White Paper

TDM Timing

By Dr. Yaakov Stein, Chief Scientist, RAD Data Communications

This paper presents a brief overview of the theory and practice of timing in pure TDM and TDMoIP networks.

The Importance of TDM Timing

TDM signals are **isochronous**, meaning that the time between two consecutive bits is theoretically always the same. This time is called the **unit interval** (UI); for T1 signals the UI is defined to be 647 nanoseconds, and for E1 the standards dictate 488 nanoseconds. In order to maintain isochronicity and to remain within tolerances specified by recognized standards, a TDM source must employ a highly stable and accurate clock.

The stringent clock requirements are not capriciously dictated by standard bodies; rather, they are critical to the proper functioning of a high-speed TDM network. Consider a TDM receiver utilizing its own clock when converting the physical signal back into a bit-stream. If the receive clock runs at precisely the same rate as the source clock, then the receiver need only determine the optimal sampling phase. However, with any mismatch of clock rates, no matter how small, **bit slips** will eventually occur. For example, if the receive clock is slower than the source clock by one **part per million** (ppm), then the receiver will output 999,999 bits for every 1,000,000 bits sent, thus deleting one bit. Similarly, if the receive clock is faster than the source clock by one **part per billion** (ppb), the receiver will insert a spurious bit every billion bits. One bit slip every million bits may seem acceptable at first glance, but translates to a catastrophic two errors per second for a 2 Mbps E1 signal. ITU-T recommendations permit a few bit slips per day for a low-rate 64 kbps channel, but strive to prohibit bit slips entirely for higher-rate TDM signals.

Temperature changes, imperfections in materials, aging, and external influences will inevitably affect a clock's rate, whether that clock is atomic, quartz crystal, or pendulum based. Hence no clock will remain at precisely the same rate forever, and no two physical clocks will run at exactly the same rate for extended periods of time. In order to eliminate bit slips, we must ensure both that the long-term average UI of source and receive clocks are identical (any rate difference, no matter how small, will eventually accumulate up to a bit slip), and that its short-term deviations from the average are appropriately bounded.

The variation of a clock's rate over time is conventionally divided into two components, jitter and wander. **Wander** expresses slow, smooth drifting of clock rate due to temperature changes, aging and slaving inaccuracies; while jitter conveys fast, erratic jumps in UI caused by phase noise phenomena and bit-stuffing mechanisms. The border between the two components is conventionally set at 10 Hz. In order to eliminate bit slips, the standards impose strict limits on tolerable jitter and wander of TDM clocks.

Timing Distribution

How are the absolute accuracy and variability (jitter and wander) limits attained? Conventional TDM networks rely on hierarchical distribution of timing. Somewhere in every TDM network there is at least one extremely accurate Primary Reference Clock (PRC) or Primary Reference Source (PRS), with long-term accuracy of one part in 10^{11} as compared to UTC, the world time standard. The characteristics of this clock, known in North America (see T1.101) as a stratum 1 clock, are described in ITU-T Recommendation G.811. The accuracy of today's atomic clocks is significantly better than that required by G.811.

The PRC is the master clock from which all other TDM clocks in the network directly or indirectly derive their timing. This hierarchy of time synchronization is essential for the proper functioning of the network as a whole. A clock that ultimately derives its rate from the PRC is said to be traceable to that PRC.

The distribution of clock information from the PRC towards other clocks exploits the isochronous nature of TDM signals. When a TDM signal clocked by the PRC is received, the receiver's local clock is compared to the observed UI of the received bits. If the observed rate is higher (lower) than the present rate, the receiver increases (decreases) its local clock's rate somewhat. This correction ensures that the long-term average UI will be correct but introduces jitter, since the local clock's rate was rapidly changed. In addition, wander is introduced, as it can take a long time until small timing discrepancies become apparent and can be compensated. Hence the secondary clock will be somewhat degraded relative to the primary one. In North America this is called a stratum 2 clock, and is historically used at tandem offices. This idea of a slave clock deriving its timing from a master clock is continued in hierarchical fashion. TDM signals whose source utilizes a stratum 2 clock are received by TDM receivers that adapt their local clocks to track the stratum 2 clock. Such stratum 3 clocks, historically used at local exchanges, already lead to appreciable bit slips. Finally, customer equipment such as channel banks may use stratum 4 clocks; such clocks may lead to bit slips every few seconds.

