Simple embedded receiver for low-cost applications

By
Keith Curtis
Technical Staff Engineer
Microchip Technology Inc.

Introduction

Not every radio link needs the sophistication of ZigBee® and a low-cost, low-data-rate link can be created using just a few components and a little RF knowledge. The combination of a simple design. with the flexibility of the microcontroller to implement different algorithms and use on-board peripherals to minimise external components, can produce a receiver suitable for compact remote controls, door openers or radio-controlled toys.

White Paper

History

To start, let's take a short trip through history. In the early 1900s, radio was still in its infancy. Marconi had demonstrated a viable telegraphy system in 1896, and then improved it with a frequency-filtering stage in 1901. However, even with this improvement, the receiver was still entirely passive, relying on the transmitter for all of its power. Next, Lee de Forest added the grid to existing diode vacuum tubes, creating the first amplifier and significantly increasing the sensitivity and range of receiver designs.

These early Tuned Radio Frequency (TRF) receivers were basically a string of tuned amplifiers. While this system was a definite improvement over the existing systems, it was expensive due to the number of tubes required, and it was difficult to tune due to the multiple individually tuned amplifiers.

In 1914, an enterprising young amateur radio operator named Edwin Armstrong patented a revolutionary design that addressed these limitations in the TRF receiver system. His idea was to eliminate most of the tuned amplifier stages with a single amplifier—the regenerative detector. He did this by using positive feedback from the output of the amplifier back to the input. This had the result of increasing the amplifier's gain from its nominal value of 15-20 dB, up to over 80 dB. The only drawback was the requirement that the feedback required periodic adjustment to prevent oscillation.

In 1922, Armstrong followed his patent on the regenerative detector with a variation on the design he called the Super Regenerative Detector. In this new design, the amplifier was again fitted with positive feedback. However, this time, the feedback was sufficient to drive the amplifier into full oscillation. Once oscillating, the circuit automatically "quenched" the oscillation, and then allowed it to build up once more. Because the time required to build to full oscillation is determined by both the gain of the amplifier and the amount of energy present at that frequency, it is possible to detect radio signals by measuring the oscillator start-up time.

This essentially makes the Super Regenerative Detector (Super Regen) a sampled oscillator that periodically measures the RF energy at the antenna at its frequency of oscillation, by using that energy as a jump start for the oscillator. This shortens its startup time. While this may sound overly complex, you must remember that it did significantly reduce the number of components required for a receiver, and it simplified tuning due to its single-tuned circuit.

The question is, how will a 90-year-old technology help build a radio link for a microcontroller? Well, we have many of the same constraints—a low-power requirement, low component cost, and the need for simple assembly. The Super Regen saved component costs in the 1920s. Its fewer tube filaments saved battery life and, with only one tune circuit, the Super Regen was simple to tune. Given these attributes, it makes sense that a Super Regen design should also meet those same requirements now. In fact, a few of the more common peripherals on today's microcontrollers should make the task that much easier.

