

Exploit Digital Advantages in an SSB Receiver

by Richard Groshong and Stephen Ruscak

**IMPROVEMENTS
IN DIGITAL
CONVERSION
TECHNOLOGY
HELP RADIO
DESIGNERS
TAKE
ADVANTAGE
OF DIGITAL
ARCHITECTURES.**

This article is the second of a two-part series on the design of digital radio receivers. The first part (ELECTRONIC DESIGN, May 23, p. 67) covered the fundamentals of digital radio and undersampling techniques, as well as selection criteria for the analog-to-digital converter. Part 2 covers the receiver design in greater detail.

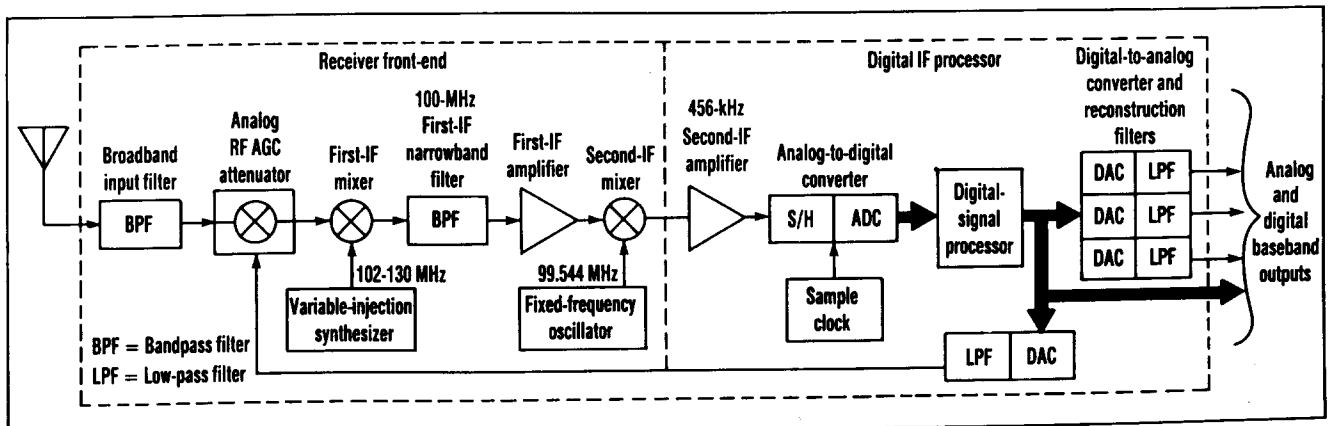
The first step in digital-receiver design is to determine the radio's front-end specifications. The process is somewhat similar to that in a conventional analog design, but it must account for the characteristics of the digital IF processor used by the radio. After the front-end is completed, the designer can proceed to the IF processor, taking advantage of highly integrated devices to simplify the design.

The design example in this article is a 2-to-30-MHz SSB receiver employing undersampling techniques. A 16-kHz front-end bandwidth was chosen to accommodate up to four 3-kHz independent-sideband channels.

This bandwidth also allows the digital-signal processor to perform fine tuning over at least a 1-kHz range, simplifying the variable-injection synthesizer (Fig. 1).

The first IF is 100 MHz, which is high enough to ensure that the broadband input filter suppresses the IF and image response by 80 dB, but low enough to allow a low-cost crystal filter to be used as the first-IF narrowband filter. Because the DSP performs the fine tuning in 1-Hz steps, the variable injection synthesizer can tune from 102 to 130 MHz in 1-kHz steps. This "high-side" first-IF mixer injection causes a passband reversal or "flip" in the mixer, which will be corrected by the sampling process.

The first-IF amplifier makes up for losses in the first-IF mixer and filter, and maintains a front-end noise figure of 15 dB. After amplification, the signal is mixed with a fixed-oscillator signal of 99.544 MHz. The result is a second-IF of 456 kHz, which is high enough to allow the second-mixer image response at 99.088 MHz to be attenuated by 80 dB in the first-IF filter, but low enough to be sampled by the analog-



1. THE CHOICE OF 100 MHz for the first IF is high enough to ensure that the broadband input filter sufficiently suppresses the IF and image response but low enough to allow a low-cost crystal filter for the first-IF narrowband filter.

