Network synchronization

Keeping your finger on the pulse

AN 1367

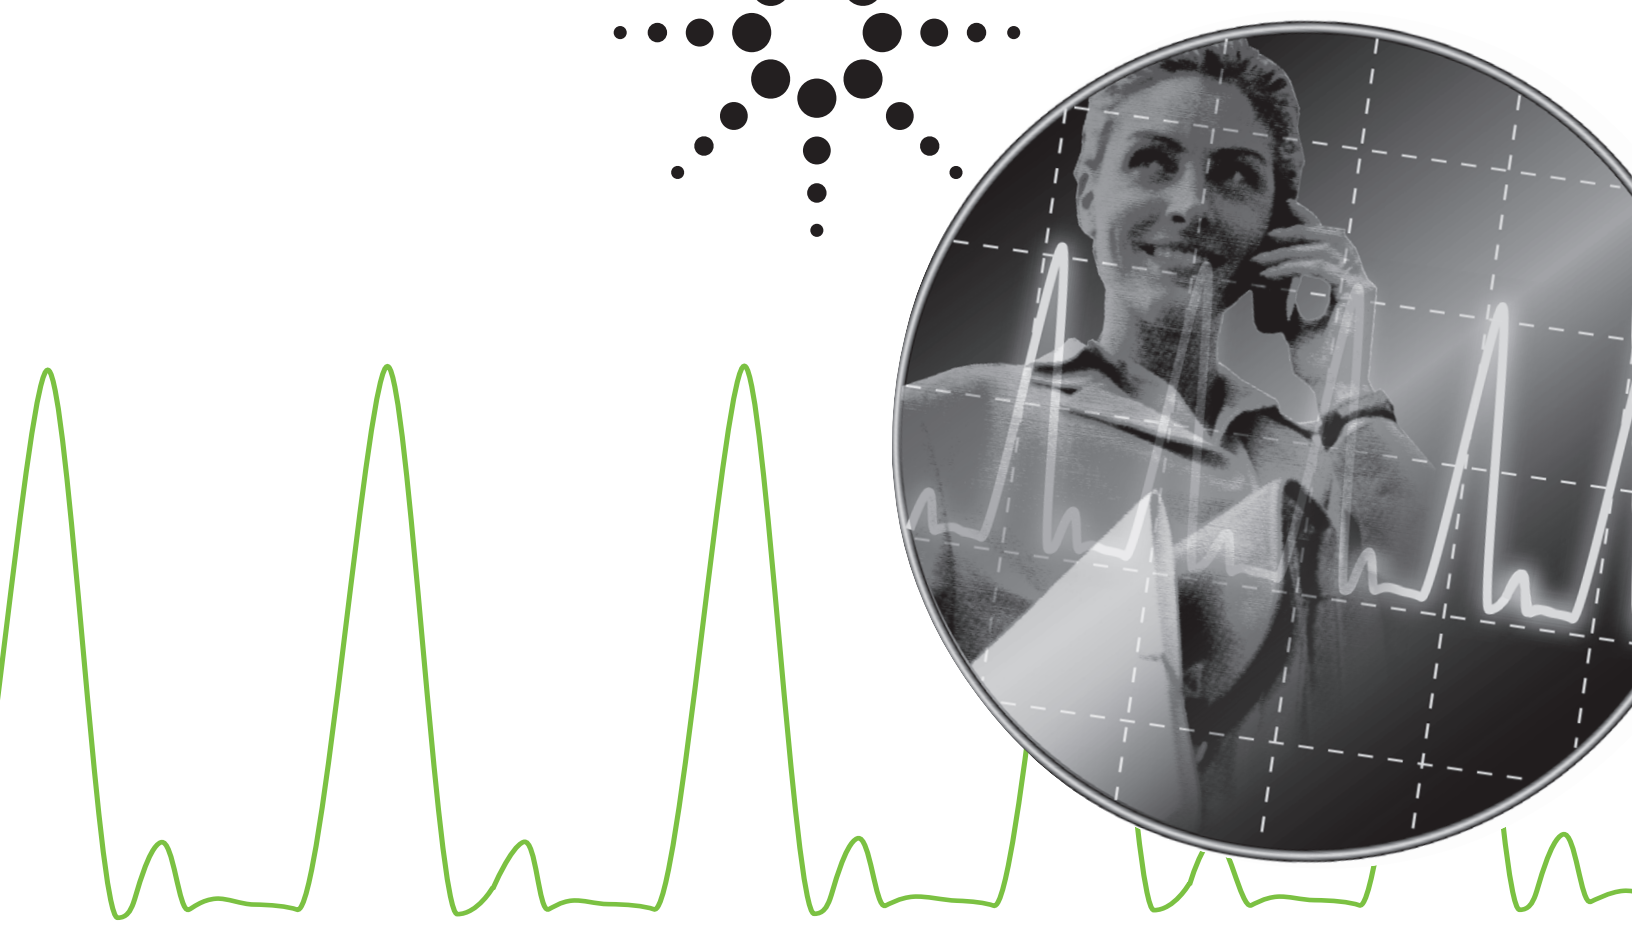

Agilent Technologies
Innovating the HP Way

Abstract

This application note provides a broad overview of the requirements and arrangements for network timing synchronization in telecom networks. It examines impairments in timing, both inherent and those due to faults or incorrect configuration. The emphasis is on faults and conditions which may not be immediately apparent but reduce margins or predispose the network to other fault conditions. It also looks at the effects of timing impairments, and how to measure and detect such impairments and their consequences.

Network timing requirements and arrangements

Today's telecom networks are hybrids of SDH/SONET and PDH/DSn. Typically, networks transport PDH signals as payloads and may consist of several SDH/SONET 'islands' interconnected at PDH interfaces (Figure 1).

