

A Next-Generation R2 Single-Signal Direct Conversion Receiver

Adapted from a paper given at FDIM, Dayton Ohio, May 2000 by Rick Campbell KK7B and Bill Kelsey N8ET

In the 8 years since publication of the “High Performance Direct Conversion Receiver” Series of articles in QST hundreds of copies of the described receivers have been built, and a number of experimenters around the world have extended the concepts in various directions. In parallel with this activity in the amateur arena, professional designers have embraced the direct conversion architecture for every application from very low-cost pager and family radio service receivers, to high-end surveillance equipment like the Rockwell-Collins 95S-1. This paper reviews the R2 and miniR2 direct conversion receiver circuits and presents some new circuitry and construction techniques to improve R2 performance beyond the 1992 benchmarks.

Review of Single-Signal Direct Conversion Receivers

The R2 receiver was developed in parallel with the R1 “High Performance Direct Conversion Receiver,” described in August 1992 QST. All of the basic circuitry from the R1 was duplicated onto the R2 circuit boards, with appropriate additions for eliminating the undesired sideband. The audio quality of the R1 remains a benchmark for amateur receivers. The R2 sounds good, but summing two channels with different time delays (as required by the image-reject circuitry) modifies the impulse response of the channel, and the receiver loses some of its presence. This is exactly the same effect one encounters with a SSB bandwidth crystal filter in a conventional superhet.

After several hundred R2 receivers had been built, a second-generation circuit was developed: the miniR2. The miniR2 circuit board is half the size of the original R2, and has only headphone output. In addition, the miniR2 circuitry is simplified and made more tolerant of component variations, so that good performance may be obtained without hand-matching the audio diplexer components. The audio filtering of the miniR2 was simplified to fit all of the parts on the small circuit board, but the audio quality is still outstanding when using headphones or an external audio power amplifier.

Many experimenters have used the basic R2 and miniR2 circuitry as the foundation for experiments using DDS frequency synthesizers and DSP audio signal processing, as suggested in the original QST articles. At KK7B, a dozen different R2 and miniR2 receivers and transceivers have been used in a wide variety of fixed and portable applications, sometimes with outstanding results, and sometimes immediately indicating directions for further work.

With all this shared experience, it was determined that the time had come to do a revision of the original R2 board. Several different versions have been built—but the requirement that the revised version work better than the original is tough to meet. The original R2 circuit and circuit board were optimized over a period of more than a year of continuous activity. No revision has been attempted to the miniR2 board—it still performs exceptionally well in its intended application, and often provides better image suppression with off-the-shelf parts than the R2 does with hand-matched components.

Updating the R2

The first task in updating the R2 circuit was to determine what needed to change. A wish list was formulated:

1. Replace the SBL-1 mixers with the TUF-3 package.
2. Replace the LM 387 audio IC with a modern low-noise dual op-amp
3. Revise the input duplexers for better tolerance to component variation
4. Improve the opposite sideband suppression
5. Improve the receiver system noise figure
6. Remove some direct conversion receiver quirks
7. Make it easier to build advanced experimental receivers
8. Design a receiver circuit that rewards component selection with performance
9. Make the receiver circuit a definite step-up from the miniR2
10. Improved muting and gain step options

The new receiver was named the R2pro. The philosophy is that the R2pro trades more expensive construction, more expensive components, component matching, design flexibility, and a higher level of builder knowledge and experience for slightly improved performance over the miniR2. The miniR2 circuit should be used for most “High-Performance Signal-Signal Direct Conversion Receiver” applications. The R2pro is for builders who want to go to the extra effort and expense required to push the receiver to the limits of the direct conversion architecture.

There is a significant problem with direct conversion receivers built on a single circuit board. Proper RF grounding and shielding are very different than the grounding and shielding techniques needed for high-gain audio amplifier circuitry. If the low-level RF signals, high-level LO signal, all the mixer conversion products, and high-gain audio amplifier are all on the same circuit board, there must be compromises in grounding and shielding. These compromises were handled on the R1, R2 and miniR2 boards by designing the ground traces such that the audio stages saw an approximate single-point-ground and the area around the mixers was an unbroken ground plane. Any of the receivers can be made to oscillate by connecting the power-supply or speaker ground wire to the wrong point on the circuit board ground, even though all of the grounds are connected together. For a review of audio grounding techniques, see Horowitz and Hill, “The Art of Electronics.”

