5.2.6.1 Group Delay Measurements

A common figure of merit for many filters is group delay; in particular group-delay ripple in some passband, and on occasion, the absolute group delay. Group-delay measurements on filters by default are only an approximation of the of the actual group delay because the analog value must be estimated from a set of discrete measurements. The definition of group-delay is classically

$$\tau_G \triangleq -\frac{d\phi_{Rad}}{d\omega} = \frac{-1}{360}\frac{d\phi_{Deg}}{df} \tag{0.1}$$

However, phase measurements of a filter are by nature discrete and the analytic slope cannot be computed through differentiation, but must be computed through a discrete finite-difference computation:

$$\tau_{G\_meas} = -\frac{\Delta\phi_{Rad}}{\Delta\omega} = \frac{-1}{360}\frac{\Delta\phi_{Deg}}{\Delta f} \tag{0.2}$$

This discrete differentiation inevitably leads to confusion as the span of the Δf , called the aperture or delay-aperture, has a very strong influence on the overall measured group-delay response.

The first issue of concern for measuring filters is that the phase shift must be less than 180 degrees per point, and for some high-order filters, it may require a large number of points to avoid aliasing issues as described in section 5.1.5. But generally, most filters require sufficiently close frequency spacing so that under-sampling is not an issue. On the contrary, often a major issue in making group delay measurements on a filter is that the data points are so closely spaced that the group delay becomes inordinately noisy due to a very small Δf .

In most legacy VNAs, the group delay was computed by simply taking the change in phase at each point divined by the frequency step. This causes two issues: first, it produces a slight skew in the data as the actual delay on each point is offset by one-half of the span, if only 2 points are used. Secondly, because there are N-1 frequency segments, there are N-1 group delay computed points for an N point sweep. In most legacy analyzers, this is handled by making one of the delay points repeated, usually the first point. One can avoid the skew issue by always using an odd number of points for the delay or smoothing aperture.

If the frequency step was very small, even very slight trace noise on the phase trace would cause very large trace noise in the group delay trace. Figure 5.41 shows the group delay response of a filter with two different frequency spans in the upper window, and for various number of points in the lower window, where the reference position is offset by 1 division for each setting to clarify the resulting traces. For each measurement, a trace statistics function is shown for the middle 60 MHz of the group delay response, and the relative value of trace noise can be seen in the Std. Dev. result.

A wide span trace has less noise than a narrow span trace, roughly in proportion to the span. Traces with fewer points show less delay noise than greater point traces roughly in proportion of the number of points. In all cases, the trace-noise on the phase trace is the same, but the divisor in equation (0.2) changes. All of these results are attributes of the group delay aperture, and the fact that it changes with span and number points.

