AT&T 62411 Design Considerations - Jitter and Synchronization

By Bob Bridge

INTRODUCTION

This application note outlines the technical requirements which must be considered when designing a system to meet the AT&T 62411 synchronization and jitter requirements. The first section discusses how the digital network is synchronized. This section is followed by a discussion of the jitter/synchronization requirements for trunk cards and system clocks.

AT&T 62411 is the network interface specification which should be adhered to at the point of demarcation between the AT&T network and the customer premises location. It applies when connecting *CPE* (Customer Premises Equipment) to an AT&T 1.544 MHz access line or private line. Note that 62411 does not usually apply to equipment sold to telephone companies, equipment used internally to a campus environment or equipment used to access an alternative long-distance carrier (e.g., MCI, US Sprint, etc.).

AT&T 62411 has gone through numerous revisions (1983, 1985 and 1988) with the most recent version published in December, 1990.

SYNCHRONIZATION OF DIGITAL NETWORKS

The digital telephone network is synchronous. "Synchronous" implies that all T1 systems in the network are designed to operate at exactly the same average frequency (1.544 MHz $\pm$ 0 ppm). Synchronous operation is required so bits of information are not dropped as data is transfered between the various sinks and sources in the network. Imagine the problem that would occur in the codec of a digital telephone set if the codec were required to perform its A/D conversions and D/A conversions at different clock rates.

Synchronization is achieved via the hierarchical distribution of a Stratum 1 clock (Figure 1). Each carrier has one or more Stratum 1 clocks. The frequency of this master clock is distributed via various transmission media (optical cable, microwave radio and satellite) through a hierarchy of switching systems (toll, tandem and central-office switches) until it finally reaches the CPE. The CPE typically recovers the frequency of the incoming T1 line, and uses that frequency as the basis for its system clock and for retransmitting back toward the network. Typically, the network switching systems have Stratum 2 and 3

clocks, and the customer system has Stratum 3 or 4 clocks. This overall timing plan is defined in ANSI T1.101-1987, Synchronization Interface Standard for Digital Networks.

The engineering of a network needs to ensure that the network will remain locked to a Stratum 1 frequency even upon the failure of one or more trunks that relay the timing information.
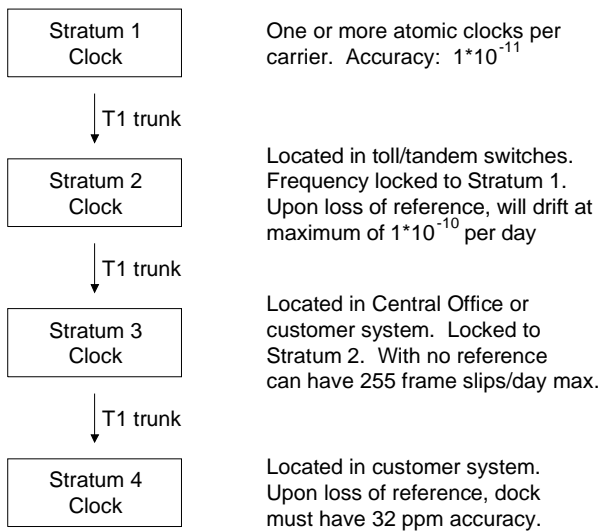
Stratum 1 Clock	One or more atomic clocks per carrier. Accuracy: $1*10^{-11}$
↓ T1 trunk	
Stratum 2 Clock	Located in toll/tandem switches. Frequency locked to Stratum 1. Upon loss of reference, will drift at maximum of $1*10^{-10}$ per day
↓ T1 trunk	
Stratum 3 Clock	Located in Central Office or customer system. Locked to Stratum 2. With no reference can have 255 frame slips/day max.
↓ T1 trunk	
Stratum 4 Clock	Located in customer system. Upon loss of reference, dock must have 32 ppm accuracy.

Figure 1 - Digital Timing Hierarchy

Figure 2 shows how a Stratum 1 traceable timing reference is passed through the T1 trunks of a private network. Upon failure of a trunk, a secondary path is used. Upon selection of an inappropriate secondary reference, an isolated timing island (with no master clock source) is created.

Figure 3 shows a typical system configuration where more than one trunk is used to input a traceable timing source into the system. Upon the failure of the primary timing source, a switch is made to a backup timing source. In the rare event that all timing sources fail, the system depends upon the ability of the synchronizer to free-run near the desired frequency.

To prevent chatter, it is important that there be hysteresis in the switch between reference sources. A switch should not be repeated any sooner than 10 seconds after the last switch, and only if one of the following conditions exists:
1) Phase hit of 1000 ns with phase
    slope $\leq 6.1 \times 10^{-5}$
2) Loss of signal (all zeros) for 0.10 seconds
3) 10^{-3} BER for ≥ 2.5 seconds
4) Excessive Input Jitter amplitude (Figure 8).
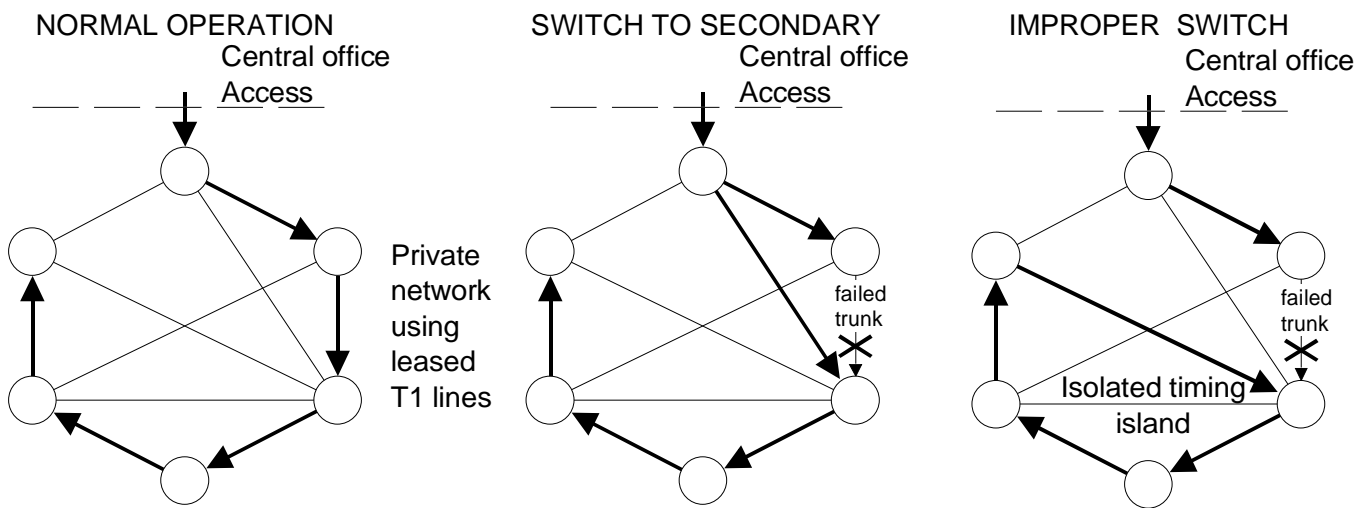
Stratum 1 tracable timing reference is transfered on trunks shown with bold lines.

