Sorting Through the Comms Infrastructure Choices for Software Defined Radio

A variety of bus and interconnect schemes are available to satisfy the demands of distributed architecture SDR systems

by Lee Pucker, Technical Lead,
Vivian Ward, Systems Engineering Manager,
and Frank Van Hooft, System Architect,
Spectrum Signal Processing

The choice of communications infrastructure is a critical part of any software defined radio (SDR) implementation. SDR enables a single reconfigurable platform to support multiple, often disparate waveforms, with a waveform being defined as an air interface standard and its associated higher-level protocols. When SDR technology is used for wireless infrastructure systems, such as cellular base stations or military tactical communications gateways, the SDR platform may need to support many tens or hundreds of simultaneous users employing a combination of both wideband and narrowband waveform applications.

To meet those requirements, a distributed transceiver architecture is often employed, consisting of multiple, heterogeneous, signal processing devices communicating over diverse parallel channels. That calls for a communications infrastructure that can accommodate this heterogeneous processing environment while maintaining the divergent data flow requirements of the waveform applications. An examination of the various available interconnect technologies helps to narrow down the options.

Architectural Requirements

The communications infrastructure in a distributed transceiver system can be segmented into five different types of data flows, each with its own requirements (see sidebar for more on this). These categories can be summarized as follows:

- High-speed input/output channels for either digitized IF or baseband data
- Inter-processor communications, with support for logical channels, for distributed processing
- Board-to-board communications for architectures that extend beyond a single printed circuit board
- Payload and control data channels that do not interfere with the primary data paths
- Chassis-to-chassis communications

A variety of communications standards are available that address these five communications types. In general, these standards break down into two main categories: legacy architectures and packet-switched architectures.

Legacy Architectures

Legacy architectures represent communications standards that have been around for several years. These include bus-based communications architectures such as PCI, and circuit-switched architectures such as Raceway.

PCI

The basic PCI bus standard is a 32-bit multiplexed address and data bus running at 33 MHz, for a theoretical aggregate bandwidth of 132 Mbytes/s. Extensions to the original standard allow the bus to run at 64-bits and 66 MHz for a maximum bandwidth of 528 Mbytes/s. With these extensions, PCI bus transfers limit the peak transfer rate through the PCI board connector to approximately 48 Mbits/s per pin. Yet another new extension to PCI, PCI-X, allows the bus to run even faster; up to 133 MHz at 64 bits for a peak data rate of just over 1 Gbyte/s peak. PCI-X also provides some new transaction types that allow packet-like operations.

PCI requires a controller to enumerate the bus, but once this is done, any device is free to transfer data. There is no concept of data types and error handling is minimal. Higher-level functions such as packet handling, which is required for logical channel support, are

Comms Infrastructure Requirements for a Distributed Transceiver Architecture

The core of a software defined radio is the digital transceiver subsystem. On the receive side, this subsystem receives either digitized IF or baseband signals, extracts user channels from these signals in the channelizer and then demodulates and decodes each user channel to extract the payload data (Figure 1) This process is reversed on the transmit side, with payload data being encoded and modulated in the channel processor and then inserted into the output signal by the channelizer for transmission. In a distributed transceiver architecture, the channelization and channel processing functions are distributed across multiple signal processing elements, with a single channelizer often supporting multiple channel processors.

The definition of the communications infrastructure for a distributed transceiver architecture begins with the interface to the RF transceiver. This input/output channel typically represents the highest overall bandwidth in the system, and must guarantee low latency and deterministic performance. Often each digitized channel must be distributed to multiple channelizer processing elements, especially for applications supporting adaptive beam forming, and so the communication infrastructure must support broadcast capabilities.

Similarly, separate payload and control input/output data paths must be provided between each of the channel processors and the rest of the SDR platform. Typically, the bandwidth per user channel for these paths is relatively low when compared to the dataflow requirements of the rest of the system, so multiple payload or control data channels will often share a common physical communication structure through the use of logical channels.

