

Unveiling Nova

Using a single Virtex-II Pro FPGA to create a reprogrammable video switching system.



By Roger Smith
Chief Engineer
Echolab LLC, Inc.
rsmith@echolab.com

Television mixers have historically been built with dedicated hardware to achieve a specific fixed functionality. As mixers have evolved from devices built with discrete transistors to more modern mixers with advanced large-scale integration (LSI) integrated circuits, a common limitation has been that these devices were built with fixed signal and data paths that pre-define the topology of the mixer. Thus, a mixer targeted for two mix/effects (M/Es) and two keyers per M/E is built specifically for that function, with limited or no future adaptability.

The Echolab Nova series completely breaks with tradition in this regard, moving to a completely reconfigurable platform based on a system-on-chip architecture.

We used a single Xilinx® Virtex™-II Pro FPGA to create a completely reprogrammable video switching system. By absorbing the interconnects of the mixer into a single FPGA, the limitations of a fixed signal path architecture have been removed, and the topology of the mixer can be redefined again and again throughout the life of the product.

Nova Family

In 2004, Echolab launched the first member of the Nova series, the Nova 1716. This product is a full program/preset mixer with 16 inputs and 16 outputs, two downstream keyers, and a full M/E upstream complete with two effects keyers. In 2005, we introduced several new members of the family. The Nova 1932 is a 32-input program/preset mixer, with two upstream M/Es and a

Xilinx parts is support for internal termination of high-speed differential inputs. Because of the high-speed nature of this video I/O, the wires must be treated as transmission lines, with great care paid to electrical termination. Historically, these termination resistors would be placed directly outside the chip, as close as possible to the end of the line. The need for a substantial amount of terminations so close

clock cannot be initially used to sample all of the incoming signals.

Typically, most SDI mixers use individual clock and data separator circuits on each input so that the hardware can recover the individual bits from each stream. After the data streams are separated and decoded, sync detector circuits are used to write these streams into FIFO memories. A common genlock clock and reference is then used to read out the video streams from the memories for effects processing downstream.

This topology is not viable for Echolab's system-on-chip architecture. The large number of clocks created by these front-end clock and data separators would overwhelm the FPGA's support for the total number of clocks, as well as the inherent need to crowd these parts near the FPGA. The video interface to the FPGA must be simpler, and require few or no parts near the FPGA. Thus, traditional clock and data separation techniques do not work here.

Echolab applied an asynchronous data-recovery technique from Xilinx application note XAPP224 ("Data Recovery") originally developed for the networking market. The technique uses precise low-skew clocks to sample the inputs in excess of a gigahertz. The samples are examined to determine the location of the data bit cell transitions. The encoded data is extracted from the stream and can cross into the genlock clock domain without ever extracting the clock from the input stream. For more information about this technique, see www.xilinx.com/bvdocs/appnotes/xapp224.pdf.

The FPGA's ability to route nets and guarantee skew performances on the order of picoseconds has enabled development and implementation of this unique SDI input.

Another design challenge in digital mixers has been the implementation of video line delays and FIFOs, which have historically been used in large numbers to time internal video paths and outputs. It is often necessary to add delay lines to AUX bus outputs to keep them in time with the primary mixer outputs. These delay lines and FIFOs have typically been implemented with discrete memory devices.

The Xilinx FPGA solution contains large numbers of video line-length memo-



Figure 1 – The Nova switcher family

full complement of six keyers. The identity4 is a 1 M/E 16-input look-ahead preview mixer. The identity4 brings tremendous advances in video layering, with four upstream and two downstream keyers and five internal pattern generators, all in a flight-pack-size panel and frame.

As shown in Figure 1, what is unique about these mixers is that they share a common frame and electronics, which are simply re-programmed to fit the target mixer design.

I/O

Major advances in FPGA design were necessary to undertake such a dramatic shift in mixer architectures. One of the first challenges that we had to overcome was how to get the video bandwidth into and out of the FPGA.

Even the smallest Nova family member (shown in Figure 2) has 16 SDI inputs and 16 SDI outputs. At 270 Mbps, this is an aggregate video data bandwidth of more than 8.5 Gbps. The Nova 1732 and 1932 family members have 32 inputs and 16 outputs, approaching approximately 13 Gbps.