Reprinted from ELECTRONIC DESIGN — June 13, 1991

to-digital converter (ADC).

If the ADC's sample-and-hold bandwidth is greater than 456 kHz, as is the case with the AD779, the receiver can undersample the second-IF signal efficiently at a sample rate of 96 kHz. Because the analog-to-digital conversion is effectively the result of a convolution of the input signal and the sample-clock impulse spectrum, the sample clock must be as "clean" and jitter-free as the local oscillator or mixer injection signal. Any phase noise present on the clock will be impressed on the signal being sampled, effectively corrupting it. If isolation or buffering is inadequate, the spurious noise picked up by the sample clock can be disastrous.

When the second-IF spectrum is convolved with the sample-clock impulse spectrum, the IF signal is frequency translated to a lower IF of one-fourth the sample frequency, or 24 kHz (Fig. 2). The convolution also creates a passband reversal that counteracts the reversal in the first-IF mixer. As expected, the sampled-IF spectrum repeats at every multiple of the sample frequency.

In this example, increasing the second-IF bandwidth will create aliasing distortion or overlap in the sampled IF spectrum because real-world filters can supply only finite

stopband attenuation (Fig. 2, again). But aliasing can be reduced to an acceptable level. In this design, the passband width of 16 kHz must be alias-protected to 60 dB.

This requirement, and a look at the IF sampling process, lead the designer to the bandpass specifications of

the first-IF filter (Fig. 3). Because of the sampling frequency and second-IF selected, the filter must have an 80-kHz-wide stopband. Given the 16-kHz-wide passband, the corresponding shape factor is 5 to 1. Attenuation of signals 40 kHz from the passband center of 100 MHz must be at least 60 dB. As noted, another 20 dB of attenuation is also needed at the image-response frequency of 99.088 MHz. A typical passband ripple specification might be 1 dB. A 4- or 5-pole crystal filter might satisfy these criteria.

THE DIGITAL STAGE

The signal from the second-IF mixer must be amplified to a level sufficient to drive the ADC. The gain required depends on the ADC and various receiver performance trade-offs, and must be determined based on a system-level analysis. For example, the theoretical quantizing noise density of an ADC, N_{0q} , in W/Hz is:

$$N_{0q} = (V_{pk-pk}/2^b)^2/6f_s R$$

where V_{pk-pk} = the peak-to-peak voltage range of the ADC in volts, b = the total number of ADC bits including sign, R = the assumed imped-


SPREADSHEET OUTPUT

Receiver performance for thermal-noise-limited example

Digital-signal-processor receiver gain distribution and noise performance

Analog RF/IF noise figure	15 dB
ADC resolution	14 bits
ADC full-scale level	5 V peak (23.98 dBm)
ADC quantization level	-60.31 dBm
ADC sample frequency	96 kHz
ADC input bandwidth (BW1)	16 kHz
Information bandwidth (BW2)	3 kHz
Analog AGC threshold	-77 dBm

Antenna signal level (dBm)	Analog RF/IF gain (dB)	ADC signal level (dBm)	Noise at ADC input in BW1 (dBm)	Noise at ADC input in BW2 (dBm)	Quantizing noise of ADC in BW2 (dBm)	Total noise in BW2 (dBm)	Output SNR in BW2 (dB)	Digital processor gain (dB)	Output signal level (dBm)	Antenna overload level (dBm)
-113.00	57.00	-56.00	-59.96	-67.23	-74.11	-66.42	10.42	60.00	4.00	-33.02
-103.00	57.00	-46.00	-59.96	-67.23	-74.11	-66.42	20.42	50.00	4.00	-33.02
-93.00	57.00	-36.00	-59.96	-67.23	-74.11	-66.42	30.42	40.00	4.00	-33.02
-83.00	57.00	-26.00	-59.96	-67.23	-74.11	-66.42	40.42	30.00	4.00	-33.02
-73.00	53.00	-20.00	-63.96	-71.23	-74.11	-69.42	49.42	24.00	4.00	-29.02
-63.00	43.00	-20.00	-73.96	-81.23	-74.11	-73.34	53.34	24.00	4.00	-19.02
-53.00	33.00	-20.00	-83.96	-91.23	-74.11	-74.03	54.03	24.00	4.00	-9.02
-43.00	23.00	-20.00	-93.96	-101.23	-74.11	-74.10	54.10	24.00	4.00	0.98
-33.00	13.00	-20.00	-103.96	-111.23	-74.11	-74.11	54.11	24.00	4.00	10.98
-23.00	3.00	-20.00	-113.96	-121.23	-74.11	-74.11	54.11	24.00	4.00	20.98
-13.00	-7.00	-20.00	-123.96	-131.23	-74.11	-74.11	54.11	24.00	4.00	30.98
-3.00	-17.00	-20.00	-133.96	-141.23	-74.11	-74.11	54.11	24.00	4.00	40.98
7.00	-27.00	-20.00	-143.96	-151.23	-74.11	-74.11	54.11	24.00	4.00	50.98
17.00	-37.00	-20.00	-153.96	-161.23	-74.11	-74.11	54.11	24.00	4.00	60.98



