Receiver Architectures

Modules: VCO (2), Quadrature Utilities (2), Utilities, Adder, Multiplier, Phase Shifter (2), Tuneable LPF (2), 100-kHz Channel Filters, Audio Oscillator, Noise Generator, Speech, Headphones

0 Pre-Laboratory Reading

The superheterodyne receiver is the classic receiver design. When implemented well, this design achieves excellent performance. However, there are circumstances under which other designs are preferable. The Weaver architecture is an alternative design that is gaining popularity.

0.1 Superheterodyne Receiver

When a sinusoid (a carrier) of frequency f_c is multiplied by a sinusoid (a local oscillator) of frequency f_{LO}, the result is the sum of two sinusoids (of equal amplitude), one having the difference frequency $|f_c - f_{LO}|$ and the other the sum frequency $f_c + f_{LO}$. In this discussion, frequencies are considered positive. The multiplication of the carrier with a local oscillator is called *low-side injection* when $f_{LO} < f_c$. It is called *high-side injection* when $f_{LO} > f_c$. A lowpass filter can be selected that will block the sum frequency. This combination of multiplier and lowpass filter is called a *downconverter*. The output frequency $f_{IF} = |f_c - f_{LO}|$ is called the *intermediate frequency* (IF).

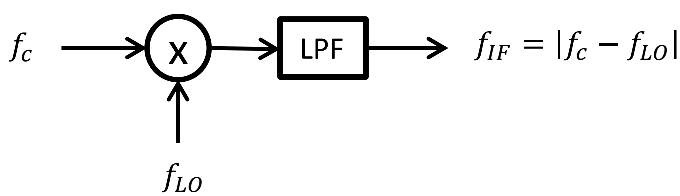

When the carrier is modulated, the downconverter shifts the spectrum of the modulated carrier so that it is centered at f_{IF} on the downconverter output.

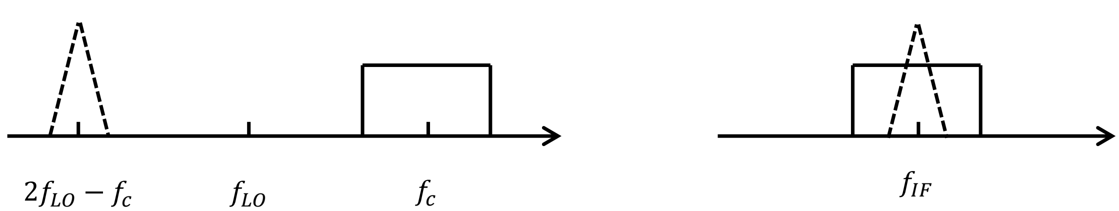

Low-side injection ($f_{LO} < f_c$): spectrum at downconverter input (left) and output (right)

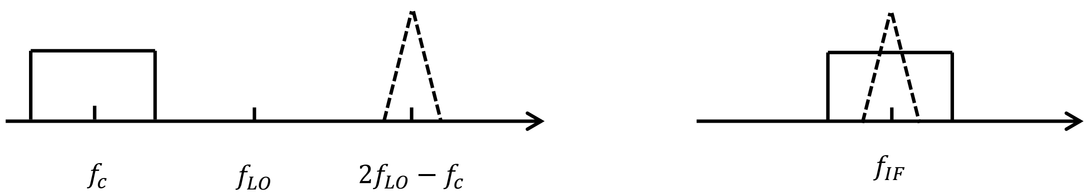

High-side injection ($f_{LO} > f_c$): spectrum at downconverter input (left) and output (right)

There is a serious problem with the simple downconverter discussed above. Any signal at the receiver input occupying the frequencies near $2f_{LO} - f_c$ will be shifted by the downconverter to the frequencies near f_{IF} at the downconverter output and will therefore overlay the desired signal. The span of frequencies near $2f_{LO} - f_c$ is known as the *image band*. For low-side injection, the image band lies below f_{LO} and the difference

$$|f_{LO} - (2f_{LO} - f_c)| = f_{IF}, \qquad \text{low-side injection} \tag{1}$$

For high-side injection, the image band lies above f_{LO} and the difference

$$|(2f_{LO} - f_c) - f_{LO}| = f_{IF}, \qquad \text{high-side injection} \tag{2}$$

Any signal in the image band at the receiver input is therefore a potential interferer.

