Implementation of Real-Time Spectrum Analysis

White Paper

Products:

| R&S®FSVR

This White Paper describes the implementation of the R&S FSVR's real-time capabilities. It shows fields of application as well as the technical implementation.

Implementation of Real-Time Spectrum Analysis

Dr. Florian Ramian

January 2011 – 1EF77

Table of Contents

1 Real-Time Analysis

1.1 What Real-Time stands for in the R&S FSVR

The measurement speed available in today's spectrum analyzers is the result of a long evolution. Traditional spectrum analyzers, like the R&S FSE, measured frequency spectra by mixing the input signal to a fixed intermediate frequency (IF) using a swept local oscillator. The signal was down converted in several mixing stages, and finally it passed the analog resolution filter, which determined the frequency resolution available at each frequency point on the screen. The measurement time was dependent on the settling time of the resolution filter and the time the first local oscillator needed to return from its end frequency to its starting point, the so-called re-trace time.
With increasing computing power, the next analyzer generation (R&S FSP, R&S FSU) was equipped with FFT filters for narrow bandwidths. Multiple narrowband FFTs were concatenated to a trace representing the selected frequency span. As the computing time for the FFTs was small compared to the settling time for narrow RBW filters, the FFT method provided a great speed advantage over the traditional sweep method.
The latest spectrum analyzer generation, the R&S FSV, makes excessive use of the FFT method for narrow resolution bandwidths. In addition, it introduces complex digital filters for wideband resolution filters, providing a factor of 25 in sweep speed increase, compared to earlier analog implementations.
The measurement speed has increased dramatically from 20 sweeps/s on the R&S FSE to more than 1000 sweeps/s on the R&S FSV. But one property has survived all evolution steps: even the R&S FSV does not detect signals between the end of one sweep and the start of the next one. This gap in data acquisition, the so-called "blind time", has decreased with each new spectrum analyzer generation, but it is still present.

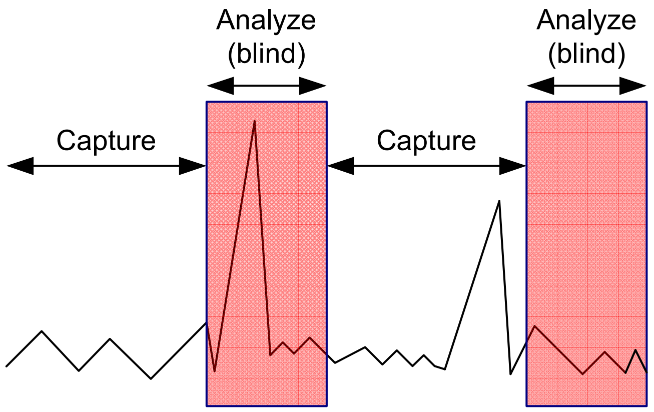

Figure 1: Sequential capture and analysis as used in e.g. FFT analyzers

Measuring signals in real time means: do not loose any signal. But how can we get rid of the blind times?
The answer comes with today's wideband, high resolution analog to digital converters (ADCs). The 16 bit ADCs available today allow capturing wide frequency ranges (e.g. 40 MHz) in a single shot with sufficient dynamic range without having to move the local oscillator (LO). Combining these wideband ADCs with fast FFT algorithms implemented in dedicated hardware (e.g. an FPGA) is the basis for the design of a real-time spectrum analyzer like the R&S FSVR.

The keys to a real-time spectrum analyzer are:

- Parallel sampling and FFT calculation:
  The data acquisition continues while the FFTs are performed.
- Fast processing of FFT algorithms:
  The computation speed must be high enough to avoid that “stacks” of unprocessed data are being built up. Slow FFT computation will result in an overflow of the capture memory and a subsequent data loss (= a new blind time).

Figure 2 shows the parallelized capture and analysis which avoids blind times. Clearly, nothing remains undetected with a real-time spectrum analyzer.

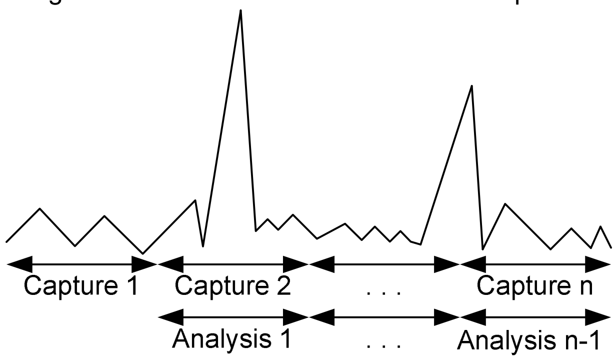

Figure 2: Parallel capture and analysis as used in real-time analyzers

1.2 Real-Time Applications

What are typical applications for real-time measurements? All measurements on short or seldom signals or signal variations, where you do not want to miss even one event.
A typical application is the analysis of a given frequency band. Assume a DUT that has a frequency hopping algorithm implemented. To analyze whether the DUT switches over the frequencies in the desired order, not a single step must be lost.
A transient event, such as the tuning of a VCO to its target frequency is another typical application for a real-time analyzer. The analyzer captures the entire tuning process without any gaps and records even the shortest glitches in frequency and level.
No matter what signals you are looking for, in most cases it is important to have a trigger possibility that allows triggering on the specific signal change of interest. A so called frequency mask trigger (FMT) in the R&S FSVR allows triggering on any spectral shape that can be displayed by the analyzer. A typical application is the analysis of a 2.4 GHz receiver. Besides the wanted signal of the system under investigation, many other signals can be found in this ISM band. To analyze the influence of disturbing signals on the system under investigation, the FMT will stop data capturing as soon as the frequency mask is violated. Without going into details, it becomes clear from Figure 3 that the persistence spectrum plot on the right hand side shows details about how a signal changes over time, whereas the Max Hold trace of a spectrum analyzer does not. Clearly by not loosing any information, the R&S FSVR is able to give precise information of a time variant signal, such as e.g. signal probability.

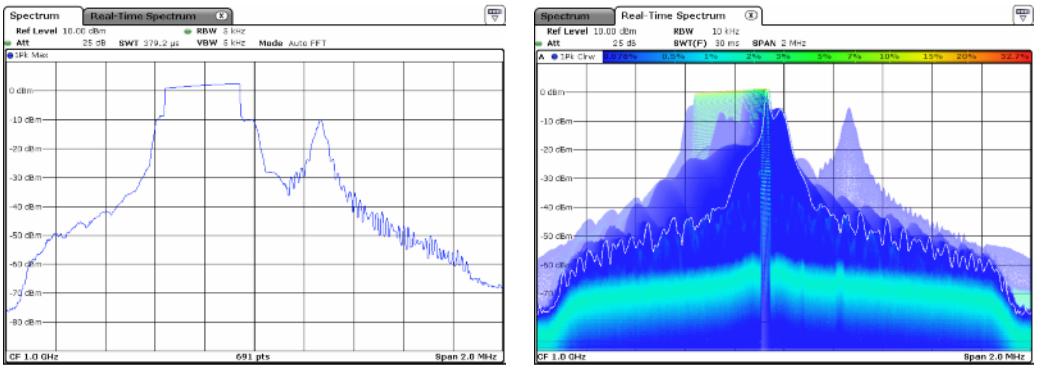

Figure 3: Comparison of a Max Hold spectrum analyzer trace and a persistence spectrum trace

The following chapters will explain the mechanisms behind data capturing without blind times and triggering on frequency masks will be explained in more detail.

2 Real-Time Implementation in the R&S FSVR

The R&S FSVR RF frontend is based on the R&S FSV signal- and spectrum analyzer. This means that the RF performance of both analyzers is almost identical. As the R&S FSVR is based on a conventional signal- and spectrum analyzer, it provides also full spectrum analyzer functionality to the user.

The core of the real-time analysis is the digital backend. As already stated earlier, the critical point behind real-time analysis is to run data acquisition and data processing in parallel. To achieve this, the digital backend of the R&S FSVR is equipped with a chain of powerful ASICs and FPGAs in combination with a large memory for captured data. This combination allows the instrument to process the data in several stages in a pipeline architecture. The last stage of the pipeline is the CPU, which reads the pre-processed data, applies the necessary scaling information and displays the resulting curve on the screen.

