

This chapter describes some of the high-speed memory selection criteria based on strengths and weaknesses, and the various Altera® FPGA devices these memories can interface with. This chapter also describes the memory component's capability and provide some typical applications where these memories are used.

The Altera IP may or may not support all of the features supported by the memory.



For the maximum supported performance supported by Altera FPGAs, refer to the [External Memory Interface Spec Estimator](#) page on the Altera website.

## Memory Overview

System architects must resolve a number of complex issues in high-performance system applications that range from architecture, algorithms, and features of the available components. Typically, one of the fundamental problems in these applications is memories, as the bottlenecks and challenges of system performance often reside in its memory architecture. As higher speeds become necessary for external memories, signal integrity gets more difficult. Newer devices have added several features to overcome this issue. Altera FPGAs also support these advancements with dedicated I/O circuitry, various I/O standard support, and specialized intellectual property (IP).

When you select an external memory device, consider the following factors:

- Bandwidth and speed
- Cost
- Data storage size and capacity
- Latency
- Power consumption

Because no single memory type can excel in every area, system architects must determine the right balance for their design.

Table 1–1 lists the two common types of high-speed memories and their characteristics.

**Table 1–1. Differences between DRAM and SRAM**

Memory Type	Description	Bandwidth and Speed	Cost	Data Storage Size and Capacity	Power consumption	Latency
DRAM	<p>A dynamic random access memory (DRAM) cell consisting of a capacitor and a single transistor. DRAM memory must be refreshed periodically to retain the data, resulting in lower overall efficiency and more complex controllers.</p> <p>Generally, designers select DRAM where cost per bit and capacity are important. DRAM is commonly used for main memory.</p>	Lower bandwidth resulting in slower speed	Lower cost	Higher data storage and capacity	Higher-power consumption	Higher latency
SRAM	<p>A static random access memory (SRAM) cell that consists of six transistors. SRAM does not need to be refreshed because the transistors continue to hold the data as long as the power supply is not cut off.</p> <p>Generally, designers select SRAM where speed is more important than capacity. SRAM is commonly used for cache memory.</p>	Higher bandwidth resulting in faster speed	Higher cost	Lower data storage and capacity	Lower-power consumption	Lower latency

## DDR, DDR2, and DDR3 SDRAM

This section describes and compares the features of the DDR, DDR2, and DDR3 SDRAM.

### DDR SDRAM

DDR SDRAM is a  $2n$  prefetch architecture with two data transfers per clock cycle. It uses a single-ended strobe,  $DQS$ , which is associated with a group of data pins,  $DQ$ , for read and write operations. Both  $DQS$  and  $DQ$  are bidirectional ports. Address ports are shared for read and write operations.

The desktop computing market has positioned double data rate (DDR) SDRAM as a mainstream commodity product, which means this memory is very low-cost. DDR SDRAM is also high-density and low-power. Relative to other high-speed memories, DDR SDRAM has higher latency—they have a multiplexed address bus, which reduces the pin count (minimizing cost) at the expense of a longer and more complex bus cycle.

### DDR2 SDRAM

DDR2 SDRAM is the second generation of the DDR SDRAM standard. It is a  $4n$  prefetch architecture (internally the memory operates at half the interface frequency) with two data transfers per clock cycle. DDR2 SDRAM can use a single-ended or differential strobe,  $DQS$  or  $DQS_n$ , which is associated with a group of data pins,  $DQ$ , for read and write operations. The  $DQS$ ,  $DQS_n$ , and  $DQ$  are bidirectional ports. Address ports are shared for read and write operations.

DDR2 SDRAM includes additional features such as increased bandwidth due to higher clock speeds, improved signal integrity on DIMMs with on-die terminations, and lower supply voltages to reduce power.




### DDR3 SDRAM

DDR3 SDRAM is the latest generation of SDRAM. DDR3 SDRAM is internally configured as an eight-bank DRAM and it uses an  $8n$  prefetch architecture to achieve high-speed operation. The  $8n$  prefetch architecture is combined with an interface that transfers two data words per clock cycle at the I/O pins. A single read or write operation for DDR3 SDRAM consists of a single  $8n$ -bit wide, four-clock data transfer at the internal DRAM core and two corresponding  $n$ -bit wide, one-half clock cycle data transfers at the I/O pins. DDR3 SDRAMs are available as components and modules, such as DIMMs, SODIMMs, and RDIMMs.

