

HP

Network Analysis

Solutions

Simplified Filter Tuning Using Time Domain

Application Note 1287-8




Table of Contents

Introduction3
Difficulties of filter tuning4
Ideal tuning method4
Basic characteristics of bandpass filters5
Time-domain response of simulated filters6
Effect of tuning resonators7
Effect of tuning coupling apertures8
Practical examples of tuning filters11
Setting up the network analyzer11
Example 1: Tuning resonators only12
Example 2: Tuning to a “golden filter”16
Example 3: Using simulated results for a template20
Effects of loss in filters21
More complex filters22
Cross-coupled filters22
Duplexers25
Conclusion26
References26
Summary: Hints for time-domain filter tuning27
Appendix A: Understanding basic bandpass filter design28
Appendix B: Using time-domain in the network analyzer for filter tuning31

Introduction

The increase in wireless communications services is forcing more and more channels into less frequency spectrum. To avoid interference, very stringent filtering requirements are being placed on all systems. These systems usually employ coupled resonator filters to handle the power levels and provide the needed isolation. The difficulty of tuning these filters quickly and accurately often limits manufacturers from increasing their production volumes and reducing manufacturing cost.

In a coupled-resonator cavity-tuned filter, the center frequency of each resonator must be precisely tuned. The couplings between resonators must also be precisely set to achieve the proper passband response, low return loss (reflection), and small passband ripple. Setting coupling coefficients and tuning the resonators are as much art as science; often a trial-and-error adjustment process. Until now, there has been no alternative.

This application note describes a method of tuning a filter using the time-domain response of its return loss, which makes filter tuning vastly easier. It is possible to tune each resonator individually, since time-domain measurements can distinguish the individual responses of each resonator and coupling aperture. Such clear identification of responses is extremely difficult in the frequency domain. Coupling coefficients may be precisely set to provide a desired filter response, and any interaction caused by adjustment of the coupling structures and resonators can be immediately determined and accounted for.

Perhaps the most important advantage of the time-domain tuning method is that it allows inexperienced filter tuners to successfully tune multiple-pole filters after only brief instruction. Such rapid proficiency is impossible with previous tuning methods. This technique also lends itself well to the automated production environment, which has always been a challenge.

Difficulties of filter tuning

The interactive nature of coupled-resonator filters makes it difficult to determine which resonator or coupling element needs to be tuned. Although some tuning methods can achieve an approximately correct filter response, final tuning often requires the seemingly random adjustment of each element until the final desired filter shape is obtained. Experienced tuners can develop a feel for the proper adjustments, but months are often required before a novice can be proficient at tuning complex filters. The time and associated cost of tuning, and the difficulty and cost in training new personnel can limit a company's growth and responsiveness to changing customer needs.

Some companies have attempted to automate the tuning process, using robotics to engage and turn the tuning screws, and an algorithmic process to accomplish the tuning. The tuning algorithms are a particular problem, especially when a filter is nearly tuned, at which point the interaction between stages can be so great that final tuning cannot be achieved. New filter designs may require entirely new algorithms, making it even more difficult for test designers to keep up with changing requirements. Manufacturing changes that affect the filter components, such as tool wear or changing vendors, may also cause algorithms and processes to become less effective.

In some cases, tuned filters go through temperature cycling or other environmental stress as part of the manufacturing process, and their characteristics may change as a result. It can be very difficult to identify which resonators or coupling apertures need to be retuned using conventional filter tuning methods.

Ideal tuning method

The solution to these difficulties would be a tuning method that is simple, flexible and deterministic. That is, one in which the individual adjustment goals for each tuning element, resonator, and coupling aperture would not depend upon the other elements in the filter. The response to each tuning screw would be easily identified, and any interactive effect would be immediately seen and accounted for. Ideally, each screw would only need to be adjusted once. Finally, the tuning method would not depend on filter type or shape, or number of filter poles.

This application note presents a technique that clearly identifies the resonator or coupling aperture that needs to be tuned, and enables the operator to see and correct for interactions. Filters can be tuned to match any filter shape within their tuning ranges. Although this technique does not meet the ideal goal of requiring only a single adjustment of each screw, it greatly simplifies and speeds up the filter-tuning process.

Basic characteristics of bandpass filters

First, let's review some basic information and characteristics about bandpass filters.

Bandpass filters are commonly designed by transforming a low-pass filter response to one that is centered about some new frequency. Coupled resonators, which may be lumped LC resonators, coaxial line resonators, cavity resonators, or microwave waveguide resonators, are used to create the upward shift in frequency. The terms resonator, cavity resonator, and cavity will be used interchangeably in this application note. More details on bandpass filter design can be found in Appendix A.

The center frequency of the filter is determined by setting the resonators. In most designs, all resonators are set exactly to the center frequency, with the effects of adjacent coupling included in the calculation of the resonant frequency.

The filter shape, bandwidth, ripple, and return loss are all set by the coupling factors between the resonators. When properly tuned, the resonators have almost no effect on the filter shape. The only exception is that the input and output resonators set the nominal impedance of the filter. Usually an input or output transformer is used to match to a desired impedance. Of course, when the resonators are not properly tuned, the return loss and insertion loss will not be at the optimal levels.

Because the resonators are coupled to each other, tuning one resonator will have the most effect on the adjacent resonators, but it will also have some smaller effect on the remaining resonators. The extent of the effect depends on the coupling factor.

With this information in mind, we are ready to explore the new time-domain tuning technique.

Time-domain response of simulated filters

To introduce this tuning method, we will use simulations to examine what happens to the time-domain response of a bandpass filter when it is tuned. We will start with a relatively simple filter: a five-pole coupled resonator filter with four coupling structures, designed for a Chebyshev response with 0.25 dB of pass-band ripple. In this example, a filter response will be simulated by HP's Advanced Design System (ADS) microwave design software, so that the exact values of constituent components are known. The frequency sweeps will be performed in the simulator, and the results will be downloaded to the vector network analyzer (VNA), where the instrument's time-domain transform application can show the effects of filter tuning. The schematic for the filter is shown in Figure 1.

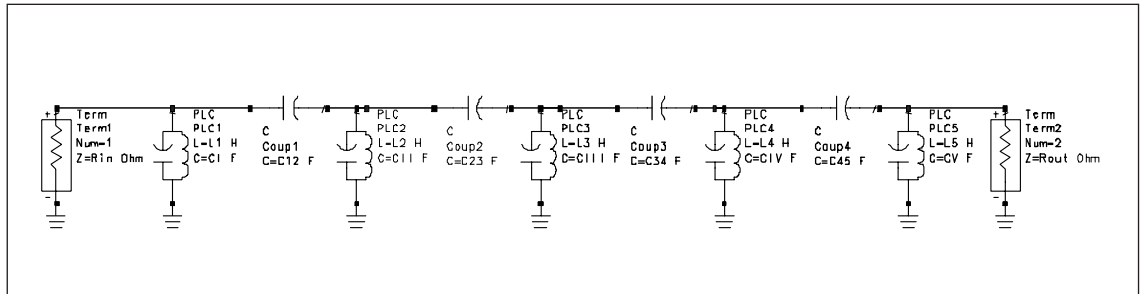


Figure 1. Schematic for five-pole coupled resonator bandpass filter

To set up the measurement for time-domain tuning, the frequency sweep **MUST** be centered at the desired center frequency of the bandpass filter. This is critical, since the tuning method will tune the filter to exactly that center frequency. Next, the span should be set to approximately two to five times the expected bandwidth.