The conflicting requirement for an RF tight enclosure and a single-point audio ground makes it difficult to package single board direct conversion receivers. Early versions of the R1 and R2 direct conversion receivers pictured in QST were enclosed in soldered-up copper-clad PC board enclosures. Other packages, particularly those made of aluminum pieces held together with screws—are prone to intermittent audio oscillations and microphonics. There are more grounding and shielding options if the PC board is broken up into functional blocks. *Then the PC board with the mixers can be completely shielded, and the PC board with the audio output amplifier can have a single point*

ground. By optimizing the gain partitioning and packaging of the receiver, hum and microphonics can be eliminated and the placement of ground connections becomes much less critical. As a fringe benefit, breaking up the PC board makes it easier to build experimental versions using DSP, different mixers, audio processors and power amplifiers etc.

The R2pro circuit description

The R2pro block diagram is shown in figure 1, and the schematics of the individual PC boards are shown in Figures 2 through 6. The R2pro system includes an RF preamp. For MF through VHF, the grounded-gate circuit is appropriate. Direct conversion receivers need low-pass filtering on the input, to attenuate signals near odd harmonics of the LO. The preamp incorporates the low-pass filtering in the input and output match to the transistor. The bias switch is part of the receiver mute circuit, and switches the amplifier gain between +13 dB and -40 dB. Unlike most amplifier circuits, this one is a strong 40 dB attenuator when it is in low-gain mode, and can be switched in as a front-end attenuator when very strong signals are present, without introducing front-end distortion. The amplifier noise figure of approximately 2 dB and the relatively low gain of the preamp stage have the effect of reducing the receiver noise figure without severely impacting two-tone third-order dynamic range. Third-order dynamic range near 100 dB is possible with standard level diode-ring mixers and a narrow CW bandwidth. High-level mixers permit better dynamic range numbers, if the LO system is quiet enough.

For wide-band systems, a broadband transformer impedance transformer can replace the tuned low-pass output on the RF preamp. This will permit coverage of multiple bands, but the low-pass function is still important and must be included somewhere in the receiver RF path. For the VHF bands, a lower noise figure is useful, and a two-stage grounded-gate RF preamp is a good choice. Two of the figure 2 circuits packaged separately with coax connectors is a high-performance construction option.

Downconverter

After the preamplifier is the downconverter block, shown in figure 3. The downconverter includes the RF in-phase splitter, the two mixers, IF port attenuators, the matched diplexer networks, and the matched pair of audio LNAs. All of the resistors in the downconverter board should be 1% metal film. The input splitter is somewhat different than earlier versions. Rather than attempting to match to 50 ohms, the splitter shown matches the mixer inputs to somewhat lower impedance—but achieves nearly perfect amplitude balance and very low loss over a very wide frequency range. The upper frequency limit is reached when the winding on T1 approaches a quarter wavelength. At the lower frequency limit, the amplitude balance is perfect, but the isolation is poor. If operation down to 50 kHz is desired, more turns on a type 71 core could be used. At 144 MHz and above, a few turns through a small ferrite bead will work. The mixers are type TUF-3, which offers better performance than the TUF-1 mixers from 150 kHz through 225 MHz, the usual operating range of R2 type systems. TUF packaged mixers are available for direct conversion applications at frequencies up to 2500 MHz.

After the mixers are a pair of matched attenuators. The 6 dB attenuators shown in the schematic should be used for most applications. If more gain is available before the mixers, more attenuation can be used. These attenuators serve three very useful purposes: they ensure textbook termination of the mixer IF ports; they attenuate mixer 1/f noise; and they provide a well-defined source impedance to drive the matched diplexer networks. Mixer IF termination has been widely discussed in the literature. Mixer 1/f noise degrades receiver noise figure. Different mixers, even matched TUF-3s with the same date code, have widely varying amounts of 1/f noise. Attenuation between the mixer and the audio preamp can't improve receiver noise figure, but it can reduce the effect of mixer 1/f noise. Advanced receiver artists are encouraged to study this. By balancing the preamp gain and post-mixer attenuation to set the receiver noise figure and dynamic range, receiver performance may be made relatively independent of mixer 1/f noise.

