



Random Interleaved Sampling (RIS)

Peter J. Pupalaiakis, Product Marketing Manager
WaveMaster Oscilloscopes
LeCroy Corporation

Introduction

Random Interleaved Sampling Mode, or RIS is a scope feature available in LeCroy oscilloscopes. It is an acquisition mode that produces waveforms at very high *effective* sample rates. While providing many specific benefits, it is important to understand what it is and how it works in order to understand the advantages and limitations of this mode of operation.

RIS mode is selectable from the timebase settings menu when the timebase is set to 20 μ s per division or below:

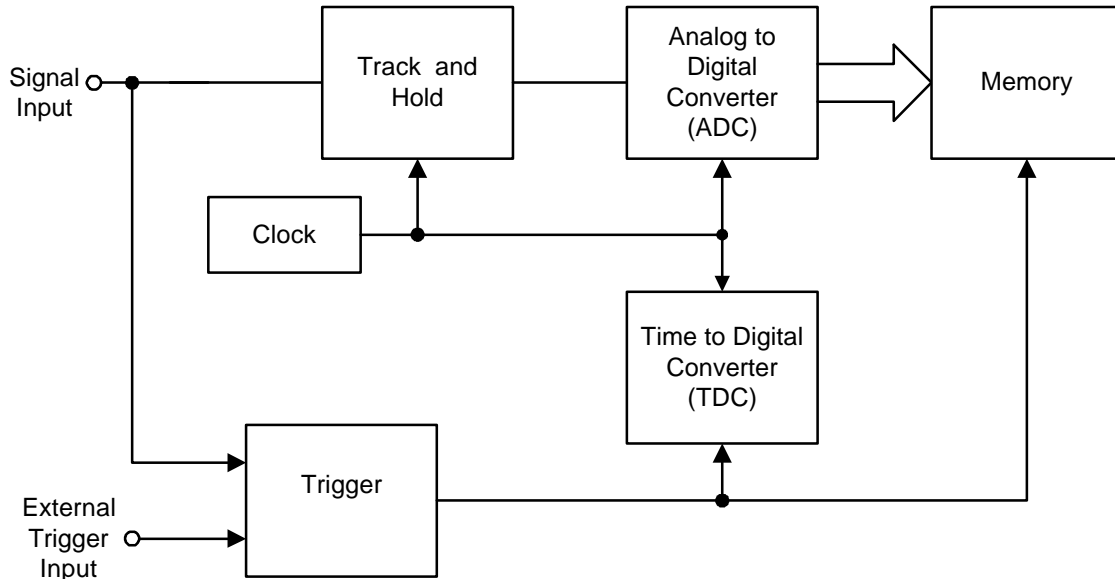


RIS is a special mode that operates differently than traditional acquisition modes and requires unique understanding to operate properly.

LeCroy scopes will not enter RIS mode unless the user specifically selects the button in the timebase menu.

How RIS Works

RIS makes use of an important piece of hardware in the DSO called the time-to-digital-converter (TDC). The TDC is the hardware that is utilized to place the trigger position correctly *between samples*. While most of the trigger and timebase hardware for trigger placement is digital and is used to determine on which sample the trigger occurred, the TDC is a combination of analog and digital hardware. It is a device that charges with constant slope from the beginning of the trigger to the next sample. Then, it discharges with constant, much lower slope, providing very fine time resolution. The discharge time is measured with a clock-counter arrangement.

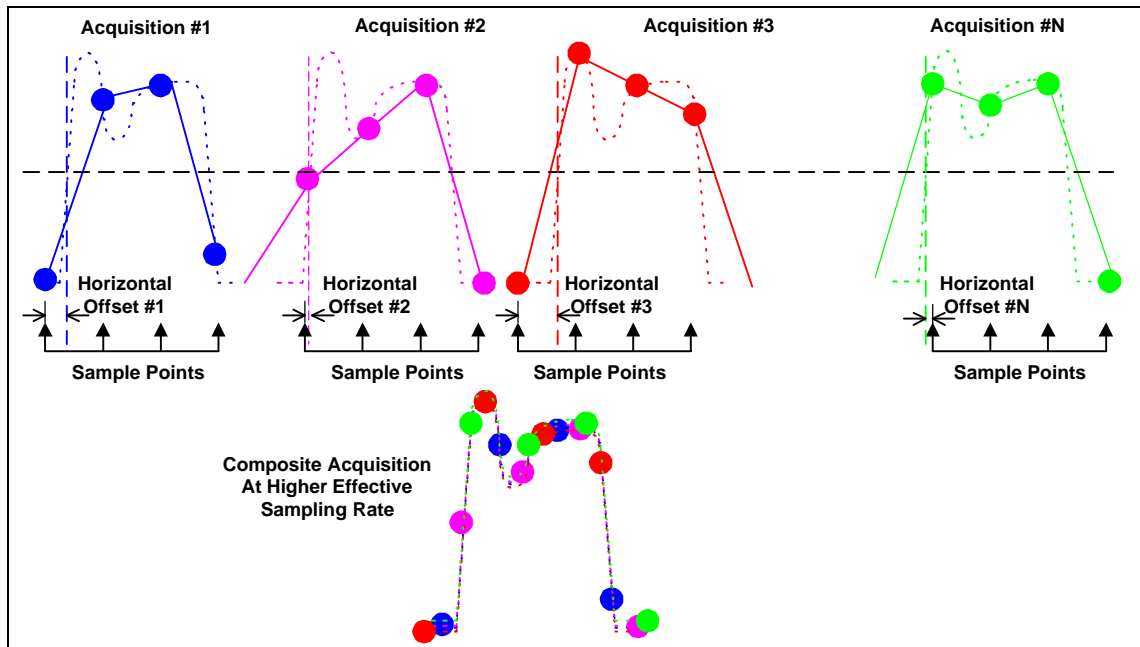


On a WaveMaster, the inter-sample time is 100 ps at 10 GS/s and the discharge time is approximately 5 us max for a trigger that occurs just after a sample is taken. The time is measured with an approximate clock rate of 100 MHz for sub-picosecond resolution.

An interesting fact is that the maximum possible RIS sample rate is dictated by the TDC resolution only!

In RIS mode, the value read from the TDC is used to *bin* successive acquisitions. Multiple acquisitions are taken of a waveform and the TDC values are used to select and arrange each acquisition for interleaving. The TDC values are completely random as long as there is no relationship between the input waveform and the sample clock of the scope.

Once the waveforms have been acquired and binned, they are interleaved together to form a single high sample-rate waveform.



There are some important characteristics of this waveform:

1. The most important - the waveform is a result of multiple scope acquisitions - in other words, the waveform segments that are interleaved to form the final waveform were acquired at far different points in time.
2. The interleaving order bears no relationship whatsoever to when the waveforms are acquired - the order of interleave depends only on the TDC value read at the time of the acquisition and is completely random.