DDR3 SDRAM is more effective at saving system power, further increases system performance, lowers power, achieves better maximum throughput, and improves signal integrity with fly-by and dynamic on-die terminations.

Read and write operations to the DDR3 SDRAM are burst oriented. Operation begins with the registration of an active command, which is then followed by a read or write command. The address bits registered coincident with the active command select the bank and row to be activated (BA0 to BA2 select the bank; A0 to A15 select the row). The address bits registered coincident with the read or write command select the starting column location for the burst operation, determine if the auto precharge command is to be issued (via A10), and select burst chop (BC) of 4 or burst length (BL) of 8 mode at runtime (via A12), if enabled in the mode register. Before normal operation, the DDR3 SDRAM must be powered up and initialized in a predefined manner.

Differential strobes DQS and DQSn are mandated for DDR3 SDRAM and are associated with a group of data pins, DQ, for read and write operations. DQS, DQSn, and DQ ports are bidirectional. Address ports are shared for read and write operations. Write and read operations are sent in bursts, DDR3 SDRAM supports BC of 4 and BL of 8.

-  The DDR3 SDRAM high-performance controller only supports local interfaces running at half the rate of the memory interface.
-  For more information, refer to the respective DDR, DDR2, and DDR3 SDRAM datasheets.
-  For more information about parameterizing the DDR2 and DDR3 SDRAM IP, refer to the *Implementing and Parameterizing Memory IP* chapter.

## DDR, DDR2, and DDR3 SDRAM Comparison

Table 1–2 compares DDR, DDR2, and DDR3 SDRAM features.

**Table 1–2. DDR, DDR2, and DDR3 SDRAM Features (Part 1 of 2)**

Feature	DDR SDRAM	DDR2 SDRAM	DDR3 SDRAM	DDR3 SDRAM Advantage
Voltage	2.5 V	1.8 V	1.5 V	Reduces memory system power demand from DDR or DDR2 by 17%.
Density	64 MB to 1GB	256 MB to 4 GB	512 MB to 8 GB	High-density components simplify memory subsystem.
Internal banks	4 (fixed number of rows and columns)	4 and 8	8	Has higher page-to-hit ratio and better maximum throughput.
Bank interleaving	—	Allows bank interleaving	Allows bank interleaving	Is extremely effective for concurrent operations and can hide the timing overhead.
Prefetch	2	4	8	Lower memory core speed results in higher operating frequency and lower power operation.
Speed	100 to 200 MHz	200 to 533 MHz	300 to 1,066 MHz	Higher data rate.
Maximum frequency	200 MHz or 400 Mbps per DQ pin	533 MHz or 1,066 Mbps per DQ pin	1,066 MHz or 2,133 Mbps per DQ pin	Higher data rate.
Read latency	2, 2.5, 3 clocks	3, 4, 5 clocks	5, 6, 7, 8, 9, 10, and 11	Eliminating half clock setting allows 8n prefetch architecture.
Additive latency (1)	—	0, 1, 2, 3, 4	0, CL1, or CL2	Improves command efficiency.