When a slave clock loses the master clock signal it is expected to go into holdover mode. In this mode it attempts to maintain clock accuracy, although it no longer has access to its reference. For this reason the standard requirements are lower in holdover mode, e.g. a Stratum 2 clock that during normal operation is required to maintain a long-term accuracy

of one part in 10^{11}, is only required to keep up one part in $1.6 * 10^{-8}$ under holdover conditions. Stratum 4 clocks need not have any holdover capabilities.

The process used by slave clocks to mimic the master clock's rate is actually more sophisticated than that described above. Due to noise and measurement inaccuracies, the receiver needs to observe the difference between its local clock and the observed TDM for some time before deciding to change its clock's rate. The task of clock recovery can thus be seen to be a kind of averaging process that negates the effect of the random variations and captures the average rate of transmission of the original bit stream. A phase locked loop (PLL) is well suited for this task because it can lock onto the average bit rate, regenerating a clean clock signal that closely approximates the original bit rate. If the incoming TDM is disrupted, and the secondary clock can no longer remain locked onto the superior source clock; it must continue in holdover mode to supply timing to yet lower-tier clocks.

ITU-T Timing Recommendations

ITU-T recommendations define a somewhat different hierarchy of clock specifications from those of T1.101. At the bottom of the pyramid are clocks conforming to G.823 (for signals belonging to the E1-hierarchy) or G.824 (for T1-hierarchy signals). These standards define the permissible output jitter and wander for two levels, the lower being traffic interface and the more stable being synchronization interface. The former may not be used as a master clock, while the latter may.

ITU-T Recommendation G.812 defines a slave clock commonly known as a synchronization supply unit (SSU). The SSU's master may be a PRC, another SSU or an SEC (see below). The SSU fulfills two objectives: it filters out jitter and relatively short-term wander, and it provides a highly accurate clock for holdover scenarios. An SSU may be designed as a Stand Alone Synchronization Equipment (SASE), known in North America as a Building Integrated Timing Source (BITS) clock, or as part of a traffic handling network element, such as a digital cross-connect. Up to 10 SSUs may be chained without overly degrading performance; when this is done, the intermediate SSUs are called transit SSUs (SSU-T), while the final one is called a local SSU (SSU-L).

ITU-T Recommendation G.813 defines a slave clock known as an SDH Equipment Clock (SEC). A SEC is required to have fairly good, but not excellent, timing accuracy while in

holdover mode (thus allowing the use of relatively inexpensive crystal local oscillators), but stringent jitter and wander generation (in order to enable chaining of multiple SECs). An SSU can be employed after a chain of SECs to counteract the accumulation of timing inaccuracies.

The newly released (first version) ITU-T G.8261 (former G.pactiming) "*Timing and Synchronization Aspects in Packet Networks*" focuses on synchronization issues that are related to transporting of synchronous signals (mostly TDM) over inherently asynchronous Packet Switched Networks or PSN (i.e. Circuit Emulation) and timing distribution over packet switched network in general. Later in this paper we discuss this new standard in more detail.

Measuring Accuracy, Jitter and Wander

We have discussed the standards that define the accuracy, jitter and wander of TDM clocks, but have yet to explain how these physical characteristics are measured. Frequency accuracy is usually specified as the fractional offset from the desired value, and is only meaningful when the time period over which it is measured is specified. Frequency accuracy is the crucial measure of PRCs, and is the main measure of a timing source's stability during holdover. Accuracy is also specified by standards for cellular transmitter stations, because frequency differences can impair hand-offs between cells. However, accuracy is not usually specified for slave clocks, as their function is to follow the timing of their master, rather than remain close to an absolute timing source.

