

40-nm FPGAs: Architecture and Performance Comparison

FPGA users are constantly looking for ways to differentiate their products in the market place and in doing so they define new systems with new requirements. The new requirements usually are increased functionality, higher processing performance, low power consumption, customizable capabilities, and reprogrammability. To achieve these requirements, the proven methodology is to double the FPGA density when moving to smaller process geometries. However, moving to smaller geometries like 40 nm does not come free. While dynamic power reduces at the 40-nm process node, static power consumption increases through unwanted leakage in transistors (shorter gate lengths). Through processing techniques and architectural innovations like Programmable Power Technology, Altera's 40-nm Stratix® IV FPGAs not only provide increased functionality and higher processing performance but also lower power consumption.

This white paper provides benchmarking data showing Altera® Stratix IV FPGAs are 35 percent faster than Virtex-5 FPGAs, and detailed architectural analysis demonstrating that the Stratix IV FPGA's basic logic element packs 1.8X more logic than that of Virtex-5 FPGAs, all in addition to the provided 8-Gbps transceivers.

Introduction

Stratix IV FPGAs offer up to 680,000 logic elements (LEs), over 22 Mbits of internal RAM, and over 1,300 18x18 multipliers for increased functionality and system integration. The core fabric is built from innovative logic units, known as adaptive logic modules (ALMs), which are routed with the MultiTrack interconnect architecture to provide the maximum connectivity with fewer hops. Coupling the power-saving techniques with efficient core architecture, Stratix IV FPGAs provide the highest system bandwidth, connectivity, and power efficiency.



For detailed information on the process and architectural innovations that enable Stratix IV FPGAs to lower power consumption, see Altera's [40-nm FPGA Power Management and Advantage](#) white paper.

Benchmark Methodology

Benchmarking FPGA performance is a complex task, and a poor benchmarking process may result in inconclusive or incorrect results. Altera has invested significantly in developing a rigorous and scientific benchmarking methodology endorsed by third-party industry experts. Altera's benchmarking methodology is used to compare FPGA performance between families from a single FPGA vendor and with those of competitive solutions. This ensures a consistent benchmarking environment when testing Altera devices and when comparing them to competitor devices.



For detailed information on the benchmarking methodology, refer to Altera's [FPGA Performance Benchmarking Methodology](#) white paper.

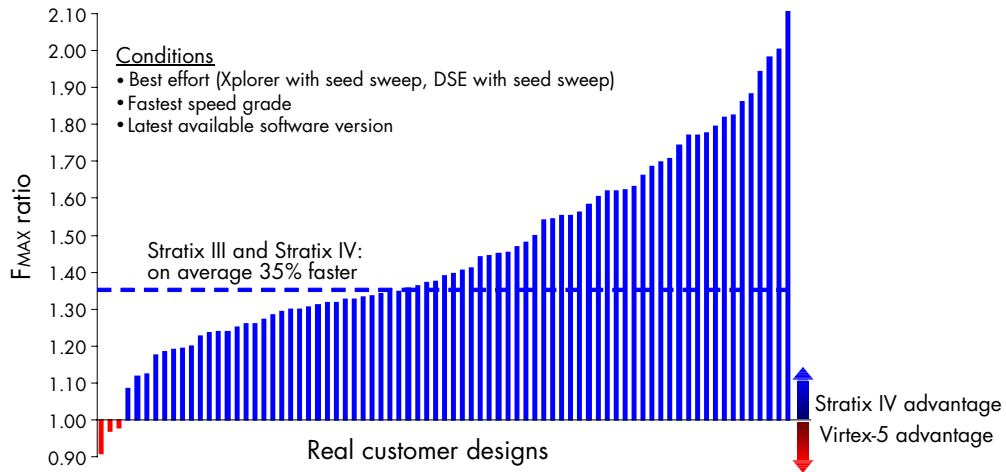
Stratix IV Performance Advantage

Stratix IV performance advantages include efficient core architecture and easier timing closure for system-level performance (with margin).

Core Architecture

The high performance of Stratix IV devices, proven to be the industry's fastest FPGAs, is achieved effortlessly through the unique core architecture fully integrated in Altera's Quartus® II development software. Designs are synthesized with successful placement and routing for faster timing closure, thus enabling designers to increase productivity while meeting their performance goals. The fast speed grade of Stratix IV FPGAs is, on average, 35 percent faster than the fast speed grade of competing Virtex-5 FPGAs (Figure 1). This average performance advantage is based on a set of real customer designs using the latest versions of publicly available design software.

Figure 1. Stratix IV FPGAs (Fast Speed Grade) vs. Virtex-5 (Fast Speed Grade) Benchmarking Results



As a direct consequence of this performance advantage, the slowest speed grade of Stratix IV FPGAs is faster than the fastest speed grade of Virtex-5. Figure 2 and Table 1 show the speed grade comparison between the two FPGA families. In comparison to Virtex-5, designers can meet their performance requirements or easily close timing in slower speed grades of Stratix IV FPGAs while lowering their costs and total power consumption.

Figure 2. Performance Comparison by Speed Grade

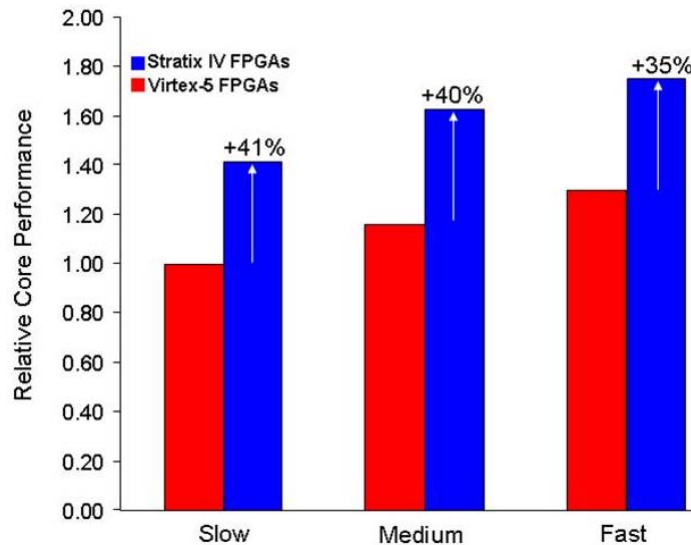


Table 1. Performance Comparison by Speed Grade

Stratix IV Speed Grade	Stratix IV FPGA Performance Advantage Over Virtex-5
Fast	35% faster than nearest competitor's fast speed grade
Medium	40% faster than nearest competitor's mid speed grade
Slow	41% faster than nearest competitor's slow speed grade
Slow	9% faster than nearest competitor's fast speed grade

High System-Level Performance

High system-level performance requires that, in addition to the high core performance, the on-chip RAM, digital signal processing (DSP) blocks, and I/Os must be optimized. Stratix IV DSP blocks are fully optimized for the

highest performance and are able to support clock rates of 550 MHz (Table 2). The TriMatrix on-chip memory offers three flavors of memory structures: 640-bit memory logic array blocks (MLABs), 9-Kbit blocks (M9K), and 144-Kbit blocks (M144K). These blocks are optimized for maximum efficiency, and the MLABs can be placed anywhere in the device, making them extremely flexible and able to support clocks rates of 600 MHz. With the TriMatrix on-chip memory and DSP blocks, Stratix IV devices are ideal for video and image processing, high-speed digital communications, and other high-performance DSP applications.

