Implementation of Real-Time Spectrum Analysis
White Paper

Products:

- R&S®FSW
- R&S®FSVR

This White Paper describes the implementation of real-time capabilities within the R&S FSW (with option B160R) and the R&S FSVR. It shows fields of application as well as the technical implementation.

ROHDE&SCHWARZ

White Paper

Dr. Florian Ramian
3.2015 - 1EF77_3e

Table of Contents

1 Real-Time Analysis

1.1 What "Real-Time" Stands for in R&S Real-Time Analyzers

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. 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 (R&S FSV, R&S FSW), makes excessive use of the FFT method for narrow resolution bandwidths. In addition, it introduces complex digital RBW filters for swept measurements. These complex digital filters can be swept up to 30 times faster than their analog counterparts.

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 and R&S FSW platforms. 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 (see Fig. 1-1).

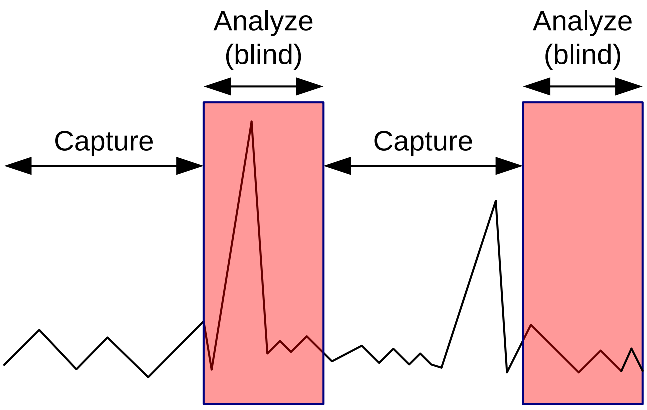

Fig. 1-1: Sequential capture and analysis as used in e.g. FFT analyzers

Measuring signals in real time means: do not lose 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. up to 160 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.

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).

Fig. 1-2 shows the parallelized capture and analysis which avoids blind times. Clearly, nothing remains undetected with a real-time spectrum analyzer.

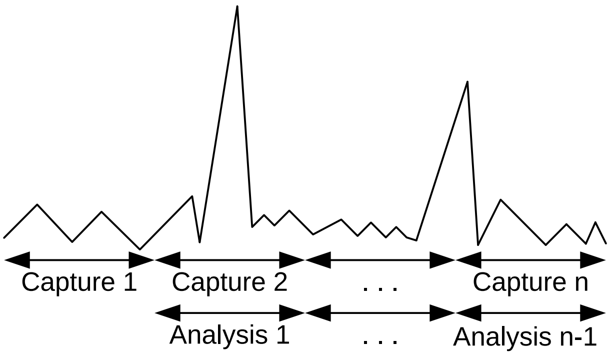

Fig. 1-2: Parallel capture and analysis as used in real-time analyzers

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 and R&S FSW-B160R 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 (interfering) signals can be found in this ISM band. To analyze the influence of these interferers on the system under investigation, the FMT can be placed around the wanted signal to start or stop capturing data as soon as an interferer violates the frequency mask. Without going into details, it becomes clear from Fig. 1-3 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 losing any information, the Rohde & Schwarz real-time analyzers are able to give precise information of a time variant signal, such as e.g. signal probability.

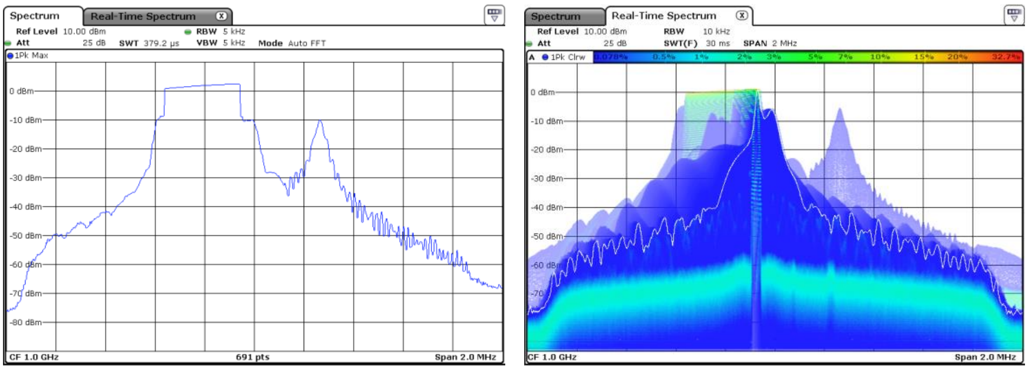

Fig. 1-3: Comparison of a Max Hold spectrum analyzer trace and a persistence spectrum trace

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

2 Real-Time Implementation in the R&S Real-Time Analyzers

The high-end signal and spectrum analyzer R&S FSW can be upgraded to a real-time analyzer by adding option FSW-B160R, whereas the R&S FSVR is a dedicated real-time analyzer that provides the functionality of a traditional signal and spectrum analyzer on top of the real-time functionality.

In terms of RF performance, the R&S FSW-B160R inherits the high end performance of the basic instrument. The R&S FSVR is based on the R&S FSV RF design, which determines its RF performance.

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 backends of the R&S FSW and R&S FSVR are 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 results on the screen.

All available real-time display modes and the frequency mask trigger run in parallel on the Rohde & Schwarz real-time analyzers. This means that all available 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 real-time analyzers.

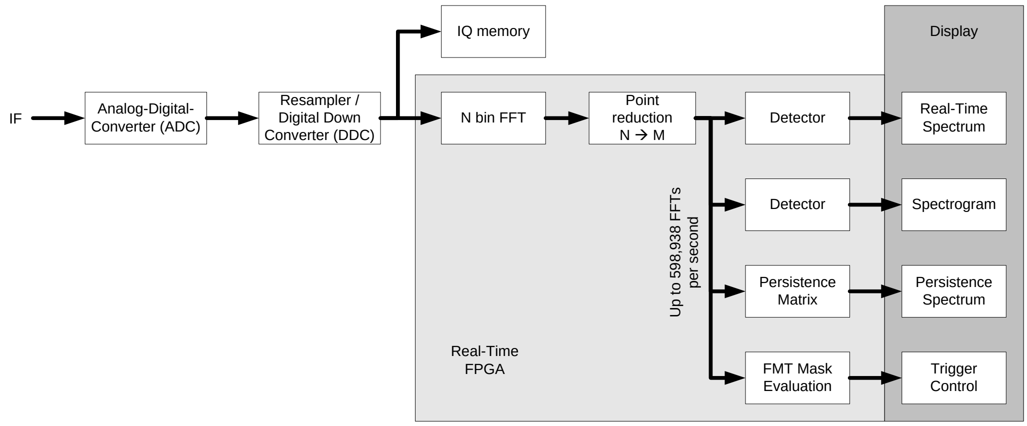

Fig. 2-1: Signal flow chart of the digital real-time part of the Rohde & Schwarz real-time analyzers

Fig. 2-1 shows the signal flow diagram from the A/D converter (ADC) to the display unit for the frequency domain displays. The ADC is operated at a constant sampling rate (128 MHz on the R&S FSVR, 1 GHz on the R&S FSW). The ADC streams raw data into the resampler and digital downconverter, which converts the input signal into the digital baseband, with a bandwidth equal to the selected frequency span and a sampling rate fulfilling the Nyquist criterion. 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, a 160 MHz span with 200 MSamples/s. 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.

On the R&S FSVR, 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. With the R&S FSW-B160R option, the FFT length is flexible, from 16,384 bins down to 1024 bins. Depending on the FFT length and operating mode, the R&S FSW reaches an FFT update rate of up to 585,938 FFTs per second.

In parallel to the FFT processing, the resampled baseband data is written into the analyzer’s I/Q memory for additional offline (non real-time) post-processing, like e.g. zooming into a captured region or reading out 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.

All Rohde & Schwarz real-time analyzers support time domain displays in real-time operation mode. Time domain displays are Power-versus-Time and Power-versus-Time Waterfall.

The R&S FSW real-time application operates in two different modes: High Resolution and Multi Domain. In High Resolution mode, the entire FPGA space is dedicated to FFT computation. In Multi Domain mode, a part of the FPGA is programmed for time domain displays. Thus, the maximum bandwidth in Multi Domain mode is limited to 100 MHz and the frequency resolution for a given span is limited compared to High Resolution mode.

2.1 FFT Windowing

Many engineers remember the frequency domain counterparts of common time domain signals and vice versa, without even thinking of it. A sine wave corresponds to two Dirac pulses in frequency domain, a pulsed signal to an Si-function and so on. These correspondences are true for the infinite continuous Fourier transform. However, all digital signal processing is time discrete and relies on a finite amount of sampled signal. Therefore a Discrete Fourier Transform (DFT) is utilized to transform time domain signals into the frequency domain. The best known DFT algorithm is the Fast Fourier Transform (FFT). DFTs deliver the same results as continuous transforms, if the finite sampled signal contains the same information as the continuous infinite signal. A sine wave for example, sampled at least with twice its frequency and for exactly an integer multiple of one period will result in a Dirac pulse when processed by an FFT algorithm. For all other cases, the resulting frequency domain signal will exhibit phenomena called spectral leakage, scalloping loss, and processing loss.

Applying a window function to the signal in time domain prior to FFT computation significantly reduces these effects, as the window forces the signal to be periodic with exactly the window length. Generally, the window length N_{Window} is equal to the FFT length N_{FFT}, but in some special cases, shorter windows may be used (see section 2.2 for an example).

Various different window functions with different properties exist. In general, the properties spectral leakage, amplitude accuracy, and frequency resolution influence the decision for a specific window. A so called rectangular window is automatically applied whenever the sampled signal is limited to a finite acquisition length. In this case the window length equals the acquisition length.

Different window functions not only influence the spectral representation of the windowed signal, but also the time domain representation. Fig. 2-2 shows the attenuating (weighting) effect of a Blackman window in time domain compared to a rectangular window.