As jitter is basically an instantaneous phenomenon, its relative magnitude is an adequate measure; in contrast, wander is time-dependent, and thus its specification must take into account the time interval over which it is measured.

Jitter is conventionally measured in UIpp, that is, the difference between maximum and minimum time intervals in units of the nominal UI. For example, for an E1 signal with a UI of 488 nanoseconds, if the maximum interval were 500 nanoseconds and the minimum 476, the jitter would be (500-476)/488 = 0.05 UIpp.

Specification of wander is more complex than jitter, as differing time intervals must be taken into account. Two measures are commonly used, namely the MTIE and TDEV. MTIE(τ) (Maximum Time Interval Error) is defined as the maximum peak-to-peak time variation of the clock measured relative to an ideal timing signal, within observation time τ for all

observation times of that length within the measurement period T. If there is a constant frequency offset during the time we measure MTIE, the MTIE will be linear in τ; if there is a linear frequency drift, the MTIE(τ) will diverge. When we are interested in how well a slave clock follows its master, we measure MRTIE (Maximum Relative Time Interval Error), which is like MTIE but with the comparison being made to the master clock, rather than to an ideal timing source.

TDEV (Time Deviation) is similar to MTIE, but rather than specifying the maximum peak-to-peak time error over a time duration τ, we calculate the expected time variation (in nanoseconds). It can be shown that TDEV is directly related to the power spectral density of the phase deviation of the timing signal. If there is a constant frequency offset during the time we measure TDEV, the TDEV is unaffected; if there is a linear frequency drift, the TDEV will behave linearly. TDEV is superior to MTIE when the spectral content of the phase noise of a signal is of interest.

Attaining Timing Goals in TDM Networks

TDM networks utilize several mechanisms in order to guarantee that their timing conforms to the requisite standards. First, the TDM physical layer signals are designed to ensure that the basic single-bit duration is readily identified, irrespective of the data being transferred. For example, coding is used to ensure transitions (e.g. B8ZS) and then a line code (e.g. AMI and HDB3) is used which has extra transitions in order to eliminate long runs without transitions. Note that conventional AMI line code may fail to have sufficient marks, *i.e.*, "1's," to permit recovery of the incoming clock, and synchronization is lost. This happens when there are too many consecutive zeros in the user data being transported. To prevent loss of synchronization when a long string of zeros is present in the user data, deliberate bipolar violations are inserted into the line code, to create a sufficient number of marks to maintain synchronization.

Second, TDM slave clocks attenuate jitter and wander that accumulate on their inputs. They do this by using sophisticated circuitry in the Line Interface Unit (LIU). The precise instant that a TDM bit is received equals the time it was transmitted plus the nominal propagation time over the physical channel plus a zero mean stochastic component attributable to temperature, oscillator noise, regenerator jitter, justification effects, etc. In order to eliminate the stochastic component, some sort of averaging process must be carried out to capture the average rate of transmission of the original bit stream. A Phase Locked Loop

(PLL) is well suited for this task because it can lock onto the average bit rate, regenerating a clean clock signal that approximates the original bit rate.

Third, the aforementioned hierarchy of TDM clocks forms a synchronization network, whereby every clock in the TDM network is traceable to the PRC. The various synchronization network elements have carefully designed jitter/wander tolerance, transfer, generation and output characteristics. Tolerance refers to the ability of an element to tolerate inaccuracies at its input, while transfer describes the cleaning up of input inaccuracies. Generation expresses the intrinsic jitter and wander contributed by the element itself, and the output jitter/wander are the result of the input deficiencies and all of the above.

Attaining TDM Timing Goals for TDMoIP

TDMoIP is a relatively new technology that enables transport of TDM traffic over Packet Switched Networks (PSNs), such as Ethernet, IP and MPLS. The TDM bit-stream is segmented and encapsulated, packets containing TDM payload traverse the PSN, and at the far end the TDM signal must be reconstructed, emulating the original TDM transport.