The third very important function of the post-mixer attenuators is to set the driving point impedance to the matched diplexer networks. In the R2, the driving point impedance to the diplexers is the mixer IF impedance, which is a strong function of LO drive level. If one mixer has more LO drive than the other (a common condition) the phase and amplitude response of one diplexer network will be slightly different than the other. These differences are typically enough that the ultimate opposite sideband suppression of R2 systems across an SSB bandwidth is about 43 dB—even with perfect audio phase-shift networks. By contrast, the miniR2 with off-the-shelf components often exhibits 50 dB of opposite sideband suppression.

An R2pro was built with all audio phase-shift network and diplexer components matched to within 0.1%. With no attenuation between the mixers and diplexers, the best sideband suppression that could be obtained across the audio band was 43 dB. With 6 dB attenuators, more than 50 dB was obtained across the audio band. With 20 dB attenuators, more than 60 dB of opposite sideband suppression was obtained—but adjustments were very critical and it is doubtful that suppression would hold over time and temperature. 50 dB of opposite sideband suppression in a receiver with no AGC means that brutally strong signals simply disappear on the opposite sideband. Most conventional receivers can't even make this measurement, because there is less than 50 dB headroom between the receiver audio output noise level and the output amplifier distortion level.

The diplexer networks are slightly simplified from the R2 networks. The R2 networks provided rapid rolloff both above and below the 300 to 4000 Hz audio band. The rolloff below the audio range does not contribute much to usable receiver dynamic range, but it does introduce rapid phase shifts in the critical 300 to 600 Hz frequency range. When R2 receivers are optimized for SSB operation, the suppression of the opposite sideband in the 300 to 600 Hz range is often right at the 40 dB spec. If the receiver is optimized for CW operation, sideband suppression usually falls off at higher audio frequencies. The miniR2 and R2pro eliminate the rapid rolloff at the low end of the audio range, which permits good performance through the CW range when the receiver is optimized for SSB.

Another change from the R2 and miniR2 circuits is the elimination of the electrolytic capacitors from the critical audio signal paths. The R2pro has only matched polypropylene capacitors in the audio path prior to the summing network.

The rolloff above the audio range is retained from the R2, with slight changes to make the receiver less sensitive to component tolerance. For good performance, it is necessary to match the diplexer components in R2pro to within 1%, just as in the original R2. If this is not done, opposite sideband suppression is likely to be poor across the audio band. By contrast, the diplexers in the miniR2 were designed to be used with standard tolerance components. The benefit of using the R2pro diplexers with matched components is that the close-in dynamic range is exceptionally good. R2pro two-tone measurements may be made at tone spacings of 10 kHz and 5 kHz.

The usual grounded-base audio preamp stages are used following the diplexer networks. There are other audio preamps that will work, but the grounded base stages have the advantage of having an input impedance that is set by the current through the transistors, which may be set up precisely using 1% resistors. The grounded base stages drive the non-inverting inputs of a low-noise dual op-amp, which provides low impedance drive to the following stages. Note that the outputs are not DC blocked. This is so that the low impedance drive from the dual op-amp can directly drive the audio phase-shift network. Because these outputs carry DC, there is the potential to short them and damage the dual op-amp. It should be mounted in a socket so it may be replaced if needed. The gain of the op-amp stages is set by the ratio of the series and feedback resistors. The audio preamp gain is an important variable that must be adjusted to optimize receiver noise-figure and dynamic range performance. The resistor values shown in the schematic are optimum for driving the analog phase-shift network. If the downconverter is used to drive a DSP card, more gain may be needed.

Since the signal is split at the input to the downconverter and the outputs are I and Q channels, it is necessary to match everything in the two channels of the downconverter block. In most cases, it is the I Q downconverter block, and not the audio phase-shift network, that sets the ultimate limitation on receiver opposite sideband suppression.