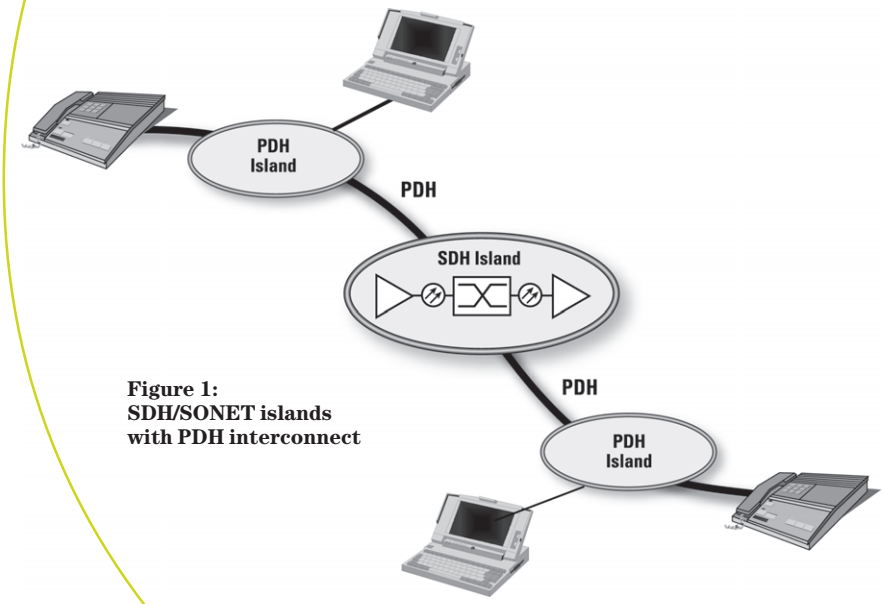

Figure 1:
SDH/SONET islands with PDH interconnect

Clearly, the ideal timing arrangement is to have all elements in the network timed from the same source, with minimal and well-controlled impairments on the clock distribution paths. The idea of a universal, single master clock for the entire worldwide telecom network is somewhat idealist: for example, where would it be located, what backup provisions would be necessary, and so on. In practice, individual national or operator networks have their own high stability master timing references (known as stratum 1 or PRC/S – primary reference clock/standards). These clocks are elaborate caesium-beam standards and their accuracy, as specified in ITU G.811 and ETS-300 462-6, is 1 part in 10^{11} over 20 years.

In order to distribute timing throughout a particular network, there are two (strictly three, in SONET) further classes of clock synchronized to the PRC in a pyramid fashion (Figure 2).

The stability specifications for each clock level reduce for two reasons: First, the physical process of transmitting the clock over a distance from one level to another introduces impairments; second, the lower level clocks have to operate autonomously for specific periods if a fault interrupts the incoming higher-level clock, and they cannot economically be as good as the PRC.

The middle-level SSU (synchronization source unit) or BITS (building integrated timing supply) clocks include memory that allows them to operate at higher stability than their inherent oscillators for a time after the PRC-aligned clock is

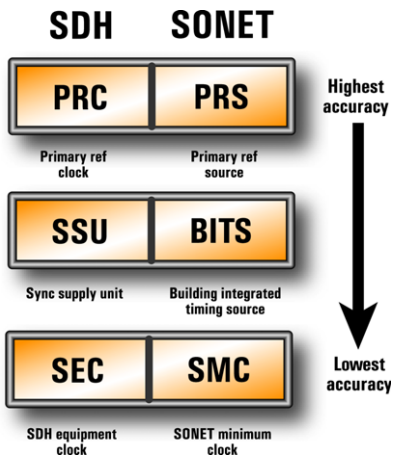

Figure 2:
Network clock levels

removed by a fault. This type of operation is known as 'holdover' mode. Once the memory based corrections become invalid, or if insufficient history is available, then the clock is said to be in 'free-running' mode.

The stability requirements on the lower level (stratum 3) clocks range from 1 in 10^{10} to 4.6 in 10^{6}. Note that all of these are considerably better than the fundamental stability requirement of 5 in 10^{5} for a basic 2 Mb/s PDH signal.

In today's environment of many different and new operators, the classical hierarchical arrangement of one (or few) PRCs is tending to break down. There are now many more chances of unintended disconnects in the timing chain, or islands of timing that need monitoring.

Timing information is physically distributed in several different ways: by the embedded PDH signals in the network, by the line-rate SDH/SONET signals between elements, or by physically separate networks to the transmission paths. Transmitting timing by embedded PDH signal is not recommended because the embedded signals are exposed to transmission impairments. Using the SDH/SONET line signal is much better from this point of view, and also allows timing 'markers' or status messages to be transmitted in the SDH/SONET overhead to help control timing. However, relying simply on checking sync status messages to check the validity of a signal for timing purposes is also not recommended. These messages depend essentially on manual configuration of the network, and since errors are always possible, even likely, direct measurement of synchronization quality remains very important (Figure 3).

Current technology has opened up some exciting new possibilities for distributing PRC signals via the GPS (global positioning system). GPS satellites contain very high accuracy clocks since precise

S1 byte b5–b8	SDH synchronization quality level description
0000	Quality unknown (existing sync network)
0001	Reserved
0010	G.811
0011	Reserved
0100	G.812 transit
0101	Reserved
0110	Reserved
0111	Reserved
1000	G.812 local
1001	Reserved
1010	Reserved
1011	Synchronous equipment timing source (SETS)
1100	Reserved
1101	Reserved
1110	Reserved
1111	Do not use for synchronization

S1 byte b5–b8	Quality level	SONET synchronization quality level description
0001	1	Stratum 1 traceable
0000	2	Synchronized – traceability unknown
0111	3	Stratum 2 traceable
1010	4	Stratum 3 traceable
1100	5	SONET minimum clock traceable
n/a	6	Stratum 4 traceable
1111	7	Do not use for synchronization
1110	user-assignable	Reserved for network synchronization use

