WP111 (v1.0) February 16, 2000

Spartan-II Family as a Memory Controller for QDR-SRAMs

Authors: Amit Dhir, Krishna Rangasayee

Introduction

The explosive growth of the Internet is boosting the demand for high-speed data communication systems. While RISC CPU speeds have exceeded clock rates of 500 MHz, static memories have been unable to keep up the pace. In order to increase memory bandwidth significantly for future high-performance communication applications, Cypress Semiconductor, Integrated Device Technology, Inc. and Micron Technology have jointly defined and developed a new SRAM architecture referred to as the Quad Data Rate™ (QDR™) SRAM technology. This architecture is aimed at the next generation of switches and routers that operate at data rates above 200 MHz, and will serve as the main memory for lookup tables, linked lists, and controller buffer memory.

This partnership of the three companies enables customers to choose these new SRAMs from multiple sources. Data throughput of 11.592 Gbits/s is possible, which is about four times the performance of comparable SRAMs in today’s market.

Any new SRAM architecture requires supporting circuitry for both interfacing and control. FPGAs are ideal to implement the control and interface logic, which ties the CPUs to the QDR SRAMs. The Spartan™-II FPGA, with its unique and extensive features is an ideal memory controller interface for the QDR SRAM.

Spartan-II FPGAs offer more than 100,000 system gates at under $10 and are the most cost-effective programmable logic devices (PLD) solution ever offered. It uses a leading 0.18 µm, six-layer metal process. The Spartan-II family addresses low cost and fast time-to-market, but more importantly integrates powerful new system-level features that provide an attractive solution for today’s system level designer. They build on the capabilities of the very successful Virtex family and incorporate all the associated features, including SelectI/O^{TM}, BlockRAM^{TM}, Distributed RAM, Delay-Locked Loops (DLLs), clock speeds up to 200 MHz, and aggressive power management. The Spartan-II family offers a solution with high performance at a low cost, hence expanding the time-to-market advantage that PLDs traditionally offer. It also increases the value of the ASSP by allowing end users to customize their solutions.

Evolution of the QDR SRAM Architecture

SRAMs are widely used in data communication systems, due to their fast, low-latency access to the CPU. Typically, the fastest available SRAMs are employed as data buffers, link-list tables, and pointer tables.Quick memory access is the key to delivering the new high bandwidth required by networking applications.

Most applications started using asynchronous SRAMs, which operated in the 10-15 ns speed range. With the increased bandwidth demand of networking applications, there came the existence of pipelined-burst SRAM (PBSRAM).

The PBSRAMs allowed networking applications to operate with a higher bandwidth and the interface was simplified by employing synchronous transfers rather than an asynchronous control. PBSRAMs were optimized for PC cache applications where the memory access is dominated by reads with much fewer writes. Hence, the wait states between reads and writes do not limit the performance of the caches. Since networking applications typically require equal amounts of reads and writes to memory, PBSRAMs offer limited incremental performance.

The no-bus-latency SRAMs (NoBL™ SRAMs) and zero-bus-turnaround (ZBT™) devices are SRAM architectures modified for networking applications. These allow operation without any wait states between reads and writes. NoBL and ZBT SRAMs enable the complete use of memory bandwidth, which significantly improves the bus utilization of networking applications.

The QDR architecture was developed by the consortium (between Cypress, IDT and Micron) to further improve the bandwidth of the interface and overcome the several limitations of PBSRAM, ZBT and NoBL SRAMs.

QDR Technology

The QDR SRAM has separate input and output ports for both read and write operations. Although those ports share address lines, separate differential clocks exist for the inputs and output ports. Data can be transferred using double-data-rate (DDR) protocols on both the input and output ports. Four words can be transferred on every clock cycle, two in and two out of the device. Hence the name, Quad Data Rate. Figure 1 shows the read and write process in a regular SRAM vs. a QDR SRAM.