**Table 1-2. DDR, DDR2, and DDR3 SDRAM Features (Part 2 of 2)**

Feature	DDR SDRAM	DDR2 SDRAM	DDR3 SDRAM	DDR3 SDRAM Advantage
Write latency	One clock	Read latency – 1	5, 6, 7, or 8	Improves command efficiency.
CAS latency	2, 2.5, 3	2, 3, 4, 5	5, 6, 7, 8, 9, 10	Improves command efficiency.
Burst length	2, 4, 8	4, 8	8	Improves command efficiency.
Termination	PCB, discrete to $V_{TT}$	Discrete to $V_{TT}$ or ODT	Discrete to $V_{TT}$ or ODT parallel termination. Controlled impedance output.	Improves signaling, eases PCB layout, reduces system cost.
ODT	—	ODT signal options of 50, 75, or 150 $\Omega$ on all DQ, DM, and DQS and DQSn signals	Parallel ODT options of RZQ/2, RZQ/4, or RZQ/6 $\Omega$ on all DQ, DM, and DQS and DQSn signals	DDR3 supports calibrated parallel ODT through an external resistor RZQ signal termination. DDR3 also supports dynamic ODT.
Data strobes	Single-ended	Differential or single-ended	Differential mandated	Improves timing margin.
Clock, address, and command (CAC) layout	Balanced tree	Balanced tree	Series or daisy chained	The DDR3 SDRAM read and write leveling feature allows for a much simplified PCB and DIMM layout. You can still optionally use the balanced tree topology by using the DDR3 without the leveling option.

**Note to Table 1-2:**

(1) The Altera DDR and DDR2 SDRAM high-performance controllers do not support additive latency, but the high-performance controller II does.

## QDR, QDR II, and QDR II+ SRAM



Quad Data Rate (QDR) SRAM has independent read and write ports that run concurrently at double data rate. QDR SRAM is true dual-port (although the address bus is still shared), which gives this memory a significantly higher bandwidth, allowing back-to-back transactions without the contention issues that can occur when using a single bidirectional data bus. Write and read operations share address ports.

The QDR II SRAM devices are available in  $\times 8$ ,  $\times 9$ ,  $\times 18$ , and  $\times 36$  data bus width configurations. The QDR II+ SRAM devices are available in  $\times 9$ ,  $\times 18$ , and  $\times 36$  data bus width configurations. Write and read operations are burst-oriented. All the data bus width configurations of QDR II SRAM support burst lengths of two and four. QDR II+ SRAM supports only a burst length of four. Burst-of-two and burst-of-four for QDR II and burst-of-four for QDR II+ SRAM devices provide the same overall bandwidth at a given clock speed.

For QDR II SRAM devices, the read latency is 1.5 clock cycles, while for QDR II+ SRAM devices it is 2 or 2.5 clock cycles, depending on the memory device. For QDR II+ and burst-of-four QDR II SRAM devices, the write commands and addresses are clocked on the rising edge of clock and write latency is one clock cycle. For burst-of-two QDR II SRAM devices, the write command is clocked on the rising edge of clock and the write address is clocked on the falling edge of clock. Therefore, the write latency is zero, because the write data is presented at the same time as the write command.

QDR II+ and QDR II SRAM interfaces use a delay-locked loop (DLL) inside the device to edge-align the data with respect to the  $\kappa$  and  $\kappa_n$  or  $C$  and  $C_n$  pins. You can optionally turn off the DLL, but the performance of the QDR II+ and QDR II SRAM devices is degraded. All timing specifications listed in this document assume that the DLL is on. QDR II+ and QDR II SRAM devices also offer programmable impedance output buffers. You can set the buffers by terminating the  $ZQ$  pin to VSS through a resistor,  $RQ$ . The value of  $RQ$  should be five times the desired output impedance. The range for  $RQ$  should be between 175  $\Omega$  and 350  $\Omega$  with a tolerance of 10%.

QDR II/+ SRAM is best suited for applications where the required read/write ratio is near one-to-one. QDR II/+ SRAM includes additional features such as increased bandwidth due to higher clock speeds, lower voltages to reduce power, and on-die termination to improve signal integrity. QDR II+ SDRAM is the latest and fastest generation. For QDR II+ and QDR II SRAM interfaces, Altera supports both 1.5-V and 1.8-V HSTL I/O standards.

-  For more information, refer to the respective QDR II and QDR II+ datasheets.
-  For more information about parameterizing the QDR II and QDR II+ SRAM IP, refer to the *Implementing and Parameterizing Memory IP* chapter.

## RLDRAM and RLDRAM II

Reduced latency DRAM II (RLDRAM II) is a DRAM-based point-to-point memory device designed for communications, imaging, server systems, networking, and cache applications requiring high density, high memory bandwidth, and low latency. The fast random access speeds in RLDRAM II devices make them a viable alternative to SRAM devices at a lower cost.