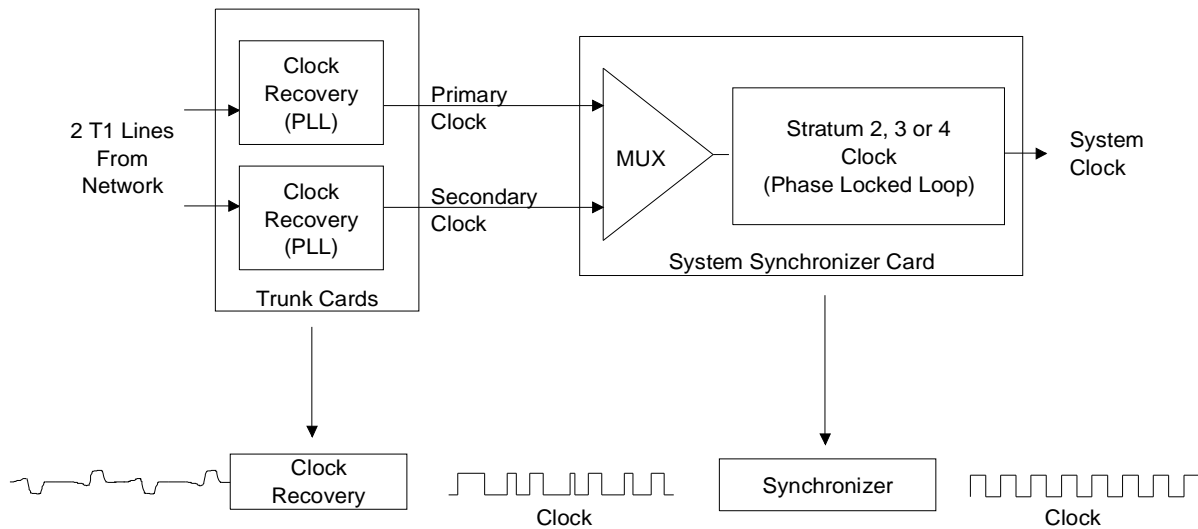**Figure 2 - Multiple Paths to Stratum 1 Clock**

Figure 3 - Multiple Timing Sources Feeding a System Synchronizer

Synchronization plans for private networks can become, in practice, quite complicated. Figure 4 shows a multi-state private network with connections to multiple inter-exchange carriers, multiple local-exchange carriers and an international carrier. Also note that Bell South has trunks to three inter-exchange carriers. In this example, the PBX in Atlanta can be the timing source for the private lines to Long Beach. The Atlanta PBX gets its Stratum 1 traceable source from AT&T via Bell South. The PBX in Long Beach uses this

AT&T timing to talk to Europe and to GTE of California.

CPE falls into one of two categories from a synchronization point of view. The CPE can either relay incoming timing to another downstream system, or it can terminate the timing chain. When a system terminates timing, failure of the system does not affect the timing of any downstream systems. This situation is refered to as *loop timing*.
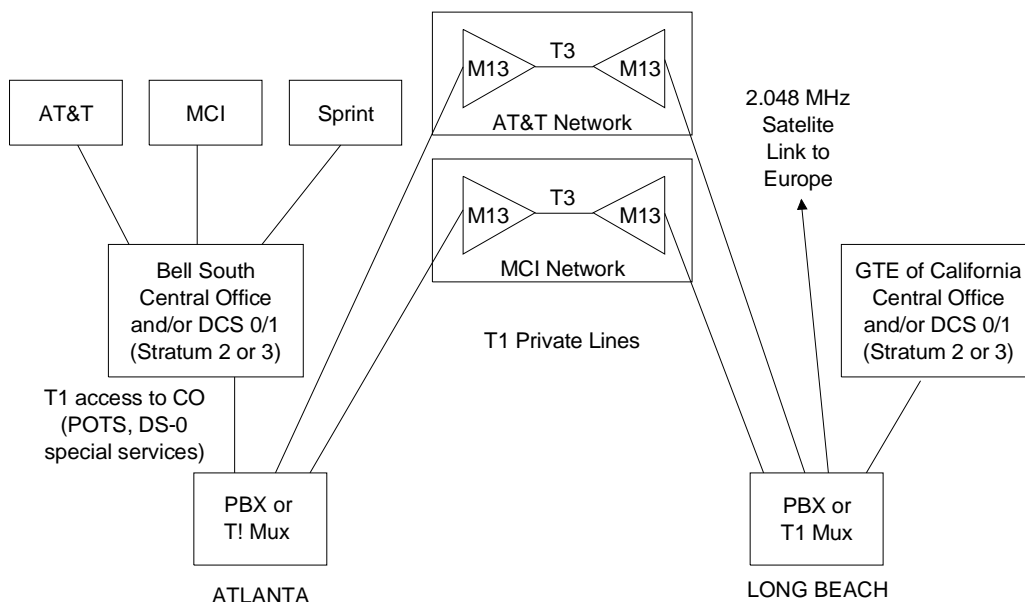

Figure 4 - Example of Private Network

A system synchronizer must meet a variety of requirements, including selecting a timing reference source to use, maintaining lock to that reference clock, tolerating impairments on the timing reference, attenuating jitter, and free-running at a specified frequency accuracy when all references are lost.

