Department of Engineering

Lecture 23: Mixers, Voltage Controlled Oscillators and Spectrum Analyzers

Matthew Spencer
Harvey Mudd College
E157 – Radio Frequency Circuit Design

Department of Engineering

Mixers

Matthew Spencer
Harvey Mudd College
E157 – Radio Frequency Circuit Design

2

In this video we're going to examine new components called mixers, which are used for frequency translation.

Department of Engineering

Mixers Multiply Signals, Results in ω Shift

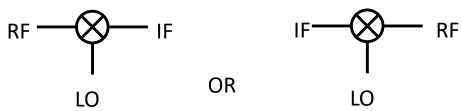

$IF = RF \cdot LO$ OR $RF = IF \cdot LO$

Usually LO must be large sinusoid

Downconversion Example

$$RF(t) = V_{RF} \cos \omega_{RF} t$$

$$LO(t) = V_{LO} \cos \omega_{LO} t$$

$$IF(t) = V_{RF} V_{LO} \cos \omega_{RF} t \cos \omega_{LO} t$$

$$= \frac{V_{RF} V_{LO}}{2} (\cos(\omega_{RF} + \omega_{LO})t + \cos(\omega_{RF} - \omega_{LO})t)$$

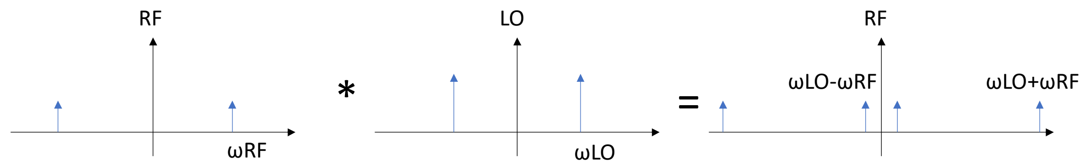

3

I’ve pictured mixers in their two common configurations here. There are three ports on a mixer, usually labeled RF, IF and LO. These stand for radio frequency, intermediate frequency and local oscillator. Regardless of whether RF is the input and IF is the output or IF is the input and RF is the output, the mixer does the same thing, which is multiply the input signal by LO. To do this, the mixer generally relies on LO being a very large signal.

CLICK This multiplication has the effect of moving frequencies in the input signal to a different place, which is the primary purpose of a mixer. Here’s an example of how that works. The RF signal is given by a sinusoid with some frequency omega_rf, and the LO signal is given by some sinusoid with a frequency omega_lo. When we multiply these to find IF, we can use the angle addition formula to convert the result from a product of sinusoids to a sum of sinusoids. We get one signal output at a frequency of omega_rf + omega_lo and one signal output at omega_rf – omega_lo.

CLICK Here’s a graphical version of that same down conversion process. The mixer multiplies the RF and LO signals in the time domain, which means that it convolves them in the frequency domain. That means we get two copies of the RF signal, one centered on omega_LO and one centered on minus omega_LO. This results in a copy of our signal close to DC for the RF and LO values we’ve chosen. The term intermediate frequency refers to any frequency that’s lower than RF, but there’s a special term for converting a signal all the way to DC. It is referred to as direct down conversion.

