How to Pack a Room of Analog FM Modulators into a Xilinx FPGA

by Fred Harris
Professor
San Diego State University
fred.harris@sdsu.edu

Dragan Vuletic
Director of Engineering
Signum Concepts
dragan.vuletic@signumconcepts.com

Wade Lowdermilk
CEO
Signum Concepts
wade.lowdermilk@signumconcepts.com

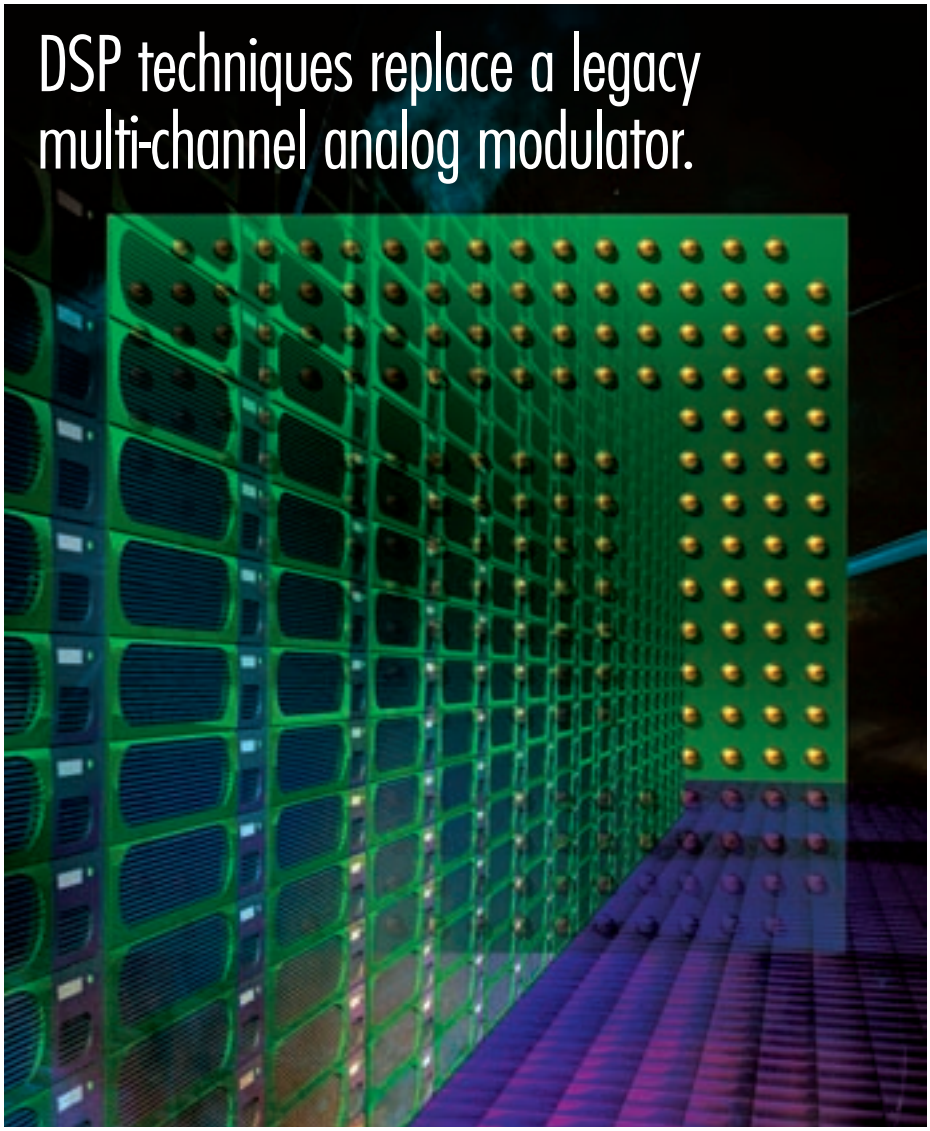

You are likely familiar with the way that digital television is transmitted from satellites as multi-channel MPEG (Motion Picture Experts Group) compressed video to a cable head end where the multiple channels are demodulated. The MPEG streams are decoded and then remodulated as channelized analog NTSC (National Television Standards Committee) or PAL (Phase Alternating Lines) television signals for insertion in a cable distribution plant.