The major technical barrier to TDM emulation is clock recovery. While isochronous TDM networks inherently deliver timing along with the data, and even ATM networks provide a physical layer clock reference, asynchronous PSNs do not transfer any timing information whatsoever.

As will be explained shortly, matters are made worse due to Packet Loss (PL) and Packet Delay Variation (PDV). In order to reconstruct TDM timing, sophisticated clock recovery mechanisms are required in order to achieve the desired timing accuracy in the presence of packet delay variation and packet loss.

TDMoIP packets injected into a PSN at a constant rate reach their destination with delay that has a random component, known as PDV. When emulating TDM on such a network, it is possible to overcome this randomness by placing the TDM into a 'jitter buffer' from which the data can be read out a constant rate for delivery to TDM end-user equipment. The problem is that the time reference of the TDM source is no longer available, and the precise rate at which the data is to be 'clocked out' of the jitter buffer is hence unknown.

In certain cases, timing may be derived from accurate clocks at both endpoints, for example, if the TDMoIP replaces a link in an otherwise isochronous network, or if atomic clocks or GPS receivers are available at both sides. However, often the only alternative is to attempt to recover the clock based only on the TDMoIP traffic. This is possible since the source TDM device is producing bits at a constant rate determined by its local clock. We receive these bits in packets that suffer PDV that can be considered a zero-mean random process. The task of clock recovery can thus, once again, be seen to be a kind of averaging process that negates the effect of the random PDV and captures the average rate of transmission of the original bit stream. As in the pure TDM case, a PLL is well suited for this task, but now the jitter and wander are orders of magnitude higher.

One conventional means of clock recovery is based on adapting a local clock based on the level of the receiver's jitter buffer. To understand the operation of the conventional mechanism, let us assume for the moment that there is no PDV but that the local clock is initially lower in frequency than the source clock. The writing into the jitter buffer occurs faster than it is emptied and thus the fill-level starts to rise. This rise is detected and compensated by increasing the frequency of the local clock. When in addition to clock discrepancy there is PDV, the jitter buffer level no longer smoothly rises or falls, but rather fluctuates wildly about its slowly changing average level. By using a PLL that locks onto the average rate, any frequency discrepancy between the source and destination clocks is eventually compensated, and the receiver's jitter buffer will settle on the level corresponding to precise frequency alignment between the two clocks.

This conventional PLL has several faults. First, the PLL must observe the sequence of level positions for a long period before it can lock onto the source clock, and hence the scheme exhibits lengthy convergence time. Second, the jitter buffer level may settle down far from its desired position at the buffer center, thus making it vulnerable to overflow and underflow conditions. Alternatively, the jitter buffer size may be increased to lower the probability of underflow/overflow, but such a size increase inevitably brings about an increase in the added latency. Finally, the low resolution of the jitter buffer level leads to unnecessarily high wander generation.

By using advanced clock recovery algorithms, recovered TDM clocks can be made to comply with ITU-T G.823 and G.824 specifications as well as the new, recently approved, ITU-T G.8261 (former G.pactiming) timing distribution aspects over packet networks standard.

Timing Requirements for Cellular Backhauling

The absolute radio frequency of the air interface plays an important role in the ability of the Universal Mobile Telecommunications Services (UMTS or 3G) to work properly. This is not new, second generation GSM base stations have strict constraints with regards to their absolute RF frequency (less than 50 parts per billion of frequency error). This is required to support the GSM handoff mechanism as mobile stations wander from one cell to the other. With migration to UMTS, we see that the old frequency error requirements have been adopted while additional requirements were added for Time Division Duplexing (TDD) based Node-Bs.

Failure to meet the timing requirements of the relevant standards would cause performance degradation for the radio access channels. In particular, this failure could compromise cell handover (especially for traveling mobile stations) and producing excess of dropped calls.