Because the audio signal levels are low and the gain is moderate, conventional RF grounding and shielding practices may be used for the downconverter block. With LO signals floating around on the same frequency as the desired RF, shielding is very important. The downconverter may be enclosed in a die-cast box, or even better, an enclosure soldered up from tin sheet or PC board scraps. A completely soldered up enclosure with good RF connectors, feedthrough capacitors for power, and shielded connectors for the I and Q outputs will result in an outstanding receiver. Be careful not to short the I and Q outputs to ground, or the cover will have to be unsoldered to replace the IC.

Analog Signal Processor

The third block in the R2pro system is the analog signal processor (ASP) board shown in figure 4. This board contains the audio phase-shift network, the summer, and a wideband passive audio filter. The audio gain is low, but the signal levels are also low, so this board should not be located where it can pick up power supply or computer noise. There are no RF signals present, so audio grounding rules apply. The single audio ground rail runs up the middle of the PC board between the ICs. The power supply line is decoupled by the 100 μ F capacitor and 100 Ω series resistor. Do not bypass the hot end of the 100 Ω resistor to ground.

There is only one change in the audio phase-shift network from the version used in the miniR2. 1.52 K is not a standard value in the 1% series. It is obtained by connecting a 1.50 K and 20 ohm resistor in series. With the audio phase-shift network resistors selected to within 0.1% of their marked value, more than 60 dB of opposite sideband suppression could be obtained—if the rest of the receiver were perfect. By selecting these resistors, the builder can be assured that the audio phase shift network is not limiting receiver performance. Some workers have suggested that a different network should be used if the receiver is only to be used for CW. This is not true. The image-reject mixer provides an attenuation band that covers the entire opposite sideband from 200 Hz to over 4000 Hz. This attenuation band is a key part of the receiver selectivity. Reducing the width of this rejection band does not improve receiver performance.

The output of the summer drives a wide bandwidth audio bandpass filter and an op-amp buffer. The filter provides a high-fidelity SSB bandwidth audio channel to the following circuitry. This is an ideal take-off point for DSP filters, FFT analyzers, home audio system stereo amplifiers, outboard audio filters, audio AGC systems, or directly into the R2pro audio output board. Two isolated audio outputs are provided from the ASP board for convenience. Output levels may be independently selected by changing the feedback resistors.

For construction hints on mounting and connecting to the ASP board, take the cover off a stereo receiver or amplifier and look at the circuitry around the magnetic phono cartridge inputs. Don't expect to find RF shielding, but a well defined single ground connection, shielded wire or twisted pair with the ground connected only at one end, and power connections directly to the big power supply capacitor are common. *This PC board should be mounted on nylon standoffs with a single wire to ground at the power supply.*

Accessory Filters

In most high-performance applications, switched bandwidths are used. Various DSP filters may be used, but the analog purist will probably want a selection of LC filters. The switched network shown in figure 5 provides SSB wide, SSB narrow, and a 500 Hz CW bandwidth, with gain equalizing resistors so the receiver output stays approximately constant.