Overview

Our design will combine the simplicity of the Super Regen Detector with the flexibility of the microcontroller to produce a simple receiver system requiring only two transistors. Figure 1 shows a general block diagram of the system. The Super Regen Detector periodically quenches its oscillation and rebuilds back to full oscillation. This results in a periodic fluctuation in its emitter/collector current. This fluctuation is buffered by the second transistor and tied to one input of the microcontroller's voltage comparator. The other comparator input is fed with a low-pass-filtered version of the buffered signal. This produces a periodic pulse tied to the detector's quench rate. By measuring the period of the resulting pulse chain on the output of the comparator, we can determine the amount of energy at the antenna, our frequency of interest, and thus demodulate the signal to retrieve the amplitude modulation of the RF signal we are trying to receive.
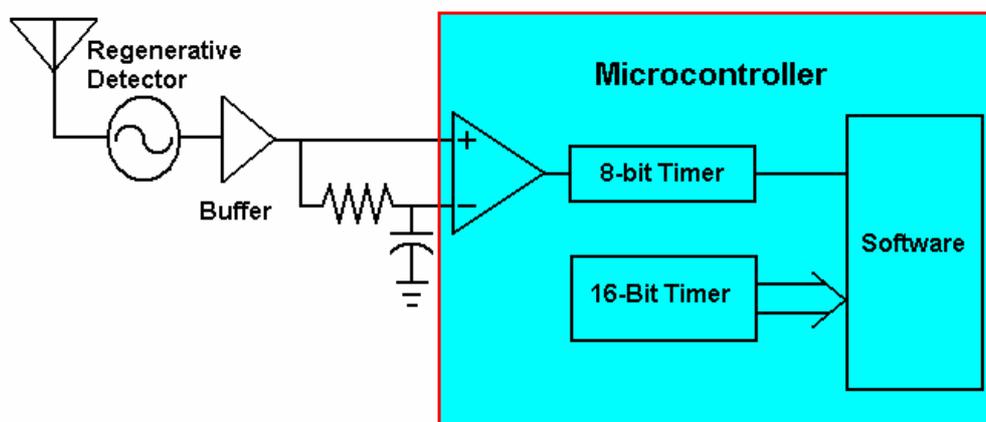

Figure 1: System Block Diagram

Detector

To begin with, let us examine the Super Regen Detector in greater detail. The basis of the circuit is a common base amplifier with a tuned circuit on the collector and capacitive feedback to the emitter (See Figure 2). The inductor on the emitter provides a DC path for the bias while providing isolating impedance for the emitter.

The antenna is coupled into either the emitter or the collector with a small series capacitance, to tune the antenna's inductor into a straight resistance so it won't pull the tuned circuit off frequency.
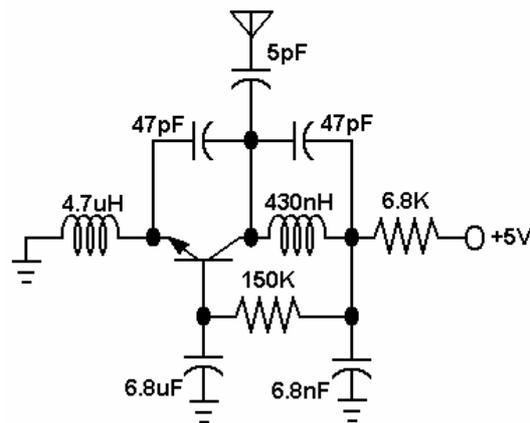
Figure 2: Schematic of a 27 MHz Super Regenerative Detector

Quenching

So far, we have seen all the elements of the oscillator, but what about this periodic quenching of the oscillation—how is that accomplished? Actually, the quenching of the oscillator is tied to the oscillation, itself. As the oscillation builds, the amount of current that the amplifier draws from the supply increases. This continues until the voltage at the collector drops so far that the oscillator literally starves itself of power. Once this happens, the oscillation quenches and the bias resistors recharge the bypass capacitor on the collector, until oscillation is once again possible. Note that some designs enhance this behavior by placing a parallel resistance and capacitance in series with the inductor on the emitter. This further squeezes the bias for the transistor by raising the emitter voltage until there is insufficient bias for oscillation to continue (See Figure 3).
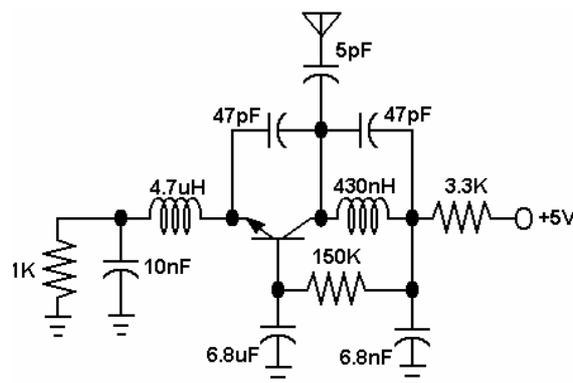
Figure 3: Alternate Schematic of a 27 MHz Super Regenerative Detectors

Buffer/Comparator

Once we have the detector oscillating and quenching, we now have to convert the bias current into a signal that we can measure for demodulation. This can be accomplished by either sampling the supply voltage at the top of the tuned circuit, or by measuring the voltage at the top of the parallel RC below the emitter inductor (See Figure 4). Regardless of the bias-current signal's source, we need to buffer the signal before we can filter it and drive the comparator.