All real-time display modes and the frequency mask trigger can run in parallel on the R&S FSVR. This means that all real-time results can be displayed in multiple diagrams at a time, and the frequency mask trigger can be used in addition to capture rare events. This flexibility is a unique feature of the R&S FSVR.

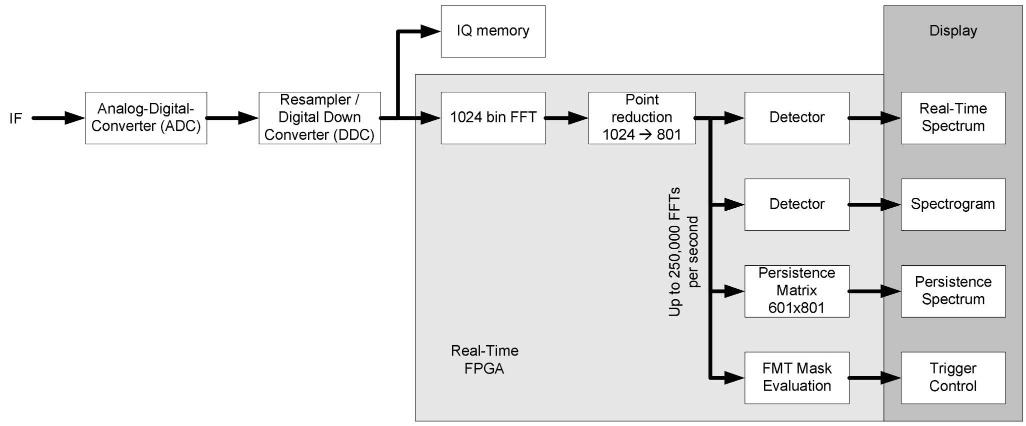

Figure 4: Signal flow chart of the digital real-time part of the R&S FSVR

Figure 4 shows the signal flow diagram from the A/D converter (ADC) to the display unit. The ADC is operated at a constant sampling rate of 128 MHz. The ADC streams raw data into the resampler and digital down-converter, which convert the input signal into a digital baseband, whose bandwidth is equal to the selected frequency span, and whose sampling rate fulfills the Nyquist criterion for this bandwidth. The ratio between complex baseband sample rate and selected frequency span is 1.2, meaning that e.g. a 40 MHz span is sampled with 50 complex MSamples per second. For smaller bandwidths, the sampling rate is automatically reduced.
The sampling rate determines the number of samples which are available for analysis. After resampling, the data stream is transformed into the frequency domain by means of an FFT. Each FFT consists of 1024 so called bins or data points. The FPGA running the FFT algorithms delivers up to 250,000 FFTs per second.
In parallel the resampled baseband data is written into the I/Q memory for additional offline (non real-time) post-processing, like e.g. zooming into a captured region or reading out the I/Q samples via LAN or GPIB. Note that the I/Q memory is implemented as a circular buffer which means that once the memory is full, the oldest samples will be overwritten.

FFT Update Rate
Consecutive FFTs are the raw spectral data that are used for all spectral displays. For a high resolution on the time axis it is a prerequisite to have the FFT update rate as high as possible. This is the precondition for implementing a frequency mask trigger, which must react extremely fast on signal changes in the frequency spectrum. With an update rate of up to 250,000 per second, the R&S FSVR calculates an FFT every 4 µs. It uses a fixed length FFT algorithm, which provides a higher speed compared to variable length FFT algorithms. The FFT length in the R&S FSVR is 1024 bins. For further processing, the FFT results are shortened to 801 usable points. The analyzer uses exclusively the 801 point FFT result for all processing steps after the FFT algorithm.

FFT Overlapping
Handling FFT results of short events (short compared to the FFT capture time) is a challenge, which must be handled properly by a real-time spectrum analyzer to avoid level errors.

To show the critical situation, let's assume that the capture time frames for two subsequent FFTs do not overlap. The energy of a short pulse, which hits the border of the two capture time frames as shown in Figure 5, will be distributed among the results of both neighboring FFTs. As a result, each of the FFT results exhibits a lower power level compared to the true power of the time domain pulse.

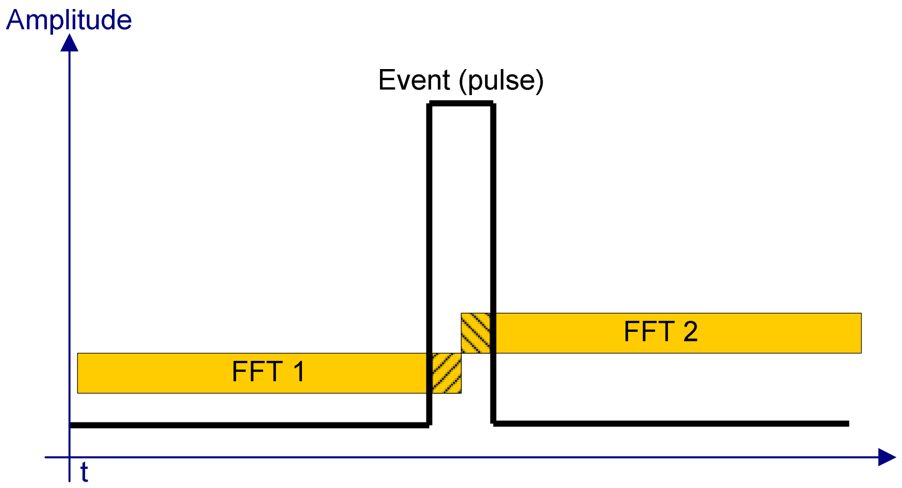

Figure 5: Pulse captured by two consecutive FFT time frames without overlapping

The R&S FSVR uses a technique called FFT overlapping to avoid this situation. Overlapping "reuses" samples that were already used to calculate the preceding FFT result. Figure 6 shows a pulsed signal that is captured by several overlapping FFT time frames.

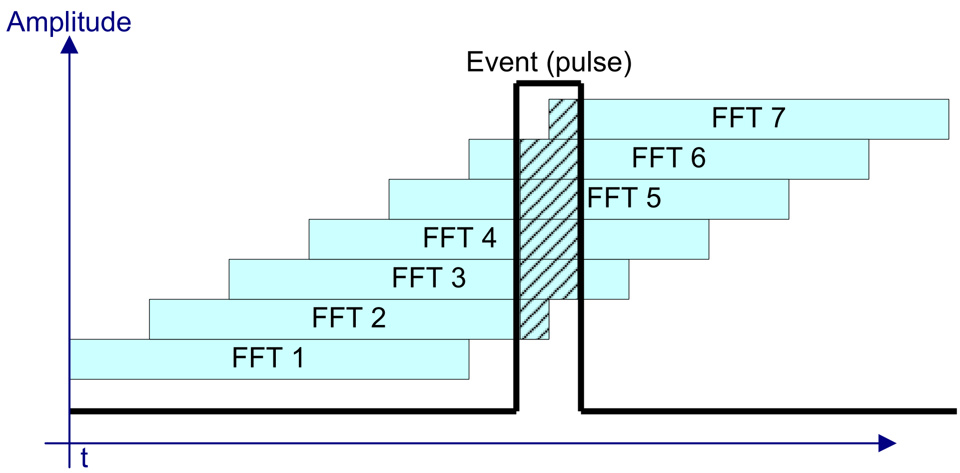

Figure 6: Pulse captured with several consecutive overlapping FFT time frames

In this example there are several FFTs that capture the entire pulse and not only fractions of it. The overlap factor describes the ratio of reused samples to the total number of samples. In the case of the R&S FSVR, an overlap factor of at least 80 % is used. In terms of samples, the R&S FSVR reuses at least 800 samples for the consecutive FFT.

Finally, a more detailed view on FFT techniques reveals another issue that requires an adequate overlapping ratio. An FFT analyzer usually applies a non-rectangular windowing function to the captured I/Q data before calculating the FFT. Clearly, without actively applying a window, the device uses a rectangular window function on the time domain samples, as it cuts them out of a real signal stream. Non-rectangular windows such as Blackman-Harris, Hanning, etc. outplay rectangular ones in the frequency domain, as they produce less side-lobes than the sin(x)/x shaping of rectangular time domain windows. The drawback is the weighting of time domain samples at the edges of the window. Figure 7 shows 3 FFT time frames that apply different weighting to the pulse. Clearly, a high overlapping ratio is suitable to handle the drawbacks of FFT analysis and at the same time make use of the advantages the FFT technique provides.
With an overlap ratio of 80 % or higher, level errors caused by the FFT can be neglected on the R&S FSVR.

