White Paper

Jitter and Wander in SyncE

Testing Jitter and Wander to ITU-T Standards

As a significant technology for digital transmission in the mobile backhaul, the development of Synchronous Ethernet (SyncE) has added physical-layer clock distribution and synchronisation to Ethernet. However, it is not without its problems. As with other transmission technologies, jitter and wander threaten to jeopardize the quality and reliability of high-bandwidth data and video services—especially over mobile networks.

Agreed international standards for jitter and wander are clearly essential for equipment and network interoperability. And the need for measureable jitter performance has recently prompted the ITU-T to add performance limits for networks and Synchronous Ethernet Equipment Clocks (EECs) to their G.8262 Standard. The existing IEEE 802.3 standard for Ethernet remains as a complementary jitter standard, but is only relevant for testing at component level, not on network equipment or networks. The ITU-T standards now contain the relevant jitter and wander limits and masks governing network performance, reinforcing that proper testing to G.8262 is both vital and unavoidable.

Measuring wander gives crucial information on synchronisation quality, especially vital in mobile backhaul. Such wander measurements are equally important irrespective of whether clocking is carried at the physical layer in SyncE, or at the packet layer in IEEE 1588v2. In addition to establishing standards compliance, wander results can aid the diagnosis of clocking problems.

With this in mind, Spirent has created this Technical Brief for engineers in the lab and those in the field who may not be familiar with the complexities of jitter and wander, and also to dispel some misunderstandings and avoid misinterpretation about measurement setup, implementation, and compliance to standards.

Jitter and Wander in SyncE

Testing Jitter and Wander to ITU-T Standards

What is Jitter and Wander?

Jitter and wander are variations in telecom signal timing instants caused by many diverse mechanisms, ranging from component behaviour and noise to inherent network action. Any digital network device that processes data or clock signals can potentially mistime the transmitted signal and induce jitter and wander effects in the network.

Jitter is also used more generally as a term, for example to describe Ethernet frame or IP packet inter-arrival time variations. This technical brief only describes the specific case of physical-layer signal timing variations.

Every network device will generate some degree of jitter, and therefore the network has to operate in its presence. When jitter becomes excessive, bit errors are produced and network performance suffers. Even if jitter itself doesn't cause bit errors, it reduces the noise margin of the system making it more prone to errors. The effects of jitter viewed on an oscilloscope are shown on right.

High-speed variations (phase variations above 10 Hz) in signal timing through a system are called jitter. Low speed variations in a signal or clock (10 Hz down to micro-Hertz) are termed wander. The 10 Hz breakpoint is a long-standing convention in telecoms.

Though jitter and wander are closely related, their effects on networks and devices are quite different. Jitter directly causes bit errors at device inputs by the eye-closure mechanism shown above. These bit errors then result in dropped packets. Wander does not impact the input stage but affects the ability of the internal clocking function to lock to and track the up-stream reference. Wander therefore impacts the distribution of clock references across networks.

