A High-Performance Digital-Transceiver Design, Part 2

Part 1 of this series looked at some of the system-level design and tradeoffs for a high-performance transceiver that translates directly from RF to digital. In this installment, we'll look at the details of the receiver front end.

By Jim Scarlett, KD7O

The first installment of this series gave an overview of a high-performance transceiver using an "almost all digital" approach.[1] While still using analog filters and amplifiers, this approach translates directly between RF and digital domains. Design goals were presented that would meet or exceed the performance of commercially manufactured radios.

Here in Part 2, we'll look at the design details for the receive side of the radio. DSP will not be covered until a later installment, but we'll look at the receive signal processor (RSP) hardware design here. Considerations for using the ADC will be covered, including the use of dither and interfacing the ADC to the RSP. The clock source will get special attention as well.

[1]Notes appear on page 9.

10625 Highstream Dr
Raleigh, NC 27614
kd7o@arrl.net

Analog Processing First

First I'll describe the analog front end. The RF amplifier uses ideas that have been covered before. I'll describe how those ideas come together in this design. I won't spend much time on the filters, since the basic design has already been covered in detail in *QEX*.[2]

I mentioned in Part 1 that my filter capacitor values are different from Bill Sabin's to account for my own areas of interest. I also made some other changes regarding components. In the interest of minimizing size and cost, I use NP0 ceramic-chip capacitors in the filters. These capacitors will be okay for up to a few watts of input power. See Fig 1 for the filter schematics. Capacitor values of "0" indicate where the board has pads for additional capacitors to adjust the values while tuning the filters.

The relays are inexpensive and can easily handle the required power. They also switch in about 5 ms, so they are compatible with rapid switching between transmit and receive—or between bands. Additional relays are included in the filter/amplifier module for transmit/receive switching of the module.

RF Amplifier Board

As I mentioned in Part 1, I rejected the use of a common-base amplifier configuration because of poor reverse isolation. While looking at the amplifier in John Stephensen's noise blanker,[3] I decided that the topology would meet

my needs in this area. The desired *IP3* performance would be achieved through a push-pull transistor pair operating at a higher bias current.

The resulting amplifier is shown in Fig 2. A Mini Circuits 1:1 transformer generates the balanced signal used by the two transistors. I used some Motorola MRF5811 transistors that I had on hand. These devices—basically an MRF581 in a SOT-143 package—provide an excellent noise figure while running at high currents. In this application, they are being operated with a bias current of about 40 mA.

The 20-Ω resistors at the emitters set the bias currents. Along with the feedback transformers, they determine the input impedance of the amplifier. The transformers are wound on binocular ferrite cores. The current levels are too high for typically available RF transformers, which should not be substituted. The transformer impedance ratio also helps set the gain of the amplifier.

The transistors can be powered with +9.5 to +12 V. The higher supply voltage will improve the *IP3* by a couple of decibels, but will force the transistors to dissipate more power. I am operating the prototype at about +10 V. Later versions will probably have a regulator on board to generate this voltage.

The 475-Ω resistors provide a proper output match for the filters. Since this resistor is unnecessary in the biasing scheme, it is ac coupled. Had it been dc coupled, a ½-W resis-

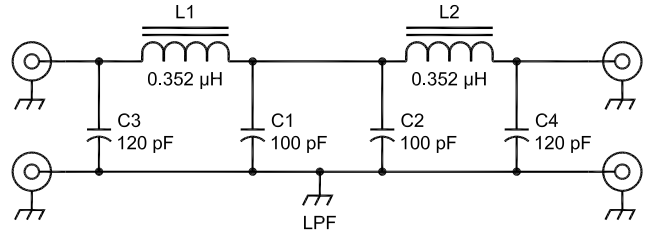

Fig 1—Filter schematics: (A) low-pass; (B) top-coupled band-pass; (C) shunt-coupled band-pass. Values shown are for the 20-meter filters. Component values for the other HF bands are available on the ARRL Web site. Resistors and capacitors are 0805 SMT unless noted. 0 pF capacitors are pads only.
LPF
L1, L2—0.352 µH, 8 turns #26 enameled wire on a T50-6 powdered-iron toroid core
Top-coupled BPF
D1, D2—DL4001
K1, K2—PC mount SPDT relay (Digi-Key #G5V-1DC5)
L1, L2—0.575 µH, 10 turns #26 enameled wire on a T50-6 powdered-iron toroid core
Shunt-coupled BPF
D1, D2—DL4001
K1, K2—PC mount SPDT relay (Digi-Key #G5V-1DC5)
L1, L2—0.657 µH, 12 turns #26 enameled wire on a T50-6 powdered-iron toroid core