One of the key aspects of the Virtex-II Pro device is its ability to support I/O bandwidths in excess of 400 MHz on all FPGA I/O pins. Another key feature of

to the chip would be a layout complication. Because Xilinx can support internal termination of these high-speed video signals, this has greatly simplified the architecture, making the support of large blocks of high-speed video I/O possible.

The next challenge was deciphering the SMPTE 259M stream. Because the SDI streams are coming from all over the studio, even when genlocked (synchronizing a video source with other television signals to allow the signals to be mixed), there can be large phase shifts of +/- half a line between these signals, and therefore, a common



Figure 2 – Nova SDI I/O

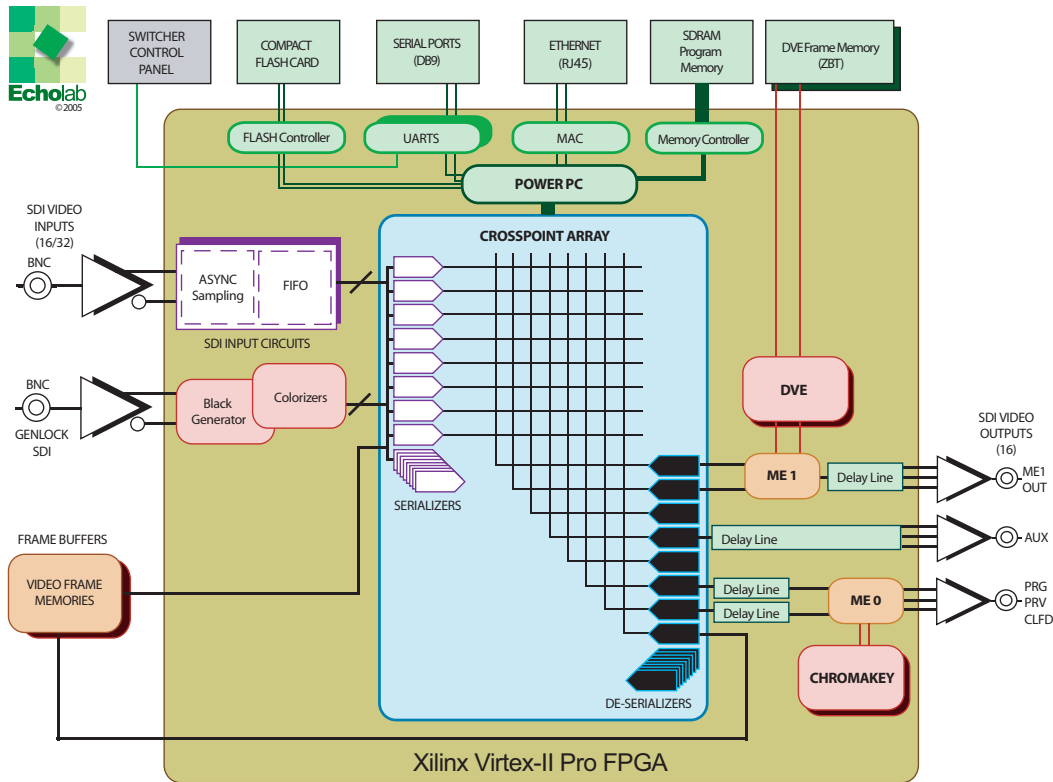


Figure 3 – Virtex-II Pro switcher implementation

ries on chip, which lend themselves naturally to delay lines. Xilinx has enough of these memories on-chip to allow all AUX bus outputs on the Nova series to be timed.

Crosspoint Array

The next challenge was to implement a crosspoint array of sufficient size to support these large mixers without overwhelming the resources of the chip. Although the crosspoint array appears to be a reasonable size from the front panel, the array is often much larger because it must support all of the internal sources and functions inside the mixer.

The Nova 1716 has 16 external inputs, but it also has 3 internal colorizers, black, three internal frame buffers, and several intermediate sources generated by the upstream M/E. For outputs from the array, each M/E requires seven video buses (A/B, key 1 cut and fill, key 2 cut and fill, as well as video for borders). There are also 12 AUX bus feeds and dedicated buses to support capture on the internal frame buffers.