Table 2. Stratix IV FPGA Block Performance

Block	Stratix IV FPGAs (40 nm)	Virtex-5 FPGAs (65 nm)
Maximum internal clock speed	600 MHz	550 MHz
On-chip RAM	600 MHz	550 MHz
DSP block	550 MHz	550 MHz

To complement the high-performance core, Stratix IV FPGAs are designed with high-performance I/Os, which allow high-bandwidth interfaces to external devices. For example, external memory devices interface with Stratix IV I/Os through dedicated memory interface circuitry that caters to high-performance memory interfaces like DDR3 and QDR II+ (summarized in Table 3). Stratix IV FPGAs are able to interface to any DDR3 SDRAM DIMM and achieve 533 MHz (1067 Mbps) by incorporating the leveling feature built directly into the FPGA I/O structure. This leveling feature first was built in the I/Os of Stratix III FPGAs and is now leveraged by Stratix IV FPGAs.

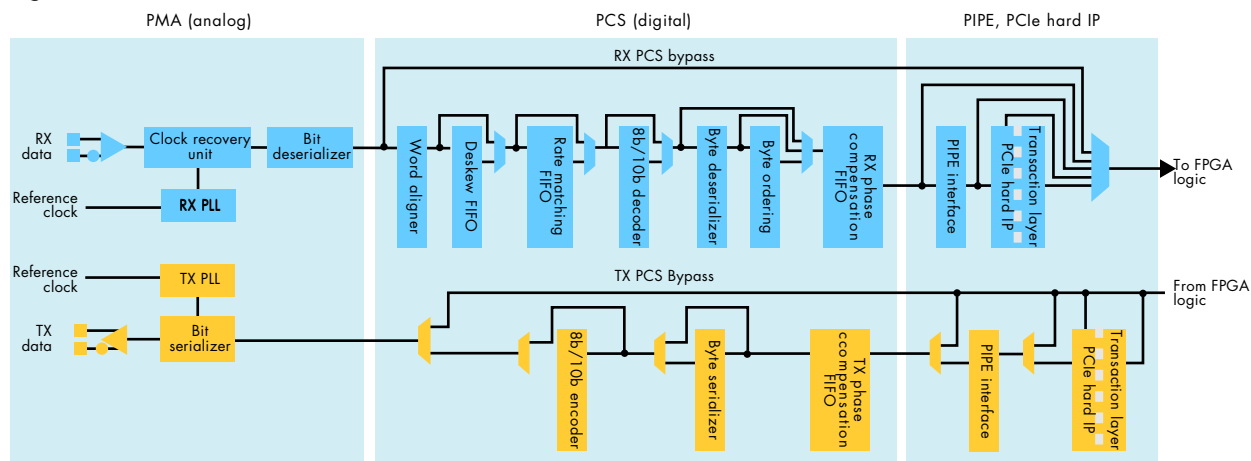
For detailed architecture information and how Stratix III and Stratix IV FPGAs achieve high DDR3 data rates, refer to Altera’s *Utilizing Leveling Techniques in DDR3 SDRAM Memories Interfaces* white paper.

Table 3. Stratix IV FPGA I/O Performance

Interconnect	Stratix IV FPGAs (40 nm)	Virtex-5 FPGAs (65 nm)
DDR3	533 MHz	No DIMM support
DDR2	400 MHz	333 MHz
QDR II +	400 MHz	Not supported
QDR II	350 MHz	300 MHz
RLDRAM II	400 MHz	300 MHz
LVDS	1.6 Gbps	1.25 Gbps

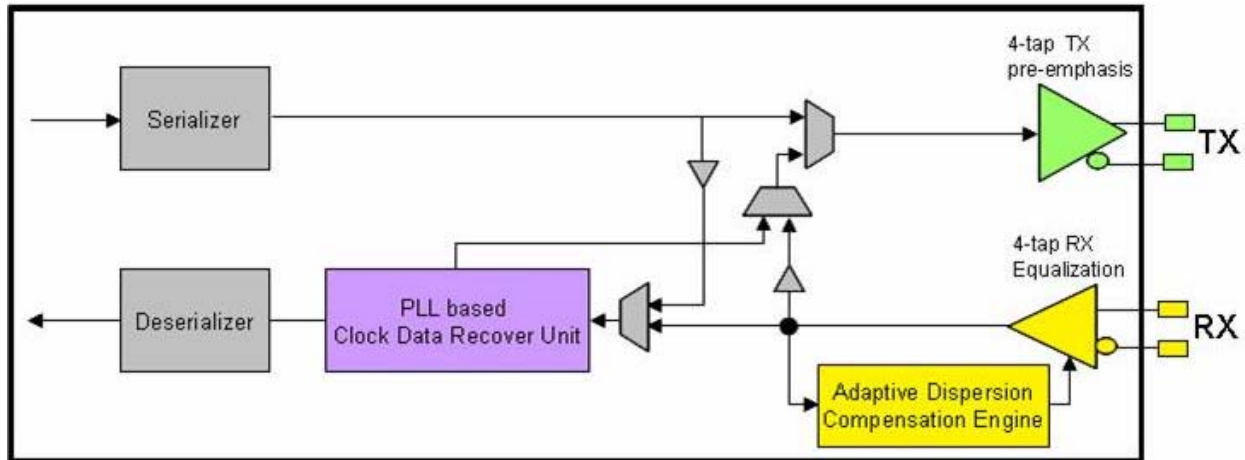
The Stratix IV GX transceiver block (Figure 3), which includes both a physical coding sublayer (PCS) and a physical media attachment (PMA), implements standard and proprietary protocols operating at data rates up to 8.5 Gbps.

Figure 3. Stratix IV GX Transceiver Block



The PMA block is an embedded macro dedicated to receiving and transmitting off-chip high-speed serial data streams. The PMA channel, shown in Figure 4, consists of full duplex paths (transmit (Tx)) and receive (Rx)) with I/O buffers, programmable output voltage, pre-emphasis and equalization, clock data recovery (CDR) and serializer/deserializer blocks.

Figure 4. Stratix IV GX PMA Channel



The transceivers in Stratix IV GX FPGAs make use of advanced power supply regulation and filtering techniques to reduce transmitter jitter and improve receiver jitter tolerance. This allows the Stratix IV GX transceivers to exhibit superior bit-error rate performance when used in real system links. On-chip voltage regulators for both Tx and Rx phase-locked loops (PLLs), careful isolation of sensitive analog circuitry, and extensive use of on-die and on-package decoupling capacitors all contribute to a robust power distribution scheme for the transceivers while delivering exceptionally clean power to the analog circuits. The eye diagram in Figure 5 shows the operation of the transmitter at 6.375 Gbps. The reduced output jitter and vertical noise displayed in the eye diagram are the result of the architectural innovations in the transceiver power-distribution network, as shown in the comparison in Table 4.