These characteristics impose serious constraints on the use of RIS. The fact that the waveform is produced out of multiple waveform segments acquired at different, random points in time limit the usefulness of this feature to specific applications. In order to use RIS, the application must meet the following requirements:

1. The waveform must be repetitive.
2. The trigger event must be identical for every sweep.

Repetition is the key. RIS makes the assumption that every waveform segment that is acquired is the exact same analog waveform. Furthermore, it assumes that this same analog waveform is triggered at that exact same place in the waveform on every acquisition. The only difference assumed is that the sampling phase (where the trigger occurs with relation to samples) is changing.

These restrictions on repetition are serious. It means that the waveform must be repetitive, the trigger must be stable, and no non-repetitive nature of a waveform, like jitter etc. can be measured properly in this mode. Failure to honor these restrictions results in severely distorted waveforms with no usefulness whatsoever and erroneous measurements.

Since these restrictions are so severe, one might ask when and why anyone would ever use RIS. Before explaining this, a history lesson is in order.

The History of RIS

The advent of the digital real-time oscilloscope revolutionized scopes. It enabled two hugely important possibilities not possible previously. First, it enabled acquisition of a single waveform in a single-shot. It also enabled measurement and processing with a computer. Never before could a single waveform acquisition be taken and analyzed. Despite this incredible capability, this capability alone would not allow it to replace the analog scope. The analog scope is capable of repeatedly triggering at a high rate on a waveform and overlaying waveforms on the screen. For this reason, persistence mode was developed. Analog persistence creates a display similar to analog scopes. But the analog persistence displays at the time suffered in the same manner as the analog scope displays in that no processing could be performed. The digital scope that is great at processing waveforms could not do anything with the persistence display. To overcome this, RIS was born. While it does not process persistence maps, it does produce waveforms from multiple, repeated acquisitions.

Subsequent to the development of RIS, LeCroy made many strides into processing of persistence maps. Despite this, RIS mode has continued to hang around for a specific reason - the scope controls when the resulting waveform is created and processed. In the processing of persistence maps, the map itself is built up over an indeterminate number of sweeps until the scope user is satisfied. Only then does the scope user process the mean value of this map. Also, in the LeCroy scope, the screen itself determines the resolution of the persistence map so the resolution is a function of the number of horizontal screen pixels and the acquisition duration.

To handle the eye-pattern and other time varying phenomena associated with many of today's waveforms, other methods have been created to analyze these non-repetitive waveforms.

Why use RIS

RIS is useful when the trigger event is asynchronous with the sample clock and generates a repetitive waveform and any variability of the waveform is not being analyzed and when a low noise, single waveform result is needed for processing. Its usefulness diminishes or disappears when any of these requirements is not needed because:

| | |
|---|--|
| <i>If the waveform is not repetitive</i> | RIS will generate a distorted, useless result |
| <i>If the waveform is repetitive, but a stable trigger cannot be generated</i> | The acquired waveform will not be repetitive and a distorted, useless result will be generated. |
| <i>If the trigger event is synchronous with the scope sample clock</i> | The TDC value will be essentially the same for every sweep making the scope incapable of generating a RIS record. |
| <i>If the waveform is basically repetitive and the variability of the waveform is being analyzed</i> | RIS will generate a waveform that hides the waveform variability. |
| <i>If a single waveform result is not needed for processing</i> | Why not use persistence display - then you can see any variability and build up a persistence map over time. You can generate the mean of the persistence map for processing, if desired. |
| <i>If the waveform variability is being studied</i> | Understand and use the wide variety of scope tools like persistence, jitter measurements, eye pattern measurements, etc. - do not use RIS. |
| <i>If you can sample the waveform in real-time (single-shot) mode at a sufficient rate and the noise performance of the single-shot acquisition is adequate</i> | Why not use SinX interpolation to generate a higher sample rate rendition of the waveform. SinX interpolation is valid when the highest frequency component possible in the waveform is lower than half the sample rate. |

Some applications where RIS is useful include:

- Signal Integrity measurements where a repeated stimulus like a TDR pulse is applied repeatedly to a circuit. Typically, the repetition rate of the stimulus is such that the circuit reaches the same relaxed state between pulses.
- Device Characterization where automatic, highly precise measurements of timing and threshold crossings are needed.


Some applications where RIS is useless include:

- Jitter measurements
- Eye-pattern analysis
- Noise measurements


A Simple Example

An example is useful in understanding how RIS works and its limitations.


Consider the processing of a clock edge where all of the requirements for RIS usage have been met. For this example, assume that the scope's real-time sample rate is 20 GS/s and it generates a 100 GS/s RIS acquisition. Here is how a RIS acquisition might go:




Sweep 0
 Time from trigger to next sample is 0 ps.
 This waveform segment is binned as segment 0.




Sweep 1
 Time from trigger to next sample is 20 ps.
 This waveform segment is binned as segment 2.




Sweep 2
 Time from trigger to next sample is 40 ps.
 This waveform segment is binned as segment 4.



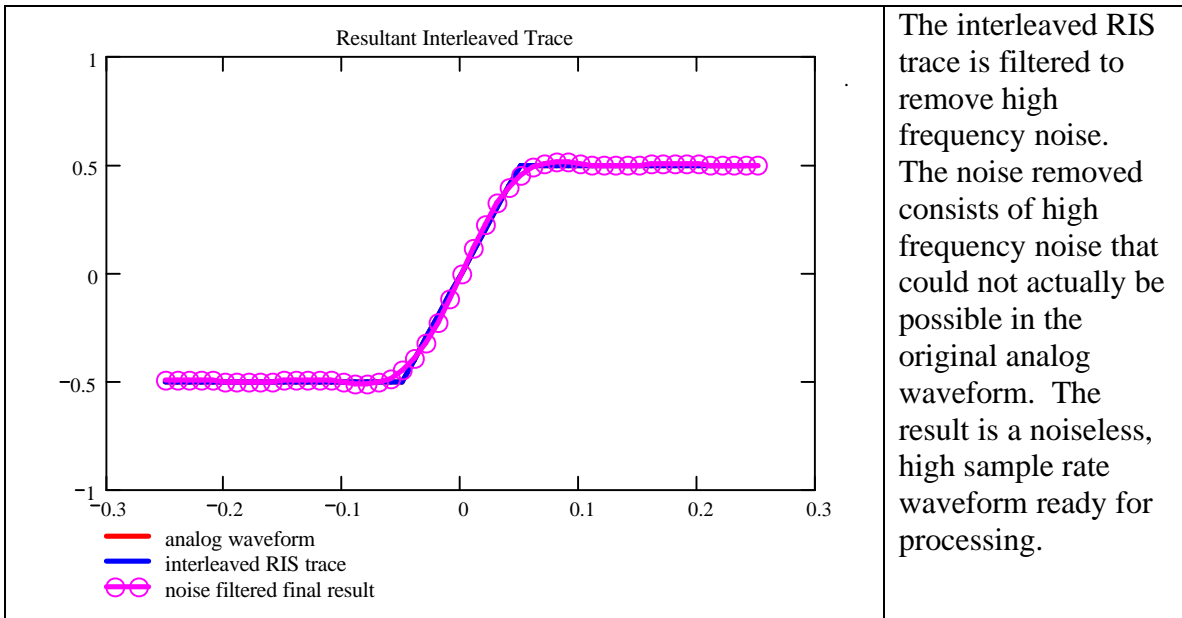
Sweep 3
 Time from trigger to next sample is 10 ps.
 This waveform segment is binned as segment 1.