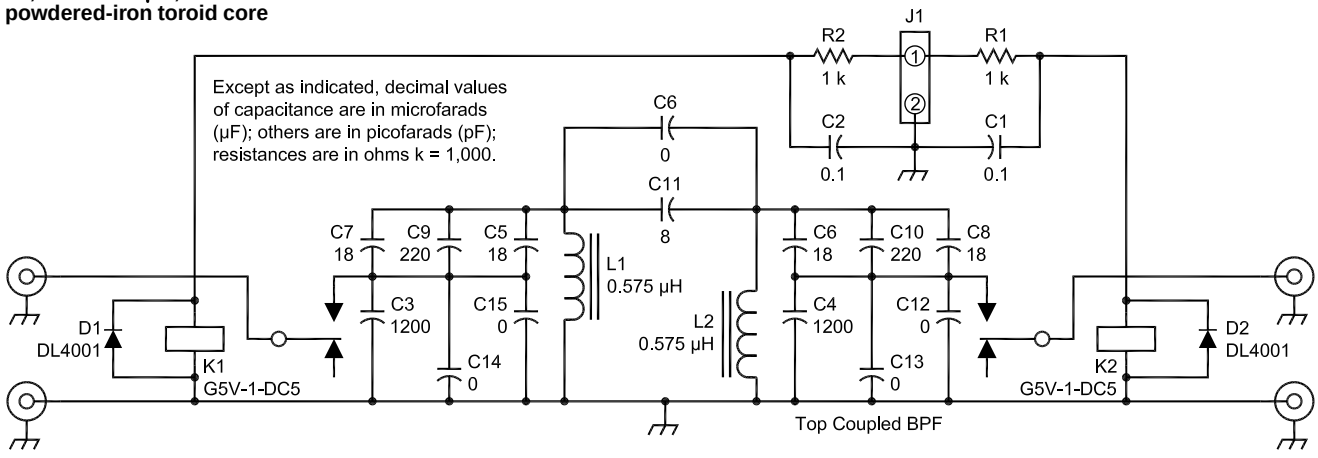

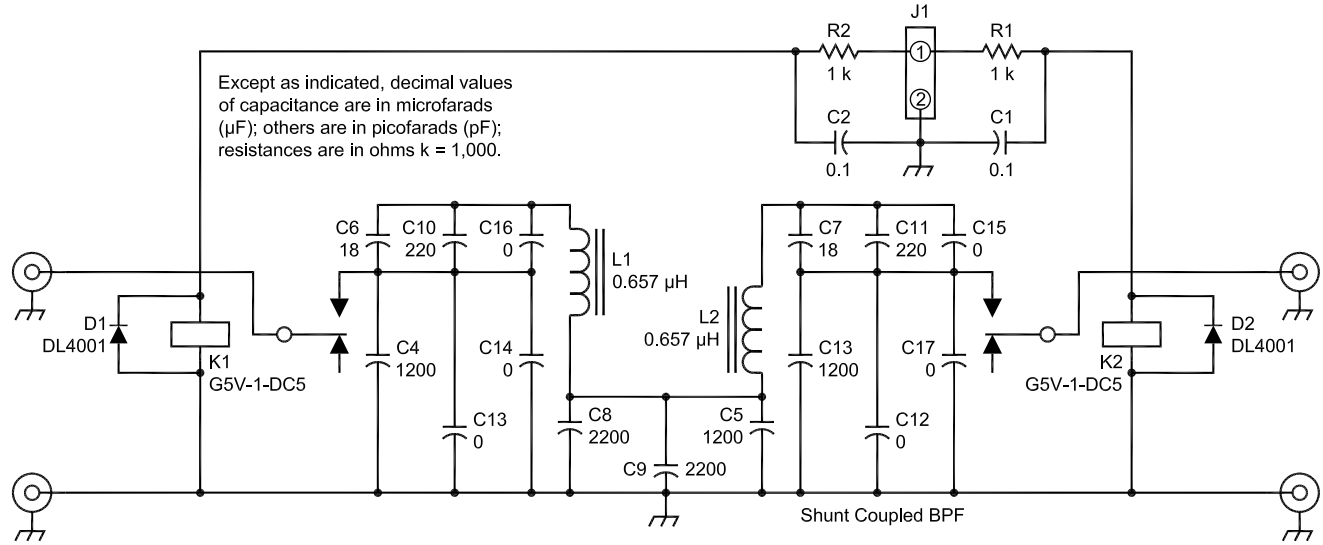

tor (expensive in chip form) would have been required. As it is, a 100-mW 0805 resistor is fine (and cheap). The output transformer matches the amplifier to 50 Ω. The impedance ratio also helps determine the amplifier gain. The output transformer is a conventional broadband transformer with a 16:1 impedance ratio (center tapped), wound on a BN3312-43 core.

The gain of the amplifier was measured to be flat within less than 2 dB across the HF spectrum, with the least gain at the high end. The input and output match also showed some degradation at the high end. Still, the SWR is 2:1 or better across the HF spectrum. One might get better performance using transmission-line transformers, but I did not try this. As is, the total gain of the front end at 20 meters was measured to be just over 6.8 dB, which is close enough to the 7 dB used in the spreadsheet in Part 1.

As shown in the block diagram (Fig 8 of Note 1), there are two RF amplifiers. One is used for the lower HF bands and both are used on the higher bands. The preamplifier used on the higher bands can be switched in or out using the relays on the board. I opted to leave the second RF amplifier in line at all times and therefore did not populate the relays or control circuitry. Jumpers were installed across the relay pads.

I have not yet measured the *IP3* of this amplifier, but my simulation in *Serenade SV*[4] gave better than +44 dBm. Because the reverse isolation is very high, this value should not suffer much even when terminated by the band-pass filters.

RF to Digital

At the heart of this receiver architecture is the AD6645 analog-to-digital converter. The design for the ADC board is quite straightforward, as shown in Fig 3. Notice that all inputs to the ADC are differential. This is necessary for all devices of this type if you are to get maximum performance. Even with much less-precise parts, the noise and distortion performance above a few megahertz will degrade if you use a single-ended input.

Local regulation is provided to reduce noise-pickup issues. Decoupling

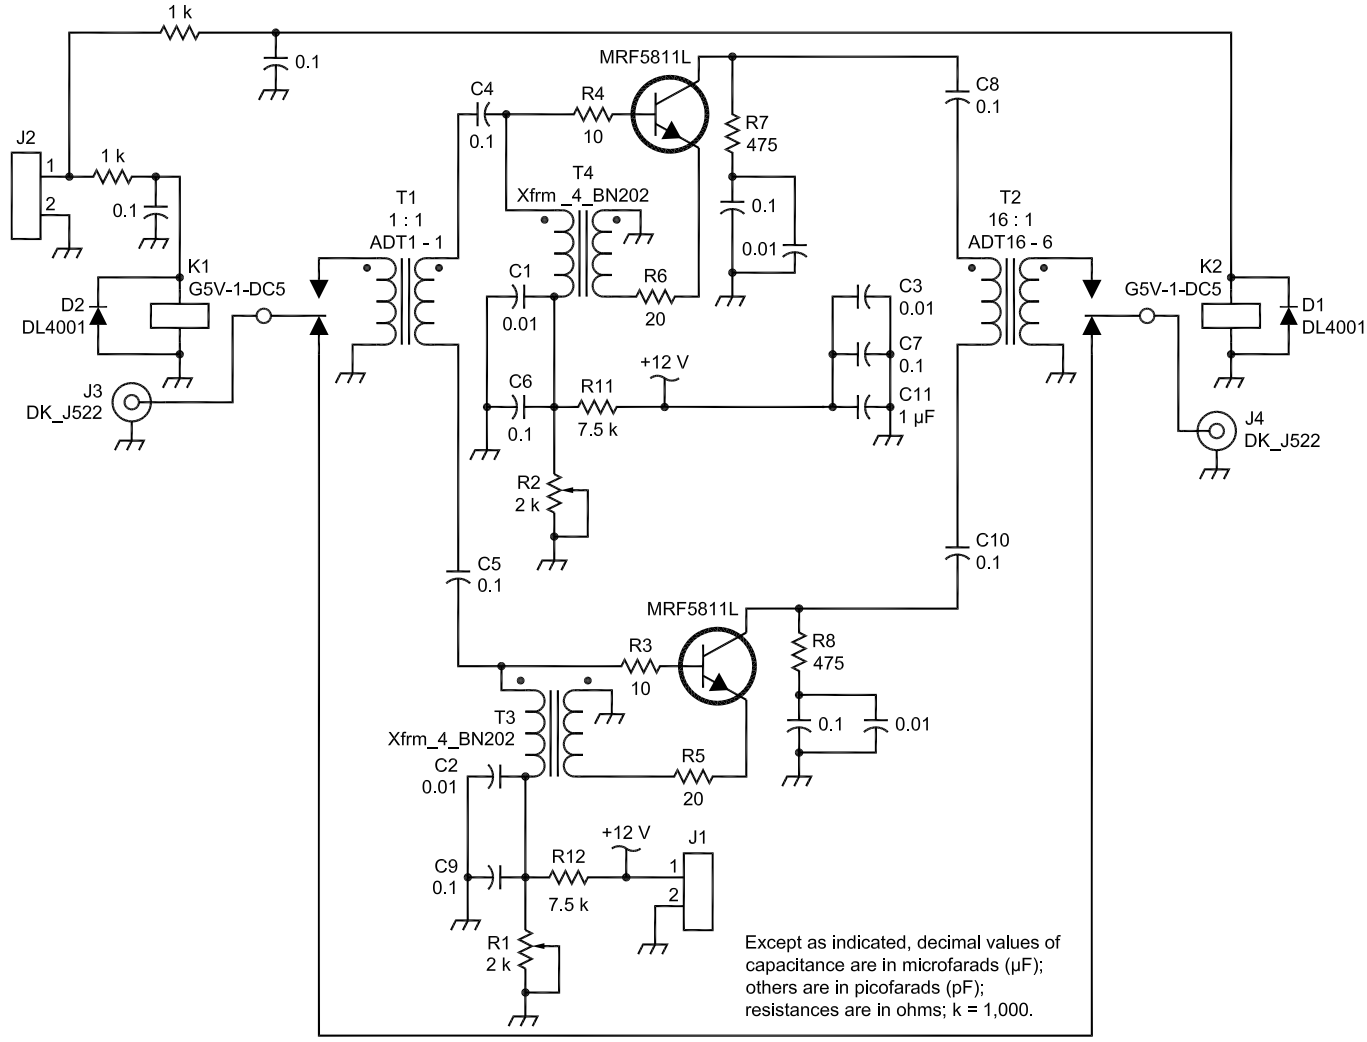

Fig 2—RF amplifier schematic diagram. Resistors and capacitors are 0805 SMT unless noted.