Figure 5. 40-nm Stratix IV Test Chip at 6.375 Gbps

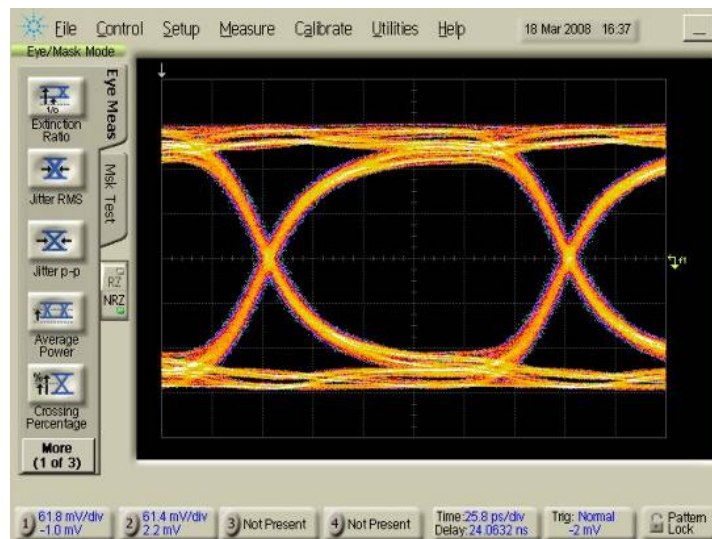


Table 4. Jitter on 40-nm Stratix IV GX Test Chip at 6.375 Gbps

Jitter Component	40-nm Stratix IV GX Transceiver Test Chip (RJ = 1.4 ps rms)	
BER	1E-12	1E-15
Random jitter at BER	19.6	22.3
Deterministic jitter	11.8	
TJ = $\sigma \times RJ$	31.4 ps	34.1 ps
TJ(UI)	0.20 UI	0.22 UI
CEI 6G spec	0.30 UI	0.30 UI

Stratix IV GX transceivers have superior signal integrity with excellent jitter performance combined with the lowest power per channel for both backplane and chip-to-chip applications. Table 5 shows a comparison of data rate and power per channel for the transceivers of Stratix IV GX FPGAs and the nearest competing device.

Table 5. Transceiver Power Per Channel (PMA Only)

Data Rate	Stratix IV FPGAs (40 nm)	Virtex-5 FPGAs (65 nm)
3.2 Gbps	100 mW	~100 mW
6.5 Gbps	135 mW	~200 mW
8.5 Gbps	165 mW	NA (1)

Note:

(1) Virtex-5 does not offer 8.5 Gbps.

Stratix IV FPGAs provide a range of complete PCI-SIG-compliant FPGA solutions for a variety of x1, x4, and x8 PCI Express Gen1 and Gen2 applications. New to Stratix IV FPGAs is the provision of up to four hard intellectual property (IP) blocks that embed a complete PCI Express protocol stack (endpoint and root port). This includes the transceiver blocks, PHY MAC, data link layer, and transaction layer, as shown in Figure 6. Table 6 summarizes the bandwidth and power of the hard IP block.

Figure 6. High-Level Diagram of the PCI Express Hard IP Block

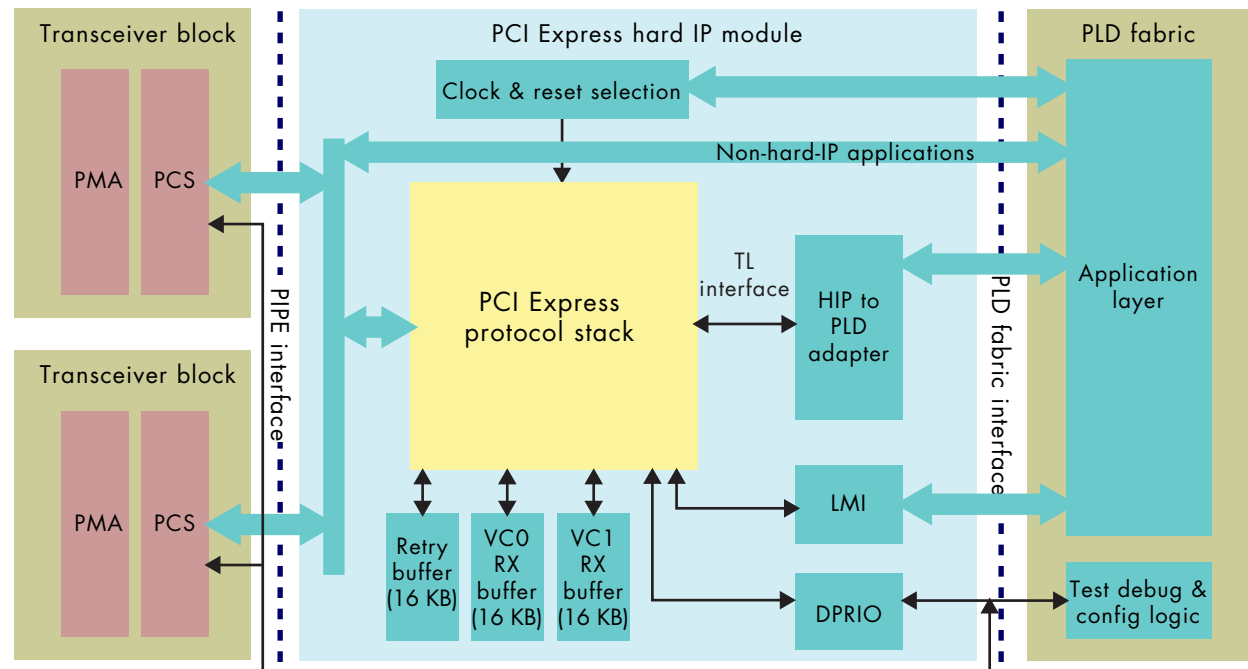


Table 6. Stratix IV GX PCI Express Gen2

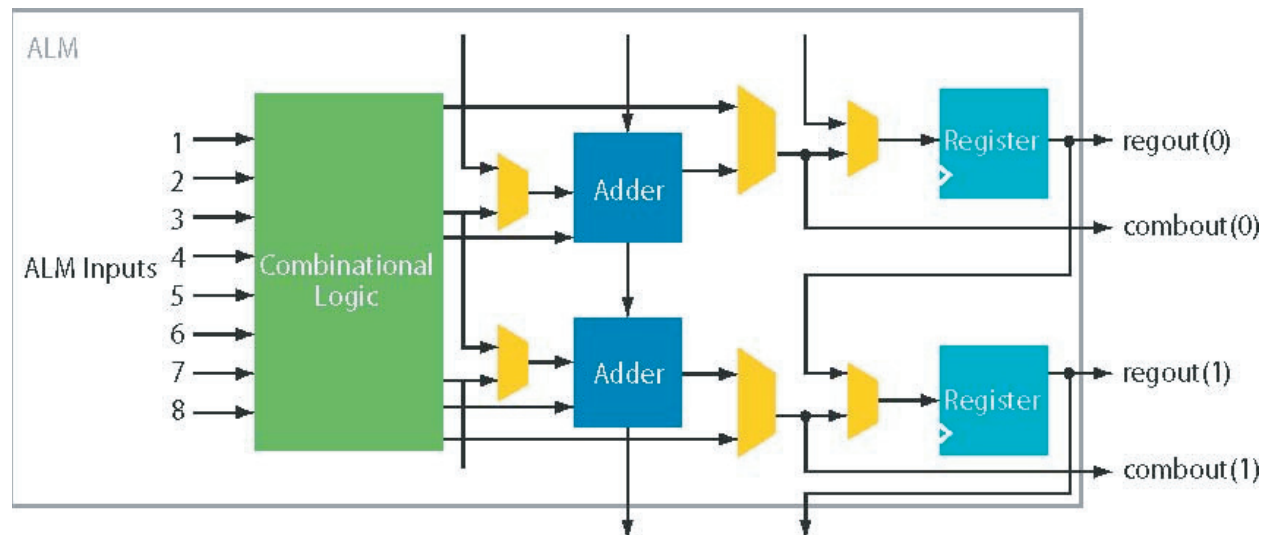
Feature	x8	x4	x2
Bandwidth	40 Gbps	20 Gbps	10 Gbps
Dynamic power	600 mW	440 mW	350 mW

Using these advanced features and customizable IP, designers integrate a wide range of applications quickly and easily into complex system designs without compromising the Stratix IV GX FPGAs' system performance.

Architecture Advantage

The key to high core performance in Stratix IV FPGAs is the area-efficient ALM. It consists of combinational logic, two registers, and two adders, as shown in Figure 7. The combinational portion has eight inputs and includes a look-up table (LUT) that can be divided between two adaptive LUTs (ALUTs) using Altera's patented LUT technology. An entire ALM is needed to implement an arbitrary 6-input function, but one ALM can implement various combinations of the two functions because there are eight inputs in a combinational logic block.