2. THE SAMPLING PROCESS IS a convolution of the second-IF spectrum (a) with the sampler spectrum (b). The resulting spectrum shows that if the second-IF bandwidth is increased, aliasing distortion will occur (c).

ance in ohms, and f_s = the sampling frequency in hertz.

For a 14-bit ADC with a 10-Vpk-pk range and a sampling frequency of 96 kHz, $N_{0q} = 1.2935 \times 10^{-14}$ W/Hz or -108.9 dBm/Hz, assuming 50- Ω impedance normalization. Other noise includes the thermally generated receiver noise. The level of this noise is the front-end noise figure, 15 dB/Hz, minus the thermal noise generated in a 50- Ω resistor, -174 dBm/Hz, or -159 dB/Hz.

The receiver's minimum weak-signal gain must be high enough so that the weakest desired usable signal plus the receiver noise is greater than at least one ADC quantizing level. If not, the signal won't be recovered. The best way to ensure a large enough signal is to provide sufficient noise at the ADC input to dither across a quantizing level. To do so, a designer can amplify the in-band receiver thermal noise or add out-of-band noise, which the digital-signal processor can filter out later. Although it may not be the better method, this article discusses the simpler approach of amplifying the receiver thermal noise until it bridges a quantizing level.

For the AD779, one quantizing level is 610 μ Vpk-pk. The equivalent sine-wave amplitude is then 215.8 μ V rms or -60.3 dBm. Therefore, amplifying the receiver noise so that its rms level at the ADC is also -60.3 dBm will guarantee adequate bridging of a quantizing level. With the front-end noise bandwidth (BW) of 16 kHz, the gain required to amplify the receiver noise to this level is:

$$\begin{aligned} \text{Gain} &= \text{quantizing level} - \text{thermal noise density} - 10 \log \text{BW} \\ &= -60.3 \text{ dBm} + 159 \text{ dBm/Hz} - 10 \log 16,000 \text{ Hz} \\ &= 56.7 \text{ dB} \end{aligned}$$

A spreadsheet program can perform these calculations, and others, to interactively examine receiver performance for different ADCs, sample rates, bandwidths, and gain distributions (see the table).

Note that in this example, the analog gain is high enough to allow thermal noise to bridge a quantization level. As a result, the receiver's noise

performance is dominated by thermal noise, not analog-to-digital quantization noise. At the sensitivity level, the thermal noise is 7 dB greater than the quantization noise when measured in the 3-kHz information bandwidth.

The output signal-to-noise ratio (SNR) is better than 10 dB at an antenna signal level of -113 dBm. Digital automatic gain control (AGC) holds the output signal level at 20 dB below full scale (+4 dBm) until the RF AGC threshold is reached. The output SNR increases with the antenna signal level until the level at the ADC reaches 44 dB below full-scale (-20 dBm). At this point, the analog AGC holds the ADC input and output levels constant. The output SNR continues to rise to 54 dB until the analog-to-digital quantization noise dominates the receiver noise.

UNWANTED SIGNALS


Another consideration is the antenna overload level, which is the antenna signal level at which an undesired received signal can pass unattenuated through the 16-kHz bandwidth front-end to the ADC and saturate it. Because this undesired signal would be outside of the 3-kHz information bandwidth, it would not affect the normal AGC setting.