Traditionally, GSM cellular networks derived the RF of their air interface from the incoming TDM link (E1/T1) delivering traffic from the base station controller. The Base Station Controller's clock is usually traceable to a Primary Reference clock (PRC or Stratum-1 clock) making it very accurate and therefore well suited for the base stations RF frequency.

As traffic volume grew, TDM link aggregation became common, first to PDH hierarchy volumes and later to SDH ones. The immediate affect was degradation of the short-term frequency stability as a result of bit justification (PDH) and pointer justification (SDH) events (although long-term accuracy was still PRC-traceable). This degradation has a detrimental effect on the air interface stability.

Another alternative (popular in North America for IS-95 and CDMA2000 cellular networks) for delivering an accurate clock to cellular base-stations is to integrate a low cost Global Positioning (GPS) device into the base station (or Node-B). This has the advantage of delivering precise PRC-quality timing and even absolute time (when needed) directly to the base station's RF clock. Nevertheless, as GPS is perceived outside the USA as an American-controlled technology with no service guarantees, its use as the primary source of timing is disallowed. Furthermore, GPS antenna installation issues may make this technique cost-prohibitive.

When the 3GPP group of standards emerged there were already well-established standards that addressed synchronization in telecommunication networks. Hence, many references to

---

these "old" standards can be found within the 3GPP. Among these "old" standards are the G.823, G.811, G.812 and G.813 of the ITU-T; EN-300-461-1 to EN-300-462-7 of ETSI and T.101 of the T1X1 standardization group. All these standards specify, in one way or another, requirements for timing and for the interfaces (that deliver the timing information and the traffic) over synchronous digital telecommunication networks (mainly PDH, SDH and SONET). The metrics used to define these requirements are practically the same in all the standards. They focus on two type of phase (time) error quantification: As described earlier, The Maximum Time Interval Error or MTIE and the Time Deviation or TDEV. One selects the metric based on the type of analysis desired. While MTIE focuses on the time domain and is used to capture and measure time transients of clocks, TDEV focuses on the frequency domain, measuring the spectral characteristics of clocks.

The synchronization requirements, as they appear in 3GPP, for the UMTS Terrestrial Radio Access Network (UTRAN) are mainly divided into two aspects. The first relates to the successful transport of data from the Radio Network Controller or RNC (BSC in GSM) to the Node-Bs (BTSs in GSM) and second to ensure sufficient frequency stability for RF channels of the Node-Bs (i.e., timing for the end application).

With regards to the latter, the classical specification for the RF accuracy of GSM base stations was .05 ppm or 5×10^{-8} and is found in ETSI TS 145 010 (TS 100 912). The equivalent UMTS specification can be found in TS 125 104 "UMTS: UTRA (BS) FDD; transmission and reception" and TS 125 105 "UMTS: UTRA (BS) TDD; Radio transmission and reception" but modifies the accuracy to $\pm 5 \times 10^{-8}$ over one timeslot. Moreover, for the case of TDD based Node-Bs, TS 125 402 "UMTS: Synchronization in UTRAN Stage 2" specifies an additional requirement for the relative phase difference between Node B's of 2.5 µs.

The above set of requirements means that in the case where the Node-B derives its RF clock from the incoming traffic link, the bitstream clock of that link must have a frequency stability of less than 50 ppb (actually, as the RF clock mechanism tends to add some frequency noise of its own, often, the actual requirement from the bitstream clock is to have frequency stability of less than 16 ppb!). In addition,, since the Node-B's internal RF PLL usually has a rather limited bandwidth. Hence, although the RF frequency accuracy is defined over time duration of 1 timeslot, the actual integration-time for measuring the input bistream frequency accuracy is usually of several seconds

As for synchronization for the transport, ETSI TS 125 402 recommends that all the UTRAN network clocks are traceable to some primary reference clock according to G.811. This clock might reside within the RNC or even further along within the MSC or even the PSTN gateway. Moreover, it recommends that the node clock implemented within the UTRAN be chosen according to the physical layer adopted for the transport network as well as the network synchronization strategy adopted. The standard gives reference to the existing ITU-T G.812, G.813 (and their ETSI counterparts).