Department of Engineering

One Mixer Output is Undesired, Called Image

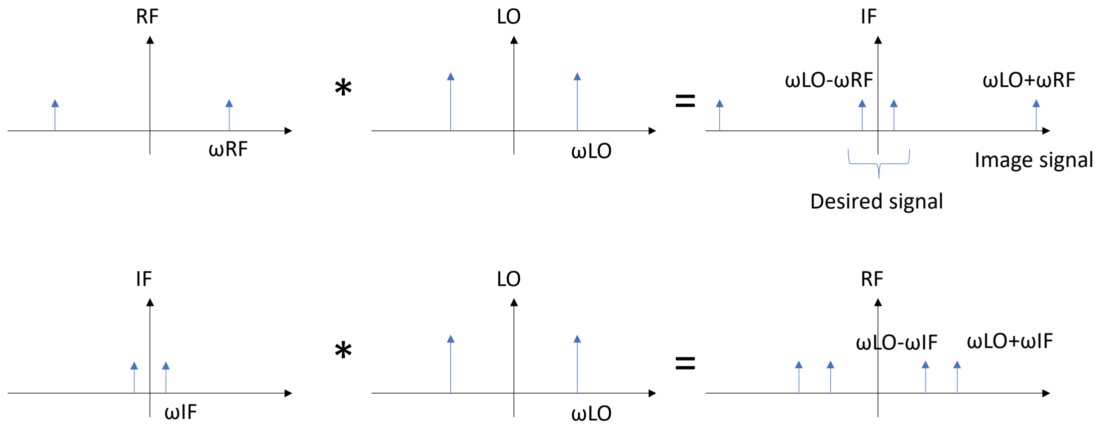

4

The down conversion process from the previous page is replicated here with some additional annotation. The signals in the middle are highlighted as a desired signal. That's the low frequency copy of the high frequency signal we started with in the RF graph. The very high frequency signal is labeled an image signal, and images are generally undesirable. Images in your signal can pose challenges for sampling because they can get aliased down into your normal signal. As a result, mixers are often followed by low-pass or band-pass filters that are called image reject filters.

CLICK The bottom half of this slide illustrates using a mixer in the opposite way, for upconversion. Upconverted signals don't have images per se, but we do have two sinusoids in our output spectrum where we had one in the input. Some very clever transmitters, called single side band transmitters, can cancel out either the higher or lower sinusoid before it is transmitted. Alternatively, one of the features of direct downconversion mixers is that they can make use of both sinusoids in the output spectrum because they stack up at the input.

Department of Engineering

Passive Mixers Have Good Linearity, Noise

Passive Mixer

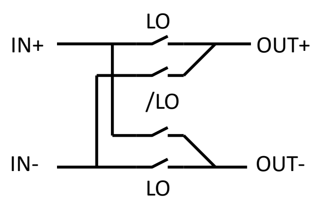

Active Mixer

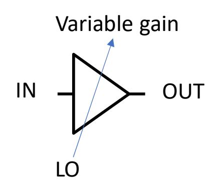

5

There are lots of ways to build mixers, but I want to draw your attention to one big category that splits mixers into two camps. One type of mixer is called a passive mixer, which works by controlling switches to flip an input signal upside down. These mixers don't require external power, and they often have excellent linearity and noise performance.

The other type of mixer is called an active mixer. It works by adjusting the gain of an amplifier with the LO signal. This type of mixer has bigger output signals, but often worse design complexity, linearity, and noise.

These blanket statements don't necessarily apply to any specific mixer. Be sure to do your research before you settle on a mixer for a design rather than just relying on this slide.

Summary

- Mixers convert signal frequencies by multiplying two signals together

$$IF(t) = \frac{V_{RF}V_{LO}}{2}(\cos(\omega_{RF} + \omega_{LO})t + \cos(\omega_{RF} - \omega_{LO})t)$$

- IF – intermediate frequency, RF – radio frequency, LO – local oscillator
- Down converting mixers have an undesired output called an image
- Mixers come in passive and active varieties

Department of Engineering

Mixer Specifications

Matthew Spencer
Harvey Mudd College
E157 – Radio Frequency Circuit Design

7

In this video we're going to talk about some of the specifications you'll find on a mixer datasheet, which will let us dig into how we evaluate mixer performance.

Department of Engineering

Mixer Gain, Isolation, Linearity, are Like Amps

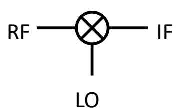

$IF = RF \cdot LO$

- Conversion Gain, $G = P_{@IF}/P_{@RF}$
- Isolation,
  - $I_{RF,LO} = P_{@RF}/P_{LO}$ w/ no input
  - $I_{IF,LO} = P_{@IF}/P_{LO}$ w/ no input
- IIP2 and IIP3 are the same as amplifiers
- Noise Figure is complicated

I’ve thrown a mixer up here for reference. It’s in a downconversion configuration.