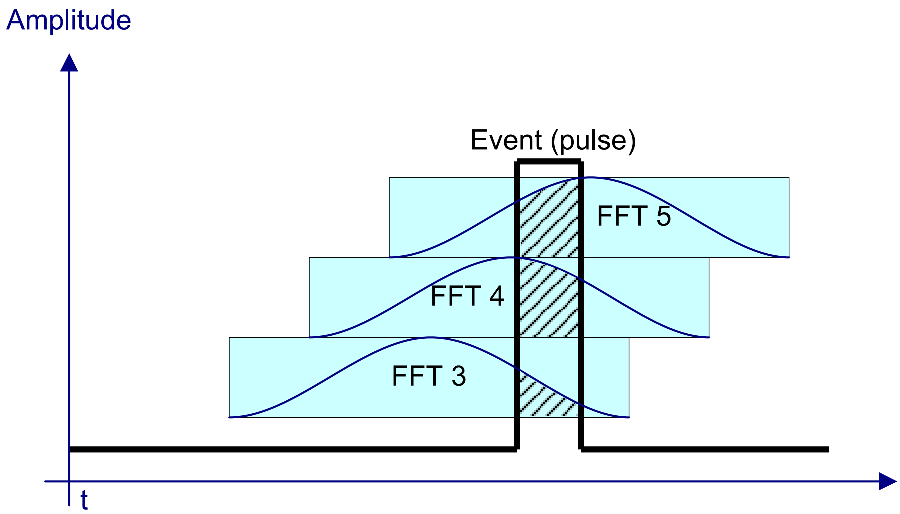

Figure 7: Overlapping compensates effects resulting from windowing function

Time Resolution of FFT Results

It is important to keep in mind that an FFT result is not the spectral representation of a single point in time, but the spectral representation of a certain time frame. This is another fundamental property of the FFT technique.
A side effect of this property is that consecutive events may be displayed in the same FFT result, similar to photograph that depicts everything that has happened within the exposure time. The R&S FSVR offers a high FFT update rate of up to 250,000 FFTs. Taking the overlap ratio into account, the effective exposure time for the R&S FSVR is roughly 20 µs.
The following example illustrates this effect. A CW signal changes in frequency. In between the frequency change from frequency 2 (f2) to frequency 1 (f1), no RF signal is present for 10 µs (see timing diagram in Figure 8).

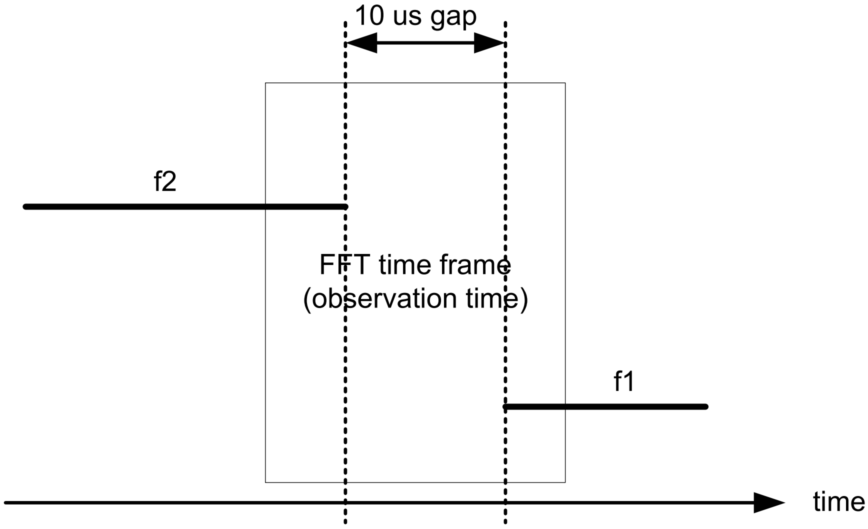

Figure 8: Timing diagram of the FFT observation time example

Without having the above principle in mind, a user might expect FFT results showing nothing but noise components. A user with knowledge about FFT processes knows what to expect: consecutive FFT results show the spectral component for f2 at first. During the 10 µs gap without a signal, the FFT result may show a spectral component of f2 at lower level as well as a spectral component of f1 with lower level. As the time interval without a signal is smaller than the above mentioned 20 µs, there won't be an FFT result showing noise components only. The spectrogram in Figure 9 shows the changing signal vs. time. The second spectrum trace from bottom in Figure 9 clearly shows the effect of the FFT time frame, i.e. all events that appear within the FFT length appear within the same FFT result, giving the impression that both frequencies were active at the same time, but with reduced power.

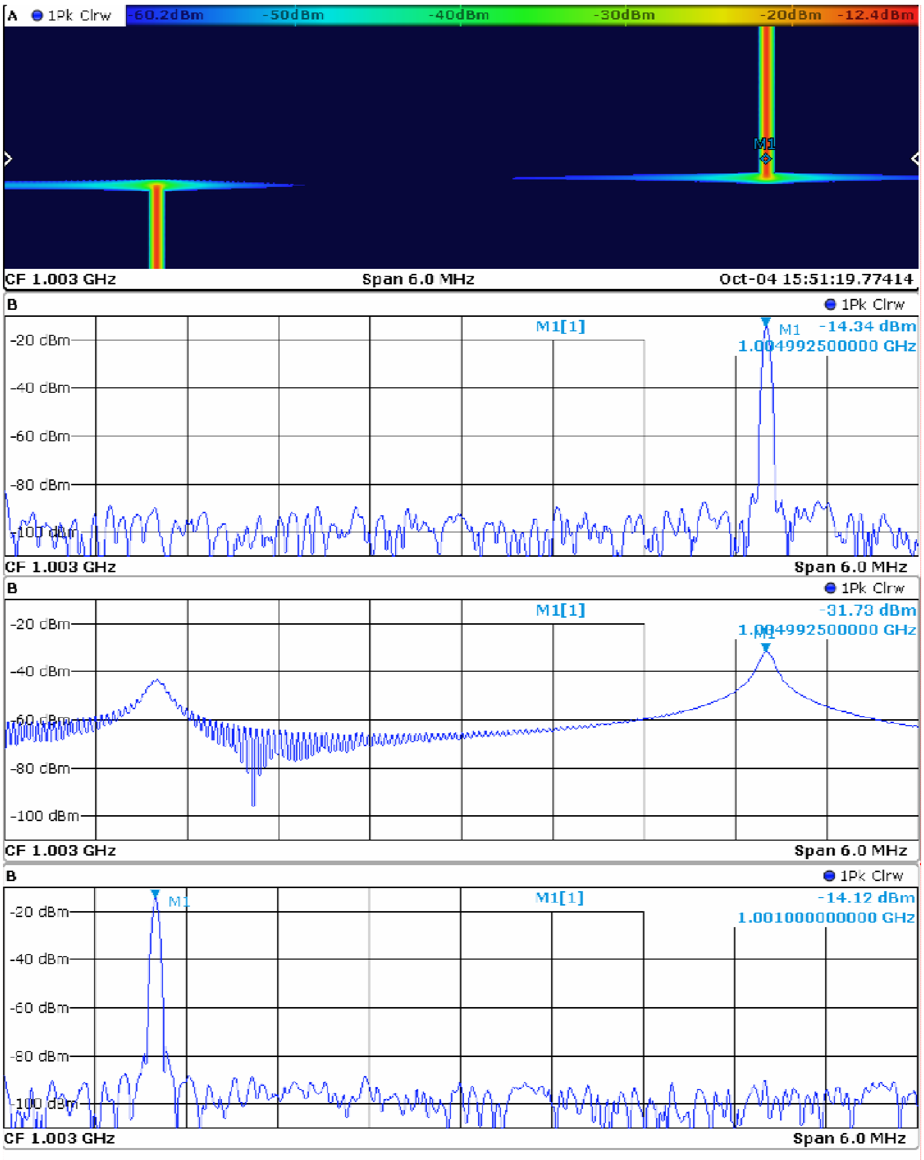

Figure 9: Spectrogram showing a frequency hop. From top to bottom, the spectrum traces show frequency f2, both frequencies, and frequency f1. The trace with both frequencies results from the FFT time frame being longer than the gap between the signals.

