

This white paper presents how the latest advances in FPGA technology support programmable state machines to support a dynamic RAID architecture using enterprise-grade flash memory. The benefits associated with programmable logic devices (PLDs)—design flexibility, modular IP integration, hardened memory controllers, and high-speed serial interfaces—provide an effective technology option in the design of flash memory array architectures. Programmable state machines provide the best performance in supporting storage subsystem requirements. The use of programmable technology as a means to support emerging and demanding memory-array architectures from prototype to volume production has proven to be very successful for innovative storage companies.

Introduction

As data center administrators address the growing demands placed upon storage resources, performance metrics, such as input-output operations per second (IOPS), must be balanced with data integrity management, system scalability, and serviceability when selecting the appropriate storage solutions.

NAND flash memory technology is becoming widely adopted in the enterprise storage industry as the highest performance and most cost-effective nonvolatile storage medium for frequently used data and applications. To improve performance beyond data caching, application architects are taking a holistic approach to shared memory types by storing application data in flash memory arrays to complement traditional spinning media. This new paradigm requires memory array architectures that support the following:

- Different memory types
- Subsystem I/O interface requirements for high-speed data routing
- RAID to support data integrity as required by the enterprise data center

PLDs are a core part of these memory array subsystems. Complementing inherent design flexibility with the modular integration of embedded processors, hardened memory controllers, and high-speed serial I/O blocks, programmable technology is an effective choice when designing a memory array subsystem for optimum performance. Data can be processed, checked for integrity, and transmitted within application specifications using these PLDs. In addition, systems can be upgraded with new capabilities while preserving the integrity of the application data.

Flash Memory Arrays

Flash memory arrays are the third generation of flash storage technologies. Solid state drives (SSDs) were developed to plug into existing hard disk drive (HDD) systems and servers. Flash PCI Express® (PCIe®) cards have been used to plug into servers for higher performance, but without the normal RAID protection. Flash memory arrays combine the performance of PCIe flash memory with the scalability and reliability of traditional storage systems.

By using memory arrays, enterprises can dramatically reduce storage footprints, the number of CPUs and software licenses needed, and, consequently, the total power, space, and cost required to operate a data center. Flash memory arrays enable both application acceleration and infrastructure consolidation.

Case Study: Violin Memory Architecture Innovation

Violin Memory, Inc. was founded in 2005 with the vision of creating memory arrays that could scale cost-effectively to meet the performance, reliability, and cost objectives of the next-generation data centers that operate 24x7. The primary attributes required by enterprise data centers are:

- **Performance**—HDDs have poor latency and low IOPS. Enterprises want the order-of-magnitude improvement that flash memory can provide with sub-millisecond latencies and very high (>200K per shelf) IOPS that can match their processors.
- **Cost**—Traditionally, it has been cost prohibitive to deploy solid state memory solutions widely. Enterprises want systems that allow them to reduce costs significantly. This requires a blend of cost/GB and cost/I/O.
- **Reliability**—Enterprise data is extremely valuable and cannot be lost. RAID protection is a given, but mirroring of flash is expensive. It is critical that systems can be serviced without downtime.

Violin Memory decided to address these objectives by re-architecting the storage array concept specifically for memory, with a special emphasis on NAND flash based on its low cost/GB and its persistence. The architecture developed has two primary levels:

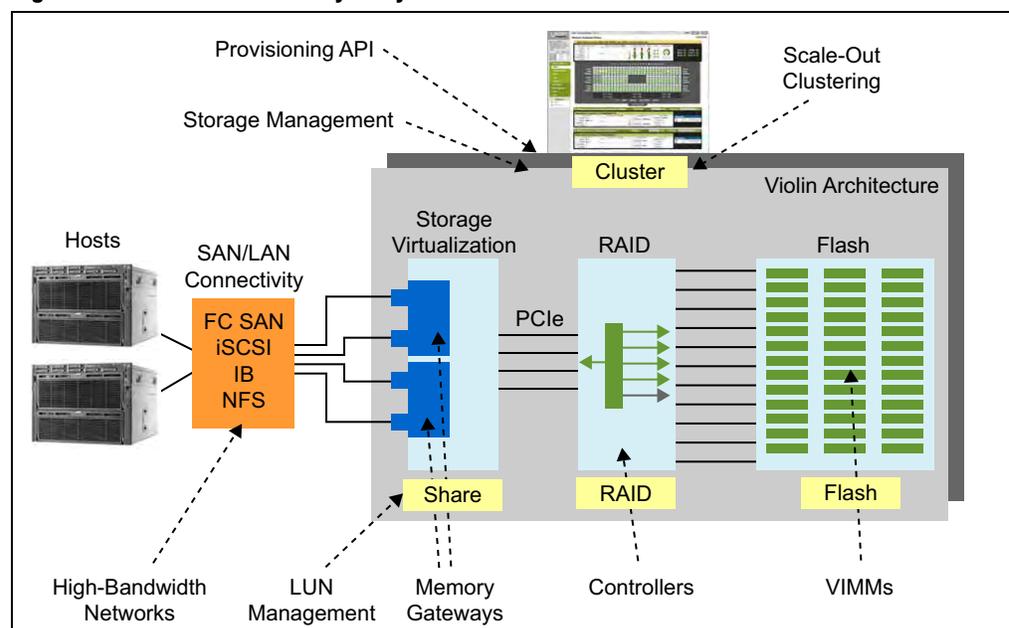
- **Flash control**—NAND flash is complex technology with read, write, and erase operations and many error conditions at bit, block, plane, and chip level. The Violin flash controller implements a complex flash translation layer (FTL) that includes a log-structured data layout with many flash management functions and the need for “garbage collection” to free up flash space for future writes.
- **Flash RAID**—Traditional RAID algorithms for HDDs (RAID-1, RAID-5) are not well suited to the unusual characteristics of flash where erases can take 1 to 10 ms and block reads occur on the same device. The Violin RAID controller implements a 4+1 parity model, which is much more efficient and has lower latency than traditional algorithms. It also has the ability to cope with chip and block failures without module replacement.

Example Architecture for Flash Memory Arrays

Altera's FPGA technology was used to implement the Flash Control (vFLASH) and vRAID functions in the Violin RAID Controller (Figure 1) for the following reasons:

- Significantly lower latency than microprocessor/software technology, achieved by implementing the key algorithms in state machines rather than software
- Greater flexibility than ASIC technology due to the rapid evolution of flash technology and the features needed
- Less capital-intensive approach for entering and growing with a fast growing market
- Straightforward transition to hard-coded, pin-compatible devices through programs such as the Altera® HardCopy® ASICs system development methodology and other competing models.

Figure 1. Violin's Flash Memory Array Architecture



The programmable technology provided by Altera not only supports both the major business objectives, but also the key technical requirements including high-speed interfaces to both memory and PCIe cards.

Memory Array Usage Examples

Memory arrays using programmable technology are gaining significant market traction with exponential market growth. Early adopters include the Web 2.0, military, and intelligence communities, where processing vast amounts of data in real time is a common requirement. More recently, the financial and Fortune 500 markets have embraced this technology. Specific applications for the technology include transaction processing, data warehousing, and virtual storage.

Transaction Processing

The high-volume transaction-processing market is growing as the use of electronic payments and mobile commerce increases. Using memory arrays, Violin has built systems that deliver over 50,000 transactions per second from a single server at a fraction of the traditional cost and space of these systems.

Data Warehousing

The amount of data collected by businesses is growing rapidly, while the demand for real-time and ad hoc queries has also increased. Traditional architectures using spinning media cannot keep pace with the requirements. Memory arrays with low random access times and high IOPS can deliver a greater-than-10X improvement in system performance at a small incremental cost.

Virtual Storage

The virtualization of servers and desktops has been a major trend over the last few years and shows no signs of abating as demand increases for both public and private clouds. Increased CPU and memory utilization leads to much lower capital and operating expenditures and is a major benefit to most organizations. Virtualization impacts storage by increasing the number of I/Os per CPU, increasing the randomness of the I/O, and adding the demand for the back-up and restore of virtual machines. Collectively, this requires more IOPS for storage and demands lower latency to keep the CPUs busy. Memory arrays have proven to be an excellent approach to the consolidation of storage infrastructure with 80% power and space savings while still delivering five times the IOPS of traditional storage media.