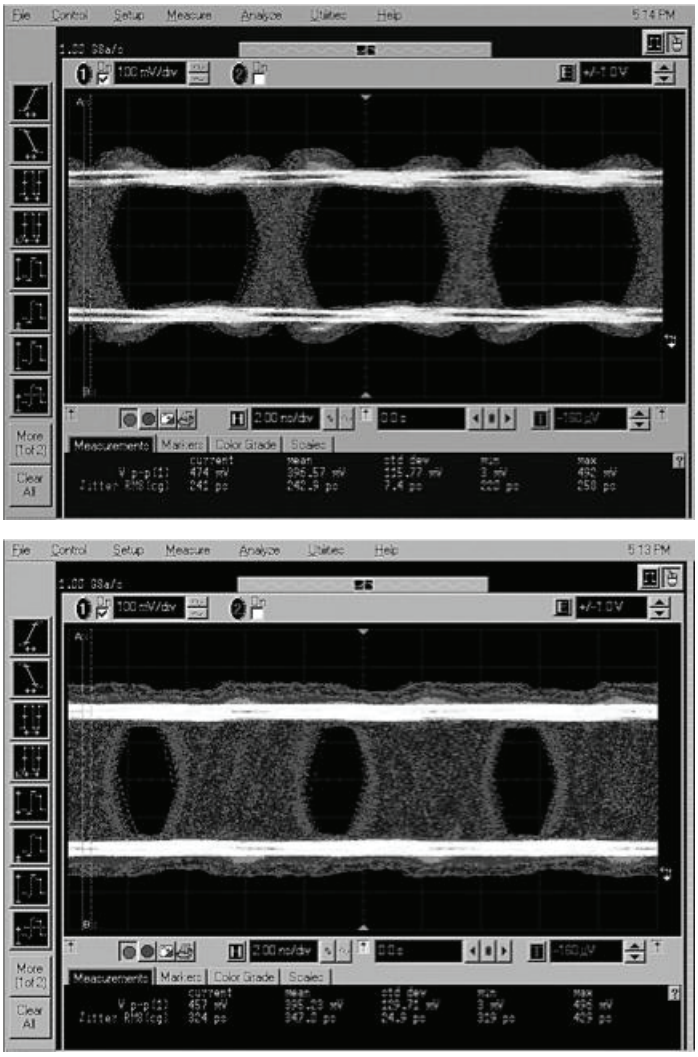

Figure 1. Jitter causes eye closure (shown above) and prevents correct sampling

Causes of Jitter and Wander

Jitter is always present at the output port of any network element (NE), even if it originates a signal or an entirely jitter-free digital signal or clock is applied to its input. This is known as (device) jitter generation, (device) output jitter, or intrinsic jitter. Intrinsic jitter arises from:

- Clock oscillator behaviour–noise, spurs, crosstalk and drift
- Pattern-dependent delay in scramblers and encoders
- Laser and modulator pattern dependency

Wander can arise at many points in a network, and numerous standards have been developed to specify wander generation, wander transfer, and wander tolerance limits. The figure below shows the synchronization architecture for SyncE, the performance of which is covered by ITU-T G.8262.

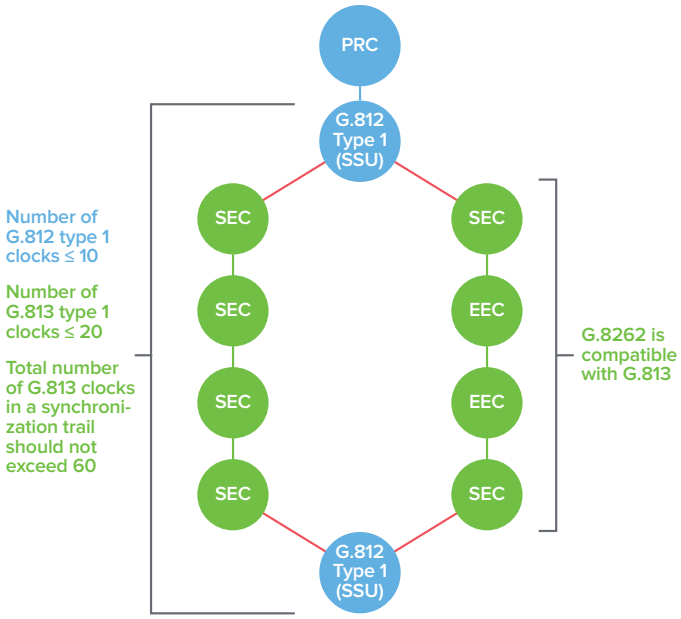

Figure 2. Synchronization architecture for SyncE

In each of the slave clock devices in a network, the device is locked to the frequency of the Primary Reference Clock (PRC) by recovering a clock from the incoming physical layer signal. The action of recovering and regenerating the clock reference is a major source of wander in a network. Wander is also caused by very slight differences in reference clocks between network segments, or by slow changes in the relative phase of two clock signals due to climatic temperature changes. It can also be caused by very low-frequency phase noise in a clock oscillator. Categories of noise include:

- White phase modulation, commonly associated with amplification stages
- Flicker phase modulation, commonly associated with amplification stages
- White frequency modulation, commonly found in passive resonators, for example Caesium Standard
- Flicker frequency modulation, which may be related to physical resonance
- Random walk frequency modulation, which may be associated with shock, vibration or temperature

Phase transients may occur when, for example, a slave clock loses its input reference and a back-up only becomes available after some period of time. During the switchover period, an accumulation of phase error can build up.

In addition, frequency drift or offset may be caused when a slave clock loses all external reference sources and enters a holdover mode. In this mode, the clock continues to operate using calibration data built up during normal locked operation. Over a period of time, the slave clock may eventually develop frequency drift. Other factors such as ageing, temperature and power supply variations may also contribute.

ITU-T Standards

To ensure interoperability between devices and minimize signal degradation due to jitter and wander accumulation across the network, international standards have been determined.