Figure 3:
Sync status messages carried in the SDH/SONET signal

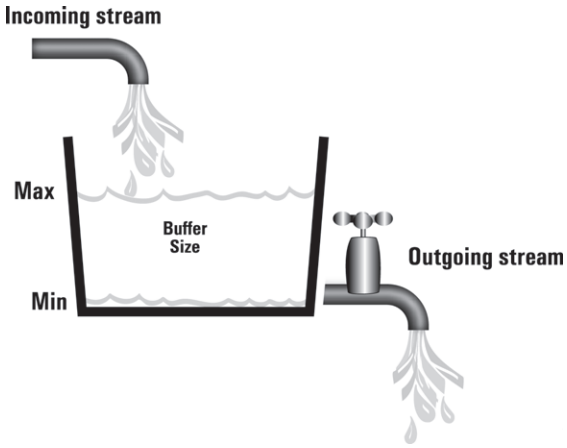

Figure 4:
In principle, a buffer is like a water tank

timing measurements are central to the primary positioning function of GPS. GPS is now well-proven, and recent US government policy changes have resulted in higher potential timing accuracy for commercial users. Deliberate 'spoiling' features that diluted positioning accuracy for public users of GPS were switched off on May 1, 2000. The primary intent was to allow for high-accuracy location of mobile phones on emergency calls, but a side result is improved timing stability. Despite that, direct-from-air GPS is not up to PRC standards, but if it is combined with local quartz or rubidium oscillators, it can provide stability up to 1 in 10^{12} which is better than a PRC. However, such timing networks need good antenna installations (so will not work in all locations), and there is some nervousness about relying a hundred percent on GPS despite high reliability and guarantees of future continuation of the service.

The consequences of timing differences

Even with such tightly-controlled timing, there are obviously small timing differences between network elements and segments. The network has design features to make it tolerant to synchronization conditions where all clock elements are at the limit of their allowed tolerances (and, some margin above that).

There are three classes of timing difference between any two parts of a network:

Steady-state timing differences

This is basically a frequency offset between one part of the network and another. The serial transmission schemes used in telecom networks require that the downstream receiver recovers its clock from the incoming serial stream, and thus adapts to variations in frequency. Gross frequency offset will ultimately cause a line signal to fall outside the clock recovery tolerance of the receiver in the downstream network element (for example, the tolerance for a 2 Mb/s PDH signal is ± 50 ppm). Faults like this are very unusual, however, even modest frequency offset is significant. Network elements usually incorporate buffers to allow for timing differences. These buffers have a fixed size (usually around one transport signal frame, or 125 µS).

A sustained frequency offset will ultimately cause a buffer to overflow or underflow, causing what is known as a 'frame slip'. This is where an entire frame of the signal is deleted or duplicated. The consequences range from trivial clicks on voice traffic down to serious disruption to video signals or packetized data. Note that two networks with single-frame buffers and timed from separate PRCs that are within specification would see a maximum slip rate of 1 every 72 days (Figure 4).

Periodic timing differences

High frequency (above 10 Hz) periodic variations – known as timing jitter – are very serious in digital transmission systems. Clock recovery in line receivers can only follow such variations up to specific values, above which errors occur. Clock recovery devices work by continuously adjusting to track the average incoming rate. There's a trade-off between the amount of jitter that can be tolerated and its frequency. The higher the frequency, the less jitter is allowed (Figure 5).

It is sometimes assumed, therefore, that low frequency (<10 Hz) periodic variations–known as wander–wouldn't be so serious. In some ways that's true since PDH

multiplexing contains a mechanism known as stuffing or justification to accommodate frequency offset (up to the maximum frequency tolerance for the particular rate) and wander.

The action of the stuffing/justification process induces some timing jitter at the demultiplexing point (because signal timing needs to be smoothed out) but it is seldom serious.

Also, because wander is periodic, it won't cause slip problems unless it is of very high amplitude. The buffers in network elements fill and empty along with the wander to accommodate it.

However, in a mixed SDH/PDH network (that is, the network as it is today), the effect of wander is more serious. In return for the elegance and simplicity of multiplexing that SDH/SONET provides, it is necessary to have a frequency offset/wander accommodation mechanism that operates in coarse steps. In the multiplexing process, the containers into which the 'tributary' signals are mapped are adjusted relative to the SDH/SONET overhead in steps as large as 24 bits at a time.

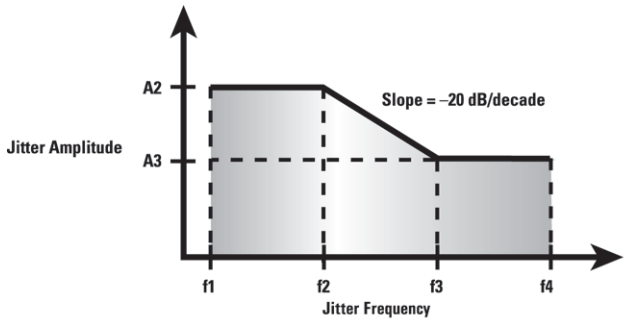

	A2	A3	f1	f2	f3	f4	Pattern
2 Mbls	1.5	0.2	20 Hz	2.4 kHz (93 Hz)	19 kHz (700 Hz)	100 kHz	2^{15}
8 Mb/s	1.5	0.2	20 Hz	40 Hz	3 kHz	400 kHz	2^{15}
34 Mb/s	1.5	0.15	100 Hz	1 kHz	10 kHz	800 kHz	2^{23}
140 Mb/s	1.5	0.075	200 Hz	500 Hz	10 kHz	3500 kHz	2^{23}

Figure 5:
Typical jitter tolerance specification

The mechanism that controls the position of the container relative to the transport signal is known as a pointer. In the case of timing differences, these pointers change over time. These so-called pointer adjustments occur at a rate in the network that depends on frequency offset and wander. For example, when a demultiplexer (strictly, a 'demapper') has to reconstitute a PDH signal from an SDH/SONET signal that includes pointer adjustments, jitter is induced depending on the design of the demapper and on the rate and nature of pointer adjustments. Because these in turn are a consequence of timing differences, the timing quality is vitally important in SDH/SONET networks.

Random timing differences

Electrical noise, and the intrinsic characteristics of devices like phase-locked loops within timing devices and clock recovery circuits can result in random timing variations. Another very common type of random change in timing is transients caused by clock nodes switching from one clock source to another.

In summary, timing variations lead either to slips and therefore data errors, or to pointer adjustments that lead to payload jitter which, in turn, leads to data errors (Figure 6).

In addition, network timing is used by cellular mobile systems extensively and impaired timing here can lead to complete network failure. This is especially the case in the TDM-based GSM mobile networks.