While the comparator has large input impedance, the low-pass RC filter that drives the other comparator input has a somewhat lower impedance. We don't want to place any circuitry on the bias that might affect its quench frequency.
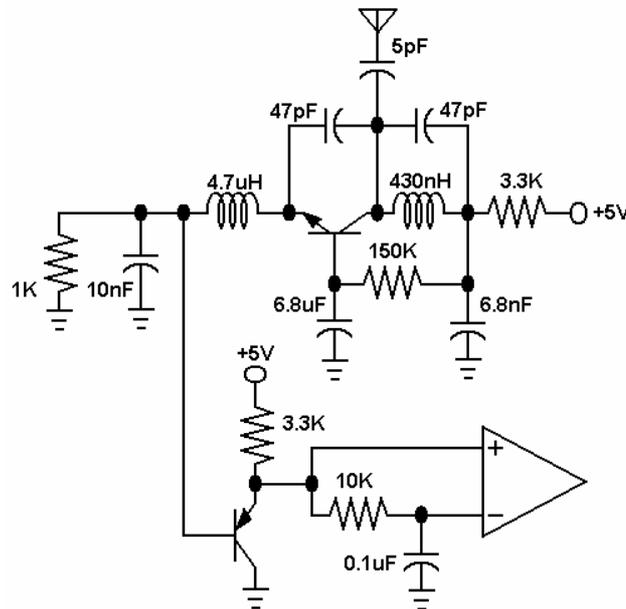

Figure 4: Buffer and Comparator

Once the bias and filtered bias are connected to the comparator, we have a means of measuring the quenching cycle's period. It is now simply a matter of accurately measuring the period and looking for deviations that indicate a signal of interest. The challenge comes from the frequency of the quench cycle. If the quench cycle is too slow, then the system's bandwidth will reduce. If the quench cycle is too fast, then we won't be able to measure the period with sufficient accuracy. While we could adjust the period back and forth until we found a reasonable compromise, there is another option due to the flexibility of the peripherals in the microcontroller.

The circuit shown in Figure 3 has a quench period of approximately 15 microseconds (µS), determined primarily by the bypass capacitor at the top of the tuned circuit and the RC time constant of the parallel RC of the network below the emitter inductor. This quench period is too fast to be measured with the requisite accuracy, if we try to measure an individual cycle of the waveform. However, if we measure the time period for multiple waveforms, we generate an averaging effect that not only provides some low-pass filtering, but also increases the measurement's resolution.

Counters

Assuming a clock frequency of 1 MHz, we would only receive 15 counts for the 15 µS quench period, giving us a resolution of only 6%. However, if we measure the period of 32 such waveforms, we derive a total of 480 counts. This gives us the ability to measure the period with an accuracy of better than 0.2%. Now, this does reduce our sample period to 480 µS, but if our incoming signal uses a minimum pulse width of 5-10 milli Seconds (mS), that still gives us an oversampling rate of 10 to 20.

To implement the system, all we need are two counters—one 8-bit counter and one 16-bit counter. The 9-bit counter is driven by the comparator's output, while the 16-bit counter is driven by the clock rate. To start a conversion, we preset the 8-bit counter to 255 minus 32, and clear the 16-bit counter. When the 8-bit counter rolls over, we halt the 16-bit counter and retrieve the quench time, multiplied by 32. The value is then stored for further demodulation, and the counters are reloaded for the next conversion.