This image-band problem is often solved by introducing a bandpass filter before the multiplier. This filter passes the desired carrier of frequency f_c and blocks the image band, which is centered at frequency $2f_{LO} - f_c$. Therefore, at the multiplier input there is no significant signal content in the image band.

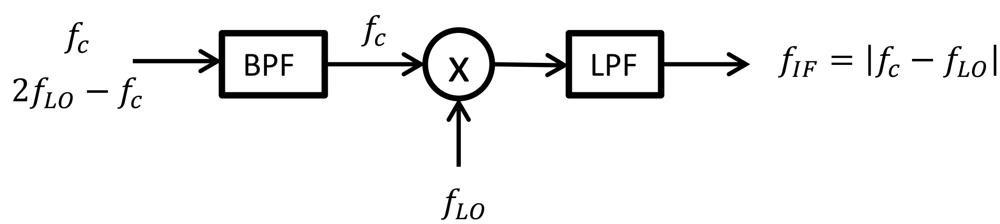

The modulated carrier at frequency f_{IF} can then be processed by a demodulator in order to extract the message signal. Demodulators are normally designed to operate at intermediate frequencies well below the carrier frequency. The downconverter places the modulated carrier within reach (in a frequency sense) of the demodulator.

Most radio receivers employ a bandpass filter (to block the image band), followed by multiplication with a local oscillator, followed by lowpass filtering, followed by a demodulator that is designed to operate at the intermediate frequency. Such a radio receiver is said to have a *superheterodyne* architecture. In some superheterodyne receivers, there are other components as well. For example, there are often amplifiers and automatic gain control circuits within the receiving chain.

0.2 Weaver Architecture

In recent years there has been considerable effort to reduce the size of radio receivers. This is largely due to the popularity of hand-held devices, such as cell and smart phones. Although superheterodyne receivers generally have excellent performance, the bandpass filter (blocking the image band) tends to be expensive and bulky. For this reason, engineers have tried other receiver architectures. The Weaver architecture, in particular, is attractive; it obviates the need for the expensive and bulky bandpass filter.

The Weaver architecture accomplishes the downconversion to an intermediate frequency without a bandpass filter and without being susceptible to interference lying in image bands. Following this downconverter would be a demodulator.

The easiest way to understand the Weaver circuit is to study an example with specific frequencies. Such an example is shown below. In a real radio, the carrier frequency would be much larger than 91 kHz. The frequencies shown below will be used in this experiment.

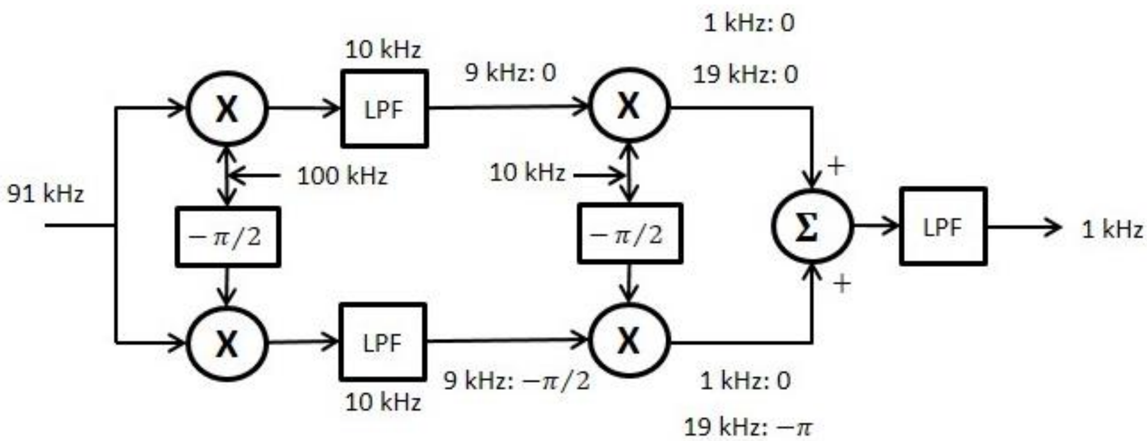

This architecture contains two channels: an upper and a lower channel. Each channel contains two stages of downconversion. The upper and lower channels are almost the same, but there is an important difference. Each local oscillator in the lower channel has a phase $-\pi/2$ radians relative to the corresponding local oscillator for the upper channel.