Figure 2 shows the frequency response and time response of the filter. Notice the distinctive dips in the time-domain S_{11} response of the filter. These are characteristic nulls that occur if the resonators are exactly tuned. The peaks between the nulls relate to the coupling factors of the filter, as we will see later. Markers 1 through 5 have been placed to show the characteristic dips corresponding to resonators 1 through 5 in the filter. Although there are some dips to the left of marker 1, those are not part of the filter response. Generally the peaks corresponding to the filter response will be much higher in magnitude than the ones in the $t < 0$ region, which are not meaningful, and usually the dip corresponding to the first resonator will occur near $t=0$.

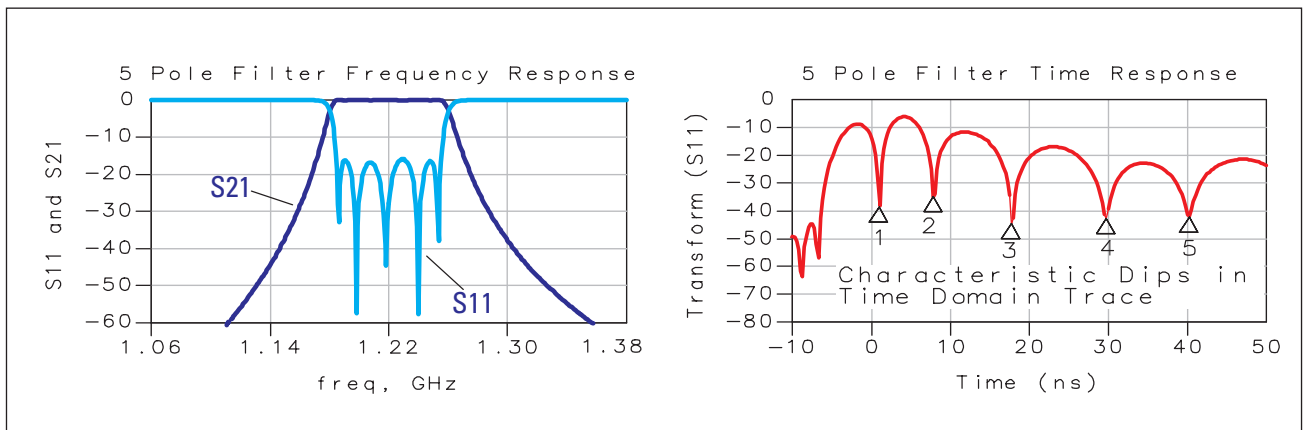


Figure 2. The frequency and time-domain response of a bandpass filter

Effect of tuning resonators

The example filter starts out with the ideal design values, which yields the desired response since it is properly "tuned" by definition. To understand the time-domain response to tuning the resonators, we will monitor the time-domain response while changing (mistuning) the resonator components in the simulation. Figure 3 shows the time-domain traces for three conditions (with the ideal response in the lighter trace). The upper plots show the filter with the second resonator mistuned 2% low in frequency. Note that the first dip has not changed, but the second dip is no longer minimized, and neither are the following dips. If a resonator is substantially mistuned (more than 1%), it will significantly mask the dips of following resonators. Therefore, to identify the mistuned resonator, look for the first dip that is no longer at a minimum. In this case, we see that mistuning resonator 2 causes the second null to move away from its minimum value.

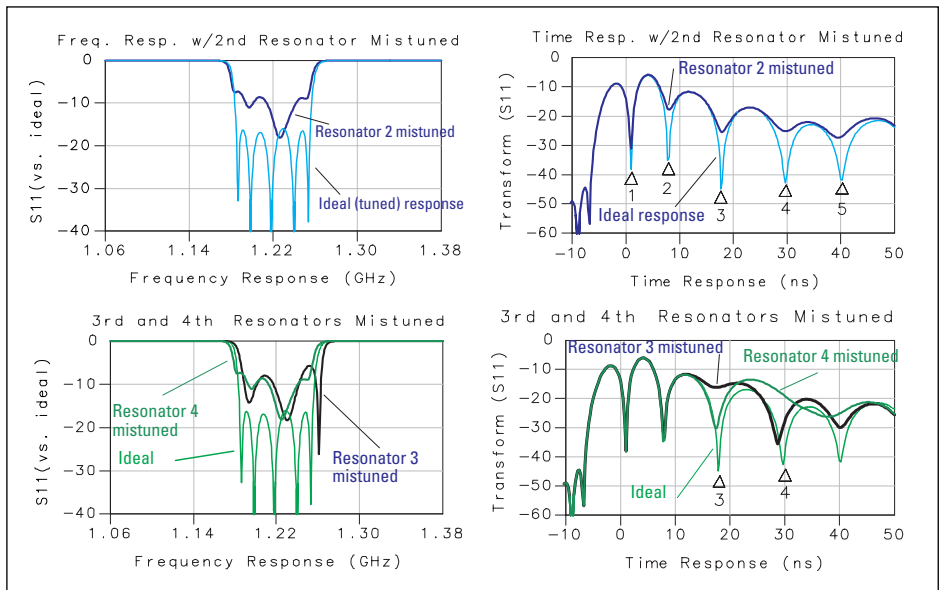


Figure 3. The response of a bandpass filter to tuning the resonators

The lower plots show one response with only the third resonator mistuned 2% high and another one with only the fourth resonator mistuned 2% low. Again, it is easy to identify which resonator is mistuned by looking for the first dip that is no longer minimized. Additional simulations have shown that the characteristic dips are minimized only when the corresponding resonators are set to their correct values. Changing the tuning in **either** direction causes the dips to rise from the minimum values.

The key to this tuning technique is to adjust the resonators until each null is as low as possible. The adjustment will be mostly independent, although if all the resonators are far from the final value the first time through, adjusting a succeeding resonator may cause the null of the previous resonator to rise from its minimum. If this occurs, the null for the previous resonator should be optimized again. Once the succeeding resonator has been tuned and the previous one optimized, additional smaller adjustment to the second resonator will have very little effect on the dip corresponding to the first resonator.

Those who are familiar with the resolution limits of time-domain measurements will know that time-domain resolution is inversely proportional to the frequency span being measured, and they may wonder how it is possible to resolve individual resonators in a filter when the frequency span is only two to five times the filter's bandwidth. Appendix B explains how the time-domain transform relates to bandpass filter measurements in more detail.