By the time you add it all up, the required internal crosspoint array is easily 30 x 30 for the 16-input mixer. To develop this crosspoint array as a 10-bit wide parallel implementation would consume a large amount of resources within the FPGA.

A more effective use requires the design of high-speed serializers and de-

serializers within the part. The Virtex-II Pro FPGA's ability to generate multiples of the genlock clock within the part with low skew allows you to create whole sections of the chip that can run at the SMPTE 259M bit-serial rate of 270 MHz. This implementation of an SDI serial rate crosspoint array (Figure 3) effectively limits the use of valuable FPGA fabric to less than 10 percent of the capacity of the target chip.

Effects Generation

Once the streams have gone through the crosspoint array and have been de-serialized, the video buses, now "timed," can be routed to the appropriate processing blocks within the FPGA to perform various video effects (see Figure 4).

The creation of video effects within the FPGA such as wipes, mixes, and keys is easily performed with the basic building blocks of the FPGA. Most basic video effects can be performed with nothing more than simple combinations of addition, subtraction, and multiplication. Hundreds of embedded high-speed multipliers within the FPGA fabric allow a variety of video effects to be performed effortlessly with very high precision.

Embedded memory can also be used for LUTs and filter coefficient storage. An example of an often-used filter would be an interpolating filter for 4:4:4 up-sampling before a DVE or Chromakey.

A high-precision circle wipe requires a square-root function. This complicated mathematical function was implemented with a CORDIC core provided by Xilinx. These blocks of pre-built and tested IP from Xilinx and other third-party developers allow you to rapidly deploy designs without having to develop and test every building block from scratch. Development with these IP cores can remain at a very high level, providing quick time to market. Also, the optimized



Figure 4 – Effects

size of these cores allows the design of complex mixers within a single chip.

Embedded PowerPC

The computer horsepower necessary to run today's large vision mixers has grown immensely over the last dozen years. The need to accurately support field-rate effects on multiple M/E banks while at the same time communicating with complex control panels and many external devices has tested the limits of earlier 8- and 16-bit microprocessors.

Most manufacturers have either used several smaller distributed processors or a larger, faster 32-bit microprocessor. Echolab's system-on-chip architecture (see Figure 5 – PowerPC™ implementation) takes advantage of two 32-bit PowerPCs

letproof as previous generations of mixers.

Echolab has chosen Micrium's μ C/OS-II real-time operating system (RTOS) for the Nova series (see www.micrium.com). This OS is a priority-based, pre-emptive multitasking kernel that has been certified for use in safety-critical applications in medical and aviation instruments. Time-critical video processing is assigned to the highest priority task. Management of the file system, console I/O, and network stack are allocated to lower priority tasks, allowing the processor to utilize spare processor cycles in the background without ever interfering with the video hardware.

The tasks communicate with each other – and synchronize their activity – with thread-safe semaphores and message queues provided by the OS.

6), designed to hold all of the firmware and software to configure and boot the Nova. With a user-accessible mode switch, you can load as many as eight different on-line configurations of Nova firmware and software from a single flash card.

The remaining storage on the card is available to store user data such as sequences and panel saves, as well as key memories and other user settings. Also, the Compact Flash is designed to hold all of the graphics and stills online. Support for cards up to 2 GB or more provides an immense storage capacity, well beyond the archaic floppy disks found in competitive products.

Given Nova's unique architecture, it is easy for the product to be extended through downloadable firmware updates. These updates can be as simple as a routine

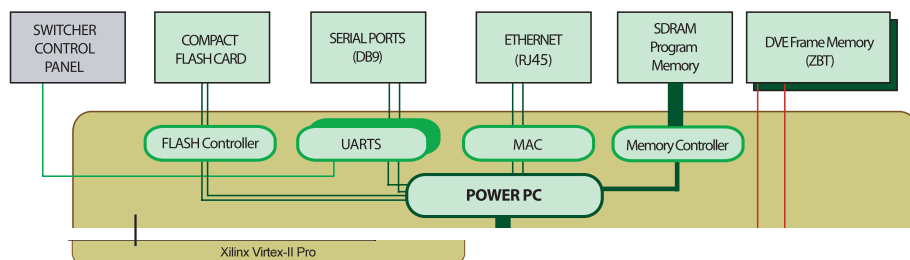


Figure 5 – PowerPC implementation

(running at 270 MHz) embedded directly in the fabric of the Xilinx FPGA. This tight coupling between the processor and the mixer hardware leads to substantial performance improvements.