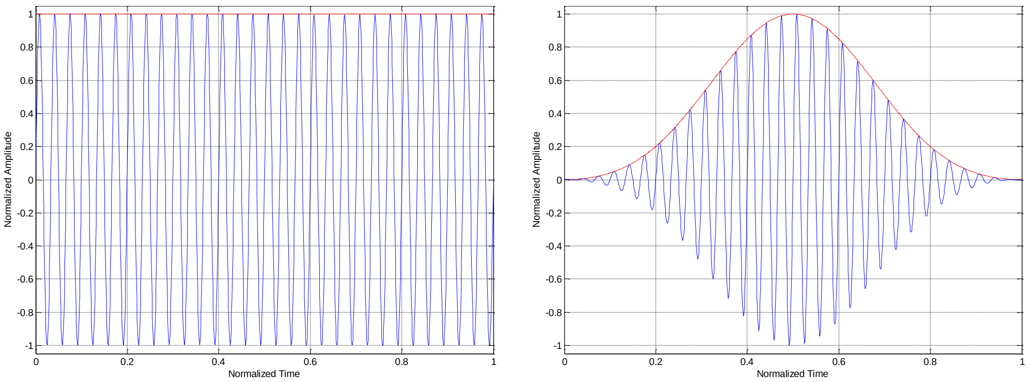

Fig. 2-2:Sine wave signal in time domain. Different window functions are applied to the same signal: left - rectangular, right - Blackman

For rare events, such as pulses, the location within the window function makes a significant difference in terms of level accuracy. Pulses located in the central part of the window function will be displayed with their correct power. Pulses located towards the edges are significantly attenuated. This dependency is minimized on a real-time analyzer with sufficient overlapping (see section 2.4).

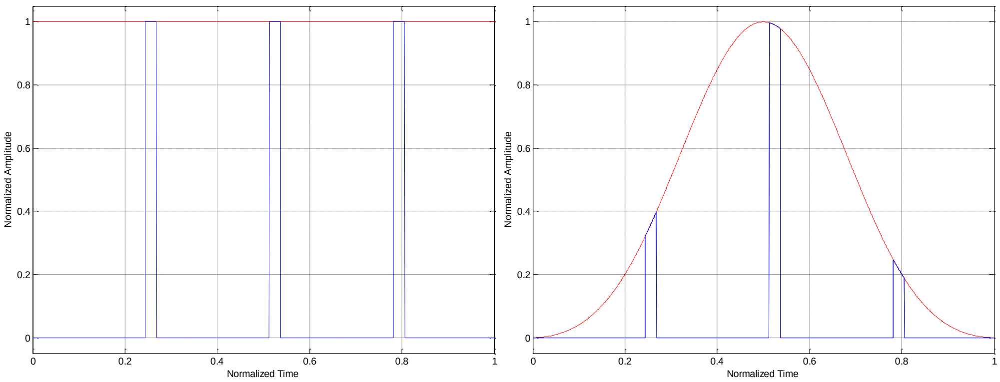

Fig. 2-3: Pulse signal in time domain. Different window functions are applied to the same signal: left - rectangular - no weighting, right - Blackman - pulses on the right and left are highly attenuated

As mentioned in the beginning, a window function is usually chosen for its characteristics in the frequency domain. Fig. 2-4 shows the spectral representation of the sine wave signal from Fig. 2-2. The figure shows the difference between a Blackman window and a rectangular window. The Blackman window clearly reduces the spectral leakage, i.e. the distribution of signal power over the entire spectral range; however, the rectangular window has the best frequency resolution, i.e. the narrowest main lobe. So given that you are looking for a very weak signal in the presence of a strong signal, the Blackman window is a good choice, since it has low spectral leakage. For many closely spaced carriers with similar power, the rectangular window is better suited to resolve all the involved carriers.

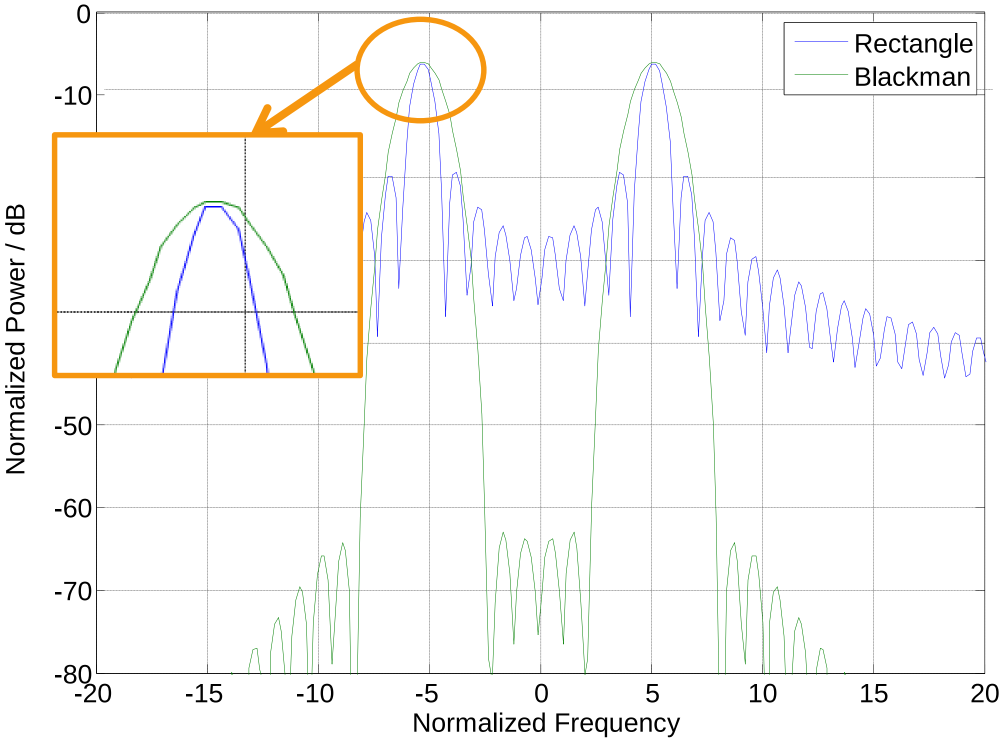

Fig. 2-4: Frequency domain view of the signal from Fig. 2-2. The Blackman window shows significantly lower spectral leakage than the rectangular window, but a wider main lobe. The inlay compares the amplitude accuracy.

Table 2-1 gives an overview of common window functions and their spectral leakage and frequency resolution capabilities. In addition, the table lists the amplitude accuracy, also called scalloping loss. Whenever a frequency component of a signal is not located on exactly an FFT bin, but somewhere in between two bins, scalloping loss occurs. The inlay of Fig. 2-4 shows the difference in amplitude accuracy for a rectangular and a Blackman window, which can be as high as 1 dB for this example.

Many FFT spectrum analyzers (including real-time analyzers) offer several different window functions to maximize their performance. Table 2-1 assists you in selecting the best window function for a specific application.

Window	Spectral Leakage	Amplitude Accuracy	Frequency Resolution
Blackman	Best	Good	Fair
Flattop	Good	Best	Poor
Gaussian	Fair	Good	Fair
Rectangle (none)	Poor	Poor	Best
Hann(ing)	Good	Fair	Good
Hamming	Fair	Fair	Good
Kaiser	Good	Good	Fair

Table 2-1: Overview of window function characteristics

2.2 FFT Length and Window Length

The length of an FFT as introduced above is defined by the number of samples that are fed into a single FFT. Most implementations of the FFT are based on an algorithm that takes only power of 2 lengths. A very common length therefore is 1024 bins.

Assuming a given bandwidth (or span) under investigation, the FFT length determines the minimum achievable resolution bandwidth (RBW) or frequency separation. Clearly, a longer FFT length is equivalent to more samples and therefore a longer capture time (assuming a constant sampling rate). A longer observation (capture) time allows better frequency resolution, i.e. smaller RBW. Therefore, it is equivalent to a lower displayed noise and thus a higher sensitivity. Shorter FFTs benefit from less computational effort and therefore higher FFT update rates.

The window length in turn describes the length of a window function that can be applied to the time domain signal, before FFT computation. The window length can therefore be equal to or less than the FFT length. Window functions can effectively shorten the FFT length, enabling the user to minimize the duration of a capture. This approach can provide 100% probability of intercept (POI), even when measuring extremely narrow pulses. Section 2.6 covers this topic in detail. The R&S FSVR uses a constant FFT (and window) length of 1024 bins. The R&S FSW provides more flexibility, allowing a selectable span-to-RBW ratio. As mentioned above, longer FFTs are equivalent to lower RBW for a given span, i.e. a higher span-to-RBW ratio. Since span and RBW are common parameters for spectral measurements, these parameters can be changed within the real-time option of the R&S FSW instead of FFT length. The setting of span and RBW is of course limited by the supported FFT and window lengths. The R&S FSW supports FFTs from 1024 bins up to 16384 bins. For span-to-RBW ratios lower than 200 (1024 bin FFT equivalent for a Blackman window), the window length can be reduced down to 32 bins. Therefore, the span-to-RBW setting allows the user to simply adjust span and RBW to his needs, without having to worry about FFT and window lengths.

2.3 FFT Update Rate

Consecutive FFTs are the raw spectral data being used for all spectral displays. It is essential for highly accurate frequency mask triggering to calculate every single FFT as fast as possible. As explained in detail in section 3.1, each FFT is checked for a mask violation. More FFTs per second therefore correspond to a higher FMT accuracy and repeatability.

In addition, the minimum signal duration for 100% POI is directly influenced by the FFT update rate (see section 2.6). In short, the minimum duration of an event for a 100% POI is similar to the minimum width of a pulse, which can be displayed at the same spectral power as a corresponding CW signal. This is also referred to as 100% true power reading. The minimum pulse width for a 100% true power reading corresponds to the FFT length, as a pulse with a duration greater or equal to the FFT length appears as a continuous wave (CW) signal to the FFT.