C11—1 µF, 25-V, X7R, 1206 SMT capacitor
D1, D2—DL4001
J1, J2—2-pin header
J3, J4—PC mount SMB bulkhead jack (Digi-Key #J522)
K1, K2—PC mount SPDT relay (Digi-Key #G5V-1DC5)
Q1, Q2—MRF5811 NPN transistor
R1, R2—2 kΩ, 10-turn trim pot (Digi-Key #3214W-202ECT)
R9, R10—56.2 Ω, 1210 SMT resistor
T1, T2—6 t primary, 3 t secondary on a BN202-61 binocular ferrite core
T3—1:1 transformer (Minicircuits ADT1-1)
T4—8 t center-tapped primary, 2 t secondary on a BN3312-43 binocular ferrite core

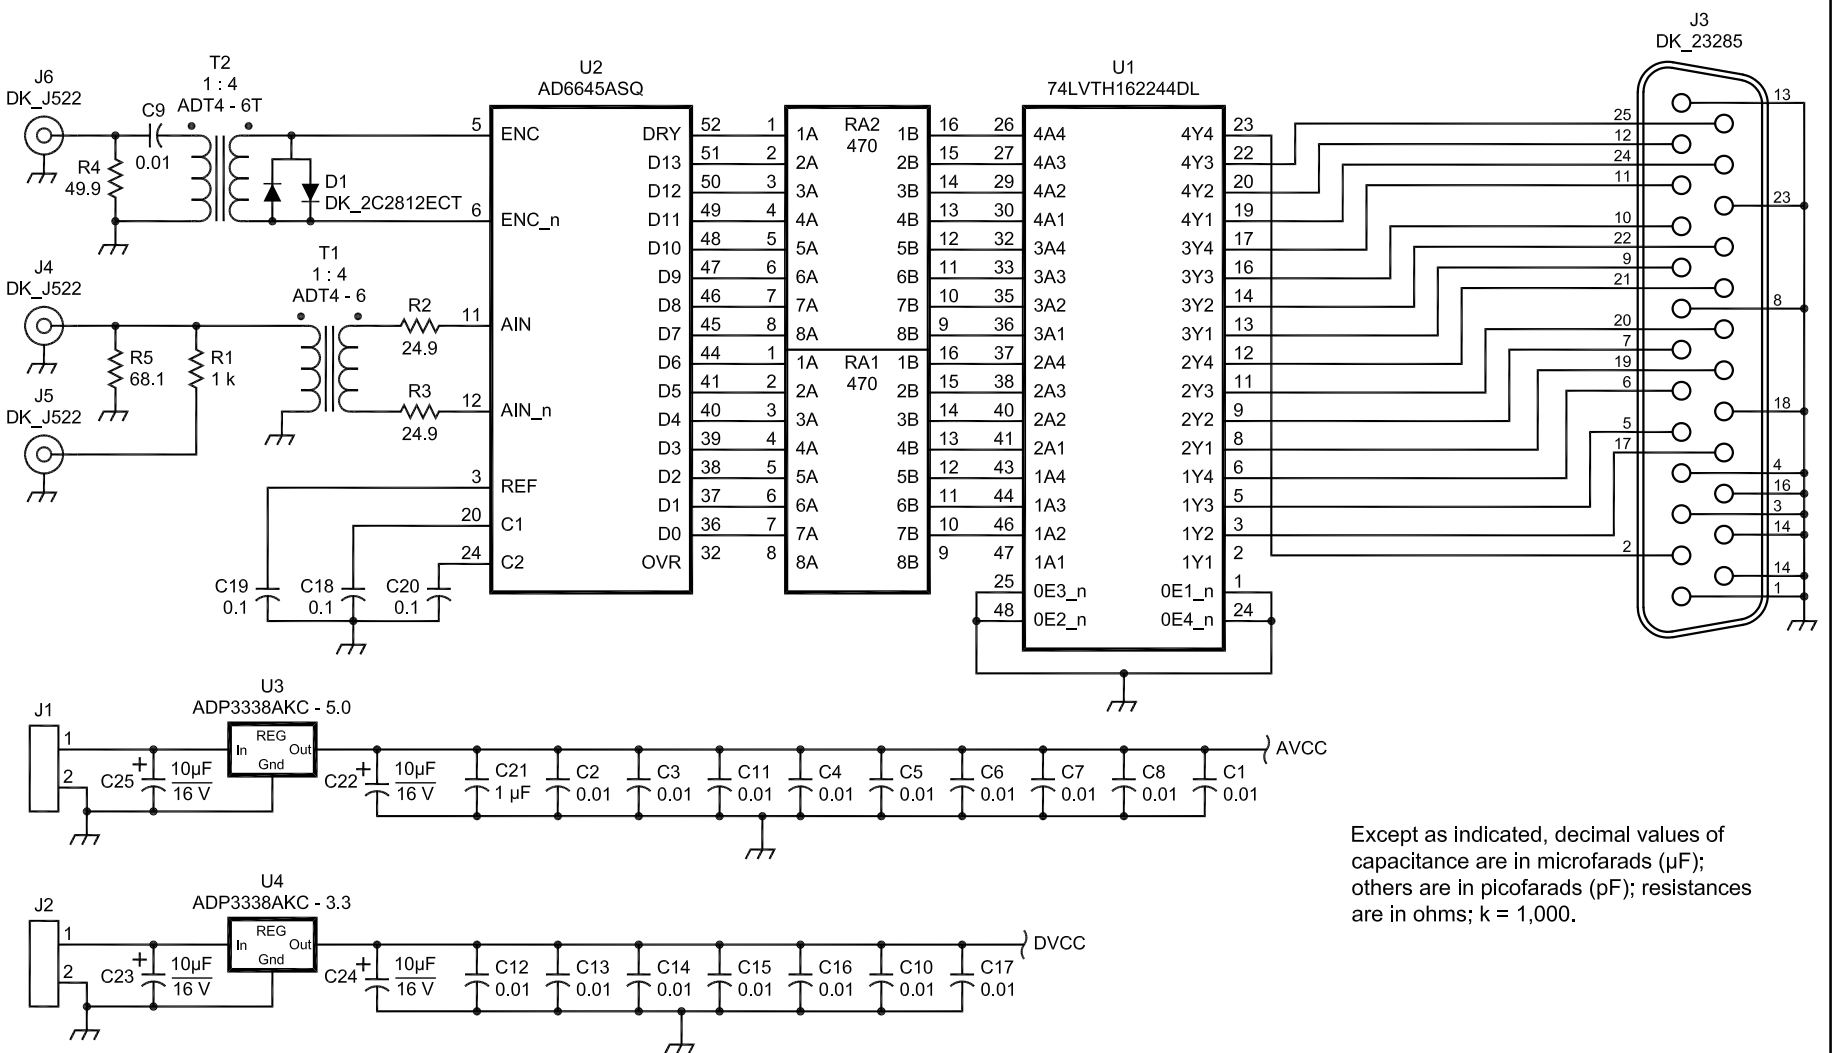

Except as indicated, decimal values of capacitance are in microfarads (µF); others are in picofarads (pF); resistances are in ohms; k = 1,000.

Fig 3—ADC schematic diagram. Resistors are 0805 SMT unless noted. Capacitors are 0603 SMT unless noted.
C9—0.01 µF X7R 0805 SMT capacitor
C21—1 µF, 16-V X7R 1206 SMT capacitor
C22-25—10 µF, 16-V tantalum capacitor
D1—Dual Schottky diode, series configuration, SOT23 (ZC2812E, BAS70-04)
J1, J2—2-pin header
J3—DB25 right-angle, PC mount (Digi-Key #A23285)
J4-6—PC mount SMB bulkhead jack (Digi-Key #J522)
RA1, RA2—Resistor array, 8×470 Ω (Digi-Key #742C163471)
T1—1:4 transformer (Mini Circuits ADT4-6)
T2—1:4 transformer, CT (Mini Circuits ADT4-6T)
U1—16-bit buffer with integrated termination resistors, SSOP, 74LVTH162244DL
U2—High-speed ADC, AD6645ASQ-80
U3—3.3-V low-dropout voltage regulator, Analog Devices ADP3338AKC-3.3
U4—5.0-V low-dropout voltage regulator, Analog Devices ADP3338AKC-5.0

capacitors are placed as closely as possible to each power pin. Small (0603) capacitors are used to allow placement on the same side of the board as the ADC. At this resolution and speed, even using vias to decoupling capacitors on the bottom of the board can degrade converter performance, according to the applications engineers at ADI.