Audio output board

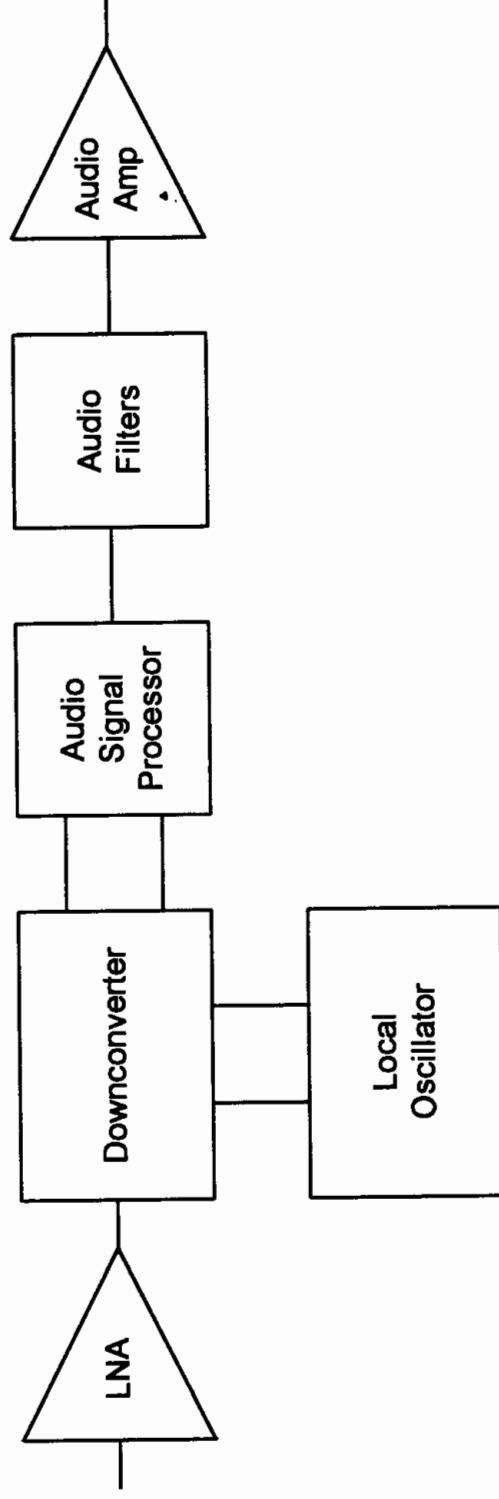
With the low level audio and RF circuitry on separate boards, the audio output section is much easier to handle. The circuit used for the R2pro experiments is shown in figure 6. The output devices have many milliamps of audio signals flowing through many milliohms of wire and ground trace resistance. This results in many microvolts of audio signals riding on the ground system. With only two gain stages and no RF grounding requirements, it is relatively easy to design an adequate ground system that will not oscillate. The single audio output board ground wire connects directly to the power supply negative lead, and also to the negative speaker lead. The +12 line connects directly to the big power supply capacitor. This audio power amplifier provides reasonable output with headphones or a small speaker in a quiet room. For more volume, an external power amplifier should be used. There is such a proliferation of external sound card amplification systems for computers that it would be difficult to justify building an audio power amplifier for an amateur receiver.

A significant change from the R2 and miniR2 designs is in the mute circuitry. The 60 dB in-band dynamic range of the R2 and miniR2 was limited by distortion in the JFET series switch. The mute switch in the R2pro circuit uses a relay across the feedback resistor of an op-amp in the ASP. When the relay is on, the op-amp goes into a low-gain state. The relay audio switch and the preamp mute circuit provide rapid, low-distortion receiver muting that may be used for graceful full-break-in keying without the usual annoying thumps at high volume levels. With the relay mute switch, the in-band dynamic range increases significantly. Careful measurements of the old and new circuits have not yet been made.

Local Oscillator

A local oscillator is not included in the R2pro receiver system, but the choice of LO in large part determines the success of the finished project. Two local oscillators that have been used to build excellent direct-conversion receivers are a well-shielded JFET Hartley and a moderately well-shielded JFET Hartley driving a balanced frequency doubler. When the diode doubler is used in a circuit with toroid inductors, open PC board construction is acceptable. The Kanga UVFO kit works well and provides additional useful features such as CW offset and a keyed auxiliary output. Because of differences in the way even and odd harmonics add, direct conversion receivers that use odd harmonic frequency multipliers must be very well shielded. Many other local oscillator systems are acceptable—but the simple elegance of a direct conversion receiver with a well-shielded on-frequency tunable local oscillator is appealing.