The available time resolution in the real-time spectrogram also depends on the available FFT update rate. However, the time resolution depends on more parameters than just the FFT update rate (see section 2.5).

The R&S FSVR provides an update rate of up to 250,000 per second, whereas the R&S FSW runs up to 585,938 FFTs per second. The limiting factor for the FFT update rate is the computing power of the FPGA.

In the end, most parameters, such as time resolution, overlapping, and minimum signal duration for 100% POI depend on the FFT rate. Therefore the power and versatility of a real-time spectrum analyzer highly depends on its FFT update rate.

2.4 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 minimize level errors.

To show the critical situation, let's assume that the capture time frames for two subsequent FFTs do not overlap.

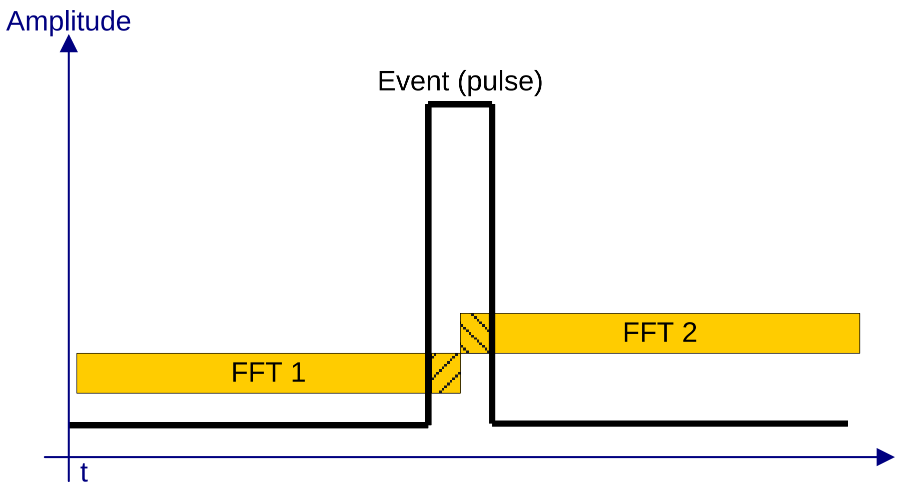

Fig. 2-5: Pulse captured by two consecutive FFT time frames without overlapping

The energy of a short pulse, which hits the border of the two capture time frames as shown in Fig. 2-5, 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.

The Rohde & Schwarz real-time analyzers utilize a technique called FFT overlapping to avoid this situation. Overlapping “reuses” samples that were already used to calculate the preceding FFT result. Fig. 2-6 shows a pulse signal that is captured by several overlapping FFT time frames.

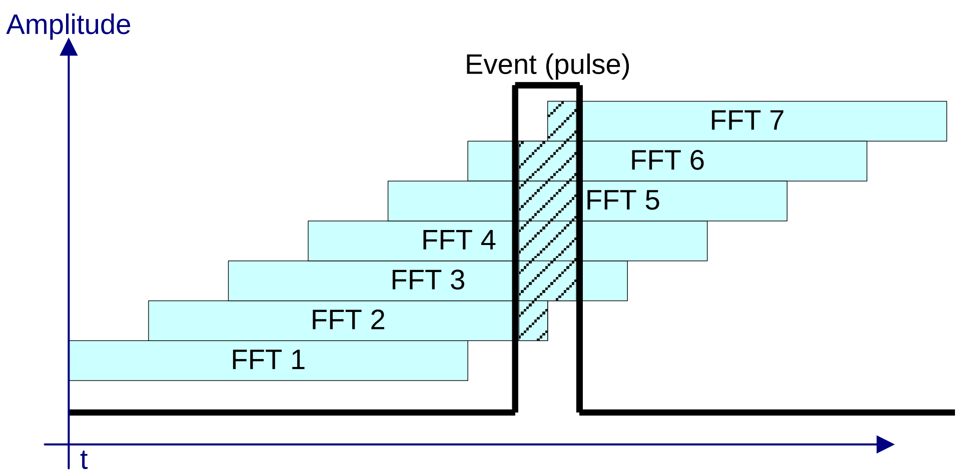

Fig. 2-6: Pulse captured with several consecutive overlapping FFT time frames

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. On the R&S FSW-B160R, overlapping depends on many factors, especially the FFT length and operating mode (High Resolution or Multi Domain), but is at least 50%, unless window lengths are below 1024 bins. Assuming an FFT length of 1024 bins and a bandwidth of 160 MHz, the FSW uses an overlapping of 2/3, i.e. 684 samples are reused. Depending on the selected bandwidth, i.e. the current sample rate, the overlapping may reach values up to (N_{Window}-1)/ N_{Window}, where N_{Window} is the window length.

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 (see section 2.1). From Fig. 2-7 it becomes evident that pulses shorter than the window length can be significantly attenuated if they are located near the window edges. Sufficient overlapping ensures that short pulses that could be attenuated if they occur at the edges of a window are also correctly measured at the center of the window in subsequent FFTs.

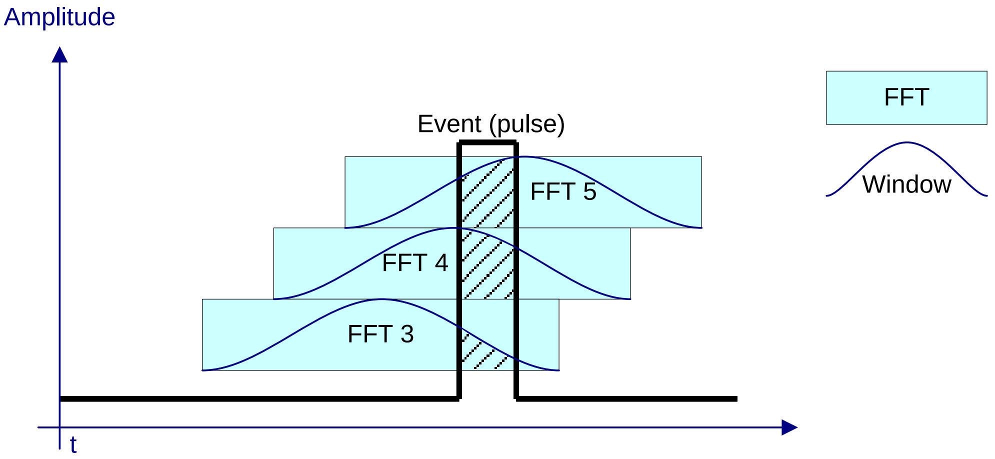

Fig. 2-7: Overlapping compensates effects resulting from windowing function

With an overlap ratio of 50% or higher, level errors caused by the window function can be neglected.

Overlapping is directly derived from the maximum real-time bandwidth and the maximum number of FFTs the analyzer can calculate per second. So for a real-time spectrum analyzer comparison, it is important to keep an eye on the specified FFT update rates, i.e. the number of FFTs per second.

2.5 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 occur in the same FFT result. The result is similar to a photograph that depicts everything that has happened within the exposure time. Taking the selected bandwidth (i.e. sampling rate) into account, the effective exposure time is calculated as

$$t_{exposure} = N_{FFT} \cdot \frac{1}{f_S}$$

with N_{FFT} the FFT length and f_S the sampling rate.

However, the exposure time is not necessarily equivalent to the time resolution, since most real-time analyzers employ overlapping (section 2.4).

The maximum theoretical time resolution is one sample, as two consecutive FFTs have to differ by at least one sample. For real-time displays, i.e. the spectrogram, the time resolution is the step width between two consecutive frames (lines) in the real-time spectrogram. As described in section 4.1.1, the sweep time parameter determines the time resolution. The minimum settable sweep time in replay zoom mode is 30 ns on the R&S FSW.

The following example illustrates the difference between time resolution and exposure time. A CW signal changes in frequency. In between the frequency change from frequency 1 (f_1) to frequency 2 (f_2), no RF signal is present for 2 μs (see timing diagram in Fig. 2-8).

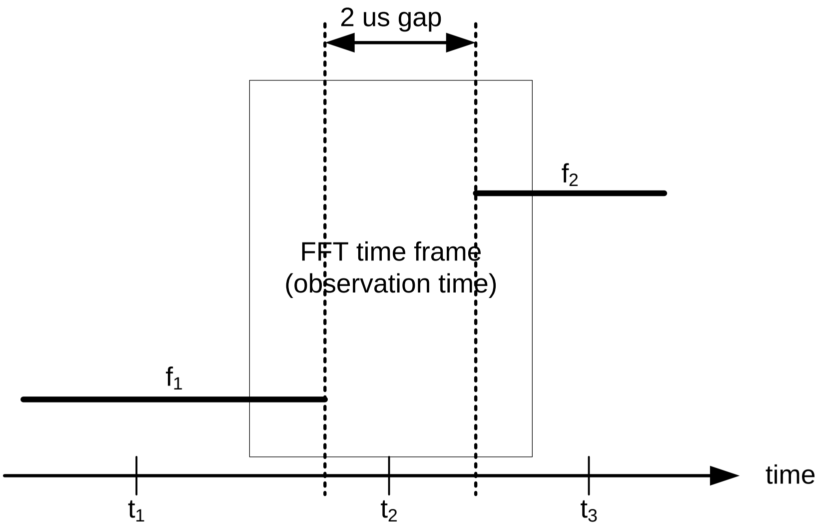

Fig. 2-8: Timing diagram of the FFT observation time example

Without having the above principle in mind, a user might expect an FFT result showing nothing but noise components in the 2 μs gap. A user with knowledge about FFT processes knows what to expect: consecutive FFT results show the spectral component for f_1 at first. During the 2 μs gap without a signal, the FFT result may show a spectral component of f_1 at reduced level as well as a spectral component of f_2 with a reduced level. As the time interval without a signal is smaller than 5.12 μs (N_{FFT}=1024 bins, f_S=200 MHz), there won't be an FFT result showing noise components only. The spectrogram in Fig. 2-9 shows the changing signal vs. time. The second spectrum trace from bottom in Fig. 2-9 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.