Another requirements with regards to the UTRAN Iu physical layer interface can be found in ETSI TS 125 411 "*UTRAN Iu Interface layer 1*", that specifies the jitter and wander requirements for the traffic interfaces to meet the ITU-T G.823, G.824, G.825, whichever is applicable.

ITU-T G.8261

The newly released G.8261 recommendation mandate was to study timing distribution aspects through packet switched networks in general. Nevertheless, the first version (that has gone to consent) is mainly focusing on Ethernet (L2) networks leaving the other packet switched technologies that are within the general scope (MPLS, IP, ATM) for future version.

In addition, this first version was mainly focusing on 2 types of timing distribution over Ethernet networks:

1. Circuit Emulation Services (CES) were the service (data) and the timing information is carried together (in the same flow). Clear representatives of this type of timing distribution are the adaptive and differential (common clock) methods.

2. Synchronous Ethernet, where the bit clocks of the Ethernet switches (100/1000 MHz) are disciplined to some PRC reference by means of a synchronous master-slave timing trail (similar to an SDH/SONET timing trail).

Other timing (and time) over packet distribution methods like the IEEE 1588 will be focused on in the next versions.

The current version of G.8261 includes the following:

- Prescribes the wander (and jitter) budget limits of a CES island for several different (synchronous) network architecture scenarios, depending on the deployment place of the CES island within the entire network plan.

- Describes four timing distribution methods for Constant Bit Rate services transported over packet networks (CES). These include the network synchronous, differential, adaptive and the trivial case where the TDM end equipments are fed with the appropriate external clock (no clock recovery is needed).

- Proposes the use of Synchronous Ethernet as a method for distributing an accurate clock (PRC) from the edge of the network to the end equipment.

- Describes the effect of impairments introduced by packet switched networks on the different timing distribution methods as well as the impact of the reference clock impairment on timing distribution and service clock recovery (for example in the case of differential timing distribution).

- Gives recommendations with regard to the preferred timing distribution method (in order to meet the wander budget requirements) for the different CES deployment scenarios.

- Surveys the TDM to Packet and Packet to TDM Interworking Function (IWF) synchronization functionality and requirements (bandwidth, settling time and more). This task will be carried from now on under a new recommendation that was invoked during the meeting.

- Stipulates reference models (at this point for Ethernet networks only) as well as equipment testing guidelines.

The most important contribution of the G.8261 standard is the formalization of the CES island wander budget. According to the standard, there are 3 different wander requirements (MRTIE masks) accounting for 3 different deployment cases for the CES island (see G.8261/clause 7):

- Deployment case 1: The CES island replaces one of the four cascaded SDH islands between the slip buffer terminators (switches). The total wander budget allocated for all 3 SDH islands and the CES one is the G.823 traffic interface. Hence, only a part of that can be allocated for the CES island alone. This deployment scenario is the most stringent one (4.3usec/1000sec).

- Deployment case 2 (application A): The CES island is located outside the switch connecting it to the end equipment. Since at the switch output the wander level is that of a G.823 sync interface, the wander budget allocated for the CES is the G.823 traffic interface minus the G.823 sync interface (16usec/1000sec) that is very close the traffic interface requirements.

- Deployment case 2 (application B): In this case the application recovers timing through the TDM signal; therefore there is no differential jitter and wander between the clock and the data other than within the bandwidth of the clock recovery since the data and clock are extracted from the same signal. The wander budget of the CES segment is only limited by the timing quality requested by the application (e.g. Base Station requirements) and not by the ITU-T G.823 specification.

- Deployment case 3: In this case a clock retimer buffers between the 3 cascaded SDH islands and the CES one. Therefore, the total wander budget allocated for the CES is the same as in deployment case 2 (application A).

