

Digital Upconversion and Downconversion

Previously, we discussed complex modulation and baseband signals. The whole point of this discussion is to create a signal that can be used to carry the information bits from one location to another, whether over copper wire, a fiber-optic cable, or electromagnetically through the air. In nearly all cases, the signal needs to ride on a carrier frequency to be efficiently sent from one location to another. To do this, the signal frequency spectrum needs to be able to be moved up and down the frequency axis at will. This is a process of (frequency) upconversion and downconversion. All early methods used analog circuits to accomplish this task. In the past couple of decades, digital circuits, particularly FPGAs, have developed the computational capacity to perform these functions in many cases and offer important advantages over analog methods. These methods are known as digital upconversion and downconversion (also known as DUC and DDC, respectively, within the industry).

The process of upconversion is to take a signal that is at baseband (the frequency representation of the signal) and move or shift that frequency spectrum up to a carrier frequency. The width of the signal's frequency spectrum does not change; it is just moved to another part of the frequency spectrum. One common area of confusion is what happens to the negative part of the frequency spectrum. In a baseband signal, the negative frequency components overlies the positive components. Positive and negative frequencies are distinguished using complex representation. When the signal is upconverted, the positive and negative components are "unfolded" from on top of each, with the negative components below the carrier and the positive components above the carrier frequency.

A common example of this type of conversion occurs with speech. When you speak into a microphone, the sound waves create an electrical baseband signal, with frequency content from near 0 to about 3 kHz. With music, the frequency range can be much greater, up to 20 kHz. When we hear, the vibrations our ears detect are within this frequency range (in fact, our ears cannot hear frequencies beyond this range). This signal can be sampled and converted into a digital baseband signal, or it can remain as an analog signal. To transmit this signal over the air using electromagnetic radio waves, we need to increase the frequency. For example, the commercial AM radio system in the United States operates between 540 and 1600 kHz. The baseband speech signal is multiplied, or "mixed," with the much higher frequency carrier signal. This process superimposes the baseband signal on top of the carrier. For example, if the carrier frequency is 600 kHz, the upconverted 3 kHz speech signal will now ideally have a frequency or spectral

content between $600 - 3 \text{ kHz}$ and $600 + 3 \text{ kHz}$. This process also allows for multiple users to transmit and receive signals simultaneously because each user can be assigned a different carrier frequency and occupy different portions of the frequency spectrum. The upconversion and downconversion process can place the baseband signal at any desired “channel” frequency, allowing many different signals to occupy a common frequency band or range without interference. In the AM radio example, the carrier frequencies are spaced 10 kHz apart. This frequency separation is sufficient to prevent interference between stations.

Traditionally, this upconversion and downconversion process was done using analog signals and analog circuits. Analog upconversion with real (not complex) signals is depicted in [Figure 11.1](#).




Figure 11.1

To see why a mixer works, let’s consider a simple baseband signal of 1 kHz tone and a carrier frequency of 600 kHz. Each can be represented as a sinusoid. A real signal, such as a baseband cosine, has both positive and negative frequency components. At baseband, these components overlie each other, so this is not obvious. But once upconverted, the two components can be readily seen both above and below the carrier frequency.

The equation for the upconversion mixer is

$$\cos(\omega_{\text{carrier}}t) \cdot \cos(\omega_{\text{signal}}t) = \frac{1}{2} [\cos((\omega_{\text{carrier}} + \omega_{\text{signal}})t) + \cos((\omega_{\text{carrier}} - \omega_{\text{signal}})t)]$$

or

$$\cos(2\pi \cdot 600,000 \cdot t) \cdot \cos(2\pi \cdot 1000 \cdot t) = \frac{1}{2} [\cos(2\pi \cdot 599,000 \cdot t) + \cos(2\pi \cdot 601,000 \cdot t)]$$

The result is 2 tones, of half amplitude, at 599 and 601 kHz.

11.1 Digital Upconversion

This upconversion process can also be done digitally. Let’s assume that the information content, whether it is voice, music, or data, is in a sampled digital form. In fact, as we covered in Chapter 9 on modulation, this digital signal is often in a complex constellation form, such as QPSK or QAM, for example.

If we want to transmit this information signal, at some point, it must be converted to the analog domain. In the past, the conversion from digital to analog occurred when the signal was in

baseband form because the data converters could not handle higher frequencies. As the speeds and capabilities of analog-to-digital converters (ADCs) and digital-to-analog converters (DACs) improved, it became possible to perform the upconversion and downconversion digitally, using a digital carrier frequency. The upconverted signal, which has much higher frequency content, can then be converted to analog form using a high-speed DAC.