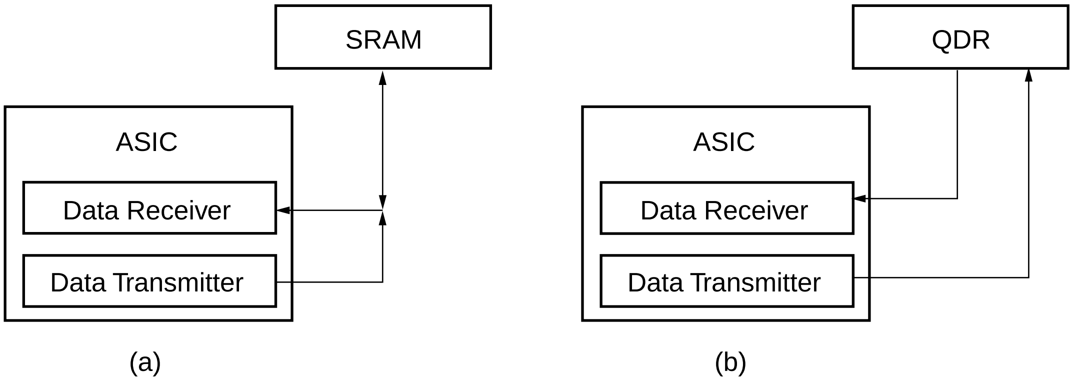

WP111_01_020900

Figure 1: **Read and Write Process in (a) Regular SRAM; (b) QDR SRAM (Courtesy QDR SRAM Consortium)**

The separate input and output ports of the QDR memory remove the possibility of bus contention and simplify the design. This facilitates its application in high frequency designs.

These SRAMs are currently available in two types: QDR2 and QDR4. The difference between the two is the number of words of data that can be obtained from the memory on a single read or write. The QDR2 and QDR4 provide two and four words of data on a single read, respectively. The consortium will support both QDR2 and QDR4. Table 1 lists the QDR devices from the consortium.

Table 1: **QDR Devices from the QDR SRAM Consortium**

	Type	Cypress	IDT	Micron
QDR2	Two Word Burst	CY7C1302	71T628	MT54V51218E
QDR4	Four Word Burst	CY7C1304	71T648	MT54V51218A

The basic block diagram of a QDR2 device is shown in Figure 2.

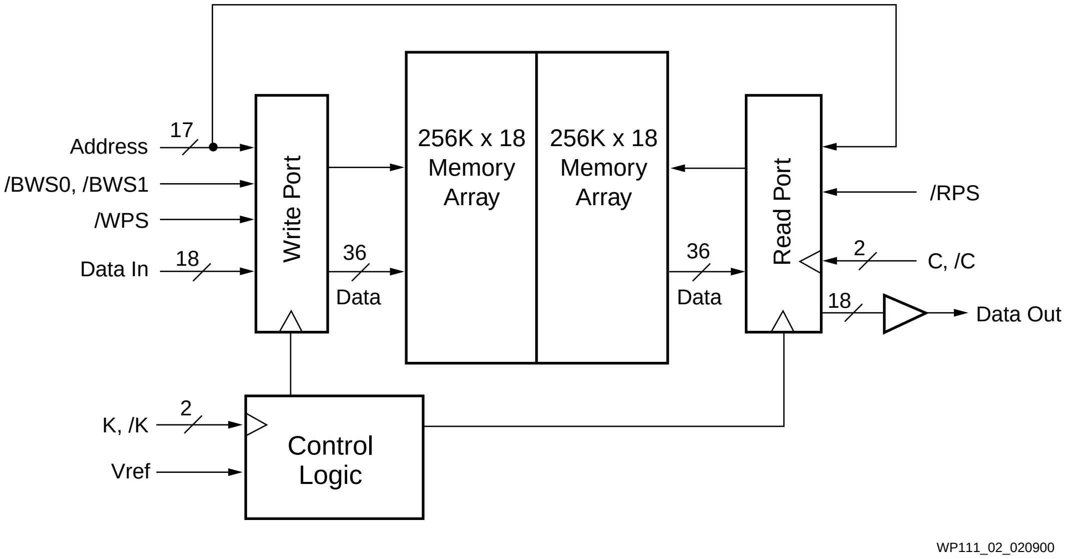

Figure 2: ***Block Diagram of the QDR-SRAM (courtesy: Cypress Semiconductor)***

The CY7C1302 QDR2 device by Cypress Semiconductor consists of a 512K word by 18-bit memory array with separate pins for the input and output data. The address lines are common, while the clocks are separately provided, for the read and write ports. Separate read and write ports distinguish the QDR SRAM from previous high-performance static memories. Simultaneous reads and writes from both ports double the throughput over standard DDR SRAMs.