One more thing to note from Figure 3 is that the S_{11} frequency response when resonator 2 is mistuned looks almost identical to S_{11} response when resonator 4 is mistuned. This illustrates why it can be difficult to determine which resonator requires tuning when viewing only the frequency-domain measurements.

Effect of tuning coupling apertures

Although simple filters may only allow adjustments of the resonators, many filters also have adjustable couplings. To understand the effects of adjusting the coupling, we will go back to our original "tuned" simulated filter. First, we will examine what happens when we **increase** the first coupling factor by 10%. Figure 4 shows the S_{11} response in both frequency and time domains, both before and after changing the coupling factor. In the frequency domain, we see that the filter bandwidth is slightly wider and the return loss has changed. This makes intuitive sense, because increasing the coupling means more energy should pass through the filter, resulting in a wider bandwidth.

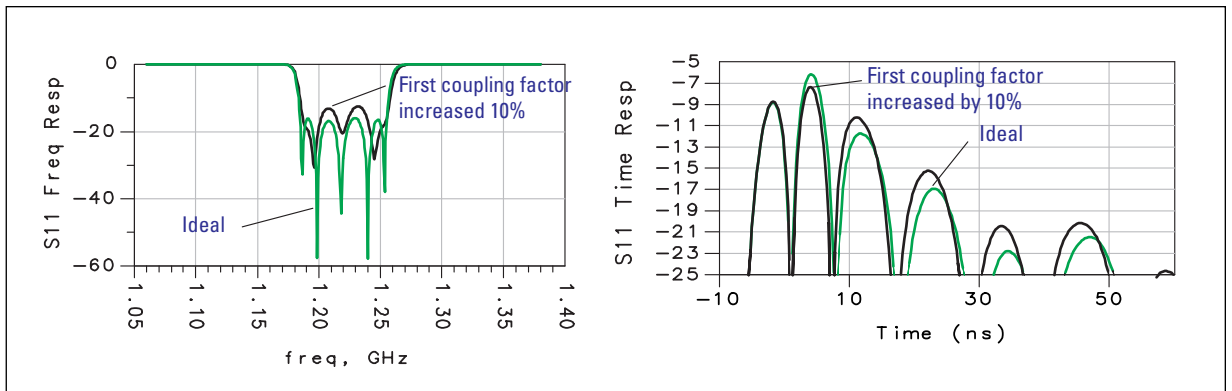


Figure 4. Effect of increasing first coupling factor (darker trace is after adjustment)

In the time-domain, there is no change in the first peak, but the second peak is smaller. While it might seem that the first peak would be associated with the first coupling factor, remember that the first coupling factor comes after the first resonator in the filter, and we have already seen that the first dip after the first peak is related to the first resonator. It turns out that the first peak can be associated with the input coupling, which has not been adjusted in this filter.

The reduction in height of the second peak when coupling is increased makes sense, because increasing the coupling means more energy is coupled to the next resonator. Thus less energy is reflected, so the peak corresponding to reflected energy from that coupling should decrease. Note that the following peaks are higher than before. More energy has been coupled through the first coupling aperture, so there is more energy to reflect off the remaining coupling apertures.

It is important to recognize that changing the first coupling factor will affect the responses of all the following peaks. This suggests that coupling factors should be tuned starting with the coupling closest to the input and moving towards those in the center of the filter. Otherwise, improperly tuned coupling near the input can mask the real response of the inner coupling factors.

Now consider what happens if we take the original filter and **decrease** the second coupling coefficient by 10%. Figure 5 shows that in the frequency domain, the bandwidth of the filter has been reduced slightly and the return loss has changed. Again, this makes sense because decreasing the coupling means less energy will pass through the filter, corresponding to a narrower bandwidth.

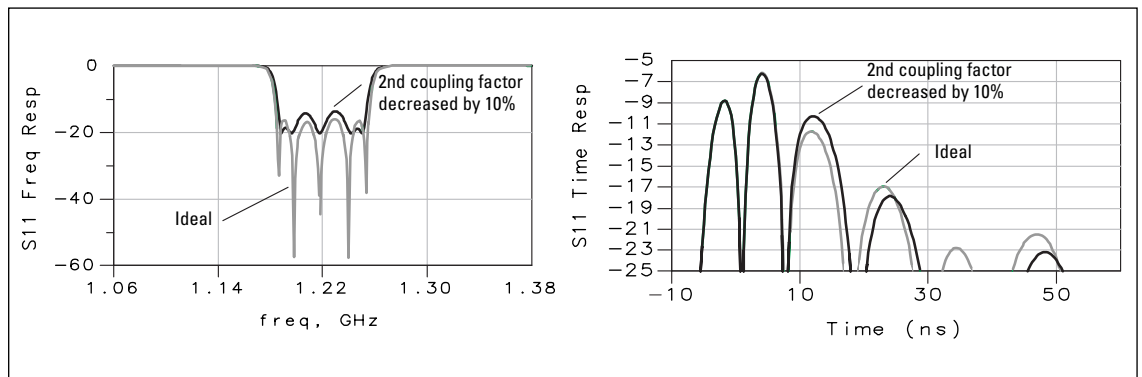


Figure 5. Effect of decreasing second coupling factor (darker trace is after adjustment)

Examining the time-domain trace, we see no change in the first 2 peaks, but the third peak is higher, consistent with more energy being reflected as a result of the decreased coupling. Since the amount of energy coupled to the following resonators and apertures is reduced, the following peaks are all lower in value. Note how well the time-domain response separates the effects of changing each coupling, allowing the couplings to be individually adjusted. In contrast, the S_{11} frequency response trace in Figure 4 is very similar to the one in Figure 5, so it would be very difficult to know which coupling changed from looking at the frequency-domain response.

Thus, we have seen that the coupling factor can be related to the height of the time-domain reflection trace between each of the resonator nulls. The exact relationship also depends on the ratio of the filter bandwidth to the frequency sweep used to compute the time-domain transform. The wider the frequency sweep (relative to the filter's bandwidth), the more total energy is reflected, so the higher the peaks.

The magnitudes of the peaks are difficult to compute because changing the coupling of one stage changes the height of the succeeding peaks. A detailed explanation of relationship between the time-domain response and coupling coefficients is beyond the scope of this application note. Even though it may not be easy to calculate these peaks simply from the coupling coefficients, once the desired values of the peaks are determined, the apertures may be tuned directly in the time domain. One method for determining the desired magnitudes of the peaks is by using a template as described in the next section.

Practical examples of tuning filters

Now that we have an understanding of the relationship between tuning resonators or coupling apertures and the corresponding results in the time-domain response, we are ready to put the theory into practice.