These requirements differ between systems that relay (or transfer) timing (Stratum 4 Type I, Stratum 3 or Stratum 2) and loop-timed systems (Stratum 4 Type II) that don't relay timing. In particular, in systems which relay timing, the maximum rate of change of the system synchronizer frequency must be limited (see definition of Maximum Time Interval Error in ANSI T1.101). Controlling the rate of change is important because if the synchronizer oscillates (i.e., overshoots then undershoots the target frequency) while acquiring lock or during degraded timing reference conditions, that frequency deviation can be relayed through all successive downstream systems.

An overview of Stratum requirements are given below:

Stratum 2: This is the highest performance level available to CPE, and is rarely used (except in very sophisticated CPE). The clock must have a free-running accuracy of 1,544,000 ± 0.025 Hz. Also a Stratum 2 clock must provide holdover to limit drift to no more than 0.0001 ppm from the last known reference frequency within 24 hours of the time the reference becomes unavailable. Redundancy is required so the synchronizer must be duplicated. The external reference clock must be Stratum 1 or 2 since the pull-in range of the synchronizer is only ± 0.05 Hz.

Stratum 3: Stratum three clocks must also be duplicated, and are required to have free-running accuracy of 1,544,000 ± 7.1 Hz. Stratum 3 requirements limit frequency drift during holdover to allow ≤255 DS1 frame slips in the first 24 hours. The external reference clock must be a Stratum 3 or better clock, since the pull in range is only ± 15 Hz.

Stratum 4 (Type I): This is normally the stratum level chosen for multi-T1 line CPE. Free-running frequency requirements are 1,544,000 ± 50 Hz, and the pull-in range is ± 100 Hz Additional pull-in range is recommended, since older equipment in the telephone company buildings are allowed to have output frequency variations of T1 ± 130 ppm (± 200 Hz). Type I does not require redundant clock hardware. However, just like Stratum 2 and 3 clocks, the Type I clock must gracefully handle rearrangements. "Gracefully" means that the rearrangements must not result in abrupt phase discontinuities in the output clock. Examples of rearrangements are:

1) a switching of timing reference (for example from T1 card "n" to card "n+1"),
2) automatic protection switch to a redundant synchronizer card, or
3) any change in clock mode

Stratum 4 (Type II): This is the lowest grade clock available, and normally would be used only in small systems (such as a single-T1 line system). Type II synchronizers cannot be used if the system transfers timing to a downstream system. The synchronizer clock output is required only during successful loop-timing. When loop-timing fails, the system no longer is required to transmit toward the network. There is no requirement on the synchronizer's free-running frequency. Nor is there any requirement to handle rearrangements. During loop-timing the frequency of the clock must be 1,544,000 ± 50 Hz. The pull-in range is 1,544,000 ± 100 Hz.

JITTER IN SYNCHRONOUS NETWORKS

Although the frequency of the T1 signal is well controlled (32 ppm), the signal itself can contain significant amount of jitter. *Jitter* is defined as a short-term variations in the position in time of a signal (or bit) from its ideal position. In other words, the bit could arrive at a receiver slightly sooner or later than expected. Jitter can cause bit errors or other impairments to occur. When jitter occurs at a 10 Hz or lower rate, it is arbitrarily defined as *wander*. Jitter and wander can be simultaneously present. Figure 5 shows high-frequency jitter superimposed upon low-frequency wander.

Due to the characteristics of the phase-lock loop technology used in trunk cards and system synchronizers, jitter can be more readily filtered than can wander (Figure 6). Normally, most wander which is input on a timing reference source will be relayed to all downstream systems.

A primary source of wander is frequency instability in a system synchronizer during a switch between timing reference sources. Why does the synchronizer need to switch between timing sources? One major source of errors is upstream clock rearrangement (typically two per day will reach Stratum 4 nodes). While acquiring lock to the new source, wander can be created and sent to all the downstream timing sources. Stratum 2 clocks can take hours/minutes to reacquire lock
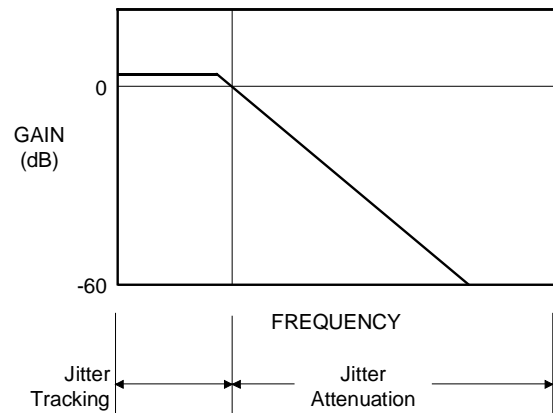

