SDR platform enables reconfigurable direction finding system

A software-defined radio (SDR) transceiver platform is used to implement a reconfigurable direction finding (DF) system. This article will describe the intended data flow and data rate requirements of a typical DF system, map the DF software components to the SDR platform, and demonstrate the availability of resources for more complex implementations.

By Jyothsna Rajan

Direction finding (DF) is the process of obtaining the angle of arrival bearings of radio signal emitters. DF has multiple military and commercial applications, such as remotely guiding an aircraft navigation system or tracing a signal interference source or locating a hostile target.

Most DF methods inherently use the phase of the incident signal wave front and its relative time-of-arrival at multiple antennae, arranged in a known configuration, to calculate its bearing.

A typical DF system must allow the user to perform spectral analysis over a predefined frequency band, select a signal of interest

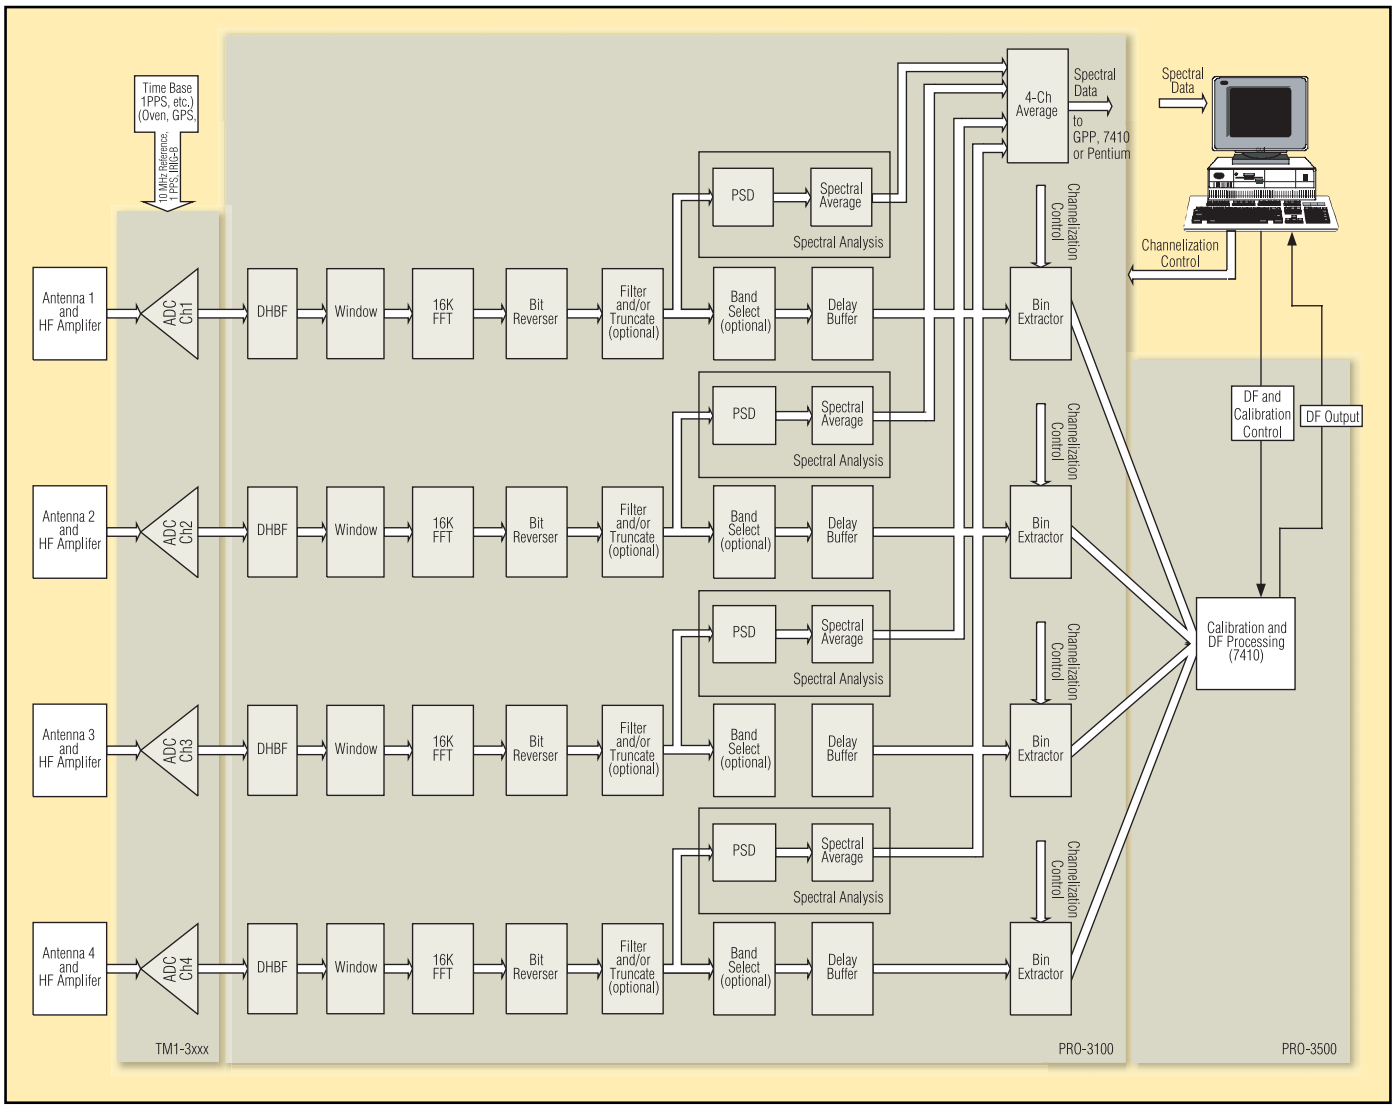

Figure1. Direction finding demonstration system processing blocks.

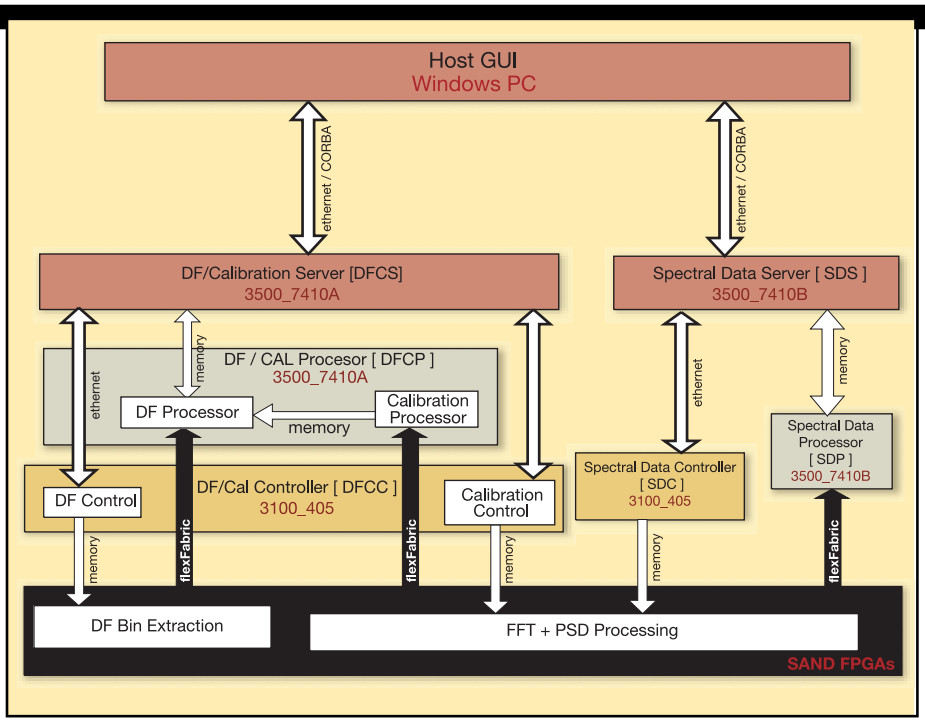

Figure 2. Software task mapping.

and then locate the source. This requires a system that can:

- digitize received analog IF signals;
- synchronize data from multiple receivers;
- calibrate data to account for component characteristic variations, environmental effects, etc.;
- perform windowing and digital filtering;
- execute FFT algorithms on the sampled data;
- channelize the data to extract the signal/s of interest:
- provide high-speed data transfer capability; and
- simultaneously process channel data from multiple receivers to deduce the angle of arrival (AOA).