ITU-T Standards (G.8261 and G.8262) specify jitter and wander performance limits for networks, NEs and EECs and are usually taken as a minimum requirement. For example, ITU-T G.8262 specifies the maximum amounts of peak-to-peak output jitter (within a defined bandwidth) permitted from an EEC. This is to ensure that the amount of jitter never exceeds the specified tolerance level for subsequent EECs. In other words, if the jitter level is excessive, the EEC's input circuitry (clock recovery circuits etc.) may not have been designed or qualified to work error-free under such conditions.

SyncE Jitter Specifications in G.8262
(in relation to IEEE 802.3 and existing SONET/SDH jitter specifications)

It is a primary requirement of SyncE that it can be connected to any Ethernet interface defined to the requirements in IEEE 802.3. SyncE differs from non-sync Ethernet only in its additional requirement to transfer network timing. Apart from network timing transfer, SyncE must be fully interoperable with non-synchronous Ethernet (which presently has a huge installed base).

ITU-T G.8261 Annex A, sub-sections A.2 & A.5, set out the interoperability requirements between Ethernet interfaces designed to IEEE 802.3 and Synchronous Ethernet interfaces designed to ITU-T G.8262. (Refer to G.8264 Sections 9.4.1 & 10.1 for additional discussion on the interoperation of Ethernet and SyncE). In summary, the interoperability requirements separate the "Port Level" from the "Clock Level". The Port Level concerns the extraction of the data and in this case, there is no difference between the two interface definitions and full interoperability is required. The difference is restricted to the Clock Level where it is required to recover the network timing from the interface and make available this recovered timing as the reference for other synchronous outputs from the network equipment it resides in. The timing can be propagated over a chain of 20 clocks as in SDH, and this is the main difference to non-sync Ethernet where all ports are timed from free running oscillators.

When it came to define the Jitter specification of SyncE interfaces, the need for interoperability was an essential requirement for the ITU-T Study Group 15, Q.13 committee (assigned responsibility for this specification). IEEE 802.3 defines a procedure for testing the high-frequency jitter performance of Ethernet interfaces, both in terms of tolerance and noise generation. The methodologies used in 802.3 are completely different to the traditional approach taken by ITU-T for defining jitter performance parameters. As it was not possible to translate one approach into an equivalent specification in the other approach, given the overriding need to ensure interoperability, it was decided only to specify the jitter up to the jitter frequency that would not overlap with the 802.3 specification and apply 802.3 specification for the remaining part of the spectrum. This why there is no High-band Jitter Generation specification in G.8262 (i.e. why it does not have a structure similar to SONET/SDH jitter specifications).

Jitter Tolerance for 1 Gb/s SyncE interfaces is defined up to 50 kHz, and for 10 Gb/s SyncE interfaces up to 40 kHz. To have extended the specification to beyond this could lead to inter-operability problems in the field between Ethernet and SyncE interfaces.

Inter-operability is a requirement stipulated in G.8261 Annex A, mainly A.2 and A.5, which defines the Synchronous Ethernet architecture. The {ITU-T, SG15, Q.13} committee used previous specifications based on G.813 to guide the development of the jitter requirements, so that all the simulation work done for SDH clock chains could be reused for SyncE.

The relevant quotation from Jean-Loup Ferrant, Rapporteur SG15, Q.13 is: "It was agreed by the Q.13 delegates that the specification should not overlap the requirements set out in IEEE 802.3, but simply exist as an additional specification due to the additional feature raised by clock chaining, that being the transfer of network timing".

The additional and non-overlapping nature of testing to G.8262 is further supported by the implied sequence of testing: IEEE 802.3 Jitter tests are done on both Ethernet and SyncE components by component suppliers, and the ITU-T G.8262 tests are subsequently done by the SyncE system vendors and network operators.

Jitter Specifications for 1G Electrical Ethernet

G.8262 does not specify jitter limits for this interface type, listing the matter in table 6, note 2 as "for further study". However, the 4-lane, multi-level physical layer scheme used for 1G electrical Ethernet has a clock rate of 125 MHz (versus 1.25 GHz for 1 Gb/s optical Ethernet), which implies a 10:1 ratio for jitter parameters. Spirent has therefore implemented 1G electrical Ethernet jitter testing using a scaled jitter tolerance mask to an upper frequency of 5 kHz and a jitter generation measurement filter of 250 Hz to 1 MHz using this implied ratio.

Test Equipment

ITU-T also specifies benchmark functionality and key parameters (range, resolution, and accuracy) for test equipment. Figures 3 and 4 below show the measurement block diagrams of jitter and wander test equipment, and identify the main functions covered in the ITU-T O.174 recommendation.