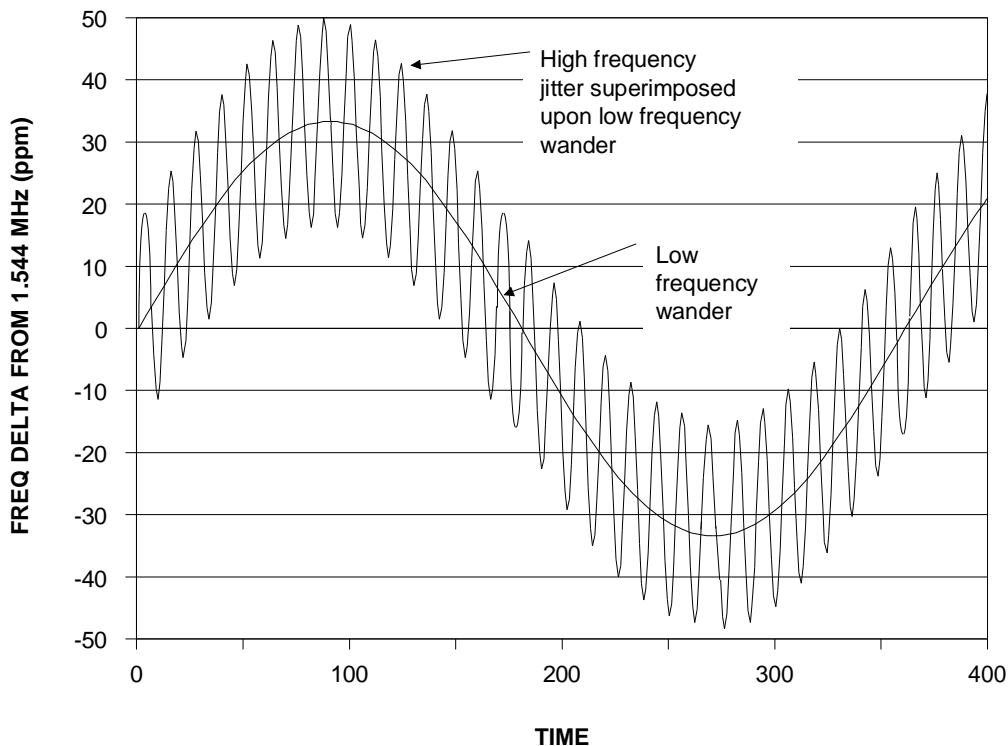**Figure 6 - Jitter Transfer of PLLs**

Figure 5 - Jitter Superimposed Upon Wander

after a switch. Stratum 4 clocks take seconds to reacquire lock. This wander can be amplified by each of the downstream synchronizers, causing an ever-larger wave-of-wander to spread out through the network.

The sources of jitter include data-dependent jitter introduced by line repeaters, jitter created by asynchronous multiplexors (such as M13 muxes) and jitter (phase) hits attributed to transient behavior of clock sources or other sudden changes in transmission facilities. Jitter is not a significant problem for equipment as long the trunk card and system synchronizer meets specified jitter tolerance requirements.

62411 JITTER REQUIREMENTS

The motivation behind 62411 jitter specifications is to ensure robust operation of CPE and the network despite the presence of jitter. The CPE is at the end of a local loop (whose line repeaters can create jitter) and is far removed from the network's Stratum 1 clock. To work robustly, CPE

must be extremely jitter tolerant, and must remove significant amounts of the jitter received over the local loop before re-transmitting back toward the network.

Before one can understand 62411 jitter requirements, it is necessary to understand some basic terminology (Figure 7). 62411 measurements are made at the NI (Network Interface) which is the demarcation line between CPE and the AT&T line. The *CSU* (Channel Service Unit) can be either a stand-alone unit (purchased from a CSU manufacturer), or may be integrated into the T1 trunk cards of the system. AT&T views the CSU and system as one entity for 62411 testing. Jitter generated or attenuated by either the CSU or the system will impact jitter testing at the NI.

From 62411's point of view, the system has two critical components, the *digital line termination of the DTE* (more commonly referred to as a T1 trunk card), and a *synchronizer (or clock recovery circuit)* (more commonly referred to as a
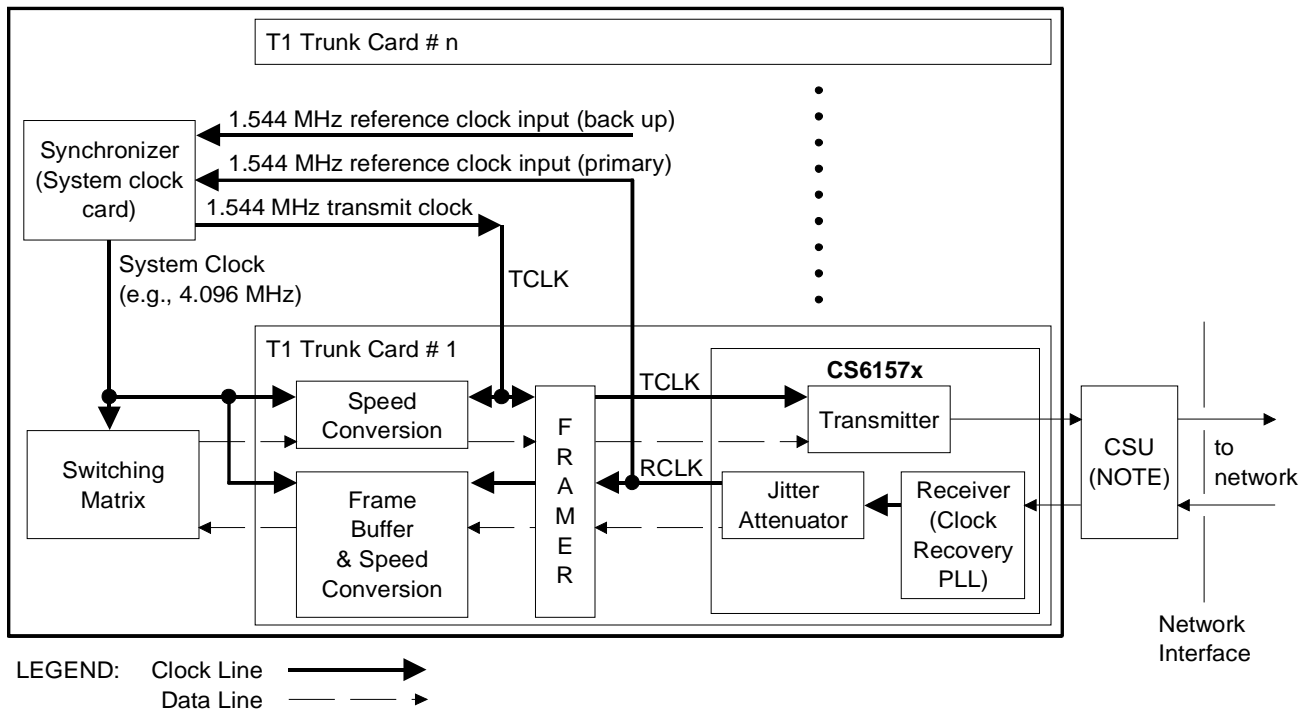

Figure 7 - Basic Terminology of 62411

system clock). 62411's use of the term "clock recovery circuit" should not cause the designer to think that 62411 is referring to a clock recovery circuit in a T1 trunk card. Rather the clock recovery circuit refers to the clock which drives the system and backplane. Typically the synchronizer outputs a system clock at 4.096 or 8.192 MHz and is phase locked to one of the incoming T1 lines. The assumption behind 62411 is that the synchronizer is the source of jitter attenuation. This assumption is generally appropriate unless the system has a 1.544 MHz backplane or only one T1 line. In this case the synchronizer may in fact not exist as a independent sub-circuit in the system, but rather may be imbedded in the T1 trunk card.