Although a variety of available processing boards and I/O modules can cater to one or more of the above functions, there are few systems that can provide all of the functionality in a seamlessly integrated environment. One system that is capable of combining the above-mentioned fundamental functions is Spectrum's *flex*Comm™ SDR-3000 integrated development system. This commercial off-the-shelf (COTS) platform was recently used to demonstrate a software-defined DF system. Using this example, this paper will discuss the architecture of the SDR-3000 system, describe the intended data flow and data rate requirements of a typical DF system, map the DF software components to the SDR-3000 hardware components and demonstrate the availability of resources for more complex implementations.

Why SDR-3000?

The SDR-3000 provides a hardware

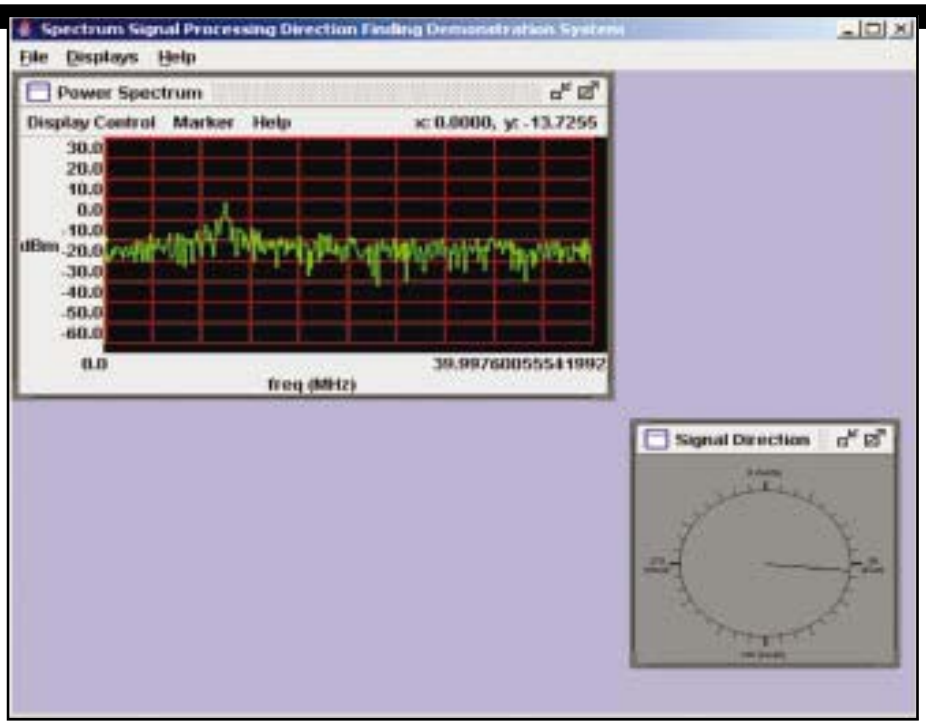

Figure 3. GUI display.

architecture and system configuration that enables users to combine all of the fundamental functions required in a DF application. The standard SDR-3000 system consists of:

- multi-input transition module (TM1-3xxx);
- signal conditioning board (PRO-3100) with four user-programmable Xilinx® Virtex-II™ XC2V6000 field-programmable gate arrays (FPGAs) and an IBM 405GP PowerPC™; and
- baseband processing board (PRO-3500) with two onboard Motorola MPC7410 PowerPCs and an IBM 405GP PowerPC.