The clock input is terminated to provide a 50-Ω match to the synthesizer. A 1:4 transformer converts the input to a differential signal, which is then clipped by the diode pair and applied to the ADC. As described in the AD6645 datasheet, the diode clipping prevents large voltage swings that can feed through to other parts of the device. It also helps reduce noise susceptibility.

Similarly, the analog inputs to the ADC are provided differentially via a transformer. The 24.9-Ω series resistors isolate the transformer from the ADC inputs. Notice that the 1-kΩ differential input impedance is not fully matched. Instead, the input signal is transformed up to 200 Ω. The result is that the full-scale input is about +4.8 dBm. Remember that the ADC is a voltage device—we don't need to worry about mismatch power losses. A termination resistor provides the proper match to the 50-Ω source.

The dither signal is also summed at the input node, via a 1-kΩ resistor. This impedance minimizes loading of the input and divides the signal down to the appropriate level. The datasheet recommends a dither power level of about –19 dBm. I thought about a simpler-looking idea where the dither would be injected at a center tap on the transformer secondary. However, as Brad Brannon, N4RGI, pointed out, this results in a balanced noise signal that negates the effectiveness of dither.

Output signals are series terminated to minimize dynamic currents at the output pins, which can reflect back and disrupt part performance. The termination resistors also help improve the quality of the signals at the buffer input. These termination resistors should be as close to the output pins as possible, which is why I used small surface-mount resistor arrays.

The output buffer was chosen for its timing characteristics: in particular, the skew between the bits (0.5 ns). The timing window between the clock and data signals at the RSP is very tight. Excessive variation between the RSP clock and the incoming data can result in bad data due to timing violations. The DRY signal from the ADC meets the timing requirements for the RSP clock, as long as additional variation is not introduced. By sending this signal through the same buffer as the data lines (since they only take up 14 of 16 bits), the variation is minimized and proper timing is guaranteed.

An important aspect of ADC performance is the quality of the encode clock and the aperture jitter of the converter itself. If the level of jitter on the clock or in the sample-and-hold circuit of the ADC is too high, it can directly affect the noise performance of the converter. This topic was discussed briefly in Part 1; but after further reflection, I felt that the explanation was somewhat confusing. Some of the feedback I received confirmed this. Therefore, a discussion is included in the sidebar "Jitter and ADC Performance."

A Low-Noise Clock

The ADC used in this architecture is subject to the same problems with a noisy local oscillator as a traditional analog design. In the traditional design, the problem is known as reciprocal mixing. Here, we have a degradation of the S/N caused by clock jitter. The result is the same, with noise sidebands appearing on the incoming signals. If the oscillator is noisy enough, these sidebands may mask a desired weak signal. This is discussed more thoroughly in the jitter sidebar.

One advantage of this architecture is that the tuning is done digitally within the RSP. Therefore, the "local oscillator" can actually be a low-noise crystal oscillator instead of the typical wide bandwidth VCO.

The oscillator design (Fig 4) is based closely on the design John Stephensen presented in *QEX*.[5] For the most part, I just altered components for the frequency difference, or for ease of procurement. The operating frequency is 64.96 MHz, which is phase-locked to a 12.992-MHz reference. The ADC frequency was chosen to allow integer decimation to both 40 kHz (for FM) and 16 kHz (for SSB and CW).

The key to John's design is the input level in the JFET. In this configuration, the impedance of the FET is fairly high (50-100 Ω). Therefore, since the same RF current flows in the resonator and FET, the oscillator input power is higher than the power dissipated in the resonator. For a given resonator power, most common oscillator designs extract much less. The higher power leads directly to lower phase noise, per Leeson's equation.

Approximately +8 dBm is available to the ADC after the 3-dB attenuator. A buffered output from the emitter follower is also attenuated to provide the correct level to the PLL chip. The PLL used is the ADF4001, which is designed specifically for clock applications below 200 MHz. There is no prescaler and there are no minimum division ratios for either the RF or reference frequency.

Other advantages of this device are the very low phase-noise floor and that the phase/frequency discriminator (PFD) can be used up to 100 MHz. In this application, we'll use it at 12.992 MHz with no reference division. The higher frequency in the PFD helps lower the noise floor within the loop bandwidth. For example, doubling the reference frequency increases the PFD noise by 3 dB, but the multiplier noise is reduced by 6 dB. Here, the PLL noise floor is about –128 dBc/Hz, and the VCXO noise will dominate outside the loop bandwidth of about 230 Hz. Close-in (inside about 100 Hz), the noise will be determined by the reference oscillator. The jitter sidebar provides additional details for the predicted clock performance.

Cheap and Easy Dither

There are two main ways to generate large-scale dither for the ADC. One that is commonly used in commercial applications—such as test equipment—involves digitally generating pseudorandom noise. This wide-band noise is then injected into the front end of the ADC, and subsequently subtracted from the digital result. The second alternative is to inject narrowband noise that is outside the bands of interest. This method is used here.

The dither circuit is shown in Fig 5. It is based on a circuit presented in the manual for Analog Devices' *High Speed Design Techniques* seminar. I opted to use op amps throughout instead of the variable-gain amplifiers presented in that discussion.[6] I also used an op amp with a higher input impedance (the AD8055), which allows me to use a larger resistor for my noise