62411 has three types of jitter specifications: input jitter tolerance (also known as jitter accommodation), output jitter generation (also known as intrinsic jitter or additive jitter), and jitter transfer function (also known as jitter attenuation). The requirements for each specification are discussed below.

All 62411 jitter tests are made with a 2^{20}-1 QRTS (Quasi-Random Test Signal) modified to ensure that no more than 14 consecutive zeros are output. The Application Note, "Jitter Testing Procedures for compliance with AT&T 62411" gives a tutorial on test procedures.

Input Jitter Tolerance: Jitter tolerance is specified for both the T1 trunk card and for the synchronizer. In fact, the tolerance requirements are more stringent for the synchronizer than the line card. The more stringent specifications ensure that the synchronizer, which is a critical system element, is extremely robust. The tolerance curves are shown in Figure 8.

Tolerance is specified for jitter frequencies from DC to 100 kHz, and is normally measured at spot frequencies. At a given frequency, the jitter amplitude is increased until bit errors are observed.
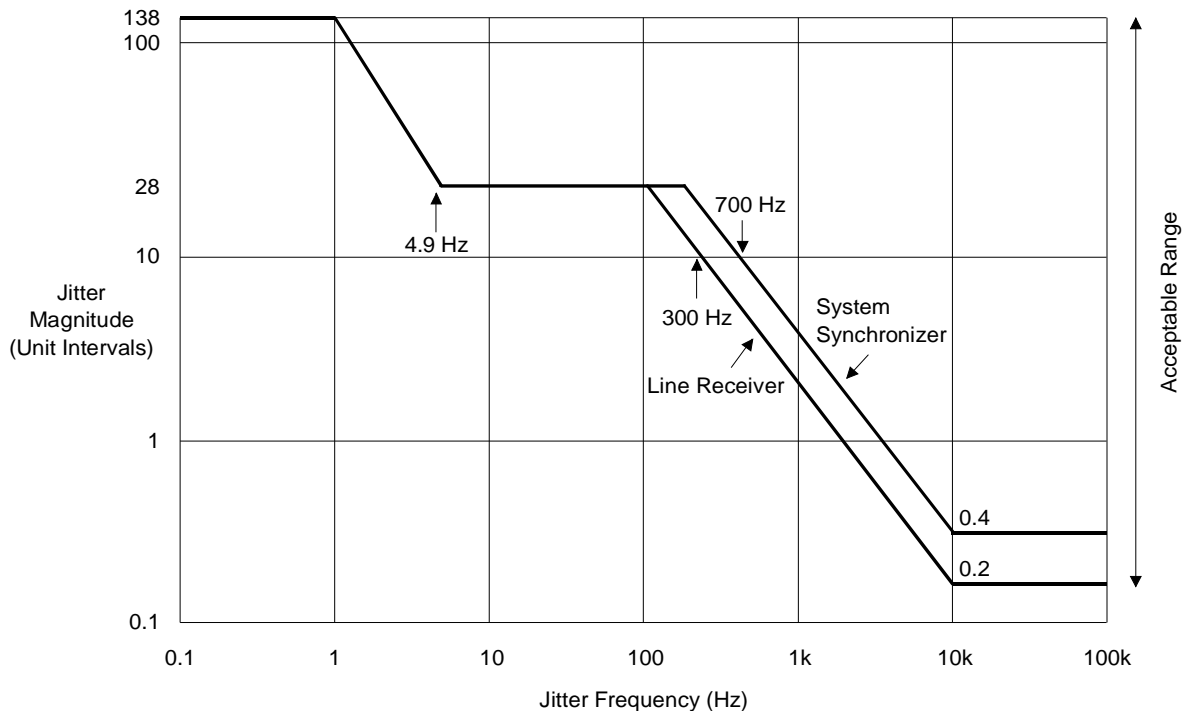

Figure 8 - 1990 AT&T 62411 Jitter Tolerance Requirement

The jitter tolerance of the synchronizer is relevant for all loop-timed systems. However, testing of the synchronizer is indirect since the synchronizer is not directly observable at the NI, and since the trunk card is not required to tolerate as much jitter amplitude as the synchronizer. Rather, the synchronizer is said to fail if loss of synchronization at the NI is observed, or if the system is observed to permanently switch to a back-up (internally-generated) clock. The switch to a back-up clock becomes apparent because of a shift to a new frequency (which may differ from the old frequency by only a few Hz).

62411 states that jitter (wander) tolerance at low frequencies can be limited by the size of a buffer. The required tolerance is one T1 frame (193 bits) plus hysteresis (138 bits). This model assumes that a frame buffer is being employed. Frame buffers allow a system to synchronize multiple, incoming, asynchronous T1 lines at the bit, channel and frame levels. Each incoming T1 line is fed into its own frame buffer (Figure 7). All of the frame buffers are emptied into the switching matrix with all frame bits aligned. If two incoming T1 lines have frequency offsets, one frame buffer will periodically output a frame twice (to compensate for a low input frequency) or will drop a frame (never outputing it to compensate for a high input frequency). This buffer adjustment is referred to as a *controlled frame slip*.

The 193 + 138 = 331 bit-length requirement applies to frame buffers making controlled frame slips. The goal is to have frame slips occur only as a result of long term frequency offsets, not due to jitter (wander). *However, only 138 bits are required in a jitter (hysteresis) buffer that does not perform frame slips*. In this case, the clock used to take data out of the buffer tracks the long term average input frequency so that the Jitter produced by modulating a QRTS with white noise bandlimited between 10Hz and 300Hz must be attenuated by 25dB when measured using an 8kHz to 40kHz filter. This performance is important. FIFO depth requirements are defined by the buffer length necessary to tolerate the jitter (wander) without overflow/underflow. The exact length needed to meet the requirements of Figure 8 depends upon the jitter transfer function of the PLL (Figure 6). The PLL will track jitter or wander at low frequencies, and only requires a FIFO long enough to contain the maximum input jitter amplitude possible in the jitter attenuation range.