The system provides the combined high-speed logic execution capability of the Virtex-II FPGAs with the algorithm execution capability of PowerPC MPC7410 G4 processors. A multi-input transition module (TM1-3xxx) enables the system to simultaneously receive inputs from different receivers, digitize and then transmit the data to the signal conditioning board (PRO-3100) via a high-speed digital IF fabric at rates up to 1.28 Gbps. Spectrum's *flex*Fabric communications architecture, based on RapidIO™ technology, facilitates interprocessor communications throughout the system for high-speed deterministic data flows, while PCI-based communications is used for "soft real-time" payload data and control. In addition, the PRO-3100 and the PRO-3500 boards are network enabled, allowing the SDR-3000 to easily operate as part of a larger system.

Application framework

A basic application was developed to demonstrate the capability of the SDR-3000 to operate independently or as part of a

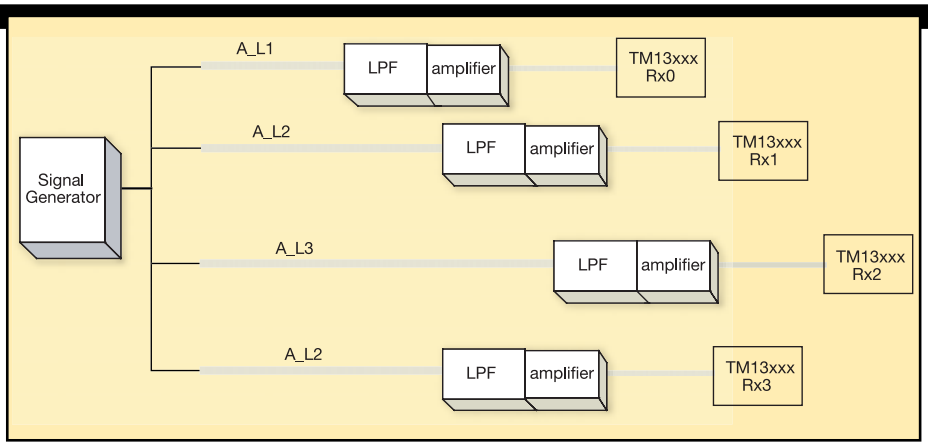

Figure 4. DFDS(S) system configuration.

larger DF system. This application enabled the user to:

- receive and process inputs coherently from four high-frequency (HF) antennae channels;
- calibrate the input channels;
- display a composite power spectrum on a host computer;
- select a signal of interest (SOI) from a spectral display; and
- calculate the AOA and display it on a compass rose.

To provide portability across OS and operation in a distributed environment, a CORBA-enabled client-server design was used. A JAVA-based GUI acted as a client, and applications running on the MPC7410 processors acted as servers.

The design emphasis was to demonstrate the ability to use the SDR-3000 system as a direction finder and not to develop a sophisticated DF system in its entirety. Accordingly, greater attention is paid to implementing typical DF application tasks and meeting functional requirements.

Concept of operation

The current design uses the IF sampling conversion technique[1] using an ADC operating at 80 MSPS. Signal amplifiers and low-pass filter at the front end limit the operational frequency range between 10 MHz and 40 MHz.

For this application, the system used FFT cores from RF Engines Ltd. (RFEL), a technology partner that specializes in advanced FPGA IP cores. RFEL's FFT core, capable of servicing two simultaneous inputs, is used to channelize the data into 2.5 kHz wide channels. The digital filters in the FFT core have a cut-off frequency of 40 MHz and an 80% passband distributed evenly about the center frequency, giving an operating frequency range of 4 MHz to 36 MHz. The software and hardware characteristics modify the net frequency range to 10 MHz to 32 MHz. Sixty four MB of SDRAM memory is used to store channel data for about 2.5 seconds for each of the four inputs. The restriction of channel data storage memory to 64 Mbytes/input mandates that the signal of interest be on air for a minimum of 2.5 seconds in order to be detected and its AOA calculated. The same channel data is used to calculate the power spectral density (PSD) for each channel. Since the aim of this first stage is merely to identify a strong signal within the operating frequency range, the data corresponding to a channel is summed across the four inputs and then displayed. Spectral averaging is used to provide data at a rate sufficient to meet the display update rates.