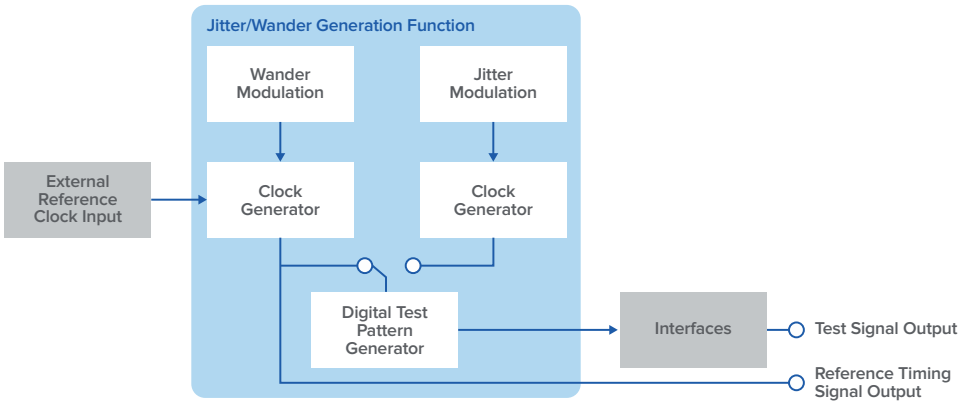

Figure 3. Jitter/wander generation function (from ITU-T 0.174)

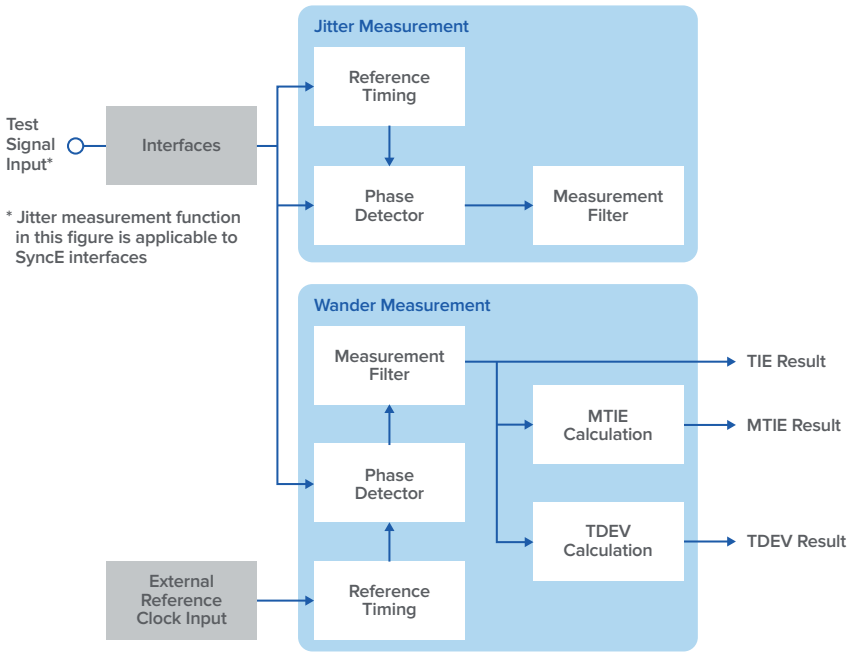

Figure 4. Jitter/wander measurement function (from ITU-T O.174)

Jitter measurement is achieved by phase-detection on the recovered clock. The phase detector produces a value that is proportional to the phase difference between the jittered clock signal and the un-jittered internal reference timing signal. The demodulated phase information is filtered in the required bandwidth, and peak-to-peak and RMS jitter values measured. It is especially important to measure pk-pk jitter correctly using advanced digital signal processing since jitter transients tend to be the most disruptive and don't register significantly in an RMS result.

It is also conventional to express jitter values (jitter amplitude) in the normalized term Unit Interval (UI), where one UI is equal to a bit period. This is more convenient than using a direct time value. As an example, for 1 GE the bit period is 800 pS, and so, a jitter amplitude of 400 pS would be expressed as 0.5 UI.

Wander analysis requires the collection of a large number of phase measurements (in this context, called Time Interval Error—TIE) over a relatively long period, followed by complex calculations to derive analysis parameters such as MTIE and TDEV defined in the ITU-T standards and discussed further below. These parameters can help to analyze the performance of reference clocks and synchronization networks and hence pinpoint the possible sources of wander. Note that (1) wander amplitudes are typically not related to bit period and so are measured in normal time units, and (2) over long periods the total TIE amplitudes can be very large.