R2pro System
Figure 1



R2pro LNA

Figure 2

revised 5/02/01 RLC


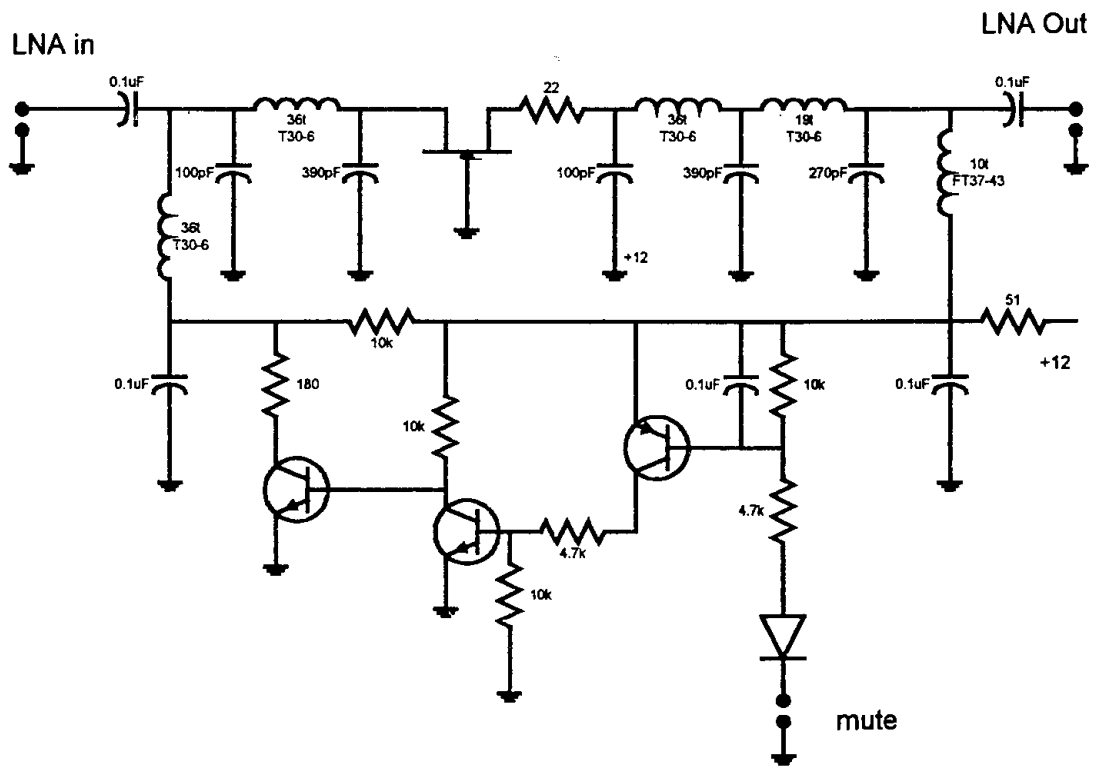