Similarly, high-quality stereo audio is transmitted from a satellite as multi-channel MP3 (MPEG Layer-3) compressed audio to a cable head end where the multiple channels are demodulated. The MP3 streams are decoded and then remodulated as channelized analog FM signals for insertion in a cable distribution plant.

Figure 1 – Equipment bay containing legacy transceiver equipment

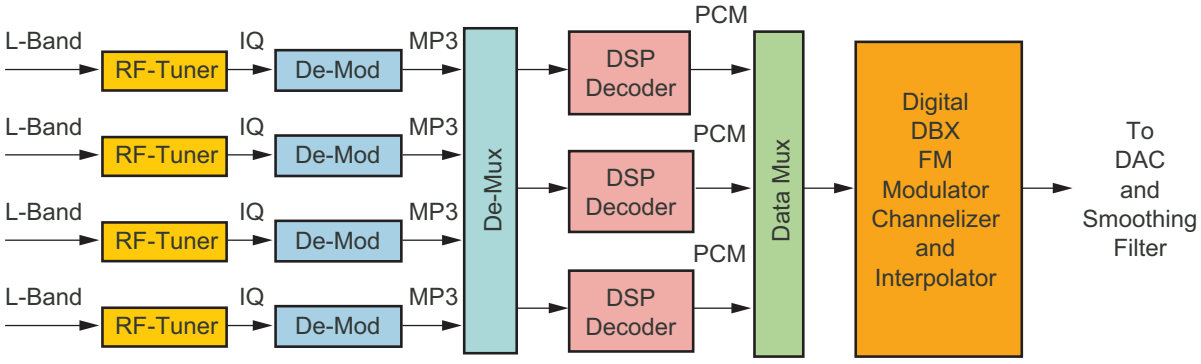

Figure 2 – High-level block diagram of transceiver

The device that combines a transmitter and receiver in a single unit is called a transceiver, while the device that performs demodulation of a collection of channels in one format and remodulates them in a second format is called a transmultiplexer. We will resist the temptation to coin a new noun from pieces of the words transmitter, receiver, translator, and multiplexer and refer to the system we describe here as a transceiver.

In this article, we'll describe the implementation of a unique digital transceiver system. The system decodes 384 MP3 stereo audio signals received from a satellite downlink. It digitally remodulates the multiple sampled data signal sets as a single analog signal made up of 384 frequency-division multiplexed (FDM) channels. Each channel contains a pair of stereo signal components as well as a third component similar to the commercial FM broadcast signal known as a subsidiary carrier authorization (SCA).

The three signals are pre-processed by a dbx (noise reduction system) and standard pre-emphasis filters before FM modulation and channelization. The entire transceiver occupies a 2U rack-mount chassis (436-mm wide [17.16"], 350-mm deep [13.78"], and 87.1-mm [3.4"] high) – replacing a legacy hardware rack containing 192 analog modulator circuit boards. Figure 1 shows the bay of legacy equipment that previously performed the 384 channels of MP3 decoding, dbx precoding, pre-filtering, and FM modulation.

The dbx encoding, digital pre-emphasis system-specific stereo FM modulation, and resampling channelizer are performed in a single Xilinx® Virtex™-4 XC4VSX55 FPGA. The FPGA outputs a single sampled data stream of channelized FDM signal spanning approximately the 20- to 80-MHz frequency band. The MP3 decoding is performed by 48 Blackfin processors from Analog Devices.

The all-DSP-based transceiver that replaces the bay of analog equipment was recently fielded by X-Digital Systems, a San Diego-based satellite communications company.

The emphasis of this article is on the processing tasks performed in the FPGA-based modulator, designed and implemented by Signum Concepts, another San Diego-based communication systems development company.

