Data Converters and DSPs Getting Closer to Sensors

As the data converters used in military applications must operate faster and at greater resolution, the digital domain is moving closer to the antenna/sensor array. In response, two different approaches to DSP architectures are being used, sometimes in combination: moving the data faster, and/or moving the DSP hardware.

By Scott Hames,Interactive Circuits & Systems

The data converters used in military applications, such as defense communications and radar systems, must operate at ever-increasing speeds and higher resolutions. As a result, the digital domain is encroaching on the antenna or sensor array. As this happens, military system designers face serious challenges when trying to move signal data in ever-increasing volumes.

This situation is driving two different approaches to DSP architectures. One approach is to employ the latest innovations in bus technology to increase the speed/bandwidth of data movement, that is, moving the data faster. The other is to move some or all of the DSP hardware toward the source of the data, a “point of load” approach; that is, moving the data processing. The two techniques are not mutually exclusive and some systems will require both.

Moving the Data Faster

System architectures based on traditional back-end DSP and front-end A/D and/or D/A converters can still be effective in many situations, as long as the ability to move the data among processing elements is at an appropriately high level. The evolution of the PCI bus is a prime example of this.

But there are caveats to using a shared bus, such as PCI, as the backbone of a data acquisition system, especially when it is built up from off-the-shelf components such as bridges, switches and endpoints. The PCI endpoint, or local bus interface, can often make

or break the performance of a data acquisition system. Bus masters must balance the need to stream data at high rates for long durations, essential to efficient throughput, with the need to share the bus among multiple users. This usually means having the ability to disable arbitration schemes when the need for maximum throughput arises.

Unfortunately, most PCI master devices seem to be designed either for maximum throughput or for fair arbitration, but not both. Arbitration among multiple bus users can substantially reduce the aggregate bandwidth from the theoretical maximums, usually by artificially limiting the amount of time any single user can hold the bus. Valuable bus cycles are then wasted on multiple bus requests and retries.

A brute force solution to the need for guaranteed PCI bandwidth has been the implementation of multiple segment architectures, which provide more data buses between nodes such as data producers, data consumers and processing engines. Often, the number of processing engines is increased as well.

There are limitations associated with increases in clock speed and bus width, and skew between data lines on a wide bus such as PCI can destroy timing margins. Because of this, it is now generally accepted that serial buses running at high speed will provide better throughput for future data acquisition systems.

Much of the PCI community is now in the midst of a migration to PCI Express. Despite the improvement over parallel PCI, this may not be sufficient for future high-speed data acquisition systems. Analog to digital converters now routinely exceed 1 GHz, and even if the resolution is only eight bits, a single channel will saturate a four-lane PCI Express link. Clearly, the strategy of moving large amounts of raw data has reached its limit, at least so far as today’s available technology is concerned.

Moving the DSP Hardware

Given that, over time, the converters have moved closer to the sensor, why shouldn’t the DSP hardware, or at least some of it, do the same? There are real opportunities for “point of acquisition” DSP hardware to extract the meaningful data from the extraneous data before moving it to a separate module for processing. In recent years, FPGAs have emerged as the de facto standard for custom hardware (Figure 1).

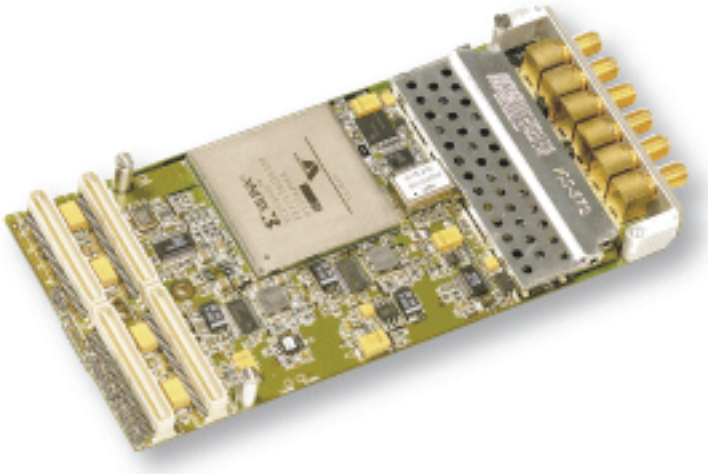

Figure 1

One approach to the faster speeds and higher resolutions in military data converters is moving some or all of the DSP hardware toward the source of the data. The latest generation of A/D and D/A converter modules, such as the ICS-572 from Interactive Circuits & Systems, features FPGA technology for pre-processing.