Figure 7. ALM Block Diagram



In addition to implementing a full 6-input LUT, the ALM can, for example, implement two independent 4-input functions or a 5-input and a 3-input function with independent inputs. (Table 7 shows a summary of combinational logic configurations supported in an ALM.) Because two registers and two adders are available, the ALM has the flexibility to implement 2.5 LEs of a classic 4-input LUT (4-LUT) architecture, consisting of a 4-LUT, carry logic, and a register.


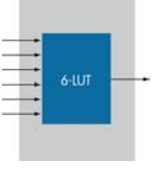
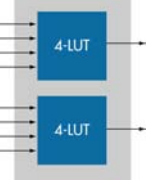
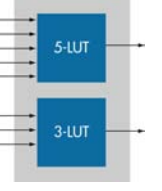
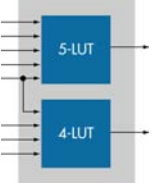
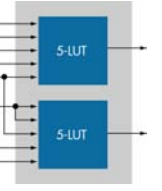
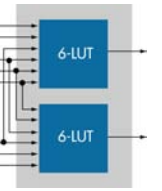
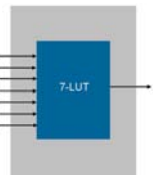
 For a more detailed architectural description, refer to Altera's *Stratix IV Device Handbook*.

Table 7. ALM Flexibility

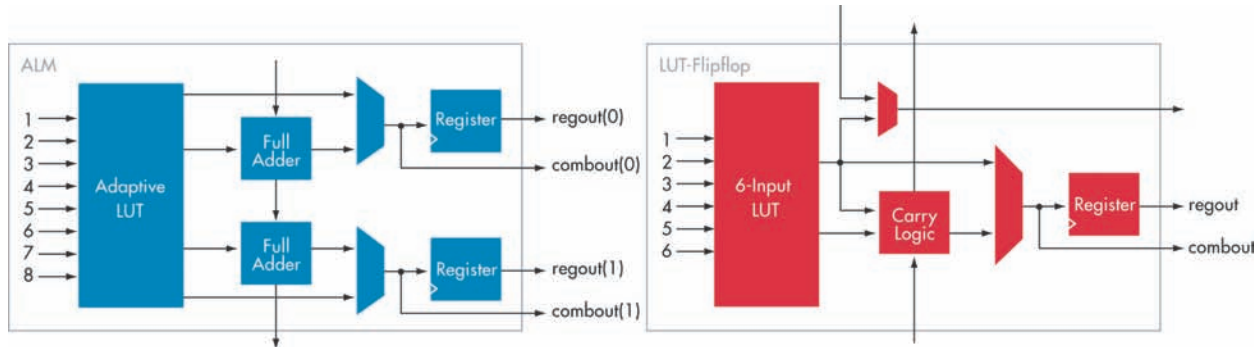
Configuration	Description
	One Stratix IV ALM can input any 6-input function.
	One Stratix IV ALM can be configured to implement two independent 4-input or smaller LUTs. This configuration can be viewed as the "backward-compatibility" mode. Designs optimized for the traditional 4-LUT FPGAs can be migrated to the Stratix family easily.
	One Stratix IV ALM can be configured to implement a 5-LUT and 3-LUT. The inputs to the two LUTs are independent of each other. The 3-LUT can be used to implement any logic function with three or fewer inputs. Therefore, a 5-LUT/2-LUT combination is also available.
	One Stratix IV ALM can be configured to implement a 5-LUT and a 4-LUT, by sharing one input between the two LUTs. The 5-LUT has up to four independent inputs, and the 4-LUT has up to three independent inputs. The sharing of inputs between LUTs is very common in FPGA designs, and Quartus II software automatically seeks logic functions structured in this manner.
	One Stratix IV ALM can be configured to implement two 5-LUTs, and commonly share two inputs. Up to three independent inputs are allowed for each 5-LUT.
	If two 6-input functions have the same logic operation and four shared inputs, the two 6-input functions can be implemented in one Stratix IV ALM. For example, an ALM can implement two 6-input AND gates with four common inputs.
	One Stratix IV ALM in the extended mode can implement a subset of a 7-variable function. The Quartus II software automatically recognizes the applicable 7-input function and fits it into an ALM. (1)

Note:

(1) Refer to the *Stratix IV Device Handbook* for detailed information about the types of 7-input functions that can be implemented in an ALM.

The ALM also contains two registers and two adders as shown in Figure 8. The extra register was added because experiments indicated many customer applications require a higher than 1:1 ratio of registers and LUTs. Two extra adders were included to enhance the arithmetic capability of the ALM, allowing for two 2-bit addition or two 3-bit addition per ALM. Thus, the ALM provides twice as much register and arithmetic capability as a basic 6-LUT, making it a superior building block.

Figure 8. Comparing the Stratix IV ALM (left) and the Virtex-5 LUT-Flipflop Pair (right)



The ALM is significantly more flexible and, as a result, is more area efficient than the Xilinx Virtex-5 LE (also called a LUT-flipflop pair). The LE consists of a basic 6-LUT, carry logic, and a single register. In comparison, the combinational logic portion of the ALM has eight inputs and supports all 6-input functions, as well as many other combinations of smaller functions using its two outputs. The combinational logic portion of the Virtex-5 LE, a basic 6-LUT, also has 64 bits of CRAM and two outputs like the ALM, but only has six inputs and a limited ability to implement more than one logic function. The two outputs—the 6-LUT output and the 5-LUT output—correspond to the lower half of the configuration RAM.

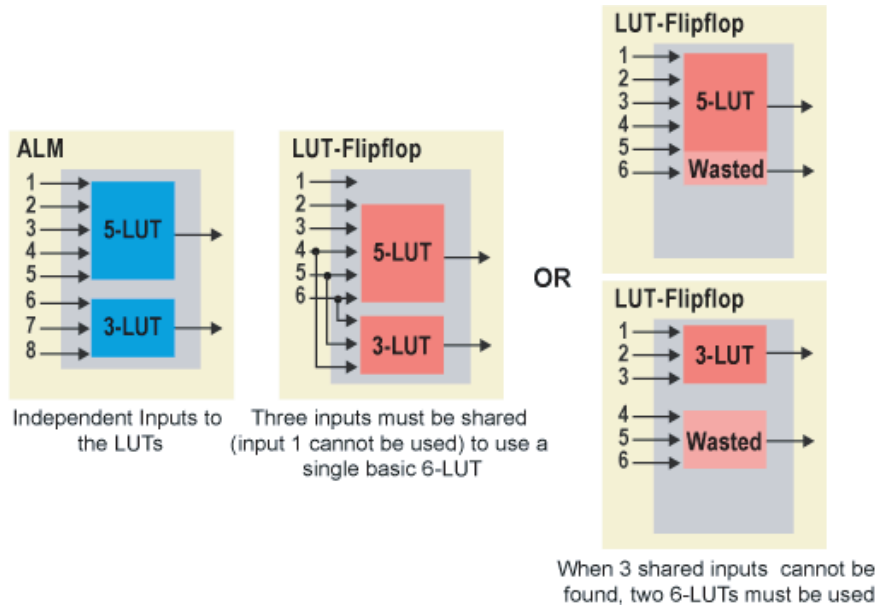
Although the basic 6-LUT has the ability to implement two smaller functions, it usually is used only as a 6-LUT. Because the LUT has only six inputs, the required number of shared inputs places severe restrictions on the types of functions that can be combined. These restrictions make using the basic 6-LUT as two 5-LUTs a rare occurrence. In contrast, the two additional inputs in the Stratix IV ALM allow it to be used as two fully functional 5-LUTs, providing a significant area advantage. Table 8 gives the number of shared inputs required for a few combinations of functions.