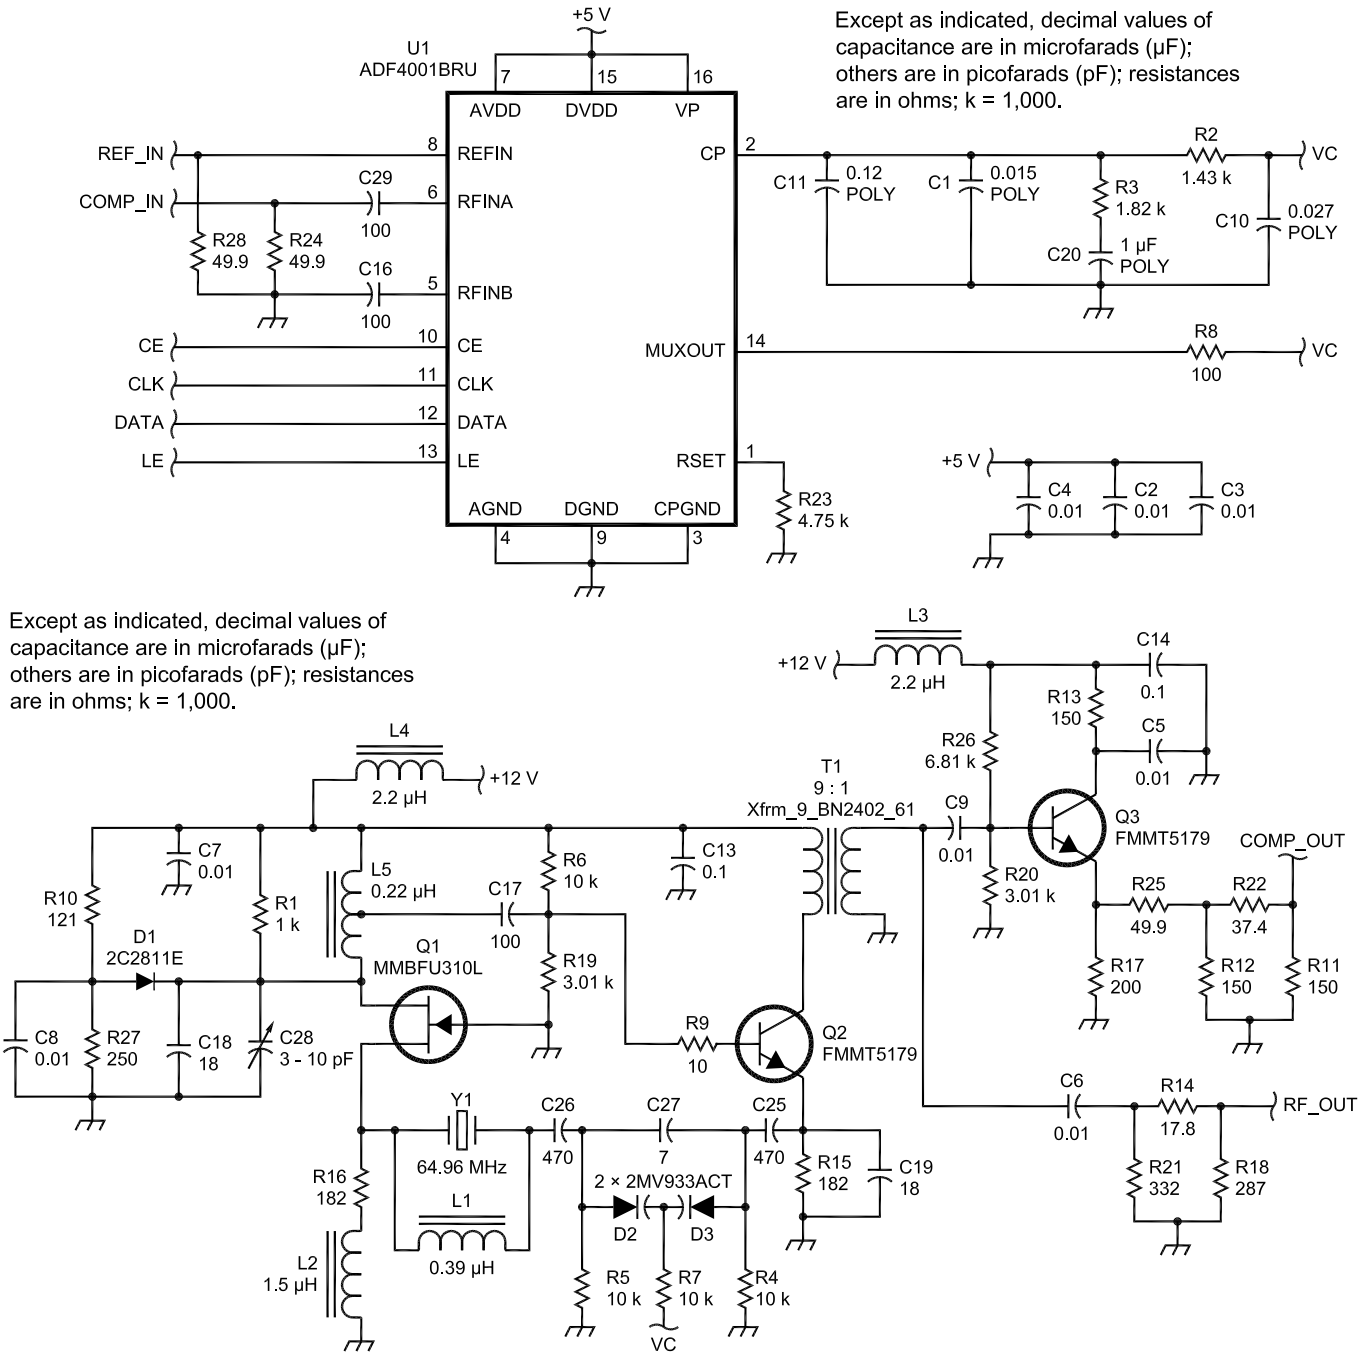

source. This requires less total gain and reduces the variation in the noise levels, as the resistor noise dominates. Additional variation in generated noise can be caused by part-to-part differences in the op amp's current noise. The third amplifier stage has adjustable gain to set the proper level.

The noise generator (first three stages) is followed by an active low-pass filter. This filter has a bandwidth of about 450 kHz; with six poles, it reduces the injected noise below the noise floor at 160 meters. Within that constraint, the bandwidth is enough to minimize the gain required to achieve the correct noise power, and makes sure the noise is random rather than sinusoidal. The output of the final filter stage drives a buffer that has a gain of two.

The AD8055 amplifier is a high-speed device, with a 3-dB bandwidth of over 300 MHz. This is helpful in two ways. First, with the high gains involved, the actual bandwidth of the amplifiers ends up being only a few megahertz. Second, the output impedance of an op amp increases with frequency. This allows the noise signal a forward path through the filter feedback capacitors, since the op amp output is no longer an ideal, zero-impedance voltage source. The result is that the ultimate rejection of the filters is limited. The AD8055 output impedance remains very low into the VHF range, thus minimizing this effect.

Receive Processing Hardware

The hardware implementation of the Receive Signal Processor (RSP) is fairly simple. See Fig 6. At this point, the sig-

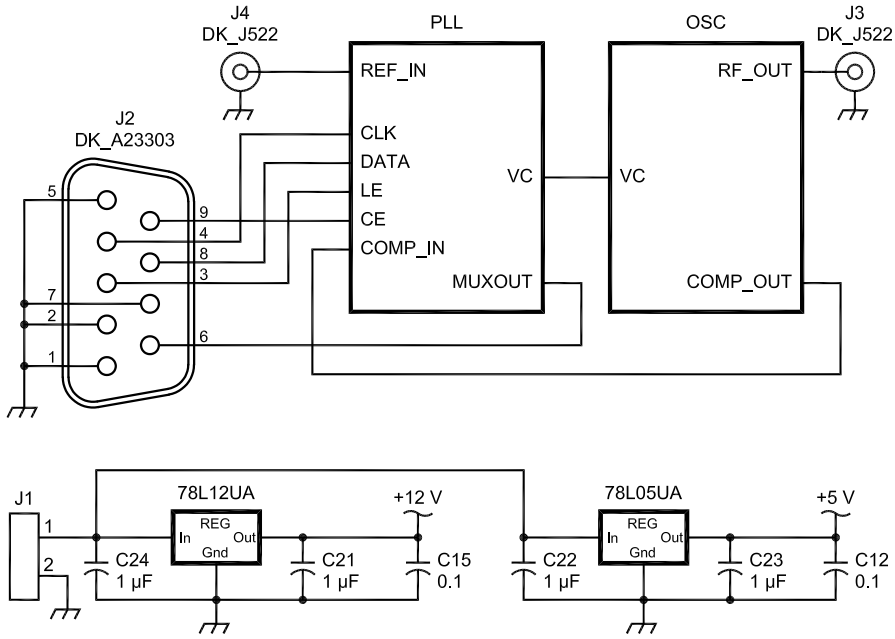

Fig 4—Receive PLL schematic diagram. Resistors and capacitors are 0805 SMT unless noted.

C1—0.015 µF polyester capacitor (Digi-Key #P10957)
C2-4—0.01 µF 0603 SMT capacitor
C10—0.027 µF polyester capacitor (Digi-Key #P10960)
C11—0.12 µF polyester capacitor (Digi-Key #P10967)
C20—1 µF polyester capacitor (Digi-Key #P10979)
C21, C24—1 µF, 25-V, X7R, 1206 SMT capacitor
C22, C23—1 µF, 16-V, X7R, 1206 SMT capacitor
C28—3-10 pF trimmer capacitor, SMT (Digi-Key #SG2002)
J1—2-pin header
J2—DB9 right angle, PC mount (Digi-Key #A23303)
J3, J4—PC mount SMB bulkhead jack (Digi-Key #J522)
L1—0.39 µH, 1210 SMT (Digi-Key #M6100)
L2—1.5 µH, 1210 SMT (Digi-Key #M6107)
L3, L4—2.2 µH
L5—12 turns #28 enameled wire on a T37-12 powdered-iron core, tap 3 turns from cold end
T1—primary 6 turns #28, secondary 2 turns #28 on a BN2402-61 binocular ferrite core
U1—PLL, ADF4001BRU
U2—5-V voltage regulator, SOT89, 78L05UA
U3—12-V voltage regulator, SOT89, 78L12UA
Y1—64.96 MHz, third-overtone (International Crystal)

nal has been digitized, so we only need be concerned with signal integrity, primarily on the input side. At the output, the serial interface only needs to be fast enough to get the I/Q words out at a maximum 40-kHz rate for DSP processing by another computer. Signal integrity is easy at this speed, but we still use proper buffering and a series termination to keep everything clean.