Most of the recently introduced data acquisition modules (Figure 2) feature some kind of onboard processing resource, either an ASIC or an FPGA. Combined with the adoption of high-speed serial bus technologies such as those previously discussed, these modules have allowed the creation of a new architecture for sensor processing systems, one that changes the meaning of the term “embedded processing.”

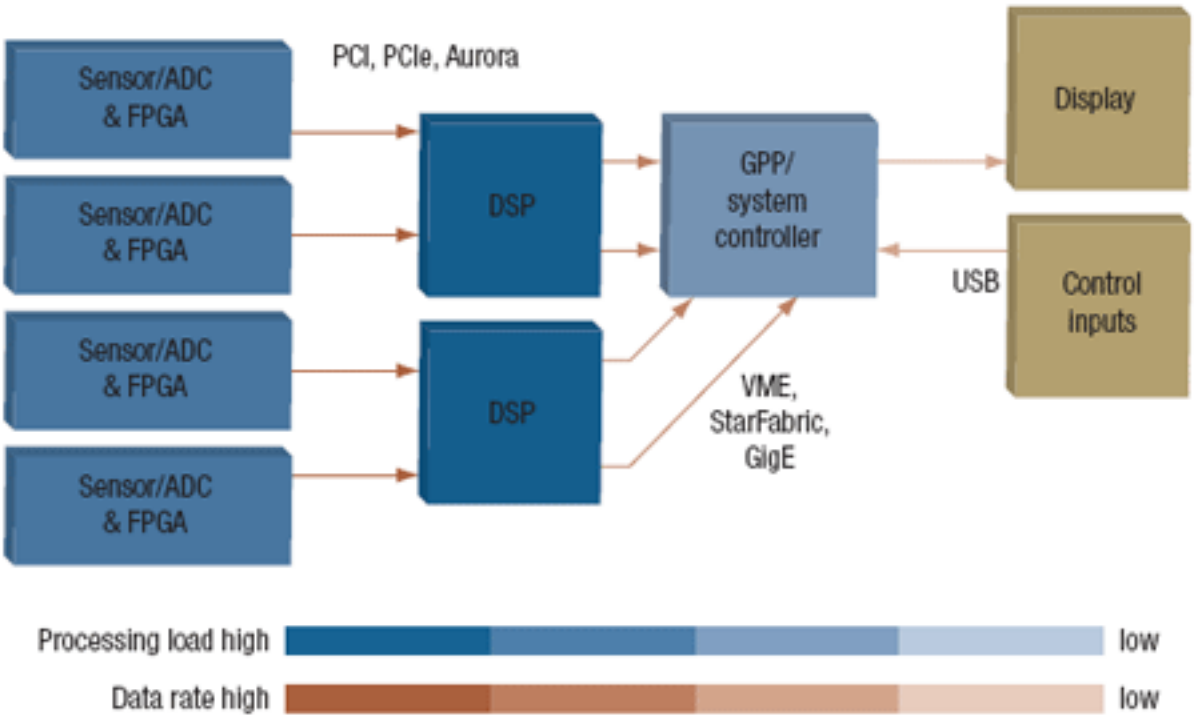

Figure 2

Block diagram of a modern data acquisition system in which the front end incorporates onboard FPGA resources to pre-process data.

One effect of the implementation of higher-speed A/D and D/A converter components is that the digital processing domain now encompasses operations that were previously executed using analog components and techniques. This is especially true of software defined radio applications, where the entire IF-to-baseband conversion stage has been replaced with DSP techniques and hardware (Figure 3).