Sweep 4
 Time from trigger to next sample is 30 ps.
 This waveform segment is binned as segment 3




The five segments acquired are interleaved in the order that they were binned. The resultant interleaved trace exactly overlays the analog waveform and forms a high effective sample rate rendition.




This is how a RIS acquisition should proceed. When used correctly, a very useful waveform is created that can be viewed, analyzed and processed using LeCroy's highly useful waveform processing functions.

Example of Unstable Trigger


Let's take the previous example and see what happens when RIS is used improperly, specifically when a stable trigger cannot be generated:






Sweep 1
Time from trigger to next sample is 20 ps.
This waveform segment is binned as segment 2.

The trigger delay to the proper trigger point is 39 ps.




Sweep 2
Time from trigger to next sample is 40 ps.
This waveform segment is binned as segment 4.

The trigger delay to the proper trigger point is 71 ps.




Sweep 3
Time from trigger to next sample is 10 ps.
This waveform segment is binned as segment 1.

The trigger delay to the proper trigger point is -42 ps.




Sweep 4
Time from trigger to next sample is 30 ps.
This waveform segment is binned as segment 3

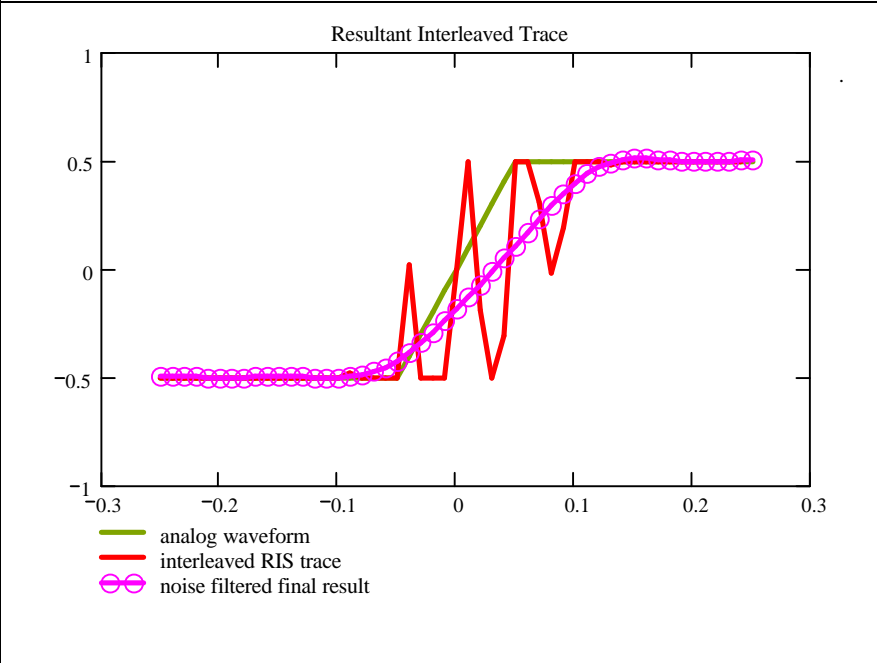
The trigger delay to the proper trigger point is 82 ps.



This is a bad situation, so let's take a look at the sweeps overlaid. This is what you would see if you repeatedly triggered the scope in single-shot mode. You would observe the unstable trigger and would not use RIS until the trigger could be stabilized.



The five segments acquired are interleaved in the order that they were binned. The resultant interleaved trace displays a non-physical waveform.

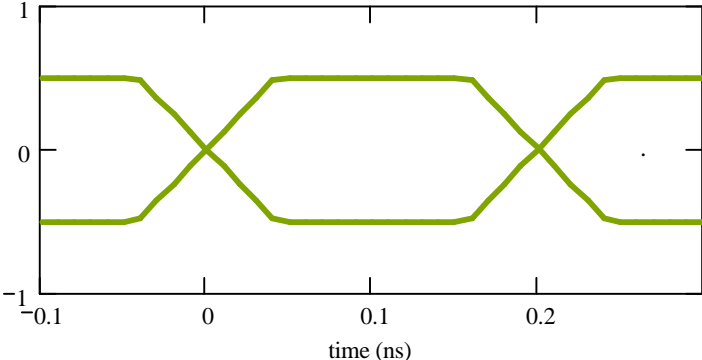
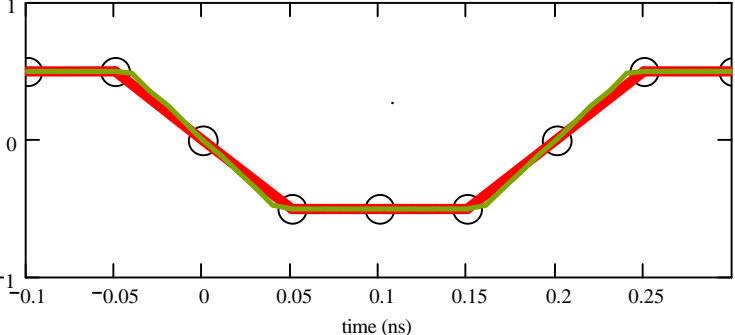
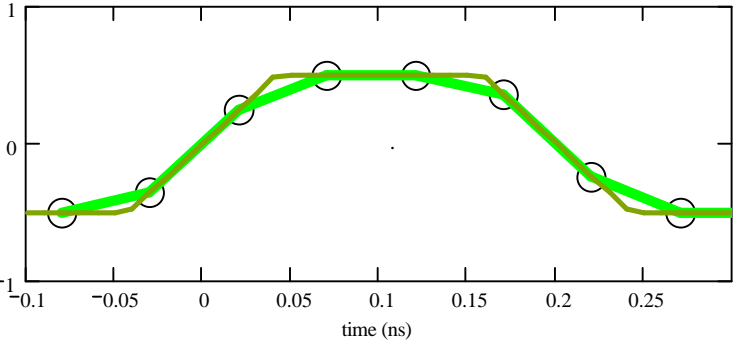



The interleaved RIS trace is filtered to remove high frequency noise. The noise removed consists of high frequency noise that could not actually be possible in the original analog waveform. The result is that the non-physical noise artifacts in the waveform are removed and the rising edge is smeared.

This example demonstrated the effect of an unstable trigger. The resultant waveform is a noise-free, artifact-free rendition of the analog waveform, but the rising edge is smeared. The resulting waveform is severely distorted and not useful. One might conclude from this that since the risetime is reduced, the RIS bandwidth is reduced. This would be an incorrect conclusion since the lower risetime results only from the non-repetitive nature of the waveform.