High-Level Description

The transceiver comprises four subsystems: the demodulator block, the demultiplexer block, the decoder block, and the FM modulator block. The demodulator block accepts four L-Band input signals as inputs to four RF tuners, which down-convert selected QPSK RF bands to IQ (in-phase and quadrature phase) baseband signals. The baseband IQ signals are processed by the demodulator block to form MP3 (MPEG audio) data or transport streams.

These demodulated streams enter the demultiplexer block, which merges the streams with optional local asynchronous digital MP3 transport streams and partitions the composite transport stream for delivery to the 48 Blackfin processors in the decoder block. The decoder block forms 16-bit pulse-coded modulation (PCM) versions of the MP3 coded input transport stream. The PCM data is delivered through a high-speed output bus to the FPGA modulation module. Figure 2 is a high-level block diagram of the system.

System Considerations

We designed the digital FM modulator and channelizer to replace an existing analog modulator system. The output of the modulator is demodulated by legacy analog receivers. These receivers contain a standard channelizer, a set of conventional FM demodulators, and analog dbx decoders. The digital modulator is required to seamlessly interact with the analog receiver, which means that the design parameters of the digital modulator have to match the design parameters of the legacy analog modulator. Thus, the gain and phase of the digital filters in the digital dbx encoders have to match those of the analog dbx encoders. This is required to assure that the analog dbx decoders perform as well with the digital encoder as they do with the analog encoder.

Similarly, the gain and phase of the pre-emphasis filters used in the digital FM modulator must match the gain and phase of the analog pre-emphasis filters. Here too, the requirement ensures that the receiver performance is the same when receiving signals formed in the digital modulator as signals formed in the analog modulator. For the same reason, the modulation index and the modulated bandwidths of the digital FM modulator chain must match the corresponding values of the analog FM modulator chain.

This is an unusual design constraint. Usually, DSP-based systems enhance the performance of an analog prototype system.