The first downconverter in the upper channel shifts the 91-kHz carrier frequency to 9 kHz. The same is true for the first downconverter in the lower channel. The lowpass filter in each channel has a bandwidth of 10 kHz and blocks the 191-kHz (sum-frequency) term. The second downconverter in each channel shifts the 9-kHz IF signal to 1 kHz. The lowpass filter that you would expect to follow the second multiplier in each channel lies after the adder. The one lowpass filter therefore serves to eliminate the sum-frequency term coming from each of two different multipliers. The output of this Weaver downconverter is 1 kHz.

In order to predict how the signals in the upper and lower channels will combine in the adder, it is essential to know their relative phases. A little trigonometry helps. If two sinusoids are placed on the input of a multiplier, the phase of the difference-frequency output depends on which input frequency is larger. If $f_1 > f_2$,

$$\mathcal{S}\{2\cos(2\pi f_1 t + \theta_1)\cos(2\pi f_2 t + \theta_2)\} = \cos[2\pi(f_1 - f_2)t + \theta_1 - \theta_2], \qquad f_1 > f_2 \tag{3}$$

$\mathcal{S}\{\cdot\}$ represents the lowpass filter (following the multiplier) that passes the difference-frequency term and blocks the sum-frequency term. If instead $f_2 > f_1$,

$$\mathcal{S}\{2\cos(2\pi f_1 t + \theta_1)\cos(2\pi f_2 t + \theta_2)\} = \cos[2\pi(f_2 - f_1)t + \theta_2 - \theta_1], \qquad f_2 > f_1 \tag{4}$$

This may be stated as a fundamental rule: *In a multiplier, the difference-frequency term on the output has a phase that equals the phase of the higher-frequency input minus the phase of the lower-frequency input.*

Consider the first pair of multipliers (each downconverting 91 kHz to 9 kHz). The higher-frequency input (the 100-kHz local oscillator) has $\pi/2$ radians less phase on the lower channel. The lower-frequency input (91 kHz) has the same phase on the two channels. Therefore, the difference-frequency output (9 kHz) has $\pi/2$ radians less phase on the lower channel.

We are only interested in the relative phases of the signals in the lower and upper channels. It is not necessary that the upper-channel 9-kHz signal have a phase of 0 radians. It is only necessary that the 9-kHz signal on the lower channel have a phase that is $-\pi/2$ radians relative to the corresponding signal on the upper channel.

Now consider the second pair of multipliers (each downconverting 9 kHz to 1 kHz). The higher-frequency input (the 10-kHz local oscillator) has $\pi/2$ radians less phase on the lower channel. The lower-frequency input (9 kHz) also has $\pi/2$ radians less phase on the lower channel. Since the difference between $-\pi/2$ and $-\pi/2$ is 0, the difference-frequency output (1 kHz) has the same phase on the lower and upper channels. In the adder, these two 1-kHz signals will reinforce each other. Hence, this Weaver circuit produces a strong output at 1 kHz in response to a 91-kHz input.

The second pair of multipliers also produces 19-kHz (sum-frequency) terms. A little reasoning similar to that done above reveals that the lower-channel 19-kHz signal has a phase $-\pi$ radians relative to the 19-kHz signal on the upper channel. If the gains on these two channels are the same, the amplitudes of these two 19-kHz signals will be the same, and they will perfectly cancel each other in the adder. If there is any 19-kHz signal left over due to imperfect cancelation, the final lowpass filter will take care of it. (In the experiment, this filter has a bandwidth of 3 kHz.)

What about the image frequency of 109 kHz at the input of the Weaver circuit? The diagram below shows what becomes of a 109-kHz input.