Upconversion is accomplished by multiplying the complex baseband signal (with I and Q quadrature signals) with a complex exponential of frequency equal to the desired carrier frequency.

The complex carrier sinusoid can be generated using a lookup table or implemented using any circuit capable of generating two sampled sinusoids offset by 90° (see Figure 11.3). If we want to do this digitally, the sample rates of the baseband and carrier sinusoid signal must be equal. Since the carrier signal is usually of much higher frequency than the baseband signal, the baseband signal has to be interpolated, or upsampled, to match the sample frequency of the carrier signal. Then the mixing, or upconversion, process results in the frequency spectrum shift depicted in Figure 11.2.




Figure 11.2

To simplify things, let's assume the output of the modulator is a complex sinusoid of 1 kHz. The equation for this upconversion process is

$$\begin{aligned} & [\cos(\omega_{\text{carrier}}t) + j\sin(\omega_{\text{carrier}}t)] \cdot [\cos(2\pi \cdot 1000 \cdot t) + j\sin(2\pi \cdot 1000 \cdot t)] = \\ & \frac{1}{2} \cdot [\cos((\omega_{\text{carrier}} - 2\pi \cdot 1000) \cdot t) + \cos((\omega_{\text{carrier}} + 2\pi \cdot 1000) \cdot t)] - \\ & \frac{1}{2} \cdot [\cos((\omega_{\text{carrier}} - 2\pi \cdot 1000) \cdot t) - \cos((\omega_{\text{carrier}} + 2\pi \cdot 1000) \cdot t)] + \\ & j \cdot \frac{1}{2} \cdot [\sin((\omega_{\text{carrier}} + 2\pi \cdot 1000) \cdot t) + \sin((\omega_{\text{carrier}} - 2\pi \cdot 1000) \cdot t)] + \\ & j \cdot \frac{1}{2} \cdot [\sin((\omega_{\text{carrier}} + 2\pi \cdot 1000) \cdot t) + \sin((-\omega_{\text{carrier}} + 2\pi \cdot 1000) \cdot t)] \end{aligned}$$

Simplifying, this becomes

$$\cos((\omega_{\text{carrier}} + 2\pi \cdot 1000) \cdot t) + j \cdot \sin((\omega_{\text{carrier}} + 2\pi \cdot 1000) \cdot t)$$




Figure 11.3

We can discard the imaginary portion. It is not needed because, at carrier frequencies, both positive and negative baseband frequency components can be represented by the spectrum above and below the carrier frequency.

The final result at the output of the DAC is

$$\cos((\omega_{\text{carrier}} + 2\pi \cdot 1000) \cdot t)$$

Note there is only a frequency component above the carrier frequency. The reason is that the input was a complex sinusoid rotating in the positive (counterclockwise) direction; there was no negative frequency component in this baseband signal.

Normally, a baseband signal has both positive and negative frequency components. For example, the complex QPSK modulator output can jump both clockwise and counterclockwise depending on the input data sequence. When upconverted, the positive and negative baseband components no longer overlap each other, but are unfolded on either side of the carrier frequency. The baseband signal with a frequency spectrum of 0-10 kHz occupies a total of 20 kHz, with 10 kHz on either side of the carrier frequency. But keep in mind that the baseband signal represents the positive and negative frequencies using quadrature form, so it is in the form of I and Q signals, each with a frequency spectrum from 0 to 10 kHz.

Frequently, there are several steps in upconverting to the final frequency used for transmission, as shown in Figure 11.4. There are several advantages to upconverting in steps. Generally, the DAC must operate at least $2\frac{1}{2}$ times the carrier frequency. For signals in the gigahertz range, this exceeds the capacity of most DACs (although very high performance DACs can now operate at several gigahertz conversion rates). Another consideration is filtering. If several upconverting steps are used, filtering can be applied at each step.

Often, multiple stages of filtering are required to meet overall system requirements, such as the spectral emission requirements in the transmitted RF signal. These requirements are designed to ensure the transmitted signal does not cause interference with other users

A typical digital upconverting circuit is shown in Figure 11.4.