A logical channel is software representation of a connection between two functional components, with multiple logical channels coexisting over a common physical communications structure. The composite data rate of this physical link must be sufficient to support not only the aggregate data rate required for

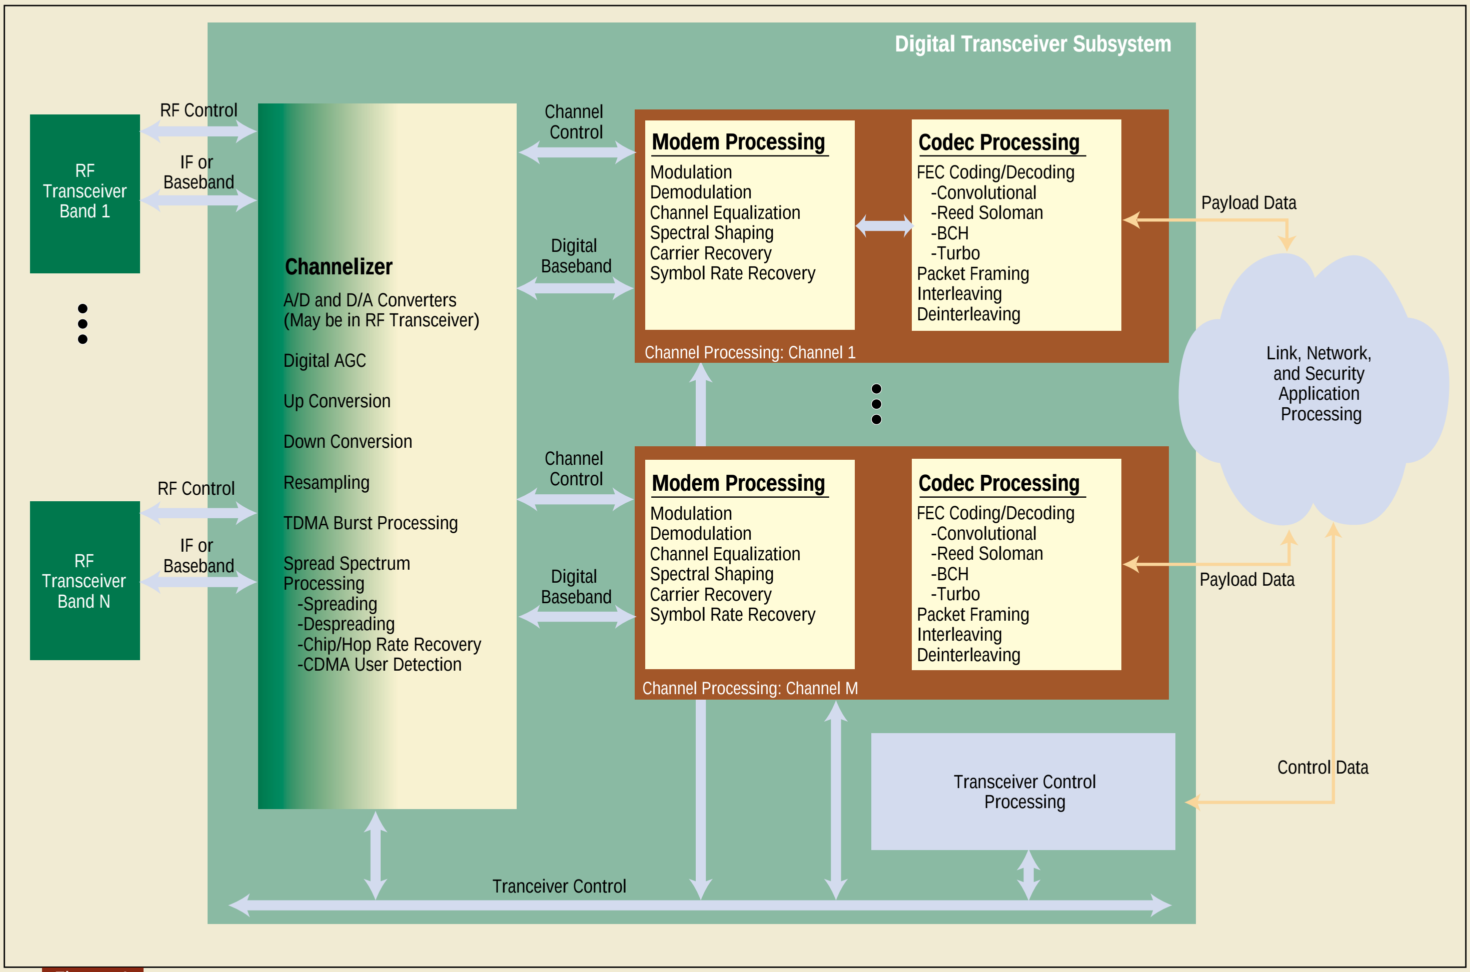

Figure 1 In a distributed transceiver architecture as shown here, the channelization and channel processing functions are distributed across multiple signal processing elements, with a single channelizer supporting multiple channel processors.