As a result, the receiver must maintain some "headroom" to accommodate signal-level peaks. Only

desired in-band signals should be present after the final narrowband filtering in the digital-signal processor. Consequently, 20 dB of headroom is sufficient at the digital-signal-processor output. At the ADC input, however, large out-of-band undesired signals may be present. Consequently, at least 40 dB of headroom is required.

The maximum headroom is limited by the SNR required for both noise- or distortion-limited cases. In the example design, the maximum SNR for large desired signals, without interference, is 54 dB. If an out-of-band undesired signal at the ADC has a level of +4 dBm, or -20 dB from the saturation level, the converter will generate harmonic and intermodulation distortion products that will alias and potentially fall inside the desired signal bandwidth. If these products are specified to be 70 dB down (-66 dBm) and the in-band desired signal level at the ADC is -20 dBm, the ultimate signal-to-noise-and-distortion, $S/(N+D)$, ratio will be limited to 46 dB (-20 dBm + 66 dBm).

The next step in the design is the digital IF processor architecture, which will depend on the desired level of functionality and the choice of hardware. In a typical implementation of a digital SSB processor, the digitized IF signal is first processed by the IF translator (Fig. 4). The



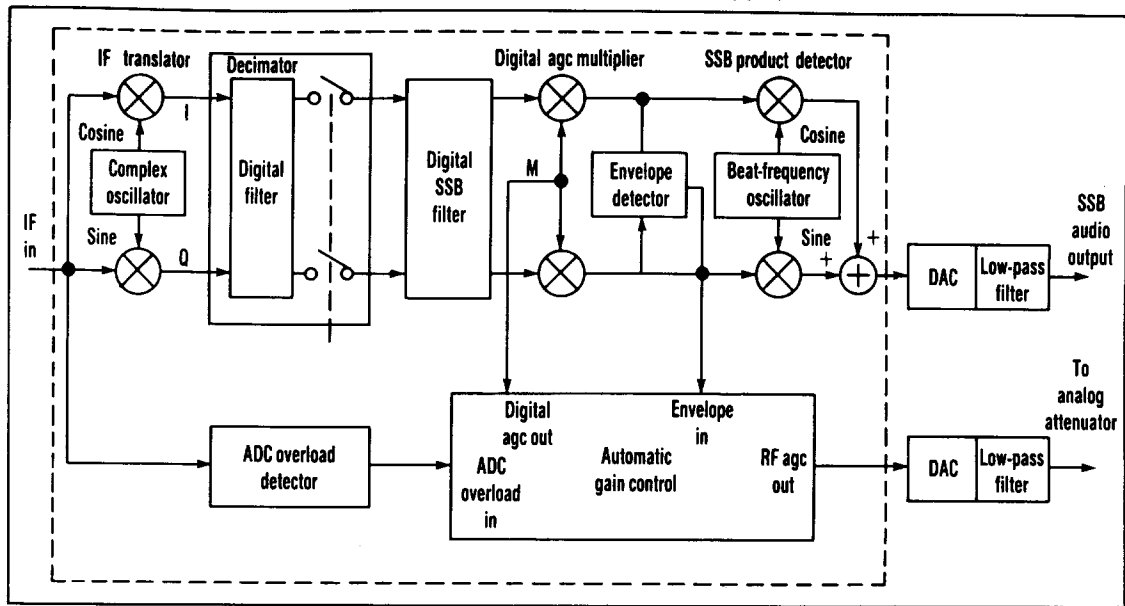
3. A DETAILED LOOK AT THE SAMPLING process helps in specifying the first-IF filter. The 456-kHz second IF (a) and 96-kHz sampling frequency (b) are the determining factors. To protect the desired 16-kHz bandwidth from aliasing, the filter must have no more than an 80-kHz-wide stopband, or a 5-to-1 shape factor (c).

translator shifts the SSB channel's center frequency to dc, or zero frequency, and at the same time converts the "real" IF signal to a complex baseband signal. The resulting real and imaginary components are often referred to as the inphase (I) and quadrature (Q) components, respectively. This representation of the signal as a rotating vector or phasor is commonly used in digital-signal-processing algorithms.¹

A complex oscillator and a multiplier perform the frequency translation and complex baseband conversion. Typically, the oscillator consists of a phase accumulator and sine and cosine algorithms (Fig. 5). The phase accumulator is a modulo 2π adder that increments its output each sample period by adding a fixed constant, or phase increment, to its previous sample output. This process produces a continuously changing output phase angle ranging from 0 to 2π radians.