Advances in FPGA Technology

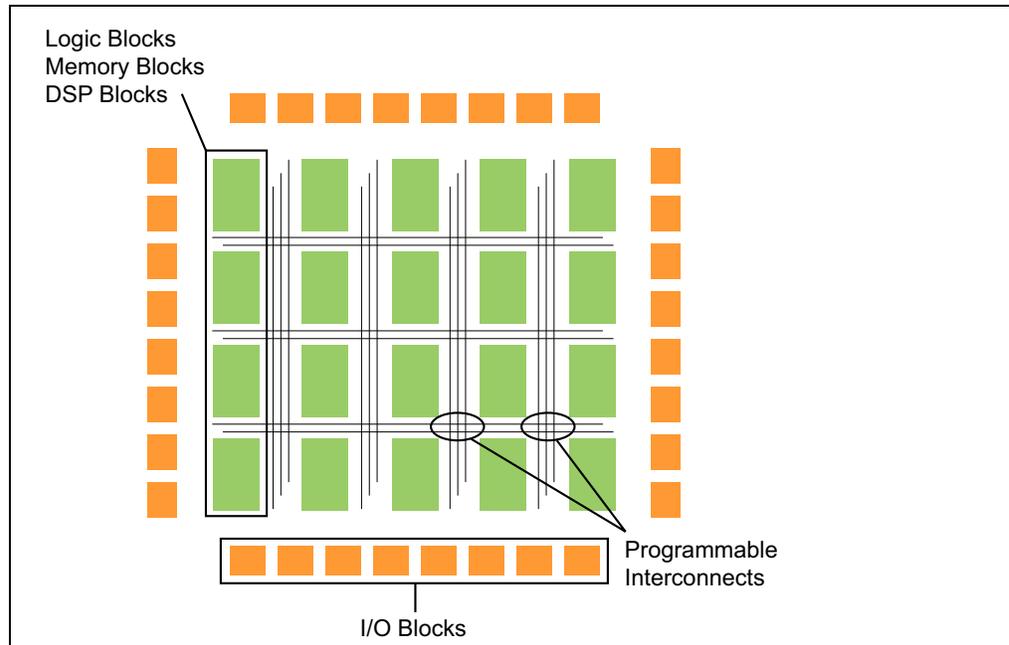
Recent advances in FPGA technology pertinent to flash memory arrays include memory control and high-speed I/O interfaces.

Memory Control

Technology advancements in PLDs support the expanding requirements in caching memory architectures with flash memory interface support and control. Performance increases in core fabric clock speed and I/O interfaces enable the support of the latest flash memory types (ONFI 3.0 and Toggle Mode 2.0).

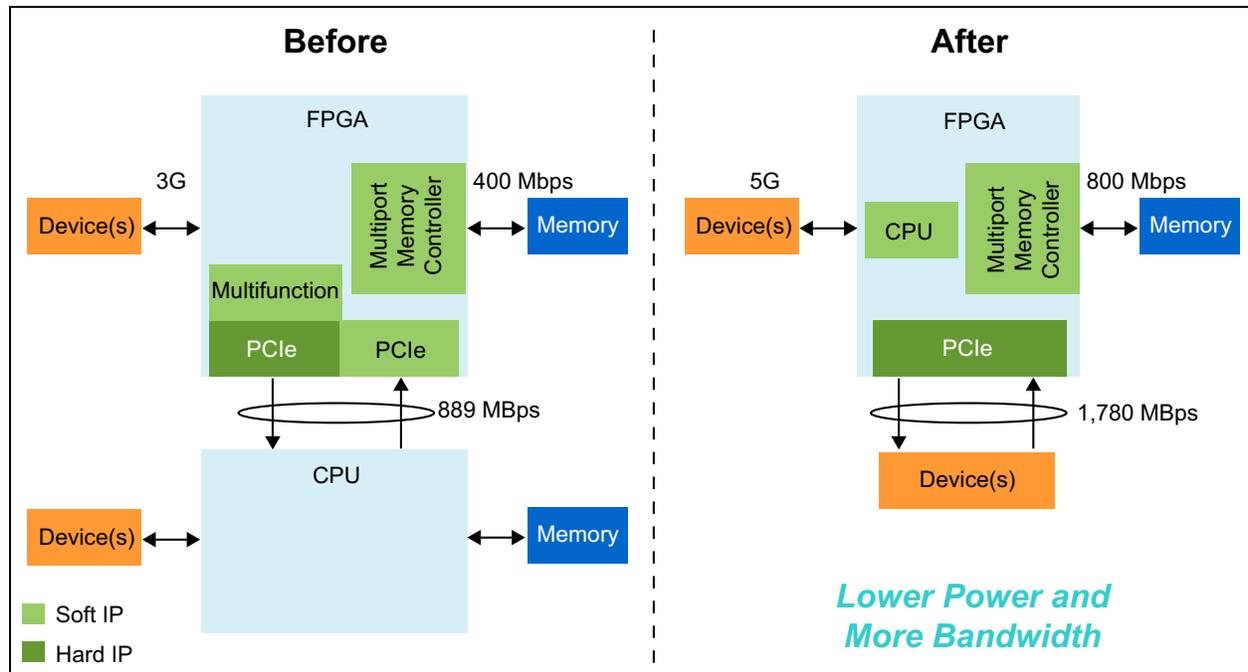
Hardened memory controller blocks (Figure 2) provide an efficient performance advantage over traditional soft IP implementations. Specifically, hardened IP enables additional otherwise constrained soft logic resources to support more efficient designs, such as the programmable state machine for flash cache described in the preceding section.