The GUI displays the PSD data for the entire range (0 MHz to 40 MHZ). Since the spectral display is only capable of displaying 800 to 1000 horizontal pixels, each pixel represents data from multiple FFT bins or frequency channels, with the highest PSD value within the set of bins is displayed. The zoom feature on the GUI allows the user to increase the bin resolution on the display. During operation, the user selects a frequency of interest from the base spectral display, and continually zooms in until the display resolution is one bin per pixel.

Using the mouse, the user can then select the signal of interest, and the GUI transmits the pixel index of the selected SOI to a processor that uses the zoom-level and the display width to identify the bin number. A bin extractor, executing on the FPGA, extracts the in-phase (I) and quadrature (Q) data corresponding to the required bin for each of the four inputs and forwards it to an MPC7410 on the PRO-3500. This processor first calculates the phase for each input and then calculates the AOA of the SOI from the phase data for the individual inputs. This AOA value is then transmitted to the GUI for display on the compass rose.

Data flow in the basic DF application

Figure 1 shows the signal data path and processing blocks implemented in the DF application. Two of the user-programmable FPGAs are used to service sampled data from four antennae and buffer the data for bin extraction in their local SDRAMs. A bin extraction module within each FPGA calculates the bin number corresponding to the selected SOI and

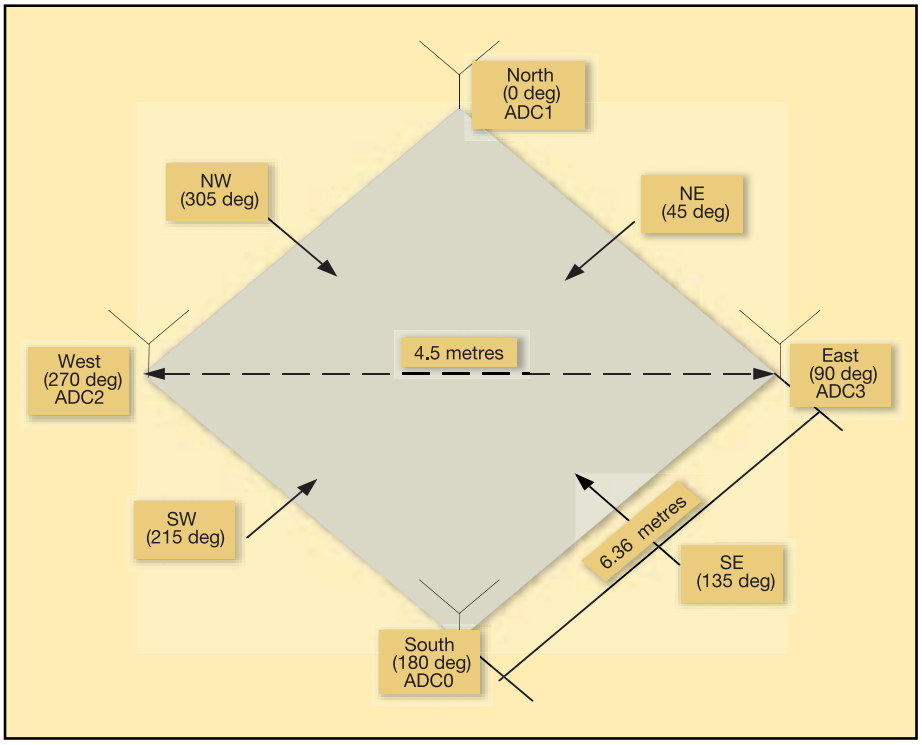

Figure 5. DFDS(T) system configuration.

Table 1.