For both jitter and wander measurement, where the reference clock comes from is important. In the case of jitter, the reference is recovered from the measurement signal using a narrowband clock recovery that removes any jitter present. This is in contrast to the clock recovery for the measurement signal, which is wideband and passes the jitter through to be measured at the phase detector. For wander, it would be possible also to recover a reference clock, but the time taken to do so to achieve the necessary stability would be impractical. So the standards prescribe an external wander reference clock. Such a clock would generally come from upstream equipment (or a house standard feeding it) in a lab situation, or from GPS or independent network timing signals in a field situation.

Jitter Measurements

Since the focus is on equipment interfaces, the two significant specifications focus on how much jitter is present at an output interface and how much can be tolerated at an input. The standards for SyncE equipment for jitter follow previous, similar standards for SDH/SONET.

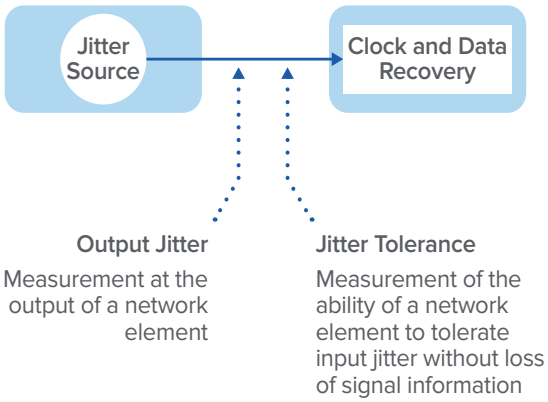

Figure 5. Measuring SyncE jitter

Output Jitter
(also known as intrinsic jitter or jitter generation)

Output jitter is a measurement on an output from a system or device. It is specified in unit intervals and only as a peak-to-peak value in the G.8262 standard. Peak-to-peak jitter relates most directly to the effect on performance.

The standards also specify a precise measurement bandwidth in terms of high-pass and low-pass filters with defined cut-off frequencies and roll-offs. The high-pass cut-off (e.g. 20 kHz for 10 Gb/s Ethernet) is determined by the narrowest cut-off frequency of any timing element expected in the transmission system, and the low-pass cut-off (e.g. 80 MHz for 10 Gb/s Ethernet) is an arbitrary value determined to be sufficiently wide to include all significant jitter components. It is specified to be consistent with predecessor standards for SDH/SONET at equivalent bit rates.

In the case of SDH/SONET output jitter measurement, results are strongly influenced by the data content of the signal being measured—in particular the presence of unscrambled overhead bytes. Test results can vary widely between a structured signal and those for a simple PRBS (Pseudo Random Binary Sequence). This has led to difficulties in comparing results between different test sets and test methods. Ethernet, in contrast, uses a more comprehensive coding arrangement that minimizes such data pattern dependencies. Nevertheless, the lessons learned from the design of jitter measuring receivers in the SDH/SONET domain remain critical in SyncE jitter testing.

Jitter Tolerance

Jitter tolerance is a measurement to check the resilience of equipment to jitter applied to its inputs. A test instrument generates a signal with added sinusoidal jitter and applies it to the device under test (DUT). To determine the actual jitter tolerance characteristics of a DUT, at each jitter frequency, the amplitude of the jitter is increased until transmission errors are detected (for example, by observing the onset of dropped packets). For Standards conformance testing, a specified level of input jitter is generated and error-free operation checked. It is recommended this procedure is repeated using jitter amplitudes above the mask defined in the Standard to check the margin of performance of the DUT. In the real world, jitter is unlikely to be sinusoidal, but it is practical to test this way and accepted across the industry as a valid approach which gives repeatable results. It also allows results for different systems to be compared and for system specifications to be written, usually in the form of a jitter tolerance mask.