Table 8. Stratix IV ALM vs. Virtex-5 LUT Flexibility

Output 1	Output 2	Virtex-5	ALM Shared Inputs (Minimum)
5-LUT	5-LUT	5	2
5-LUT	4-LUT	4	1
5-LUT	3-LUT	3	0
4-LUT	4-LUT	3	0
4-LUT	3-LUT	2	0
3-LUT	3-LUT	1	0

Figure 9 shows an example: the ALM can implement a 5-input and a 3-input function without any shared inputs, while the Virtex-5 LUT requires three shared inputs. It is difficult to find functions that can be packed into a Virtex-5 LUT, resulting in functions with less than six inputs being implemented in 6-LUT resources. By combining the 8-input fracturable LUT and 2:1 register-to-LUT ratio, the ALM can effectively implement two independent functions that may require a flipflop each without the need to share inputs or resources. The result is that, in virtually any design, implementation requires fewer ALMs than Virtex-5 LEs.

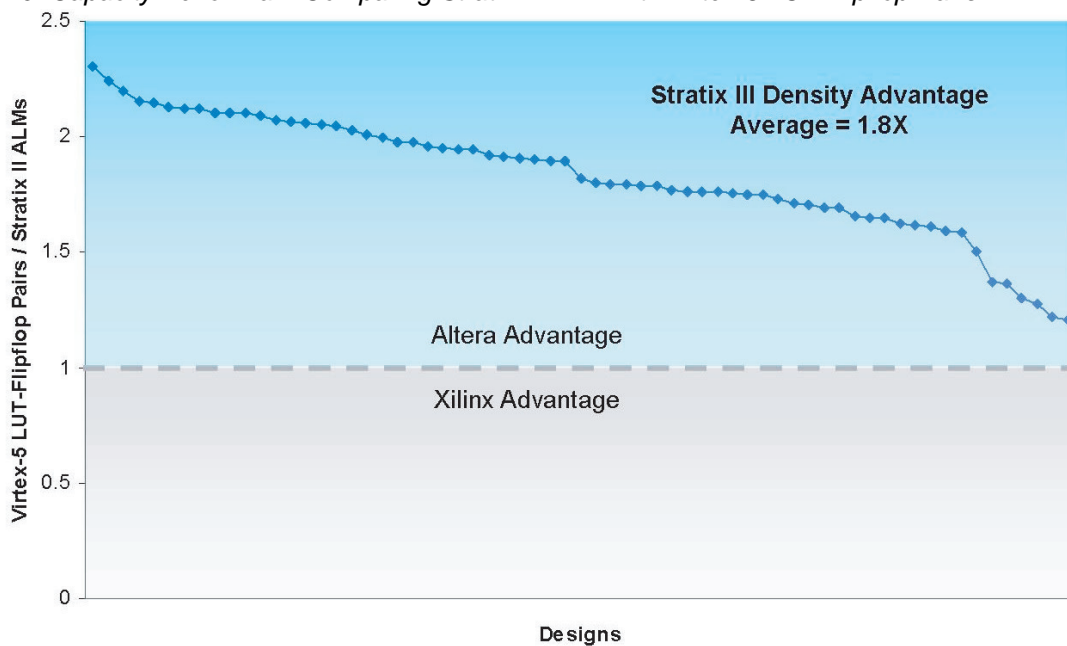
Figure 9. Implementing 5-and 3-Input Functions in Stratix IV ALM and Virtex-5 LUT-Flipflop Pair



Capacity Benchmarking

Capacity benchmarks on over 65 real-world designs (optimized for minimum area) were run to measure the capacity of Stratix IV devices compared to Virtex-5 devices. The capacity benchmarks provide a relative measure of capacity for Altera’s ALM and Xilinx LUT-flipflop pair. Because of efficient fracturability, the Stratix IV ALM has, on average, a 1.8X advantage over the Virtex-5 LUT-flipflop pair, and on certain designs, the advantage can be as high as 2.3X, as shown in Figure 10. The horizontal black line at the “1” mark indicates a point at which the number of LEs for Virtex-5 (LUT-flipflop pairs) and Stratix IV (ALMs) are the same.

Figure 10. Capacity Benchmark Comparing Stratix IV ALM With Virtex-5 LUT-Flipflop Pairs

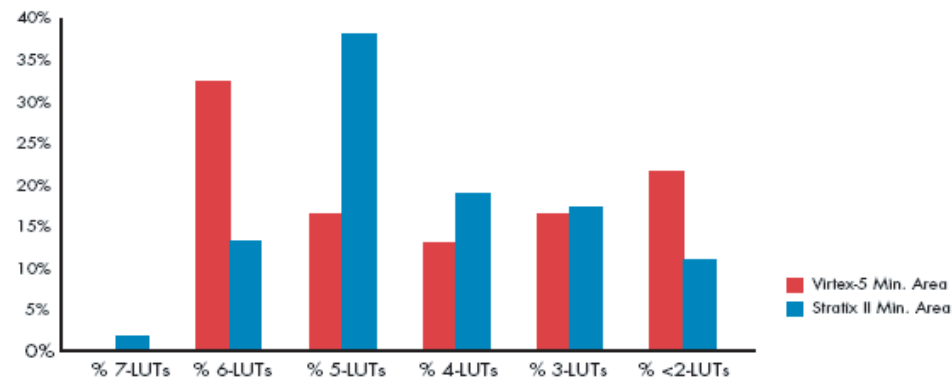


Note:

(1) 10 out of 75 customer designs for Virtex-5 resulted in errors when optimized for area.

To better understand these results, a breakdown of LUT sizes generated by synthesis is shown in [Figure 11](#) illustrates that synthesis generates a much larger percentage of 6-LUTs for Virtex-5 devices than for Stratix IV devices, 32 percent versus 13 percent. The reason for this difference is that, when using a basic 6-LUT such as Virtex-5, it is desirable to use as many inputs as possible because the entire LUT is used in most cases regardless of whether the function requires six inputs or fewer. Since Virtex-5 implements only 6-LUTs efficiently, synthesis attempts to generate as many 6-LUTs as possible. Creating smaller functions does not make sense because it is unlikely that two can be packed given the number of inputs that need to be shared. With a fixed 6-input architecture like Virtex-5, any LUT size smaller than 6 inputs results in inefficient use of the FPGA architecture. Since Xilinx ISE only creates 6-input LUTs 32 percent of the time, the remaining 68 percent of the design uses the underlying architecture inefficiently.

Figure 11. LUT Sizes Generated During Synthesis



Because of LUT fracturability, the synthesis tool alters the distribution of LUT sizes to produce the right mix of large and small functions, ensuring efficient packing and thus resulting in fewer ALMs being used. Specifically, any function of five or fewer inputs uses only half the ALM, making it more important to use the 6-input functions only for speed-critical logic.

The ALM provides flexibility for software optimization. [Figure 12](#) shows the mixture of LUT sizes produced by Quartus II Integrated Synthesis (QIS) when optimizing for three different goals: speed, area, or a balanced approach. The mixture of LUT sizes varies depending on the goal. Optimizing for speed generates the largest number of 6-LUTs, while optimizing for area generates a distribution that packs in the smallest number of ALMs. This flexibility is unique to Altera and is the result of intensive research on the interaction between software and hardware during architecture development to achieve optimal results.