A vast library of pre-built processor peripherals from Xilinx and third parties makes it easy for functions like serial ports, memory controllers, and even Ethernet peripherals to be dropped right into the design. Custom peripherals are also easy to design.

Several large switcher manufacturers have gone to commercial operating systems like Windows or Linux to improve their software productivity. Although the early benefits can be appealing, the downside to this transition is a loss of control over the reliability of the switcher's code base, making a device that is not as robust and bul-

System Connectivity

All Nova series switchers have multiple ports for broadcast studio interconnectivity. An industry-standard RS-422 port allows for the implementation of industry-standard editing protocols. A standard RS-232 port is available for PC connectivity. An Ethernet port allows the switcher to be directly connected to a network.

Under control of Nova's μ C/OS-II, several servers are running concurrently that provide an integrated Web server for remote status and display, an embedded XML-RPC server for remote control, as well as a full TFTP server for remote upload and download of graphics and stills.

Compact Flash and Re-Configurability

At the heart of the Nova system is an industry-standard Compact Flash card (Figure



Figure 6 – Compact Flash

software patch, as complex as adding a keyer to an AUX bus output, or restructuring the internal video flow within an ME for a specialized application. Architectures as different as the 2 ME Nova 1716 (with its program/preset architecture) and the Nova identity4 (with a six-keyer look-ahead preview structure) can be loaded into the same hardware.

Roadmap to the Future

Next-generation Virtex-4 FPGA technology from Xilinx will allow Echolab to move its system-on-chip architecture to support high-definition products. Virtex-4 FPGAs bring a

higher level of performance to the embedded logic, as well as the embedded peripherals. Some of these enhancements include:

Fabric enhancements

- Larger arrays
- Faster
- Lower power

I/O enhancements

- General-purpose I/O speeds to 1 GHz
- Dedicated Rocket IO™ transceiver speeds beyond 10 GHz

Dedicated hardware resources

- Multiple tri-mode Ethernet MACs
- 500 MHz multipliers with integrated 48-bit accumulators for DSP functions
- Block memories now have dedicated address generators for FIFO support

Conclusion

Television mixers have grown larger and more complex in the last dozen years. More and more, they are the focal point for the interconnection of a wide range of studio equipment.

One of the primary benefits of the system-on-chip architecture is reduced parts count. This reduction in parts count contributes directly to lower power, reduced PCB complexity, higher reliability, and reduced cost.

Another major benefit of the system-on-chip architecture is that it is almost entirely reconfigurable. This has allowed multiple products with different video architectures to be built on a common platform. This also lends itself to easy customization for specialized applications or specific vertical markets.

As the computer network continues to play more of a role in today's modern television studio, the Nova series will be ready with support for streaming video over Gigabit Ethernet. H.264 and WMV9 codecs will drop right into the Nova's system-on-chip architecture, providing future features on today's hardware.

For more information, please feel free to see Echolab's complete line of television mixers at www.echolab.com.



Onward to Glory

Complete Software Radio Solution on a single 6U card

Features

- ▶ 6 Million gate Virtex-II FPGA
 - Scalable Software Defined Radio
 - Embedded Processing via MatLab
- ▶ TMS320C6416 DSP
- ▶ 105 MHz, 14 bit 2 Ch. Analog I/O
- ▶ 32 MB RAM
- ▶ 32/64-bit cPCI bus
- ▶ PMC expansion site
- ▶ STAR Fabric interface







Get your data sheets now!
www.innovative-dsp.com/quixote
 sales@innovative-dsp.com
 805.520.3300 phone • 805.579.1730 fax

FREE Multimedia Demo

Seamless FPGA for Xilinx Virtex-II Pro



During this demo, you will learn the following:

- Simplifying the embedded processor verification flow
- Easily importing designs from Xilinx Platform Studio EDK to Seamless FPGA
- Verification with Seamless FPGA: accelerating the simulation process
- Performance profiling and analysis

Watch the demo today:
http://www.mentor.com/products/fv/hws_w_coverification/seamless_fpga/demo_request.cfm

For additional details about Seamless FPGA,
visit us at www.seamlessfpga.com