3 Triggering on Real-Time Spectra

This section focuses on a trigger mechanism which is only available with real-time spectrum analyzers: the frequency mask trigger (FMT). It is a reliable and powerful tool that helps the user to capture exactly the data needed for a quick analysis. The FMT is available with all real-time display modes as it is evaluated in parallel to persistence spectrum and spectrogram calculations (see Figure 4).

3.1 Frequency Mask Trigger

One way to analyze rare events in a given frequency range is to capture real-time data over a very long time. This method requires large amounts of fast memory. As a consequence post-processing the bulk of stored data to find the event may be extremely time consuming.
Another way is to trigger on the event in the frequency spectrum and to acquire exactly the data of interest. This method reduces the necessary memory size dramatically, and in addition keeps the time to spot the event of interest in the acquired memory low. The question is: how can the analyzer trigger on events which show up in a certain frequency range only now and then?
The answer is the Frequency Mask Trigger (FMT). Speaking graphically, the FMT is a mask in the frequency domain, which is checked with every calculated FFT. In case of the R&S FSVR this happens up to 250,000 times per second. Taking the overlap factor of 80% into account this allows to resolve events at intervals down to 12 µs.
The frequency mask can consist of up to 801 interpolation points and may have any shape.
The R&S FSVR offers 4 scenarios for triggering the data capture. It can start or stop data acquisition if

- the signal enters the mask area (Entering)
- the signal leaves the mask area (Leaving)
- the signal returns from outside the mask, i.e. it was in the mask area, left it and re-entered it (Outside)
- the signal returns from inside the mask area, i.e. it was outside the mask area, entered it and left it afterwards (Inside).

All of the above criteria apply to a configurable lower limit line as well as to an upper limit line. In addition, the criteria can also be applied to both lines (lower and upper) at the same time.
The FMT can be selected as a trigger source for all displays in real-time operation. As it is evaluated in parallel to the selected display modes, there is no influence on the real-time capabilities of the R&S FSVR.
The FMT is a trigger source which exceeds the capabilities of standard spectrum analyzers. To allow other instruments in a test system to make use of it, the R&S FSVR provides a special port (Trigger Out) as part of its option Additional Interfaces. The trigger out port provides a trigger pulse with a pulse width of 1 µs and a level of 5 V every time the FMT triggers the R&S FSVR. This trigger pulse may be provided to a system setup as an external trigger source.

3.1.1 Setting up an FMT trigger

A typical RF frequency band with a lot of interfering signals is the 2.4 GHz ISM band. Besides Bluetooth and WLAN, a variety of other services operate in that band. For this example, a Bluetooth receiver is assumed. The receiver looses its link to the corresponding transmitter in a lab environment, as the example Bluetooth link uses a single channel only. To analyze the interferer that leads to a disturbed Bluetooth link, an FMT is set up around the known Bluetooth signal. The trigger condition for the assumed example is:

- stop data acquisition if a significant amount of power is measured next to the Bluetooth channel.

This condition will trigger on all frequency hopping signals that cross the active channel and may cause the loss of connection.

A trigger mask that fulfils this requirement can be easily set up with the FMT mask editor (Figure 10). It is equipped with a live update of the signal as well as with an automated mask generator, making it very easy and intuitive to create the necessary mask.

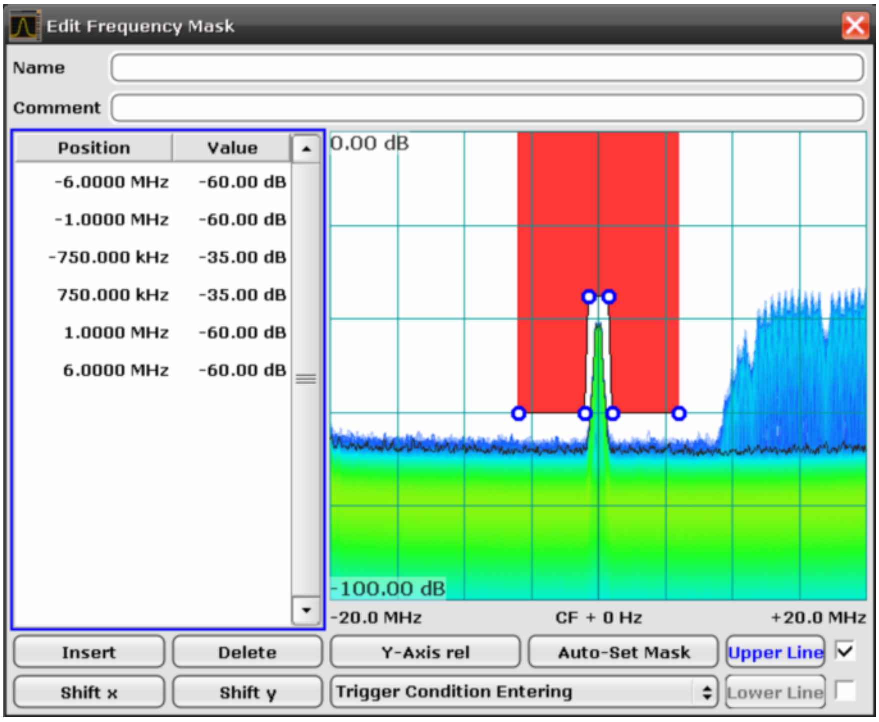

Figure 10: FMT dialog box

To indicate an active FMT, the trigger mask appears in the current persistence or real-time spectrum display as a red background mask (see Figure 11). Make sure that the R&S FSVR is in Run Continuous mode. To make sure that the displayed data contains all necessary information for your analysis, adjust the pre- and post-trigger time

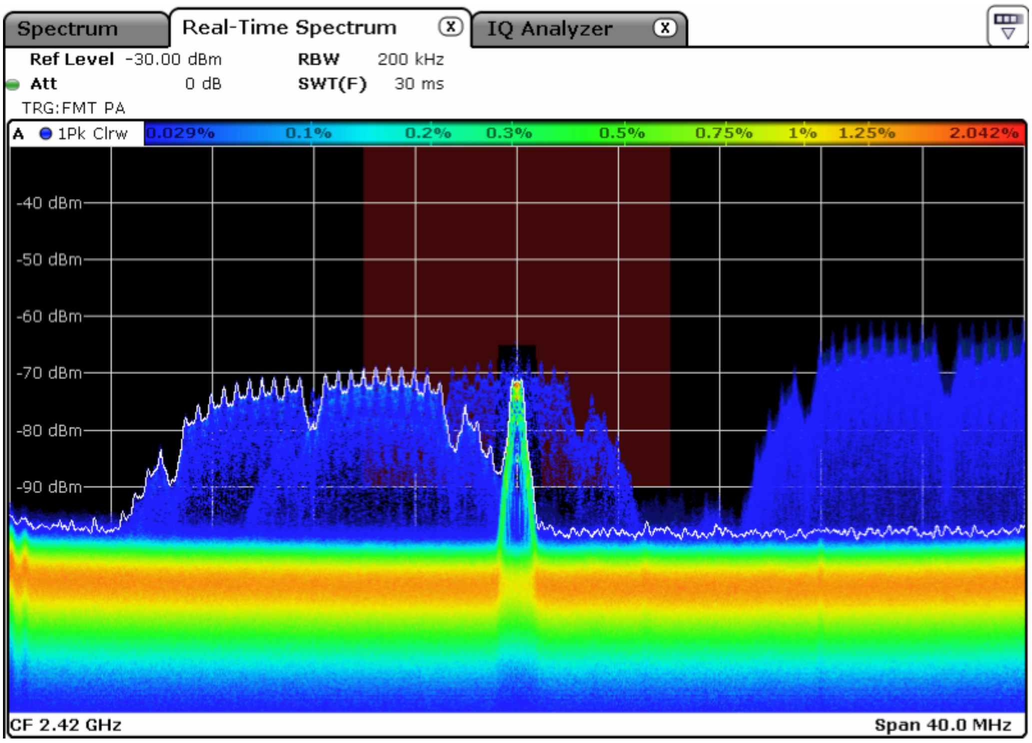

Figure 11: Active FMT triggered by frequency change of a WLAN signal

Once the trigger condition appears, the R&S FSVR will stop data acquisition after the trigger event plus the set post-trigger time. Please note: The time period recorded before the trigger event may be shorter than the specified pre-trigger time. The FMT is a real-time trigger, and the first trigger event it recognizes defines the moment when the post-trigger time starts – no matter whether the pre-trigger time has expired or not. With a real-time analyzer, you should not miss any single event.