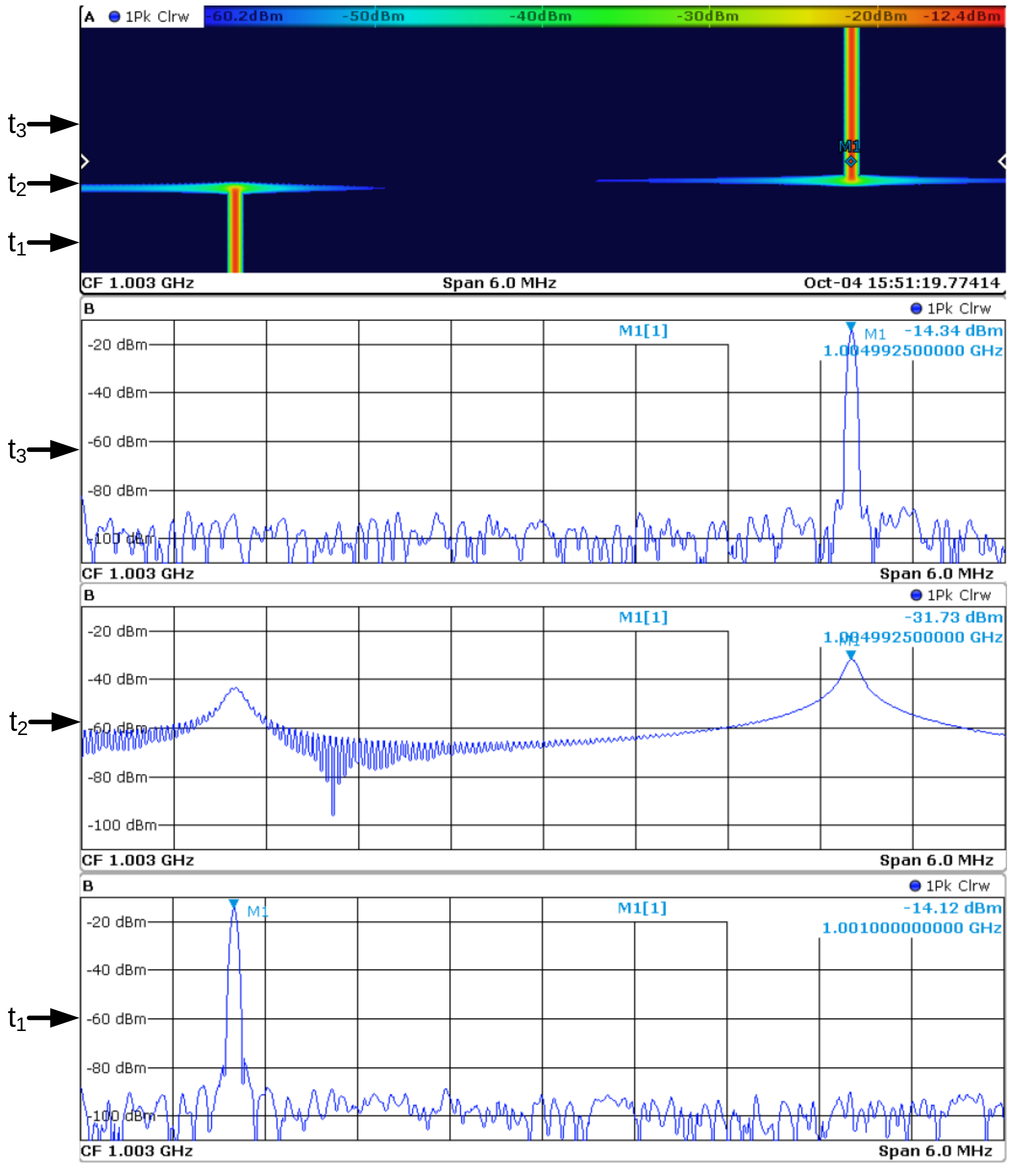

Fig. 2-9: Spectrogram showing a frequency hop. From top to bottom, the spectrum traces show frequency f2, both frequencies, and frequency f1. The trace with both frequencies results from the FFT time frame being longer than the gap between the signals.

2.6 Probability of Intercept (POI)

The minimum signal duration for 100% Probability of Intercept (POI) is a parameter that is specified for most real-time analyzers. By definition, the signal or event has to be present during at least one full FFT to reach 100% POI. For the standard operation mode of a real-time analyzer (window length equals FFT length), the minimum time for 100% POI is calculated as

$$t_{100\%\,POI} = \left(2 - P_{overlap}\right) \cdot \frac{N_{FFT}}{f_S}$$

with $P_{Overlap}$ the overlap ratio.

Clearly, if the duration is two times the FFT duration (minus the overlap), there will be at least one FFT that is completely filled with the signal, no matter if signal and FFT are aligned or not. In Fig. 2-10 an event with exactly the specified minimum duration for 100% POI is displayed. It will always be displayed with a 100% true power reading in frequency domain.

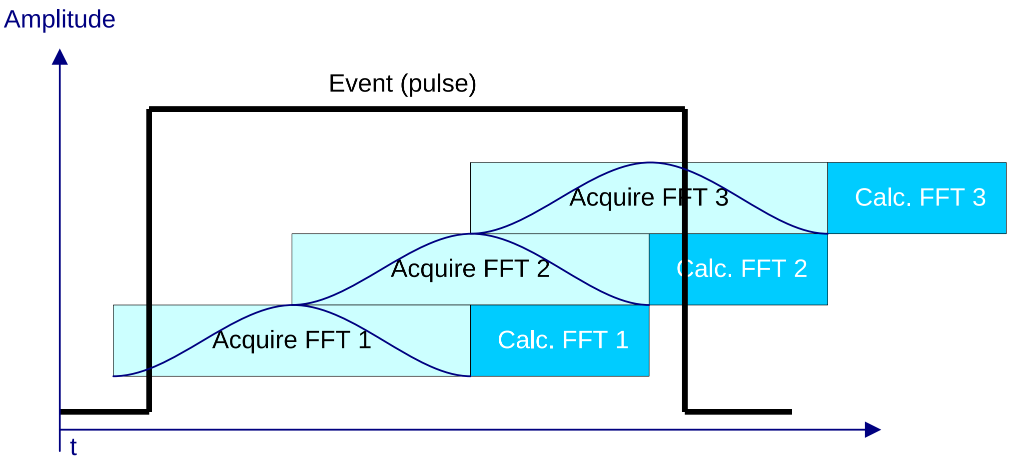

Fig. 2-10: Real-time configuration with overlap; the measured event is longer than the POI time and therefore displayed with 100% true power

For the scenario depicted above in Fig. 2-10, even the shortest pulse will be displayed in the frequency domain. No event can be missed due to the overlapping. However, the displayed spectral power decreases for pulses shorter than the minimum duration for 100% POI.

In order to decrease the minimum signal duration for 100% POI according to the above definition, many real-time analyzers can use windowing functions, shorter than the FFT length. Fig. 2-11 shows a schematic for this case. The benefit is that the FFT calculation can start earlier, since for a 1024 bin FFT and a 32 bin window function, only 32 values have to be sampled. The remaining 992 values are set to zero. So after 32 acquired samples, FFT computation and acquisition of the samples for the next FFT may already start. However, the computation time for the FFT is longer than the acquisition time for 32 values at full sampling rate. Therefore, some values acquired during the computation time are discarded. The result is blind time. So in order to fully describe the capabilities of a real-time analyzer, the minimum signal duration for 100% POI is one figure of merit. In most use cases, it is at least as important to not miss anything. Therefore the maximum length of an event that can be completely missed should also be mentioned.

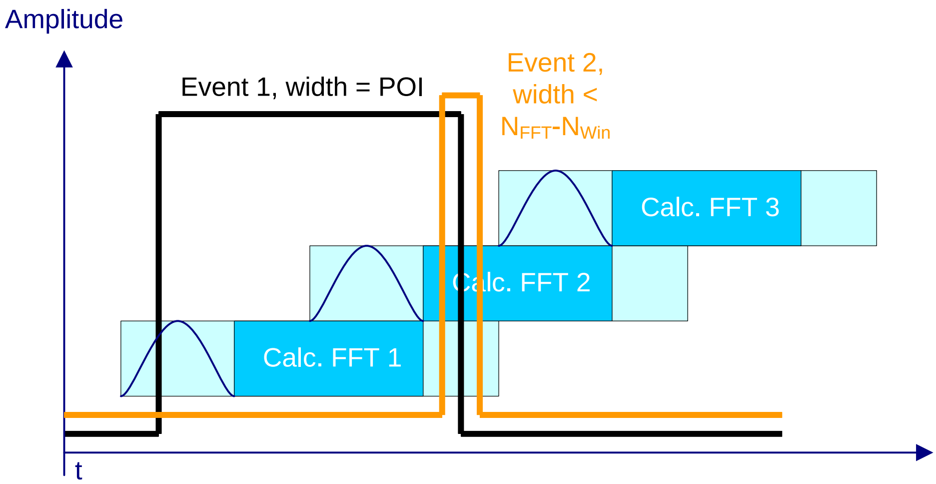

Fig. 2-11: Real-time configuration with no overlapping; event is short enough so it can completely disappear between two consecutive window functions

In default operation, both Rohde & Schwarz real-time analyzers operate with an overlap of more than 50%, i.e. they have no blind time. However, the R&S FSW allows configurations, where the window size is smaller than the FFT length in order to obtain very low minimum signal duration values for 100% POI. The R&S FSW does not only provide scalable FFT lengths (1k bins up to 16k bins) but also window functions smaller than 1024 bins (down to 32 bins). Table 2-2 gives an overview over real-time parameters for the default window function (Blackman). Note that the FFT update rate depends on the operation mode (High Resolution or Multi Domain). The minimum signal duration for 100% POI depends on the FFT rate and the selected span.