If a read and write operation is started to the same address in the same cycle, the SRAM forwards data from the read port to the write port and ensures that valid data is driven out on the data bus. This guarantees data coherency in all cycles.

QDR Architecture Advantages

The QDR SRAM architecture was designed to overcome the problems of the PB-, NoBL and ZBT SRAMs. The several benefits of the QDR are described in this section.

Separate I/O

QDR SRAMs have separate input and output data ports, which solves issues that have dogged all common I/O devices. Bus Contention is one such problem which occurs frequently during a read or write operation in a networking environment. This happens when the SRAM drives out data on a read operation faster than the data source can remove data from the bus after a write. Most common I/O devices are susceptible to this since the same bus has to be used for reads from and writes to the SRAM. As the operating frequencies increase, the chance of bus contention also rises dramatically. Separating the read and write data buses ensures that bus contention never occurs.

Common I/O devices also suffer from a constant flow of data that requires the bus to be turned around for a read and write. This causes a non-uniform flow of data between the controller and the SRAM. With its separate I/O, the QDR SRAM permits a constant data flow, hence providing a higher throughput than the SRAMs with common read and write I/Os.

Maximum Frequency/Bandwidth

QDR SRAMs permit operation in frequency ranges which previously could not be achieved with earlier generations of SRAMs. Standard PBSRAM, ZBT and NoBL SRAMs can also be made to operate at very high frequencies, but the inherent limitation of operating a common I/O data path at such high levels limits the maximum bus frequency. Because they lack the overhead limitations of turning data buses around on read/write transactions, QDR SRAMs can operate at the native frequency of the SRAM.

While standard SRAMs are all single data rate devices, QDR SRAMs allow double data rate (DDR) transfers on both the inputs and outputs, thereby significantly improving the throughput over standard SRAMs.

Voltage Migration

The signal levels used on the I/O pins of QDR SRAMs operate at HSTL (High-Speed Transceiver Logic) voltage levels. This lower voltage level swing is required to support the high-speed operation of the inputs and outputs. QDR SRAMs provide users a quantum leap in bandwidth with a migration path to higher frequencies and lower voltages. Using 166 MHz QDR devices in conjunction with the 166 MHz capable Spartan-II FPGA family, QDR memories can significantly improve network system performance between a cache to a networking type of application. The QDR SRAM has an 18-bit interface while synchronously pipelined, NoBL and DDR SRAMs have 36-bit interfaces.

Spartan-II FPGAs as a Programmable Memory Controller Solution for QDR SRAMs

PLDs provide users with reduced time-to-market and a low-risk means to implement new designs. Since the networking market changes rapidly, both these factors are important to manufacturers of networking equipment. PLDs are often used as memory controllers to simplify the process of designing the interface to the QDR SRAM. Currently no commercial memory controller ASSPs are available, because QDR SRAM is a recently developed architecture. The Spartan-II FPGA with its flexibility through reprogrammability and extensive features can be used to implement the control logic which can be customized to the needs of the memory interface. The block diagram of the Spartan-II memory controller is shown in Figure 3.