One More Example - the Eye Pattern

As discussed earlier, the eye pattern should not be acquired in RIS mode. Here's what happens if it is:


| | |
|--|--|
|  <p>The plot shows two waveforms crossing each other. The x-axis is labeled 'time (ns)' and ranges from -0.1 to 0.2. The y-axis ranges from -1 to 1. One waveform starts at approximately 0.5 and the other at -0.5. They cross at time 0.</p> | <p>Here is an eye pattern. The scope is set to trigger such that the trigger is stable and the waveform segments that are acquired are either positive or negative pulses.</p> |
|  <p>The plot shows a negative pulse waveform. The x-axis is labeled 'time (ns)' and ranges from -0.1 to 0.25. The y-axis ranges from -1 to 1. A red segment highlights the negative portion of the pulse between approximately 0.05 ns and 0.15 ns. White circles mark the signal levels at various time points.</p> | <p>Sweep 0 Time from trigger to next sample is 0 ps. This waveform segment is binned as segment 0. The pulse is negative.</p> |
|  <p>The plot shows a positive pulse waveform. The x-axis is labeled 'time (ns)' and ranges from -0.1 to 0.25. The y-axis ranges from -1 to 1. A green segment highlights the positive portion of the pulse between approximately 0.05 ns and 0.15 ns. White circles mark the signal levels at various time points.</p> | <p>Sweep 1 Time from trigger to next sample is 20 ps. This waveform segment is binned as segment 2. The pulse is positive.</p> |



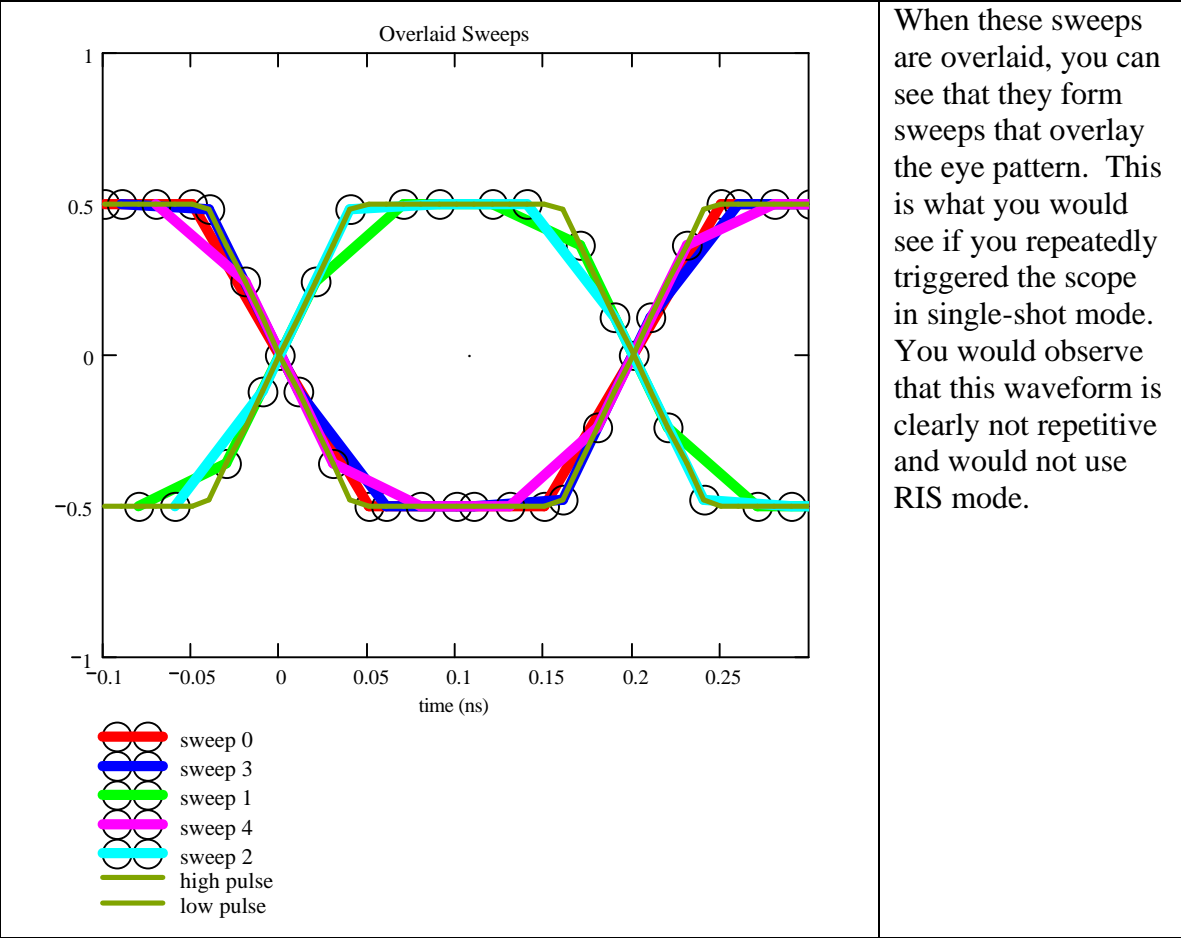
Sweep 2
 Time from trigger to next sample is 40 ps.
 This waveform segment is binned as segment 4.
 The pulse is positive.



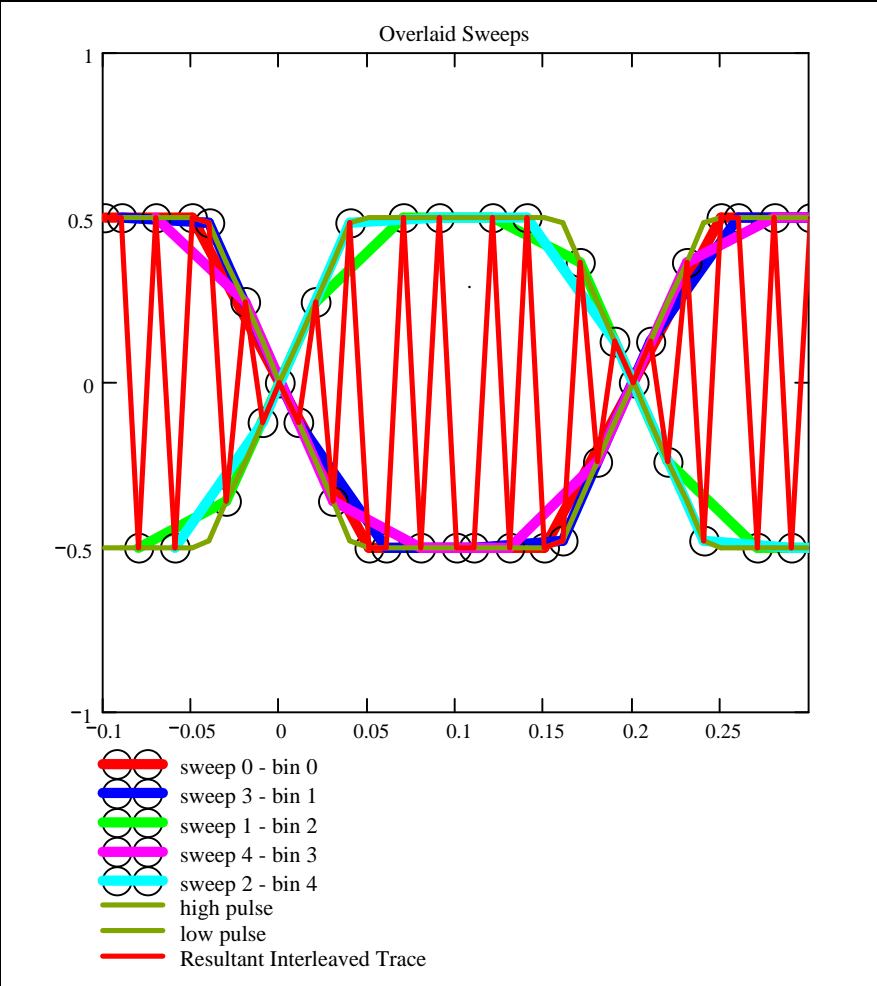
Sweep 3
 Time from trigger to next sample is 10 ps.
 This waveform segment is binned as segment 1.
 The pulse is negative.