Blackman Window Function						
Span/RBW	**FFT Length (N_{FFT} in bins)**	**Window Length (N_{Window} in bins)**	**FFT Update Rate f_{FFT} (FFTs/s)**	**Overlap in % of window**	**Min. signal duration for 100% POI (µs)**	**Longest Event that can be 100% missed (µs)**
3200	16384	16384	36,621	66.7	109.22	0
1600	8192	8192	73,242	66.7	54.61	0
800	4096	4096	146,484	66.7	27.30	0
400	2048	2048	292,969	66.7	13.65	0
200	1024	1024	585,938	66.7	6.82	0
100	1024	512	585,938	33.4	4.26	0
50	1024	256	585,938	0	2.99	0.43
25	1024	128	585,938	0	2.35	1.07
12	1024	64	585,938	0	2.03	1.39
6	1024	32	585,938	0	1.87	1.55

Table 2-2:Overview of real-time parameters for the default Blackman window function in high precision operating mode, span 160 MHz

The corresponding minimum signal duration for non-overlapping operation can be calculated as

$$POI = \frac{1}{f_{FFT}} + \frac{N_{Window}}{f_S}$$

The maximum length of an event that may completely be missed by the analyzer, due to a window function shorter than the FFT length can be written as

$$T_{MaxMiss100} = \frac{1}{f_{FFT}} - \frac{N_{Window}}{f_S}$$

For applications with a need for 100% true power reading, e.g. frequency mask triggering on well-known signal levels, short windows are recommended. However, most real-time applications rely on “not missing any signal”. So for the majority of applications, the window length should match the FFT length and the real-time analyzer should reach a significant overlap ratio.

2.7 Replay Zoom Mode

The Rohde & Schwarz real-time analyzers provide a so called replay zoom mode. It allows the user to post-process the I/Q data which is still in memory. In many cases, the real-time analyzer will run at its full bandwidth. In a debugging scenario, we could assume the frequency mask trigger is configured to trigger on a certain side-band level. After having captured a trigger event, it turns out that the selected resolution bandwidth is not sufficient to clearly analyze the interferer.

The R&S FSVR as well as the R&S FSW allow to reprocess the captured I/Q data with different FFT parameters, such as center frequency, RBW, or span and save the user from having to wait for another interferer to appear.

Clearly, the FFT parameters can only be changed within the limits of the original I/Q capture. For ease of use, the spectrogram display will automatically turn on replay zoom mode, if the magnifying icon of the toolbar is used. In addition, the spectrogram will indicate which area can be zoomed into by keeping it light, whereas the area that cannot be zoomed is displayed with a dark shade. By placing the zoom rectangle on the spectrogram, the replay zoom parameters, such as center frequency and span are automatically set.

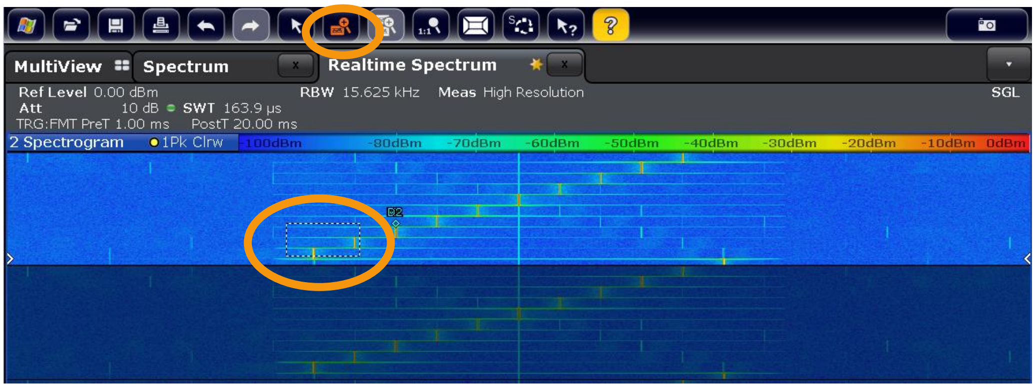

Fig. 2-12: Spectrogram with magnifying tool selected. Light area in spectrogram indicates zoom able area, dark area: not zoom able

3 Triggering on Real-Time Spectra

This section focuses on a trigger mechanism which is only available on 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 Fig. 2-1).

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 data 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. That means a limit check 585,938 times per second (R&S FSW). This translates to a time resolution of 1.71 µs.

The frequency mask can consist of up to 1001 points (801 on the R&S FSVR) and may have any shape.

The Rohde & Schwarz real-time analyzers offer 2 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)

Both 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 analyzer.

A trigger event on a normal spectrum analyzer starts data acquisition. A frequency mask trigger can exactly do the same. In addition, many real-time applications require a trigger event to stop data acquisition, once a certain event has happened. This trigger mode is called “Stop on Trigger”, whereas the default operation mode is “Auto Rearm”.

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, all

Rohde & Schwarz real-time analyzers provide a special trigger output. The trigger out port provides a trigger pulse with a pulse width of a level of 5 V every time the FMT triggers the instrument. This trigger pulse may be provided to a system setup as an external trigger source.

3.1.1 Setting up a Frequency Mask Trigger (FMT)

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 loses 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 fulfills this requirement can be easily set up with the FMT mask editor (Fig. 3-1). 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.

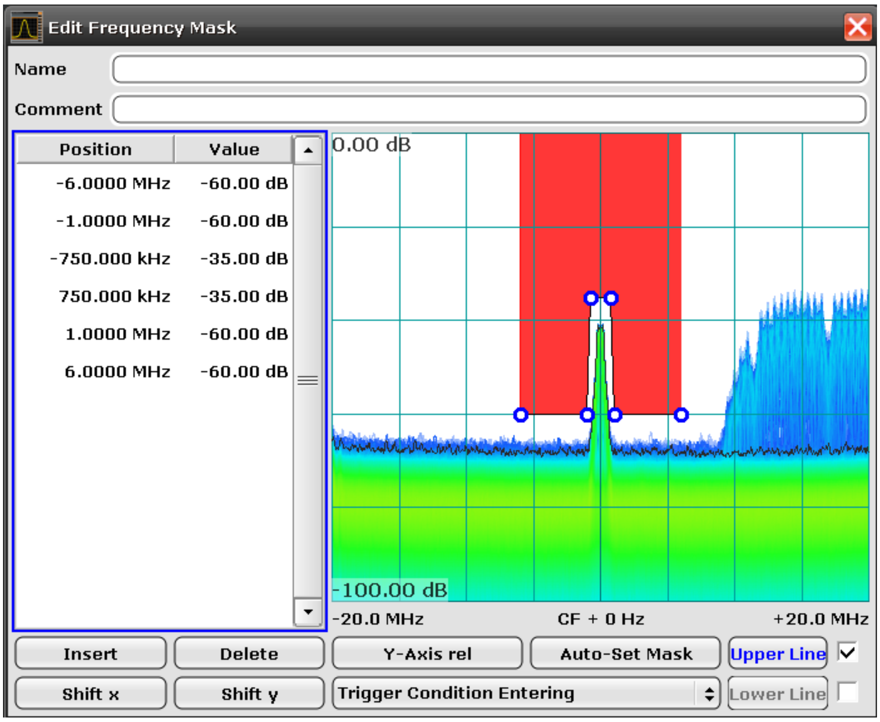

Fig. 3-1: FMT dialog box

To indicate an active FMT, the trigger mask appears in the current persistence or real-time spectrum display as a red background mask (see Fig. 3-2). Make sure that the analyzer is in Run Continuous mode, when using “Stop on Trigger” mode. To make sure that the displayed data contains all necessary information for your analysis, adjust the pre- and post-trigger time.

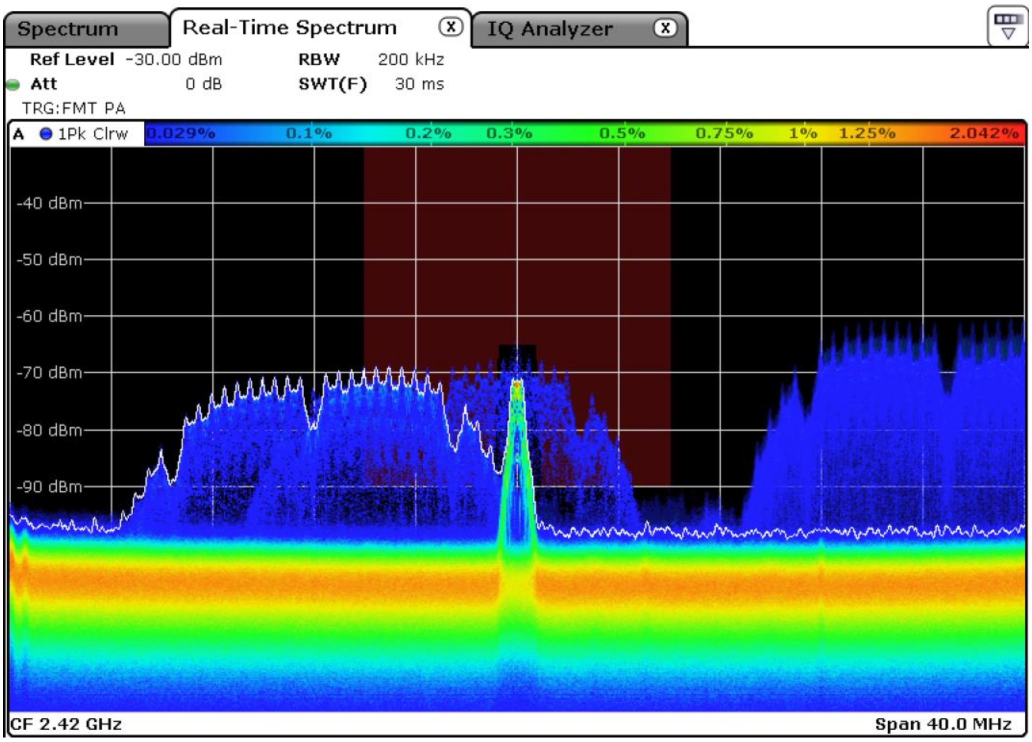

Fig. 3-2: Active FMT triggered by frequency change of a WLAN signal

Once the trigger condition appears, the analyzer 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.