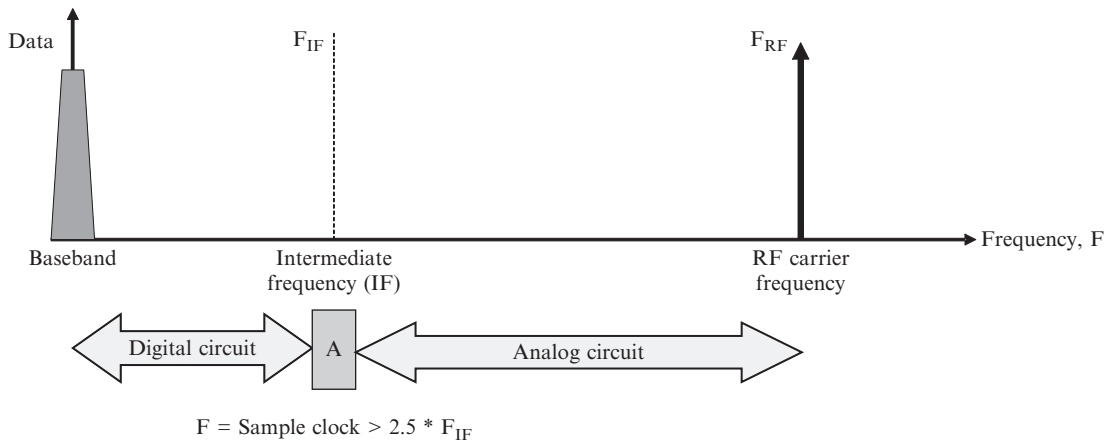


Figure 11.4

11.2 Digital Downconversion

Digital downconversion is the opposite of upconversion. The circuit diagram for this process looks similar, as you can see in Figure 11.5.

Digital downconversion involves sampling, so naturally aliasing needs to be considered. The Nyquist sampling rule states that the sampling frequency must be at least twice the highest frequency of the signal being sampled. In practice, usually the sampling frequency is at least $2\frac{1}{2}$ times the rate of signal frequency, to allow an extra margin for the transition band of the digital low-pass filters following.

But this is not quite true. It is possible to sample a signal at a frequency lower than its carrier frequency. The Nyquist rule applies to the actual bandwidth of the signal, not the frequency of the carrier.




Figure 11.5

11.3 IF Subsampling

Using a technique called IF subsampling, we are able to sample at a much lower frequency than the carrier frequency. The term *IF subsampling* is used because the frequencies typically used in this technique lie somewhere between baseband and RF. The term *IF* refers to intermediate frequency. In this case, we are going to deliberately take advantage of an alias of the signal of interest. This is best illustrated using an example.

Let's use a 4G (fourth generation) wireless example. The signal of interest lies at 2500 MHz. Analog circuits are used to downconvert the signal to an IF, or intermediate frequency. Let's assume that we have an ADC sampling at 200 MHz. Further, let us also assume we have 20 MHz BW IF signal centered at 60 MHz, as in [Figure 11.6](#), which we are trying to sample and downconvert for baseband processing. By the Nyquist rule, we can sample up to $\frac{1}{2}$ the sample rate, or 100 MHz. To allow easier post-sampling filtering, we may want to limit this amount to 80 MHz instead. Our signal here lies between 50 and 70 MHz, so these conditions are met.




Figure 11.6

This process should be familiar from Chapter 3 on sampling. Now let us consider what happens if the IF signal is centered at 460 MHz, as in [Figure 11.7](#), rather than 60 MHz.

As we also discussed in Chapter 3 on sampling and aliasing, there is no way to distinguish between a signal and an alias at a multiple of the sample frequency. When IF subsampling is performed, this aliasing can be used to our advantage.




Figure 11.7

In our examples, any signal that aliases to a baseband frequency from -80 MHz to $+80$ MHz can work. Remember, the digital downconversion multiplies by a complex exponential, which can be rotating either counterclockwise or negative, thereby either shifting the spectrum left or right. Figure 11.8 provides another example, showing downconversion from a 340-MHz carrier frequency.

The signal at 340 MHz is aliased as if it were at -60 MHz and then can be shifted to baseband. The only areas of the spectrum that cannot be properly aliased down are those that alias to the region near the Nyquist frequency—in our case, 100 MHz. So signals close to 300 MHz, 500 MHz, and 700 MHz cannot be sampled in this way. But the rest of the spectrum can be sampled, so long as the sampling frequency meets the requirement of being more than twice the bandwidth of the signal of interest. In practice, the sampling frequency is often much higher than the signal bandwidth, which offers the additional advantage of allowing an increase in signal-to-noise ratio when the baseband signal is decimated to a lower frequency for baseband processing.