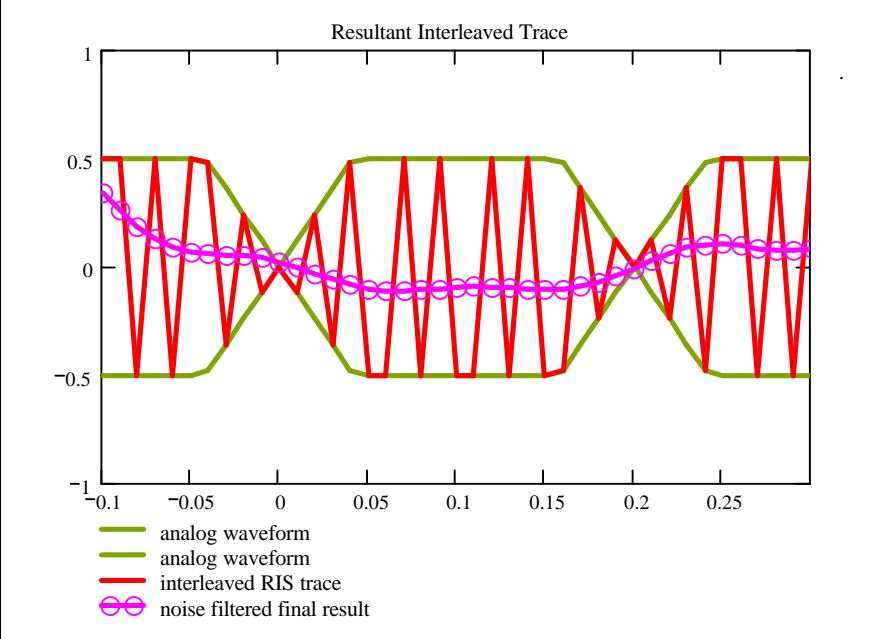
Sweep 4
 Time from trigger to next sample is 30 ps.
 This waveform segment is binned as segment 3.
 The pulse is negative.



When these sweeps are overlaid, you can see that they form sweeps that overlay the eye pattern. This is what you would see if you repeatedly triggered the scope in single-shot mode. You would observe that this waveform is clearly not repetitive and would not use RIS mode.



The five segments acquired are interleaved in the order that they were binned. The resultant interleaved trace displays a non-physical waveform.



The interleaved RIS trace is filtered to remove high frequency noise. The noise removed consists of high frequency noise that could not actually be possible in the original analog waveform. The result is that the non-physical noise artifacts in the waveform are removed and the eye pattern is destroyed.

How RIS Implementations Vary Among Scopes

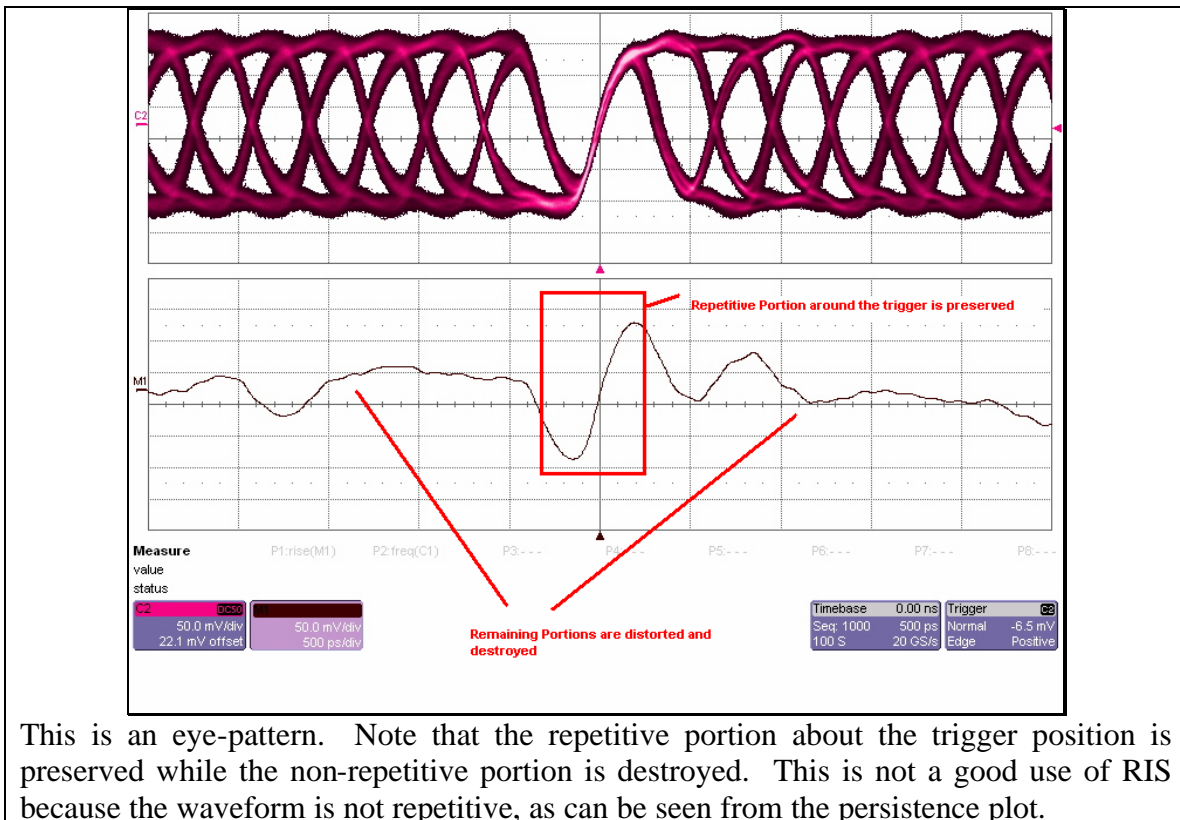
RIS is a word coined by LeCroy. Usually, on other vendor's scopes, this mode goes by names like Equivalent Time (ET), Repetitive Equivalent Time (RET) or Sequential Sampling. Some vendors have stopped offering this mode in their high-end scopes.