Timing impairment measurements

The most obvious measurement to make is simply to measure the frequency of a line or tributary signal directly. Network test equipment like the Agilent OmniBER family of analyzers include a built-in frequency measurement, either against the built-in stratum 3 clock, or against an external reference. This test is quite insensitive to short-term variations, and is more useful in determining gross faults like clock loops–a situation in which two

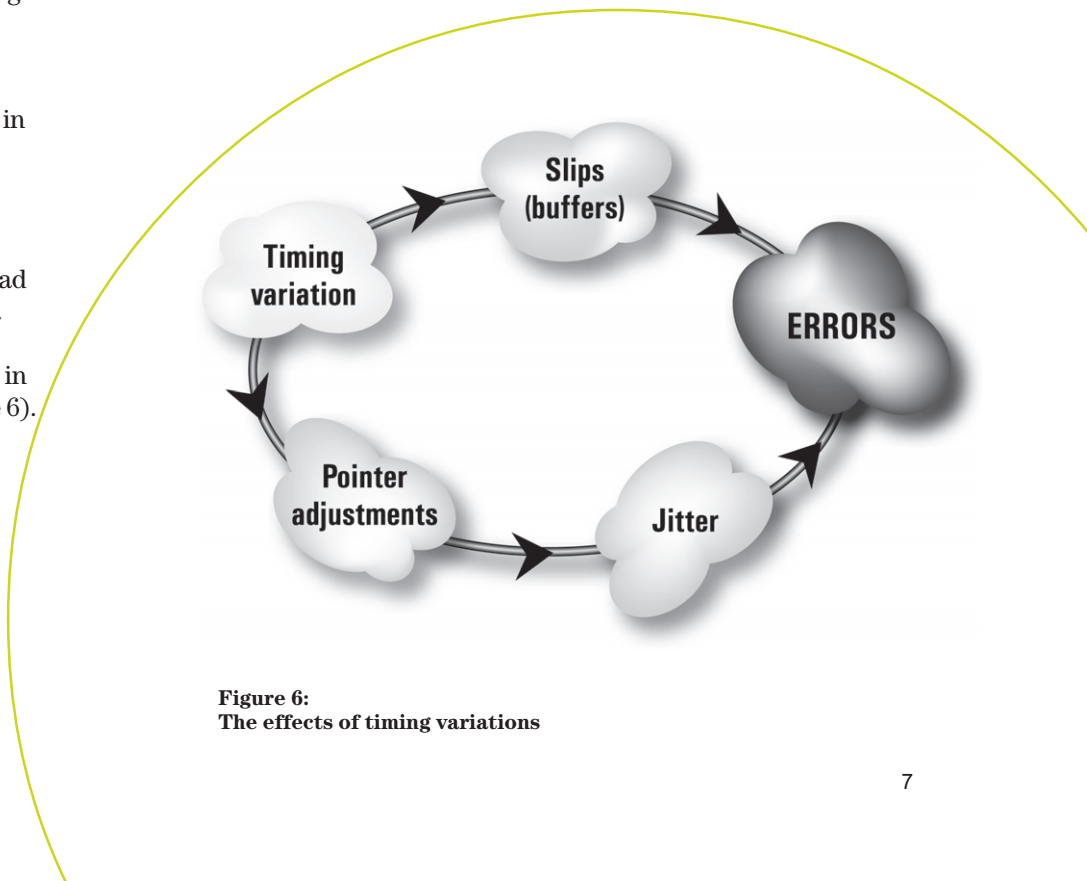

Figure 6:
The effects of timing variations

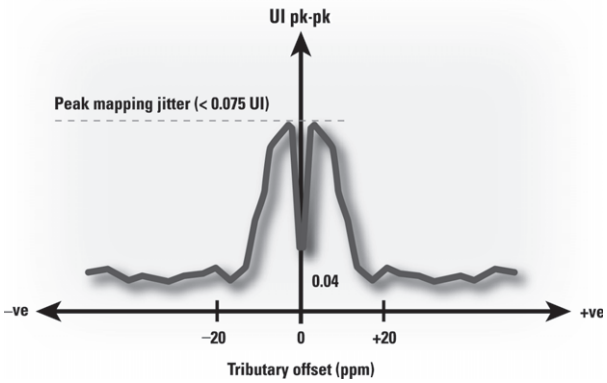

Figure 8:
How tributary jitter varies with frequency offset

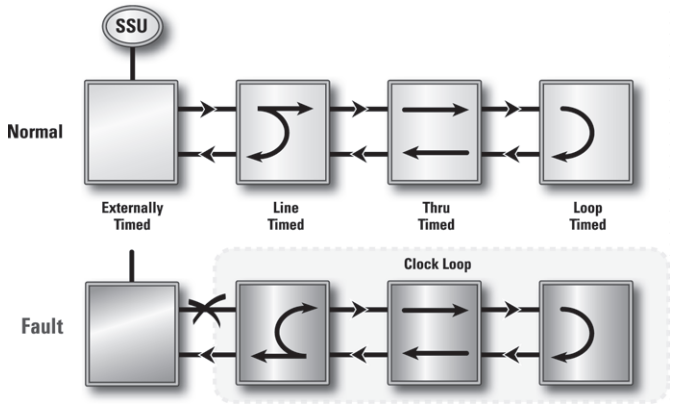

Figure 7:
Unintended clock loops

devices connected together are both set to synchronize to each other's timing without either using a master clock or internal reference. Sometimes in such situations the frequency can be substantially high or low (Figure 7).

A related test to frequency measurement is to offset the clock of a test set transmitter to see how tolerant the receiver of the network element is to clock offset, and to see how the multiplexing process behaves in the presence of frequency offsets. The offset can be applied either directly to a signal at a physical interface, or indirectly to a signal prior to mapping it within the test set into a higher-rate SDH/SONET line signal. In this case, the offset can be varied to see how the demapping jitter of a demultiplexer varies. Note that quite fine resolution is required in offset for such tests because the demapping jitter has a general characteristic of very narrow peaks close to zero offset (Figure 8).

Testing the phase of one network clock relative to another–jitter and wander testing–is much more sensitive than frequency measurement, both to short-term and long-term variations. All jitter and wander testing involves measuring the phase variations over time between the clock signal at the point being measured, and a reference clock (Figure 9).