RLDRAM is partitioned into eight smaller banks. This partitioning reduces the parasitic capacitance of the address and data lines, allowing faster accesses and reducing the probability of random access conflicts. Also, most DRAM memory types need both a row and column phase on a multiplexed address bus to support full random access, while RLDRAM supports a non-multiplexed address, saving bus cycles at the expense of more pins. RLDRAM utilizes higher operating frequencies and uses the 1.8V High-Speed Transceiver Logic (HSTL) standard with DDR data transfer to provide a very high throughput.

There are two types of RLDRAM II devices—common I/O (CIO) and separate I/O (SIO). CIO devices share a single data I/O bus which is similar to the double data rate (DDR) SDRAM interface. SIO devices, with separate data read and write buses, have an interface similar to SRAM.

Compared to DDR SDRAM, RLDRAM II has simpler bank management and lower latency inside the memory. RLDRAM II devices are divided into eight banks instead of the typical four banks in most memory devices, providing a more efficient data flow within the device. Each bank has a fixed number of rows and columns. Only one row per bank is accessed at a time. The memory (instead of the controller) controls the opening and closing of a row, which is similar to an SRAM interface. RLDRAM II offers up to 2.4 Gigabytes per second (Gbps) aggregate bandwidth.

RLDRAM II uses a DDR scheme, performing two data transfers per clock cycle. RLDRAM II CIO devices use the bidirectional data pins (DQ) for both read and write data, while RLDRAM II SIO devices use D pins for write data (input to the memory) and Q pins for read data (output from the memory). Both types use two pairs of unidirectional free-running clocks. The memory uses DK and DK# pins during write operations, and generates QK and QK# pins during read operations. In addition, RLDRAM II uses the system clocks (CK and CK# pins) to sample commands and addresses and generate the QK and QK# read clocks. Address ports are shared for write and read operations.



The RLDRAM II SIO devices are available in  $\times 9$  and  $\times 18$  data bus width configurations, while the RLDRAM II CIO devices are available in  $\times 9$ ,  $\times 18$ , and  $\times 36$  data bus width configurations. RLDRAM II CIO interfaces may require an extra cycle for bus turnaround time for switching read and write operations.

Write and read operations are burst oriented and all the data bus width configurations of RLDRAM II support burst lengths of two and four. In addition, RLDRAM II devices with data bus width configurations of  $\times 9$  and  $\times 18$  also support burst length of eight.

The RLDRAM devices have up to five programmable configuration settings that determine the row cycle times, read latency, and write latency of the interface at a given frequency of operation.

RLDRAM II also offers programmable impedance output buffers and on-die termination. The programmable impedance output buffers are for impedance matching and are guaranteed to produce 25- to 60-ohm output impedance. The on-die termination is dynamically switched on during read operations and switched off during write operations. Perform an IBIS simulation to observe the effects of this dynamic termination on your system. IBIS simulation can also show the effects of different drive strengths, termination resistors, and capacitive loads on your system.

RLDRAM II devices use either the 1.5-V HSTL or 1.8-V HSTL I/O standard. You can use either I/O standard to interface with Altera FPGAs.

-  For more information, refer to the RLDRAM II datasheets.
-  For more information about parameterizing the RLDRAM II IP, refer to the *Implementing and Parameterizing Memory IP* chapter.

## Memory Selection

One of the first considerations in choosing a high-speed memory is data bandwidth. Based on the system requirements, an approximate data rate to the external memory should be determined. You must also consider other memory attributes, including how much memory is required (density), how much latency can be tolerated, what is the power budget, and whether the system is cost sensitive.

Table 1-3 lists the memory bandwidth, the features and target markets of each technology.

