

WHICH ARCHITECTURAL MODEL SHOULD I ADOPT IN MY WIRELESS COMMUNICATIONS SYSTEM: WELL, ONCE AGAIN, IT DEPENDS...

Lee Pucker, Spectrum Signal Processing, Inc.

In the September “Trends in DSP,” I explored the selection of signal processing technologies for use in wireless applications [1]. Many of the technologies discussed have an inherent ability to process the signals associated with multiple intermediate frequency (IF) channels on a single device while maintaining real-time performance. This capability raises an interesting question in creating wireless infrastructure systems supporting multiple concurrent channels: “Even though a processing resource can be shared by multiple channels, should it be?” For simple air interfaces, where the processing of these channels may be entirely contained within a single device, the answer is likely “of course.” However, in many wireless systems the supported air interfaces are much more complex, requiring a combination of signal processing technologies to optimize the system for size, weight, power, and cost. The use of shared resource processing in this class of systems requires a number of advanced capabilities within the wireless architecture. These capabilities may increase overall system cost, and as such the decision to support shared resource processing can only be made by weighing the advantages accrued in sharing the processing resources against the

cost impact associated with developing and deploying these advanced capabilities.

To explore these trade-offs, let us begin by examining the dedicated resource architecture inherent in many wireless infrastructure systems. This type of architecture dedicates radio frequency (RF) and signal processing resources within the wireless system on a per channel basis, with field programmable gate arrays (FPGAs) or application-specific standard processors (ASSPs) generally used for front-end signal processing such as network synchronization and control, multicarrier processing, and chip rate processing [2–6]. Back-end transmit and receive channel processing, including the modem and channel coding functions, are then typically supported by either a general-purpose processor or digital signal processor (GPP or DSP). Multiple channels in a system following the dedicated resource paradigm are supported by duplicate RF/signal processing subsystems as illustrated in Fig. 1.

The size, weight, power and cost of a signal processing subsystem following a dedicated resource architecture are generally managed by scaling the processing resources and

(Continued on page S6)