the combined data paths, but also for the overhead required to route the channel within the communications infrastructure while guaranteeing that the latency and determinism requirements of the data paths are maintained.

Finally, an inter-processor communications structure must be provided to support the dataflow requirements between the processing elements. Since a single processing element will often host multiple simultaneous digital radio functions, each with independent data and control channels, the inter-processor communications structure must provide a logical channel interface for each processing element.

Specifying a Hardware Platform

The requirements for the communications infrastructure are further refined when the selected processing elements are mapped to a specific hardware platform. For a simple software defined radio architecture, supporting at most a few channels, all of the signal processing elements may be mapped to a single printed circuit board, with additional channels supported by combining these boards in a common chassis. For these types of systems, the primary constraint on the communications infrastructure is in the input/output channels, which are limited by the number and type of pins and the supported data rates of the board connectors.

For more complicated architectures, however, where the digital transceiver may be distributed across several boards, the inter-processor communications structure within each board must be extended to provide board-to-board communications that maintain the necessary data flows between processing elements. These communications paths include both base card-to-mezzanine and slot-to-slot backplane communications. For high-availability systems, the slot-to-slot communication infrastructure is further constrained in that it cannot result in a single point of failure for the overall system and may require support for hot-swap.

The communications infrastructure in a distributed transceiver architecture may also extend outside of the chassis housing the digital transceiver subsystem. For example, several digital transceivers housed in a single chassis may share a common high-speed data link into and out of the chassis for payload data. When I/O communications extend beyond the chassis, the communications type must provide adequate drive to transmit and receive data over the required chassis-to-chassis distance, and must support adequate bandwidth for aggregation of the various data types.

generally implementation-specific. The PCI bus is not terminated; bus timings are kept to very tight tolerances limiting PCI to relatively short distances.

Since PCI is a shared bus, it can easily exhibit non-deterministic latencies. The PCI bus also limits system reliability since a single faulty card can capture the bus, resulting in the bus being unable to function until the faulty card is removed. As there can only be a single controller on the bus, the controller also acts as a single point of failure. These issues limit the usefulness of the PCI bus to primarily payload and control data paths that are not mission-critical.

Raceway and Race++

Raceway is a circuit-switched communications architecture that assumes a node-to-node or classic “star” topology. Each Raceway endpoint device connects directly to a Raceway crossbar, with the crossbars routing data from endpoint to endpoint. Originally Raceway utilized a 32-bit data path operating at 40 MHz for a maximum of 160 Mbytes/s per link.

Raceway Interlink++ (usually called Race++) increases the clock speed to 66.667 MHz for a maximum of approximately 267 Mbytes/s per link. Depending upon the number of node-to-crossbar and crossbar-to-crossbar links in the system, the aggregate system data rate can grow into the gigabytes-per-second range. Race++ uses 38 wires per link, resulting in peak transfer rate through a board connector of 56 Mbits/s per pin.

Raceway provides a data link layer with support for priorities and routing capabilities in the crossbar switches, improving determinism and latency. The crossbars provide some isolation between boards in the system, which can improve system reliability. However, for board-to-board communications across a backplane, Raceway typically uses an active backplane architecture consisting of one or more active crossbar devices. The failure of a crossbar could require shutting down the entire system to replace the backplane. This may limit the use of Raceway in most high-availability systems.

Packet-Switched Architectures

The requirement to support large numbers of logical channels within the communications infrastructure is causing a move away from legacy bus-based and circuit-switched architectures toward packet-based switched communications fabrics. Like Raceway, the endpoint devices in a packet-switched architecture connect to a switch. In a packet-switched architecture, however, the data is routed based on a destination address embedded in each transmitted packet, with the switched fabric providing a transport layer capable of end-to-end routing and multiple links. As such, packet-switched communications infrastructure technologies provide for greater efficiency in supporting logical channels. Packet-switched communications fabrics may be divided by physical layer implementations into parallel and serial architectures, each of which have a common set of strengths and weaknesses.

