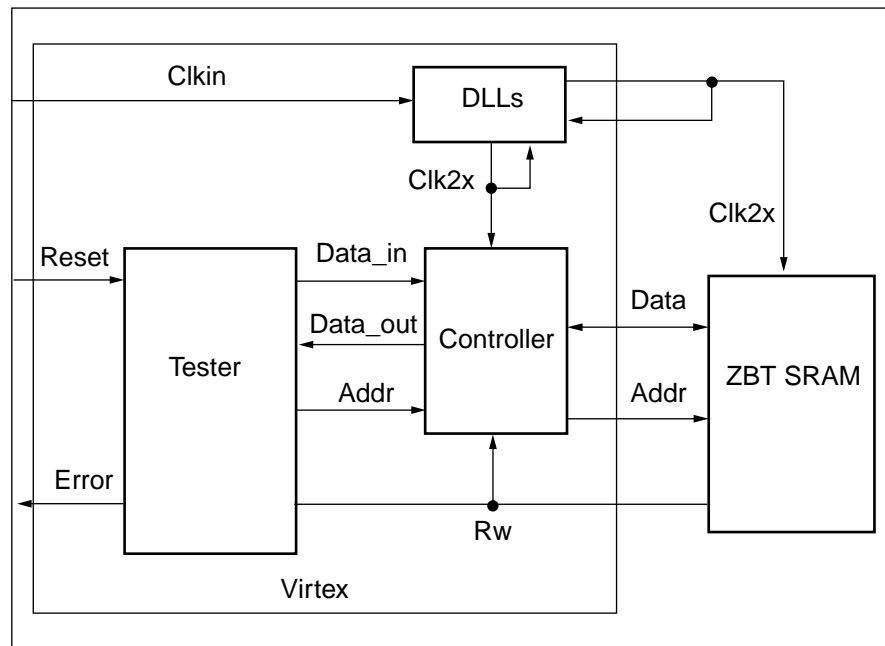
x134_05_010800

Figure 5: Block Diagram of the Verilog Design

Bus Contention

ZBT SRAMs maximize bus utilization by eliminating turnaround cycles. This enhances system performance, but the probability of "bus contention" increases. Bus contention occurs when two devices on the bus attempt to drive the bus to opposite logic values. Consider the example of an FPGA interfacing with a ZBT SRAM. During the read-write transition, the FPGA can actually start driving the bus before the ZBT SRAM device stops driving the bus. If the FPGA is driving a logic High onto the bus and the ZBT SRAM is driving a logic Low, a current path is created from the FPGAs V_{CC} to the ZBT SRAMs Ground through the "on" transistors of the two devices, as shown in Figure 6a. Likewise, if the FPGA drives a logic Low and the ZBT SRAM drives a logic High, a current path is also created (as shown in Figure 6b).

When both the devices drive the same value onto the bus, there is no current path between V_{CC} and Ground and, therefore, no contention current (as shown in Figure 7).


Figure 7: No Contention Current

Bus contention can also occur during device turn-off or turn-on times. Turn-off time is the time it takes to drive device outputs to a high impedance (High-Z) state, where both pull-up and pull-down transistors are off with no current path through the output drivers. Conversely, Turn-on time is the time it takes to drive device outputs to a logic High or Low state. During the read-to-write transition, bus contention can occur if the turn-on time of the FPGA is less than the turn-off time of the ZBT SRAM. Similarly, during the write-to-read transition, bus contention can occur if the turn-on time of the ZBT SRAM is less than the turn-off time of the FPGA.

Worst-case contention time for the read-to-write transition is derived by calculating the difference between the slowest turn-off time for the ZBT SRAM and the fastest turn-on time for the FPGA. Similarly, the worst-case contention time for the write-to-read transition is derived by figuring the difference between the fastest turn-off time for the FPGA and the slowest turn-on time for the ZBT SRAM. The fastest and the slowest switching times are observed at opposite extremes of temperature. Arguably, two chips on a board would never be at opposite temperature extremes, so any contention time would actually be less than predicted by worst case analysis.

Bus Contention Analysis

This worst-case contention analysis uses an XCV300-6 Virtex FPGA with LVTTL fast slew-rate 16 mA drivers in a BG432 package. The ZBT SRAM timing numbers are from Micron data sheets. The Micron -7 speed grade is used for the pipelined version and the -10 speed grade is used for the flowthrough design.

Pipelined ZBT SRAM and a Virtex Device

The slowest turn-off time for the Virtex device is 1.1 ns using the CLKDLL. The fastest turn-on time for the pipelined ZBT SRAM is 1.5 ns. Therefore, during a write-to-read transition bus contention will not occur. The slowest turn-off time for the pipelined ZBT SRAM is 3.5 ns and the fastest turn-on for the Virtex device with CLKDLL is 1 ns. During a read-to-write transition, there would be worst-case contention current of 30 mA for 1.5 ns in addition to 1 ns of current as part of the regular operation. Assuming that the read-to-write transition occurs only half of the time, average contention current for a 7 ns clock period is 3.2 mA per data pin. Also assuming a voltage drop of 1.5V across the Virtex transistor, there is additional power dissipation of 4.82 mW per data pin. Figuring on average that only half the data pins are in contention, the average power dissipation due to contention is 86.7 mW. For a BG432 package, the temperature increase would be 1.1°C.

Flowthrough ZBT SRAM and a Virtex Device

The slowest turn-off time for the flowthrough ZBT SRAM is 5 ns and the fastest turn-on time for the Virtex device is 4 ns. Again, there is no bus contention during a write-to-read transition. During the read-to-write transition, the worst-case bus contention is 3 ns, in addition to 1 ns of current as part of the regular operation. For a 10 ns clock, the average contention current is 4.5 mA per data pin with power dissipation of 6.75 mW per data pin. For a BG432 package, the additional increase in temperature would be 1.5°C. (Please note again that the worst case contention would be observed at the temperature and voltage extremes. It is not expected that two devices on the same board would be at opposite temperature extremes. In reality, any contention current that may occur would be less than worst-case analysis predictions.)

Reference Design

The reference design is available at:

ftp://ftp.xilinx.com/pub/applications/xapp/

The Verilog code: xapp136.tar.gz & xapp136.zip

The VHDL code: xapp136_vhdl.tar.gz & xapp136_vhdl.zip

Revision History

Date	Version #	Revision
04.01.99	1.0	Initial release.
04.06.99	1.1	Modified ucf statements
06.25.99	1.2	Modified source code for both Flowthrough and Pipelined controller designs.
06.29.99	1.3	Updated Figures 4, 5, & 6 and added corporate signature block
07.20.99	1.4	Added a missing ';' to the first line of code on page 5: `NET dq* NODELAY ;`
09.23.99	1.5	Removed code and updated for Virtex-E. See Reference Design for details.
1.10.00	2.0	Reformatted text, corrected some errors and added Spartan-II device support.