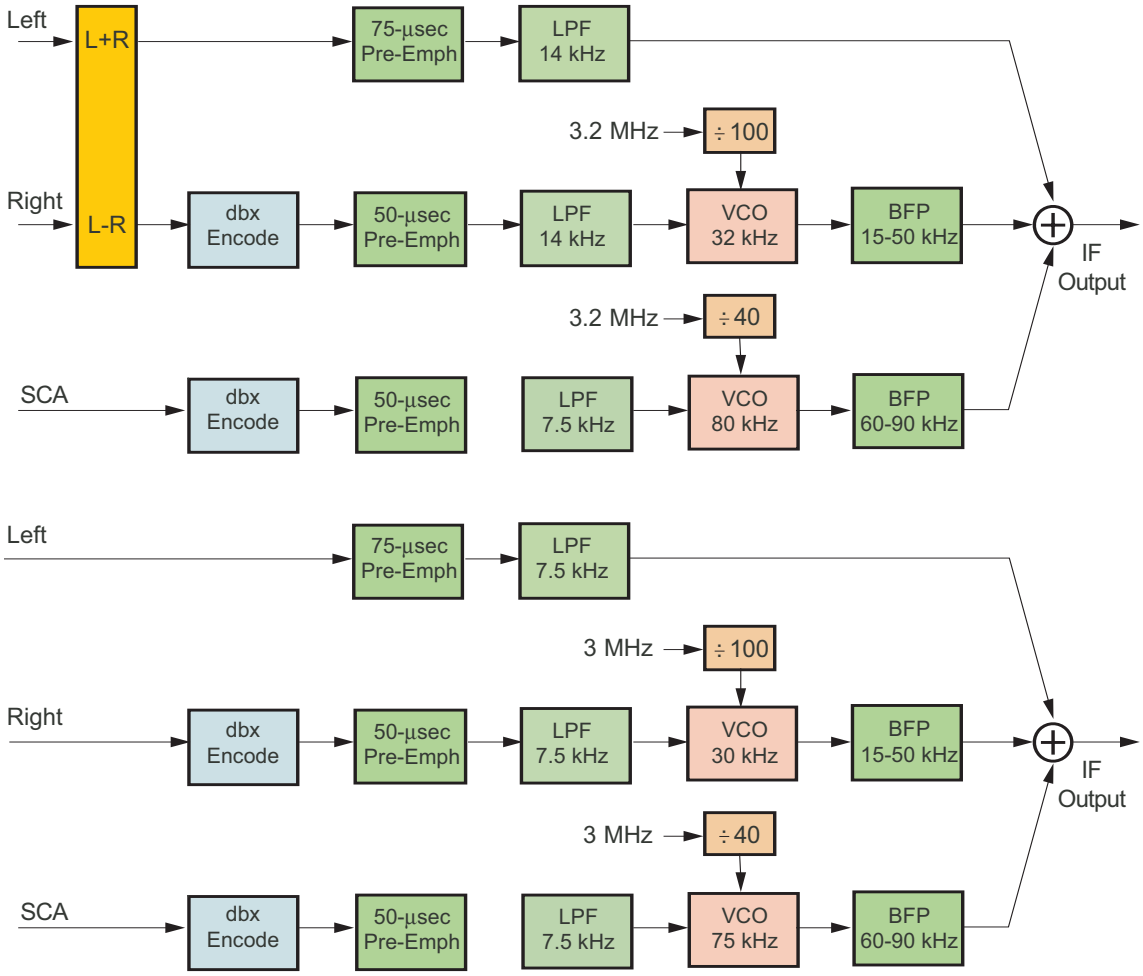

Figure 3 – Signal processing chain for each channel of the FM modulator: stereo and non-stereo options

Here, the DSP-based system is required to emulate the analog system. This is particularly important because the acceptance test procedure for the signals formed by the digital modulator chain is the procedure written for the analog modulator chain.

Figure 3 shows a simplified block diagram of the analog processing chain for each FM signal set. There are three input signals in each set: L (left), R (right), and SCA. Also note that the modulator can operate in two modes, stereo and non-stereo. The two modes differ in filter bandwidths (14 kHz or 7.5 kHz) and center frequency of the two VCOs (32 kHz and 80 kHz or 30 kHz and 75 kHz, respectively).

Of course, the digital modulator must mimic both modulation modes. In this modulator, narrowband FM is used to modulate and spectrally offset the baseband L-R signal component as well as the baseband SCA signal component. The composite signal is then frequency modulated and up-converted to its appropriate RF center frequency.

This modulation format is markedly different from the modulation format used in the commercial FM band. There, the compatible stereo uses double side-band suppressed carrier (DSB-SC) modulation for the L-R and SCA signals and further contains a pilot signal to enable demodulation of the suppressed carrier signals.

The composite signal is then FM modulated to the RF band center. Remember not to confuse the FM modulation described in this transceiver with the commercial FM modulation scheme with which you are familiar.

A DSP Perspective

As explained in the previous section, the digital implementation of the modulator segment must mimic the performance parameters of the legacy analog implementation. This does not mean that the digital implementation must emulate the analog implementation. The DSP systems rule is not to emulate the analog process but rather to return to first principles and avoid inheriting legacy compromises. Although there are some processing blocks common to both realizations, others are purely DSP related.

For example, the analog system uses a conventional VCO (voltage controlled oscillator) to implement the two FM modulators required for each signal set. The digital system uses complex baseband phase modulation of the up-sampled and digitally

integrated input signal to form the digital FM sequence. The block diagram of the DSP-based FM modulator is illustrated in Figure 4.

This process forms a complex baseband sequence whose phase derivative is proportional to the input signal. The sample rate of this sequence is raised by a second digital interpolator and digitally heterodyned to shift the baseband spectrum to the desired center frequency. The complex exponentiation and the DDS are implemented by two versions of the well-known CORDIC (Coordinate Rotation Digital Computer) algorithm.

