CONNECTING WIRELESS INTERNET DEVICES USING DSP SOFTWARE

With the latest DSP devices, you can design a software-configurable, low-cost modem that's very efficient and small.

By Andrew Bateman

It's a very exciting time if you're designing a wireless modem. The Internet is the fastest-growing sector of the IT market, and high-speed wireless connection for business and residential users is a huge business opportunity. You also have unprecedented flexibility in the range of wireless functions that can be implemented in the digital domain, and new algorithms that fully exploit this advantage are appearing .

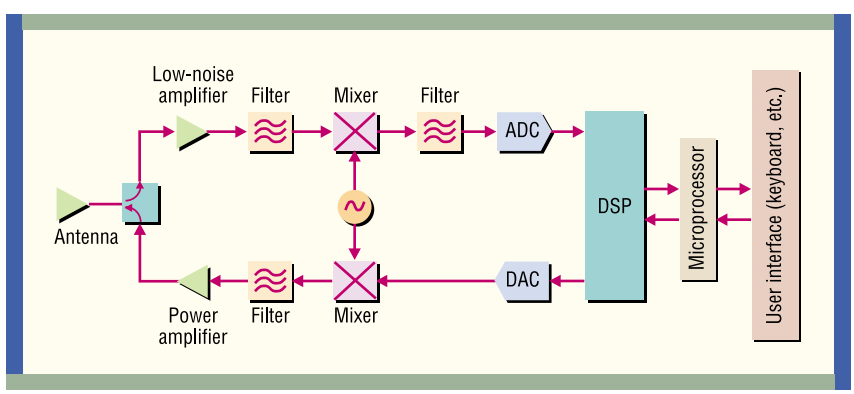

Figure 1. The architecture of a modern software radio is surprisingly simple, comprising only a few core building blocks. Linearity in the power amplifier and a good low-noise front end and synthesizer are essential, as of course are a high-speed, low-cost, low-power DSP engine and data converters (highlighted).

Still, wireless modems present a serious technical challenge, and that's even truer for wireless LAN (WLAN) devices. Cost-effectively distributing the ultrahigh capacity of a fiber node (multiple gigabits per second) to disparate users often requires a radio or free-space optical transmission link. Data rates in excess of 10 Mb/s are called for, with upward of 100 Mb/s a target for larger business users. In addition, a point-to-multipoint network or a distributed network is needed for distribution from a fiber hub.

For the WLAN terminal, your design challenge is to achieve robust communications at those high data rates, with good spectral efficiency (to maximize the number of customers that can be served for a given frequency allocation). Of course, minimal setup time, low maintenance, and a competitive price are assumed. From a manufacturing standpoint, the wireless platform must be easily tailored to new frequency bands as they become available and be able to flexibly exploit and manage the characteristics of the channel and variable data transfer demands.

Those goals are now achievable. By harnessing the ever increasing performance/power potential of modern DSP devices such as the Texas Instruments TMS320 DSP family, you can build a highly configurable, low-cost solution with minimal frequency-selective components, high efficiency, and small size.

Unlike the cellular phone market, which is tightly controlled by standards—GSM/IS95 and third-generation (3G) Universal Mobile Telephony Service (UMTS)—the WLAN market has no dominant air interface standard. Literally dozens of proprietary systems are in play. This situation gives you the flexibility to work with the latest modulation, coding, access, and equalization formats, thus squeezing every last drop of capacity out of the channel. However, the manufacturer has to build a degree of flexibility into its design (or develop multiple versions) so that the product can operate with multiple air interface standards across a range of service providers and network types.

Achieving flexibility in a wireless platform requires three key features: minimal frequency-selective components, high linearity to minimize self-induced signal distortion, and maximum software-defined functionality.