**Table 1-3. Memory Selection Overview (Part 1 of 2)**

Parameter	DDR3 SDRAM	DDR2 SDRAM	DDR SDRAM	RLDRAM II	QDR II/+ SRAM
Bandwidth for 32 bits (Gbps) <sup>(1)</sup>	34.1	25.6	12.8	25.6	44.8
Bandwidth at % Efficiency (Gbps) <sup>(2)</sup>	23.9	17.9	9	17.9	38.1
Performance / Clock frequency	400–1,066 MHz	200–533 MHz	100–200 MHz	200–533 MHz	154–350 MHz
Altera-supported data rate	Up to 2,133 Mbps	Up to 1,066 Mbps	Up to 400 Mbps	Up to 2132 Mbps	Up to 1400 Mbps
Density	512 Mbytes–8 Gbytes, 32 Mbytes – 8 Gbytes (DIMM)	256 Mbytes–1 Gbytes, 32 Mbytes – 4 Gbytes (DIMM)	128 Mbytes–1 Gbytes, 32 Mbytes – 2 Gbytes (DIMM)	288 Mbytes, 576 Mbytes	8–72 Mbytes
I/O standard	SSTL-15 Class I, II	SSTL-18 Class I, II	SSTL-2 Class I, II	HSTL-1.8V/1.5V	HSTL-1.8V/1.5V
Data width (bits)	4, 8, 16	4, 8, 16	4, 8, 16, 32	9, 18, 36	8, 9, 18, 36
Burst length	8	4, 8	2, 4, 8	2, 4, 8	2, 4
Number of banks	8	8 (>1 GB), 4	4	8	N/A
Row/column access	Row before column	Row before column	Row before column	Row and column together or multiplexed option	N/A
CAS latency (CL)	5, 6, 7, 8, 9, 10	3, 4, 5	2, 2.5, 3	4, 6, 8	N/A
Posted CAS additive latency (AL)	0, CL-1, CL-2	0, 1, 2, 3, 4	N/A	N/A	N/A
Read latency (RL)	RL = CL + AL	RL = CL + AL	RL = CL	RL = CL/CL + 1	1.5, 2, and 2.5 clock cycles
On-die termination	Yes	Yes	No	Yes	Yes
Data strobe	Differential bidirectional strobe only	Differential or single-ended bidirectional strobe	Single-ended bidirectional strobe	Free-running differential read and write clocks	Free-running read and write clocks
Refresh requirement	Yes	Yes	Yes	Yes	No
Relative cost comparison	Presently lower than DDR2	Less than DDR SDRAM with market acceptance	Low	Higher than DDR SDRAM, less than SRAM	Highest


**Table 1-3. Memory Selection Overview (Part 2 of 2)**

Parameter	DDR3 SDRAM	DDR2 SDRAM	DDR SDRAM	RLDRAM II	QDR II/+ SRAM
Target market	Desktops, servers, storage, LCDs, displays, networking, and communication equipment	Desktops, servers, storage, LCDs, displays, networking, and communication equipment	Desktops, servers, storage, LCDs, displays, networking, and communication equipment	Main memory, cache memory, networking, packet processing, and traffic management	Cache memory, routers, ATM switches, packet memories, lookup, and classification memories

**Notes to Table 1-3:**

- (1) 32-bit data bus operating at the maximum supported frequency in a Stratix® IV FPGA.
- (2) 70% efficiency for DDR memories, which takes into consideration the bus turnaround, refresh, burst length and random access latency and assumes 85% efficiency for QDR memories

Altera supports these memory interfaces, provides various IP for the physical interface and the controller, and offers many reference designs (refer to Altera’s [Memory Solutions Center](#)).