For multi-pole cavity filters that have fixed apertures, it is only necessary to tune for the characteristic dips in the time domain in order to achieve optimal tuning of the filter. To tune a filter with variable coupling coefficients, it is easiest to tune the coupling to a target time-domain trace or template. This target time-domain response for any filter type may be determined in several ways. One method is to use a "golden" standard filter that has the same structure and is properly tuned for the desired filter shape. This filter can be measured and the data placed in the analyzer's memory. Each subsequent filter can be tuned to obtain the same response.

An alternative is to create a filter from a simulation tool, such as HP's Advanced Design System. The simulated response can be downloaded into the network analyzer and used as a template. This is a very effective approach, as there is great flexibility in choosing filter types. The only caution is that each real filter has limits on the Q of the resonators and the tuning range of the coupling structures and resonators. It is important to make the attributes of the simulation consistent with the limitations of the structures used in the real filters.

In this section, we will begin with a discussion of how to set up the network analyzer to tune bandpass filters in the time domain, and then we will show three examples to illustrate how to tune both resonators and coupling apertures in real filters.

Setting up the network analyzer

It is essential to set the center frequency of the analyzer's frequency sweep to be equal to the desired center frequency of the filter, since tuning the filter in the time domain will set the filter's center to this frequency. Choose a frequency span that is 2 to 5 times the bandwidth of the filter. A span that is too narrow will not provide sufficient resolution to discern the individual sections of the filter, while too wide a span will cause too much energy to be reflected, reducing the tuning sensitivity.

The primary parameter to be measured is S_{11} (input match). However, for time-domain responses more than halfway through the filter, the responses often get more difficult to distinguish. Even in low-loss filters, there can be significant return loss differences between the input and output due to loss in the filter. In addition, there is a masking effect that tends to make reflections from couplings and resonators farther from the input or output appear smaller, since some of the incident energy has been lost due to earlier reflections in the device. For these reasons, the most effective way to tune is to look at both sides of the filter at once, so a network analyzer with an S-parameter test set is recommended. To aid in tuning, the instrument's dual-channel mode can be used to measure the reverse return loss (S_{22}) on a second channel. With this setup, you will tune the first half of the resonators and couplings using the S_{11} response, and tune the remaining ones using the S_{22} response. Keep in mind that you need to count resonators and coupling apertures starting from the port where the signal is entering the filter for that measurement. Thus for S_{11} , the first dip would correspond to the resonator closest to the input port of the filter. For S_{22} , the first dip would correspond to the resonator closest to the output port of the filter.

For the network analyzer time-domain setup, the bandpass mode must be used. The start and stop times need to be set so that the individual resonators can be seen. For most filters, the start time should be set slightly before zero time, and the stop time should be set somewhat longer than twice the group delay of the filter. If the desired bandwidth is known, the correct settings can be approximated by setting the start time at $t = -(2/\pi BW)$ and the stop time at $t = (2N+1)/(\pi BW)$, where BW is the filter's expected bandwidth, and N is the number of filter sections. This should give a little extra time-domain response before the start of the filter and after the end of the filter time response. If you are tuning using both the S_{11} and S_{22} responses of the filter, you can set the stop time to a smaller value, since you will use the S_{22} response to tune the resonators that are farther out in time (and closer to the output port).

The format to use for viewing the time-domain response is log magnitude (dB). It may be helpful to set the top of the screen at 0 dB.

Example 1: Tuning resonators only

The first example is a simple five-pole cavity filter with fixed apertures, so only the resonators can be tuned to adjust the center frequency. This filter has a center frequency of 2.414 GHz and a 3 dB bandwidth of 12 MHz. The network analyzer is set up for this same center frequency and a span of 50 MHz. Dual channel mode is used to display both S_{11} and S_{22} . The time-domain response is set up to sweep from -50 ns to 250 ns.

Experience has shown that it is best to begin tuning from the input/output sides and move toward the middle. Figure 6 shows the time-domain response after the first and fifth resonators have been tuned to obtain the lowest dips. Note that the first resonator closest to the input corresponds to the first dip in S_{11} , while the fifth resonator, which is the first one when looking in the reverse direction, corresponds to the first dip in S_{22} . These responses are good illustrations of masking. Even though the fifth resonator is correctly tuned, you cannot see that from looking at the S_{11} response. Similarly, you cannot see that the first resonator is tuned by looking only at the S_{22} response.




Figure 6. Time-domain response of 5-pole filter after tuning resonators 1 and 5

Next, we tune the second resonator, readjusting the first one as needed to keep its dip minimized. Then we go back to the output side and tune the fourth resonator, readjusting the fifth one as needed. Finally, we tune the third resonator in the middle, readjusting the second and fourth resonators as needed. It may be necessary to go back and readjust each of the resonators again to fine-tune the response. Figure 7 shows the time-domain response after the filter has been tuned. Figures 8 and 9 show the frequency domain reflection and transmission responses. Note that the center frequency has been set precisely to 2.414 GHz without looking at the frequency domain while tuning. With frequency domain tuning methods, it is often possible to tune the filter to have the correct shape while the center frequency is slightly off. The time-domain tuning method centers the filter very accurately.




Figure 7. Time-domain response of 5-pole filter after tuning all resonators




Figure 8. Final reflection frequency response




Figure 9. Final transmission frequency response

Now, what if we want to change the center frequency of the filter, for example to 2.42 GHz? We simply need to repeat the tuning process with the analyzer's center frequency set to the new frequency. Figure 10 shows the time-domain response (in bold) that results from measuring the 2.414 GHz filter after changing the network analyzer's center frequency to 2.42 GHz. The original time-domain response is shown in the lighter trace. It is clear that the resonator dips are no longer at their minimums, so the resonators need to be retuned. Adjusting the resonators to minimize the dips again will result in a filter tuned to a center frequency of 2.42 GHz.




Figure 10. Time-domain response with center frequency changed

Example 2: Tuning to a “golden” filter

The second example uses a filter that has eight poles with seven tunable inter-stage coupling structures, along with input and output coupling. In the discussion that follows, we use a "golden" filter that was tuned by an experienced engineer to obtain the desired frequency response and return loss. A second, untuned test filter, shown in Figure 11, was used as a test example. Figure 12 shows the time-domain and frequency-domain plots of both filters. A four-parameter display mode is used to show both the S_{11} and S_{22} (input and output return loss) in both the time and frequency domains.




Figure 11. Eight-pole, seven-aperture filter used for Examples 2 and 3




Figure 12. The response of a “golden” filter (lighter trace) and an untuned filter of the same type (darker trace)

The test filter was pre-tuned by arranging the coupling screws (the long screws in the picture) to about the same height as the "golden" filter. Such pre-tuning is commonly done to get the coupling apertures closer to the correct value before beginning to tune, but it doesn't work for situations where a previously tuned filter is not available.

The first step in tuning this filter is to assume that the inter-stage coupling is close to correct, and adjust the resonators to optimally tune the filter without adjusting the coupling screws. The setup for this filter is a center frequency of 1220 MHz and a span of about 320 MHz. The filter bandwidth is about 80 MHz, so the time domain is initially set up from about -8 ns ($-2/\pi BW$) to about 70 ns ($((2N+1)/\pi BW)$). After the first tuning, -20 ns and 80 ns are determined to be a good choice for time settings.