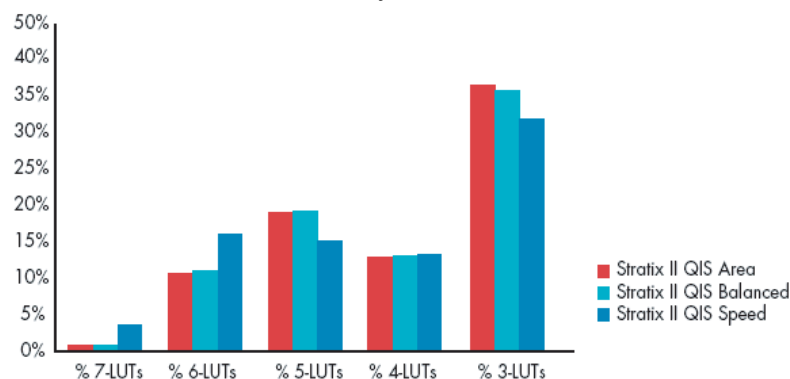
 For an in-depth description of the ALM design and research, read Altera's [Stratix II FPGA Architecture](#) (the foundation of Stratix IV architecture) white paper.

Figure 12. Distribution of Functions Generated by QIS



Based on extensive benchmark analysis performed using real customer designs and a full synthesis, placement, and routing design flow targeted for minimum area, the Stratix IV ALM “is equivalent to” or “can hold as much logic as” 1.8X Virtex-5 LEs (see [Table 9](#)). Another way to look at this is that the ALM is 1.8X more efficient than the LUT-flipflop pair in Virtex-5.

Table 9. Normalized Relative Logic Capacity

Stratix IV ALM	Virtex-5 LUT Flipflop Pair
1.8	1

Taking the 1.8X factor into account, an EP4SE680 device has 2.4X more logic than XC5VLX330 (see [Table 10](#) and [Figure 13](#)).

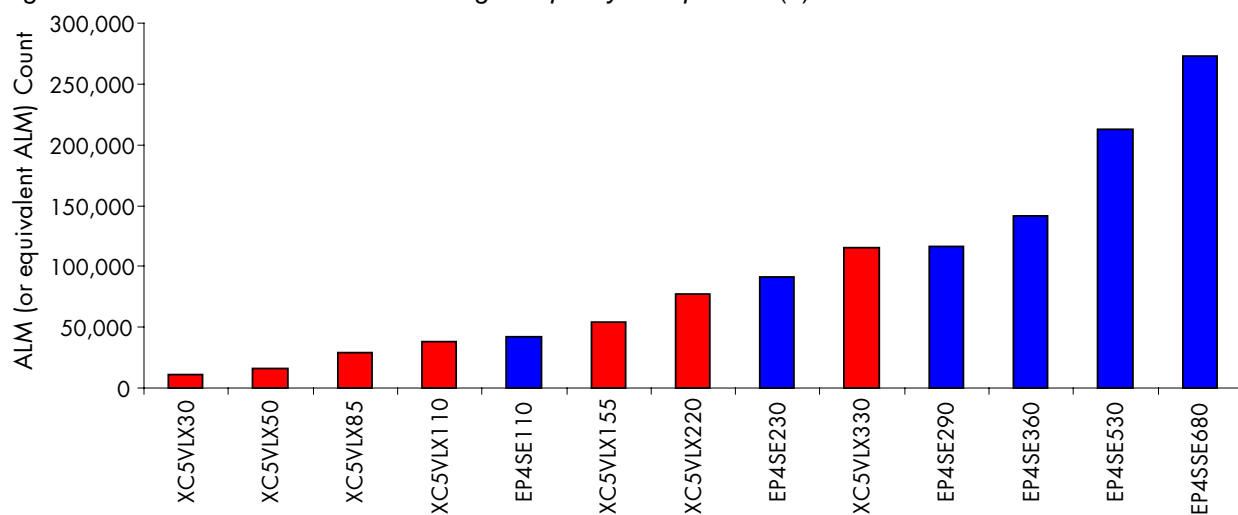
Table 10. Stratix IV and Virtex-5 Equivalent Device Match-Up

Altera Device	ALMs	LEs	Registers	Xilinx Device	Equivalent ALMs (1)	LEs	Registers
				XC5VLX30	10,667	26,667	19,200
				XC5VLX50	16,000	40,000	28,800
				XC5VLX85	28,800	72,000	51,840
				XC5VLX110	38,400	96,000	69,120
EP4SE110	42,240	105,600	84,480				
				XC5VLX155	54,044	135,111	97,280
				XC5VLX220	76,800	192,000	138,240
EP4SE230	91,200	228,000	182,400				
				XC5VLX330	115,200	288,000	207,360
EP4SE290	116,480	291,200	232,960				
EP4SE360	141,440	353,600	282,880				
EP4SE530	212,480	531,200	424,960				
EP4SSE680	272,440	681,100	544,880				

Note:

(1) Converting Virtex-5 6-LUT flipflop pair to equivalent ALM count using 1.8X factor

Figure 13. Stratix IV L vs. Virtex-5 LX Logic Capacity Comparison (1)



Note:

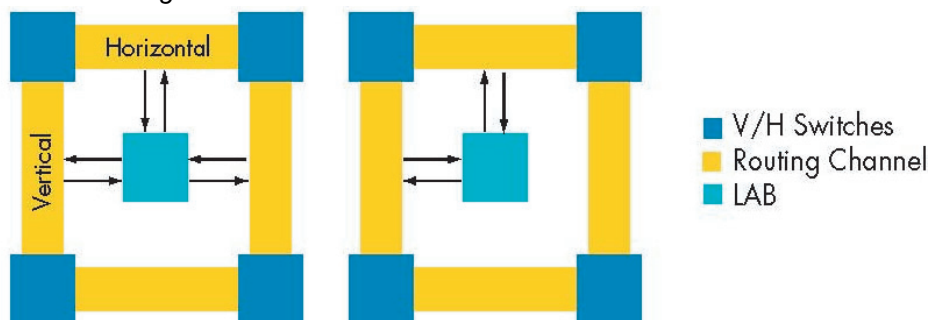
(1) Converting Virtex-5 6-LUT flipflop pair to equivalent ALM count using 1.8X factor

Routing Architecture

In addition to logic block architecture, another key FPGA feature is the routing architecture. The Stratix series of devices introduced the MultiTrack interconnect to maximize connectivity and performance. The routing architecture provides the connectivity between different clusters of LABs, and can be measured by the number of “hops” required to get from one LAB to another. The fewer the number of hops and the more predictable the pattern, the better the performance and the easier the CAD tool optimization.

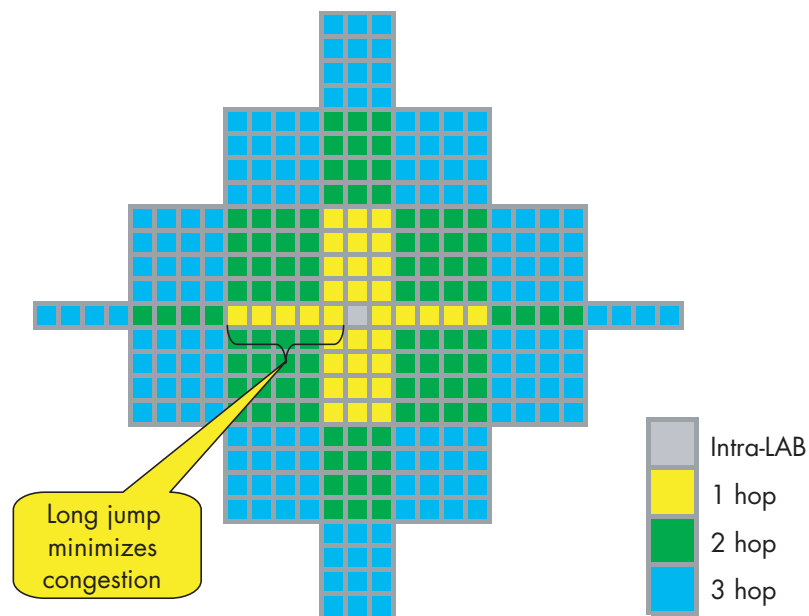
Routing is organized as wires in a number of rows and columns. The Stratix series uses a three-sided routing architecture, shown in Figure 14. This means a LAB drives or listens to all of the wires on one horizontal (H) channel above it and two vertical (V) channels to the left and right sides of it. The channels contain wires of length 4, 8, 16, and 24, and signals can get off at any LAB along the length of the wire.