The units used to measure the phase variation are either phase units (1 unit interval or UI, equal to one clock period or 360 degrees of phase) or time units (usually in nS). The use of the unit interval is normal in jitter measurement and in some simple wander measurement situations, whereas direct measurement in nanoseconds is normal in more sophisticated wander measurements.

The phase reference is very important in all phase-based measurements. In jitter measurements, a clock recovered from the point of measurement itself can be processed to provide a local reference. There is some settling time associated with that, but it is generally only a few seconds. Jitter measurement is limited to frequencies above 10 Hz, so it is only necessary to remove components above this frequency from the reference. In wander measurement, the situation is more difficult.

Wander frequencies can be extremely low (in the µHz range and below–for example the frequency of a wander component that repeated once per day would be about 12 µHz), and so the settling time of an internal reference is very long. Wander measurements are, therefore, only really meaningful where an external reference of known good quality is used, for example an SSU or BITS clock.

As previously discussed the possibilities of GPS-tuned quartz timing sources mean that direct wander measurements against a PRC-quality clock are now possible.

Let's now look at how the phase measurements are processed and interpreted. First, jitter:

In real network situations, jitter has various frequency components that depend on source. It's also important to be able to measure the magnitude of the jitter at different frequencies because, as discussed earlier, the disruptive effect of jitter depends on its frequency. Therefore, the magnitude of the phase variations is measured after bandpass filtering, so that the magnitude in different frequency bands can be measured separately. Jitter test sets like the Agilent OmniBER also provide an electrical output of the detected (or demodulated) jitter signal that can be separated into specific frequency components using a spectrum analyzer. Think of the in-built filters in a jitter test set as a coarse spectrum analyzer.

Jitter magnitude is then measured and displayed as a pk-pk value and as an RMS value. It is also useful to measure the number of times pk-pk jitter exceeds a threshold value, known as jitter hits measurement (Figure 10).

There are three standardized jitter measurement situations:

Jitter tolerance

Here, a test set is connected to an isolated network element (the device-under-test, DUT) through a specific physical interface (for example, an STM-64 or OC-48 port). The test set receiver is connected to an output port such that the test set can measure the error performance of the through connection. The test set then adds jitter at a range of frequencies and amplitudes to find the points at

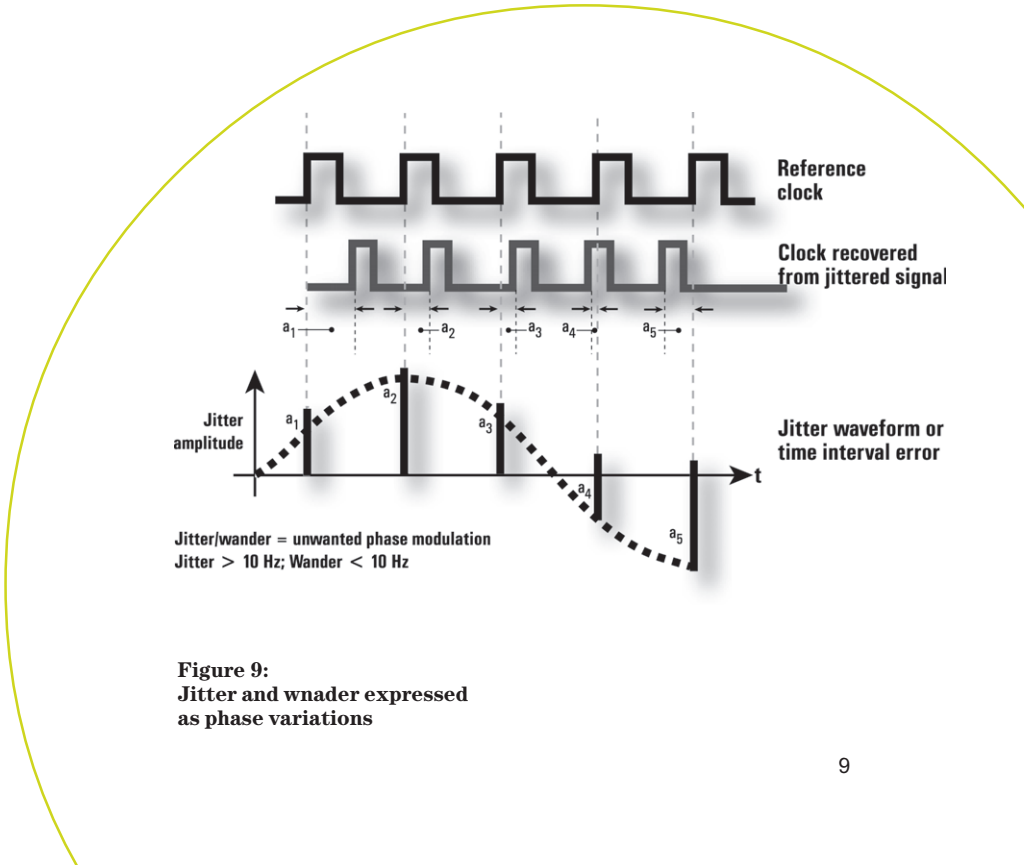

Figure 9:
Jitter and wnader expressed as phase variations

which errors first occur. The network element should have a jitter tolerance that exceeds specific masks (Figure 11).

Jitter transfer

Here, the ratio of jitter at an output port to the amount of jitter applied at an input port is calculated, for different jitter frequencies. It is vital to measure with a narrow filter on the receiver that tracks the transmitted jitter frequency to exclude jitter from outside sources.

Output (intrinsic) jitter

This is a straightforward measure of the amount of jitter on a DUT output, with no applied jitter to any input port. Network equipment specifications on output jitter are very tight, so jitter test sets must measure such low values repeatably and accurately.

Pointer sequences

As mentioned above, SDH/SONET networks tolerate timing differences by allowing the containers carrying service traffic to float relative to the transport signal frame (Figure 12).