The core building blocks of a modern software-defined digital radio are a linear power amplifier (PA); a wideband, low-noise front end; a good synthesizer; and, of course, a high-speed DSP engine and data converters (Figure 1). For maximum software configurability, the analog signals should ideally be converted into digital form at the RF carrier frequency. Unfortunately,

the conversion isn't viable at frequencies much above 500 MHz; with most WLAN spectrum allocations in the 5-GHz-plus range, an analog intermediate-frequency stage is still required. The digitization frequency used is governed by component costs; the linearity, speed, and dynamic range of D/A and A/D converter technology; and power consumption constraints.

Both the linearity and the speed of converters have improved significantly, with 14-bit, 100-megasample-per-second, 80-dB SNR converters readily available. These devices can easily support 40-MHz subsystems using direct sampling and can extend toward 300-MHz IF solutions using subsampling methods.

The power consumption of fast converters is significant (several hundred milliwatts), and further reduction of IC feature size or other advances in process technology are required to realize the power savings needed to allow IF sampling in handheld equipment. For fixed WLAN installations, however, power consumption is less critical; instead, linearity and sampling rate are likely to dominate the choice of frequency.

DSP tasks for a high-speed wireless Internet modem comprise core modem functionality (source and channel coding, pulse shaping, modulation, demodulation) and more advanced software radio management tasks (quadrature frequency translation, PA linearization, receiver dynamic range extension, calibration of the in-phase and quadrature [I-Q] components of the signal, automatic power and frequency control, and direct digital synthesis, for example). We'll focus here on pulse shaping, demodulation, and PA linearization.

PULSE SHAPING

To maximize the data transmission rate over a wireless link with finite bandwidth, you must shape the data pulses modulating the carrier signal. For modems using frequency shift keying (FSK), shaping traditionally involved Gaussian filtering; for the more advanced quadrature amplitude modulation (QAM) modems, a root raised cosine (RRC) filter is commonly employed. Whatever the pulse shaping, there are two core algorithms for implementation: the classical filter and the lookup table (LUT) method.

The filter approach is normally realized using an infinite impulse response (IIR) filter for a Gaussian filter shape and a finite impulse response (FIR) filter for the RRC shape. The FIR filter realization allows the synthesis of a near-perfect RRC transfer function with linear phase and controlled stopband attenuation. (IIR approximation of an RRC filter can be used, but in the