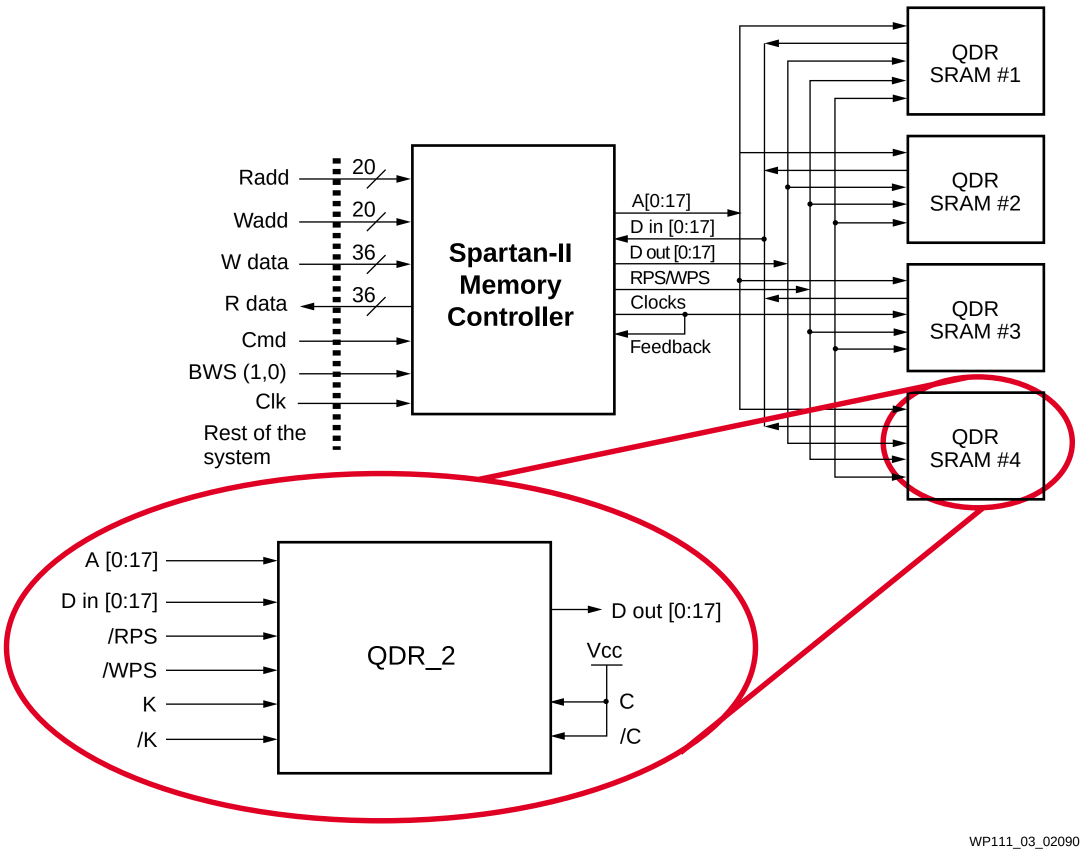

Figure 3: **Block Diagram of the Spartan-II Memory Controller**

The basic memory-control system for QDR SRAMs can be implemented in a Spartan-II FPGA that ties the host system to a bank of memories. The memory controller circuit (shown in Figure 3) is designed to interface four QDR SRAMs, which are connected in a depth-expanded mode to a host CPU. Each of the QDR SRAM receives separate control signals for the read and write ports, while the address and data ports are common for all the SRAMs. The SRAMs are configured as a
2-Mword by 18-bit storage array. The controller generates all the signals for the memory bank and makes the complete SRAM bank look like a unified memory bank. It also supports concurrent DDR operations on all of the input and output signals and lets byte-write operations into the memory bank.

Operating in the single-clock mode, the controller significantly simplifies the memory interface. At 100 MHz, it boasts a bandwidth of 7.2Gbits/s. The controller employs a command based interface with a 2-bit command input (01 read, 10 write, 11 read/write) and has independent read and write state machines. The other inputs to the controller are the clock (CLK), write address (W_{ADD}), read address (R_{ADD}), write data (W_{DATA}), read data (R_{DATA}) and the byte-write control (BWS[0,1]) signals.

From the controller's point, all signals are relative to the SRAM clocks. It provides the address, data and control inputs within the setup/hold time window requirements of the SRAM. The state machine provides the addresses and data on certain clock edges, while the SRAM latches the addresses and data on those edges.

The memory controller polls the command signal (CMD) on its input port on every clock. Depending on the clock, read, write or read/write operations are completed on the bank of memory. The read port selects (/RPS) and the write port selects (WPS) for the different SRAMs are generated by the memory controller depending on the state of the command (CMD [0,1]) inputs and the higher order address lines.

Traditionally, designing memory systems that operate at speeds above 100 MHz has required external devices to minimize clock skew. The Spartan-II FPGA incorporates four on-chip Delay-Lock Loops (DLLs) which are used to deskew the internal global clock network, and deskew clocks fed off-chip to other system components. The on-chip DLLs eliminate the need for external clock-management devices, hence simplifying system design.