Despite elaborate re-timing arrangements in terminals that demultiplex the service payloads, the disruptive effect of pointer movements (or 'adjustments') leaks through to the tributary outputs as spikes of large-amplitude, low-frequency jitter. To evaluate the quality of network elements, use a test set to generate various sequences of pointer adjustment to stress the retiming function. The test set measures the amount of jitter induced in the presence of pointer adjustments. There are standardized sequences of pointers (Figure 13).

Pointer behavior

In network monitoring and maintenance situations, it's very useful to monitor and graph the behavior of pointers over long periods. Pointer activity shows if significant timing differences

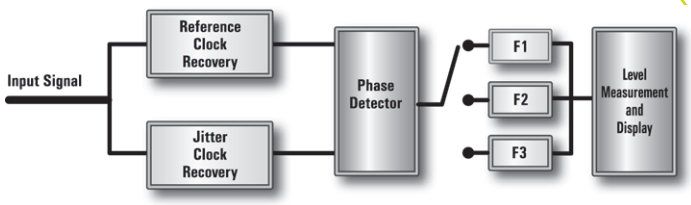

Figure 10:
Jitter measuring arrangement

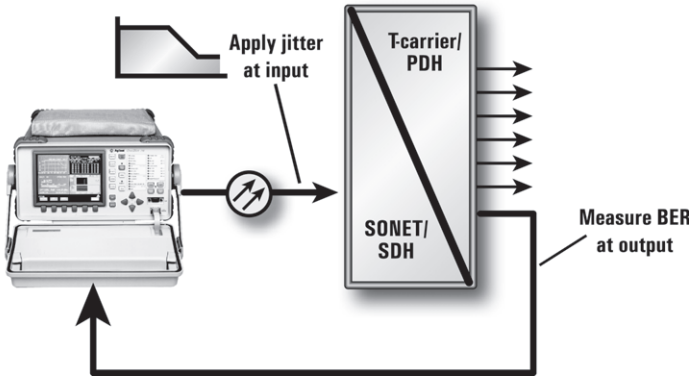

Figure 11:
Jitter tolerance arrangement

exist in the network, and the approximate nature of these differences over time (Figure 14).

Explicit wander measurement

Let's now turn to the direct and explicit measurement of timing differences. The need for a reliable reference against which to measure, and possible solutions, has already been discussed. Assuming a viable reference, the root measurement is to compare the phase of the timing signal at the point being measured with the phase of the reference clock. The value of this phase difference (usually expressed in nS) as a function of time is known as time interval error (TIE). TIE is usually initialized to zero at the start of the measurement interval. As we've seen previously, timing is usually carried at SDH/SONET line rates between sites, so it is vital for test equipment to source the clock to be measured from a line-rate input prior to comparing it to the reference (Figure 15). A TIE plot shows the overall timing situation at the point of measurement.

There are two particular aspects of a TIE plot to interpret: First, the gross TIE behavior. The most obvious gross element is a steadily increasing (or decreasing) TIE with time. In this case, the clock at the point of measurement is running at a different frequency