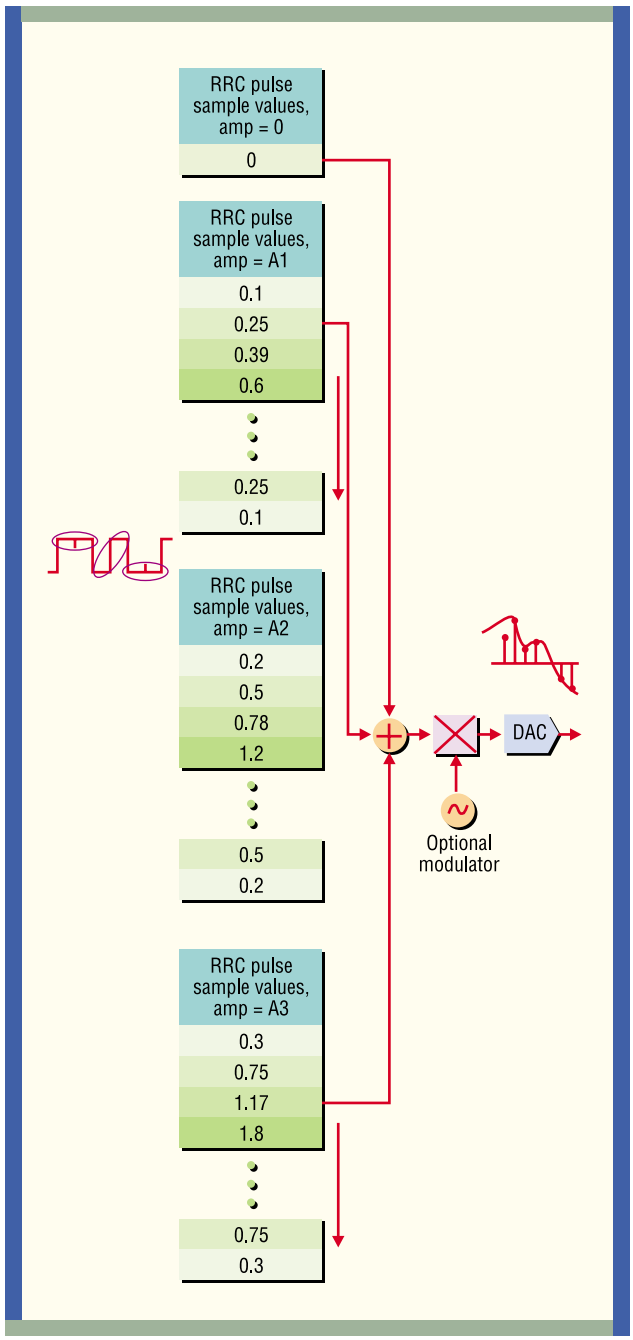

Figure 2. Although appearing more complex at first glance, the lookup table (LUT) method of symbol pulse shaping, using the data transitions to index the stored RRC pulse shape, imposes little processing overhead. The output waveform is synthesized from the summation of the pulse samples for several preceding data bits.

majority of cases, the FIR structure yields a lower-overhead algorithm than the compensated IIR design.) The RRC filter shape is always an approximation of the true Nyquist filter response, trading off filter length (and hence delay and processor load) for stopband attenuation and roll-off rate. Most filter design packages offer raised cosine (RC) and RRC filter options, making it easy to explore the trade-off between the two parameters.

In many cases, it may be preferable to cascade an RRC filter with a second, FIR filter (possibly half-band for ease of implementation). The second filter achieves the desired level of out-of-band attenuation to meet a given spectral mask while relaxing the requirements on the first—and more processor-intensive—filter.

The alternative approach to pulse shaping makes use of a lookup table, The table holds the precalculated values for the pulse response of the desired filter, based on all possible input state transitions (Figure 2). The state transition is used to index the correct stored pulse response from the chosen filter, which is then summed with the pulse responses from previous transitions to form the composite pulse-shaped waveform.

The lookup table method is preferred when execution time is at a premium, as lookup table indexing carries very low overhead. Conversely, the real-time filter realization of pulse shaping is used when memory space is at a premium and storage of the multiple filter pulse responses is impractical.

DATA DEMODULATION

Efficient algorithms for data demodulation are key to the success of software-defined wireless modems. A broad range of demodulation algorithms are in widespread use, mirroring the multiple modulation formats, coding strategies, carrier and symbol timing recovery mechanisms, and equalization methods deployed.

We'll look at one demodulation algorithm for frequency discrimination that has widespread application and is a good example of a DSP-optimized solution. Frequency discrimination has two primary uses:

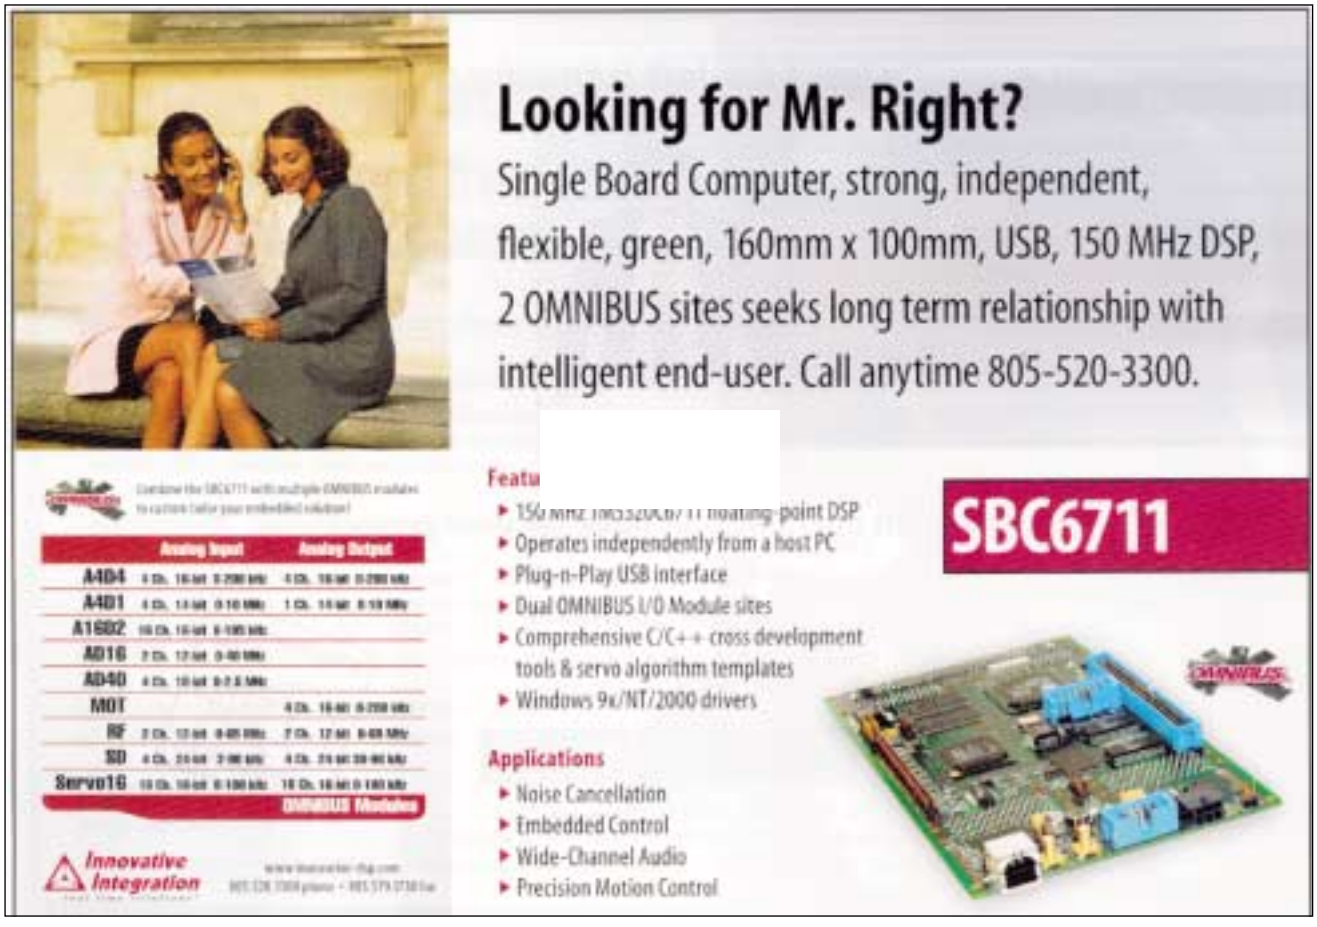

demodulation of the FSK family of waveforms, such as Gaussian minimum shift keying (GMSK) used in GSM cellular communications, and automatic frequency control loops. The DSP frequency discriminator algorithm uses differentiation and cross multiplication to generate an output that directly corresponds to the instantaneous frequency of the input signal samples. An optional envelope normalization block removes fading components present in the input signal (Figure 3).

Mathematically, the operation of the quadrature discriminator is as follows: For general complex input signals of the form:

$$I(t) = \tfrac{1}{2}\, r(t) \times \sin \phi(t)$$
$$Q(t) = \tfrac{1}{2}\, r(t) \times \cos \phi(t)$$

where $r(t)$ is the signal envelope and $\phi(t)$ is the angular phase/frequency, the signals at the outputs of the two differentiators can be represented as:

$$\dot{I}(t) = \tfrac{1}{2}\, r(t) \times \dot{\phi}(t) \times \cos \phi(t) + \tfrac{1}{2}\, \dot{r}(t) \times \sin \phi(t)$$
$$\dot{Q}(t) = -\tfrac{1}{2}\, r(t) \times \dot{\phi}(t) \times \sin \phi(t) + \tfrac{1}{2}\, \dot{r}(t) \times \cos \phi(t)$$