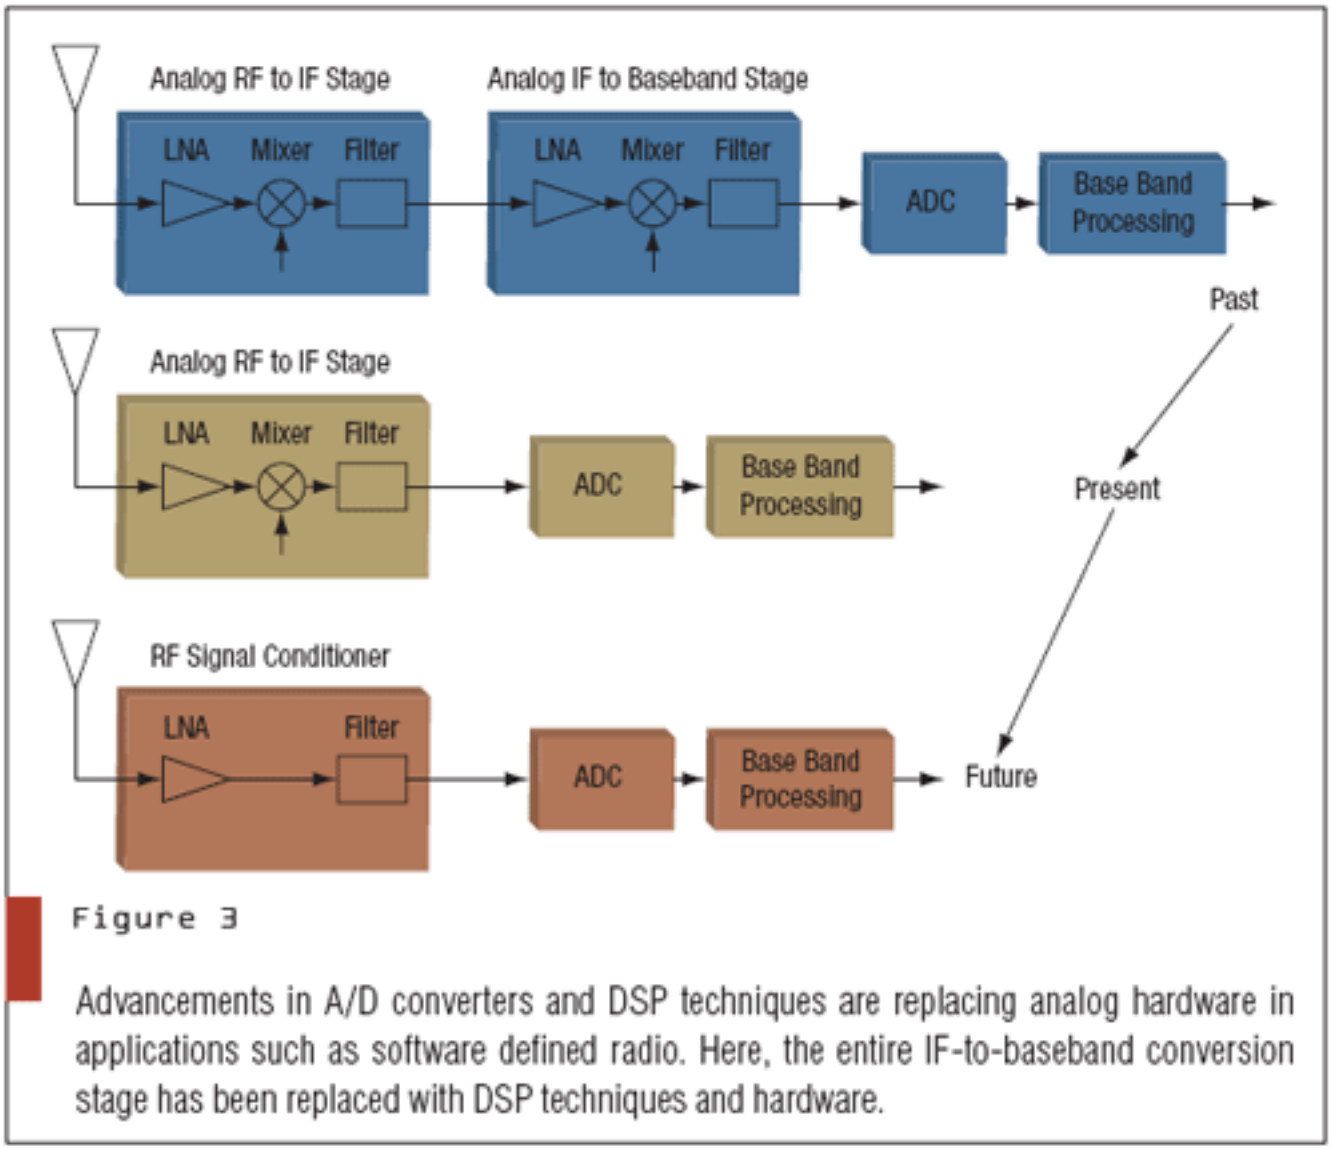

Figure 3

Advancements in A/D converters and DSP techniques are replacing analog hardware in applications such as software defined radio. Here, the entire IF-to-baseband conversion stage has been replaced with DSP techniques and hardware.

Isolating Signals of Interest

A key technique for reducing the data rate is to focus on the signals of interest and eliminate redundant data. Many applications—including radar, communications and high-frequency sonar—use carrier frequencies much higher than the actual bandwidths of interest. If the band of interest can be shifted to a lower carrier frequency, ideally all the way to DC, the data rate required to represent the signals digitally can be reduced substantially, regardless of the techniques used to downconvert the signal.

The Nyquist sampling theorem states that the data rate required to represent the signal is twice the highest frequency of interest. A 100 kHz band located on a 70 MHz carrier requires a data rate of about 140.1 MHz to meet the Nyquist criteria, but the same 100 kHz band requires only a 200 KHz data rate if it is located at baseband (DC) (Figure 4).