Please note:
The FMT works with all display modes, which means it can be used in parallel to any combination of real-time diagrams.

3.1.2 Technical background

Basically the frequency mask trigger (FMT) is an extended limit line check: the FMT mask is compared to every FFT spectrum calculated by the real-time hardware.

The R&S FSVR performs this mask check up to 250,000 times per second according to the FFT update rate. To ensure a real-time trigger, i.e. a given reaction time, the FMT is evaluated by the real-time hardware.
Figure 12 shows the element wise comparison of a real-time FFT with an FMT mask. The FFT-result is subtracted from the FMT-Mask value. If one result is negative, the R&S FSVR triggers.

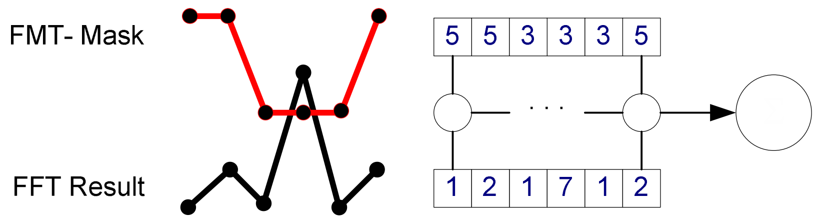

Figure 12: Element-wise comparison of FMT mask with current FFT result

Extended limit check means that the FMT can link a complex condition to the limit line violation, such as entering, leaving, inside (enter – leave – enter the marked region), and outside (leave – enter – leave the marked region).
As already mentioned, the FMT mask may contain up to 801 points, but may also be as short as 2 points. Shorter FMT definitions will be extended to 801 points by interpolation within the firmware. The FMT trigger therefore always compares 801 FFT points to 801 FMT mask definition points. If the mask is violated at a single point, the FMT will trigger.

Evaluation of the FMT equals the comparison of power levels. As mentioned in section 2, power levels in the displayed FFT are only comparable to the time domain power level, if the signal, e.g. the pulse fills an entire FFT. The spectral power level of a short pulse depends on the ratio event duration to FFT length. With a pulsed signal, where each pulse rises to a level of e.g. 0 dBm, the minimum pulse duration for a resulting spectral component reaching also 0 dBm is 24 µs. This figure is derived from the FFT length plus an additional 200 samples for the next overlapping FFT. If a pulse lasts 24 µs or longer, there is at least one FFT fully located within the pulse. Within the observation time of that FFT, the (unmodulated) pulse is equivalent to a continuous wave signal, as the edges of the pulse are not located within the FFT time frame. The result is a spectral component that reaches the same level as the pulse has in time domain. For shorter pulses, the so called pulse desensitization describes the dependency between time domain power level and the power level of the main spectral component. For more details on pulse measurement, see Application Note 1EF48.
In order to get a reliable FMT trigger with very short events, it is preferable to set the mask limit levels lower than the expected spectral power levels.

4 Display Modes for Real-Time Signals

Modern conventional spectrum analyzers offer the possibility to capture I/Q data. I/Q data capturing itself is real-time, meaning no information is lost. This statement is valid, as long as the I/Q memory of the analyzer is sufficient to cover the observation time. The stored I/Q data is post-processed. During post-processing no new information can be captured. Even analyzers that provide digital streaming interfaces, such as the R&S FSV and R&S FSQ require post-processing, as they have no means to process data in real-time.
The R&S FSVR does not only process data in real-time, but it also offers several display modes that help the user to analyze the data as it is displayed. The human eye has a limited capability of detecting changes – therefore real-time displays visualize the time axis, i.e. the changes of as signal over time. Display modes with information on past and present spectra at the same time allow a quick analysis of changes for human eyes.

4.1 Spectrogram

The spectrogram is a way of displaying multiple consecutive spectra over time. The power, or more exactly the power level, which is usually displayed over frequency is displayed over frequency and time. Graphically, time and frequency represent the vertical and horizontal axes of the display plane. Each coordinate (frequency f, time t) of the plane is filled with a color representing the level for the respective frequency and time.

At the beginning of a measurement, the plane is empty. As the measurement advances, the graph is filled line by line from top to bottom. Lines in the spectrogram are called frames, as each frame represents one spectrum that contains several FFTs.

As the graph fills from top to bottom, the latest spectrum is always the topmost line, whereas older FFTs move towards the bottom.

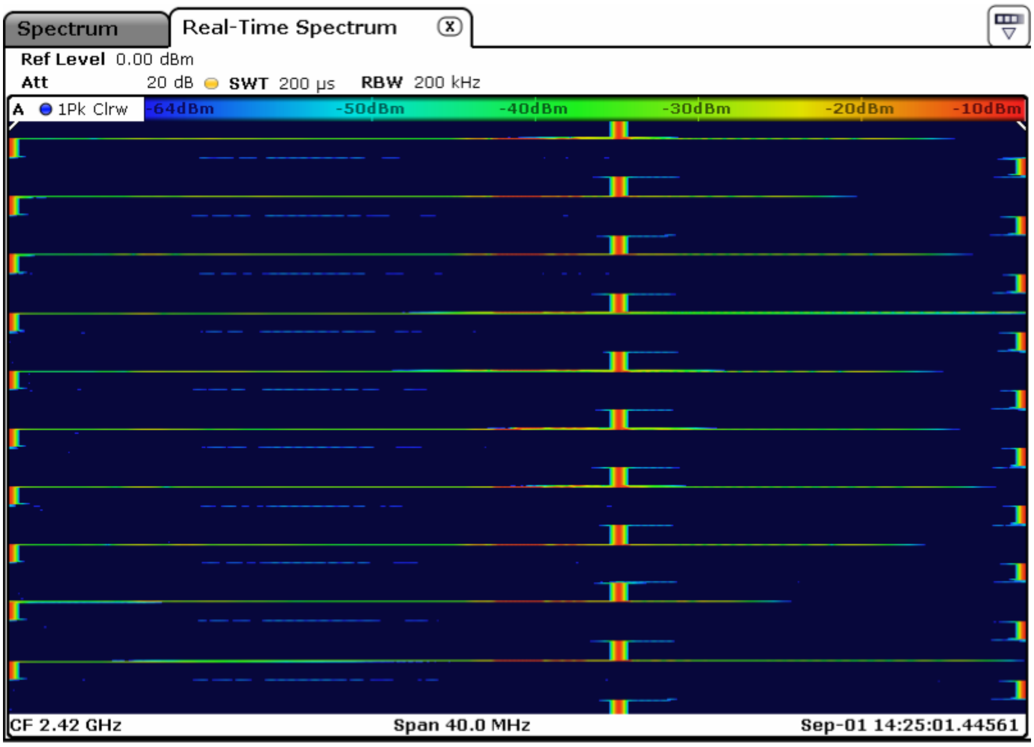

Figure 13: Frequency hopper exhibiting a transition with significant RF level from lowest to intermediate frequency

The spectrogram is a powerful tool to analyze time variant spectra. Typical applications are the transient oscillation of a VCO and the analysis of frequency hopping signals.

Figure 13 shows a frequency hopper, regularly hopping between 3 frequencies. It is clearly visible that the signal is not completely off during the first hop (lowest frequency to middle frequency), whereas no significant RF level can be observed during the second hop.

4.1.1 Parameters

The spectrogram offers various parameters that help the user to optimize the display for his specific application. As already mentioned, the user interface originates from the user interface of a spectrum analyzer, allowing an intuitive usage. The same is valid for the parameter denotation.