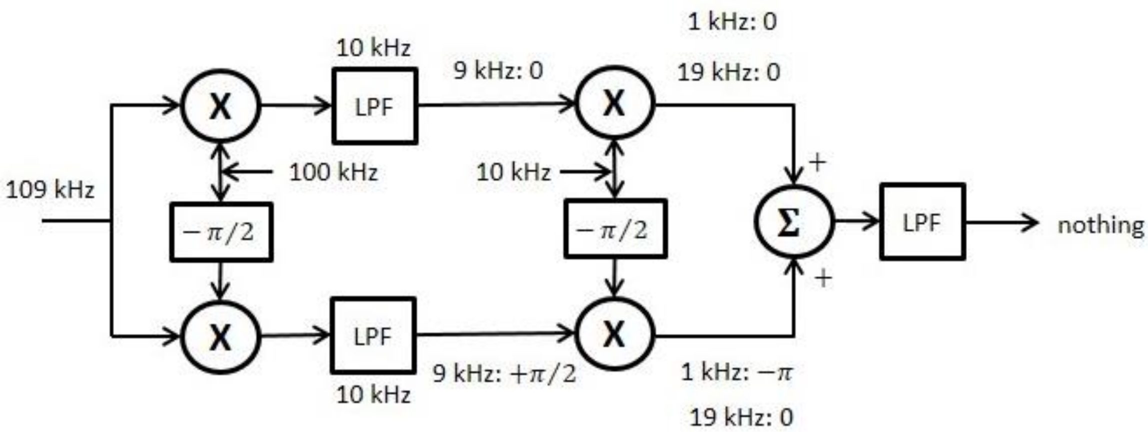

Consider the first pair of multipliers. The higher-frequency input (109 kHz) has the same phase for both multipliers. The lower-frequency input (the 100-kHz local oscillator) has $\pi/2$ radians less phase on the lower channel. Therefore, the difference-frequency output (9 kHz) has $\pi/2$ radians *more* phase on the lower channel.

Now consider the second pair of multipliers. The higher-frequency input (the 10-kHz local oscillator) has $\pi/2$ radians less phase on the lower channel. The lower-frequency input (9 kHz) has $\pi/2$ radians *more* phase on the lower channel. Since $-\pi/2$ minus $\pi/2$ is $-\pi$, the difference-frequency output (1 kHz) has π radians less phase on the lower channel. If the gains are equal in the upper and lower channels, these two 1-kHz signals will cancel each other in the adder. Hence, this Weaver circuit produces no output at 1 kHz in response to a 109-kHz input. This is accomplished without an expensive and bulky bandpass filter at the front of the receiver.

The 109-kHz input to the receiver also gives rise to 19-kHz (sum-frequency) outputs from the second pair of multipliers. With an analysis similar to that done above, you can see that the two 19-kHz signals are in phase. They will reinforce each other in the adder, but the final lowpass filter will block this 19-kHz signal.

We have shown that the image frequency of the first pair of multipliers is nulled in the Weaver circuit. What about the image frequency of the second pair of multipliers? This image frequency is 11 kHz (because 11 kHz minus 10 kHz equals 1 kHz). At the front of the receiver a frequency of either 89 kHz or 111 kHz could give rise to a difference-frequency term of 11 kHz. However, the lowpass filter in each channel will block this 11-kHz signal before it gets to the second pair of multipliers.

It might seem that the Weaver architecture is a complicated and expensive alternative to the superheterodyne architecture. The idea of the Weaver circuit has been around for a long time, and it was once considered expensive and complicated. However, modern electronics technology has rendered many complicated designs inexpensive. There are a lot of components in the Weaver circuit, but none of them are particularly expensive when implemented with modern technology. There are three filters. However, all three filters are lowpass, do not have wide bandwidths, and do not have to be tuned. These filters can be small and inexpensive. The superheterodyne receiver's bandpass filter is expensive and bulky because it is a high-frequency device that must have a sharp roll-off and, if multiple carrier frequencies are to be received, must be tuned. The Weaver circuit also requires multipliers, an adder and two phase shifters. The latter are narrowband devices because each only has to change the phase the correct amount at only one frequency (that of the local oscillator); they are inexpensive. (A wideband phase shifter, such as a Hilbert transform, is another story.) All of the components in the Weaver circuit can be small.

1 Superheterodyne Receiver

Build a modulator for AM. Use an audio signal from the Speech module as the message signal. Use the Quadrature Utilities module for both the adder and the multiplier. The DC component will come from the Variable DC panel, and the knob on this panel should be set clockwise from the vertical position, so that a positive DC value is supplied. Connect the output of the adder to a Buffer Amplifier and the output of this amplifier to a multiplier. The other input to the multiplier will be a 100-kHz sinusoid (Master Signals).