User-programmable FPGA 1 (processes the spectral data)		
Device utilization summary:		
Number of External IOBs	184 out of 824	22%
Number of LOCed External IOBs	184 out of 184	100%
Number of MULT18X18s	4 out of 144	2%
Number of RAMB16s	45 out of 144	31%
Number of SLICEs	2862 out of 33792	8%
Number of BUFGMUXs	8 out of 16	50%
Number of DCMs	4 out of 12	33%
User-programmable FPGA 0 and 2 (implements the decimation/FFT core)		
Device utilization summary:		
Number of External IOBs	203 out of 824	24%
Number of LOCed External IOBs	203 out of 203	100%
Number of MULT18X18s	84 out of 144	58%
Number of RAMB16s	136 out of 144	94%
Number of SLICEs	14787 out of 33792	43%

extracts the FFT data for that bin from the delay buffer. A third user-programmable FPGA is used to calculate the PSD for spectral analysis. A spectral averaging module on this FPGA allows the user to control the rate at which the PSD data is sent to the PPC7410 processing nodes.

On the PRO-3500, one PPC7410 node is used to execute the DF and calibration algorithms while the second node is used to scale spectral data corresponding to the zoom level selected on the spectral display.

Task classification and mapping

Figure 2 shows the major software task breakdown for the basic DF application, data flow, task to processor mapping and the protocols used for intertask communication.

The graphical user interface (GUI) executes on a Windows PC and displays the frequency spectrum and bearing of the received signal. The GUI is written in JAVA and operates as a CORBA-enabled client interacting with the CORBA-enabled DF and spectral data servers on the PPC7410 nodes. Figure 3 shows a typical display on the GUI.

The DF and calibration processor (DFCP) task reads the selected bins of interest from the bin extraction module on the FPGA and calculates the AOA.

The DF and calibration server (DFCS) services DF and calibration-related data requests and configuration commands from the GUI. The spectral data server (SDS) services spectral display-related data requests and configuration commands.

The spectral data processor (SDP) continuously fetches spectral data from the user-programmable FPGAs and scales it for display using the most recent zoom index and display area pixel count.

The DF/calibration controller (DFCC) and spectral data controller (SDC) are used to configure the FPGA application resources in response to commands from the GUI. The controller blocks communicate with the servers over the network using TCP/IP sockets.

The processing blocks described earlier can be implemented on any of the available 405GP or PPC7410 processing nodes. Different placement schemes were evaluated, each requiring different intertask data communication methods. The primary metrics used for comparison included, but are not limited to, the following:

- data throughput (high, low);
- frequency of data transmission (frequent, continuous, intermittent on request);
- protocol or transport medium; and
- scope (within same processor, different processor on same board, processor on a different board).

The mapping scheme shown in Figure 2 provides judicious use of resources and allows implementation of data manipulation techniques that eliminate the need for unnecessary high-speed bulk data transfers without compromising system performance. Many different software partitioning and mapping schemes were evaluated. Of these, the current software task mapping scheme, shown in Figure 2, provides the maximum processor idle times and free bandwidth on the interboard and intraboard communication bus.

"Many different software partitioning and mapping schemes were evaluated. Of these, the current software task mapping scheme, shown in Figure 2, provides the maximum processor idle times and free bandwidth on the interboard and intraboard communication bus through judicious use of resources. This allows implementation of data manipulationtechniques that eliminate the need for unnecessary high-speed bulk data transfers without compromising system performance."

DF and calibration algorithms

The DFCP incorporates a DF algorithm suitable for use with fixed antenna systems. It uses the phase of the incident signal at an antenna to detect the AOA. This simple algorithm calibrates the DF data but does not compensate for multipath effects. The DFCP also has the ability to average a user-defined number of sequential DF measurements before passing it on to the host.

The data from the four antennae traverse through different cables and hardware components (filters, amplifiers, ADC/DAC) prior to the DDC stage. The system must be calibrated to compensate for relative phase and amplitude imbalance caused by physical and electrical properties of the cables and hardware components. Calibration is mandatory the first time the system is powered on and must be repeated whenever there is a change in hardware, cables, antenna configuration or at extended time periods to account for component degradation due to aging. The current design assumes that calibration and DF processing are mutually exclusive functions and consequently uses a different set of firmware images for the user-programmable FPGAs for performing off-line calibration. Each time the DF algorithm is executed on a selected data bin, run-time calibration uses the calibration data to compensate for inherent differences introduced by the different components in the input data paths.