CLICK Arguably, the most important specification for a mixer is the conversion gain. This measures how much of the power at the RF frequency gets converted to the IF frequency. However, conversion gain is something of a misnomer because this “gain” is often less than one, and it’s guaranteed to be less than one in passive mixers.

CLICK Isolation describes how much of the powerful LO signal leaks into other ports. This is an extremely important specification because LO appearing at other ports can cause great mischief. LO on the RF port can re-radiate out of the antenna, and LO on the IF port can cause intermodulation with the IF signal.

CLICK The linearity of mixers is described by IIP2 and IIP3, just like amplifiers. Passive mixers have very good linearity, particularly IIP2, so they are popular in many modern receivers.

CLICK The noise figure of mixers is somewhat tricky, so we’re going to give it another slide.

HARVEY MUDD COLLEGE
Department of Engineering

Mixer Noise is Complicated

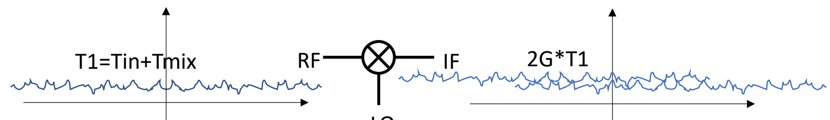

$SNR_{in} = P_{in}/k(T_{in} + T_{mix})B$

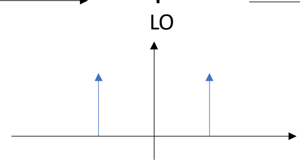

$$SNR_{out,SSB} = \frac{GP_{in}}{2Gk(T_{in} + T_{mix})B} \rightarrow nf_{SSB} = 2\left(1 + \frac{T_{mix}}{T_{in}}\right)$$

$$SNR_{out,DSB} = \frac{2GP_{in}}{2Gk(T_{in} + T_{mix})B} \rightarrow nf_{DSB} = 1 + \frac{T_{mix}}{T_{in}}$$

Assumes signal in both side bands

$$SNR_{out,IEEE} = \frac{GP_{in}}{k(T_{in} + T_{mix})B} \rightarrow nf_{IEEE} = 1 + 2\frac{T_{mix}}{T_{in}}$$

Assumes band-pass knocks out side-band noise

9

To analyze the noise behavior of a mixer, we're going to imagine that we have noise with a temperature T1 at the input, which represents a combination of the input noise temperature and the mixer noise temperature. We're going to get two copies of the input noise spectrum at the output, one from the positive LO sinusoid and one from the negative LO sinusoid, that are both scaled by the conversion gain. Getting multiple copies of noise like this is called noise folding.

The value of Tmix for a lossy mixer is (1/L-1)T just like any other lossy passive. For an active mixer, it's going to be related to the noise factor, but you need to read the noise factors for mixers carefully, as we will see on this slide. As a preliminary to that, I've calculated the input SNR near the input spectrum.

CLICK The signal to noise ratio at the output will reveal that our SNR is degraded by affected by both copies of the noise if we have a single desired RF signal for downconversion. This produces a noise factor that includes a factor of two in front of it. This noise factor is called the single sideband noise factor, and it is widely used by RF designers.

CLICK Another definition of noise factor presumes that we have desirable signal at both omega_lo+omega_rf and omega_lo-omega_rf. This doubles our signal power at the output because both copies land at the same place in the IF signal. As a result, our noise factor is halved. This is called the double sideband noise factor.

CLICK Finally, the IEEE has a third definition of noise figure, which doesn't penalize mixers for

Sufice to say, I think it's easiest to think about the noise temperature at the output of a mixer rather than dabbling in noise figure.