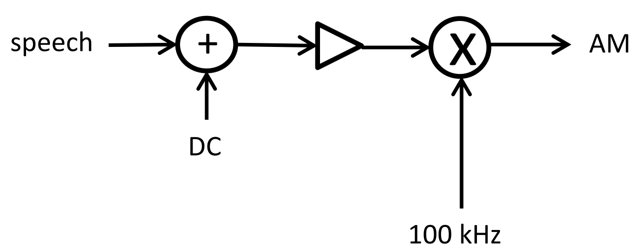

The adder in the Quadrature Utilities module is a weighted adder. There are two knobs on the PCB for controlling the weighting factors (gains). The exact settings of these gains are unimportant for this experiment. You should set these two gains for mid-range values. Both of these gains are negative, but the Buffer Amplifier following the adder will cancel the negative signs.

Place the Buffer Amplifier output on Channel A and the AM modulator output on Channel B. Switch to an XY view. Adjust the DC voltage in order to achieve something close to 100% modulation.

Channel A: Buffer Amplifier output
Channel B: modulator output

Observe the modulator output as a function of time on the oscilloscope. You will not achieve a stable display, but the signal should clearly look like an AM carrier.

 Channel A: modulator output

Connect the output of the modulator to one input of the Adder module. Connect the output of a VCO (with mode switch set to “VCO”) to the other input of the Adder. Set the VCO frequency to approximately 71 kHz; this sinusoid will represent an image-band interferer in this experiment. The output of the Adder should go to a Buffer Amplifier, whose negative gain will cancel the negative gains in the Adder.

You will use a Tuneable LPF in the downconverter. Before placing it in the downconverter, adjust its bandwidth to approximately 20 kHz. (The Tuneable LPF’s clock output has a frequency equal to 100 times the bandwidth.) Use the Noise Generator module to get a quick display of $|H(f)|$.

 Channel A: Tuneable LPF output

Build a receiver consisting of a downconverter followed by an envelope detector. The downconverter will consist of a multiplier plus a low-pass filter. In the first instance, do not use a bandpass filter before the multiplier. The local oscillator for the downconverter will come from a VCO (a different VCO than the one that serves as the interferer), with mode switch set to “VCO” and tuned to a frequency of approximately 85 kHz. This is an example of low-side injection. For the filter in the downconverter use the Tuneable LPF whose bandwidth has been set to 20 kHz. The IF signal produced by the downconverter will have a center frequency of approximately 15 kHz. This IF signal should be placed at the input of a rectifier (on the Utilities module). The output of the rectifier should go to the low-pass filter in the Headphone Amplifier module. This filter has a bandwidth of approximately 3 kHz.

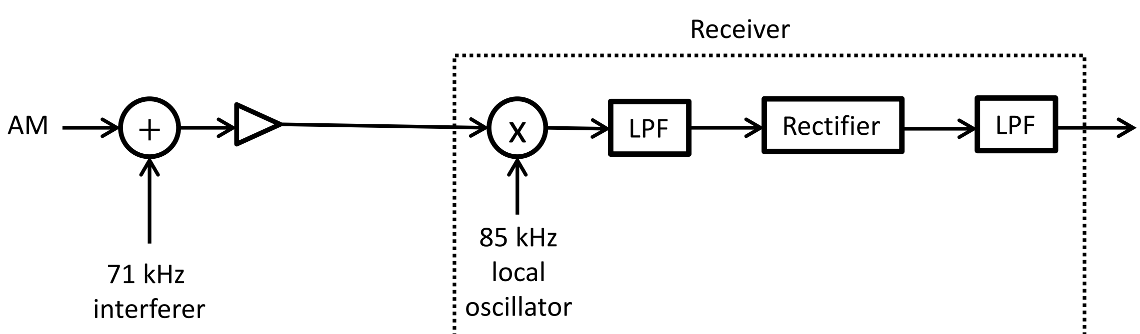

When the AM carrier is 100 kHz, there is no genuine need to downconvert the carrier before passing the signal to an envelope detector. In a real application, however, the carrier frequency would typically be higher than 100 kHz and it is better to downconvert the carrier to an IF before envelope detection. Downconversion is demonstrated here for a 100-kHz carrier for the simple reason that the TIMS instrument does not permit the generation of carriers that are much higher in frequency than 100 kHz.

Initially set the gain within the (weighted) Adder for the 71-kHz interferer to 0. (The gain knob should be fully counter-clockwise.) With no interference present in the receiver, make the following observations. First, observe the spectrum at the receiver input. Second, observe the spectrum at the envelope detector input. Third, observe the spectrum at the receiver output (that is, the output of the 3-kHz low-pass filter). Listen to the detected audio signal.