Figure 14. Number of Routing Architecture Sides



Considering only wires of length four for simplicity, Figure 15 shows the number of hops required to connect to LABs from a given LAB located at the location denoted by the gray box.

Figure 15. Number of Routing Architecture Sides



Virtex-5 FPGA uses a two-sided routing architecture because a configurable logic block (CLB) can connect to all of the wires in a single vertical channel and a single horizontal channel (with connectivity to half of the wires above and to half of the wires below the CLB). In addition, it uses wires that can only connect to CLBs at select points along the

length of the wire. Both of these factors place restrictions on connectivity and placement. With Virtex-5 devices, a CLB can still talk to two channels, but has also included L-shaped (referred to in Xilinx material as diagonal) wires.

Table 11 compares the connectivity of the Stratix IV family with Virtex-5 in terms of the number of LABs/CLBs reachable in a given number of hops. In Stratix IV devices, many more LABs (34) are reached in one hop than CLBs in Virtex-5 devices. If the numbers are scaled by the greater efficiency of the ALM, the results are even more favorable to Stratix IV devices. Because a LAB contains the equivalent of 25 4-LUT-based LEs versus the approximately 11 of Virtex-5 (using the 1.8X factor), if we scale the amount of logic that can be reached within a given number of hops by these factors, the improved routing connectivity in terms of logic capacity is even greater.

Table 11. Stratix IV vs. Virtex-5 Connectivity

Hops	Number of LABs/CLBs Reachable		Number of LEs Reachable		Ratio of Stratix IV LEs to Virtex-5 LEs
	Stratix IV	Virtex-5	Stratix IV (1)	Virtex-5 (2)	
1	34	12	850	132	6.4
2	96	96	2,400	1,056	2.3
3	160	180	4,000	1,980	2.0
Total	290	288	7,250	3,168	2.3

Notes:

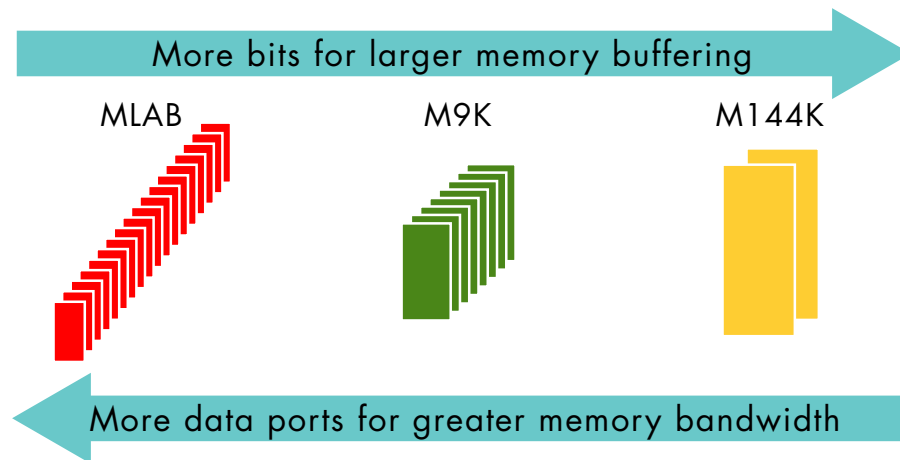
(1) 1 ALM = 2.5 LEs and each LAB = 10 ALMs

(2) 1 CLB = eight 6-LUT-flipflop pairs and converting Virtex-5 6-LUT to equivalent ALM count using 1.8X factor

TriMatrix Memory and DSP Blocks

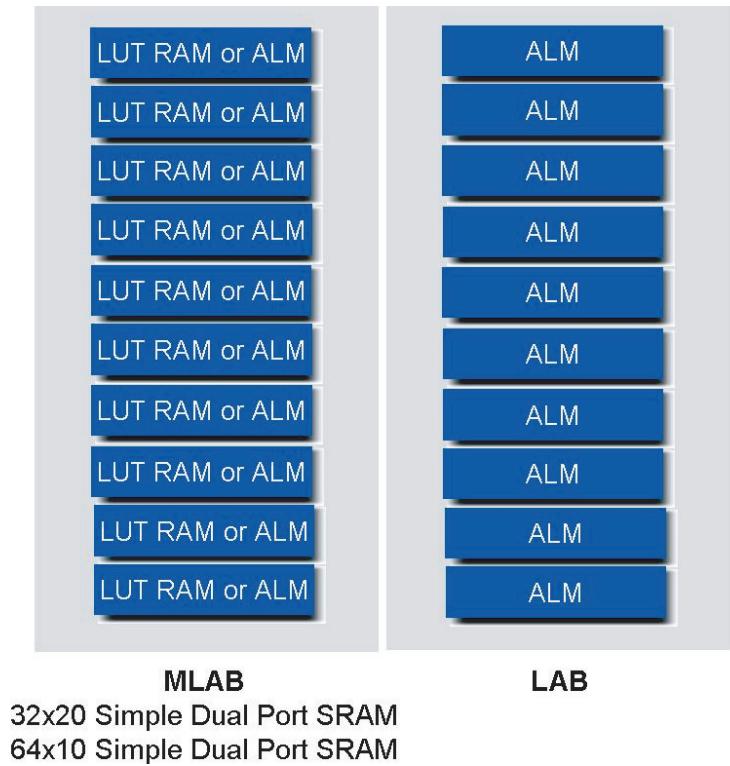
Stratix IV FPGAs offer three “flavors” of on-chip memory, collectively known as TriMatrix memory. By offering three different memory block sizes, Stratix IV FPGAs allow the designer to select the best fit for an application. The three types include MLABs (640 bits), M9K blocks (9K bits), and M144K blocks (144K bits). As shown in Figure 16, TriMatrix memory significantly improves memory utilization and reduces the need for memory cascading.