In our example, the downconverted signal has a spectrum from -10 to $+10$ MHz. This requires a sampling frequency of at least 20 MHz for both I and Q. If a decimate-by-4 FIR filter is used to low-pass-filter the downconverted signal, then further complex symbol processing can take place at 25 MHz. If we sample at a much higher rate and then low-pass-filter, a much greater percentage of the sampling quantization noise can be eliminated. Quantization noise, as discussed in Chapter 3, is due to the effect of the




Figure 11.8

signal being sampled and mapped to specific amplitude levels that are limited by how many bits of resolution are available in the ADC. This noise is broadband, meaning that it is distributed evenly across the frequency spectrum. The effect of the low-pass filter is to attenuate much of this quantization sampling noise, along with unwanted signals, in the stopband region. The net effect is the equivalent of adding 1 bit of precision or 6 dB to the SNR, for every factor of two in decimation and filtering.

In our example, if we use a 12-bit ADC at 200 MHz F_s and perform the digital downconversion, low-pass-filter, and then decimate by 4, it is the equivalent of sampling at 50 MHz F_s using a 14-bit ADC. The decimation filter has, in essence, added 2 bits of resolution to the ADC, or 12 dB to the SNR of the sampled signal. This concept can be taken to extremes. There is a class of ADCs, called sigma-delta converters, that run at very high frequencies relative to the signal they are sampling, but use only a 1-bit sampler. An effective 10- or 12-bit ADC can be built using a single-bit sampling front end running at extremely high frequencies followed by decimation filtering stages.

Now let us go back to IF subsampling. In theory, we could sample a signal at any arbitrary frequency, but there must be some practical limitations. So far, we have discussed only the limitation of *not* sampling signals that alias to near the Nyquist frequency. But there are other limitations, and we discuss two of them here.

First, we must consider the performance of the ADC. Two principal characteristics are of concern here. Of course, the first is the maximum sampling rate at which the ADC operates.

In our example, we assumed an ADC that could sample at 200 MHz. When we use IF subsampling, we must also consider the analog bandwidth of the sampling circuit in the ADC. Ideally, the ADC samples the signal for an infinitely small instant of time and converts that measurement into a digital number. In practice, this circuit has a sampling window, or period of time in which it samples. The narrower this window, the higher signal frequencies it can sample. This specification is given in the datasheets provided by ADC manufacturers. In our example, we sampled a signal at 460 MHz. So we should check that our ADC has an analog front-end bandwidth of 500 MHz or higher.

Another factor that must be considered is clock jitter. Clock jitter is the amount of timing variability of the edge of the ADC sampling clock from cycle to cycle. It can be readily seen on an oscilloscope of sufficient quality. It is more easily seen as clock phase noise on a spectrum analyzer. Jitter shows up as spectral noise tapering off on either side of the clock frequency component. The less jitter, the more closely the clock appears as simply a vertical line in the frequency response. The effect of clock jitter is proportional to the frequency of the signal being sampled. It limits the SNR according to the following relationship:

$$\text{SNR}_{\text{jitter}} = 20 \log (1/(2\pi \cdot F_{\text{signal}} \cdot t_{\text{jitter}}))$$

An example may make this concept clearer. In our example, the ADC clock is 200 MHz, which has a period of 5 nanoseconds (ns). Let us assume this clock has jitter of 5 picoseconds (ps). This doesn't sound like too much—only 0.1% of the clock period. First, we compute the SNR limitation due to clock jitter with an input signal at 60 MHz and then at 460 MHz:

60 MHz IF signal:

$$\text{SNR}_{\text{jitter}} = 20 \log (1/(2\pi \cdot 60 \cdot 10^6 \cdot 5 \cdot 10^{-12})) = 54.5 \text{ dB}$$

460 MHz IF signal:

$$\text{SNR}_{\text{jitter}} = 20 \log (1/(2\pi \cdot 460 \cdot 10^6 \cdot 5 \cdot 10^{-12})) = 36.8 \text{ dB}$$

This level of clock jitter limits the ADC performance for a 60 MHz signal at any level of precision beyond 10 bits. For a 460 MHz signal, the clock jitter limits the SNR to 36 dB, or about 6 bits. Clearly, this level of clock jitter is excessive in this IF subsampling application.

An analogy can be the same strobe light concept used in Chapter 3 on sampling and aliasing. The strobe light is assumed to flash at exact intervals. Any variance, or jitter, in the flashing intervals causes distortion in the position of the red dot on the wheel. The IF subsampling process would be as if the wheel made multiple rotations between strobe flashes, though it would appear that the wheel was rotating at a slower speed. But because the wheel is actually rotating much faster, any sampling errors due to jitter in the strobe flash are magnified, making the red dot appear blurred.

A very reasonable question to ask is if we can do something similar with a DAC to utilize aliased versions of the digitally upconverted signal. The answer is yes, but with an important limitation that discourages use of this technique in practice.