- Lossless mixer adds noise temp Tm=T
- Lossy mixer adds Tmix=(1/L-1)T
- Lossless mixer has noise factor of 2 ← called NF_SSB
- NF_DSB and NF_IEEE also exist, ideal mixers have NF=1

If passive, noise temp is (1/L-1)T, like any other passive
If active, have a noise figure
Noise figure can involve two side bands … blarg

Summary

- Conversion gain: $G = P_{IF}/P_{RF}$
- Isolation: $I_{LO-RF} = P_{RF}(\omega_{LO})/P_{LO}, I_{LO} \quad = P_{IF}(\omega_{LO})/P_{LO}$
- IIP2 and IIP3 are the same. IIP2 counts in direct down conversion.
- Mixers capture noise temperature from the desired band and the image, which leads to multiple NF definitions.

10

Department of Engineering

Voltage Controlled Oscillators

Matthew Spencer
Harvey Mudd College
E157 – Radio Frequency Circuit Design

11

In this video we’re going to learn about another crucial communication component called the voltage controlled oscillator

Department of Engineering

Output Sinusoid Frequency Changes …

$$V_{out} = V_{zp}\cos(f(V_{ctrl})t)$$
$$\approx V_{zp}\cos(K_{VCO}V_{ctrl}t)$$

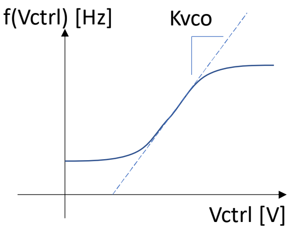

12

Voltage controlled oscillators are elements that take a relatively slow varying input called the control voltage and produce an output sinusoid whose frequency depends on the control voltage, which I've indicated with the notation f(Vctrl), to indicate it's some function. The amplitude is set by the oscillator, and generally not controllable.

The function f is called a tuning curve and I've included an example tuning curve on the right. Tuning curves are generally very non-linear, but it's common to approximate them as linear for small control voltage deviations, which leads to the approximation that the frequency is linearly dependent on the control voltage. The slope relating voltage to frequency in this approximation is called K_VCO

LC VCOs Have Clean Output Spectra

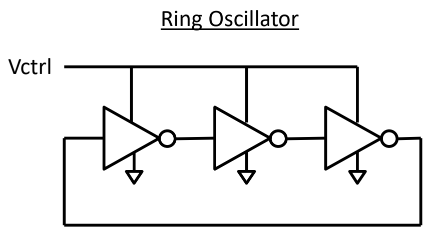

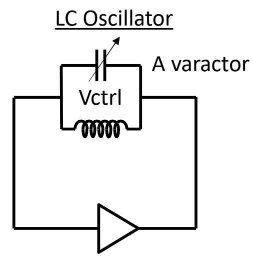

13

There are many ways to make VCOs, but we're going to break them into two broad categories again. Ring oscillators or other digital oscilators are very easy to build – you just hook up a few logic gates, usually inverters, such that they're never satisfied with their internal state, which causes oscillation. Changing the voltage on the power rail of the logic gates changes how fast they propagate signals and, accordingly, the frequency of oscillation. The output signals from a ring oscillators aspire to be square waves, so they tend to have lots of harmonic content.

LC oscillators are built by using an LC network in feedback around an amplifier. If the feedback is designed right, the amplifier will be unstable, which will create an oscillation at the natural frequency of the tank. A variable capacitor called a varactor can be controlled by Vctrl to tune the oscillation frequency. The filtering properties of the LC tank make LC oscillator spectrums very clean, but they tend to be power hungry and tricky to design.

Summary

- Voltage controlled oscillators have a frequency controlled by an external voltage (and usually a fixed amplitude)
- Tuning curves are non-linear, but often approximated as linear
- Ring oscillators make noisier spectra than LC oscillators

Department of Engineering

Spectrum Analyzers

Matthew Spencer

Harvey Mudd College

E157 – Radio Frequency Circuit Design

15

In this video we’re going to assemble the components from the last few videos into a common piece of benchtop equipment.