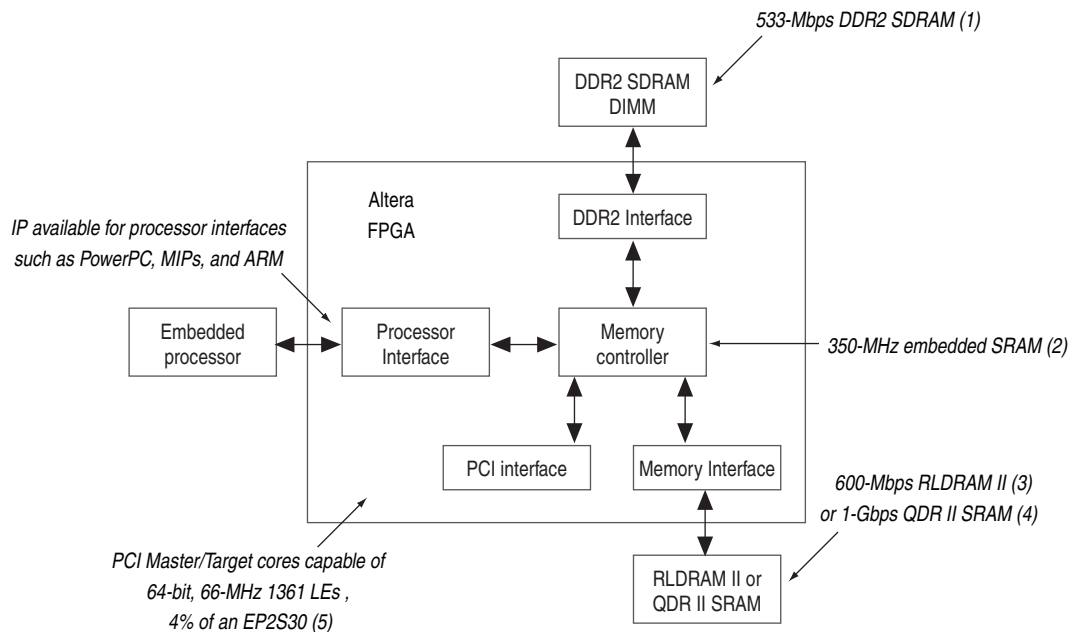
 For Altera support and the maximum performance for the various high-speed memory interfaces, refer to the [External Memory Interface Spec Estimator](#) page on the Altera website.

## High-Speed Memory in Embedded Processor Application Example

In embedded processor applications—any system that uses processors, excluding desktop processors—DDR SDRAM is typically used for main memory due to its very low cost, high density, and low power. Next-generation processors invest a large amount of die area to on-chip cache memory to prevent the execution pipelines from sitting idle. Unfortunately, these on-chip caches are limited in size, as a balance of performance, cost, and power must be taken into consideration. In many systems, external memories are used to add another level of cache. In high-performance systems, three levels of cache memory is common: level one (8 Kbytes is common) and level two (512 Kbytes) on chip, and level three off chip (2 Mbytes).

High-end servers, routers, and even video game systems are examples of high-performance embedded products that require memory architectures that are both high speed and low latency. Advanced memory controllers are required to manage transactions between embedded processors and their memories. Altera Arria® series and Stratix series FPGAs optimally implement advanced memory controllers by utilizing their built-in DQS (strobe) phase shift circuitry. Figure 1-1 highlights some of the features available in an Altera FPGA in an embedded application, where DDR2 SDRAM is used as the main memory and QDR II SRAM or RLD RAM II is an external cache level.

**Figure 1-1. Memory Controller Example Using FPGA**



**Notes to Figure 1-1:**

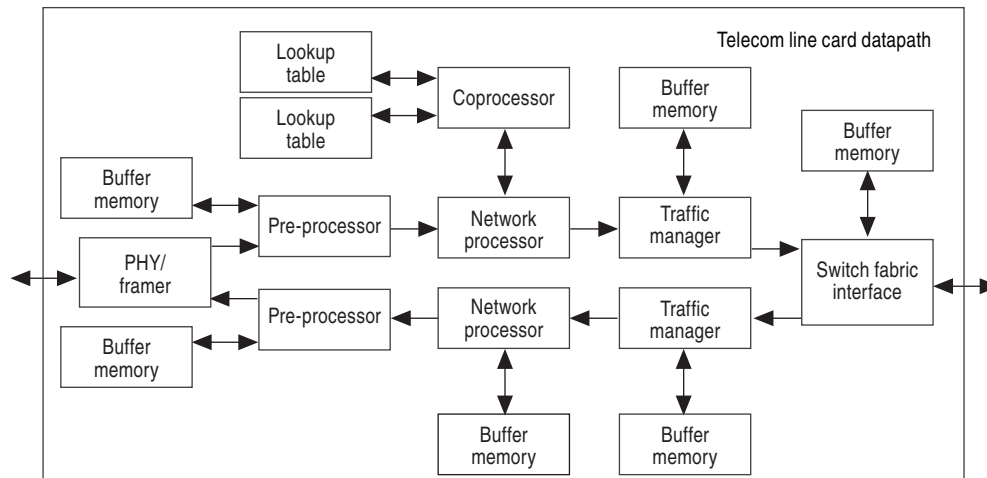
- (1) 533-Mbps DDR2 SDRAM operation using dedicated DQS circuitry, post-amble circuitry, automatic phase shifting, and six registers in the I/O element: 790 LEs, 3% of an EP2S30, and four clock buffers (for a 72-bit interface).
- (2) High-speed memory interfaces such as QDR II SRAM require at least four clock buffers to handle all the different clock phases and data directions.
- (3) 600-Mbps RLD RAM II operation: 740 logic elements (LEs), 3% of an EP2S30, and four clock buffers (for a 36-bit wide interface).
- (4) Embedded SRAM with features such as true-dual port and 350-MHz operation allows complex “store and forward” memory controller architectures.
- (5) The Quartus® II software reports the number of adaptive look-up tables (ALUTs) that the design uses in the FPGA. The LE count is based on this number of ALUTs.