Demodulating

Once the value has been retrieved from the counter, two things need to happen for the value to have meaning as a received signal. First, the signal must be averaged with previous values to determine the value corresponding to the system noise floor. Second, the new value must be compared to the ongoing average, to determine if the new value represents noise or an actual signal (data slicer).

Let's start with a quick note on averages and averaging systems. If an averaging system operates too quickly, it can "pump" or track incoming noise and signals, providing an inaccurate value for the system's noise floor. If an averaging system operates too slowly, then it can miss changes in the noise floor and report an elevation in the noise as a signal. To counter these problems, we need to do two things. First, we need to slow the average such that even a full pulse from the transmitter will not significantly shift the average. Second, we need to include a sanity check in the final decoding logic for the receiver such that extended pulses, or pulses that do not fit the transmitter's format, are discarded.

Averaging

The averaging system can be accomplished using one of two different methods. In the more traditional method, the last X number of samples are retained in a software shift register, which automatically shifts out the oldest value when a new value is received. This system then re-averages the entire contents of the shift register after each new value is shifted in. An alternative method simplifies the process by only retaining an ongoing average value and adding in only a portion of the new value, like an RC low-pass filter. While the system holding the last X values is traditional, it wastes processing power and data memory. This paper will employ the alternative system, instead.

To implement the system, it is necessary for the ongoing average variable to include at least 1 decimal value to the right of the decimal point. Floating-point numbers can easily accomplish this, but they tend to be large and slow. Instead, we will simply use the convention that the least significant byte of the average variable is to the right of the decimal point. Basically, assume that the average variable is in increments of $1/16^{th}$. When a new value is received, it is multiplied by 16. The average is subtracted from the multiplied value; the result is divided by 16 and then added back to the average variable (See Code Snippet 1.0).

$$New = New * 16;$$
$$Diff = New - Average;$$
$$Average = Average + (Diff/16);$$

Code Snippet 1.0: Averaging system

This system has several advantages. The multiplication and division can be accomplished by simple shifts, and the only storage requirement is the Average variable. This system also behaves better over time.

Using the traditional method, the response to a sudden jump in value is just a straight line as shown below, in Figure 5. However, the new system actually behaves like a low-pass RC filter, as shown in purple.
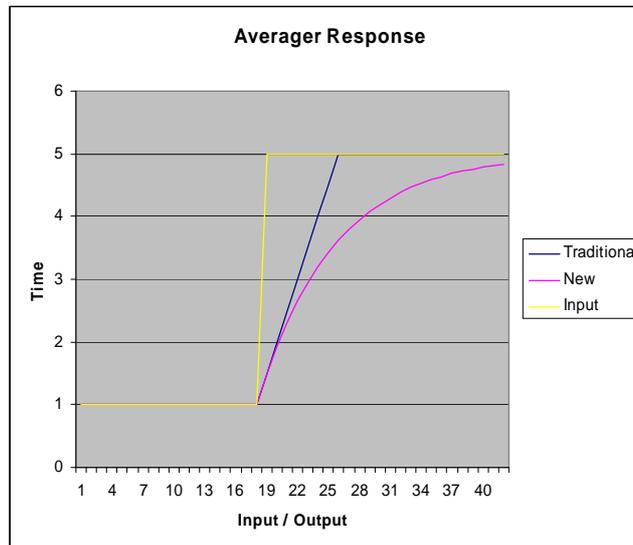

Figure 5: Average Function Responses

This slows the averaging routine's response to noise and incoming signals, keeping it from being pumped above the noise floor.

Data Slicer

The final demodulator function required is a software data slicer. This function is actually quite simple. A threshold value is subtracted from the average and the new valued received from the timer is compared to it. If the new value is below the threshold, then a signal at the frequency of interest is sufficiently strong to be decoded as a one. If the new value is greater, then it is assumed to be noise and the decoded output is left as zero. This string of ones and zeros is then fed into a decoder routine, which measures and verifies that the received ones are of a valid length, with a proper spacing, and decodes any message present in the resulting data.