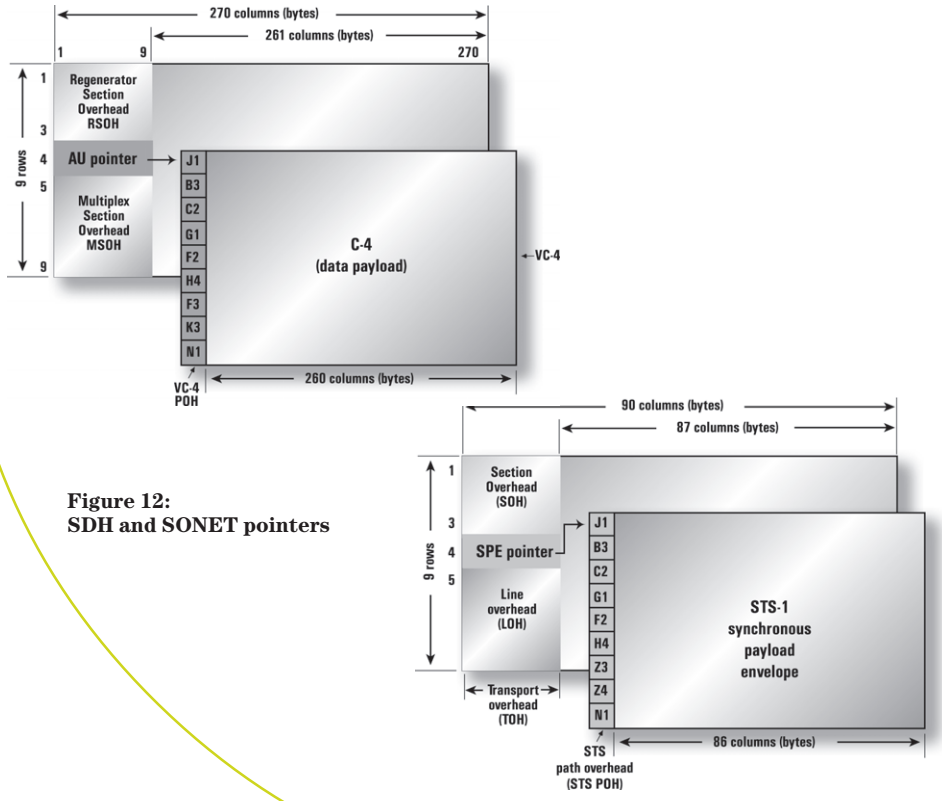

Figure 12:
SDH and SONET pointers

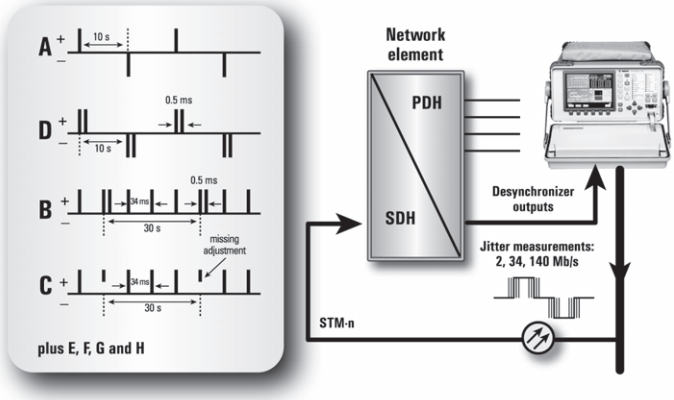

Figure 13:
Pointer adjustment sequences

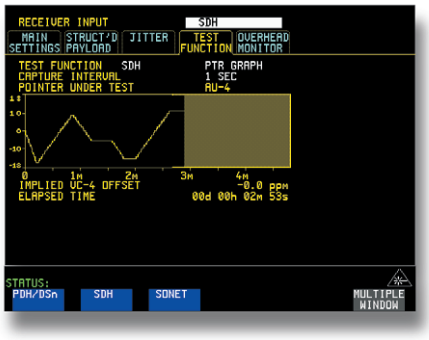

Figure 14
OmniBER screen shot

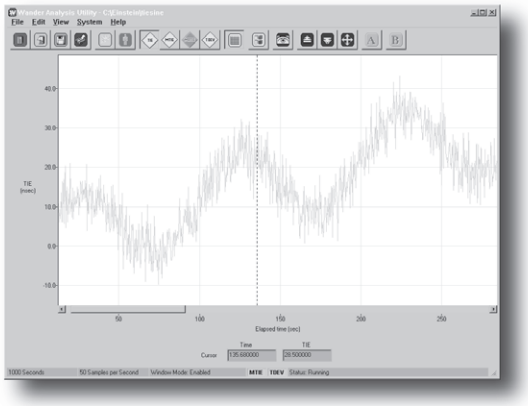

Figure 15:
TIE vs time

to the reference clock. This could be because an intended clock path does not actually exist (sometimes, despite the sync status messages appearing okay), or a timing node has gone into free-run or holdover mode because of a fault or configuration error.

MTIE

The gross TIE performance is also more rigorously quantified by a calculation from the TIE plot known as MTIE (maximum time interval error). In this situation, a variable window or 'observation interval' is moved through the TIE values, and the measurement latches the highest 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.

Note that network standards give several algorithms for the calculation of MTIE. Some are designed to give more approximate answers but are easier to calculate. In practice, none of the computation techniques allow manual calculation, so the only practical approach is to use a PC-based package like the Agilent Technologies E4547A wander analysis software to control an Agilent OmniBER analyzer to measure TIE over extended periods and process the results (Figure 16).

The MTIE measurement shows the long-term behavior and stability of the timing at the point being measured. Generally, it is interpreted relative to standardized masks in the same way as jitter. Results within the mask are acceptable, those above it are not. It is also sometimes useful to benchmark actual timing performance and to look for changes with time, even if results are still within the mask. You can also compute a very closely related result called MRTIE (maximum relative time interval error) in which the underlying frequency offset component is subtracted out, leaving only the periodic and random parts of MTIE. This is particularly useful if you have to rely on a reference clock that has adequate long-term stability but is still essentially free-running relative to the network node being measured. An example could be using a GPS-tuned local clock.

The Agilent wander analysis software also adds an important enhancement to MTIE measurement. Sometimes, a sudden change (transient) can occur in the phase of reference clocks. This could, for example, be due to switching from

one clock to another if an interruption occurs. As the diagram shows, the effect of a transient is to mask all other MTIE information of lower amplitude other than the transient itself. This is because the transient is fully registered by even very small observation window values, and the MTIE plot only becomes sensitive to behavior away from the transient once the observation window becomes larger.

The Agilent package includes an optional second measurement interval that sweeps through a restricted range of TIE values, thus suppressing the effect of transients. The measurement package latches then displays the short-term MTIE values alongside the full MTIE values. From this modified plot, it is then easy to spot MTIE trends in the area that is normally masked by transients, and also to see if any significant transients are present (if no transients, full MTIE and short-term MTIE are similar, or if there are transients, then the two lines will be further apart) (Figure 17).

TDEV

The short-term characteristic (that is, phase noise) performance of timing signals are also very significant. Qualitatively (that is, just by looking), it is difficult to analyze the noise on a TIE plot. Also, the MTIE result is focused on revealing the long-term behavior 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, and is calculated by taking the average standard deviation of TIE for varying window sizes that 'move' through the TIE data. Even more so than in MTIE, the calculation is complex and requires automation (Figure 18).

Once again, results are compared against masks, or benchmarked. TDEV is also, in some ways, a 'signature' of a particular clock 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: s^{0} [constant]
White noise frequency modulation: $s^{0.5}$
Flicker frequency modulation: s^{1}.
Random walk frequency modulation: $s^{1.5}$.

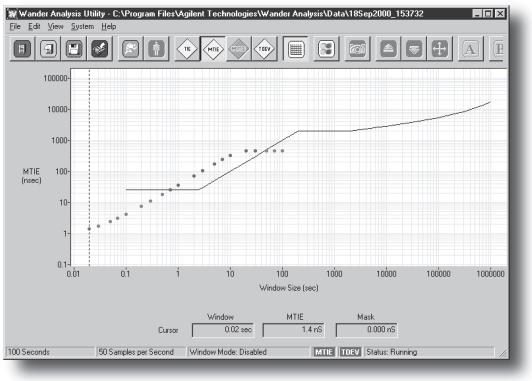

Figure 16:
MTIE vs window size

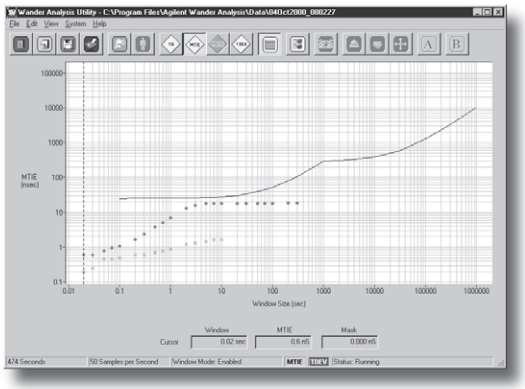

Figure 17:
Full MTIE and short-term MTIE vs window size

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 characterizing and differentiating between different clock types.

MTIE and TDEV testing is performed both at equipment design/verification stages and on working networks in the field. There are a large number of masks defined for each quantity, and new masks appear regularly. In manufacturing or design, nonstandard masks are sometimes used. For these reasons, the Agilent wander analysis package includes easy ways to specify new masks with simple equations rather than having to input many specific mask points.