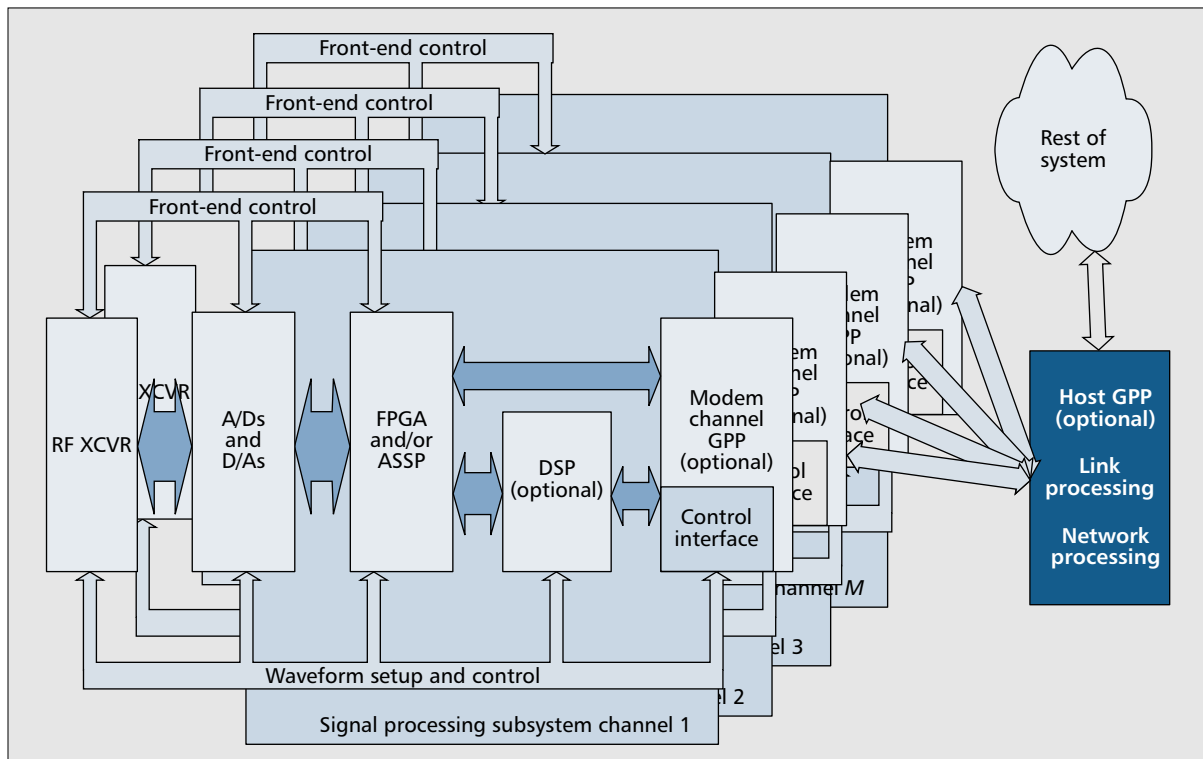


FIGURE 1. Typical front-end radio architecture for a traditional multichannel wireless infrastructure system following a dedicated resource model.

(Continued from page S4)

supporting infrastructure down to the minimal level necessary to provide for the target air interface standards. As technology advances, however, the performance of commercial off-the-shelf signal processing devices has increased disproportionately with the functional requirements of many of these air interfaces [7]. As such, the ability to find an off-the-shelf device with “just enough” processing capacity for a particular wireless application is often limited, leaving unused capacity available in each deployed channel. For many systems, this effect may be somewhat offset by the concurrent reduction in “dollars per MIP” often associated with newer signal processing devices, but for wireless infrastructure systems supporting many simultaneous channels, the cumulative “excess processing capacity” from each dedicated channel may be sufficient to support multiple additional channels, and therefore represents a significant inefficiency in the overall architecture.

These inefficiencies in processor utilization can often be addressed through the adoption of an architecture where the signal processing resources are shared by multiple channels. In this type of architecture, processing devices are generally contained within a “processor pool” or “DSP farm,” with the number and type of processors selected based on their cumulative ability to support the total number of target channels (Fig. 2).

Inherent in this architectural model is the need for a communications infrastructure that will establish connectivity between various processing devices as required to support the target air interfaces. This infrastructure represents increased cost over the equivalent dedicated resource architecture, but this may be offset by an overall reduction in the number of signal processing devices required due to increased efficiency.

While the shared resource model may offer clear cost advantages in certain system architectures over the dedicated resource model in terms of efficiencies achieved in processor resource utilization, the practical realization of a shared resource system requires a number of advanced capabilities, each with its own cost impact. The first of these lies in the communications infrastructure required in the shared resource model. The simplest approach to establishing this infrastructure is to utilize a mesh network, with each processing

element maintaining a direct connection to the other processing elements within the pool. This type of architecture can be prohibitively expensive, requiring hardware for $N - 1$ data transports at each of the N processing devices.

The alternative approach, most commonly used, is to incorporate a switched fabric interconnect such as RapidIO between the devices [8]. In this model each device supports a connection to a fabric switch that routes signal data as appropriate between devices based on the requirements of the instantiated air interfaces. While often more cost effective than a mesh fabric, the disadvantage of the switched fabric interconnect is that the switch acts as a single point of failure in the system. This issue can be overcome through the use of a redundant switch that will take over the routing function when an error is detected in the primary switch; however, there will likely be some interruption in service during the failover process.

Another advanced capability required when adopting a shared resource model over a dedicated resource model is the assured partition in each of the processing elements. In a shared resource model memory leaks and other issues in the application code for one channel can spill over into processing resources that have been allocated for another channel operating on the same device, effectively “stealing” those resources from the other channel, ultimately causing hard-to-diagnose crashes. This issue can be addressed through assured operating partitions, as defined in ARINC-653, which divides the processing resources of a device across a number of “partitions,” each with dedicated memory and CPU cycles [9]. These partitions prohibit the processing associated with one channel from “accidentally” accessing the resources assigned to another, but can represent a significant increase in cost over a dedicated resource model.

For high availability, there are additional trade-offs that must be explored between dedicated and shared resource architectures. In a dedicated resource model, the failure of any signal processing device will typically only take down the channel associated with that specific device. This implies higher per channel availability than in a shared resource model, where the loss of a signal processing device may interrupt the multiple channels it actively supports. However, when adopting a redundancy model in supporting high availability, the shared resource model may offer some significant cost advantages.