Design verification

Two configurations, varying only in the way the received signals are fed to the SDR-3000 system, were used to verify the system design. The DFDS(S) configuration, shown in Figure 4, uses a single signal source but different cable lengths to simulate the phase difference between signal incidents at the different antennae.

DFDS(T) configuration, uses antennae arranged in a rectangular array as shown in Figure 5, to receive CB radio frequency signals emitted by a COTS CB radio. The four COTS HF omni antennae were mounted at optimum spacing for operation up to 32 MHz.

All functions and operational features of the DFDS(S) were satisfactorily tested in the presence of the end user. Cables of pre-calculated lengths were used to simulate signal source at cardinal points (N, S, E, W) and at ±45 from them (NE, NW, SE, SW). Testing covered the entire operating frequency range from

10 MHz to 32 MHz. The simulated system displayed AOA to an accuracy of ±5.

DFDS(T) was tested in a parking lot setting exhibiting a minimal level of multipath. The spectral display successfully showed peaks at frequencies equal to that of received signals. The AOA display for this case exhibited ±25 accuracy. This ambiguity was expected, as the test environment was not completely multipath-free and reflection-free while the implemented DF algorithm assumed such ideal conditions.

Capability demonstration and system acceptance

All system requirements were either met or exceeded. Despite using a simple DF algorithm, the DFDS(T) was able to display AOA to within ±5 when tested out in the open. System resource usage statistics clearly depicted that sufficient resources were available to expand the system to include more inputs or use more complicated algorithms. The delivered system, configured as a stand-alone unit, could be easily transported, set up and operated.

Resource usage statistics for the user-programmable FPGAs are shown in Table 1.

The majority of user-programmable FPGA 0 and 2 slices are used up by RFEL's FFT core with very little usage by application processing code. Most of FPGA 1 is still unused. This leaves an entire user-programmable FPGA and associated SDRAM for any additional processing and delay buffer implementation. Depending on the size of the FFT/DDC core used, more optimal designs can be developed to implement additional processing on these FPGAs, such as IFFT, mod/demod or decoding.

In the current design, the 7410s continually extracts bin data and spectral data from the PRO-3100. With this processing, the 7410 used for DF and calibration is lightly loaded while the 7410 used for spectral data processing is occupied most of the time. However, with minor software modifications that allow the two processors to process data only upon request from the user, the processor idle time can be considerably increased.

Conclusion

A direction finding system demands a high degree of functionality. Today, COTS platforms are sophisticated enough for DF applications as successfully demonstrated with the *flex*Comm SDR-3000. The platform provided the necessary processing power, speed, and design and operational flexibility required to perform the complex functions. Furthermore, the team was able to design, develop, procure equipment, and test and validate the DF system on this platform in only four months.

The system can be expanded to accommodate more input channels by adding additional transition module and PRO-3100 pairs. Channel data must be synchronized across all the transition modules. To handle the increased data volume and channelization requirements, additional FPGA cores and software tasks may need to be incorporated and the software task mapping either re-partitioned or new software tasks added.

This platform can also be used to efficiently perform multiple simultaneous DFs on multiple signals of interest. This may require the use of different channelization techniques and probably, implementation of wideband and narrowband channelization. As an example, a user could implement a beam former to scan a wide frequency range, automatically detect signals with power above a preset threshold, and then use narrowband channelization to extract the signal of interest. RFD

References

1. L. Pucker, "Channelization Techniques for Software-Defined Radio," Proceedings of the SDR Forum Technical Conference, November 2003.

ABOUT THE AUTHOR

Jyothsna Rajan is an Applications Engineer with Spectrum Signal Processing Inc. in Burnaby, B.C., Canada. With more than 13 years of experience, her areas of expertise include digital signal processing and real-time embedded systems software for distributed communications and control. She received her Bachelor of Engineering from Karnataka Regional Engineering College in India and a Masters in Electrical Engineering, specializing in DSP/communications from the University of Saskatchewan in Saskatoon, Canada. Prior to joining Spectrum, she worked under the aegis of Telecommunications Research Lab at Saskatoon developing a portable wideband microseismic system for fault location in underground mines.