A pair of DLLs on the Spartan-II FPGA, as shown in Figure 4, achieve zero clock skew between the FPGA on-chip clock and the QDR SRAM clock. In addition to clock deskewing, the Spartan-II DLLs provide valuable features like phase adjustment, frequency division, and frequency multiplication. When working with DDR and QDR memory devices, the availability of a double frequency clock that is phase-locked to the system clock is particularly important. With the on-chip global clock distribution network, high-speed synchronous I/O resources, and the SelectI/O lines on the Spartan-II family that can be programmed to accommodate different signaling standards, the QDR SRAM interface achieves a phenomenal data throughput of 7.2 Gbits/s.

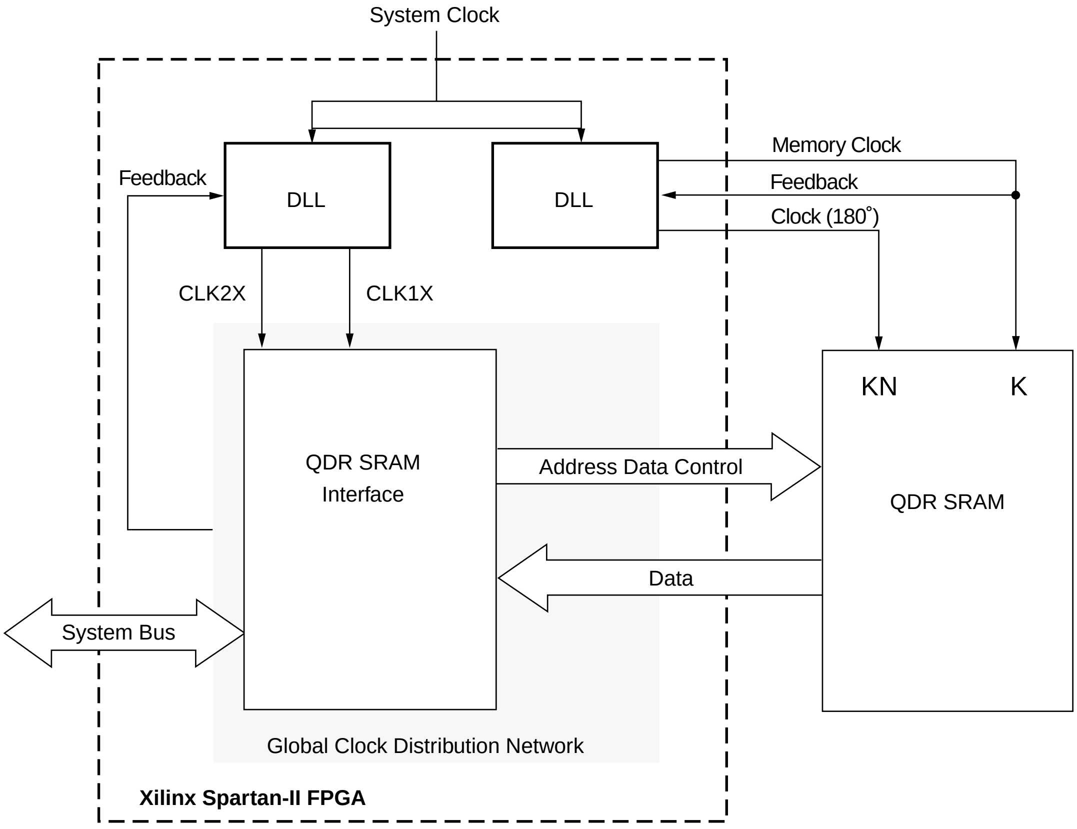

WP111_04_020900

Figure 4: **Dual DLLs Allow the Spartan-II FPGA-based Controller to Minimize Clock**

The Spartan-II based memory controller supports DDR transfers on all I/Os, byte-level operations on the memory bank, and concurrent reads and writes to all SRAM blocks. It also provides HSTL-compatible interfaces to the QDR SRAMs. At a system level, the controller must have 20 write-address lines, 20 read-address lines, 36 read-data lines, 36 write-address lines, four command signals, four byte-write signals, and one each for the clock, reset, and data-ready signals. On the QDR memory side, it has to provide 18 write-data lines, 18 read-data lines, 18 address lines, two clock lines, two byte-write lines, and one write-port and one read-port select line.