Availability of FMT

The FMT works with all real-time display modes and in addition is supported for a number of non-real-time options such as e.g. the vector signal analysis option.

Stop on Trigger vs. Auto Rearm

As mentioned above, both modes are supported. Stop on Trigger will fit for most applications that are looking for rare occurrences of e.g. interferers, whereas the Auto Rearm trigger is best suited for applications looking at periodically repeating signals. Due to this concept, the Stop on Trigger mode delivers best results in continuous sweep mode. Auto Rearm mode, similar to an external trigger, has many applications that can use it in either single sweep or continuous mode.

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 FSW performs this mask check up to 585,938 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.

Fig. 3-3 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 analyzer triggers.

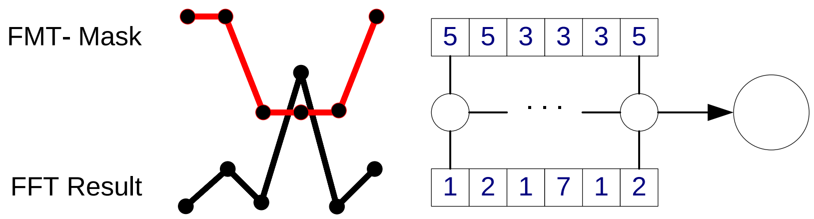

Fig. 3-3: Element-wise comparison of FMT mask with current FFT result

Extended limit check means that the FMT can link a complex condition to the limit line violation, such as entering or leaving.

As already mentioned, the FMT mask contains up to 1001 points (801 on the R&S FSVR), but may also be as short as 2 points. Shorter FMT definitions will be extended to 1001 points by interpolation within the firmware. The FMT trigger therefore always compares 1001 FFT points to 1001 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 2.6, 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 1.87 µs (R&S FSW with 32 bin window in high resolution mode). This figure is derived as explained in the POI section (2.6). 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 FSW require post-processing, as they have no means to process data in real-time.

The Rohde & Schwarz real-time analyzers do 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.

The R&S FSW-B160R has two different operating modes: the High Resolution mode and the Multi Domain mode. The Multi Domain mode offers power-versus-time displays, but is limited to 100 MHz of bandwidth and to a span/RBW ratio of 800 (for a Blackman window).

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.

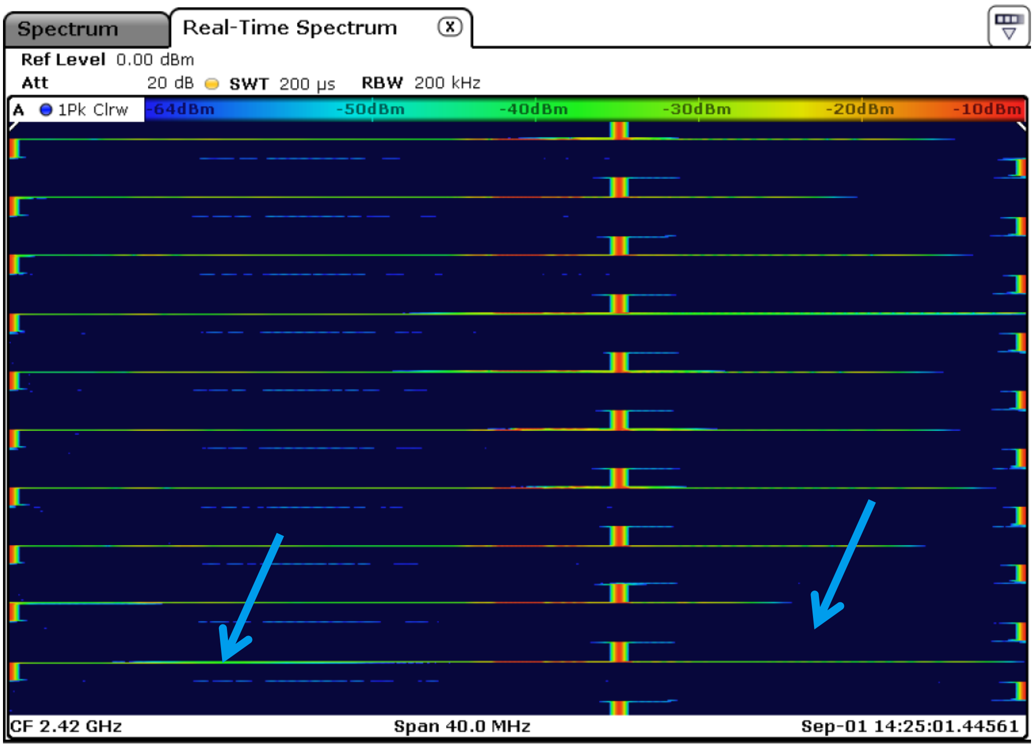

Fig. 4-1: Frequency hopper regularly hopping between 3 frequencies. First transition shows significant RF level, second transition shows almost no level during hop. Transitions highlighted by arrows.

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.

Fig. 4-1 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 real-time analyzer is at first set to the VCO target frequency setting the center frequency parameter. A band around the center frequency, 40 MHz wide on the R&S FSVR, up to 160 MHz wide on the R&S FSW, passes through the IF stages of the analyzer and is finally digitized. Due to its FFT concept, the maximum IF bandwidth is the maximum span that can be displayed in real time mode.

Within the real-time FPGAs, the resolution bandwidth (RBW) is tied to the FFT length and the span. On the R&S FSVR with its fixed FFT length, the RBW cannot be set explicitly. The R&S FSW with B160R allows different FFT lengths. By selecting a certain Span/RBW ratio, the RBW can be changed for given span. Since the available span/RBW ratios vary with the selected window function, Table 4-1 gives an overview of the maximum ratios available, i.e. the span/RBW ratio that corresponds to a 16k FFT. Clearly, the FFT or window length scales down with the selected ratio.

Maximum span/RBW ratio (16k FFT)						
Blackman	Flattop	Gaussian	Rectangle	Hann(ing)	Hamming	Kaiser
3200	1600	3200	6400	4000	4000	3200

Table 4-1: Maximum selectable span/RBW ratios for a different window functions

Sweep Time, Detector

As already mentioned in section 4.1, the spectrogram displays consecutive spectra, where each spectrum consists of multiple FFTs. The parameter sweep time therefore determines, along with the FFT length and the sampling rate the number of FFTs in a spectrum.

The nomenclature is derived from traditional spectrum analyzers. In a swept spectrum analyzer, the sweep time specifies how fast the local oscillator (LO) is swept over the selected span, i.e. how long it takes to acquire the full spectrum. Modern signal- and spectrum analyzers also use this parameter, even if they calculate FFTs instead of sweeping an LO. It is important to keep in mind that for an FFT the sweep time is the data acquisition time, not the measurement time which would also include computation time.

Taking this understanding of sweep time into a real-time analyzer consequently means: a real-time spectrum consists of data taken during the selected sweep time. Since FFTs are computed without interruption, the parameter sweep time determines how many FFTs are combined into one spectrum.

Combining several FFTs into one spectrum during the selected sweep time offers several possibilities of weighting 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 bin, or selecting an arbitrary FFT result to represent the entire sweep time. For combing the FFTs into one spectrum, the so called detector is used. In a traditional spectrum analyzer, the detector combines multiple measurement points into one trace point using a certain algorithm. The detector in a real-time analyzer combines multiple FFTs on a per bin basis, i.e. is equivalent to the detector on a traditional spectrum analyzer. Available detectors are: Average, Positive Peak, Negative Peak, and Sample. Positive Peak is the default selection to make sure that even the shortest events can be analyzed. A detector is only used for the spectrogram and real-time spectrum displays, as shown in Fig. 2-1.

In a summary, the parameters detector and sweep time describe the data reduction from multiple FFTs to a single spectrum, as shown in Fig. 4-2 for a peak detector.

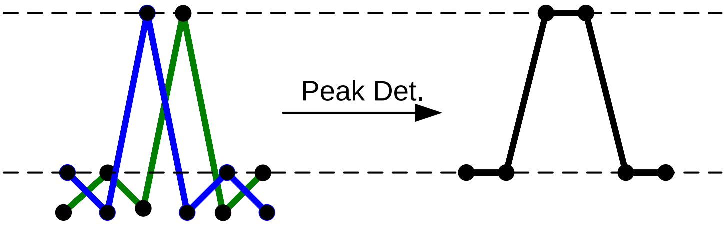

Fig. 4-2: Peak detector combining two FFTs into one spectrum

History Depth

The Rohde & Schwarz real-time analyzers keep the displayed spectra in their memory, as well as the I/Q data. Obviously, keeping the spectrum traces in memory requires in general significantly less memory than keeping the I/Q data, as the spectrum has only 801 to 1001 points per sweep time interval. The device memory is sufficient to save up to 100,000 spectra. With the sweep time from the above example (200 µs), the maximum spectrogram history depth is 20 seconds, whereas a sweep time of 30 ms allows the spectrogram to cover 3000 s, almost an hour. For continuous operation of the real-time analyzer, the history depth can be directly converted into a maximum display time by multiplying the depth of 100,000 frames with the selected sweep time. However, if the analyzer is in triggered mode or in case single sweeps are performed, the displayed spectrogram has gaps indicated by black lines. Therefore, the history depth cannot be converted to time (see Fig. 4-3).