Following example 1, each of the eight resonators are tuned, starting with the two outside resonators and continuing until the center resonators are tuned. Each is tuned by minimizing the response (making the deepest dip). Again we begin by first tuning the two outside resonators (numbers one and eight), looking at both S_{11} and S_{22} , then retuning them after the next inside resonators (two and seven) are tuned. After the third set of resonators are tuned (three and six) the second resonators (two and seven) are re-tuned. This continues one more time for the fourth and fifth resonators. After this initial tuning is complete, the filter exhibits a very nice frequency response (Figure 13), but does not match the desired response. Now it is time to tune the coupling structures.




Figure 13. The response of a “golden” filter (lighter trace) and a filter where only the resonators are tuned (darker trace)

To tune the coupling structures, the scale can be changed so that it is easy to see the peaks of the time-domain response. For this example, four-parameter display capability is used to show the time domain in full scale with a close-up view of the peaks. With this display it is easier to adjust both the peaks and the dips. To tune the coupling, start by tuning the coupling apertures that are closest to the input and output of the filter and work towards the center, to avoid masking effects from improperly tuned outer couplings. Turn the screw in to increase the coupling (reduce the peak). After each coupling screw is adjusted, readjust the resonators on each side to make the dip as low as possible, starting from the outside and working in. Figure 14 shows the result after the first pass of adjusting the coupling structures and resonators from the outside in.




Figure 14. The response of a “golden” filter (lighter trace) and another filter with both couplings and resonators tuned (darker trace)

This filter response is nearly identical to the template filter. The coupling (and hence return loss) is not symmetrical for input and output, but it is also not symmetrical for the "golden" filter used as a template. If the filters had no loss, the input and output match would be the same. The loss in the filter causes the input match to be different from the output match. It is possible to tune this filter to have exactly the same input and output match, but with a lossy filter, one match may be improved only at the expense of the other.

Also, note that the filter tuned in the time domain has better return loss than the "golden" filter, and that from the time-domain trace, we can see that the first resonator is not optimally tuned according to the time-domain tuning process, even though the filter has been tuned by an expert.

Example 3: Using simulated results for a template

Using a simulated filter response to create a template for tuning the filter is the basis for the final example of tuning. An ideal eight-pole Chebyshev filter is simulated, and any value can be chosen for bandwidth or ripple. For this filter, a wider bandwidth with larger ripple was chosen. We will attempt to tune the same filter used in example 2 to yield this new filter shape. Since the example filter does not have adjustable input and output coupling, there are limits on the filter shape that can be achieved. In this case, the bandwidth was fixed, and a return loss value that yields the same value for input coupling in the time domain as that of the example filter was chosen.

The frequency response of the simulation was downloaded into the network analyzer and used as a template. In the simulation, loss was added to the resonator structures to approximate the total loss of the real filter. This allows the S_{11} and S_{22} from the simulation to better match the actual time-domain response of the filter. The effects of loss are discussed in more detail in the next section.

Each coupling aperture and resonator is tuned to achieve the same time response as the simulated template, following the procedure described in Example 2. The last coupling structure is not tunable, but it is close enough to avoid distorting the overall response.

Figure 15 shows the result with the simulated trace, and the final tuned filter. The results are remarkably close, considering that the filter was tuned only in the time domain, and that the simulation used capacitively-coupled lumped elements, while the real filter had magnetically coupled distributed elements. Using this technique, virtually any filter shape that can be simulated can be used as a template for a real filter that can be easily and deterministically tuned, as long as the filter elements have the tuning flexibility. Even inexperienced tuners can follow this simple tuning technique because each coupling and resonator structure can be distinguished in the time domain.

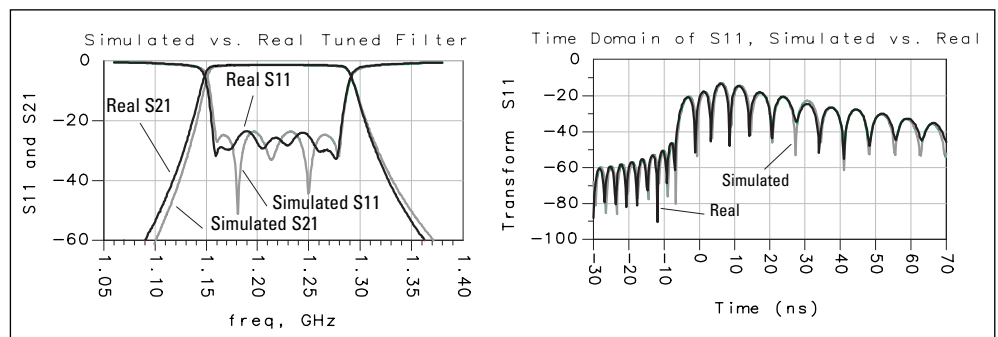


Figure 15. An example of a simulated filter and a real filter tuned to match the time-domain response.

Effects of loss in filters

Earlier, there was a caution about considering the effects of loss when using simulation to generate the time-domain trace. A lossy filter has peaks in the time-domain trace that are lower than those of a lossless filter, and the differences in the peak levels are greater for the apertures that are farther into the filter. Therefore, tuning a lossy filter to a template based on the simulation of a lossless filter will probably result in incorrect settings of the coupling factors.

Trying to set the coupling apertures in the lossy filter to match the template of a lossless filter requires increasing the peaks in the time-domain trace higher than the proper value, so the coupling must be reduced to get more reflection. Usually it will not be possible to match all of the peaks, especially the ones for the apertures that are farther into the filter, because as we observed earlier in Figure 5, decreasing one coupling factor will cause the corresponding peak to increase, but the following peaks will all decrease.

In the frequency domain, the result is that you may be able to achieve a similar return loss, but the filter will be narrower due to the higher reflection, as shown in Figure 16.