Channel A: receiver input (no interference)

Channel A: envelope detector input (no interference)

Channel A: receiver output (no interference)

Connect the receiver input to Channel A. Adjust the gain within the (weighted) Adder for the 71-kHz interferer while observing the spectrum of the receiver input. Set this gain so that the power in the interferer has a line height approximately equal to that of the residual carrier. With a significant interference now present in the receiver, make the following observations. First, observe the spectrum at the receiver input, noting the line height of the interfering tone. Second, observe the spectrum at the envelope detector input, noting the line height of the interfering tone. Third, observe the spectrum at the receiver output.

Channel A: receiver input (interference present)

Channel A: envelope detector input (interference present)

Channel A: receiver output (interference present)

You should note that the 71-kHz interferer finds its way into the envelope detector because it lies in the image band for this downconversion. The fact that the interference line height matches (approximately) that of the residual carrier confirms that you are observing the interferer as it passes through the receiver. Listen to the detected audio signal. You should note that the interference has a deleterious effect on the audio output from the receiver.

Modify the above experimental arrangement by introducing a bandpass filter to the receiver, preceding the multiplier. For this filter, use the 100-kHz Channel Filters module with the position switch set to 3 (BPF). The purpose of this filter is to pass the 100-kHz carrier but block the image band centered at 70 kHz. Use the Noise Generator module to get a quick display of $|H(f)|$ for this BPF.

Channel A: BPF output

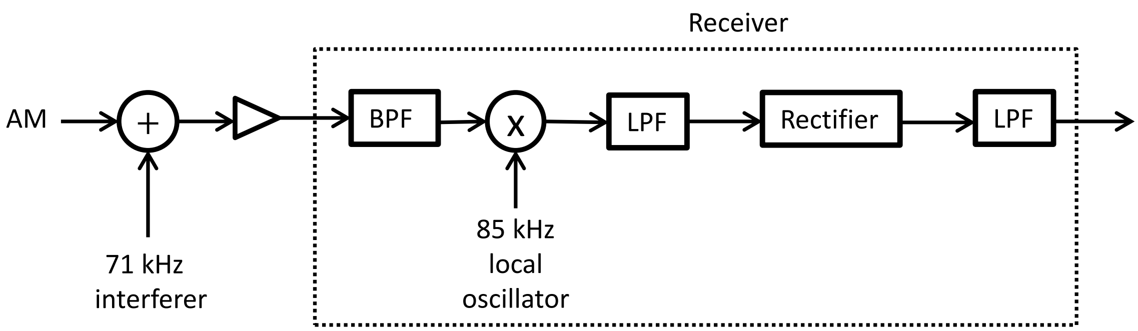

Connect the Buffer Amplifier output (the BPF input) to Channel A. Adjust the gain within the (weighted) Adder for the 71-kHz interferer while observing the spectrum of the BPF input. Set this gain so that the power in the interferer has a line height approximately equal to that of the residual carrier, as seen at the BPF input.

With the image-band-blocking BPF in place, make the following observations. First, observe the spectrum at the receiver input. Second, observe the spectrum at the envelope detector input. Third, observe the spectrum at the receiver output.

Channel A: receiver input (image-band-blocking BPF present)

Channel A: envelope detector input (image-band-blocking BPF present)

Channel A: receiver output (image-band-blocking BPF present)

Listen to the envelope detector output to make sure the interference no longer degrades the demodulated audio.

2 Weaver Architecture

You can take apart the AM modulator, superheterodyne receiver, and envelope detector. You won't need them for the remainder of this experiment.

For the Weaver architecture, you will need two Phase Shifters. Set one of these (with frequency range set for "HI" on the PCB) to shift the phase of a 100-kHz (one-hundred kilohertz) sinusoid by $-\pi/2$ radians. Use the oscilloscope for this purpose, making sure the sign of the phase shift is correct. (The Phase Shifter output should lag the input.) Set the other Phase Shifter to shift the phase of a 10-kHz (ten kilohertz) sinusoid by $-\pi/2$ radians. Make sure the Phase Shifter output lags the input.