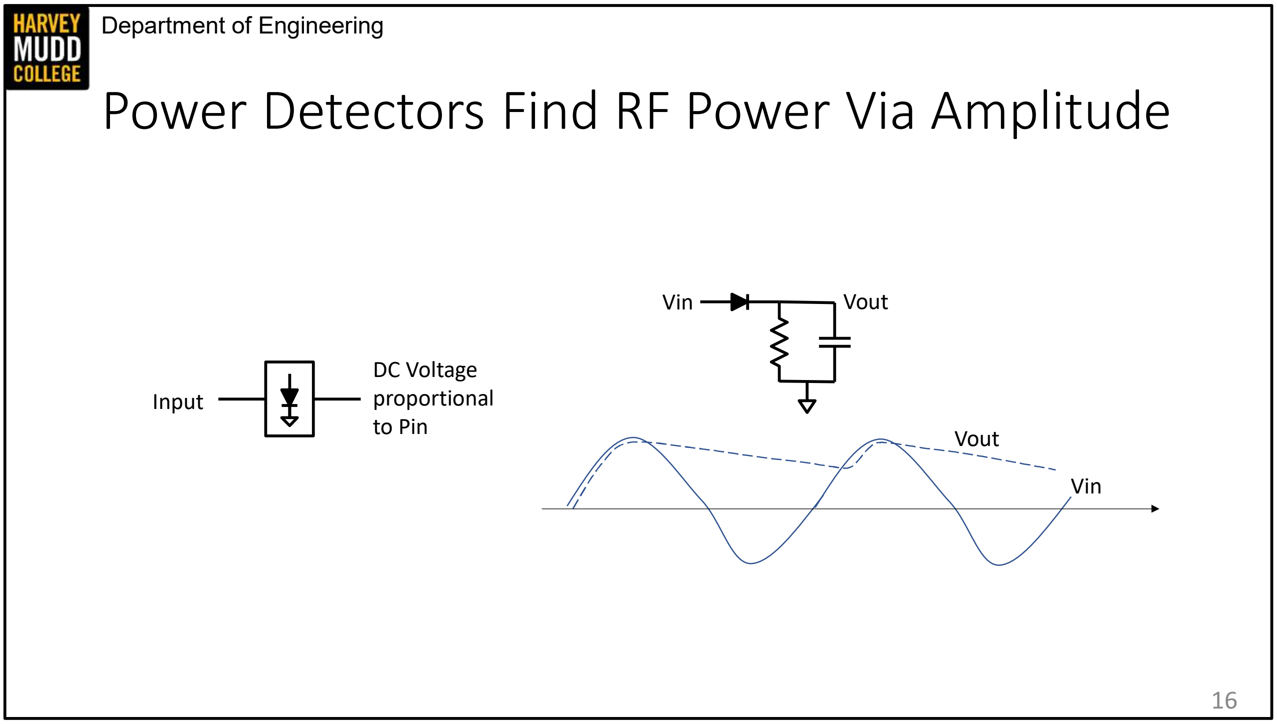

We need one more component to make a spectrum analyzer and it's called a power detector. Power detectors convert signals of any frequency into slowly varying voltages that are proportional to the signal power. The symbol for a power detector includes a diode because most common ways to make a power detector involve a diode. One very common example of a power detector is the envelope detector circuit, pictured on the right. This circuit consists of a diode driving an RC low pass filter, and it's operation is illustrated in the waveforms below. The output voltage follows the input voltage until it peaks, and then the output begins a slow decay based on the time constant of the low pass filter. When the input rises above the decaying output, the output is restored back to the amplitude of the input. If the low pass is slow enough, then the output will closely approximate the input amplitude.