The RSP is initialized and controlled using the parallel *Microport* interface. The serial interface can be used for control after setup, but cannot be used for initialization. Since the parallel interface is necessary anyway, I thought it would be simpler to use it exclusively and to use the serial port only for the signal path. I also wanted to use the DSP for signal processing only, not for housekeeping. The Microport interface is set to mode 0 in hardware by grounding the **MODE** pin.

Grounding the **PAR/SER** pin selects the serial interface for output data. The serial word length is set to 24 bits, which transfers all of the data but minimizes overhead. The RSP is the serial master, so control signals are generated in the RSP and sent to the DSP serial port. The serial clock is set to its minimum value, which meets the bandwidth requirements for a 40-kHz output using 24-bit I/Q words.

The actual use of the RSP functions is more closely related to the DSP than to the analog front end. Therefore, I will spend more time in the DSP installment discussing RSP operation than I will here. The RSP will be set up to provide I/Q data at a sampling rate of 16 kHz for SSB/CW and 40 kHz for the wider-bandwidth modes (AM/FM). The goal was to get better than 100-dB out-of-band rejection. The high decimation rates make this a challenge in the RSP, but it can be done with the chosen output rates. I'll go into my thoughts in this area more in the DSP installment.

Summary

In this installment, we looked at the design details for an analog front end that should meet the requirements laid down in Part 1. Some of the ideas have been seen in these pages before. The key, as always, is to ensure that the tradeoffs are managed to allow maximum performance.

The AD6645 ADC is at the heart of the receiver architecture. This part has excellent noise and distortion performance, which allows this architecture to work. Low aperture jitter helps to minimize the ADC equivalent of reciprocal mixing, further enhancing performance. Further improvements of ADCs in the future will allow even better system performance, and the modular design will allow us to take advantage of this. For now, however, the performance available from the AD6645 is the state of the art.

In the next installment, we will look at the transmit side of the transceiver. Just as the receiver has been helped by improvements in ADCs, newer high-speed DACs allow excellent transmitter performance. The details of integrating them into the transmitter will be covered in Part 3.

Acknowledgements

I would like to thank Doug Smith, KF6DX; Paul Smith and Gary Hendrickson for their input on topics related to this article.

Notes

[1]J. Scarlett, KD7O, "A High-Performance Digital Transceiver Design, Part 1," *QEX*, Jul/Aug 2002, pp 35-44.
[2]W. Sabin, WØIYH, "Narrow Band-Pass Filters for HF," *QEX*, Sep/Oct 2001, pp 13-17.
[3]J. Stephensen, KD6OZH, "The ATR-2000: A Homemade, High-Performance HF Transceiver, Part 2," *QEX*, May/Jun 2000, pp 39-51.
[4]Available at **www.ansoft.com**.
[5]J. Stephensen, KD6OZH, "A Stable, Low-Noise Crystal Oscillator for Microwave and Millimeter-Wave Transverters," *QEX*, Nov/Dec 1999, pp 11-17.
[6]Analog Devices, *High Speed Design Techniques*, 1996, pp 5.41-5.48.

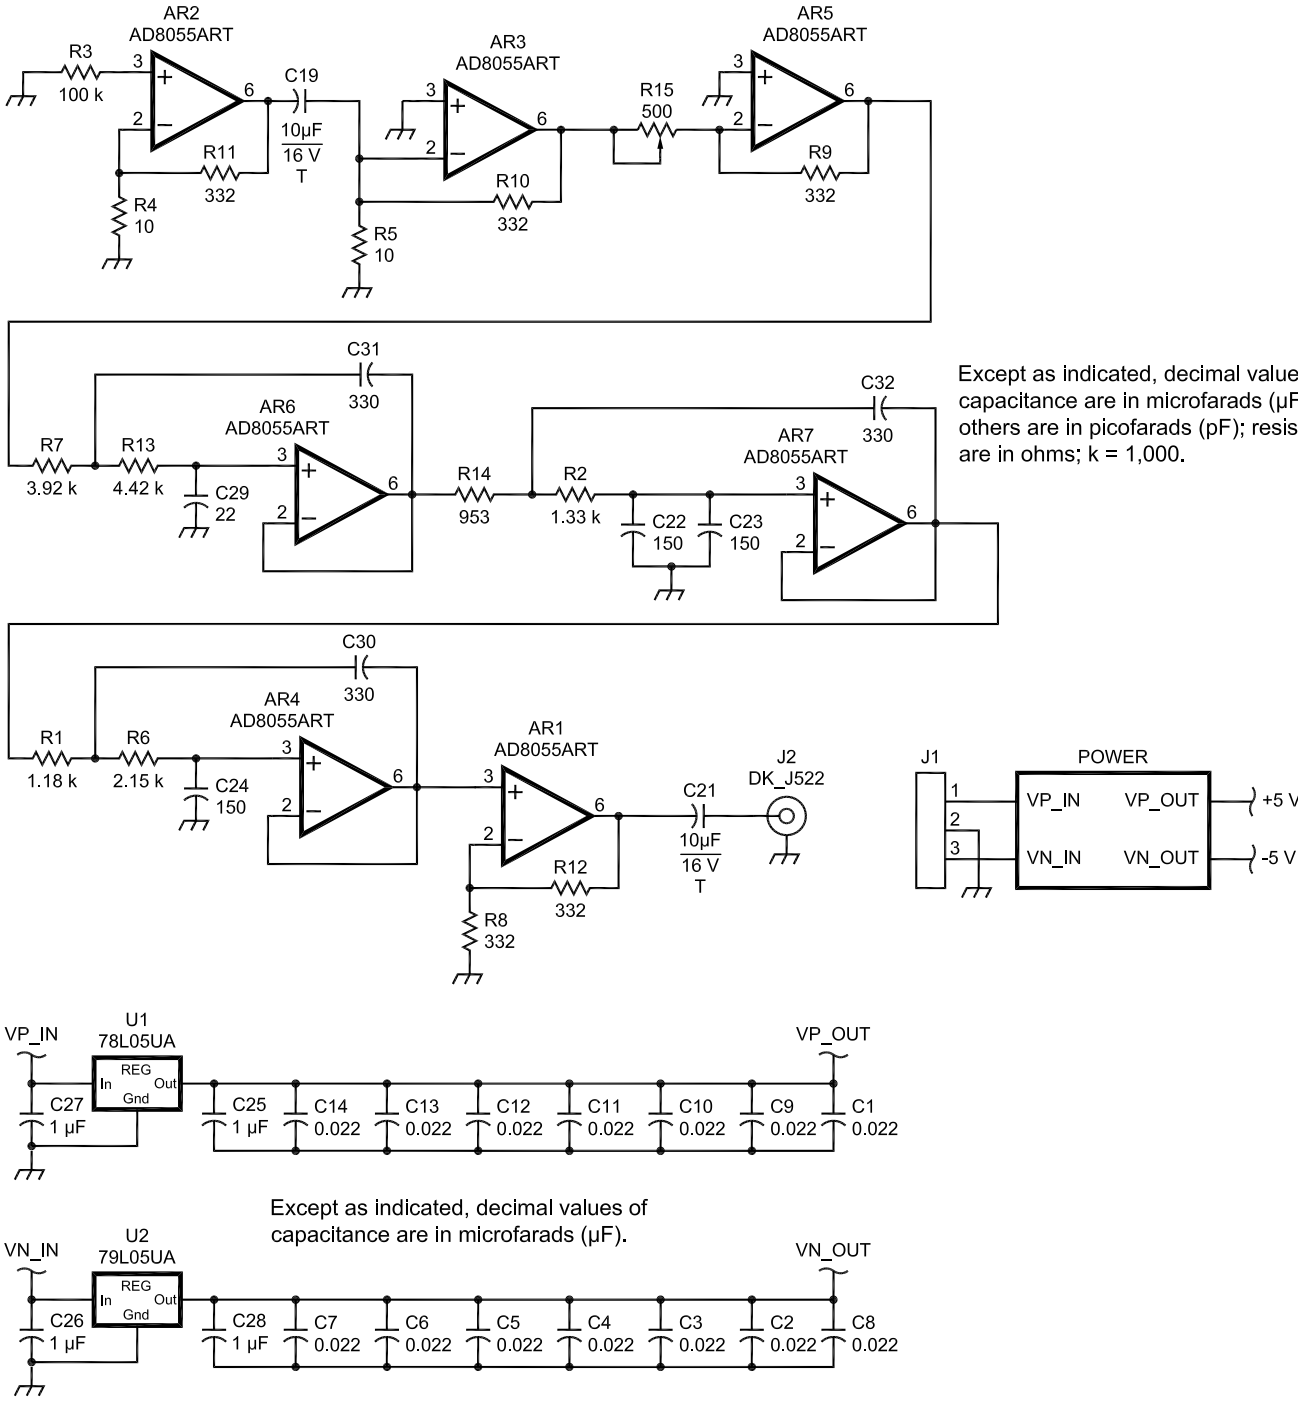

Except as indicated, decimal values of capacitance are in microfarads (µF); others are in picofarads (pF); resistances are in ohms; k = 1,000.