Another important achievement within the G.8261 is the definition of the packet networks (currently Ethernet) reference models (Appendix V) and the testing methodologies for the CES segment (measurement guidelines).

RAD's Timing over Packet Solutions

RAD has been a major player in the Circuit Emulation world for the past decade. Its portfolio includes many products taking advantage of this technology. Among these are the IPmux and Gmux products that provide a combination of Pseudowire (TDMoIP) and Ethernet switching capabilities over IP, Ethernet and MPLS networks, the ACE product line that provides a combination of ATM, Pseudowire and Ethernet over TDM, ATM, IP, Ethernet and MPLS networks as well as the Vmux, LA and Megaplex product lines that also provide a variety of TDM, Pseudowire and Ethernet capabilities.

RAD is investing significant research and development resources into achieving precise timing solutions over packet networks to support these product lines and is considered to be a leader in that field. Most of the above-mentioned products base their timing recovery on RAD's TDMoIP RJ020 and RJ021 family of ASICs, developed especially for that purpose.

RAD's first generation of timing recovery solutions was based on a combination of software and discrete hardware components. The second-generation algorithms are built into RAD's RJ020 TDMoIP ASIC, whose advanced adaptive clock recovery algorithms were designed to meet ITU-T G.823/G.824 jitter and wander traffic interface masks. Each RJ020 contains four independent TDM clock recovery mechanisms, one for each TDM interface.

The newly released RJ021 is considered to be a major step forward, as it includes innovative clock recovery algorithms allowing it to meet the more stringent jitter and wander G.823/G.824 synchronization interface masks as well as the masks of the recently released G.8261 (former G.pactiming) Recommendation.

The RJ021 extends the features of the RJ020, allowing it to better tolerate impairments introduced by packet switched network and still deliver precise timing. Among these are:

- Enhanced packet loss concealment capabilities.

- Advanced detection and treatment mechanisms for network link-breaks, outages and constant delay changes.

- Improved frequency synthesis resolution of 0.5 ppb.

- A bandwidth adaptation mechanism that optimally tunes the clock recovery PLL bandwidth according to an on-going characterization of the input noise (PDV).

- A fast initial frequency acquisition phase offering a frequency capture-range of +/- 90 ppm.

These ASICs have been incorporated into a number of products. Most popular today is the IPmux-14 that is available with several timing options:

	Ipmux-14	Ipmux-14/T	Ipmux-14/A
Frequency accuracy	Better than 1 ppm	Better than 100ppb	Better than 16ppb
Jitter&Wander output	G.823/4 Traffic	G.823/4 Traffic	G.823/4 Sync* G.8261 (section 7)

To achive synchronization interface adequate traffic engineering is required.

In order to select the timing option that makes most sense, one must examine the needs of the application and the type of network available. For example, toll bypass and CDMA cellular backhaul applications work well with the standard IPmux-14 even if the network is sub-optimal (e.g. fixed wireless or through routers without strict QOS). On the other hand, GSM and UMTS cellular backhaul applications work best with IPmux-14/A over well engineered networks since inaccurately recovered clock could result in out-of-phase timing between BTS cell sites (when GSM/UMTS BTS are out-of-phase, they will have a slower resync time that will not cause calls to drop but will result in a voice interruption of about 200ms during hand-off between towers). Note that in the case of CDMA, all towers have local GPS timing and therefore timing derived from the T1/E1 circuits need not be as precise. On the other hand, CDMA does require low differential delay (under 16ms) between towers to avoid dropped calls during soft-handoff. The IPmux-14 excels in this regard with less than 0.5ms of intrinsic delay and jitter buffers starting at 0.5ms.

Summary

RAD has developed powerful ASICs that not only improve performance by reducing latency and power consumption, but also make it possible to deliver precise timing over packet switched networks that introduce impairments such as wander, jitter and packet loss. Accurate timing and low latency are especially important when designing transport networks for cellular backhaul. For further information, including MTIE/TDEV and G.8261 test results, please contact market@rad.com .