Output Jitter Generation: Output jitter generation is a measurement of how much jitter is output by the system when no jitter is input to the system (using a QRTS). The requirements are measured using band-pass filters to measure jitter in various frequency bands, as shown below:

Filter applied	Maximum output jitter
None (Broadband)	0.05 UI peak-peak
10 Hz to 40 kHz	0.025 UI peak-peak
8 kHz to 40 kHz	0.025 UI peak-peak
10 Hz to 8 kHz	0.02 UI peak-peak

Jitter Transfer Function: Jitter transfer describes the ratio of input jitter and output jitter at a given jitter frequency using the following equation:

$$\text{Jitter Transfer(dB)} = 20 \log (\text{jitter output/jitter input})$$

The jitter input and output are measured in peak-to-peak UIs.

If a system contains a Stratum 4 clock, the jitter transfer function must meet the requirements of Figure 9. Figure 9 is irrelevant for systems with Stratum 2 or 3 clocks. Those clocks will attenuate jitter to the "output jitter generation" levels defined in the preceding section.

Since the system synchronizer has its basic frequency locked to the T1 rate of one particular
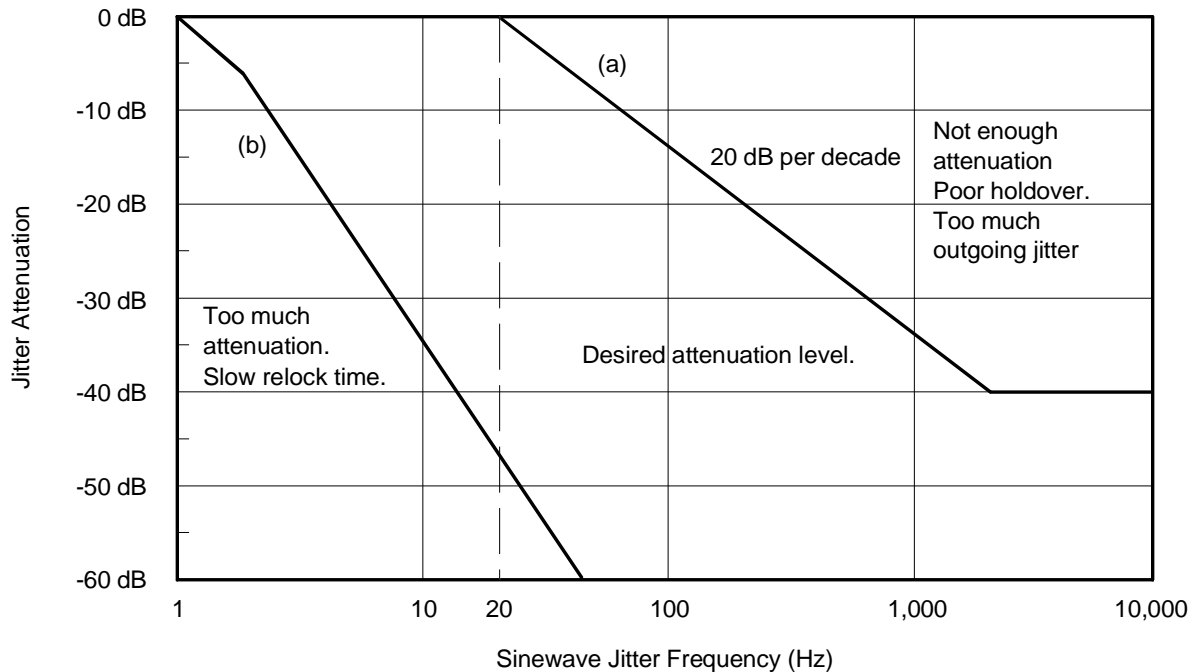
Figure 9 - 1990 AT&T 62411 Jitter Transfer Requirement

line card, the attenuation test is at the NI associated with that particular line card. QRTS input jitter is applied at spot frequencies in the band 10 Hz to 40 kHz. The input jitter level is three-quarters of the system's jitter tolerance. Output jitter is measured using narrow frequency windows (1 Hz in 10-100 Hz range, 4 Hz in 100-1000 Hz range and 10 Hz above 1kHz). Jitter Transfer must be between the two limits, (a) and (b), shown in Figure 9. In the vicinity of 10 kHz, the output jitter level, after attenuation, must typically be in the range of 0.001 to 0.002 UIs (or 0.65 to 1.3 ns).

The curve (b) indicates an attenuation level which must not be exceeded. Too much attenuation can be harmful. The negative effects can include increasing the number of frame slips that occur in the central office, and increasing the time needed to reacquire lock after a loss of signal condition.

Jitter peaking is not allowed above 10 Hz. Peaking is defined as a higher amplitude of output jitter than input jitter (at any frequency, not just

at the frequency under test). Peaking is especially undesirable at low frequencies which lie within the tracking range (and not the attenuation range) of downstream PLLs. Analog PLLs can easily have peaking at the "knee" of the jitter transfer curve (just before attenuation starts).

In digitally controlled PLLs, another phenomenon, jitter aliasing, can often be observed. With jitter aliasing, jitter which appears at one frequency is attenuated and shifted to another frequency. Jitter aliasing is especially undesirable when the frequency is shifted down to the tracking range of the downstream PLLs. The 1990 version of 62411 has added language to address aliasing.

T1 TRUNK CARD DESIGN CONSIDERATIONS

A number of architectural choices face the trunk card designer. One is whether to integrate the CSU onto the trunk card.