HARVEY MUDD COLLEGE Department of Engineering

Spectrum Analyzers Shift Spectra Into Filter

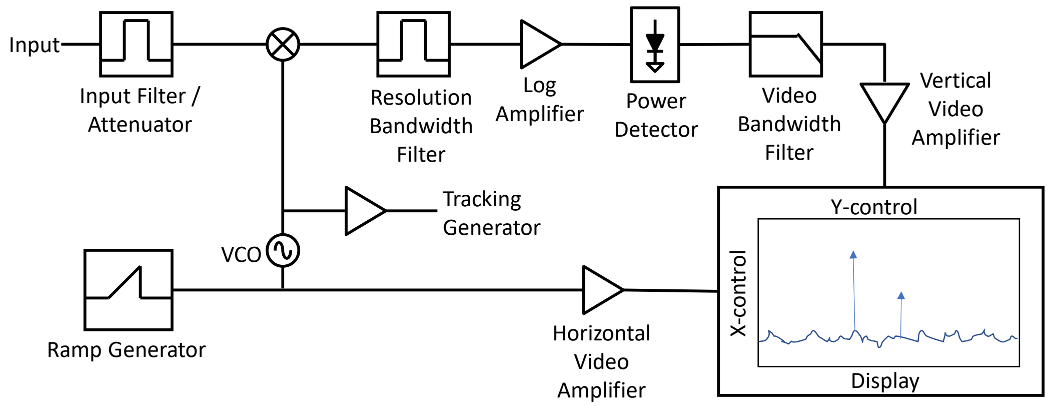

17

A spectrum analyzer is designed to produce a frequency domain picture of an input signal. The ramp generator in the lower left is the main controller for the spectrum analyzer. It produces a ramp the controls the position on the x-axis of the output display that we are currently drawing. It also influences the y-value of what to draw by acting as the control for a VCO. The VCO output is tapped off to an optional output called a tracking generator, which can be used to make Bode plots and a few other things, but it also drives a mixer LO. The RF input of the mixer is the input signal. This means the ramp will result in the VCO shifting the input signal by different amounts of frequency as the LO increases, which is a sentence we'll examine in greater detail on the next slide.

The down-mixed input signal is put through a narrow band-pass filter called the resolution bandwidth, or RBW, filter. This RBW filter sets the total noise power of the signal. The signal is scaled by a log amplifier to convert it to a dB scale, then the power is detected. A little more filtering, including a filter called the video bandwidth or VBW filter, ends with the signal controlling the y-value of the display. So the y-value of the display is set by the output of a mixer being passed through a narrow bandpass filter, and the changing LO on the mixer means that bandpass filter is effectively swept over every frequency in the input.