Parallel Packet-Switched Technologies

Parallel packet-switched architectures are distinguished by multiple data lines running in parallel with a separate data clock line. They typically support high link bandwidth with low protocol overhead, making them very attractive for distributed transceiver architectures. Additionally, they often have scalable bus widths and clock rates that allow them to meet both current and future system requirements.

Parallel RapidIO

RapidIO is a packet-switched protocol developed specifically for embedded systems by the RapidIO Trade Association to emphasize high-bandwidth, low latency communications. RapidIO utilizes low voltage differential signaling (LVDS) with a source-synchronous protocol to provide a full-duplex communications link at data widths of either 8 or 16 bits. Data is signaled on both edges of a transmitted clock at rates of 250 MHz, 500 MHz or 1 GHz. RapidIO 8 uses 40 wires to provide a peak transfer rate of 400 Mbits/s per pin for both the transmit and receive links, or 800 Mbits/s per pin aggregate. RapidIO 16 uses 76 wires to provide twice RapidIO 8's throughput.

RapidIO was designed with a relatively small protocol stack to allow for simple endpoint devices. RapidIO efficiency is then a function of packet size, data width and whether the system in question includes the protocol's logical layer. Efficiencies of greater than 90% are possible with large packet sizes. Its low protocol overhead and high bandwidth make RapidIO very appealing for inter-processor communications, but the relatively high pin count makes it less useful for board-to-board links.

HyperTransport

HyperTransport is a bus-like fabric originally defined by AMD. This is a source-synchronous protocol that uses differential signaling to support data widths of 4, 8, 16 and 32 bits at signal rates of up to

Standard	*Pros*	*Cons*	*Best used for ...*
PCI/PCI-X	Low cost. Simple.	Poor QoS. Low BW/pin	Payload and control data, within the chassis.
RACEway/ RACE++	Simple. Scalable.	Moderate QoS. Low BW/pin.	Interprocessor and board to board communications.
Parallel RapidIO	High BW/pin. Simple Protocol.	High pin count.	High Speed I/O, Interprocessor Communications, Board to Board Communications.
Hyper-Transport	High BW/pin.	High pin count. Market positioning.	Interprocessor Communications, Board to Board Communications.
InfiniBand	High BW/pin. Powerful protocols.	Large protocol stack. High power consumption.	High Speed I/O, Payload and Control Data outside the chassis.
Serial RapidIO	High BW/pin. Few pins. Simple protocol.	High power consumption.	Digitized IF, payload and control data for board to board communications.
Switched Ethernet	Common & well understood.	Limited bandwidth, latency.	Payload and Control Data, inter-board & inter-chassis.

Table 1 A combination of RapidIO and Serial RapidIO offers the best mix of support for inter-board and inter-processor communications. Switched Ethernet and PCI are the best choices for slower speed data such as payload data and control. The three serial packet- switched fabrics are suited for communications that extend beyond the chassis.

1600 Mbits/s per differential wire pair. AMD has positioned HyperTransport primarily as a front-side bus for chip-to-chip communications, with support provided for bridging HyperTransport to inter-board communications fabrics. Like RapidIO, HyperTransport utilizes 40 pins for an 8-bit link, providing a peak transfer rate through a board connector of 320 Mbits/s per pin.

The Deficiencies of Parallel Packet-Switched Technologies for Distributed Transceiver Architectures

Limited I/O pins and signal skews become significant considerations when signals are routed across a backplane. Unfortunately, parallel-based connections do not address either of these factors well. The high pin counts severely limit the number of links that may be connected to a board. In addition, parallel protocols are susceptible to differential timing skew between data signals. Minimizing such skews across a PCB is a reasonable task but becomes much more difficult when routing through a backplane and associated connectors. Ultimately, high bandwidth per pin across inter-board connections can only be obtained by obviating the need for clock/data skew control, which leads to serial packet-switched protocols.