Except as indicated, decimal values of capacitance are in microfarads (µF).

Fig 5—Dither-generator schematic diagram. Resistors and capacitors are 0805 SMT unless noted otherwise.

AR1-7—High-speed op amp, AD8055ART
C19, 21—10 µF 16-V tantalum capacitor
C25-28—1 µF, 16-V, X7R, 1206 SMT capacitor
J1—3-pin header
J2—PC mount SMB bulkhead jack (Digi-Key #J522)
R15—500 Ω, 10-turn trim pot (Digi-Key #3214W-501ECT)
U1—5-V voltage regulator (positive), SOT89, 78L05UA
U2—5-V voltage regulator (negative), SOT89, 79L05UA

Fig 6(right) —Receive Signal Processor schematic diagram. Resistors and capacitors are 0805 SMT unless noted.
C9, C10—10 µF, 16-V tantalum capacitor
J1—2-pin header
J2, J3—DB25 right angle, PC mount (Digi-Key #A23312)
J4—DB9 right angle, PC mount (Digi-Key (#A23303)
RA1—Resistor array, 4x100 Ω (Digi-Key #742C083101)
U1—Quad buffer, 74ALVC125PW
U2—RSP, AD6620AS
U3—3.3-V low-dropout voltage regulator, Analog Devices ADP3338AKC-3.3

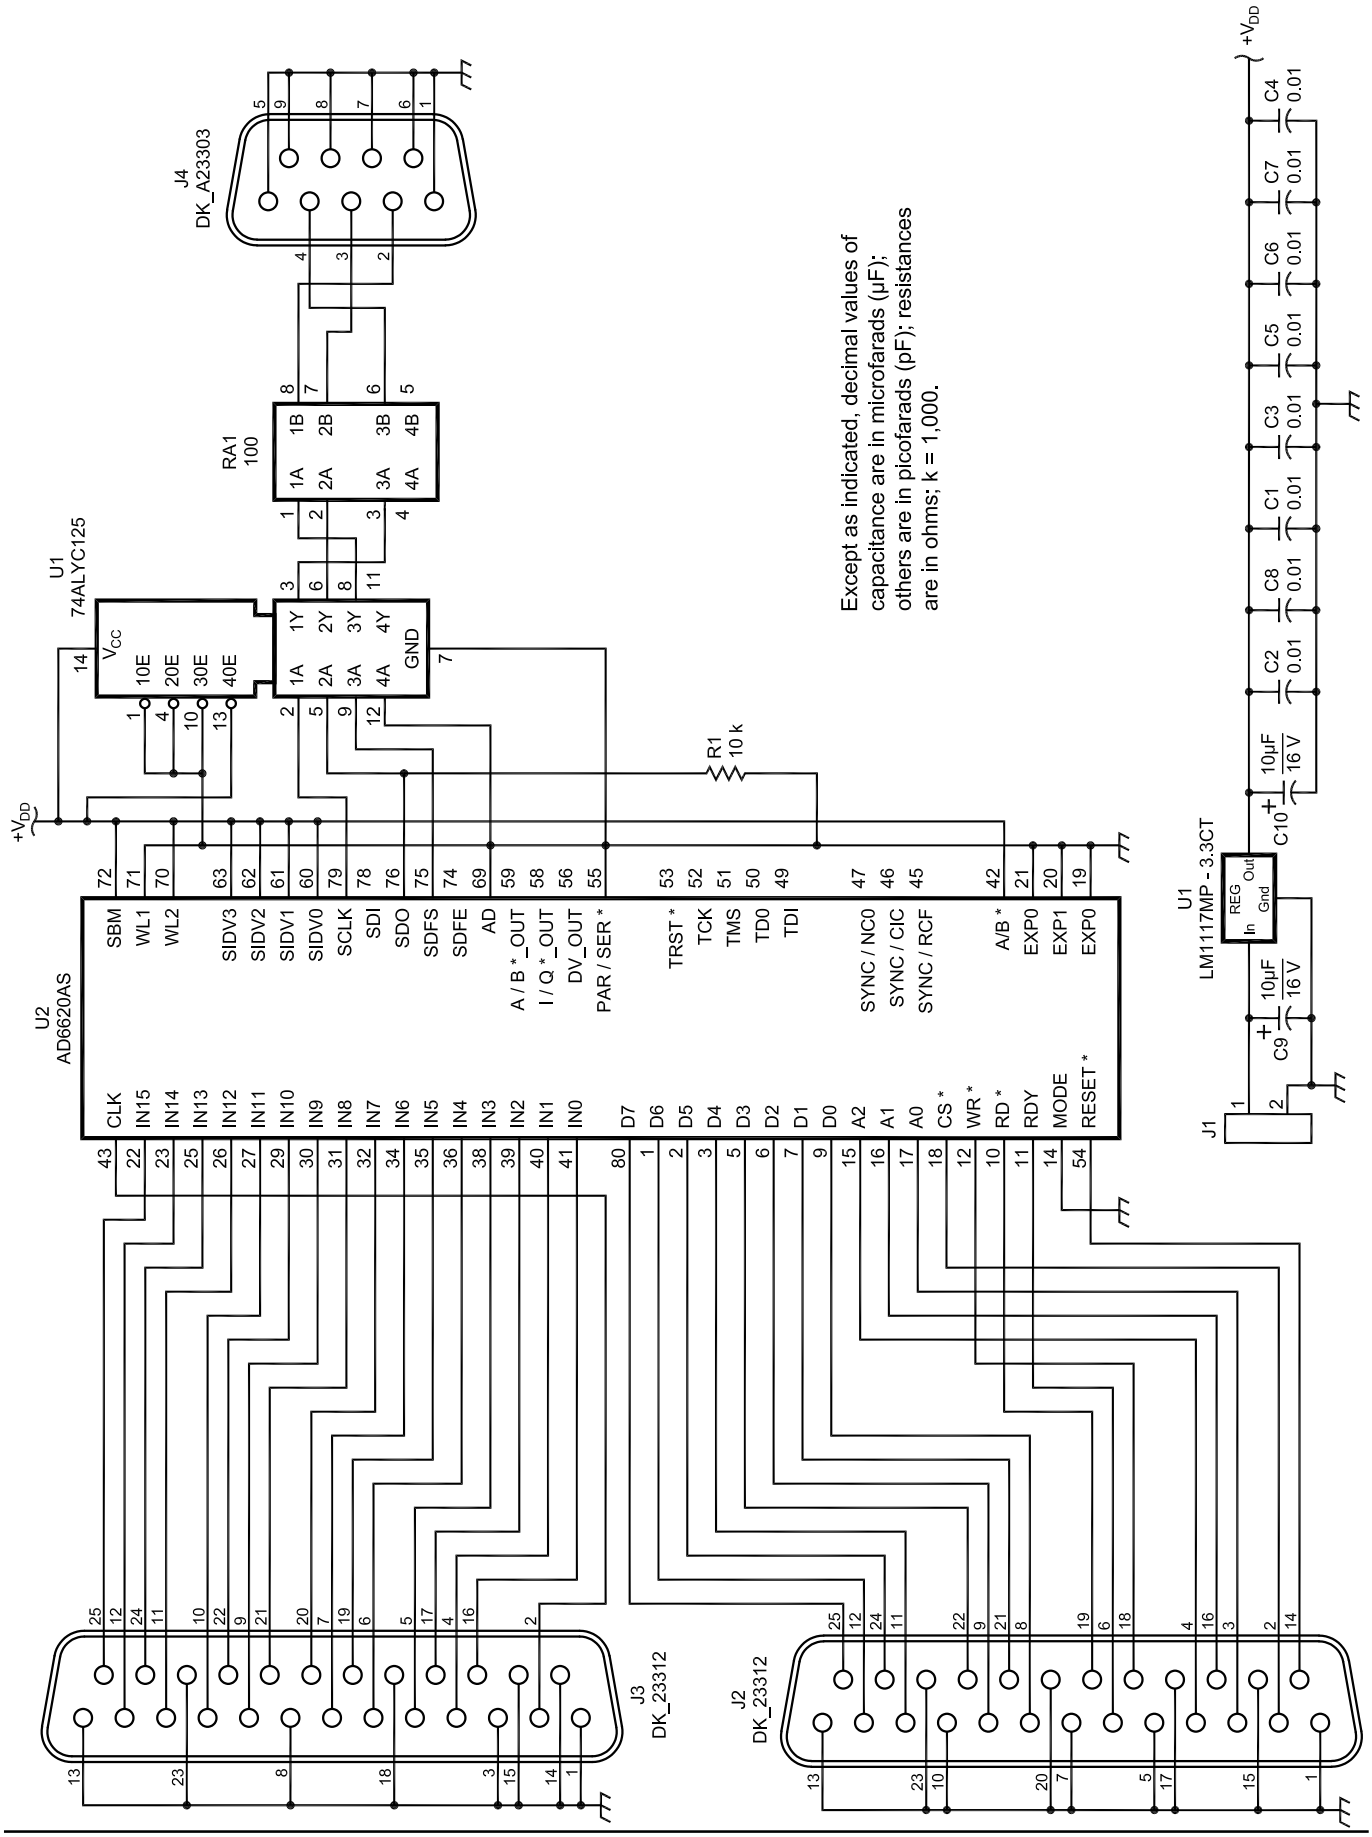
Except as indicated, decimal values of capacitance are in microfarads (µF); others are in picofarads (pF); resistances are in ohms; k = 1,000.
U2
AD6620AS
U1
74ALYC125
U1
LM1117MP - 3.3CT
RA1
100
R1
10 k
J1
J2
DK_23312
J3
DK_23312
J4
DK_A23303
+VDD
C1
C2
C3
C4
C5
C6
C7
C8
C9
C10
10µF
16 V
0.01