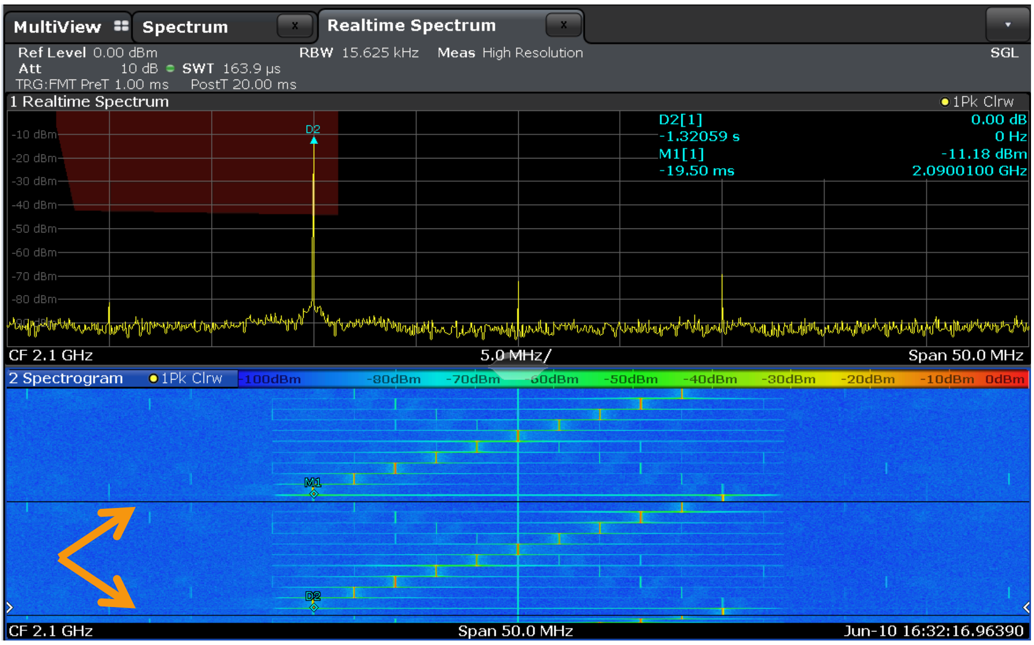

Fig. 4-3: Spectrogram with two black lines, indicating non-contiguous captures, i.e. gaps, due to triggered mode or single sweep operation

Color Mapping

A key element of the spectrogram is the color mapping, i.e. the conversion of numeric power values into a corresponding color. In order to help the user getting the desired information from the spectrogram, the Rohde & Schwarz real-time analyzers have 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 real-time analyzers assist 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.

The default setting is a linear color mapping curve ranging from reference level to the minimum displayable level.

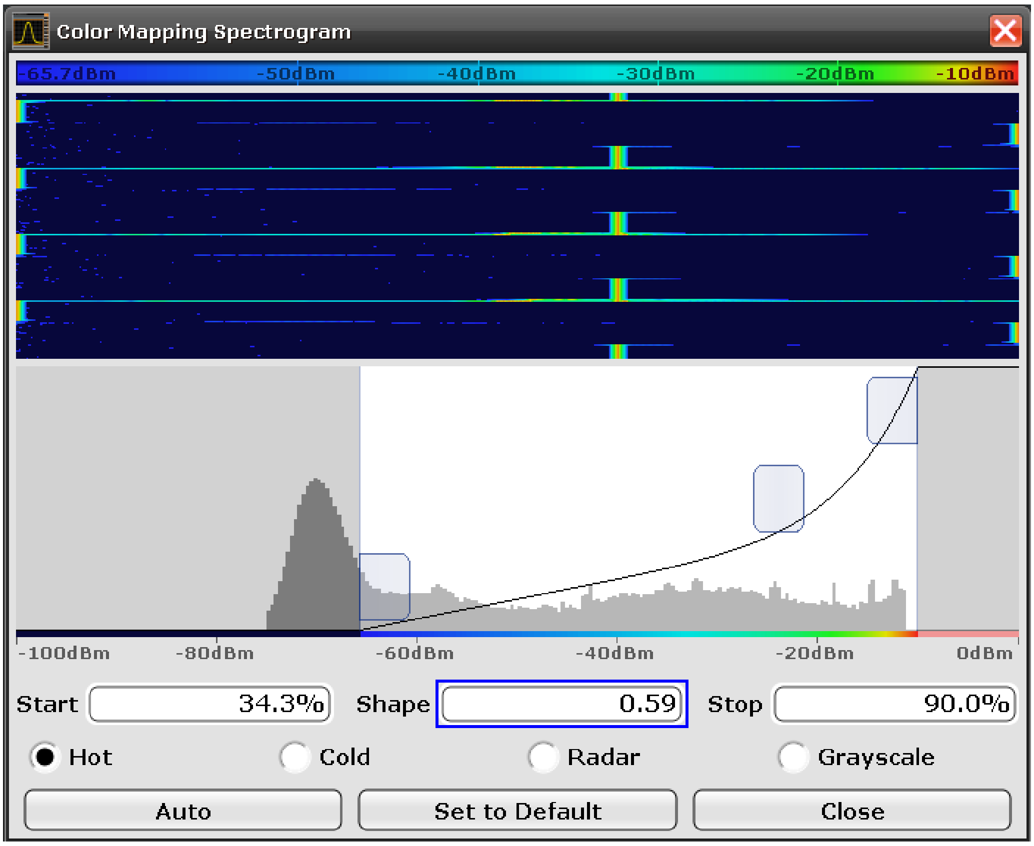

Fig. 4-4: Color mapping: the lower end of the color map is increased to fade out noise – the default straight line was also modified to highlight signal levels above -20 dBm

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 Rohde & Schwarz real-time analyzers allow marker navigation and positioning in the spectrogram using their 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. The default search axis is the frequency axis. A maximum search along both axes at the same time is also available (2D search). 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. Fig. 4-5 shows the result of the timing measurement.

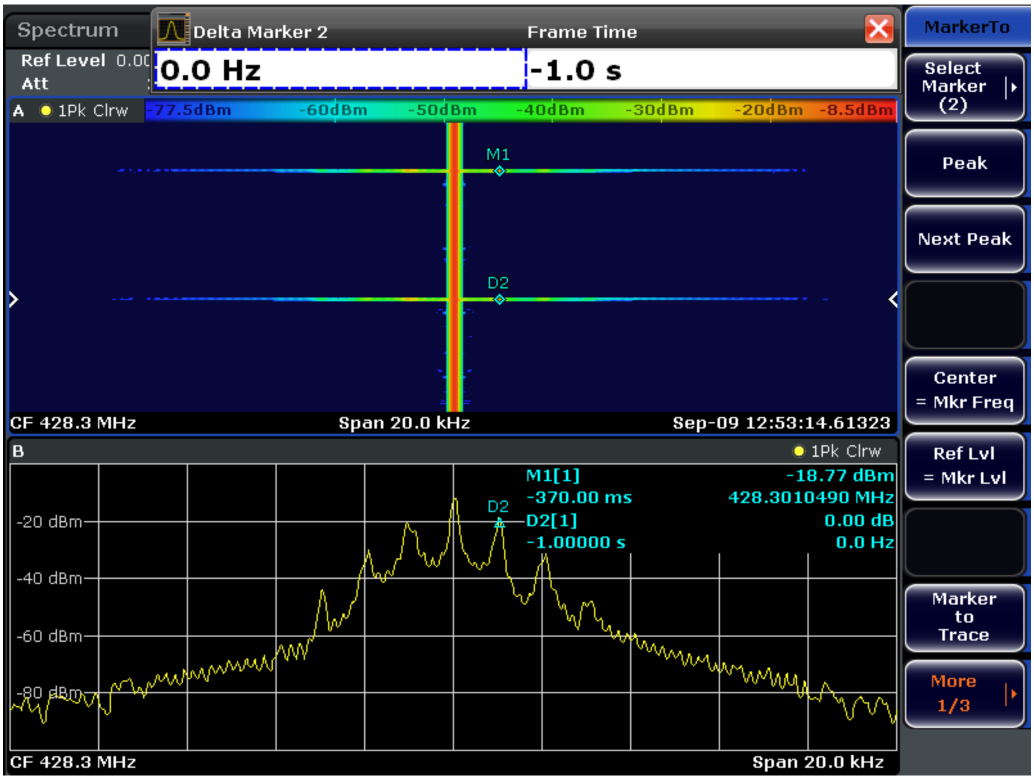

Fig. 4-5: Timing measurement with split screen spectrogram and real-time spectrum

4.3 Persistence Spectrum

Probably the best known real-time display is the so called persistence spectrum. It 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. The histogram information allows the comparison of signals on the same frequency but with different modulation.

So the persistence spectrum color codes signal statistics (probability of appearance for each displayed level - frequency pair). At the same time, it uses the fading effect to make events stay longer on the display than they actually last.

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. Fig. 4-6 shows two persistence spectra, one with a frequency agile DUT in the 2.4 GHz ISM band, and a second one in the 5 GHz band. At the moment this screenshot was taken, the signal was located on the right side of the spectrum. However, the persistence makes it clear that either the same or a different signal was located in the center part of the spectrum before.

The R&S FSW provides the unique feature to operate several real-time options sequentially and to display all options at the same time. Therefore, two or more different bands can be visualized at the same time - no matter if the bands are contiguous or not.