Note that interpolators are used throughout the modulation process. We do this to permit signal processing at the lowest sample rate consistent with the signal bandwidth. The interpolators are applied to the sampled data signal at appropriate points in the processing flow to enable bandwidth expansion in successive processing stages.

Single-Channel Processing Block

Figure 5 is a block diagram of a single-channel processing block that handles a stereo signal set and an SCA signal. Each of the three input signals is delivered to this block as 16-bit-wide data at a 48-kHz sample rate. The three signals are subjected to three different processing paths that have limited commonality. Interestingly, only a single-channel processor was programmed into the system; this single processor is sequenced through the 384 signals to be channelized.

The stereo components are initially processed by a butterfly to form the standard L+R and L-R components. The L-R and SCA signals are encoded in dbx dynamic range enhancement filters. All three channels are then pre-emphasized by first-order lead/lag filters, with the L+R pre-emphasis zero at 2.1 kHz and the L-R and SCA pre-emphasis zeros at 3.2 kHz. All three paths are

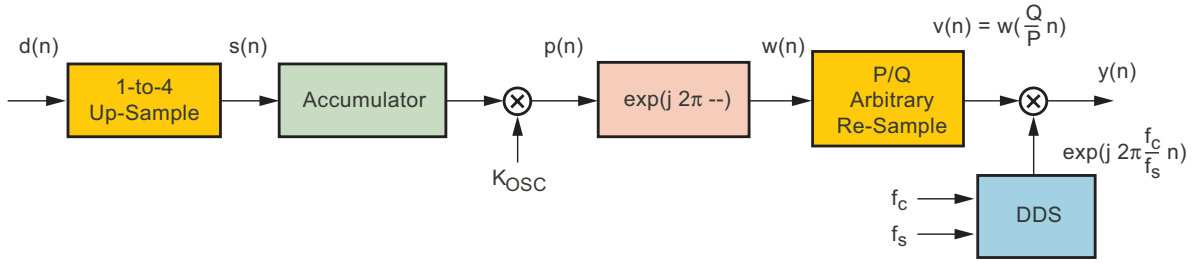

Figure 4 – Signal flow diagram: DSP-based complex baseband FM modulator

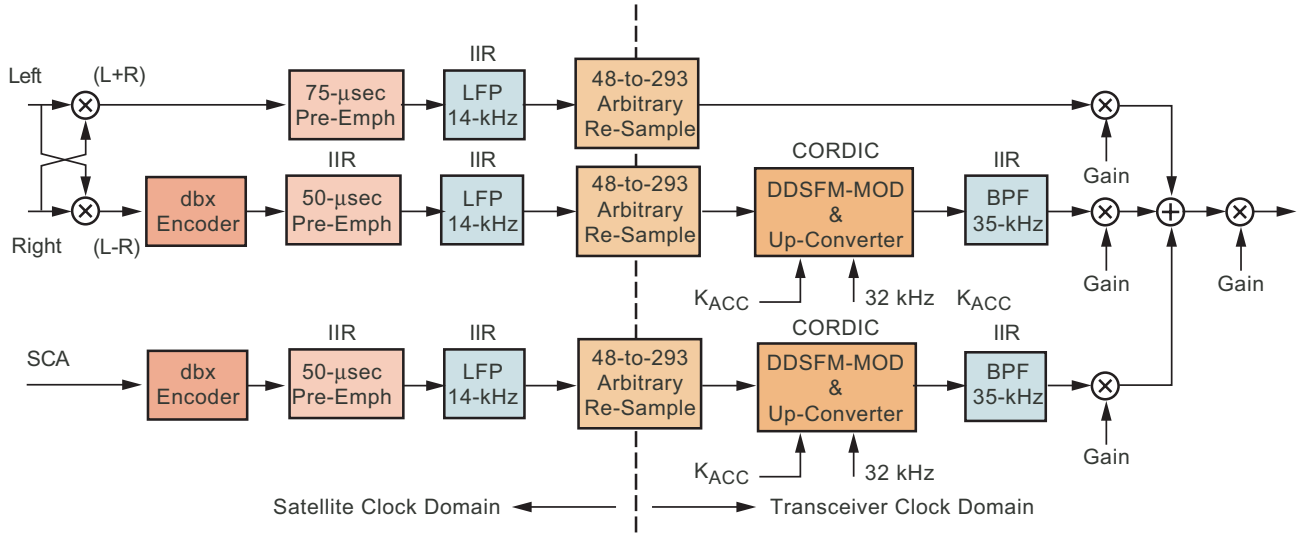

Figure 5 – Block diagram of single-channel processing block: stereo and SCA signal

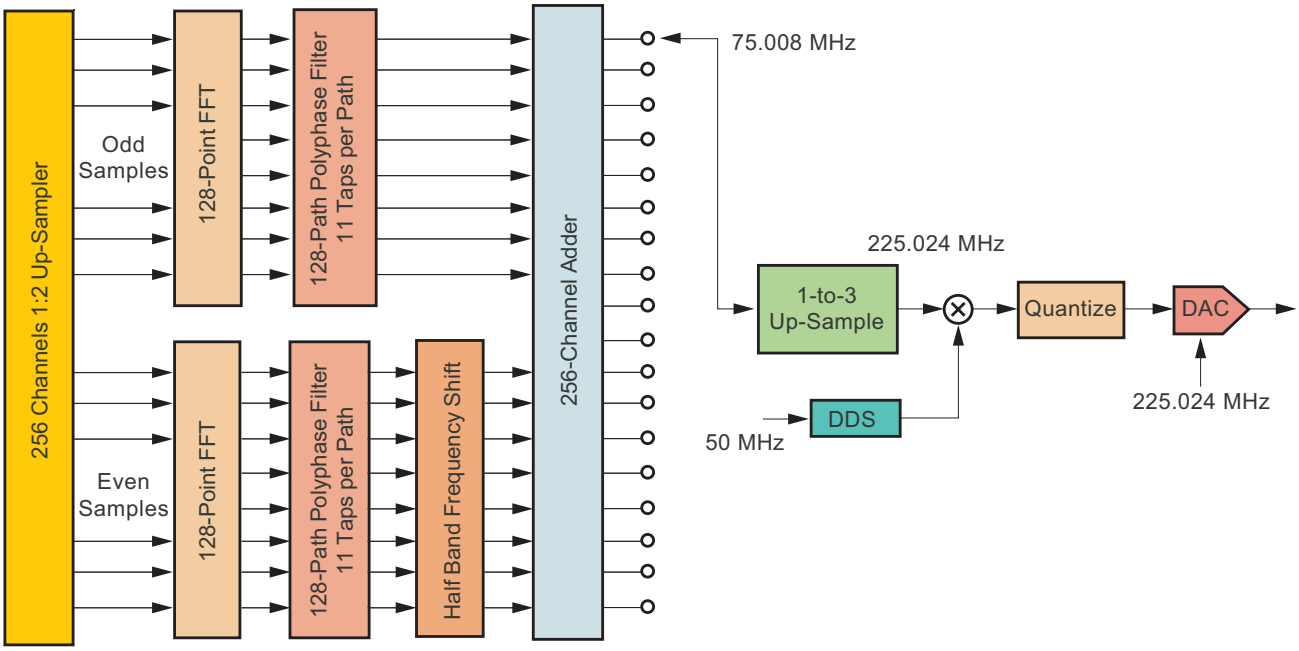

Figure 6 – Polyphase channelizer and final post-processing tasks to operate output DAC

filtered to limit their 3-dB bandwidth with the L+R and L-R band edges at 14 kHz and with the SCA band edge at 7.5 kHz.