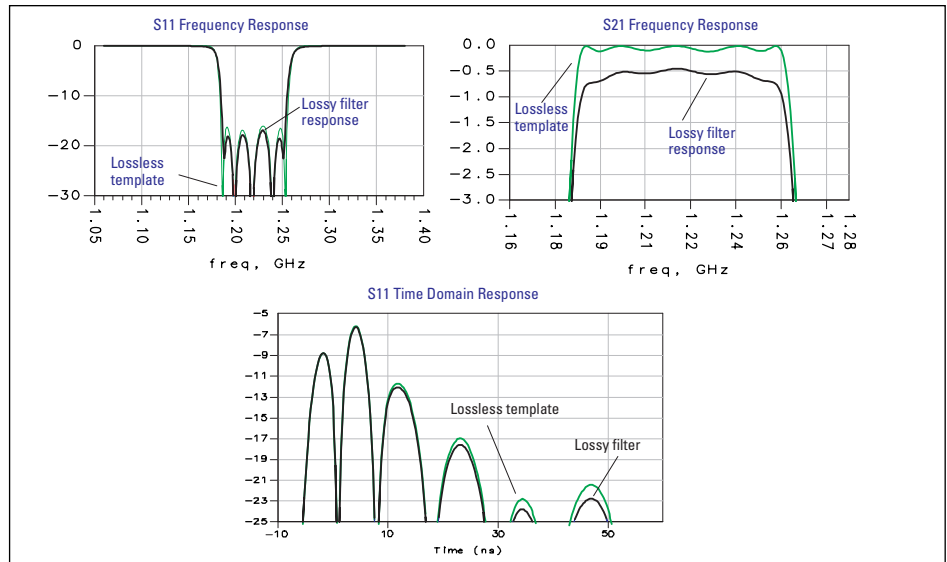


Figure 16. Comparison of lossless and lossy filters

For many cases, filter loss may be ignored, but for higher-order filters, it may be necessary to include the loss of each resonator in the model. Further, while many simulators allow loss to be applied to filter shapes, they do not distribute the loss throughout the filter. Thus, to properly account for loss, it may be necessary to create a filter structure using lossy resonators with discrete coupling in between.

To match a filter's return loss to a lossless filter simulation, it may be necessary to tune a lossy filter primarily from the S_{11} (input) side. The loss of the filter will cause the S_{22} time-domain response to differ from the S_{22} of a lossless simulated filter. Since the forward reflection and transmission (S_{11} and S_{21}) are more important in most cases, tuning from the S_{11} side will provide better results.

If a template for a lossless filter must be used, you may need to adjust the coupling apertures so they don't completely match the peaks; that is, allow them to be a little lower to account for the loss in the filter.

More complex filters

Cross-coupled filters

Finally, many filters are more complex than the traditional all-pole filters. Cavity-resonator filters often have "cross-coupling" that effectively adds one or more transmission zeros, similar to an elliptic-filter response. If these zeros, which create very narrow isolation regions in the transmission response, are close to the filter passband edges, they can distort the time-domain filter response so that it no longer shows a deep null associated with the resonator near the structure that creates the zero. In general, the resonators that are not cross-coupled can still be tuned using the nulling technique described earlier. But what about the cross-coupled resonators?

Some filters have transmission zeros that are symmetrical as shown in Figure 17; the response from the zeros can be seen on both sides of the passband. These filters can usually be tuned with the methods previously described. The symmetry of the zeros keeps the cross-coupled resonators at approximately the same frequency as the other resonators, so all of the resonators can be tuned close to their proper values by tuning for deep nulls in the time-domain response.

Some fine-tuning may be necessary, either by tuning in the frequency domain, or by using the techniques described in the next section.




Figure 17. Transmission and reflection responses of a filter with symmetrical transmission zeros

For filters that have asymmetrical zeros as shown in Figure 18, the resonators that are cross-coupled do not have the same frequency as the other resonators, so the dips in the time-domain response that correspond to these resonators will not be minimized when viewed with the network analyzer's center frequency set to the filter's center frequency. Tuning the resonators to a template may not yield the correct response, because there is more than one setting of the tuning screw that can yield the same amplitude response. Recall that when we were discussing the time-domain response of a simulated filter, we found that tuning the resonator either too high or too low will both cause the dip to rise from the minimum value. The setting is unique only when you are tuning for a null. However, we can modify the time-domain filter tuning technique to account for this.

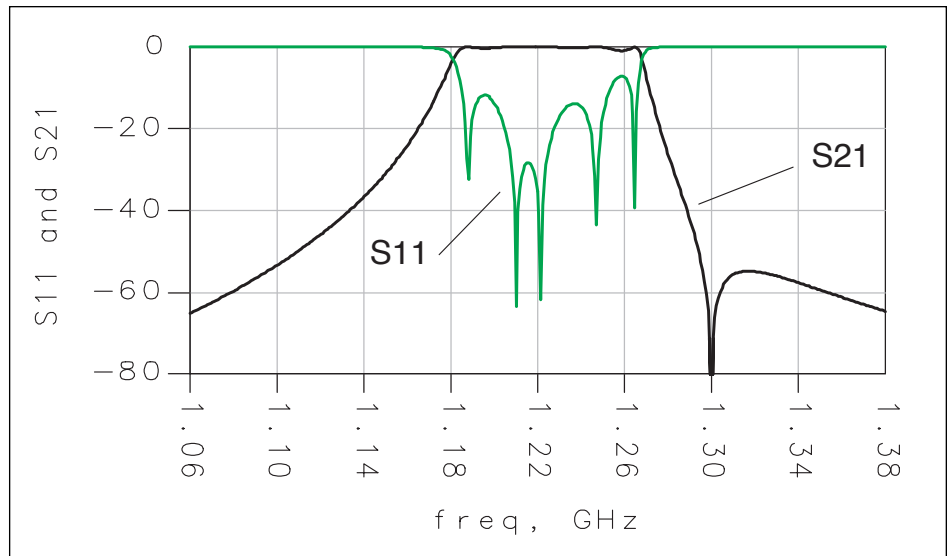


Figure 18. Transmission and reflection responses of a filter with asymmetrical transmission zeros

Recall that the all-pole filters we've been examining have resonators that are all tuned to the same frequency, with the effects of coupling included. We set the network analyzer's center frequency to that frequency, and when we look at the reflection response in the time domain, we get nulls corresponding to each resonator when that resonator is set to the analyzer's center frequency. For filters with asymmetrical responses, if we can determine the correct frequency of the cross-coupled resonators, we should be able to set the analyzer's center frequency to that new value, and tune the dip corresponding to the cross-coupled resonator to its minimum value to properly tune the resonator. Now the challenge is how to determine the correct frequency of the cross-coupled resonators.

One way is to calculate the correct frequency mathematically based on the filter design. Simulation tools can be very useful for doing this.

An alternative method is to derive the information empirically using a "golden" or template filter. You can set up the analyzer for a frequency sweep on one channel and the time-domain response on another channel. Identify the dip in the time-domain trace corresponding to the cross-coupled resonator. Watch the change in this dip as you slowly vary the center frequency of the analyzer's sweep. You should see the dip reach a minimum when the analyzer's center frequency is set to the correct frequency for that resonator. Use this information to set up a new instrument state for use in tuning that particular resonator. All of the resonators that are not cross-coupled will probably still need to be tuned with the analyzer's center frequency set to the filter's center frequency. However, depending on the coupling, a cross-coupled resonator may also pull the frequency of its adjacent resonators slightly off from the filter's center frequency, so you may need to find the correct frequencies for some of the neighboring resonators using this method as well.

In general, cross-coupling will not have much impact on tuning the coupling apertures, since the amount of cross-coupling tends to be light and has minimal effect on the peaks in the time-domain response corresponding to the coupling apertures.