Implementing the Memory Controller in Spartan-II FPGAs

To implement the memory controller, the design can be divided into several sections. In a few lines of VHDL code (see Application Note XAPP183: Interfacing the QDR to the Xilinx Spartan-II FPGA) the internal FPGA clock can be phase locked with the system clock to generate the second system clock. Similarly, the QDR SRAM interface can be implemented with a DDR interface. The system interface is a straightforward setup of the register read, write and section operations.

The controller shows the memory interface with its three 18-bit buses and the host interface with the dual 36-bit data buses and 18-bit address buses. The Spartan-II FPGA operates internally at 200 MHz. Externally, the buses need to operate at 100 MHz onlysince the DDR interface transfers data on both the leading and trailing edges. The 36-bit read data path from the host is internally split into two 18-bit sections and latched by separate registers. These registers are clocked at 200 MHz, allowing one to send or receive data on both edges of the clock.

The four on-chip DLLs available on the Spartan-II FPGA family can deskew either the internal global clock network or clocks fed off chip to other system components. The two DLLs permit the controller to achieve zero clock skew between the FPGAs on-chip clock and the QDR SRAM clock.

The Spartan-II FPGA family implemented as a memory controller requires two DLLs, two global clock buffers, 119 I/O buffers, and with a 2.5 ns clock-to-out timing permits a good margin for operation at 100 MHz. Features like DLLs, HSTL I/O buffers, I/O and core performance, and reprogrammability enables the Spartan-II product family to be a unique fit as a memory controller.

The Spartan-II FPGA Advantage

The Spartan-II family capitalizes on the very same architectural advantages of the original Virtex family. The Spartan-II family couples eight advanced DLLs in every device with the Xilinx exclusive SelectI/O and Block SelectRAM memory to deliver superior performance. The following individual features make the Spartan-II FPGA family an ideal memory-controller for QDR-SRAMs.

High-Speed Transceiver Logic (HSTL)

The Spartan-II Series of FPGAs feature the exclusive SelectI/O technology integrating support for 17 different single-ended and differential I/O standards. HSTL is one of the single-ended I/O interfaces supported by every Spartan-II device (and on every single I/O), eliminating the need for external level translators to interface with high-speed memories and reducing overall system design complexity and cost. This interface standard was developed for voltage scalable and technology independent I/O structures. Since HSTL compliance does not specify device supply voltages, it is a process-independent I/O standard. The HSTL nominal logic switching range is 0.0V to 1.5V, resulting in faster outputs with reduced power dissipation, and minimized EMI radiation. The SelectI/O technology gives system designers enhanced flexibility in optimizing system performance with adjustable trip-point (V_{REF}) and output power supply voltage (V_{CCO}).

In computing, slow memory access times have traditionally hindered fast processor operations. In the mid-frequency range (between 100 MHz and 180 MHz), the I/O interface options for all single ended signals are; HSTL, GTL/GTL+, SSTL, and LVTTL. Beyond 180 MHz, the HSTL standard is the only single ended I/O interface available.

With HSTL speeds, faster I/O interface significantly improves overall system performance. HSTL is the I/O interface of choice for high-speed memory applications, and is ideal for driving address buses to multiple memory banks. Spartan-II FPGAs are the leading PLD solutions with integrated HSTL I/Os for memory intensive designs at very economical prices.

Delay-Locked Loop

The Spartan-II family provides four fully digital dedicated on-chip DLL circuits, allowing for very precise synchronization of external and internal clocks. These circuits provide zero propagation delay, low clock skew between output clock signals distributed throughout the device, and advanced clock domain control.

Each DLL can drive up to two global clock routing networks within the device. The global clock distribution network minimizes clock skews due to loading differences. By monitoring a sample of the DLL output clock, the DLL can compensate for the delay on the routing network, effectively eliminating the delay from the external input port to the individual clock loads within the device.

In addition to providing zero delay with respect to a user source clock, the DLL can provide multiple phases of the source clock. The DLL can also act as a clock doubler or it can divide the user clock by up to 16. Clock multiplication gives the designer a number of design alternatives. It can simplify the board design because the clock path on the board no longer distributes such a high-speed signal. A multiplied clock also provides the designers the option of time-domain multiplexing, using one circuit twice per clock cycle, consuming less area than two copies of the same circuit.

The DLL can also act as a clock mirror. By driving the DLL output off-chip and then back again, the DLL can be used to deskew a board level configuration process until after the DLL achieves lock.