One of the target markets of RLD RAM II and QDR/QDR II SRAM is external cache memory. RLD RAM II has a read latency close to SSRAM, but with the density of SDRAM. A sixteen times increase in external cache density is achievable with one RLD RAM II versus that of SSRAM. In contrast, consider QDR and QDR II SRAM for systems that require high bandwidth and minimal latency. Architecturally, the dual-port nature of QDR and QDR II SRAM allows cache controllers to handle read data and instruction fetches completely independent of writes.

## High-Speed Memory in Telecom Application Example

Because telecommunication network architectures are becoming more complex, high-end network systems are running multiple 10-Gbps line cards that connect to multi-shelf switch fabrics scaling to Terabits per second. Figure 1-2 shows an example of a typical system line interface card. These line cards offer interfaces ranging from a single-port OC-192 to multi-port Gigabit Ethernet, and consist of a number of devices, including a PHY/framer, network processors, traffic managers, fabric interface devices, and high-speed memories.

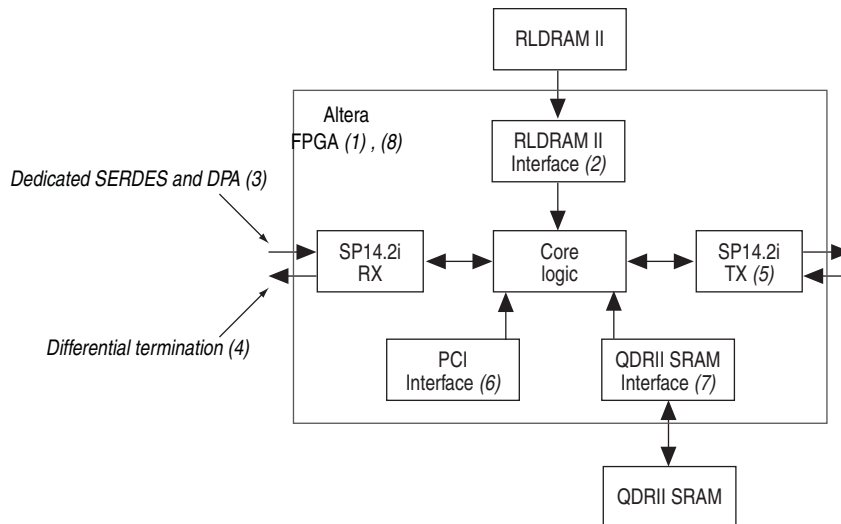
Figure 1-2. Typical Telecom System Line Interface Card



As packets traverse from the PHY/framer device to the switch fabric interface, they are buffered into memories, while the data path devices process headers (determining the destination, classifying packets, and storing statistics for billing) and control the flow of packets into the network to avoid congestion. Typically DDR/DDR2/DDR3 SDRAM and RLDRAM II are used for large buffer memories off network processors, traffic managers, and fabric interfaces, while QDR and QDR II SRAMs are used for look-up tables (LUTs) off preprocessors and coprocessors.