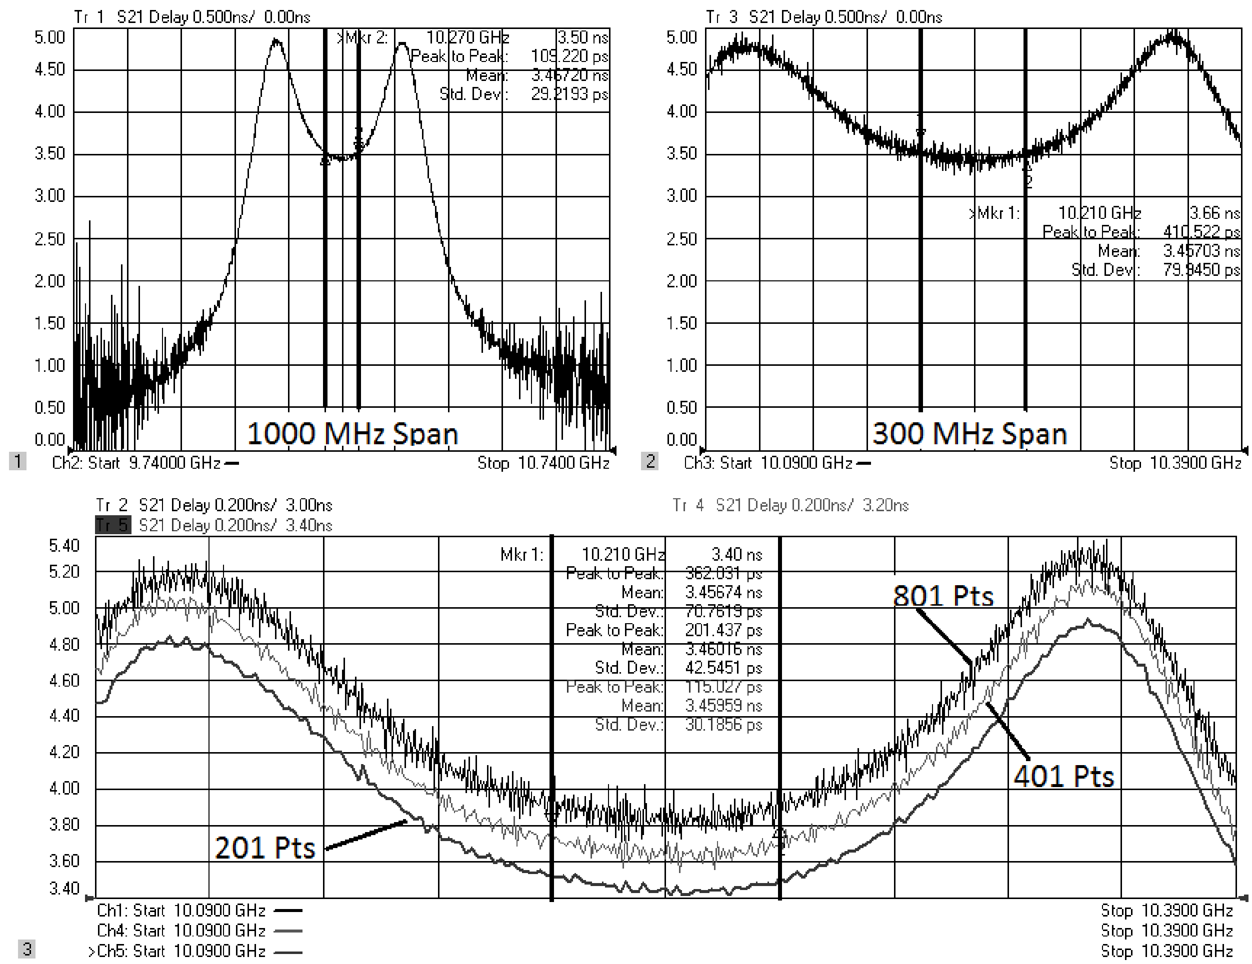

Figure 5.41: Group delay on a filter with various number-of-points, and various spans.

While the trace noise in the phase responses is essentially a constant level, the changing number of points changes the aperture with a corresponding increase in the trace noise of the group delay trace. One way to avoid this issue is to use less points on the group delay trace, but in many instances a large number of points is required to ensure that there are no issues in the amplitude response, and it would be convenient to have the group delay computed on the same trace. Most legacy VNAs incorporate a smoothing function that applies a moving-average smoothing on a trace, and this can be used to set a wider effective aperture on a group delay trace. The smoothing function operates on a moving average window where

$$Y_{n\_Smooth} = \frac{\left[Y(n-m)+Y(n-m+1)+...+Y(n)+...Y(n+m-1)+Y(n+m)\right]}{2m+1} \tag{0.3}$$

Where 2m+1 is the smoothing -aperture, in points. Smoothing is often given in percent of span, and this is converted to smoothing points by

$$m = \mathrm{int}\left[\frac{N \cdot (smoothing\ \ percent)}{2}\right] \tag{0.4}$$

Where N is the number of points.

When this is applied to most traces, each point forms an average of the surrounding points, and is similar to the video bandwidth function of a spectrum analyzer. In ***most*** cases, smoothing is an invalid way to reduce noise, as it can also remove important structure in the response. In the case of an amplitude response, smoothing can be used to eliminate mismatch ripple due to a poor calibration, but it can also hide the fact that a response has some excess ripple to due to the DUT response.

In group delay traces, however, a quirk of mathematics makes using smoothing ***identically equal*** to reducing the number of points around the target point, and thus any intermediate values have no effect on the delay; a simple example of applying the definition of group delay from equation (0.2) to the definition of smoothing from equation (0.4) for a 5 point smoothing (m=2) illustrates this point

$$D(n)_{Smo} = \frac{\left[\frac{\cancel{\varphi_{n-1}} - \varphi_{n-2}}{\Delta f} + \frac{\cancel{\varphi_{n}} - \cancel{\varphi_{n-1}}}{\Delta f} + \frac{\cancel{\varphi_{n+1}} - \cancel{\varphi_{n}}}{\Delta f} + \frac{\cancel{\varphi_{n+2}} - \cancel{\varphi_{n+1}}}{\Delta f} + \frac{\varphi_{n+3} - \cancel{\varphi_{n+2}}}{\Delta f}\right]}{360 \cdot 5}$$

$$D(n)_{Smo} = \frac{\varphi_{n+3} - \varphi_{n-2}}{360 \cdot 5\Delta f} \tag{0.5}$$

From which it is clear that none of the intermediate points contribute anything to the computation of the smoothed delay. Unlike smoothing in other traces, the values of the measurements of intermediate points have no consequence on the result, such that the value after smoothing is identical to the value that one would see if only the phase of the endpoints of the smoothing aperture were measured.