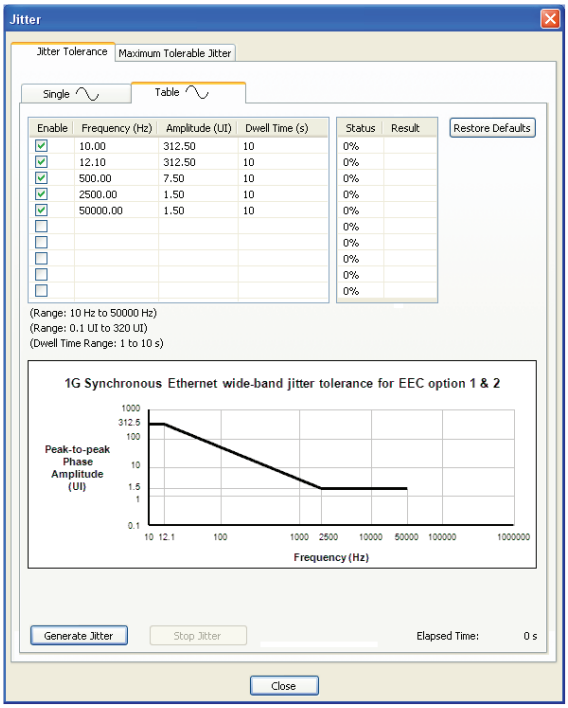

Figure 6. Jitter tolerance measurement application on the Spirent Paragon-X

Wander Measurements

A completely different set of measurements is used to characterize wander. While jitter is normally measured with reference to a clock extracted from the data signal, wander is measured against an external reference clock. Ideally, the external reference should be a direct connection to the master reference clock. The fundamental measurement is of Time Interval Error (TIE). This represents the time deviation of the clock signal under test relative to the reference source. An example of a TIE wander measurement is shown below.

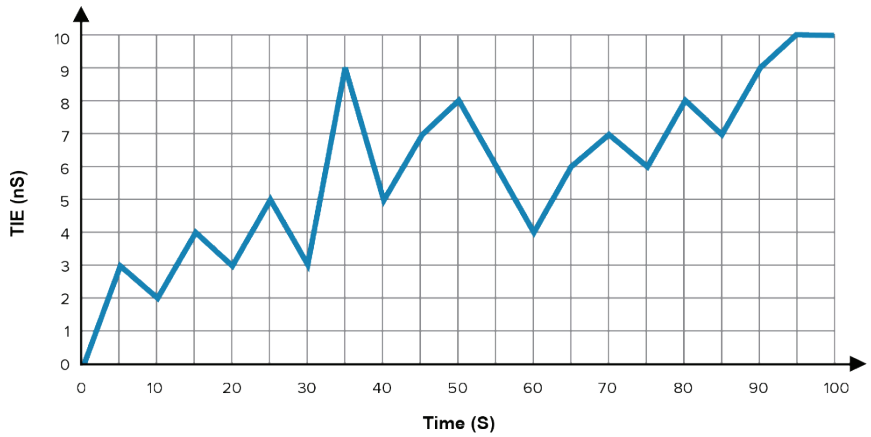

Figure 7. TIE wander measurement. The overall increasing TIE value in this example is due to a frequency offset.

Several results requiring significant computation can be calculated from TIE according to the ITU-T G-series recommendations:

- MTIE (Maximum Time Interval Error): The peak-to-peak variation of TIE within a defined observation interval .
- TDEV (Time Deviation): A measure of the spectral content of wander as a function of the observation interval .
- Frequency Offset: A measure of the degree to which the clock frequency deviates from its ideal value.

These wander results are obtained by processing the TIE samples from the measurement equipment. This can be done during the measurement or accumulated results post-processed at a later date and/or in a different location.

MTIE

The gross TIE performance is also more rigorously quantified by a calculation from the TIE plot known as MTIE (maximum time interval error). A variable window or „observation interval is moved through the TIE values, and the measurement records the highest pk-pk value of TIE for each observation interval. The derived plot of these latched maxima, as a function of the size of the observation interval, gives the MTIE result. This is plotted as a graph of MTIE against observation interval. An MTIE result always increases monotonically with observation interval as a consequence of the latching process.

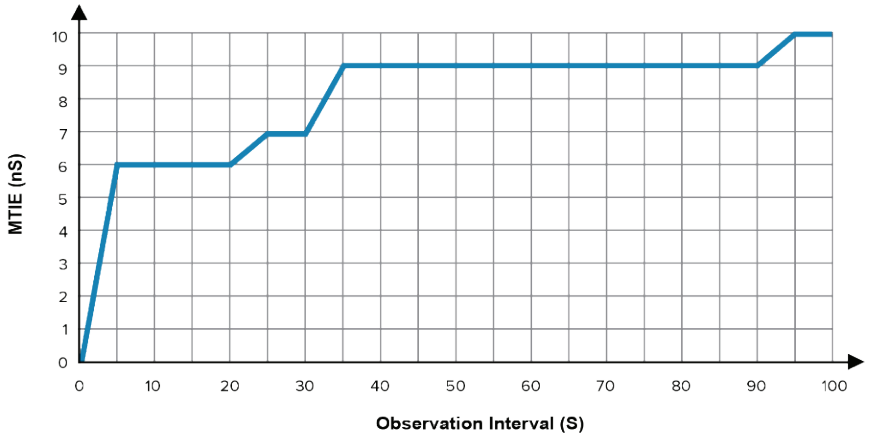

Figure 8. MTIE based on TIE values in Figure 7.
The steep rise to 6 ns is caused by the transient at the 30 second point in the TIE data.