In many designs, FPGAs connect devices together for interoperability and coprocessing, implement features that are not supported by ASIC devices, or implement a device function entirely. Altera Stratix series FPGAs implement traffic management, packet processing, switch fabric interfaces, and coprocessor functions, using features such as 1-Gbps LVDS I/O, high-speed memory interface support, multi-gigabit transceivers, and IP cores. [Figure 1-3](#) highlights some of these features in a packet buffering application where RLD RAM II is used for packet buffer memory and QDR II SRAM is used for control memory.

**Figure 1-3. FPGA Example in Packet Buffering Application**



**Notes to [Figure 1-3](#):**

- (1) As an example, 85% of the LEs still available in an EP2S90.
- (2) 600-Mbps RLD RAM II operation: 740 LEs, 1% of an EP2S90, and four clock buffers (for a 36-bit wide interface).
- (3) Dedicated hardware SERDES and DPA circuitry allows clean and reliable implementation of 1-Gbps LVDS.
- (4) Differential termination is built in Stratix FPGAs, simplifying board layout and improving signal quality.
- (5) SPI 4.2i core capable of 1 Gbps: 5178 LEs per Rx, 6087 LEs per Tx, 12% of an EP2S90, and four clock buffers (for both directions using individual buffer mode, 32-bit data path, and 10 logical ports).
- (6) PCI cores capable of 64-bit 66-MHz 656 LEs, 1% of an EP2S90 for a 32-bit target
- (7) 1-Gbps QDR II SRAM operation: 100 LEs, 0.1% of an EP2S90, and four clock buffers (for an 18-bit interface).
- (8) Note that the Quartus II software reports the number of ALUTs that the design uses in Stratix II devices. The LE count is based on this number of ALUTs.

SDRAM is usually the best choice for buffering at high data rates due to the large amounts of memory required. Some system designers take a hybrid approach to the memory architecture, using SRAM to store the packet headers and DRAM to store the payload. The depth of the memories depends on the architecture and throughput of the system.

The buffer memory for the packet buffering application of an OC-192 line card (approximately 10 Gbps) must be able to sustain a minimum of one write and one read operation, which requires a memory bandwidth of 20 Gbps to operate at full line rate (more bandwidth is required if the headers are modified). The bandwidth requirement for memory is a key factor in memory selection (see [Table 1-3](#)). As an

example, a simple first-order calculation using RLDRAM II as buffer memory requires a bus width of 48 bits to sustain 20 Gbps ( $300 \text{ MHz} \times 2 \text{ DDR} \times 0.70 \text{ efficiency} \times 48 \text{ bits} = 20.1 \text{ Gbps}$ ), which needs two RLDRAM II parts (one  $\times 18$  and one  $\times 36$ ). RLDRAM II also inherently includes the additional memory bits used for parity or error correction code (ECC).

QDR and QDR II SRAM have bandwidth and low random access latency advantages that make them useful for control memory in queue management and traffic management applications. Another typical implementation for this memory is billing and packet statistics, where each packet requires counters to be read from memory, incremented, and then rewritten to memory. The high bandwidth, low latency, and optimal one-to-one read/write ratio make QDR SRAM ideal for this feature.

## Document Revision History

Table 1-4 lists the revision history for this document.

**Table 1-4. Document Revision History**

Date	Version	Changes
November 2011	4.0	Moved and reorganized “Selecting your Memory” section to Volume 2:Design Guidelines.
June 2011	3.0	Added “Selecting Memory IP” chapter from Volume 2 Section I.
December 2010	2.1	<ul style="list-style-type: none"> <li>■ Moved protocol-specific feature information to the memory interface user guides in Volume 3.</li> <li>■ Updated maximum clock rate information for 10.1.</li> </ul>
July 2010	2.0	<ul style="list-style-type: none"> <li>■ Added specifications for DDR2 and DDR3 SDRAM Controllers with UniPHY.</li> <li>■ Streamlined the specification tables.</li> <li>■ Added reference to web-based Specification Estimator Tool.</li> </ul>
January 2010	1.1	Updated DDR, DDR2, and DDR3 specifications.
November 2009	1.0	First published.