A major difference between an ADC and DAC is that while an ADC converts the signal to a digital representation, the DAC converts the digital samples to analog form *and* performs a “sample and hold” function of the analog signal between clocks. It is this sample and hold function that is a critical difference.

The sample and hold function is a rectangular filter in the time domain. Each digital sample input is an impulse of a given magnitude. The output is a rectangular shape of the input magnitude. The DAC impulse response, just like in any filter, defines the frequency response. In this case, the rectangular impulse response yields a sinc function $[\sin(x)/x]$ response in the frequency domain. This is shown in Figure 11.9, with the nulls in the frequency response corresponding to multiples of the DAC conversion or clock frequency (F_s).

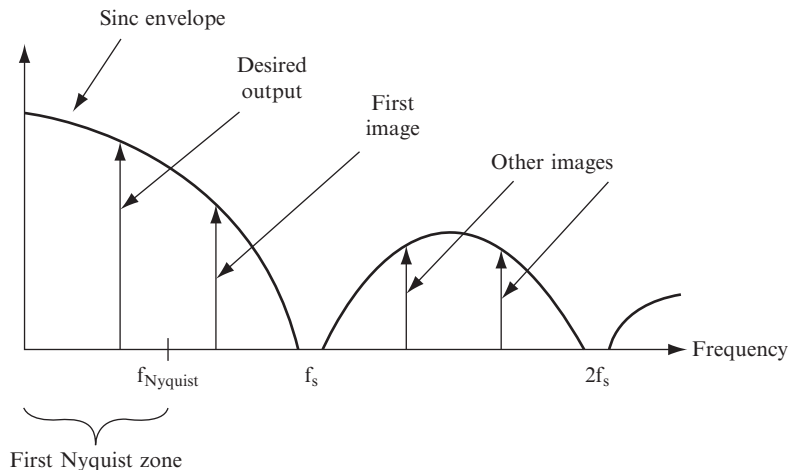


Figure 11.9

What this means in practice is that there will be a reduction in signal level at higher frequencies. By the time the DAC frequency response reaches the Nyquist frequency ($\frac{1}{2} F_s$), it will have a droop of nearly 4 dB. The peak of the first lobe is 6 dB below the DC response, and the second lobe peak a further 6 dB lower. Due to this attenuation, usefulness of the DAC output at frequencies well above the DAC F_s is limited. Moreover, because the frequency response is not flat, it often needs to be compensated for. The closer the IF signal lies to multiples of the DAC F_s , the more distorted the frequency response is. Therefore, the IF frequency is usually limited to the first Nyquist zone. Furthermore, it is usually limited to about 80% of the Nyquist frequency or to 40% of the DAC F_s . For a DAC being clocked at 250 MHz, the upconverted IF signal should be at 100 MHz or less. For example, if the

IF carrier is chosen to be 70 MHz with a complex baseband signal extending to 10 MHz, the IF signal will occupy a spectrum from 60 to 80 MHz. After the DAC, the analog output will need to be filtered using an analog filter to remove the higher frequency DAC images and the DAC clock harmonics. Often, a surface acoustic wave (SAW) bandpass filter is used, and the choice of IF frequency is often determined in part by the frequencies where the SAW filters are commercially available.

Many systems, particularly if the IF signals have fairly wide bandwidth (like our 4G wireless example with an IF signal BW of 20 MHz), need to compensate for the frequency response droop due to the sinc response of the DAC even at these low IF frequencies. To do this, these systems use a sinc compensation FIR filter, after the NCO and complex multiplier performing the digital upconversion.

Figure 11.10 shows the frequency response of a typical sinc compensation FIR filter, calculated using a popular digital signal processing design tool called MATLAB[®] (available from The MathWorks). It compensates up to 80% of the Nyquist frequency. This digital filter would immediately precede the DAC and be placed after the digital upconversion process. It would have a data throughput rate equal to the DAC conversion rate. The associated coefficients are shown in Figure 11.11.




Figure 11.10

```
Numerator:  
-0.0065468956647895373  
 0.014409998368132634  
-0.0291511111437169457  
 0.053017946956878462  
-0.091122876704257547  
 0.15519455343211477  
-0.28282698843194176  
 0.68492126102237871  
 0.68492126102237871  
-0.28282698843194176  
 0.15519455343211477  
-0.091122876704257547  
 0.053017946956878462  
-0.0291511111437169457  
 0.014409998368132634  
-0.0065468956647895373
```

Figure 11.11
Sinc filter coefficients