The main implementation difference between LeCroy and other vendors is the filtering of the final interleaved RIS trace. In other words, some vendors simply acquire the waveform segments and interleave them together. There are pros and cons to both methods: On one hand, the noise that is filtered out of the RIS waveform can only come from one source - the non-repetitive nature of the waveform or trigger and serves to degrade the rendition of the true analog waveform. On the other hand, the removal of this noise hides any non-repetitiveness that can lead to confusion unless the operation is understood.

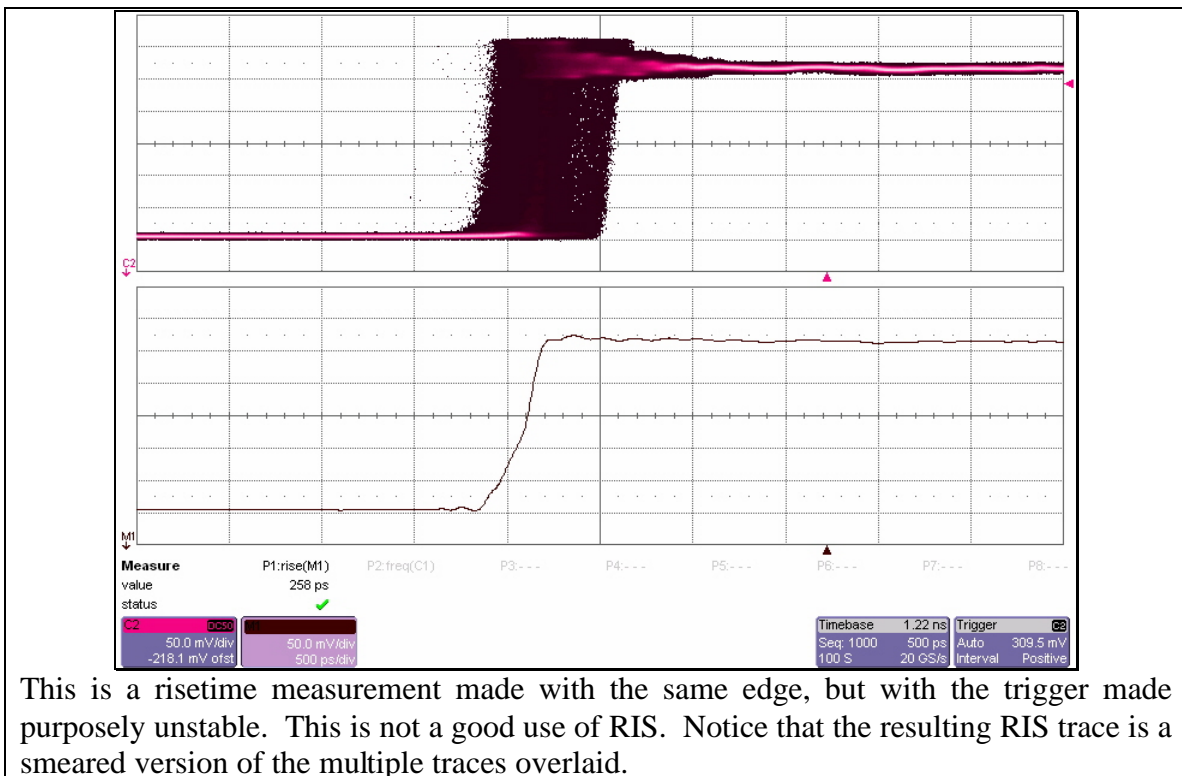
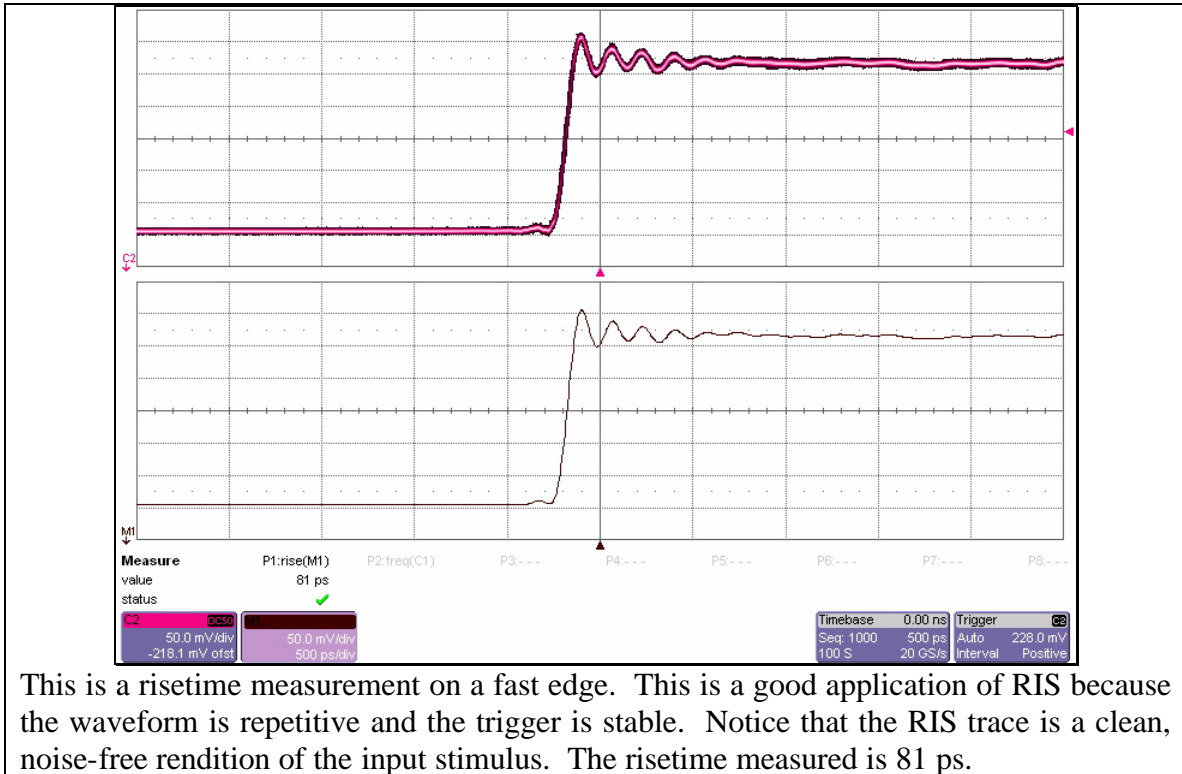
LeCroy believes that the removal of the noise and resulting higher performance and better signal fidelity is more important than the preservation of the non-repetitive variations. Confusion is avoided by understanding the scope operation. LeCroy believes this because the resulting RIS waveform is confusing even if this filtering is not applied if the waveform is not repetitive. The only solution to confusion in RIS mode is an understanding of how the mode works regardless of the scope vendor.


Simple Methods to Determine the Suitability of RIS Mode

Based on the RIS requirements stated previously, it should be easy now to determine whether RIS is useful for a particular application. Here are some actual measurements with comments:




This is an eye-pattern. Note that the repetitive portion about the trigger position is preserved while the non-repetitive portion is destroyed. This is not a good use of RIS because the waveform is not repetitive, as can be seen from the persistence plot.