Figure 3
Downconverter

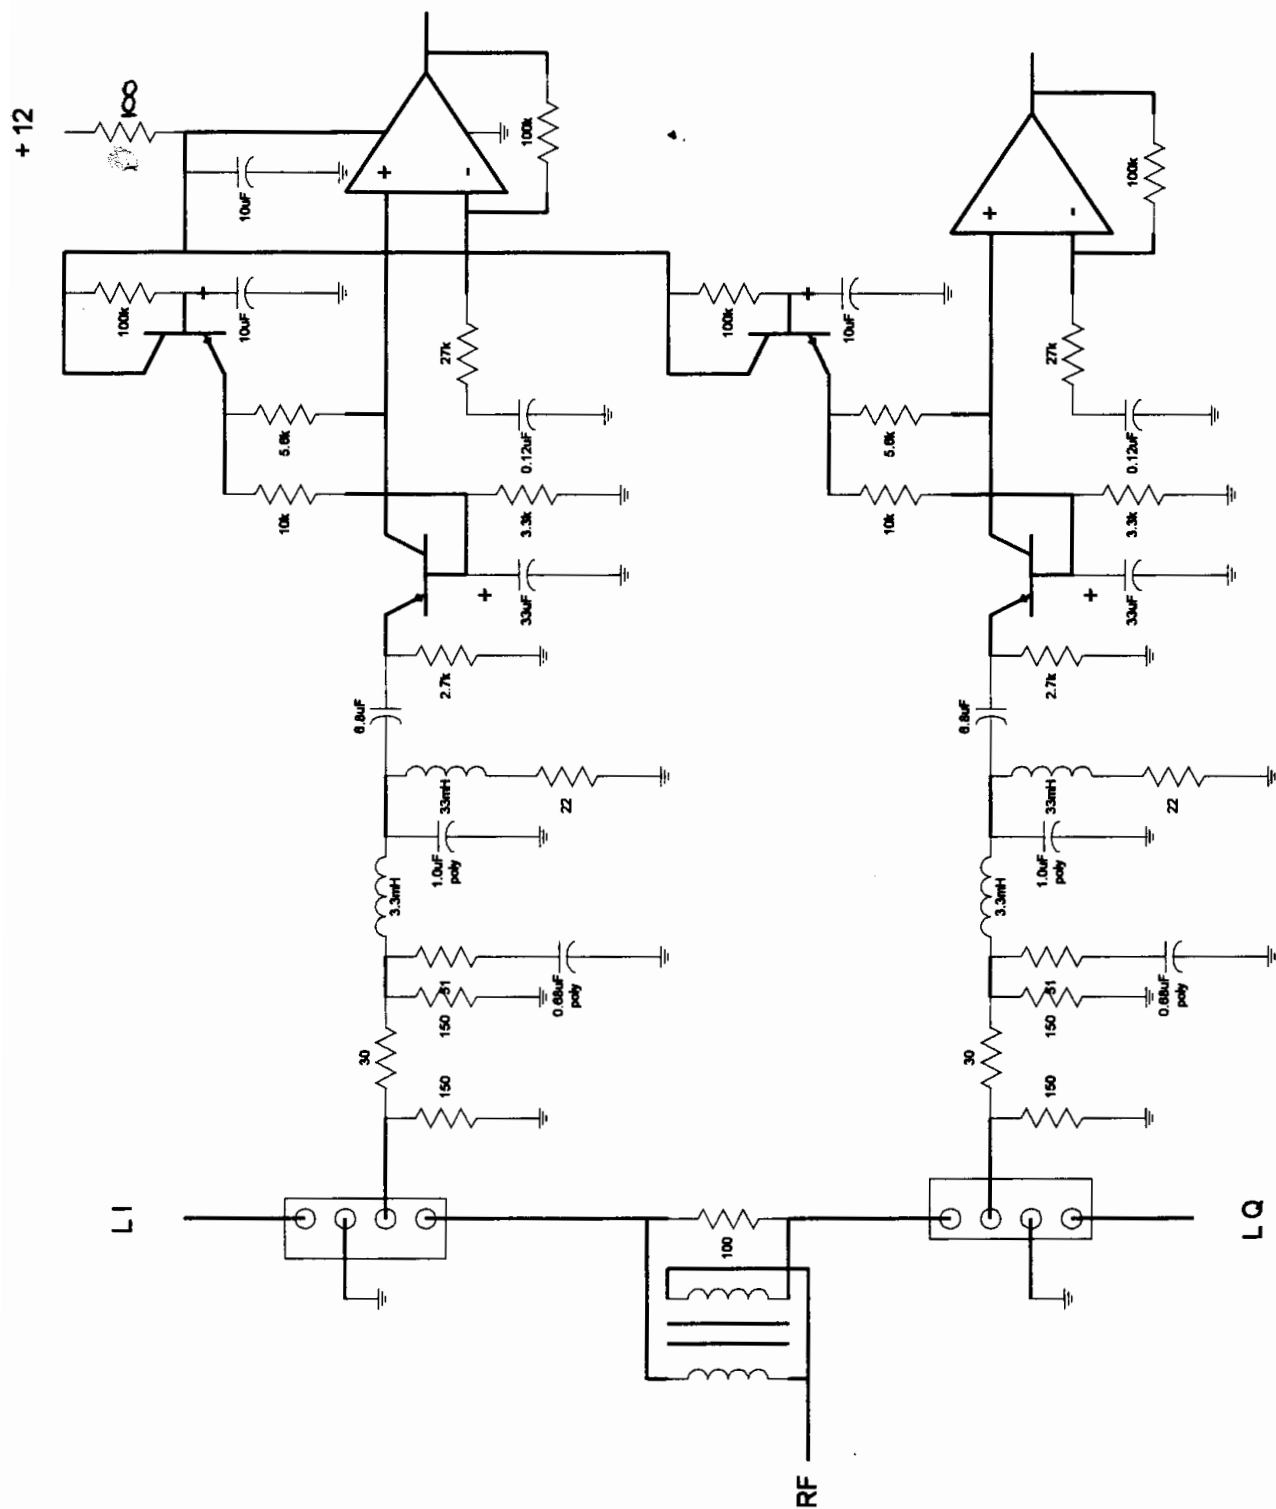
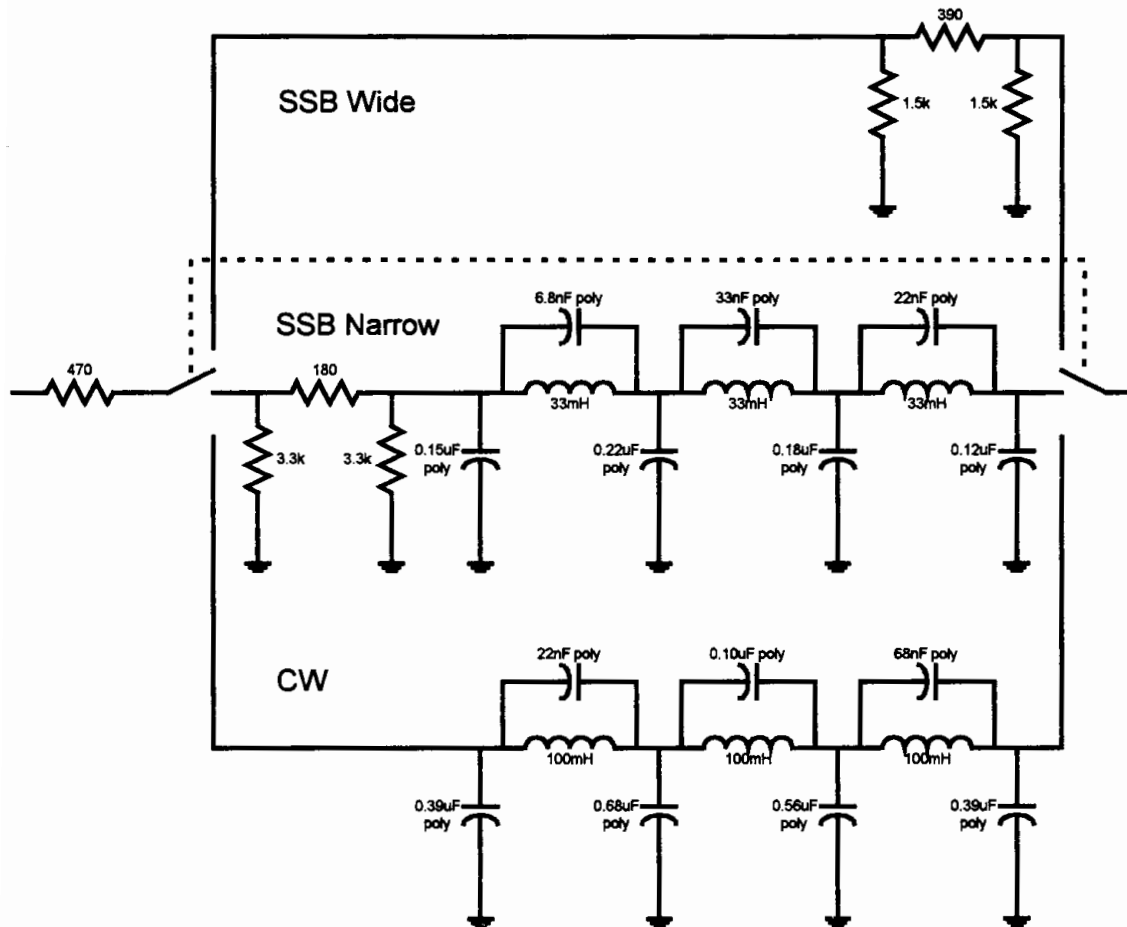


Figure 4

Accessory Filters
Figure 5



revised 8/31/00 RLC
revised 5/02/01 RLC