In a dedicated resource architecture, a $2 * N$ redundancy model supports high availability, with each channel failing over to a dedicated redundant channel when the signal processing subsystem’s built-in test function detects an error. This model inherently doubles the hardware required and therefore doubles the cost. Conversely, the switched fabric interconnect in a shared resource model may allow the same availability to be achieved with an $N + M$ redundancy model, where M is less than N , by providing redundant processing resources that can be accessed via the fabric and dynamically loaded to support any of the active channels. This potentially represents a significant cost savings over the dedicated

(Continued on page S9)

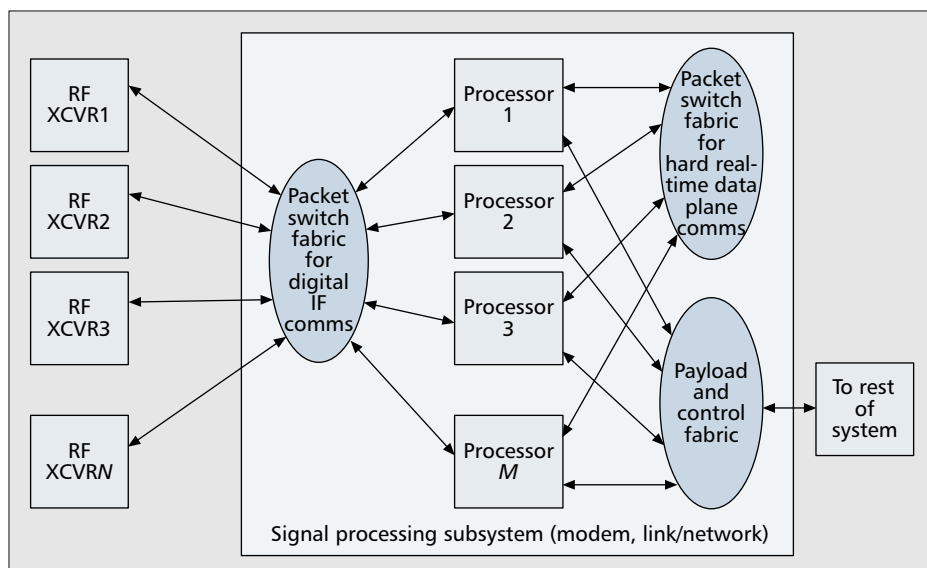


FIGURE 2. Use of switched fabric interconnects to enable shared resource processing.

edge about radio wave propagation mechanisms. Without such a background the result of a network design can be unpredictable. *Introduction to RF Propagation* by John S. Seybold gives the reader all the information required to understand the propagation of electromagnetic waves and the commonly used modeling techniques. It treats a wide range of wireless applications including mobile communications, radar, and satellite communications. After studying this book the task of designing a wireless network becomes much easier.

The book is organized in 12 chapters. Chapter 1 introduces terms related to RF propagation. The rules for frequency designation, the propagation modes (i.e., line of sight [LOS], non-LOS, indirect, tropospheric, ionospheric), as well as basic terms related to propagation modeling are explained. In Chapter 2 the physical background of radio wave propagation is described. Terms related to an electric and magnetic field are introduced in a smooth way, without exaggerated math analysis.

Chapter 3 provides the fundamentals of antennas. The basic antenna parameters like gain, effective area, radiation pattern, polarization, impedance, and voltage standing wave ratio (VSWR) are defined. Additionally, the properties of some common antennas are analyzed.

In Chapter 4 the author introduces the reader to the problem of radio link budget analysis. First, the impact of path loss, noise, and interference on link performance is analyzed, and later the link budget for an example link is calculated. Chapter 5 is devoted to radar systems in which the author has great experience. The equation for the radar range as well as several radar measurements are analyzed. The problem of clutter is also described. At the end of the chapter, atmospheric impairments are considered.

In Chapter 6 the atmospheric effects of radio wave propagation such as atmospheric refraction, multipath propagation, attenuation, and losses from moisture and precipitation are described. The problem of proper selection of the propagation model is deeply analyzed in Chapter 7. Several near-Earth models for foliage and terrain as well as for propagation in built up areas are introduced. At the end of the chapter the characteristics of the models are compared.