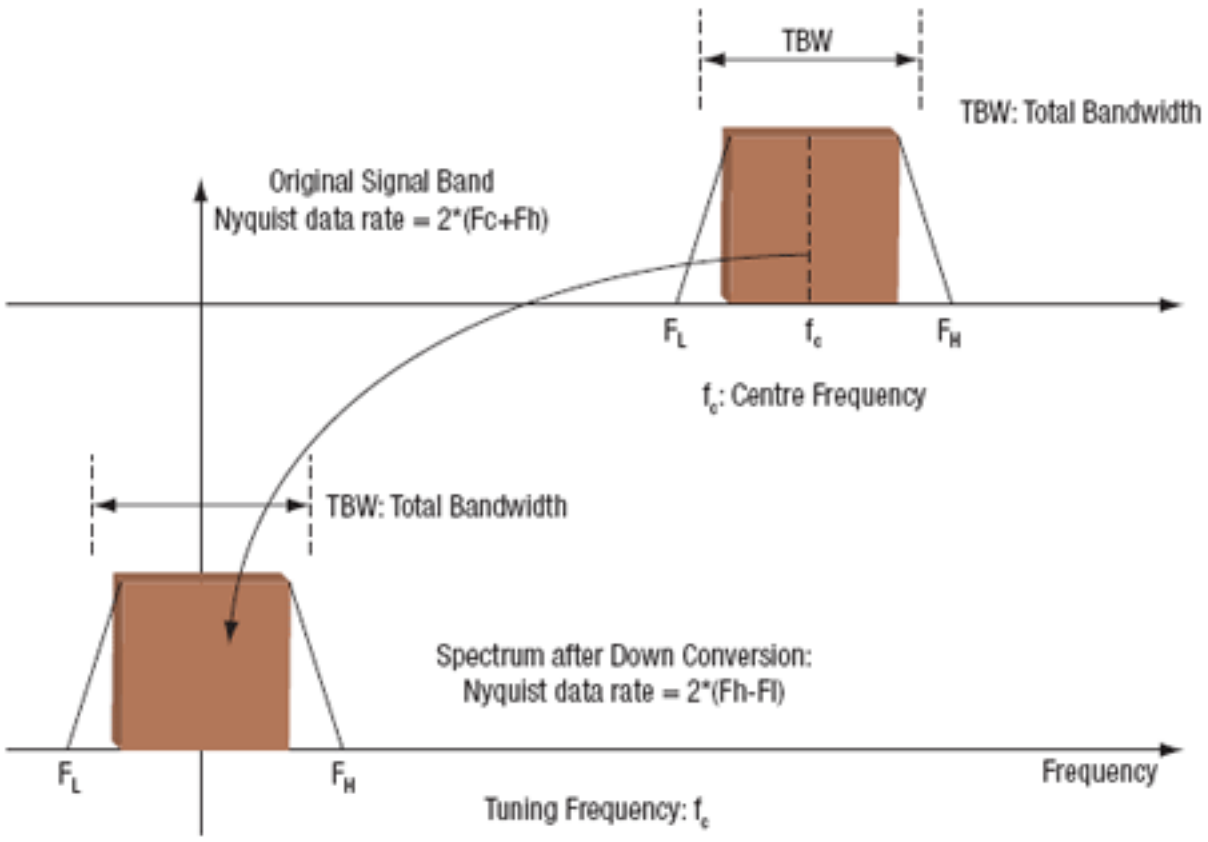

Figure 4

In some defense applications such as radar and high-frequency sonar, which use carrier frequencies higher than the actual bandwidths of interest, band-shifting techniques allow reduced data rates.

Until the introduction of high-speed, high-resolution A/D converters, band-shifting functions such as downconversion were done using analog components. But there are huge benefits to performing operations such as downconversion, or band shifting, using digital techniques. Even the best analog components add noise to signals. Compared to their analog counterparts, digital circuitry such as numerically controlled oscillators (NCOs) provide an unmatchable combination of stability and flexibility. In addition, where digital processing techniques can allow exact execution of mathematical functions, analog circuitry often provides the closest approximation of a transfer function rather than what was actually required.

Design Considerations

Despite these benefits, serious risks exist when moving the boundaries of the digital domain in this manner. When replacing analog hardware with a DSP technique, it is

imperative that the operation of the analog circuitry be understood exactly, and that the DSP hardware and techniques replacing the analog circuitry provide identical functionality.

This is particularly true of band shifting operations such as upconversion and downconversion. Careful application of DSP hardware—general-purpose, FPGA and ASIC—can greatly reduce the amount of data to be moved back to a host processor. But, despite its appeal, this approach requires careful component selection to ensure that the desired DSP functionality can be properly executed at the converter I/Os. Multi-channel and multi-event synchronized applications have special requirements that must be respected in order to make this approach work.