For filters with cross-coupled resonators, the recommended order of tuning is:

1. Start out with the coupling screws pre-tuned (to match the physical settings of a "golden" filter), as described in Example 2.
2. Set the analyzer's center frequency to the filter's center frequency and tune all of the resonators to minimize the dips to get all of the resonators close to the proper settings, ignoring the error for the cross-coupled resonators for now.
3. Tune the coupling apertures to match the time-domain response to the template values.
4. Go back and fine-tune the cross-coupled resonators and any other resonators that need to be tuned to a frequency other than the filter's center frequency.

Duplexers

Tuning duplexers using the time domain can be a problem if the passbands are too close together. If the passbands are separated by at least one bandwidth, and you can set up the analyzer for a span of at least two times the bandwidth without seeing the other filter, you should be able to tune the duplexer using the techniques described in this application note. If the passbands are closer than one bandwidth apart, you will get interference from the response of the other filter, and you may not be able to clearly distinguish the responses due to individual resonators in the time domain. In this case, you may be able to partially tune the filter using time domain, but you will need some other method to complete the tuning.

Many duplexers have common elements (one or more resonators) in the antenna path that will form part of the response for both the Tx-Ant and the Ant-Rx paths. To tune these resonators, it may be necessary to set their frequencies to the center frequency **between** the Rx and Tx bands, instead of tuning them to the center frequency of either passband.

Both cross-coupled resonator filters and duplexers are more advanced topics that require more research. Further refinement of time-domain filter tuning techniques for dealing with such filters is currently under development.

Conclusion

While various techniques to simplify the process of filter tuning have been tried, until now, none have succeeded fully because coupled-resonator filters are inherently resistant to techniques that cannot account for characteristics such as coupling interaction. The method described in this application note goes a long way toward solving this problem. It allows coupling apertures to be tuned to match any filter shape within their tuning ranges, and resonators to be adjusted to provide a perfectly matched filter, with interaction immediately seen and corrected.

While a better understanding of some types of filters such as cross-coupled filters is needed, this technique already shows enough promise in allowing filters to be tuned easily that the current trend to automate filter tuning on the production line may not be needed. Alternately, this time-domain tuning may allow automation to become practical for the first time. It certainly makes it easier to train inexperienced filter tuners quickly. These attributes alone make the technique worthy of implementation and further study.

References

The Fourier Transform and Its Applications, second edition, Ronald N. Bracewell, McGraw-Hill, 1978.

Electronic Filter Design Handbook, (Chapter 5), Williams and Taylor, McGraw-Hill, 1988.

Filtering in the Time and Frequency Domains, Blinchikoff and Zverev, John Wiley & Sons, 1976.

“Simplify Filter Tuning in the Time Domain,” Joel Dunsmore, *Microwaves and RF*, March 1999, pp. 68-84

Summary: Hints for time-domain filter tuning

- Set the center frequency of the network analyzer equal to the desired center frequency for the filter.
- Set the frequency span to be 2 to 5 times the bandwidth of the filter.
- Use 201 points in the sweep for a good compromise between sweep speed and resolution.
- Measure S_{11} on one channel and S_{22} on the other channel. If desired, 4-parameter display can be used to view both the frequency- and time-domain responses at once. Viewing both domains while tuning may provide better insight for optimizing the filter's response.
- Select the bandpass time-domain transform.
- In the time domain, choose the start limit to be about one resonator's delay on the minus side; approximately $t = -(2/\pi BW)$. Choose a stop limit of about 2 to 3 times the full filter's delay; approximately $t = (2N+1)/(\pi BW)$, where N is the number of filter sections (resonators) and BW is the filter's 3 dB bandwidth in Hz.
- Use log magnitude format (dB), and set the reference position to 10 (top of the graticule) and the reference value to 0 dB.
- If the filter has tunable apertures, set the coupling screws approximately correct; for example, by adjusting them to the same physical height as those on a "golden" filter.
- Tune the resonators first, adjusting for deepest dips in the time-domain trace. Start with the resonators at the input and output sides and work towards the middle.
- Tuning one resonator may cause the previous resonator to become slightly untuned. In this case, go back and retune the previous resonator, then optimize the current resonator again.
- Tune the coupling apertures from the input and output sides first and work towards the middle. After adjusting each coupling screw, readjust the resonators on each side to make the dips as low as possible.
- If the filter has cross-coupled resonators, fine-tune the cross-coupled resonators to their correct frequencies.
- Repeat the tuning process at least once to fine-tune, or as needed to achieve desired response.

Appendix A: Understanding basic bandpass filter design

Many bandpass filters are designed by starting with a low-pass prototype that has the desired characteristics, such as passband ripple, input return loss, or stop-band rejection. The values for the prototype low-pass filter elements that are necessary to obtain these characteristics may be found in most filter design books (see References). This prototype low-pass filter can be transformed into a bandpass filter by changing the inductors and capacitors into LC circuits, with the center frequency of each LC circuit at the desired bandpass filter center frequency. Figure 19 shows an example of a prototype 3-element low-pass filter with the corresponding bandpass filter structure. The equations for calculating the values of the filter elements are also found in most filter design books.




Figure 19. 3-element prototype low-pass filter and corresponding bandpass filter

This design technique results in filters that approximately retain the desired filter shape. However, many narrowband (less than 10% bandwidth relative to the center frequency) bandpass filters designed with this method end up with LC elements that cannot be realized. For these narrowband filters, an alternative design technique has been developed that uses coupled resonators as the main elements. With this technique, each resonator is tuned to the filter's center frequency, with the effects of the adjacent coupling elements included. The resonator's center frequency is calculated by treating the adjacent coupling capacitors as though they were shorted to ground, so that the capacitances will be in parallel with the capacitance in the resonator. Figure 20 shows the bandpass filter from Figure 19 transformed into its equivalent coupled-resonator structure.

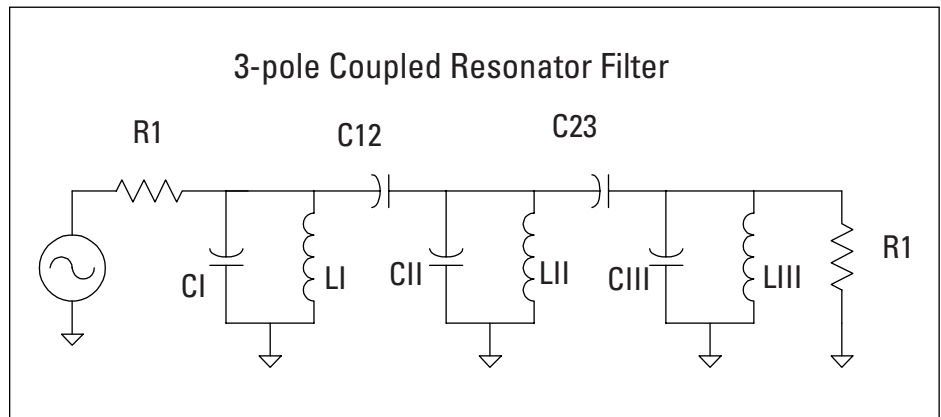


Figure 20. Equivalent 3-pole coupled resonator filter

A second aspect of the coupled-resonator design technique is that any changes in filter type and order affect only the coupling factor between the resonator structures. Thus the filter shape, bandwidth, ripple, and return loss depend only on the coupling between resonator sections, when the resonators have been properly tuned. These filters retain the shape factors of the prototype low-pass filter.