By cross-multiplying and subtracting the signals as shown in Figure 3, you obtain an output signal, given by:

$$\dot{I}(t) \times Q(t) - I(t) \times \dot{Q}(t) = \tfrac{1}{4}\, r^2(t) \times \dot{\phi}(t)$$

Further division by the (envelope)^2 term yields a normalized real-time measure of the instantaneous frequency variations of the input signal. (In practice, it's much more efficient to use a lookup table to generate $1/r^2(t)$, which is multiplied with the top path signal.)

Unlike conventional frequency discriminators based on a phase-locked loop (PLL), the algorithm doesn't involve a feedback process. It also introduces little or no bandwidth expansion into the signal, thus ensuring that the Nyquist sample rate limit is not violated.

Perfect for your environment...

Blackhawk™ USB-JTAG Emulator

Now Supporting C2000 DSP processors

- *USB (bus powered) Peripheral • Portable and Compact (1.1 x 2.6 x 5.5 inches)*
- *Works with TI Code Composer Studio™*
- *TMS320C5000 and C6000 Support*
- *Quick Installation*

DSP Enabling Technologies™

- *Development Hardware*
- *Operating Systems*
- *Bundled Toolsets*
- *Design Services*
- *Consulting*

Blackhawk™

by EWA Inc.

© 2001 EWA

For more information please visit: **www.blackhawk-dsp.com** *or phone:* **1-877-983-4514**

Code Composer Studio™ is a registered trademark of Texas Instruments

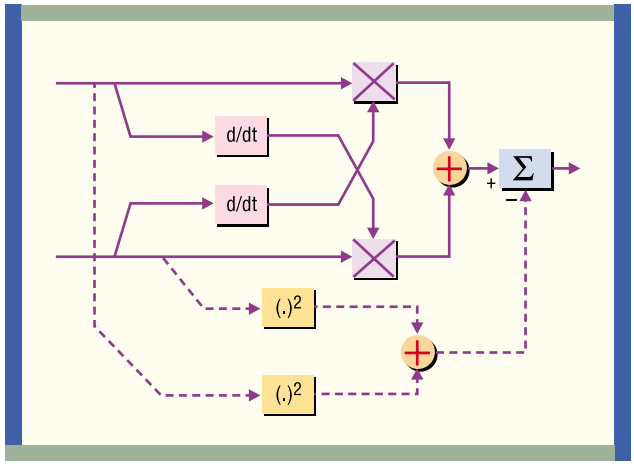

Figure 3. Obtaining an instantaneous measure of the frequency of a signal is very difficult using analog methods. In contrast, the algorithm for frequency discrimination shown here uses quadrature processing to recover the instantaneous (normalized) frequency of the input signal with very low processing overhead.

PA LINEARIZATION

Pulse shaping and multisymbol modulation are wasted if the waveform is overly distorted after passing through the analog transmitting or receiving functions. Designers often sacrifice linearity because they're trying to maximize PA efficiency and power output, so some distortion often results. The conventional solution—backing off the PA drive to operate within a linear portion of the PA characteristic—wastes considerable power. It's much more efficient to use DSP techniques to predistort the waveform of the signal driving the PA, in a complementary manner to the PA nonlinearity.

The source waveform is passed through a lookup table, which stores the correction factor for the amplitude (and phase) of the waveform at any given PA drive level (Figure 4). For most applications, the PA characteristic isn't stable enough (for example, temperature, supply voltage, output load) for the process to be purely open-loop. A feedback path from the PA output is therefore usually employed to allow the residual PA distortion to be measured and the lookup table coefficients to be updated.

For the high-bit-rate solutions required in WLANs and the LUT-specific nature of the processing tasks involved, it's often best to implement adaptive predistortion using dedicated ASICs or FPGAs.