Figure 2. Memory Controller Blocks in the FPGA Architecture



FPGA technology advancements support the integration of distributed block functions such as memory controllers, embedded processors, and high-speed serial interfaces. Figure 3 shows how FPGAs can provide a more efficient solution with higher bandwidth, power savings, and a reduced board footprint.

Figure 3. Subsystem Integration with FPGAs

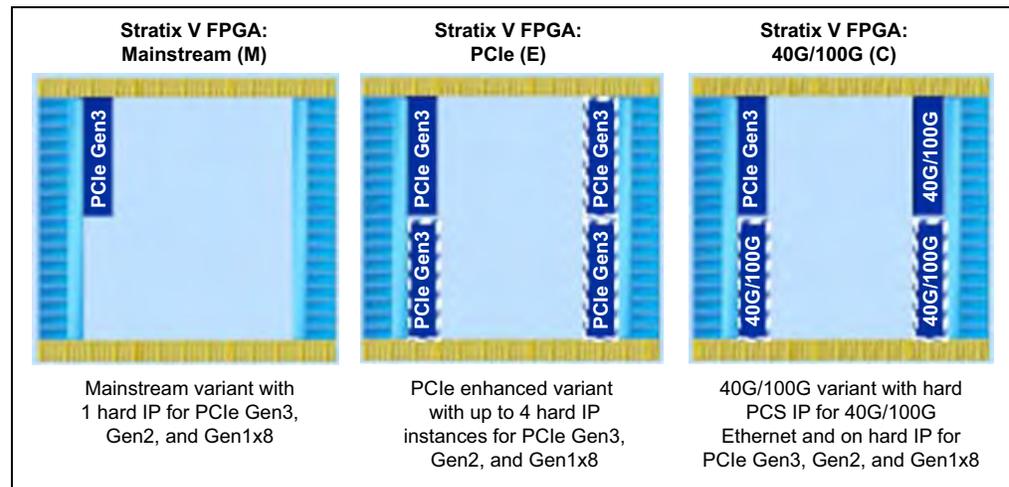


High-Speed I/O Interfaces

PLDs continue to share the leadership in semiconductor process node technology advancements where the current generation of FPGAs is at the 2X nanometer node. High-speed interfaces in PLDs are used to transfer data across high-speed data traffic hubs, where transmission speeds of up to 28 Gbps are now realized in 28-nm PLDs. At this advanced process technology node, the electrical- and physical-layer performance requirements of the highest speed serial protocols—such as the third-generation PCIe, SAS/SATA, and Fibre Channel technologies—can be supported with hardened transceiver blocks. Hardened transceiver blocks provide a stable and reliable configuration for optimum transmit and receive performance with minimum signal jitter. Hardened transceiver blocks also enable more programmable resources than soft logic implementations.

Advances in PLD packaging technology support an increased number of high-speed I/O ports along with increases in general-purpose I/O pin counts. These process-node advancements support high-speed I/O interfaces as well as faster memory control. Figure 4 shows the Altera Stratix® V FPGA variants that are supported with this process technology.

Figure 4. 28-nm Programmable Technology—Altera Stratix V Product Option



Conclusion

The benefits associated with PLDs—design flexibility, modular IP integration, hardened memory controllers, and high-speed serial interfaces—provide an effective technology option in the design of flash memory array architectures. Programmable state machines provide the best performance in supporting storage subsystem requirements.

Major paradigm shifts, such as the move from disk technology to flash memory storage, require a revamping of traditional storage architectures with a cost-effective way to both innovate and then grow capacity as demand increases. The memory arrays with integrated flash-specific RAID and flash controllers are an excellent example of what is required to address these newer markets.

Where innovation is mostly a software function, microprocessors can provide the necessary platform. Where real-time, low-latency, and high-bandwidth solutions are required, the business will naturally involve more application-specific semiconductors. Through advances in both interfaces and core logic, PLDs have met this need and are a much less capital-intensive solution.

The use of programmable technology, such as Altera FPGAs, as a means to support emerging and demanding memory-array architectures from prototype to volume production has proven to be very successful for innovative storage companies such as Violin Memory.

Further Information

- Computer and Storage:
www.altera.com/end-markets/computer-storage/cmp-index.html
- White paper: *Providing Battery-Free, FPGA-Based RAID Cache Solutions*
www.altera.com/literature/wp/wp-01141-raid-cache.pdf
- Violin Memory:
www.violin-memory.com

Acknowledgements

- David McIntyre, Senior Business Unit Manager, Altera Corporation
- Morgan Littlewood, Vice President of Product Management, Violin Memory, Inc.

Document Revision History

Table 1 shows the revision history for this document.

Table 1. Document Revision History

Date	Version	Changes
August 2011	1.0	Initial release.