Build a downconverter that will convert the range of frequencies 90 kHz $< f \leq$ 93 kHz to the range 0 kHz $< f \leq$ 3 kHz using a Weaver architecture. Use a Quadrature Utilities module for the first pair of multipliers. For the first local oscillator, use a 100-kHz sinusoid (Master Signals). For the two IF low-pass filters, use Tuneable LPFs (toggle switch set to "wide") with the bandwidth of each adjusted to 10 kHz. (The Tuneable LPF's clock output has a frequency equal to 100 times the bandwidth.) Use the other Quadrature Utilities module for the second pair of multipliers. For the second local oscillator, use an Audio Oscillator adjusted for a frequency of 10 kHz. For the adder, use a (weighted) Adder. For the final lowpass filter, use the 3-kHz LPF on the Headphone Amplifier module.

The fact that this weighted Adder has negative gain on both channels is not a problem. If you want, you can follow the Adder with a Buffer Amplifier to cancel the negative signs. But this is not really necessary since only the relative phases of the upper and lower channels are important in this experiment.

Use the VCO (PCB slide switch set to "VCO" and the toggle switch set to "HI") to generate a sinusoid that will be downconverted. In the first instance, adjust the VCO frequency to 109 kHz. Ideally, with this input frequency, the output of the Weaver-architecture downconverter will be zero. However, this will only happen if the gain in the upper channel exactly equals the gain in the lower channel. Adjust the gains of the (weighted) Adder to minimize the output of the downconverter (the output of the 3-kHz lowpass filter). If you have already properly adjusted the delays of the two Phase Shifters, it should be unnecessary to make any further phase adjustments. However, you might want to experiment with slight changes in the phase delay of the 100-kHz Phase Shifter in order to see if you can further minimize the downconverter output. At most, only a slight phase change will be required, so use the fine control on the Phase Shifter.

Connect the output of the upper-channel LPF to Channel A and the output of the lower-channel LPF to Channel B. With the VCO frequency still set for 109 kHz, you should observe that the 9-kHz sinusoid on the lower channel *leads* the 9-kHz sinusoid on the upper channel by 90°.

Channel A: upper-channel LPF output (109-kHz downconverter input)
Channel B: lower-channel LPF output (109-kHz downconverter input)

The downconverter input (the VCO signal) should still be set for 109 kHz. Place copies of the two Adder inputs on the oscilloscope, with the upper-channel (Adder input) on Channel A and the lower-channel on Channel B.

Channel A: upper-channel Adder input (109-kHz downconverter input)
Channel B: lower-channel Adder input (109-kHz downconverter input)

Each of these signals should look like a 19-kHz variation superimposed on a lower-frequency (approximately 1 kHz) variation. For a downconverter input of 109 kHz, the upper and lower 1-kHz variations should be 180° out of phase. The higher-frequency (approximately 19-kHz) variations, on the other hand, should be in phase. The lowpass filter at the output of the Adder blocks the 19-kHz sinusoid that would otherwise appear at the downconverter output.

Adjust the VCO frequency to 91 kHz. There should now be a relatively large-amplitude sinusoid of frequency 1 kHz on the downconverter output.

Connect the output of the upper-channel LPF to Channel A and the output of the lower-channel LPF to Channel B. With the VCO frequency still set for 91 kHz, you should observe that the 9-kHz sinusoid on the lower channel *lags* the 9-kHz sinusoid on the upper channel by 90°.

Channel A: upper-channel LPF output (91-kHz downconverter input)
Channel B: lower-channel LPF output (91-kHz downconverter input)

The downconverter input (the VCO signal) should still be set for 91 kHz. Place copies of the two Adder inputs on the oscilloscope, with the upper-channel (Adder input) on Channel A and the lower-channel on Channel B.

Channel A: upper-channel Adder input (91-kHz downconverter input)
Channel B: lower-channel Adder input (91-kHz downconverter input)

Each of these signals should look like a 19-kHz variation superimposed on a lower-frequency (approximately 1 kHz) variation. For a downconverter input of 91 kHz, the upper and lower 1-kHz variations should be in phase. The higher-frequency (approximately 19-kHz) variations, on the other hand, should be 180° out of phase.

For each of six different VCO frequencies (109, 108, 107, 93, 92, and 91 kHz), you will measure the amplitude at the output of the Weaver-architecture downconverter, placing the data in an Excel worksheet.

Weaver architecture