A/D AND D/A CONVERSION

Many manufacturers, including Texas Instruments, are competing aggressively in the area of high-speed digital IF converter solutions, and new devices are appearing on the market almost weekly.

Digital up and down conversion. The initial task for the digital processing unit is to convert the IF signal into a complex baseband form (down conversion) and from complex baseband to IF (up conversion). Down conversion ensures minimum sampling rate processing for the remaining radio functions. The tasks involve mixing (multiplication) with quadrature versions of a digital oscillator. In addition, a process of interpolation and decimation is needed in the up converter and down converter, respectively, to optimize the sampling rate between the digital IF requirements and the complex baseband requirements.

Because the two functions (frequency mixing and sample rate conversion) are common to all digital IF solutions, various manufacturers have produced dedicated ICs optimized for digital up and down conversion. The high-sample-rate processing associated with the digital-to-analog interface can be accommodated in these hardwired devices, allowing the (comparatively) slower digital signal processing of the wireless signal content to be undertaken in cheaper, lower-power, software-programmable DSPs.

An alternative route that maintains full flexibility of design is to implement the mixing and sample rate

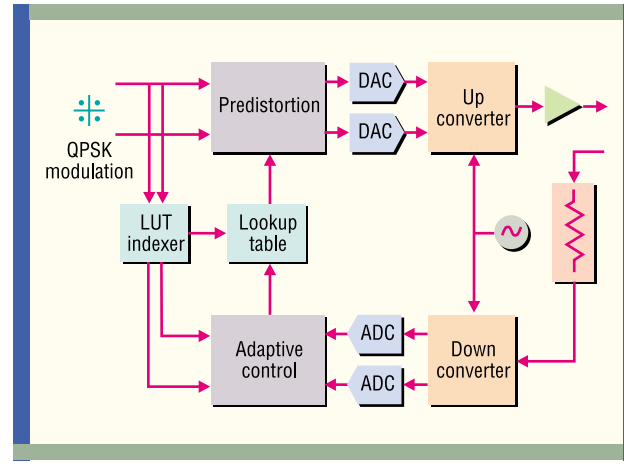

Figure 4. DSP-based adaptive baseband predistortion can greatly simplify the design of the RF power amplifier and enhance its efficiency. The solution shown here uses feedback from the PA output to update a lookup table operating on the baseband samples.

conversion processing in software, using high-speed DSPs or user-programmable gate arrays. Some companies offer FPGAs suited to that task, and TI and others offer DSPs with sufficient processing speed, such as the TMS320C6000 and C5000 DSP platforms. Additionally, numerous third-party suppliers, including TI Third Party Network Members, are providing custom algorithms or development tools.

Digital signal processing engines. When the IF signal is in complex baseband form, you can process the signals using a dedicated ASIC (very limited flexibility), an off-the-shelf DSP solution (or a DSP core), or one or more FPGAs. Again, the maximum flexibility sought in our software-programmable WLAN solution is achieved using the DSP or FPGA option. A growing number of high-speed DSP devices can handle data rates of several megabits per second, including the C5000 platform of ultralow-power devices for portable use and the C6000 platform for very fast applications. ◆

REFERENCES

Bateman, Andrew, *Digital Communications: Design for the Real World*, Addison-Wesley, London, 1998.

Bateman, Andrew, and Iain Paterson-Stephens, *The DSP Handbook*, Prentice-Hall, Upper Saddle River, N.J., forthcoming.

Kenington, Peter B., *High Linearity RF Amplifier Design*, Artech House, Boston and London, 2000.

Andrew Bateman is the CEO of Avren Ltd., a design and consultancy company that advises communications and IC companies on technologies, standards, and partnerships for next-generation wireless systems, and of DSPStore.com. Previously, he was a professor of communications and signal processing at Bristol (U.K.) University and cofounded Wireless Systems International Ltd., where he was the business development director. He is also the author of three books on digital communications and digital signal processing.