Center Frequency, Span, RBW

To optimize the spectrogram for the analysis of a VCO transient analysis, the R&S FSVR is at first set to the VCO target frequency setting the center frequency parameter. A band, 40 MHz wide around the center frequency, passes through the IF stages of the analyzer and is finally digitized. Due to its FFT concept, the maximum IF bandwidth of 40 MHz is the maximum span that can be displayed in real time mode.
Within the real-time FPGAs, the FFT length is fixed at 1024 bins. A fixed FFT length implies a coupling between span and resolution bandwidth (RBW). Reducing the span in real-time mode therefore automatically reduces the RBW and vice versa.

Sweeptime, Detector

As already mentioned in section 3, the spectrogram displays consecutive spectra, where each spectrum consists of multiple FFTs. The parameter sweeptime determines the amount of time used to sample data for one spectrum. One spectrum containing several FFTs is sometimes also called a frame. In a conventional spectrum analyzer, the parameter sweeptime describes the amount of time needed to sweep over the selected frequency range. As the effect is the same, i.e. it takes the sweeptime to complete one spectrum, the real-time parameter is also called sweeptime.
Combining several FFTs into one spectrum during the selected sweeptime offers several possibilities of weighing each single FFT result: averaging is an obvious one. Other possibilities of combining several FFTs are picking either the maximum or minimum for each frequency point, or selecting an arbitrary FFT result to represent the entire sweeptime. The combination of FFTs is called detection and a detector is available for each of the mentioned methods: Average, Positive Peak, Negative Peak, and Sample. Positive Peak is the default selection to make sure that even the shortest events can be analyzed. So as a summary, the parameters detector and sweeptime describe the data reduction from multiple FFTs to a single spectrum, as shown in Figure 14 for a peak detector. A detector is only used for the spectrogram and real-time spectrum displays, as shown in Figure 4

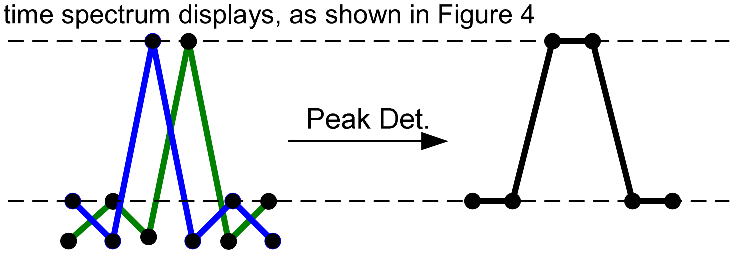

Figure 14: Peak detector combining two FFTs into one spectrum

History Depth

The R&S FSVR keeps the displayed spectra in its memory, as well as the IQ data. Obviously, keeping the spectrum traces in memory requires significantly less memory than keeping the IQ data, as the spectrum has only 801 points per sweeptime interval. The device memory is sufficient to save up to 100,000 spectra. With the sweeptime from the above example (200 µs), the maximum spectrogram history depth is 20 seconds. The amount of time covered by the spectrogram history depth increases with the sweeptime. The parameter History Depth controls the amount of spectra kept in the spectrogram memory.

Color Mapping

A key element of the spectrogram is the color mapping, i.e. the conversion of the numeric power values into a corresponding color. In order to help the user getting the desired information from the spectrogram, the R&S FSVR has 4 major parameters for color mapping that can be adjusted. These are:

- The color map: it determines the set of colors for level encoding. "Hot" ranges over the entire color spectrum, with blue representing low levels and red high levels. "Cold" is the same range but assigned vice versa, i.e. red corresponds to low power levels. "Radar" ranges from black through the entire range of greens, from dark to light green. "Grayscale" is black and white, ranging from black for low levels through grey to white.
- The *Start* value: it determines the threshold at which the color starts to change. All power levels lower than the *Start* value will appear in the same color, dark blue in the example. The percentage given as a numeric value is relative to the reference level.
- The *Stop* value: it determines the upper threshold. All power levels larger than this value will appear in the same color, light red in the example.
- The *Shape* value: a numeric value between -1 and 1. A value of 0 describes a linear distribution of colors between the lower and upper thresholds. Values larger than 0 result in a steeper slope of the curve for higher power levels. Higher power levels are therefore resolved with more color grades than levels close to the lower threshold. For values smaller than 0, levels close to the lower threshold are resolved with more colors grades.

The R&S FSVR assists the user during color mapping settings with a probability distribution of power levels. The display below the live update of the spectrogram displays probability over power level. On the left hand side of this graph, the Gaussian shaped probability distribution of the noise floor is clearly visible. The *Start* value may be modified in order to exclude the noise floor, to clearly display the signal under investigation.

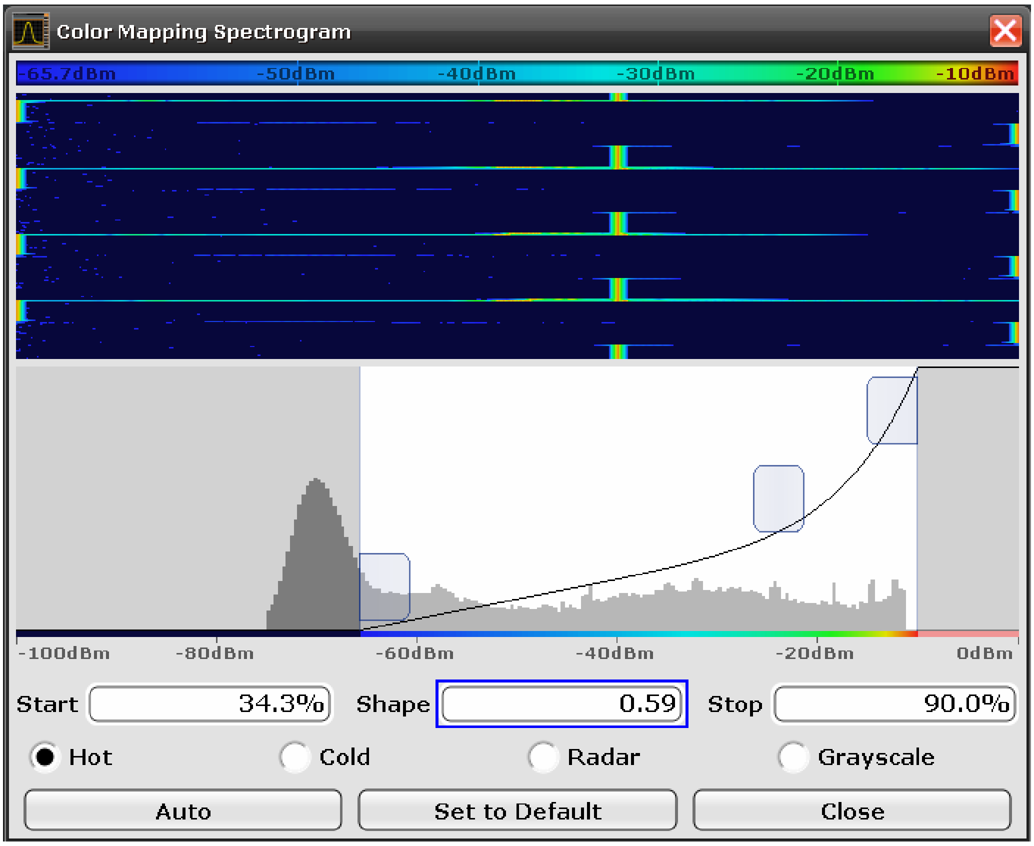

Figure 15: Color mapping: the lower end of the color map is increased to fade out noise – the default straight line was also modified to highlight signal levels above -20 dBm

4.2 Spectrogram with Real-Time Spectrum

For detailed analysis of a spectrogram especially during post processing, it is often helpful to additionally activate the Real-Time Spectrum display. The Real-Time Spectrum always shows one spectrum, i.e. one line or frame of the spectrogram. The time position of the active maker determines which particular spectrum is displayed in the real-time spectrum screen. In case no marker is active, the latest spectrum, i.e. the top most line, is displayed

The spectrogram together with the real-time spectrum is ideal for detailed timing and spectral analysis. Due to the color coding of levels, it is hard to position markers on exactly the desired peak in the spectrogram. The signal under investigation for this example is a CW signal with short sections of frequency modulation (FM) applied. In order to analyze the time in between two consecutive FM sections, a pair of markers is used. A double input box will appear indicating the time and frequency position of the current marker.