By taking advantage of the DLL on the Spartan-II family to remove on-chip delay, the designer can greatly simplify system level designs involving high-fanout, high-performance clocks.

Block SelectRAM

The Spartan-II FPGAs contain dedicated blocks of true dual-port RAM, known as Block SelectRAM memory. The dedicated memory provides a cost-effective use of resources without sacrificing the existing distributed SelectRAM memory or logic resources. The Block SelectRAM memory is fully synchronous for easy timing analysis and is easily initiated at configuration. This integrated capability makes the Spartan-II family ideal for cost-sensitive applications.

The Block SelectRAM memory in the Spartan-II family provides four to 12 blocks of 4K bits of dual-port RAM. Both ports are configurable to any size from 4K x 1 to 256 x 16, allowing built-in bus width conversion. Each port is completely independent and fully synchronous, allowing for the creation of flexible RAM structures. The dedicated RAM serves as a very high speed memory sourece comparable to discrete memories or ASIC memory cores.

This memory allows the Spartan-II FPGAs to handle on-chip memory requirements, which tend to grow faster than logic densities. It can be used for FIFOs to buffer data on and off chip, caches for high-speed parallel searches, ATM packet buffers, or sample bus-width converters. The integration of RAM and logic opens new application possibilities such as shift registers, state machines, and register stacks. Consolidating these memory-based functions into Spartan-II FPGAs further enhances the logic integration and cost effectiveness of the family.

Conclusions

QDR SRAMs are designed to greatly increase memory bandwidth compared to existing SRAM solutions in applications such as switches and routers. These SRAMs will serve as the main memory for look-up tables, linked lists and controller buffer memory to enhance bandwidth in future data communication systems. A family of high-performance QDR SRAMs is defined to ensure customers have the security of consistent multiple supplier roadmaps. In DDR and QDR memories the double-frequency clock is phase locked to the system clock. Used in conjunction with the on-chip global clock-distribution network, high-speed synchronous I/O resources, and SelectIO flexible signaling standards, the Spartan-II FPGA-to-QDR SRAM interface achieves a data throughput of 7.2 Gbits/s. All QDR signals are registered in the I/O buffers and use HSTL buffers.

The flexible and powerful features of the Spartan-II FPGA makes it make it an ideal solution for high-speed QDR SRAM memory controllers. The Spartan-II family addresses low cost and fast time-to-market, but more importantly integrates powerful new system-level features that provide an attractive solution for today's system level designer. They build on the capabilities of the very successful Virtex family and incorporate all the associated features, including SelectI/O, BlockRAM, Distributed RAM, DLLs, clock speeds up to 200 MHz, and aggressive power management. This family of FPGA supports 17 different I/O standards and supports V_{CCIO} ranging from 1.8V to 5V. It also supports four DLLs that can be used for many different applications. Achieving the full benefits of the QDR SRAM relies upon the performance and features of the Spartan-II FPGA family. The SRAM utilizes all of the Spartan-II family's features like HSTL voltage levels, Fast I/O Performance, DLLs, Embedded and Distributed memory.

References

1. *"Quad-Data-Rate SRAM Subsystems Maximize System Performance,"* Krishna Rangasayee and Rajesh Manapat, Electronic Design, February 7, 2000.
2. QDR SRAM Consortium – www.qdrsram.com
3. Application Notes for Spartan-II:
   - HSTL: High Speed Transceiver Logic www.xilinx.com/products/virtex/techtopic/hstl.pdf
   - Select BlockRAM: Using Block SelectRAM+ Memory in Spartan-II FPGAs—www.xilinx.com/xapp/xapp173.pdf
   - DLLs: Using Delay-Locked Loops in Spartan-II FPGAs — www.xilinx.com/xapp/xapp174.pdf
   - SelectI/O™: Using SelectI/O Interfaces in Spartan-II FPGAs — www.xilinx.com/xapp/xapp179.pdf
   - The Spartan-II Family—The Complete Package www.xilinx.com/products/spartan2/wp106.pdf
   - Application note XAPP183: Interfacing the QDR to the Xilinx Spartan-II FPGA

Revision History

The following table shows the revision history for this document.

Date	Version	Revision
02/16/00	1.0	Initial release.