TDEV

The short-term characteristic (that is, phase noise) performance of timing signals is particularly significant. Qualitatively, it is difficult to analyze the noise on a TIE plot. Also, the MTIE result is focused on revealing the long-term behaviour and stability of a clock and is insensitive to noise.

Therefore, a measurement result known as TDEV (time deviation) is calculated from TIE. TDEV is effectively a spectral analysis of the TIE plot as a function of observation interval, and is calculated by taking the RMS value of TIE after applying a filter bandwidth that varies inversely with the observation interval. Even more so than in MTIE, the calculation is complex and absolutely requires automation. Once again, results are compared against masks, or benchmarked. To get accurate TDEV results, G.8262 specifies that the source TIE measurement period must be at least 12 times the maximum observation interval required.

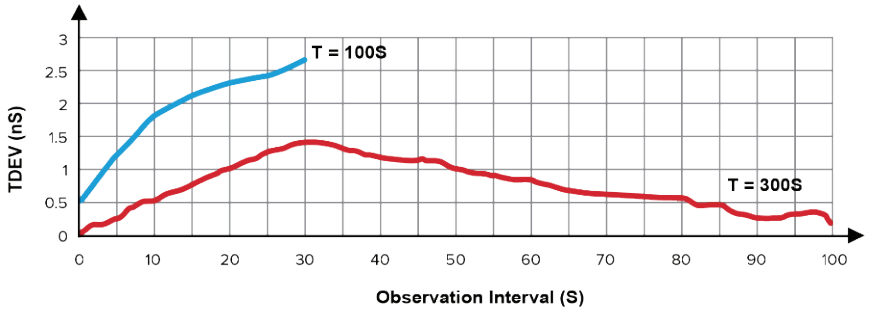

Figure 9. TDEV based on TIE values in Figure 7 (T = 100 S curve).
T = 300 S curve shows how TDEV reduces if TIE period extended and no significant further transients occur.

TDEV is also a 'signature' of an oscillator type, or of particular noise mechanisms. For example, ETS 300 462-1 associates particular gradients of the TDEV curve with particular noise mechanisms in clock oscillators:

- White noise phase modulation: s-0.5
- Flicker phase modulation: s0 [constant]
- White noise frequency modulation: s0.5
- Flicker frequency modulation: s1
- Random walk frequency modulation: s1.5

In all these gradient examples, s is the observation time (the width of the window that is used to interpret the TIE data).

In many situations, several noise mechanisms are present and the complex TDEV plot exhibits a range of different gradients at different points. Nevertheless, the overall characteristics of the TDEV signature are generally more useful in characterising and differentiating between different clock types.

Wander Tolerance

G.8262 specifies wander tolerance masks for stimulus wander characterised both in TDEV and MTIE terms. It also provides in some cases an explicit sinusoidal wander mask equivalent to the MTIE tolerance specification, since sinusoidal wander is particularly simple to generate. Nevertheless, the ability to generate wander based on TDEV is required by the standards and so is essential in a wander test set. Tolerance is assessed as a combination of all the following criteria:

- Maintaining the clock within prescribed performance limits (the exact performance limits are for further study)
- Not causing any alarms
- Not causing the clock to switch reference
- Not causing the clock to go into holdover

Wander Transfer

Wander transfer is also specified in G.8262. In the case of EEC-1 devices, the wander gain should be less than 0.2 dB in the band 1 Hz to 10 Hz, measured using a sinusoidal wander stimulus. For EEC-2, wander transfer is specified in TDEV terms for both stimulus and measurement.

Phase Transient Measurements

SyncE network schemes allow for EECs to switch clock source in order to optimise network timing accuracy. They can do this based either on signal quality considerations (e.g. if input clocks drift outside permissible frequency limits), or based on clock quality messages from upstream clocks. The clock quality is signalled by a specific Ethernet slow protocol packet forming an Ethernet Synchronisation Message Channel (ESMC). The phase transient response sections in G.8262 set limits on both short-term and long-term phase excursions when a clock source switch is forced. In the case of EEC-2, there are also limits on fractional frequency offset (rate of change of phase) and on fractional frequency drift (rate of change of fractional frequency offset) in the period after a forced clock source switch.