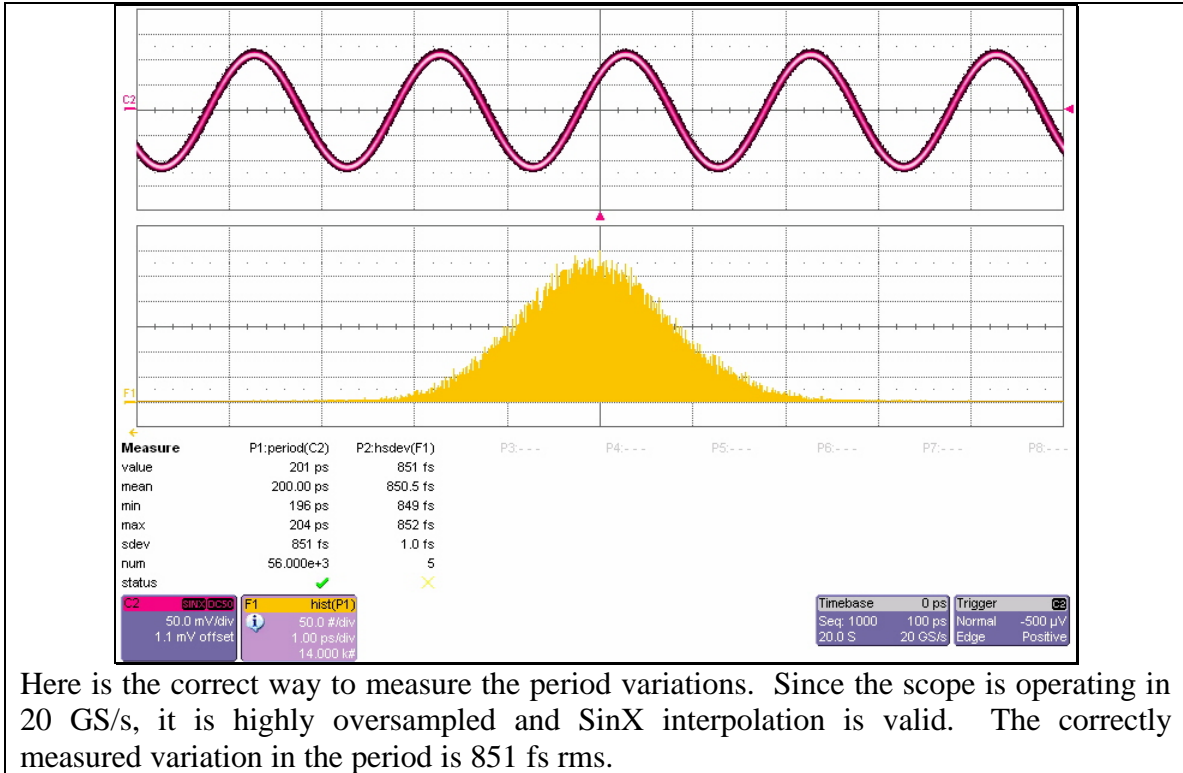




This is a frequency measurement of a sinewave. Note that the sinewave is repetitive and triggered properly. The measurement measures properly the 5 GHz frequency. *Note that not all scopes are even capable of triggering on this signal!*



Here is an example of an invalid measurement. The standard deviation of the period is measured as 273 fs - but the waveform is acquired in RIS! Despite the fact that the waveform appears repetitive and a stable trigger is generated, the jitter is a measurement of waveform variations! Since RIS hides any non-repetitive nature of the signal, it should certainly not be used to measure variations in a waveform. It is extremely important to remember this fact.

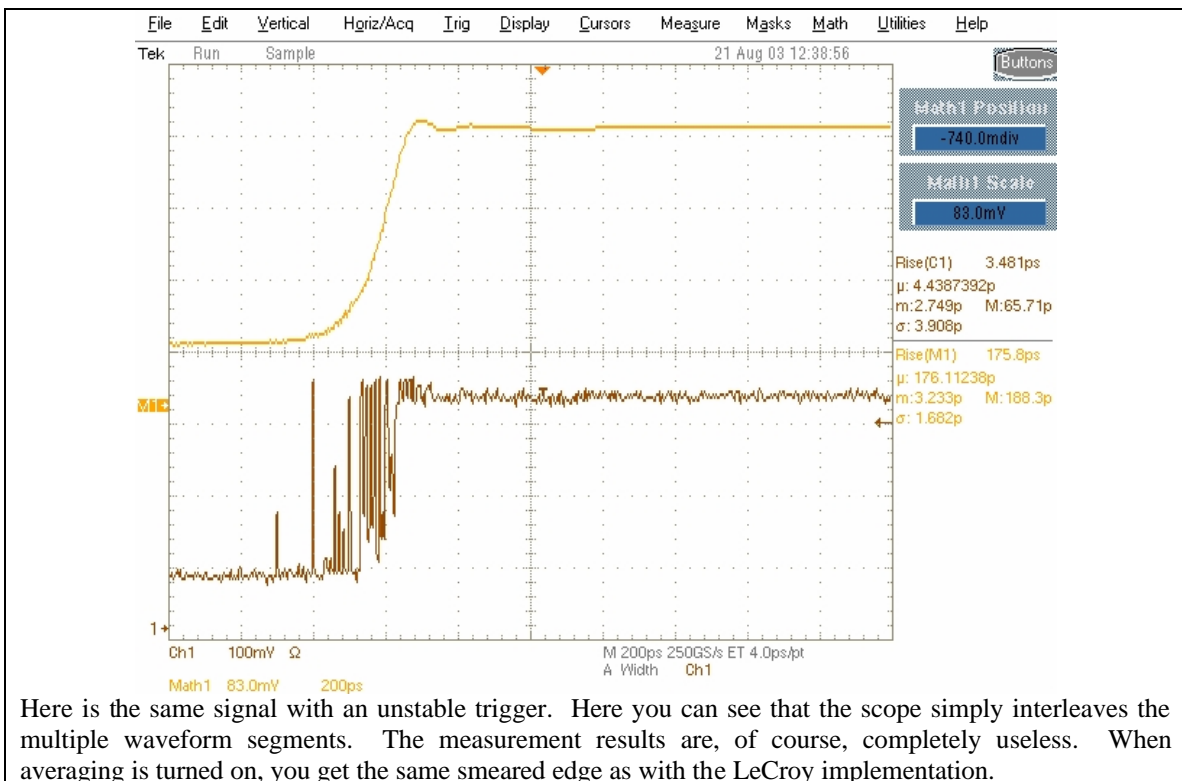
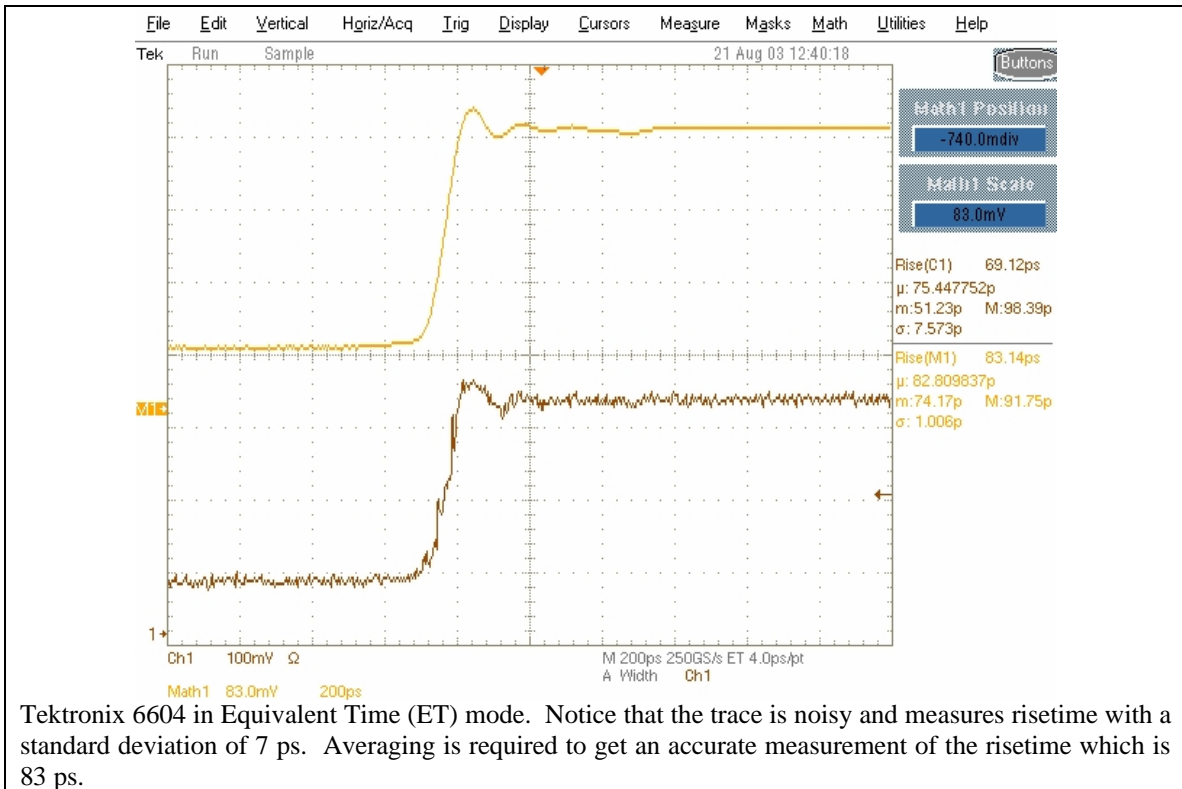