The L+R signal and the SCA signal are up-sampled 1-to-4 to 192 kHz by polyphase FIR filters in preparation for the complex baseband FM modulators described in the previous section. The outputs of the two baseband FM modulators are again filtered to control their 3-dB bandwidths, which have experienced bandwidth expansion caused by the nonlinear FM modulation. The L-R FM-modulated low-pass filter bandwidth is 35 kHz, while the SCA FM modulated low-pass filter bandwidth is 30 kHz. Both FM-filtered signal components are up-sampled 192:293 by polyphase FIR filters to an output sample rate of 293 kHz. In turn, the outputs of the interpolators are spectrally translated by an IQ heterodyne to form real signals at their design center frequencies. The L-R is translated to 32 kHz and the SCA is translated to 80 kHz.

The filtered L+R signal is up-sampled 48:293 in a polyphase FIR filter to 293 kHz – the same sample rate formed at the output of the two paths performing the FM modulation and spectral translation. An important design consideration was embedded in the 192:293 and 48:293 up-samplers. These re-samplers are arbitrary interpolators that transfer data samples from the input satellite clock domain to the local transceiver clock domain.

The three signal paths are separately amplitude scaled to control their relative spectral densities and are merged by addition. The summed signal is scaled once again and delivered as one of 200 inputs to a 256-path polyphase channelizer that performs the complex baseband channelization and raises the sample rate to 75.008 MHz. This final scaling is used to control the spectral densities of the separate channels across the composite modulated bandwidth.

Output Interpolator, Channelizer, and Up-Converter

We have examined the signal processing blocks required to implement a single channel of composite baseband signal with offset FM sub-carriers. A total of 200 of these channels are formed with data streams obtained from the MP3 decoders. These baseband channels are merged in a polyphase up-sampling channelizer. The channelizer raises the sample rate by a factor of 1:256 from 293 kHz to 75.008 MHz. The two-sided bandwidth of each composite baseband channel is approximately 180 kHz. The spacing between centers of the channelized signal is 293 kHz.

From the channel width and the channel spacing, we find that the transition bandwidth of the prototype low-pass filter used in the channelizer is 113 kHz. To achieve 96-dB isolation between channels, the stop-band attenuation of the prototype filter has 96-dB out-of-band attenuation levels. Using standard equations found in filter design texts, we can estimate the length of the prototype low-pass filter that meets the specified sample rate, pass-band bandwidth, transition bandwidth, and out-of-band attenuation levels. This length is approximately 2,896 taps. To better fit into the 256 paths of the polyphase partition, the

filter length is reduced slightly to 2,816 taps. The polyphase partition forms 256 paths of 11 taps each.

Figure 6 is a block diagram of the 256-path polyphase channelizer implemented as a cascade of a 256-point fast Fourier transform (FFT) and the polyphase partition of the prototype low-pass filter. The output data rate of the polyphase filter is 75.008 MHz. This time series is resampled in the final interpolator to 225 MHz. The up-sampled series is heterodyned to move the spectral center of complex baseband