The algorithms then compute the phase angle's cosine and sine to generate the complex oscillator's real and imaginary outputs. Next, the real input signal is multiplied by the oscillator's outputs to produce the I and Q components of the complex baseband signal.

The next step is to implement the SSB filter, AGC, and demodulator in the digital-signal-processor chip. But the system's sample rate must be reduced as much as possible before the digital-signal processor can perform these computationally intensive algorithms. For example, if the DSP's instruction rate is 12 MHz and the sample rate remains at 96 kHz, only 125 instruction cycles will be available to run the algorithms on each point. This figure is insuffi-



4. THE RADIO'S SSB FILTER, AGC, and demodulator are implemented by the digital-signal-processor chip. The system's sample rate is first reduced by a factor of 10 in the decimator to produce enough instruction cycles between samples to perform the required algorithms.

cient. Therefore, the receiver decimates the sample rate by a factor of 10. The resulting 9600-Hz rate is enough to support the typical 3-kHz SSB bandwidth. It also enables 1250 instruction cycles between the decimated samples, enough to perform the required algorithms.

FILTER THE INPUT

Before the sample rate can be reduced, the input signal must be filtered to attenuate out-of-band signals. Otherwise, these signals can alias back in-band, causing distortion. This decimation filter usually consists of two identical finite-impulse-response (FIR) filters in the I and Q paths. Computer programs for designing FIR filters with arbitrary frequency response are readily available from many commercial sources. The most popular filter design algorithm, used in many of these programs, was distributed by the IEEE Press in 1979.²

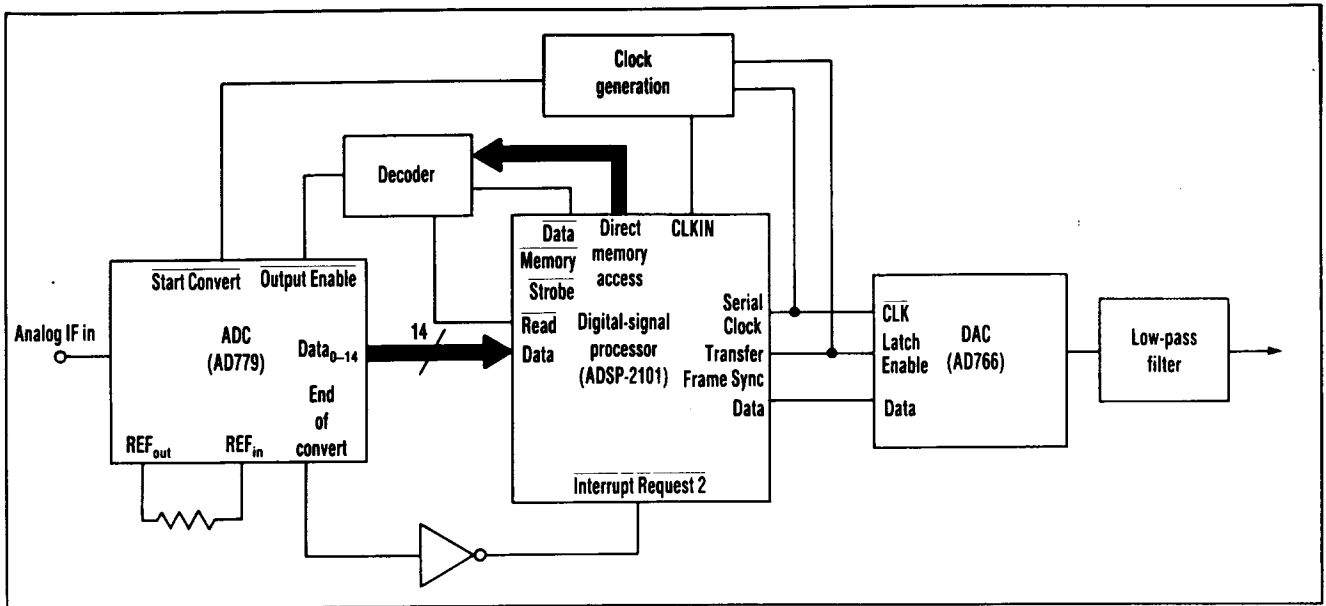
After decimation, the signal passes through a high-performance receiver IF filter. This SSB filter is an FIR or infinite-impulse-response device. If necessary, the SSB filter's output level is adjusted by the digital portion of the AGC algorithm. The AGC controls the digital gain by mul-

tiplying the I and Q components by a positive scalar value, M (Fig. 4, again).