Limitations

So far in the design, all that has been discussed are the advantages associated with Super Regen Detector designs. For completeness, the limitations of the systems should also be reviewed. The limitation is that the Super Regen system has only a limited ability to filter out adjacent channel interference. If another system is transmitting on a relatively close frequency, it can interfere with the receiver, thereby degrading its sensitivity to the desired frequency or even completely blocking the reception of the desired frequency.

The second limitation is that the Super Regen system does put out an interfering signal on the frequency it is trying to receive. This occurs as the oscillation builds—some of the energy passes out through the antenna and could interfere with other receivers. In fact, part of the reason for keeping the transistor's bias so small is to limit the amount of energy that leaks out. The following two modifications can be made to reduce this noise:

1. Use an RF amplifier between the antenna and the detector. The reverse isolation of the amplifier will reduce the amount of oscillator energy that can leak back through the amplifier to the antenna.

2. Use an autodyne circuit between the antenna and the detector. The autodyne is a single-transistor amplifier, mixer and oscillator that translates the desired frequency at the antenna down to an intermediate frequency (IF) at which the detector can operate. This reduces the amount of oscillator energy present at the desired frequency at the antenna. It does not completely eliminate the energy, as the autodyne will translate some of the IF frequency back up to the antenna, but it is much lower than the RF amplifier's reverse isolation.

   This modification has two additional advantages. First, an intermediate frequency filter can be used in front of the detector to reduce adjacent channel interference. Second, the system can now be tuned by adjusting the oscillator frequency, making the system much easier to tune.

Conclusions

Even with all its limitations, the Super Regen detector is a useful method for receiver designs and still enjoys widespread use in systems such as garage-door openers, ultra-small remote controls, and even low-cost radio-controlled electric toys. By utilizing peripherals within the microcontroller, the typical Super Regen circuit can be further simplified to produce a truly low-cost design. Because the microcontroller also handles decoding as well as the demodulation in software, it opens up the possibility of using different algorithms based upon system needs and the power level received through the detector. Figure 6 shows a picture of a prototype 27 MHz receiver built using the techniques described in this article.
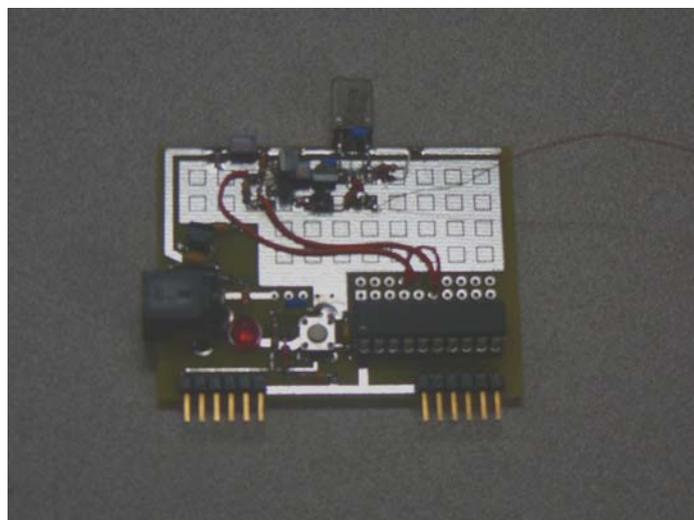

Figure 6: Prototype Receiver Built for this Paper

###

About Microchip Technology

Microchip Technology Inc. (NASDAQ: MCHP) is a leading provider of microcontroller and analog semiconductors, providing low-risk product development, lower total system cost and faster time to market for thousands of diverse customer applications worldwide. Headquartered in Chandler, Arizona, Microchip offers outstanding technical support along with dependable delivery and quality. For more information, visit the Microchip website at http://www.microchip.com.

Additional Information
"Designing Super-Regenerative Receivers"
By Dr. Eddie Insam
Published in Electronics World, April 2002

Applications circuits
PT8A97 Datasheet
Pericom Technology Incorporated

PIC16F690 Datasheet
Microchip Technology Inc.

MCA481wp