Serial Packet-Switched Technologies

Serial packet-switched architectures have a physical layer that was specifically designed to cope with inter-board skew issues. These architectures use a single differential transmit and receive pair to provide the physical layer connection between the end points. The clock is embedded into the data signal, typically utilizing an encoding scheme such as 8B/10B, to ensure that the clock can be recovered in the serial receiver. Serial packet-switched technologies also typically support transmission over a longer range, making them ideal not only for board-to-board communications, but also for chassis-to-chassis links.

Switched Ethernet

With the recent development of the PICMG 2.16 standard, switched Ethernet is now available as a backplane communications fabric that can be used in distributed transceiver architectures. Its primary benefit is that the physical connections and protocol are widely used and understood. Although the protocol stack results in relatively high latency communications, the determinism allowed by multiple dedicated links is invaluable when compared with conventional Ethernet. The recent extension of Ethernet to gigabit speeds allows it to remain competitive, however even at gigabit speeds, switched Ethernet supports less than 100 Mbytes/s per link for sustained transfers, which makes it unsuitable for many primary data paths. Switched Ethernet, however, is an outstanding choice for control and payload data because it can connect seamlessly to other systems.

Serial RapidIO

Serial RapidIO is a serial adaptation of the parallel RapidIO protocol. This technology supports serial links at 1.25, 2.5 and 3.125 Gbits/s. However, current connector technology limits Serial RapidIO to the 1.25 Gbit/s rate over many standard connector types, including the 2 mm connectors used in CompactPCI backplanes. Higher data rates are supported by combining four serial transmit and receive pairs (1X links) together into a single 4X link.

Serial RapidIO is not an attractive choice for most intra-board links, since the high clock speeds consume significantly more power than parallel alternatives. However, the protocol overlap with RapidIO allows bridge chips to easily connect both parallel and serial links. This overlap suggests an approach of using parallel RapidIO for intra-board links and serial RapidIO for inter-board links, extracting the benefits of each approach in the appropriate domain.

InfiniBand

InfiniBand developed from a merger of Intel's Next Generation I/O (NGIO) and Future I/O. It is intended to replace Ethernet and Fiber Channel as the high-bandwidth systems interconnect for computing clusters and server I/O. The strengths of InfiniBand come from its high speed and the richness and flexibility of its protocol stack. InfiniBand uses differential signaling at 2.5 Gbits/s with 1X, 4X or 12X data links supported to provide raw bandwidths of 250 Mbytes/s, 1 Gbyte/s and 3 Gbytes/s. The 2.5 Gbit/s rate is too fast for most existing backplane connectors, and so work is under way to develop a new connector to allow these speeds to be used reliably in a standard architecture.

InfiniBand was designed to communicate over fairly long distances, which at 2.5 Gbits/s requires a significant amount of power. In addition, the complexity of the InfiniBand protocol requires significant resources, in both logic and software. These elements make InfiniBand, in general, unsuitable for use as an embedded fabric for low latency communications. InfiniBand is therefore limited, in a distributed transceiver architecture, primarily to high-speed I/O for communications beyond the chassis.

Future Alternatives

Several new packet-switched communications technologies are under development that could have significant impact on future distributed transceiver architectures. Chief among these is 3GIO. 3GIO has been heavily promoted by Intel as the successor to PCI. Like many of the new high-speed fabrics described in this article, 3GIO uses multiple LVDS transmit and receive pairs to achieve high levels of point-to-point bandwidth.

Each LVDS pair is run at an initial 2.5 Gbits/s (derived from Infiniband), with 8B/10B encoding, with the resulting full-duplex link called a "lane". Up to 32 lanes can be combined to increase the throughput of the resulting 3GIO link. 3GIO is currently in the specification drafting stage. It is expected that an initial spec will be available for review in 2002, with first silicon supporting 3GIO anticipated in 2003.

For today and the near future, a combination of RapidIO and Serial RapidIO seems to provide the best mix of support for inter-board and inter-processor communications (Table 1), primarily because these two technologies share a common, efficient, protocol stack. Switched Ethernet and PCI are the best choices for slower speed data such as payload data and control, and all three serial packet-switched fabrics make good choices for communications that extend beyond the chassis. ◢

Spectrum Signal Processing
Burnaby, B.C., Canada.
(604) 421.5422.
[www.spectrumsignal.com].