Measurement Considerations

It is possible to measure or estimate jitter using general-purpose test equipment such as oscilloscopes and spectrum analyzers. While these may give useful information, the inability to perform a true peak-to-peak measurement conforming to the test conditions required by the jitter standards means Standard compliance testing generally isn't possible. To do that, a purpose designed jitter measurement system, which provides wideband phase demodulation of the input jittered signal, is required. The demodulated signal is then filtered by standards-defined high-pass and low-pass filters, and processed to give peak-to-peak and RMS jitter measurements. The measurement of actual instantaneous peak values is critical as these are the ones specified in G.8262 because of their direct effect on the equipment under test.

Jitter measurements also require the clock to be recovered transparently from the signal to include all the jitter, followed by removal of jitter from the recovered clock to create the reference for jitter measurement in the phase detector. However, a common source of intrinsic jitter is in the clock recovery process itself. To ensure optimum performance and measurement accuracy, a test set must be designed and calibrated to minimize intrinsic noise in the measurement circuitry. Note that simply connecting a test set's transmitter to its receiver back-to-back with zero transmit jitter selected and noting the apparent noise floor is of limited use. It is tempting to subtract this value from jitter measurements, but incorrect since a back-to-back measurement includes the transmitter intrinsic jitter as well as the receiver contribution. In other words, a usable jitter test set needs to have minimal and well-quantified receiver intrinsic jitter levels.

To obtain quantitative measurements that are traceable to international standards, a dedicated jitter/wander test is strongly recommended. The advantages of this are:

- Ethernet data rates with clock recovery and built-in measurement filters.
- Specifically designed peak-to-peak detector.
- Accumulates and displays output jitter amplitude statistics in real time and displays against masks.
- Measures data—specifically designed clock-recovery circuits.
- Automated jitter tolerance measurement with pre-set masks, based on dropped packet detection.
- Measures wander against all common reference clock formats
- Displays TIE and calculated MTIE/TDEV results against masks
- Generates wander specified in both sinusoidal and TDEV (noise) terms
- Ability to generate and capture ESMC clock-quality messages
- Designed in accordance with O.174 (the ITU-T recommendation for jitter and wander test equipment applied to SyncE).

About Spirent Communications

Spirent Communications (LSE: SPT) is a global leader with deep expertise and decades of experience in testing, assurance, analytics and security, serving developers, service providers, and enterprise networks.

We help bring clarity to increasingly complex technological and business challenges.

Spirent's customers have made a promise to their customers to deliver superior performance. Spirent assures that those promises are fulfilled.

For more information, visit: www.spirent.com

Conclusion

Testing jitter and wander is a challenging exercise, particularly considering the variables involved. As has been explained, both devices and network behaviours generate jitter and wander. To ensure the safe delivery of signal and timing quality to base stations, the performance of the transmitted signal and clock must remain within the limits set by the ITU-T jitter and wander performance standards.

Spirent engineers have unsurpassed experience and expertise in the field of jitter and wander, both in terms of how it affects the network and how to build test equipment that accurately measures these parameters in practical situations. They also participate in developing industry standards and have contributed to the development of the ITU-T standards discussed here. Spirent measurement equipment offers very low intrinsic jitter and high accuracy so that jitter and wander measurements made on networks and devices are a true and repeatable reflection of their compliance to standards.

Glossary of Acronyms

EEC	Ethernet equipment clock
TIE	Time interval error
MTIE	Maximum time interval error
TDEV	Time deviation error
PRC	Primary reference clock
SDH	Synchronous digital hierarchy
SEC	SDH equipment clock
SSU	Synchronisation supply unit
UI	Unit interval

References

G.8262	Timing characteristics of a synchronous Ethernet equipment slave clock, International Telecommunication Union, July 2010
O.174	Jitter and wander measuring equipment for digital systems that are based on synchronous Ethernet technology: Amendment 1, ITU, April 2011
G.825	The control of jitter and wander within digital networks which are based on the synchronous digital hierarchy (SDH), International Telecommunication Union, March 2000

spirent™ Promise. Assured.

Contact Us

For more information, call your Spirent sales representative or visit us on the web at www.spirent.com/ContactSpirent.

www.spirent.com

Americas 1-800-SPIRENT
+1-800-774-7368 | sales@spirent.com

US Government & Defense
info@spirentfederal.com | spirentfederal.com

Europe and the Middle East
+44 (0) 1293 767979 | emeainfo@spirent.com

Asia and the Pacific
+86-10-8518-2539 | salesasia@spirent.com

Rev B | 01/19