Wander stimulation

Just as in the jitter testing discussed earlier, it is useful to subject a network element to wander, and to study the amount and characteristics of the wander transferred to the element output. Another useful test is to subject an SSU to wander, such that the built-in memory 'learns' the applied wander, then to remove the input signal (forcing the SSU into hold-over mode) and study subsequent behavior of the output. Note that these tests are normally performed only in design or verification of network elements.

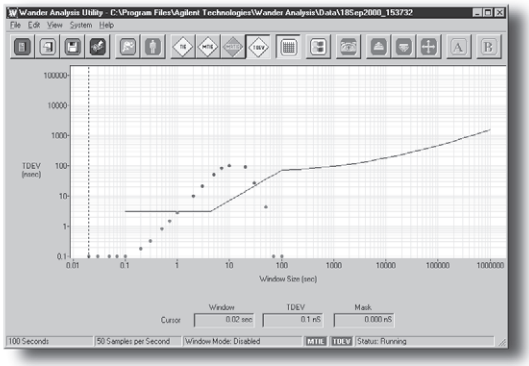

Figure 18:
Typical TDEV vs window size

Conclusion

Distributed timing is the heartbeat of the telecom network, and just as monitoring the pulse of a person indicates state of health, so monitoring timing in the network also gives health and diagnostic information. Poor timing can exist for many reasons, and increasing deregulation of the telecom network makes the problem more challenging. Good design and tight control of configuration minimizes the chance of timing problems, but only regular monitoring and checking can ensure that no latent problems exist.

Standards and references

Document Number	Title	Revision/Issue
ITU G.810	Definitions and terminology for synchronization networks	(08/96)
ITU G.811	Timing characteristics of primary reference clocks	(09/97)
ITU G.812	Timing requirements of slave clocks suitable for use as node clocks in synchronization networks	(06/98)
ITU G.813	Timing characteristics of SDH equipment slave clocks (SEC)	(08/96)
ITU G.803	Architectures of transport networks based on the synchronous digital hierarchy (SDH)	(03/93)
ITU G.823 Draft	The control of jitter and wander within digital networks which are based on the 2048 kbit/s hierarchy	(Geneva, Feb 99)
ITU G.824 Draft	The control of jitter and wander within digital networks which are based on the 1544 kbit/s hierarchy	(Geneva, Feb 99)
ITU G.825 Draft	The control of jitter and wander within digital networks which are based on the synchronous digital hierarchy (SDH)	(Geneva, Feb 99)
ITU G.707	Network node interface for the synchronous digital hierarchy (SDH)	(03/96)
ITU G.783	Characteristics of synchronous digital hierarchy (SDH) equipment functional blocks	(04/97)
ITU G.822	Controlled slip rate objectives on an international digital connection	1993
ITU O.172	Jitter and wander measuring equipment for digital systems which are based on the synchronous digital hierarchy (SDH)	(03/99)
EN 300 462-1-1	Definitions & terminology for synchronization networks	(05/98)
ANSI T101-1	Synchronization interface standard	T1X1.3/99-020R2

Glossary

Allen Deviation	A quantity calculated from TIE data using a similar algorithm to TDEV and yielding similar analysis of clock noise
BITS	Building Integrated Timing Source – a slave clock
GPS	Global Positioning System
MTIE	Maximum time interval error (see text for details)
PDH	Plesiochronous Digital Hierarchy – In North America the DS-1 through DS-3 legacy network, in the rest of the world the 2/8/34/140Mbit/s legacy network
Ppm	parts per million
PRC	Primary Reference Clock
SSU	Synchronization Source Unit – a slave clock
TDEV	Time Deviation (see text for details)
UI	Unit Interval – jitter & wander measurement unit equal to one complete clock period

---

Agilent Technologies'
Test and Measurement Support, Services, and Assistance

Agilent Technologies aims to maximize the value you receive, while minimizing your risk and problems. We strive to ensure that you get the test and measurement capabilities you paid for and obtain the support you need. Our extensive support resources and services can help you choose the right Agilent products for your applications and apply them successfully. Every instrument and system we sell has a global warranty. Support is available for at least five years beyond the production life of the product. Two concepts underlie Agilent's overall support policy: "Our Promise" and "Your Advantage."

Our Promise
Our Promise means your Agilent test and measurement equipment will meet its advertised performance and functionality. When you are choosing new equipment, we will help you with product information, including realistic performance specifications and practical recommendations from experienced test engineers. When you use Agilent equipment, we can verify that it works properly, help with product operation, and provide basic measurement assistance for the use of specified capabilities, at no extra cost upon request. Many self-help tools are available.

Your Advantage
Your Advantage means that Agilent offers a wide range of additional expert test and measurement services, which you can purchase according to your unique technical and business needs. Solve problems efficiently and gain a competitive edge by contracting with us for calibration, extra-cost upgrades, out-of-warranty repairs, and on-site education and training, as well as design, system integration, project management, and other professional engineering services. Experienced Agilent engineers and technicians worldwide can help you maximize your productivity, optimize the return on investment of your Agilent instruments and systems, and obtain dependable measurement accuracy for the life of those products.

By internet, phone, or fax, get assistance with all your test & measurement needs

Online assistance:
www.agilent.com/find/assist

Phone or Fax
United States:
(tel) 1 800 452 4844

Canada:
(tel) 1 877 894 4414
(fax) (905) 206 4120

Europe:
(tel) (31 20) 547 2323
(fax) (31 20) 547 2390

Japan:
(tel) (81) 426 56 7832
(fax) (81) 426 56 7840

Latin America:
(tel) (305) 267 4245
(fax) (305) 267 4286

Australia:
(tel) 1 800 629 485
(fax) (61 3) 9272 0749

New Zealand:
(tel) 0 800 738 378
(fax) 64 4 495 8950

Asia Pacific:
(tel) (852) 3197 7777
(fax) (852) 2506 9284

Product specifications and descriptions in this document subject to change without notice.

Printed in USA February 1, 2001
5988-0360EN

Agilent Technologies
Innovating the HP Way