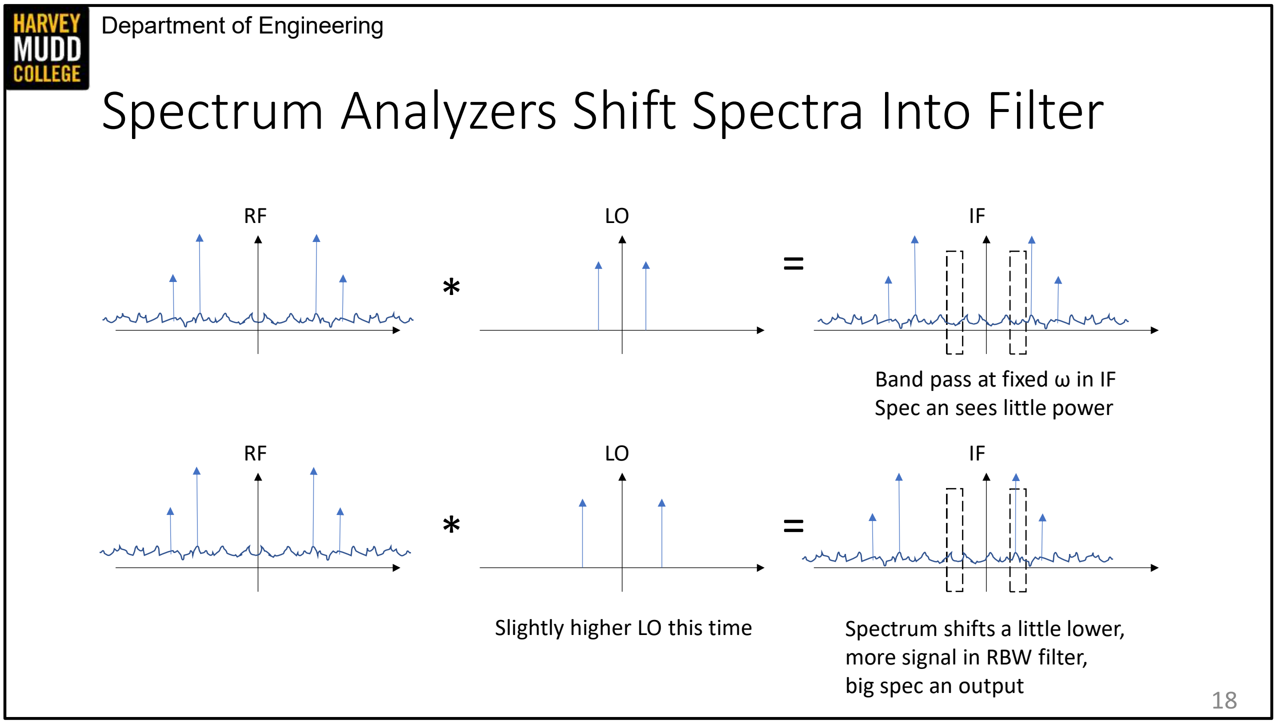

We can look the input and output spectra of the mixer to understand how this works. The input spectrum is a signal with some noise, and we would get two copies of that input spectrum, one at each value of LO. I've only drawn one of those copies on the IF spectrum here, but you can see that this copy has been shifted lefon the IF spectrum by looking at the pattern that used to be centered at zero. I've also indicated the part of the IF spectrum captured by the RBW filter. We can see that this spectrum doesn't have any interesting values in the RBW filter.

CLICK Hoewver, if I increase the LO a little bit, then the copy of the input signal will be at a slightly different frequency, which results in a large signal landing in the RBW filter. As we vary the LO, we effectively sweep every part of the input signal through the RBW filter, which means that we can build up a measurement of each frequency component of our input.

Department of Engineering

RBW sets Noise Floor Level, VBW Sets “Noise”

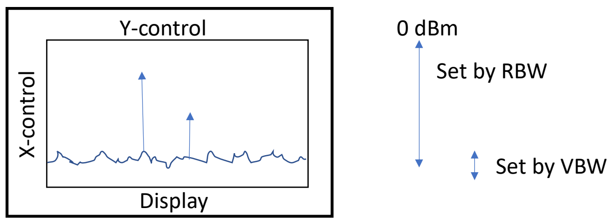

Number of steps = Span/RBW
Time per step α 1/RBW

19

The RBW filter is usually the narrowest bandwidth in our system, which means that it sets the thermal noise power that gets into the spectrum analyzer. The RBW filter is immediately followed by a power detector, so the thermal noise power is converted to a voltage level in the spectrum analyzer plot. The level that the thermal noise lives at is called the noise floor. You can reduce the noise floor in a spectrum analyzer by making the RBW smaller. That will increase the sweep time to measure a spectrum because the total number of measurements the spectrum analyzer takes is related to the number of resolution bandwidths required to cover the entire measurement range. Further, the settling time of a narrow filter is related to the inverse of its bandwidth.

The noise floor will still have a little bit of noise on top of it introduced by the rest of the signal chain. The video bandwidth determines how much this noise floor jumps up and down.

The fact the noise floor changes with RBW gives us some tools to measure very small signals. Shrinking the resolution bandwidth, which often happens automatically when you zoom in on a signal on the spec an, will lower the noise floor. That lets you observe less powerful signals.

Department of Engineering

Summary

- Spectrum analyzers use a mixer and a VCO to slide an input spectrum across a band pass filter, measuring power in each filter band.
- Resolution bandwidth determines the noise floor (and the sweep time). To see small signals, zoom in on their frequency.