The leveled I and Q outputs are then applied to the envelope detector and SSB product detector. The envelope detector computes the magnitude of the phasor represented by the I and Q components. The output envelope is simply the square root of the sum of the squared I and Q values. The AGC algorithm uses the envelope to adjust the receiver output level by a combination of analog (or RF) and digital gain control. To prevent saturation by strong out-of-band signals, the algorithm also senses any ADC overload and adjusts the RF and digital gain distributions accordingly.

The SSB product detector is a "half-complex" frequency translator that shifts the SSB carrier to its proper frequency. Just as in the IF translator, the circuit uses a complex oscillator, commonly called a beat-frequency oscillator. First, the I and Q signal values are multiplied by the cosine and sine outputs of the oscillator. Then the products are added to produce the receiver's audio output.

The hardware needed to implement the digital IF processor depends on several factors, but nearly



5. BECAUSE TYPICAL DIGITAL IF PROCESSOR architectures are highly integrated, most functions can be implemented in the DSP software. In this example circuit, only a few ICs are needed to perform the functions, as described in Figure 4.

all systems will include certain common components (*Fig. 6*). The total number of ICs can be minimal, keeping overall system cost low.

Modern communication receivers would typically use a high-performance digital-signal processor, such as the ADSP-2101. This processor is well-suited to many of the operations of a digital receiver and can easily link to the 14-bit parallel output of the AD779. Conversion back to the analog domain is simplified by digital-to-analog converters (DACs), like the AD766. These low-cost, 16-bit, serial-input DACs require no additional glue logic or interface circuitry, have an on-board reference, and are completely tested and specified over temperature.

Additional circuitry performs decoding, clock generation, and output filtering. The decoder handles handshaking between the AD779 and the digital-signal processor. Depending on the number of peripherals connected to the processor, the decoder may be as simple as a few logic gates. Clock signals are typically generated by a high-stability crystal with low phase noise. The crystal can serve as the processor's master clock and can also be divided by counters or other means to generate the ADC and DAC sampling clocks.

On the falling edge of the sampling clock (SC), the AD779 digitizes the incoming IF signal (*Fig. 6, again*). Approximately 8 μ s later, the conversion is complete and the AD779 asserts End of Convert (EOC). The falling edge of the inverted EOC signal interrupts the digital-signal processor. The control and address lines from the processor are decoded to generate the Output Enable (OE) pulse for the ADC. The ADC then places the digitized output on the data bus to be read by the processor.

The ADSP-2101 offers several features tailored to communications systems. The ability to fetch two operands (typically a coefficient and a data point), multiply these operands, and then sum the results with previous products in one processor cycle makes FIR filter algorithms extremely efficient. And an internal 40-bit accumulator enables full-precision products to be calculated without round-off noise caused by truncation in the filter computations. Interprocessor communication between multiple DSPs, which is often needed in SSB equipment, is easily implemented with one of the processor's two integrated serial ports.

The digital IF processor's high integration level is indicative of the fu-

ture of SSB equipment. As the performance of the ADC improves, it will be placed closer to the receiving antenna in the analog front-end. The ultimate goal is to digitize the incoming RF signal directly from the antenna and implement the remaining digital functions in custom ASICs. □

References:

- ¹ W.E. Sabin and E.O. Schoenike, *Single-sideband systems & circuits*, (McGraw-Hill Book Co., 1987).
- ² J.H. McClellan, T.W. Parks, and L.R. Rabiner; DSP Committee of the IEEE Acoustics, Speech, and Signal Processing Society, *Programs for digital signal processing*, (IEEE Press, 1979), Chap. 5.