Figure 16. TriMatrix Memory Structure



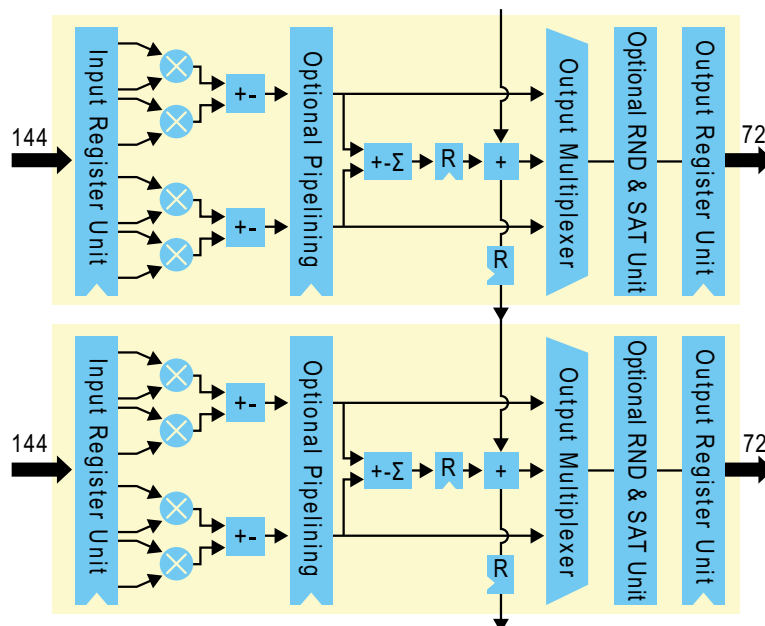
The M9K and M144K blocks are dedicated memory resources. Each MLAB is implemented using one LAB containing 10 flexible ALMs, which in turn are configurable to implement logic functions, register functions, and complex arithmetic functions. Half the LABs in a Stratix IV FPGA can be configured as MLABs. The LAB and MLAB coexist as pairs in Stratix IV FPGAs and the MLAB ALM can be used as a regular LAB ALM or can be configured as a dual-port SRAM, as shown in Figure 17. This LUT-RAM capability supports a maximum of 640 bits, which can be configured as 64x10 or 32x20 simple dual port SRAM blocks, as compared to the CLB of Virtex-5 that only supports a 64x4 configuration. The SRAM blocks are optimized to implement filter delay lines, small FIFO buffers, and shift registers with maximum performance of 600-MHz clock speeds.

Figure 17. Stratix IV MLAB



The Stratix IV DSP block is a high-performance dedicated resource with programmability that delivers optimized processing across many applications. Each block provides eight 18x18 multipliers, as well as registers, adders, subtractors, accumulators, and summation unit functions frequently required in typical DSP algorithms. The DSP block supports variable bit widths with various rounding and saturation modes to meet the exact requirements of the application efficiently. (See Figure 18.)

Figure 18. DSP Block Architecture

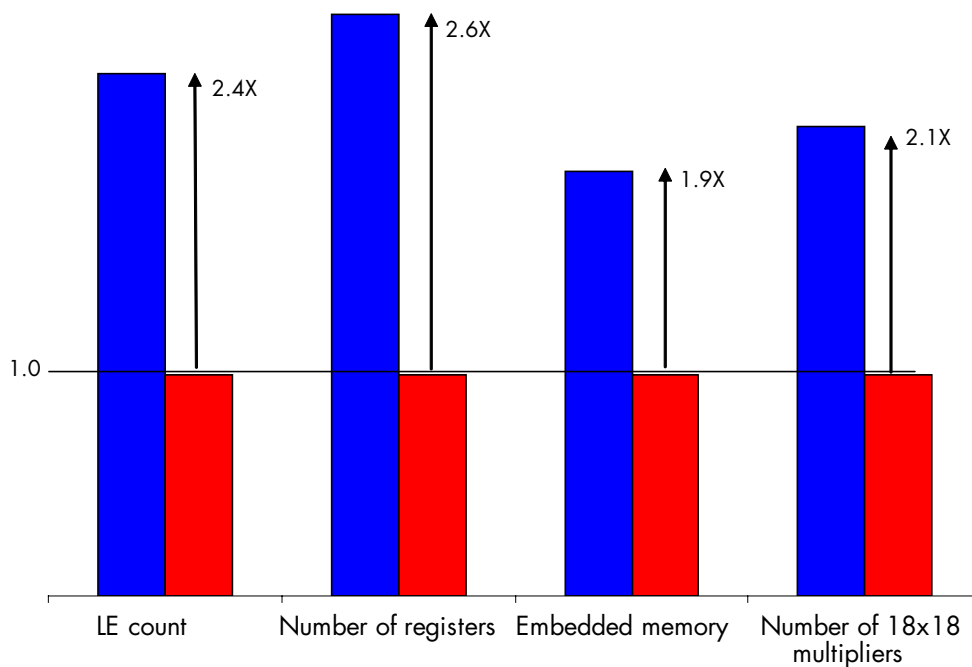



At the 40-nm process node, Stratix IV FPGAs provide unmatched resources for integrating next-generation, high-performance designs. Table 12 and Figure 19 compare the largest device in the Stratix IV FPGA family to that of the Virtex-5 FPGAs.

Table 12. Resource Comparison on Largest Parts Between Stratix IV FPGAs and Virtex-5

FPGA Vendor	Device	LEs	Embedded Memory	18x18 Multipliers	Registers
Altera	EP4SE680	681,100	22,400	1,360	544,880
Xilinx	XC5VLX330	288,000	11,664	640	207,360

Figure 19. Resource Comparison on Largest Parts Between Stratix IV FPGAs and Virtex-5



 For detailed information on the resources available, refer to Altera's [Stratix IV FPGA Handbook](#).

Conclusion

Migrating to smaller geometries requires FPGA vendors to get creative in finding ways to minimize power while delivering higher integration and performance. Through architectural innovations and processing techniques, Stratix IV FPGAs deliver the highest performance at the lowest power. Overall, at the 40-nm process node, Stratix IV FPGAs not only provide unprecedented amount of resources like logic, on-chip memory, and DSP blocks, but also deliver 50 percent lower core power and 35 percent higher performance than Virtex-5 FPGAs. This higher core performance is well matched with optimized I/Os that support DDR3 modules at 533 MHz (1067 Mbps), 1.6-Gbps LVDS channels, and 8-Gbps transceivers.

Further Information

- *FPGA Architecture:*
www.altera.com/literature/wp/wp-01003.pdf
- *Stratix II Performance and Logic Efficiency Analysis:*
www.altera.com/literature/wp/wpstxiiple.pdf
- *Stratix II Logic and Routing Architecture* (foundation of Stratix IV core architecture):
www.altera.com/literature/cp/cp-01005.pdf
- *Improving FPGA Performance and Area Using an Adaptive Logic Module:*
www.altera.com/literature/cp/cp-01004.pdf
- *Fracturable FPGA Logic Elements:*
www.altera.com/literature/cp/cp-01006.pdf
- *Benchmarking Methodology:*
www.altera.com/products/devices/performance/benchmark/per-benchmarkmeth.html
- *Altera at 40 nm: Jitter-, Signal Integrity-, Power-, and Process-Optimized Transceivers:*
www.altera.com/literature/wp/wp-01057-stratix-iv-jitter-signal-integrity-optimized-transceivers.pdf
- *Leveraging the 40-nm Process Node to Deliver the World's Most Advanced Custom Logic Devices:*
www.altera.com/literature/wp/wp-01058-stratix-iv-40nm-process-node-custom-logic-devices.pdf
- *40-nm FPGA Power Management and Advantages:*
www.altera.com/literature/wp/wp-01059-stratix-iv-40nm-power-management.pdf

Acknowledgements

- Seyi Verma, Senior High-End Technical Analysis Staff, High-End FPGA Products, Altera Corporation
- Paul McHardy, Supervising Member of Technical Staff, Software, and Systems Engineering, Toronto Technology Center, Altera Corporation
- Alexander Grbic, Manager, Software and Systems Engineering, Toronto Technology Center, Altera Corporation



101 Innovation Drive
San Jose, CA 95134
www.altera.com

Copyright © 2008 Altera Corporation. All rights reserved. Altera, The Programmable Solutions Company, the stylized Altera logo, specific device designations, and all other words and logos that are identified as trademarks and/or service marks are, unless noted otherwise, the trademarks and service marks of Altera Corporation in the U.S. and other countries. All other product or service names are the property of their respective holders. Altera products are protected under numerous U.S. and foreign patents and pending applications, maskwork rights, and copyrights. Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera Corporation. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.