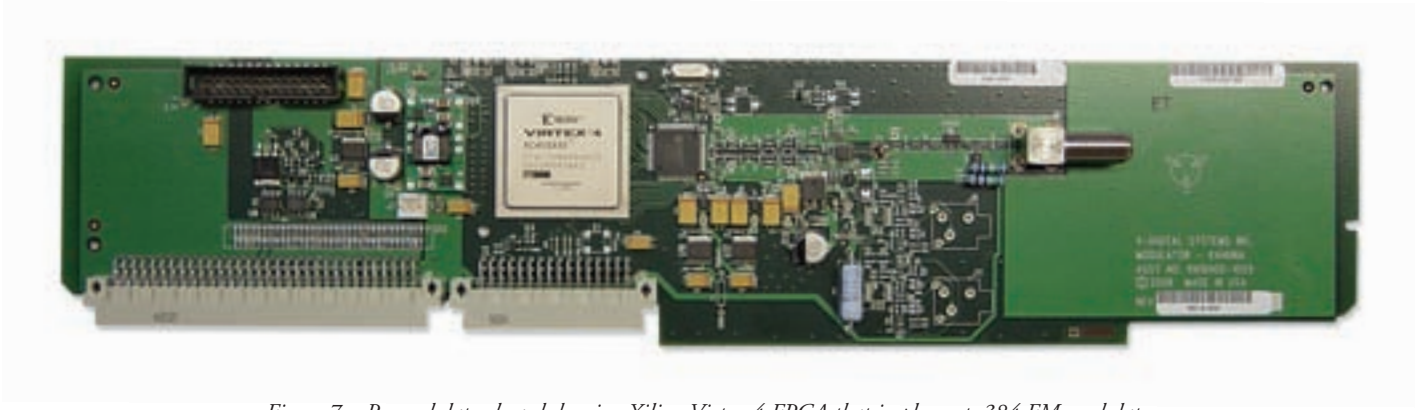

Figure 7 – Re-modulator board showing Xilinx Virtex-4 FPGA that implements 384 FM modulators implemented on 192 modulator boards in the equipment bay shown in Figure 1.

channelized time series to approximately 50 MHz. The real valued offset series is finally re-quantized to 14 bits and delivered to the 225-MHz DAC. The output of this DAC is passed through an analog smoothing filter to suppress the spectral remnants following the DAC's sin(x)/x filtering.

FPGA Resources

The Xilinx Virtex-4 XC4VSX55 FPGA contains a DSP cornucopia of 512 multiplier accumulator slices (18 x 18) and 320 blocks of 18-Kb RAM operating at 500 MHz. The modulator board is shown in Figure 7. Remember, this board implements the function and tasks of the 192 analog FM modulators mounted in the bay shown in Figure 1. The modulator implemented 200 sets of dbx encoders, pre-emphasis filters, multiple low-pass filters, FM modulators, multiple interpolators at 48-kHz input and 192-kHz output rates, clock domain interpolators and re-samplers to 293 kHz, the final 256-point FFT and polyphase channelizer at a 75-MHz output rate, and the final interpolator at a 225-MHz output rate. This impressive body of processing employed less than 60% of the FPGA resources.

FPGA Design Flow

We used the Xilinx System Generator for DSP model-based design flow to realize the transceiver. This is a tool chain that extends The MathWorks's Simulink framework with FPGA hardware-generation capabilities. Although Simulink has not traditionally been associated with hardware design, we found that System Generator was an excellent vehicle for developing FPGAs. System Generator is a visual design environment that allows you to work at a suitable level of abstraction from the target hardware and to use the same computation graph not only for simulation and verification but for FPGA hardware implementation. System Generator blocks are bit- and cycle-true behavioral models of FPGA intellectual property (IP) components, or library elements. The library-based approach results in design cycle compression, in addition to generating area-efficient high-performance circuits.

In addition to providing a natural development environment for developing FPGA signal processing implementations, System Generator has a rich set of features that support the development of heterogeneous applications comprising not only the FPGA element but a processor. The processor could be the embedded PowerPC™ 405 hard IP block, the MicroBlaze™ soft-processor core, or a processor external to the FPGA.

In the past, the Signum engineering team developed high-performance FPGA-based DSP systems using conventional HDL-based design flows for synthesis and simulation. In comparison, we estimate that our System Generator-based design flow compressed the development time to the final transceiver by a factor of roughly 10x compared to HDL-based flows.

Conclusion

The transceiver hardware described here is already replacing the fielded legacy analog system hardware. Because this hardware is DSP-based, we can easily incorporate additional capabilities and enhancements in the near future. Our next goal is to replace the analog legacy FM receivers (that demodulate our remodulated signal) with DSP-based systems with matching enhanced capabilities. With inserted enhancements at both ends of the communication link, the service provider will be able to offer additional revenue-generating communication services.

For more information, contact Fred Harris at *fred.harris@sdsu.edu.* For Signum Concepts FPGA-based activity, contact Wade Lowdermilk at *wade.lowdermilk@signumconcepts.com.* For X-Digital Systems activity, contact Ian Lerner at *ilerner@xdigital.com.*