In Chapter 8 the origins of multi-

Introduction to RF Propagation gives the reader all the information required to understand the propagation of electromagnetic waves and the commonly used modeling techniques. After studying this book the task of designing a wireless network becomes much easier.

path fading are described and the fading phenomena is analyzed. The difference between large-scale and small-scale fading is explained. Chapter 9 is devoted to indoor propagation models. The rain attenuation of microwave and millimeter-wave signals is considered in Chapter 10. The procedure for incorporating fading due to rain into the link budget analysis is described. At the end of this chapter, in Annex 10A the reader can find useful tables with the values of parameters for rain attenuation models. Interesting information about satellite communication can be found in Chapter 11. The author smoothly introduces the reader to this complex subject and analyzes the

Trends in DSP/from page 56

model, and is another factor that must be evaluated when defining the system architecture.

Other factors play into selecting between shared and dedicated resource architectures as well, including the complexity of the software operating environment and a potential need to support partial reconfiguration on signal processing devices hosting air interfaces with multiple complex operating modes. So, given these factors, which model is the best architecture for use in a wireless infrastructure system? Well, as with the selection of device technology, the answer is, of course, "It depends." The additional cost of supporting the advanced capabilities required in the shared resource model may outweigh the savings in cost accrued through the reduction in the number of signal processing resources required in the dedicated resource model. Over time it is expected that technology advances will obviate these issues, but for now a trade-off analysis must be performed on a case-by-case basis to define whether a shared resource, a dedicated resource, or some combination of both is appropriate for a given wireless system.

problems presented in previous chapters (i.e., path free space-loss, rain attenuation, rain fades, antenna parameters) in a new perspective. The last chapter in the book is devoted to RF safety aspects including the biological effects of RF exposure. Additionally, FCC recommendations are presented.

The index at the end of the book will help readers to find their way through the wide range of topics covered in the text. The exercises and references at the end of each chapter can be helpful for self-studying students as well as teachers. Additionally, an accompanying Wiley ftp site provides supporting MathCad files for selected figures in the text.

Introduction to RF Propagation by John S. Seybold should be a good starting point into the world of wireless communications. It can be recommended for upper-level undergraduate and graduate students as well as for researchers and engineers interested in acquiring knowledge about electromagnetic wave propagation. Due to comprehensive coverage of models and applications, even a practicing engineer can find this book useful.

REFERENCES

- [1] L. Pucker, "Radio Communications Trends in DSP: Which Signal Processing Technologies Should I Use in My Wireless Application? Well, It Depends..." *IEEE Commun. Mag.*, vol. 43, no. 9, Sep. 2005.
- [2] R. Baines and D. Pulley, "A Total Cost Approach to Evaluating Different Reconfigurable Architectures for Baseband Processing in Wireless Receivers," *IEEE Commun. Mag.*, vol. 41, no. 1, Jan. 2003.
- [3] Texas Instruments, "TMS320C6000E DSP Multiprocessing With a Crossbar Switch," Application Note SPRA725, <http://focus.ti.com/lit/an/spra725/spra725.pdf>
- [4] L. Pucker, "Design Considerations for Distributed Architecture Software Defined Radios," *Proc. of the Commun. Des. Conf.*, Sep. 2002.
- [5] M. Lopez and J. Stephens, "Implementing WiMAX PHY Processing using a Programmable DSP," *CommsDesign*, Feb. 2005, <http://www.commsdesign.com/showArticle.jsp?articleID=59200057>
- [6] P. Cook, "Network Oriented Base Stations," SDR Forum Document Number SDRF-01-I-0057-V0.00, http://www.sdrforum.org/MTGS/mtg_25_sep01/01_i_0057_v0_00_nobs_09_25_01.pdf
- [7] L. Pucker, "Radio Communications Trends in DSP: Has FPGA Technology Peaked in Wideband Wireless Applications?" *IEEE Commun. Mag.*, vol. 42, no. 9, Sep. 2004.
- [8] http://focus.ti.com/pdfs/bcg/rapidio_wp.pdf
- [9] ARINC, "653-1 Avionics Application Software Standard Interface," https://www.arinc.com/cf/store/catalog_detail.cfm?item_id=496