A circuit simulation program has been used to model the response for the mathematically simple three-pole Butterworth low-pass filter. Examining this filter's response using the time-domain transform shows that the characteristic nulls in the time-domain transform are indeed a consequence of the filter design. Repeating this simulation with a bandpass filter shows that the bandpass filter has exactly the same time-domain reflection impulse magnitude response as the low-pass prototype. Since the low-pass prototype's impulse response has the characteristic dips, and this filter has optimal circuit element values since it has no tunable components, we can conclude that the dips must also be present in a properly tuned bandpass filter.

The actual values of the elements used in the resonator are of little consequence, except that they affect the input and output impedances, so input and output coupling often include an impedance transformer to ensure a 50-ohm match.

These couplings can be capacitive, which is frequently the case in lumped-element filters, or inductive (sometimes called magnetic or B field coupling) which is often the case in cavity-tuned filters. In the latter, the coupling structure is an opening in the wall between sections that permits the circulating magnetic fields to couple. These openings or apertures can be made adjustable by narrowing the width of the opening, which reduces coupling, or adding a shorted tuning element, such as a machine screw, which increases coupling.

For many filters, the coupling factor changes only slowly with frequency, so that the center frequency of the filter can be changed over a substantial range without changing the basic shape of the filter. This is because the center frequency of the filter is determined only by tuning the center frequency of each resonator.

An intuitive way to think about this is that the coupling of other sections is what slightly pulls the center frequency of different resonators to move the poles about the necessary amount to produce the desired filter response. So, if a tuning technique can assure that each resonator is properly tuned, the total filter response will be correct.

In a simple cavity resonator filter, all resonators have the same center frequency, with the effects of the resonator coupling included in the calculation of resonator frequency. This frequency is also the center frequency of the filter. However, this is not true for filters with transmission zeroes, where cross-coupling between resonators will cause the cross-coupled resonators to be at a different center frequency than the other resonators. These cross-coupled resonators may pull the adjacent resonators slightly off from the center frequency of the filter as well. Thus, in tuning these filters, we need to determine the correct center frequency of the cross-coupled resonators (and possibly some of the adjacent resonators), and tune those resonators for that frequency, while tuning the remaining resonators to the filter's center frequency. A better understanding of using time-domain filter tuning for cross-coupled filters is still needed, and more research is being done on this topic.

Appendix B: Using time-domain in the network analyzer for filter tuning

To understand how to set up the network analyzer for time-domain filter-tuning measurements, it is helpful to review some basics of the time-domain transforms.

Normal time-domain reflectometers (TDRs) are inherently broadband and low-pass in nature. This means they are only useful for measuring DC-coupled circuits. They cannot be used for measuring bandpass filters, since the filters will appear to be almost totally reflective. However, a special mode of the network analyzer time-domain transform called bandpass mode can be used on band-limited devices.

In this mode, the center frequency of the frequency sweep is effectively translated to DC, and the inverse Fourier transform is applied from minus one-half of the frequency span to plus one-half of the span. This is important when looking at a bandpass filter with a frequency response that is the same as a low-pass filter response translated up in frequency to the center of the bandpass filter.

The time-domain transform represents the return loss as a function of length through the device under test. For time-domain transforms to be useful, they must have enough resolution to resolve the distinguishing characteristics of the network being measured. In general, the resolution of a transform is inversely proportional to the frequency span, although in bandpass mode the resolution is reduced by half because half the span is for negative frequencies and half for positive frequencies.

Looking at measurements of bandpass filters with a broad frequency sweep causes the same problem as in a low-pass TDR measurement: you see a near-total reflection at the input, and almost no other reflections. A normal network analyzer sweep of the bandpass filter, perhaps over two or three times the filter's bandwidth, would be a narrow sweep and was previously thought to have insufficient resolution to determine any characteristics of the filter. However, if the measurement is properly set up, the resolution limitation does not apply in measuring filters.

When a filter is examined in the time domain, each filter section has substantially more delay than its physical size would suggest. This is because the delay of a filter is inversely proportional to its bandwidth. The narrower the bandwidth, the longer the delay. For multiple-section filters, the transmission delay is approximately $N/\pi BW$, where BW is the bandwidth in Hz and N is the number of sections. Each section can be considered to add about $1/N$ of the delay. Thus the reflection delay of each section is about $2/\pi BW$, and the total delay for reflection is about $2N/\pi BW$ (twice as much as the transmission delay because the signal must go through the filter and back).

If the frequency bandwidth used to sweep the filters is at least two times the filter bandwidth, there will be sufficient resolution to discern the individual sections of the filter. The frequency span should not be too wide, or too much of the energy will be reflected, and tuning sensitivity will be reduced. Depending upon the filter, a frequency span of two to five times the filter bandwidth can be used.



For more information about Hewlett-Packard test and measurement products, applications, services, and for a current sales office listing, visit our web site, <http://www.hp.com/go/tmdir>. You can also contact one of the following centers and ask for a test and measurement sales representative.

United States:

Hewlett-Packard Company
Test and Measurement Call Center
P.O. Box 4026
Englewood, CO 80155-4026
1 800 452 4844

Canada:

Hewlett-Packard Canada Ltd.
5150 Spectrum Way
Mississauga, Ontario L4W 5G1
(905) 206 4725

Europe:

Hewlett-Packard
European Marketing Centre
P.O. Box 999
1180 AZ Amstelveen
The Netherlands
(31 20) 547 9900

Japan:

Hewlett-Packard Japan Ltd.
Measurement Assistance Center
9-1, Takakura-Cho, Hachioji-Shi,
Tokyo 192, Japan
Tel: (81) 426-56-7832
Fax: (81) 426-56-7840

Latin America:

Hewlett-Packard
Latin American Region Headquarters
5200 Blue Lagoon Drive, 9th Floor
Miami, Florida 33126, U.S.A.
(305) 267 4245/4220

Australia/New Zealand:

Hewlett-Packard Australia Ltd.
31-41 Joseph Street
Blackburn, Victoria 3130, Australia
1 800 629 485

Asia Pacific:

Hewlett-Packard Asia Pacific Ltd.
17-21/F Shell Tower, Times Square,
1 Matheson Street, Causeway Bay,
Hong Kong
Tel: (852) 2599 7777
Fax: (852) 2506 9285

**Data Subject to Change
Copyright © 1999
Hewlett-Packard Company
Printed in U.S.A. 6/99
5968-5328E**