Another design choice is whether to use a Line Interface IC that provides jitter attenuation. For Stratum 4 Type II systems that don't relay timing and utilize a 1.544 MHz backplane clock, the answer is simple. By using a CS61574, CS61574A, or CS61575 on the trunk card, the Type II clock is provided on the trunk card by the line interface RCLK output, and no separate synchronizer card needs to be designed. If the Type II system has multiple T1 lines, then the line interface IC's may be connected as shown in Figure 10. In this configuration, the master 1.544 MHz system clock is provided by the top CS61574A or CS61575 in the figure. The system clock will normally be recovered from the T1 line connected to the top line interface until the line interface enters the Loss of Signal state. Upon Loss of Signal (175 ±75 zeros), the CS61574A and CS61575 will substitute ACLKI (if present) for the recovered clock at the input of the jitter attenuator. The jitter attenuator filters any phase hits resulting from the switch before the clock is output on RCLK. In this configuration the RCLK output of the top line interface will always remain locked to one of the T1 recovered clocks. Each line interface that experiences Loss of Signal will pass on the recovered clock from the line interface preceeding it in the chain. When the master experiences Loss of Signal, the line interface nearest the master that is not in the Loss of Signal State will provide the 1.544 MHz system clock. If all line interfaces experience Loss of Signal, the reference clock is used instead. This configuration is advantageous for multi-line loop timed applications because it provides a jitter-free 1.544 MHz line recovered timing reference (ideally traceable to Stratum 1) even if there is only one functioning T1 line reference available. This capability helps to maximize the system's ability to successfully maintain loop timing on as many T1 lines possible.

However, for three reasons, the configuration shown in Figure 10 is not appropriate for systems which transfer timing (i.e., Stratum 4 Type

I or Stratum 2 or 3). First, the CS61574A and CS61575 switch to a secondary reference clock upon Loss of Signal and not upon the conditions outlined in 62411 (summarized previously in the "Synchronization of Digital Networks" section) for systems that transfer timing. Furthermore, the CS61574A and CS61575 may not meet the maximum time interval error (MTIE $\leq 1\mu s$) and phase change slope (81 ns in any 1.326 ms) requirements which apply during rearrangements. Systems which transfer timing must utilize a system synchronizer and the control circuitry requried to judiciously switch between reference sources, and the synchronizer's 1.544 MHz PLL is responsible for insuring the MTIE and phase slope requirements are met. Finally, the configuration in Figure 10 does not necessarily enforce a well defined network synchronization plan. Such a plan is essential in a network throughout which timing is transferred. A synchronization plan defines an appropriate manner in which timing rearrangements may be made without creating network timing problems such as self-locked timing islands (refer back to Figure 2). The configuration in Figure 10 will not necessarily prevent the fallback to an inappropriate timing reference (eg., a line already loop timed to
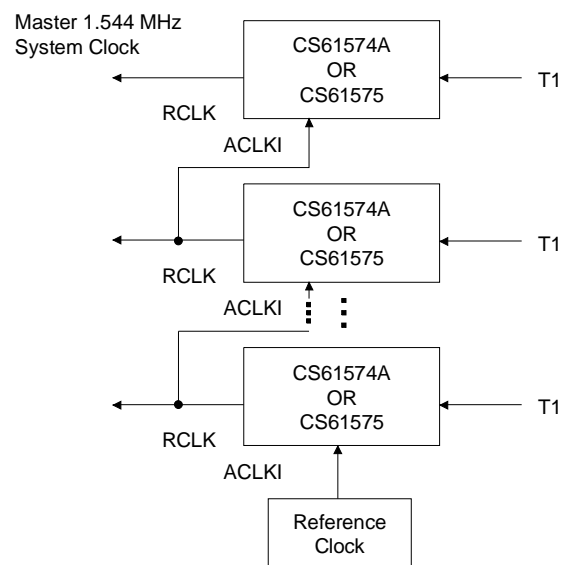

Figure 10 - Cross Coupling of Recovered Clocks (Stratum 4 Type II)

this node). Therefore, when using the configuration shown in Figure 10, it is important that there be a field-installable definition of how the chain is configured.

The use of a Line Interface IC with jitter attenuation also provides advantages for Stratum 4 Type I systems and Stratum 2 or 3 systems. By using a CS61574, CS61574A or CS61575 on the trunk card, the synchronizer becomes easier to design. Since the CS61574 exceeds jitter tolerance requirements and outputs an essentialy jitter-free recovered-clock, the synchronizer has reduced jitter tolerance requirements, and will see phase hits only when the switch is made between primary and secondary references. This frees the synchronizer designer to concentrate on meeting the holdover frequency, MTIE and reference switching criteria requirements.

SYSTEM CLOCK CARD DESIGN CONSIDERATIONS

As shown in Figure 7, the synchronizer provides the master clock for the system. It typically runs at a multiple of 8 kHz (the PCM frame rate), and generates a backplane clock at the rate of 2.048 MHz, 4.096 or 8.192 MHz.