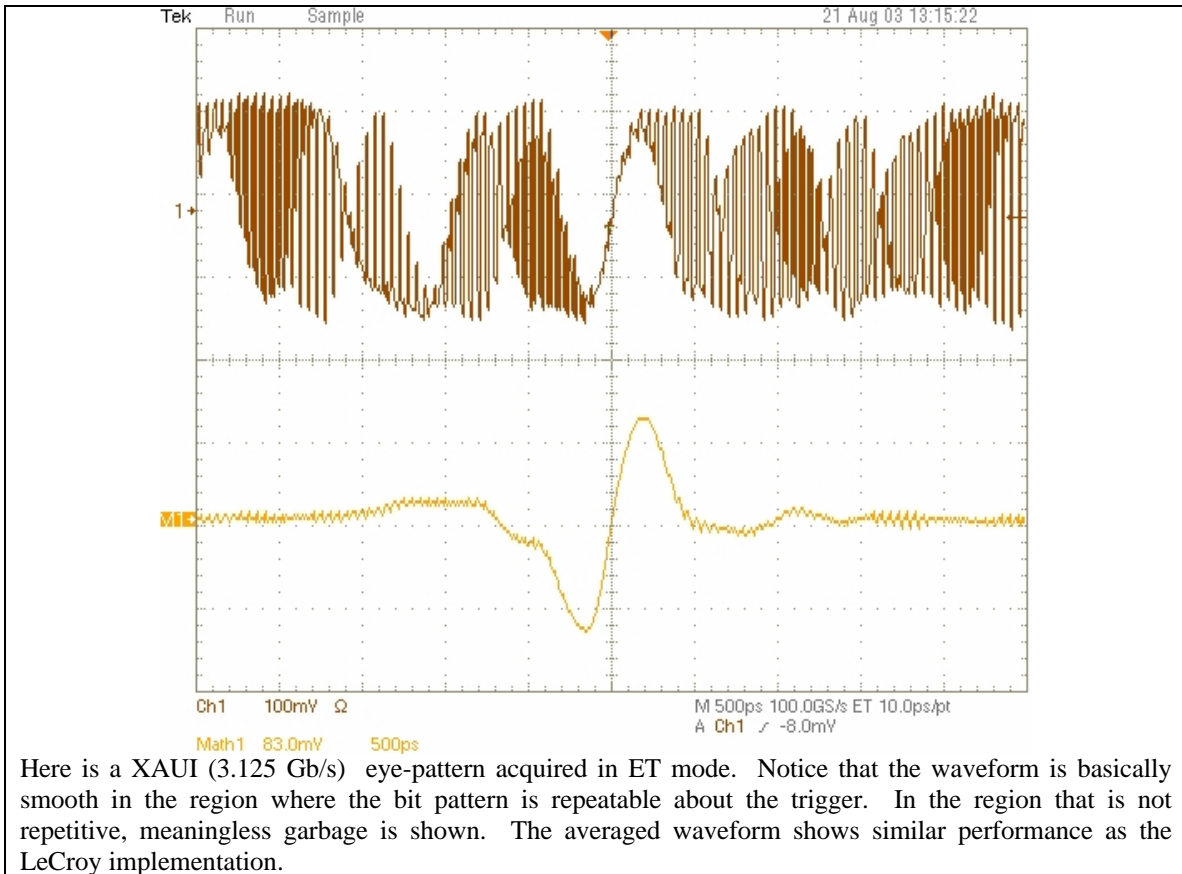
Some Comparisons

It is useful to compare the implementation for the LeCroy WaveMaster DSO to the other top competitor's products - the Tektronix 6604 and the Agilent 54855A.

The Agilent 54855A is easy - Agilent does not offer this mode or any equivalent in this scope.

Here are some screen shots showing the Tektronix 6604 implementation:





Summary

RIS is a fundamentally different acquisition mode than the single-shot acquisition mode most commonly used within the scope. It is used to generate a high effective sample rate waveform for processing. It is useful only when all of the following conditions are met:

- the trigger event is asynchronous with the sample clock
- the trigger generates a repetitive waveform
- Variability of the waveform is not being analyzed
- A low noise, single waveform result is needed for processing.

It should not be used if any of these conditions are not met.

Improper use of RIS mode will result in distorted, unusable waveforms with incorrect measurement results.

Used properly, the result is an incredibly clean waveform that provides highly precise measurement results.

RIS is particularly useful for signal integrity measurements and device characterization.

RIS is should not be used for jitter measurements, eye-diagram analysis, and noise measurements.