In the example above, the modulating signal is known. It is a 1 kHz CW signal. With a 1 kHz modulating signal, the corresponding FM will exhibit a significant peak at 1 kHz offset at both sides of the RF carrier signal. Navigate the marker to an active FM section and position it on either 1 kHz side lobe. The R&S FSVR allows marker navigation and positioning in the spectrogram using its touch screen functionality or the time and frequency input boxes. Use the real-time spectrum display to control, whether the marker is properly positioned on the side lobe.

Marker peak searches can be performed along either the time or frequency axis. The default search axis is the frequency axis. A maximum search along both axes at the same time is also available. The marker search function can either be performed over the entire displayed spectrogram, or the entire spectrogram that is in memory (up to 100,000 lines). The result of the timing measurement can be directly taken from the Frame Time box of the delta marker. The position in frequency and time of a delta marker is always relative to its corresponding reference marker. Figure 16 shows the result of the timing measurement.

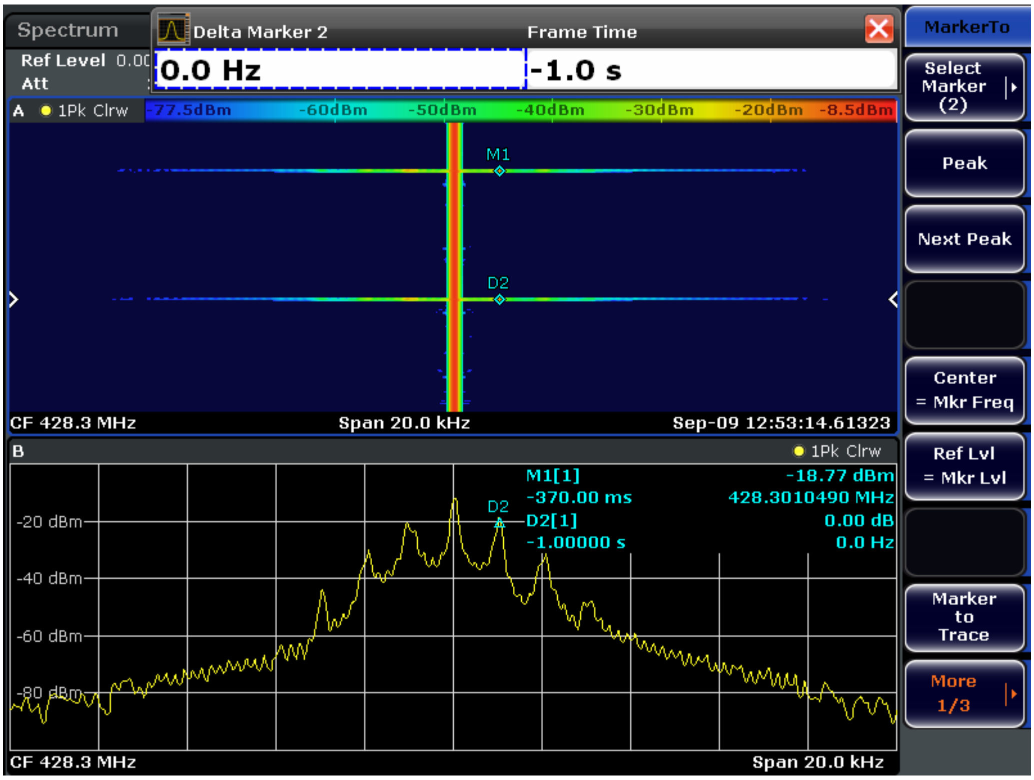

Figure 16: Timing measurement with split screen spectrogram and real-time spectrum

4.3 Persistence Spectrum

The R&S FSVR offers not only the spectrogram as a mode of display in real-time mode, but also the so called Persistence Spectrum. This mode is also referred to as spectral histogram. The two names also highlight the main features of this display mode: persistence and histogram information. Persistence helps to view even very short events that the human eye could not capture otherwise. Moreover, it also allows comparison between two events that are separated in time but both within a time frame called persistence granularity. This time frame specifies the amount of time it takes for a singular event to fade completely.
Histogram information is basically a counter that sums up the appearance of a certain frequency – level pair within a certain amount of time. Instead of displaying the total of a counter, the persistence spectrum displays the counter result normalized to the maximum achievable count, which yields a probability of appearance for each frequency – level pair.

The persistence spectrum is made up of a horizontal frequency axis and a vertical level axis just as a normal spectrum display. The color of each dot in the persistence spectrum contains the histogram information, i.e. the probability information.
A typical application for the persistence spectrum is the analysis of time varying signals. It is an especially powerful tool to give the user a first idea of a signal, before it can be analyzed in detail. Fast frequency hops can be clearly distinguished from amplitude drops with the persistence spectrum, whereas conventional analyzers may mislead the user. Opposite to the spectrogram display, the persistence spectrum offers a higher level resolution, as it does not employ color coding. In addition, the persistence spectrum achieves a higher time resolution as the real-time spectrogram, as it does not use detectors. Another application for the persistence spectrum is the separation of superimposed signals if they can be distinguished in terms of probability distribution of frequency – level pairs.
Figure 17 shows a persistence spectrum of a noise-like signal resulting from a motor with brushes. Clearly visible is a weak GSM signal in the center of the span. A standard spectrum analyzer cannot resolve the two different signals, as it does not display probabilities for each signal point.

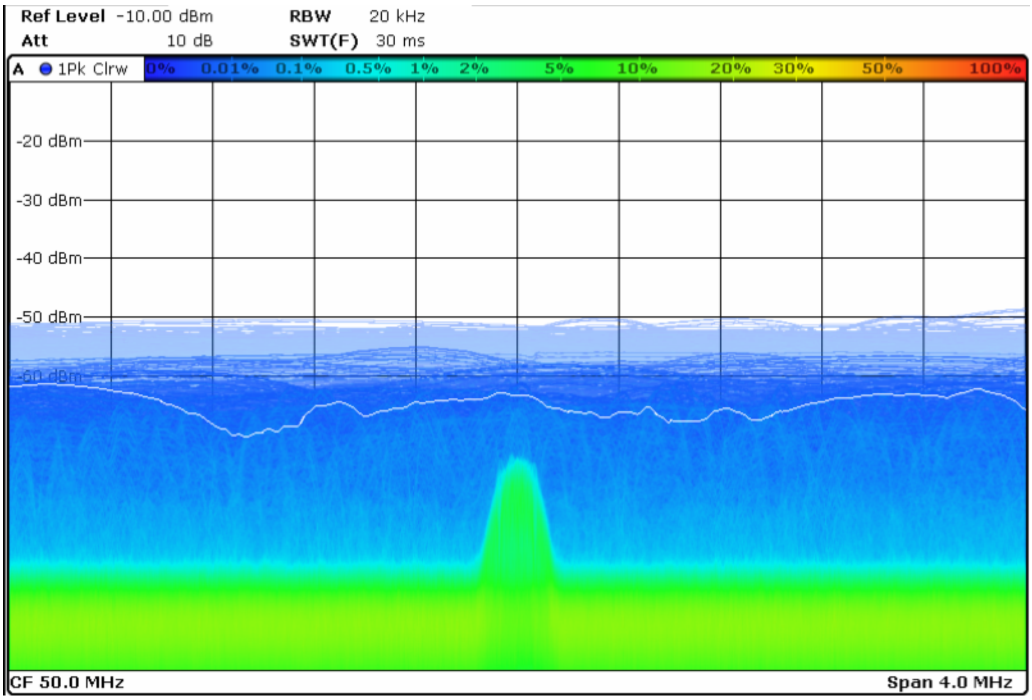

Figure 17: Wideband noise-like signal covering a GSM signal.

4.3.1 Parameters

Center Frequency, Span, RBW

These parameters specify the frequency and bandwidth setting of the R&S FSVR. They are identical throughout the real-time displays.

Persistence Granularity

Persistence Granularity is the timer for the sum operation already mentioned. A single histogram image is calculated during the persistence granularity time. The initial zero matrix with 600 by 801 elements represents 600 discrete power levels and 801 discrete frequency steps. Assuming a full 40 MHz span for the example, the R&S FSVR runs an FFT update rate of 250,000 per second, thus providing an FFT every 4 µs. With the default persistence granularity of 30 ms, the maximum count for one cell of the matrix is 7500. For each new histogram, i.e. every time the 30 ms interval is completed, the matrix is reset to zero for each element.