In more modern VNAs, the group delay aperture can be sent independently from the trace smoothing, so it is no longer necessary to turn smoothing on and off when changing the trace format from delay to other formats such as log magnitude. In some VNAs, the aperture can be set either as a point aperture, or a percent of span, or a fixed delta-frequency aperture. In these cases, it is necessary for the delay to be computed from interpolated phase points if the fixed delta frequency or percent of span does not lie exactly on measurement points. Using a fixed delta frequency is a very convenient way to specify aperture as the resulting group delay trace, and noise on the trace does not change with changes to frequency span or the number of points. . In Figure 5.42 the aperture of 3 MHz was used, which matches approximately the default aperture for the 201 point trace above. Looking at the Std. Dev. numbers, one sees that they are nearly identical for all traces. Thus specifying the aperture in terms of delta frequency yields the most consistent results.

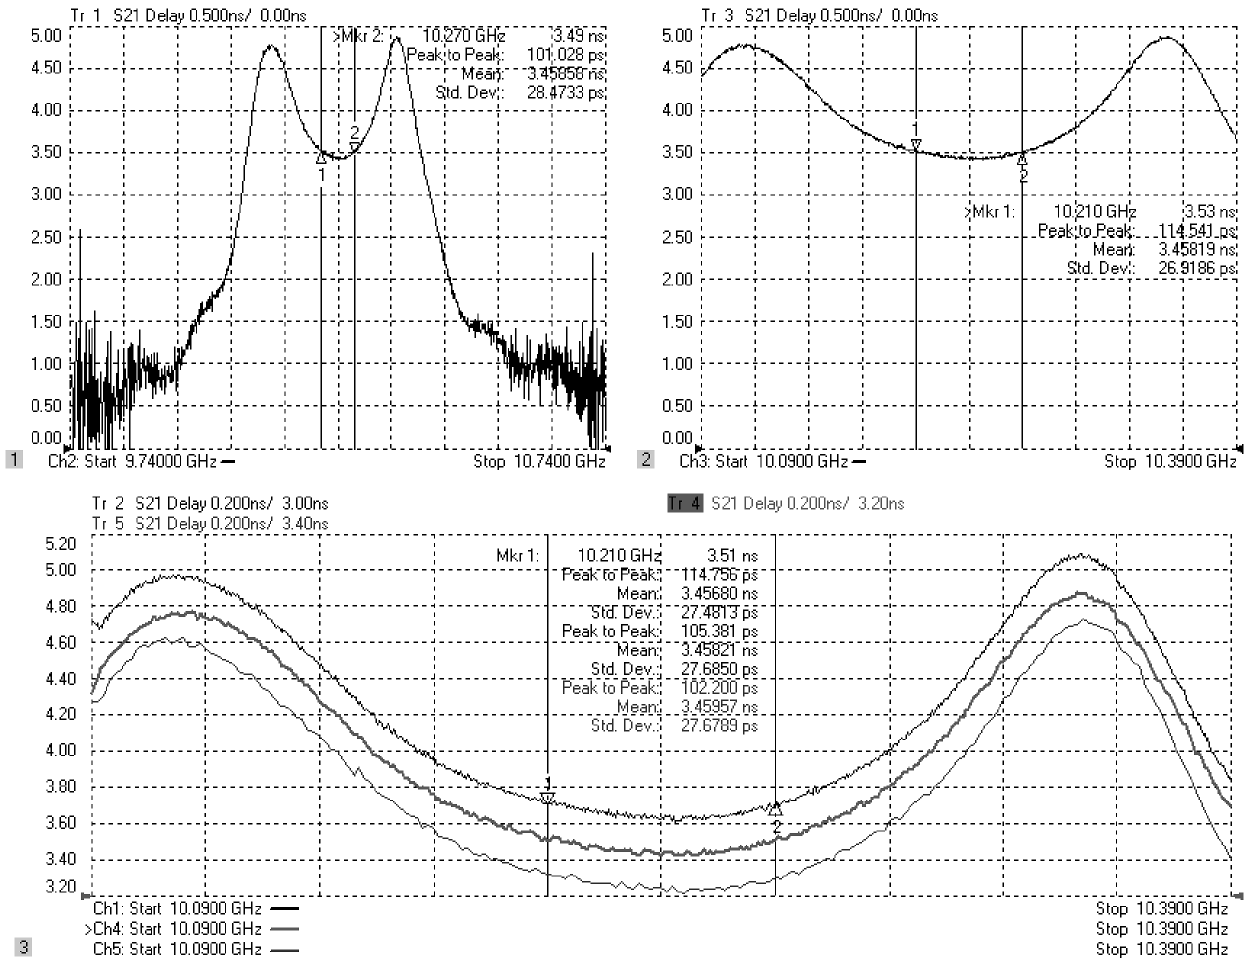

Figure 5.41 Group delay results from applying a fixed delay aperture to the previous figure's response.

In most filters, the absolute value of the group delay is not important; only the group delay ripple matters to the communications channel. Group delay ripple can cause distortion in modulated measurements. However, in some classes of filters, the absolute group delay is critically important, and these are called, not surprisingly, “delay” filters. These types of filters are often found in feed-forward amplifiers to provide a fixed delay to a signal applied to an error amplifier that is eventually combined at the output and used to cancel the distortion of the main power amplifier. These filters must be tuned for a precise delay, as well as other characteristics.