A block diagram of a synchronizer is shown in Figure 11. This diagram shows multiple T1 lines being available as reference clocks for the synchronizer. A MUX is used to select one of the lines. The output of the MUX is fed through a CS61574 used only as jitter attenuator (accomplished by placing the CS61574 in local loop back mode). This jitter attenuation stage shields the 1.544 MHz Phase Locked Loop (PLL) from significant phase and frequency discontinuities
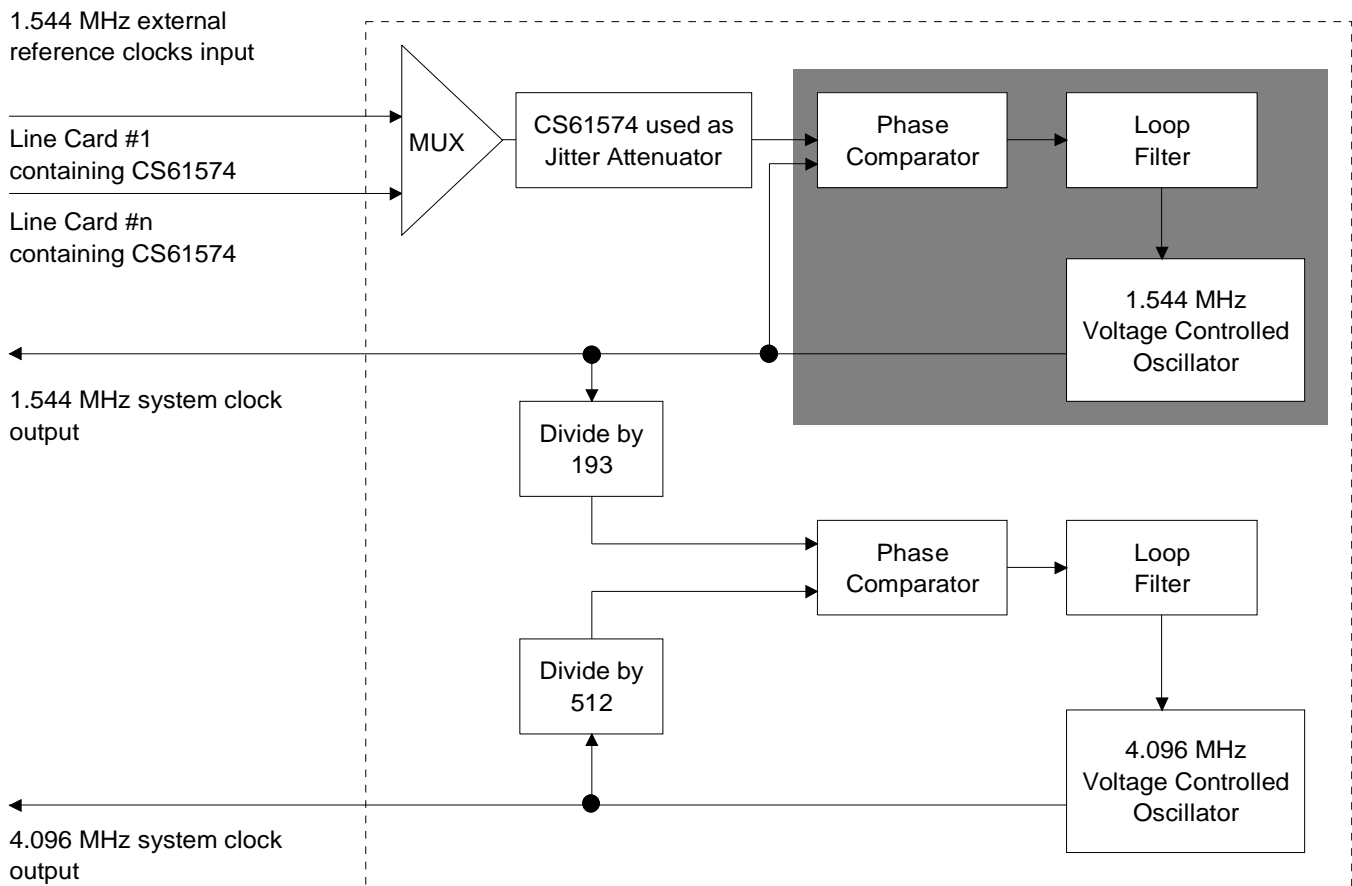

Figure 11 - Block Diagram of System Synchronizer

which can result from either the loss of signal on the active reference clock input, or the switching of the MUX. The 1.544 MHz PLL provides a clean 1.544 MHz output clock for transmitting back toward the network. The system backplane clock is also locked to the 1.544 MHz clock using a second PLL. This system backplane clock does not need to meet 62411 jitter transfer specifications, since it is not used to clock data back to the network (Figure 7).

The system synchronizer clock card design is greatly simplified in systems which employ CS61574's, CS61574A's or CS61575's on the T1 trunk cards. In such a system, jitter performance requirements (jitter transfer, jitter tolerance and intrinsic jitter) are met by the line interface IC. The 1.544 MHz PLL does not need to be designed with the narrow closed loop bandwidth required for 62411 (i.e., jitter attenuation starting at 6 Hz). Wider closed loop bandwidth simplifies the design of the PLL and promotes stability. For example, VCO (Voltage Controlled Oscillator) phase noise is suppressed in a system which has wider closed loop bandwidth. The fact that VCO stability is improved by wider closed loop bandwidth (feedback) means that it is often possible to use a VCO instead of a more expensive VCXO (Voltage Controlled Crystal Oscillator) in the 1.544 MHz synchronizer PLL of a Stratum 4 (Type I) system. VCXO based designs offer excellent stability even in low bandwidth applications and are recommended for use in Stratum 3 systems. Although the synchronizer does not need to be the primary source of jitter attenuation in a system utilizing line interface IC's with jitter attenuation, the synchronizer PLL should not exhibit jitter peaking (jitter gain). It is

easier to design a PLL which does not exhibit jitter peaking when a significant amount of jitter attenuation is not required.

Not only does the use of a CS61575 based trunk card design relax the jitter performance constraints imposed on the synchronizer's 1.544 MHz PLL, but it also eliminates the holdover (free-running) frequency accuracy requirement. Because the CS61574A and CS61575 will gracefully transition from the line recovered clock to a reference clock upon Loss of Signal, the 1.544 MHz PLL will always have an appropriate timing reference to stay phase locked to and will never need to free-run. The holdover frequency accuracy of the system will be that of the 1.544 MHz reference clock supplied to the line interface IC's.

With the other design constraints greatly relaxed, the PLL may be designed specifically to meet the dynamic performance required by 62411. To do this, the PLL must be designed with adequate damping to meet the Maximum Time Interval Error (MTIE) and phase change slope requirements discussed earlier.

There are several types of components from which the synchronizer PLL's may be constructed. Monolithic PLL IC's like the 74HC4046A offer excellent performance and minimal circuit board space requirements. For more information on PLL design using the 74HC4046A, refer to (2) - (4). The free design program (3) offered by Signetics is an excellent tool for evaluating and optimizing a 74HC4046A PLL Design.

Recommended References on PLL design are:

1. Gerdner, Floyd M: <u>Phaselock Techniques</u>, John Wiley & Sons, New York, 1979. (ISBN: 0-471-04294-3)

2. Volgers, B.: "Phase-Locked Loop Circuits: 74HC/HCT4046A & 74HC/HCTR7046A HCMOS Designer's Guide", Signetics/Phillips Components. (Ordering Code: 98-2908-350)

3. Signetics/Phillips Components: <u>HCMOS Phase-Locked-Loop Design Program</u> . (For IBM compatible computers)

4. Austin, W. M.: "CMOS Phase-Locked-Loop Applications Using the CD54/74HC/HCT 7046A", ICAN - 8823, Harris Semiconductor.

• Notes •