Figure 18 shows this process with a 6 by 8 elements matrix and FFT time to granularity ratio of 2, instead of a 600x801 matrix and a FFT to granularity ratio of 7500. Two FFTs are calculated. Both FFTs contain the same signal and varying noise neighboring the signal. Figure 18 illustrates the conversion of an FFT into a matrix of frequency – level pairs. The two matrices are summed up into the result matrix. The result matrix determines the color of the result trace. In this example, red corresponds to a high count or probability, whereas the noise band is displayed in blue for a lower probability.

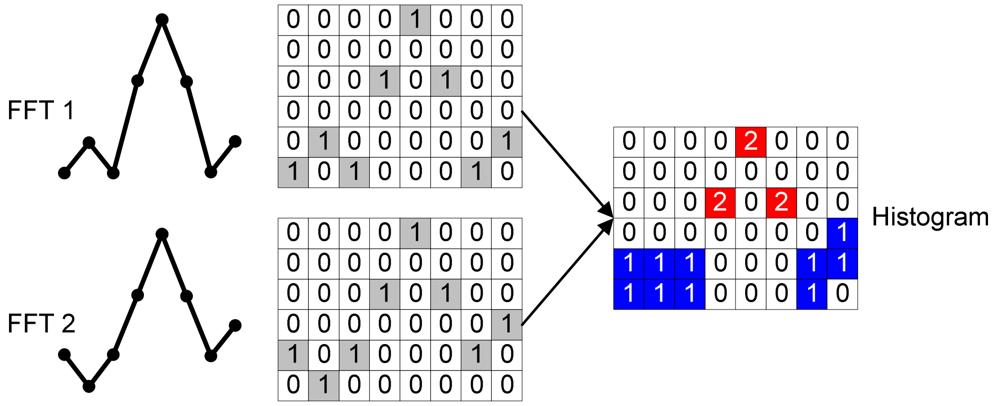

Figure 18: Schematic of histogram calculation (dot style)

Persistence

The persistence parameter determines the amount of time until a trace has completely faded. During this interval, the trace looses intensity. As each new trace is plotted with full intensity, the fading allows the determination between consecutive histograms, although multiple histograms are plotted within the same display. The persistence feature simulates the persistence that can be observed on cathode ray tube instruments.

MaxHold Intensity

During analysis of a time varying signal, level variations are usually of great interest. In detail, the ratio between the current signal and the maximum occurred signal. The so called MaxHold trace allows a worst case estimation of signal-to-noise-ratios (SNR), when talking about noise or interferers. For useful signals, it allows an estimation of amplitude variation. The persistence spectrum display can hold a MaxHold trace on top of the persistence spectrum. As already mentioned, the persistence traces will fade in intensity. The MaxHold trace in contrast is assigned a transparency value to allow determination between MaxHold trace and persistence spectrum. The MaxHold Intensity parameter specifies the level of transparency. The default value of 100 is a good choice for standard applications. To switch off the MaxHold, a level of 0 may be set, whereas the maximum value of 255 keeps the trace non-transparent. With no transparency, the trace can no longer be distinguished from the current histogram trace.

A MaxHold trace is cleared with every new setting on the R&S FSVR or with the Reset MaxHold button.

Style

The FFT matrices in Figure 18 contain only a single value per frequency column. This is the level value returned by the FFT. The example corresponds to the Dot style, i.e. the matrices are filled with dots only. In contrast, vector style mode forces each element with a 1 entry to have at least one neighboring 1 element. Two consecutive frequency points are always connected with 1-elements, independent of the level difference. To derive the matrices in Figure 19 from those in Figure 18, additional “1” elements are inserted to connect the “1” in column 4 to the neighboring “1” in columns 3 and 5. Figure 19 shows the vector style representation for exactly the same example that was used in Figure 18 for dot style.

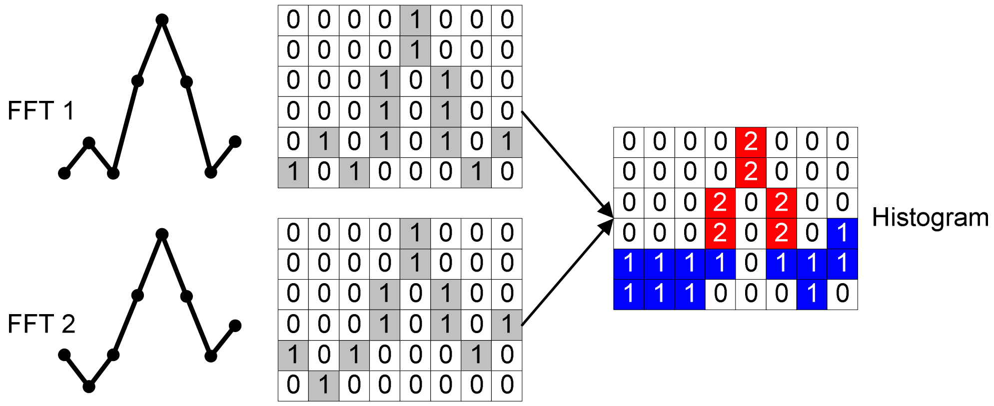

Figure 19: Histogram calculation using vector style

The additional “1” elements result in an increase in probability levels when changing from dot to vector mode. The increase is especially visible in areas with noise like signals, i.e. large level fluctuations.

Real-Time FFT trace: Detector, Sweeptime

The persistence spectrum display creates persistence and histogram information directly from the FFT results. There is no need to use detectors for data reduction as in the spectrogram, as the histogram algorithm already reduces data to a rate that can easily be displayed.
The detector setting in persistence spectrum mode affects only the real-time spectrum that trace can be plotted on top of the persistence spectrum. It assists the user in a fast recognition of the latest signal shape. For the FFT plot, the common parameters, such as detector, sweep time, and trace mode are used. These parameters influence only the real-time FFT trace, but not the persistence or histogram display.

Color Mapping

Color mapping for the persistence spectrum is identical to color mapping for the spectrogram. It provides the same dialog box and behaves in the same way. The probability distribution in the bottom part of the dialog provides information on the distribution of the color coded probability.
The dialog box offers a checkbox *Truncate*. Once activated, all values below the *Start* value and above the *Stop* value will no longer be shown with the lowest or highest color, but in black, i.e. they will be invisible. This feature is especially useful is only spectral components of a certain probability shall be displayed.

A new color mapping is usually necessary after changing the persistence style from vector to dot or vice versa, as the resulting probabilities may vary largely as explained above.

5 Ordering Information

R&S FSVR7	Real-Time spectrum analyzer, 10 Hz to 7GHz, 40MHz bandwidth	1311.0006.07
R&S FSVR13	Real-Time spectrum analyzer, 10 Hz to 13GHz, 40MHz bandwidth	1311.0006.13
R&S FSVR30	Real-Time spectrum analyzer, 10 Hz to 30GHz, 40MHz bandwidth	1311.0006.30
R&S FSVR40	Real-Time spectrum analyzer, 10 Hz to 40GHz, 40MHz bandwidth	1311.0006.40

About Rohde & Schwarz
Rohde & Schwarz is an independent group of companies specializing in electronics. It is a leading supplier of solutions in the fields of test and measurement, broadcasting, radiomonitoring and radiolocation, as well as secure communications. Established more than 75 years ago, Rohde & Schwarz has a global presence and a dedicated service network in over 70 countries. Company headquarters are in Munich, Germany.

Environmental commitment

- Energy-efficient products
- Continuous improvement in environmental sustainability
- ISO 14001-certified environmental management system

Certified Quality System
ISO 9001

Regional contact

Europe, Africa, Middle East
+49 89 4129 12345
customersupport@rohde-schwarz.com

North America
1-888-TEST-RSA (1-888-837-8772)
customer.support@rsa.rohde-schwarz.com

Latin America
+1-410-910-7988
customersupport.la@rohde-schwarz.com

Asia/Pacific
+65 65 13 04 88
customersupport.asia@rohde-schwarz.com

Rohde & Schwarz GmbH & Co. KG
Mühldorfstraße 15 | D - 81671 München
Phone + 49 89 4129 - 0 | Fax + 49 89 4129 – 13777

www.rohde-schwarz.com