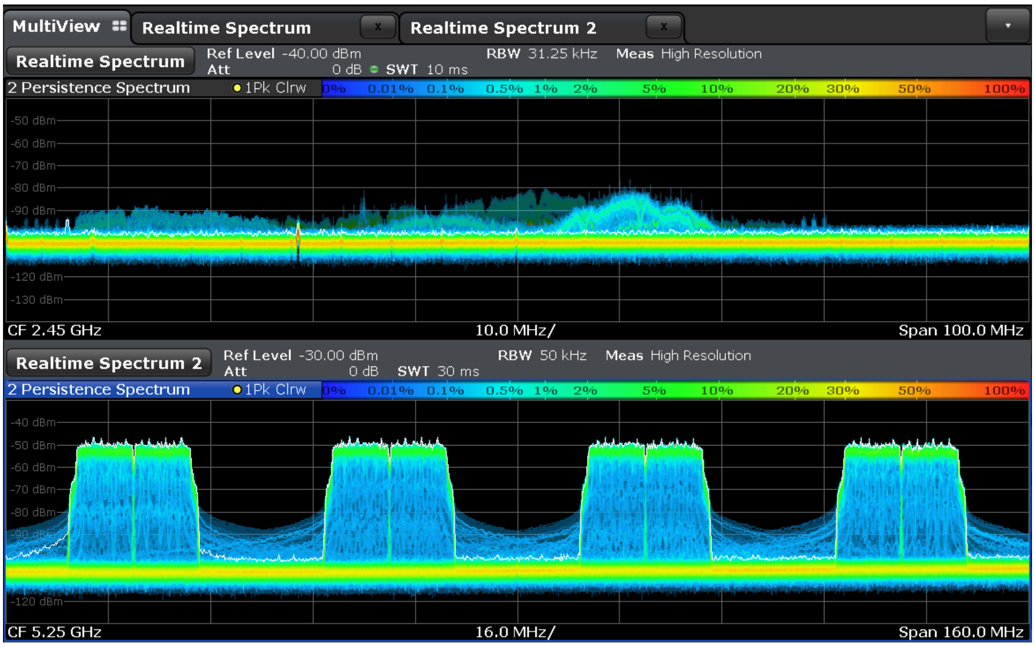

Fig. 4-6: Frequency agile DUT in the 2.4 GHz ISM band and 5 GHz band. Persistence shows signal longer than its duration

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. Fig. 4-7 shows a persistence spectrum of a noise-like signal resulting from a motor with brushes. Clearly visible is a weak WCDMA 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.

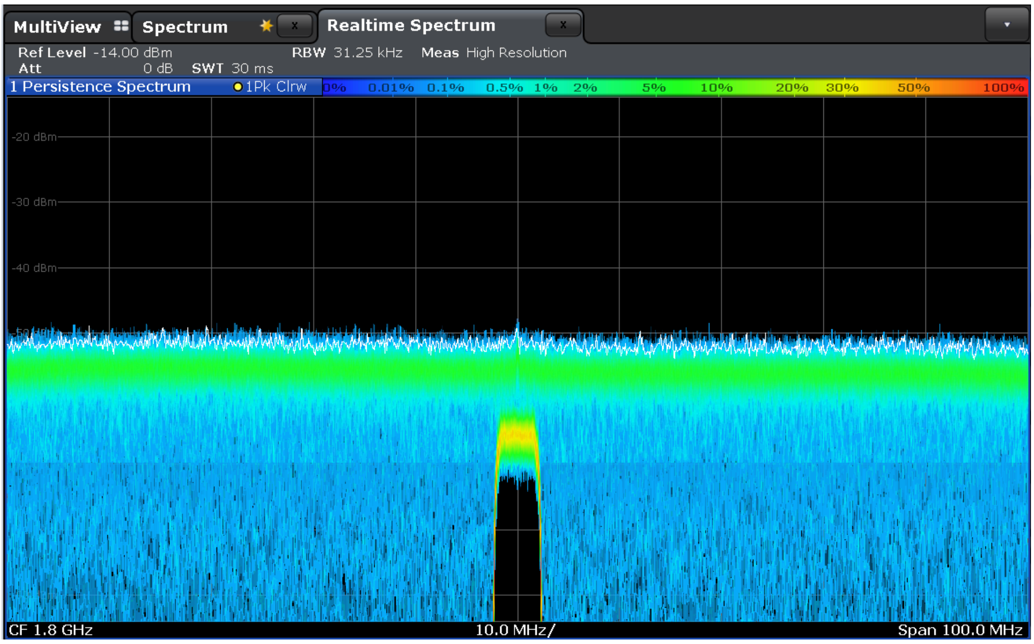

Fig. 4-7: Wideband noise-like signal covering a WCDMA signal

4.3.1 Parameters

Center Frequency, Span, RBW

These parameters specify the frequency and bandwidth setting of the real-time analyzer. They are identical throughout the real-time displays.

Persistence Granularity

A single histogram image is calculated during the persistence granularity time. The initial zero matrix with 600 by 1001 elements (600x801 on the R&S FSVR) represents 600 discrete power levels and 1001 discrete frequency steps. Assuming a full 160 MHz span and an FFT length of 1024 for the example, the R&S FSW runs an FFT update rate of 585,938 per second, thus providing an FFT every 1.71 µs. With the default persistence granularity of 100 ms, the maximum count for one cell of the matrix is almost 60,000. For each new histogram, i.e. every time the 100 ms interval is completed, the matrix is reset to zero for each element.

Fig. 4-8 shows this process with a 6 by 8 elements matrix and an FFT time to granularity ratio of 2, instead of a 600x1001 matrix and a FFT to granularity ratio of almost 60,000. Two FFTs are calculated. Both FFTs contain the same signal and varying noise neighboring the signal. Fig. 4-8 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.

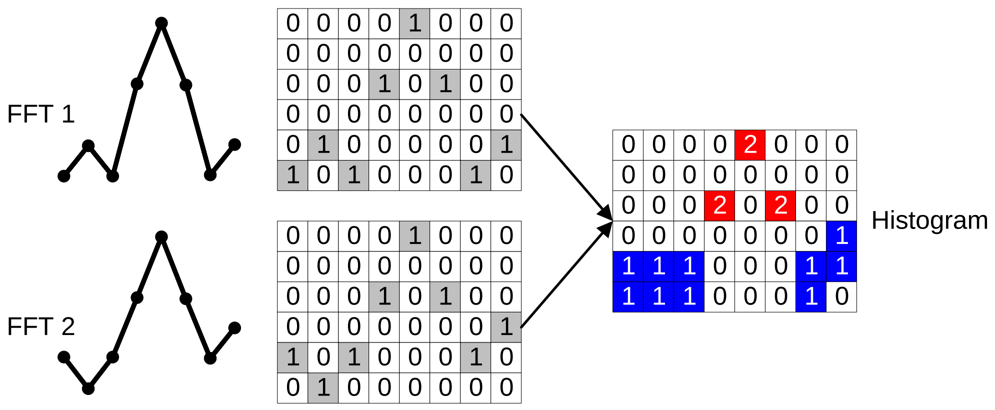

Fig. 4-8: Schematic of histogram calculation (dot style)

Persistence

The persistence parameter determines the amount of time until a trace has completely faded. During this interval, the trace loses 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 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. 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 Fig. 4-8 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 Fig. 4-9 from those in Fig. 4-8, additional “1” elements are inserted to connect the “1” in column 4 to the neighboring “1” in columns 3 and 5. Fig.

4-9 shows the vector style representation for exactly the same example that was used in Fig. 4-8 for dot style.

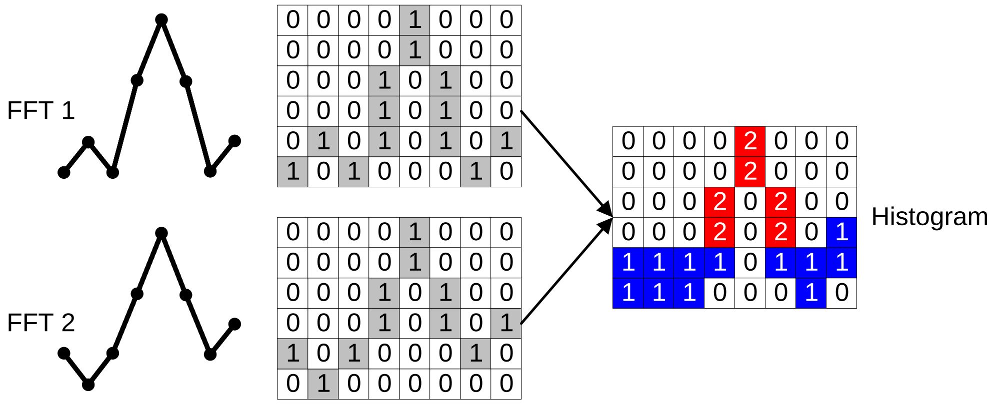

Fig. 4-9: Histogram calculation using vector style

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, Sweep Time

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 if 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

Designation	Type	Order No.
Signal- and Spectrum Analyzer, 2 Hz to 8 GHz	R&S®FSW8	1312.8000.08
Signal- and Spectrum Analyzer, 2 Hz to 13.6 GHz	R&S®FSW13	1312.8000.13
Signal- and Spectrum Analyzer, 2 Hz to 26.5 GHz	R&S®FSW26	1312.8000.26
Signal- and Spectrum Analyzer, 2 Hz to 43 GHz	R&S®FSW43	1312.8000.43
Signal- and Spectrum Analyzer, 2 Hz to 50 GHz	R&S®FSW50	1312.8000.50
FSW-B160R Real-time Spectrum Analyzer, 160 MHz (HW opt.)	R&S®FSW-B160R	1313.1668.06
Real-Time Spectrum Analyzer, 10Hz - 7GHz, 40MHz bandwidth	R&S®FSVR7	1311.0006.07
Real-Time Spectrum Analyzer, 10Hz - 13GHz, 40MHz bandwidth	R&S®FSVR13	1311.0006.13
Real-Time Spectrum Analyzer, 10Hz - 30GHz, 40MHz bandwidth	R&S®FSVR30	1311.0006.30
Real-Time Spectrum Analyzer, 10Hz - 40GHz, 40MHz bandwidth	R&S®FSVR40	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.

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

China
+86-800-810-8228 /+86-400-650-5896
customersupport.china@rohde-schwarz.com

Environmental commitment

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

Certified Quality System
ISO 9001

PAD-T-M: 3573.7380.02/02.00/CI/1/EN/

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

www.rohde-schwarz.com