Digital upconversion algorithms are excellent examples of the need for proper execution in hardware. There are many interpolating D/A converters on the market that offer extremely high conversion rates and reasonable data input rates. However, these may lack the ability to properly phase synchronize multiple events—an absolute prerequisite for pulsed operations such as coherent radar—or multiple channels, required for phase-sensitive operations such as beamforming. Mistakes in component selection can prevent proper system operation, even if the DSP algorithm is well thought out and carefully designed.

Digital downconversion (DDC) applications often have an extra requirement as well: the need to phase lock, or “track,” an incoming signal for the purposes of carrier recovery. Real-time adjustments to the NCO operation will be required to fine-tune the operating frequency so it exactly matches that of the asynchronous transmit system. This will usually require the DDC to be implemented in an FPGA, which will allow the designer to optimize the NCO performance for the target system.

The need to process rapidly increasing amounts of sensor-captured data is at a level that stretches the capabilities of the most sophisticated hardware. One approach is to try to move the data faster, but theoretical speeds are not always achievable in practice. An alternative, or complementary approach is to pre-process the data, identifying and isolating only data that is of interest and discarding the remainder, which results in a potentially significant reduction in the amount of data to be moved. Whichever approach is followed, there are a number of factors to be taken into account in the design of the solution. However, if the tradeoffs are appropriately made and their implications precisely understood, solutions can be designed that will deliver the required levels of performance.

Interactive Circuits & Systems,
Part of Radstone Embedded Computing

Ottawa, Canada.
(613) 749-9241.
[www.ics-ltd.com].