Phase Noise and ADC Performance

There was some confusion from the discussion of clock jitter in Part 1 of this series. This led me to reevaluate how the material was presented and is the reason for this sidebar. One of the problems was making a connection between phase noise and clock jitter. Therefore, this discussion will be primarily focused on phase-noise requirements, and I will only discuss the relationship with clock jitter as necessary.

Rather than generating arbitrary jitter specifications, let's look at actual system requirements and translate this to a clock phase-noise requirement. Measurements are usually termed "noise-limited" when reciprocal mixing has increased the system noise figure by 1 dB. In our case, this would happen if the SNR of the converter increases by 1 dB.

Initial measurements with an evaluation board for the AD6645 showed a value for the SNR at 30 MHz of about 75 dB. We'll use this number for our calculations. We'll also do the calculations for 10 meters, since Eq 2 in Part 1 shows that jitter has a greater effect on higher frequency signals. For this discussion, we'll ignore that jitter on the clock or input signal during the SNR measurement may have had a significant contribution to the 75 dB SNR and assume this measurement was "jitterless."

In order to keep the SNR from increasing by 1 dB, the SNR contribution from jitter must be better than 81 dB. Thus, over the Nyquist bandwidth ($^1/_2$ of the sampling rate), the integrated noise due to jitter must be less than –81 dBc. For the purposes of this discussion, the carrier is assumed to be full-scale (approximately –13 dBm at the receiver input on 10 meters).

Using Eq 1 from Part 1 (processing gain), the required noise applied by the clock to the incoming signal must be less than –156.1 dBc/Hz. The clock rate is 64.96 MHz in this calculation. However, this is not the required clock noise performance, because of the relationship

$$SNR = 20 \log[\omega_{clk} / (\sigma_{clk} \omega_{sig})] \qquad \text{(Eq 1)}$$

where σ_{clk} is the frequency jitter of the clock and equals $\sigma_t \omega_{clk}$. This relationship shows that if the potential interfering signal is lower in frequency than the clock, the noise sidebands are improved based on the ratio. Likewise, noise sidebands applied to a signal that is higher in frequency than the clock are worse than the clock itself.

Thus, for a 64.96-MHz clock and a 28.5-MHz incoming signal at full scale, the clock noise requirement is –149 dBc/Hz to meet our noise specification. That is, the average noise level integrated over our desired bandwidth at a given offset from the interfering signal must be better than –149 dBc/Hz. Assuming a constant sideband slope (not true at PLL corner frequencies), that would mean that the noise must meet this specification at the center of the desired passband. We can use this to determine how our clock's predicted performance measures up.

Also, notice that the above requirements are for 10 meters. The requirements would be less stringent at 20 meters (about 6 dB), because jitter has a smaller effect at lower frequencies. Likewise, the requirements would be much more difficult at 2 meters (about 14 dB). This is offset by the fact that close-in signals tend to be much stronger at HF than on 2 meters. This is important as I consider whether to adapt this radio for 2 meters (undersample the third Nyquist zone) or to simply use a transverter.

On a separate note, by solving Eq 1 for jitter, we find that the jitter requirement is 0.5 ps at 10 meters, not the 0.1 ps example used in Part 1. At 2 meters, meeting the same specifications would require 0.1 ps jitter.

Clock Performance

Table 1 shows the expected performance of the receive PLL. The VCXO and reference were modeled using 'typical' expected values for crystal parameters (except the maximum *R*s of 40 Ω for the VCXO crystal and 25 Ω for the reference crystal) and the maximum specification for FET resistance. This should give a reasonably conservative model. I did the model using Leeson's equation on an *Excel* spreadsheet.

The noise floor within the PLL was determined using the ADF4001 datasheet. The specified noise floor for the PFD is –153 dBc/Hz at a comparison frequency of 1 MHz. Using a reference frequency of 12.992 MHz (with a 10 log relationship) and a multiplication of 5 (with a 20 log relationship) results in a PLL noise floor of approximately –128 dBc/Hz.

At offset values above the loop frequency of 230 Hz, the VCXO noise dominates, and below about 100 Hz, reference noise does. Between these offsets, the PLL is the primary noise source.

The table shows that we meet the required noise specifications at an offset of about 1200 Hz. This means that a *full-scale* signal outside of the 2400 Hz SSB passband will have no significant reciprocal mixing effects. For a 500-Hz CW bandwidth, a full-scale signal offset less than 1200 Hz from the center of the passband can cause reciprocal mixing, but an S9+40 signal outside the passband would not cause noticeable degradation. For the lower HF bands, these results are even better, while at 2 meters, a 5 kHz offset would be required for no significant effects from a full-scale signal.

This example shows a distinct advantage to this receiver architecture with respect to reciprocal mixing. Since the tuning is done digitally, we can use a low-noise crystal oscillator that drastically reduces reciprocal mixing effects. Between this and the outstanding noise and jitter performance of the AD6645 converter, we can virtually eliminate reciprocal mixing as a problem in our HF receiver. This example also demonstrates why reciprocal mixing from close-in signals is so difficult to tackle in receivers with wide-band, general coverage VCOs.

Table 1—Predicted Receive PLL Phase-Noise Performance

Offset (Hz)	*Phase Noise (dBc/Hz)*
10	–95
20	–104
50	–115
100	–125
200	–128
500	–138
1 k	–147
2 k	–155
5 k	–163
10 k	–167
20 k	–169
50 k	–170
100 k	–171