CXR Larus White Paper –

Next Generation Network Timing, Wireless Backhaul, and the Old West

June 15, 2009

Westward Ho!

Back in the 1840's there existed a mass migration of eager pioneers forever trekking westward in search of cheap land, jobs, and of course, the lure of gold. It was a migration primarily driven by local and regional economics, and things are really not a whole lot different it seems some 170 years later, at least in the wireless backhaul business. Telecom service providers and carriers are migrating "westward" from traditional Time Division Multiplex (TDM) (T1/E1) based networks to the lure of Internet Protocol (IP) transport for wireless backhaul applications. The economics are clear; at the same price as a single T1, one high speed ethernet connection at the basestation can replace all TDM circuits, provide voice offload and the required bandwidth for new 3G/4G applications. However, they also demand carrier-class TDM transport, support of legacy voice and other applications, timing and sync for handoffs, and Key Performance Indicator (KPI) monitoring for Service Level Agreement (SLA) enforcement. These requirements cannot be satisfied using legacy timing and sync methodologies, as packet based networks break traditional timing distribution models. Next generation timing is required from the Core to the Edge to support important functions, such as higher spectral efficiency to provide higher bandwidth and more subscribers per cell, and timing precise enough to prevent poor hand-offs between cell sites resulting in dropped calls. Also, carriers deploying Femtocells are finding that a method of delivering precise distributed network timing is required, as GPS is not viable in indoor environments, nor very cost effective. In addition, it is also nearly impossible to track one-way latency and jitter measurements required for KPI reporting without ubiquitous and accurate time. What are the alternatives?

Delivering Practical and Cost Effective Timing Solutions for Wireless Backhaul

There are several technologies that can be considered to provide the frequency, phase and Time of Day (TOD) required for Next Generation networks. Some of the choices that backhaul carrier engineers consider include Network Time Protocol, (NTP), Precision Time Protocol, (PTP), also known as IEEE-1588v2, Pseudowire Emulation (PWE) using Differential Mode or Adaptive Clock Recovery, and Synchronous Ethernet (Sync-E). The question of the moment is; which of these alternative technologies are ideal for wireless backhaul applications? Let's take a closer look at each of the options;

Pseudowire Emulation is designed to encapsulate and transport TDM (e.g. T1/E1) over packet-based (IP) networks, usually with a proprietary (vendor specific) protocol and coding scheme. There are two viable methods for getting timing data through PWE circuits, Differential Mode (DM), and Adaptive Clock Recovery (ACR). While one can provide frequency with some level of accuracy, only one practical solution provides accurate frequency, phase and Time of Day.

Most PWE services generally use *Adaptive Clock Recovery* as a method to recover frequency when no external reference clocks are available. ACR primarily uses two components, packet arrival rate and fill level of the jitter buffer, as the mechanism to recovery and distribute frequency synchronization. However, ACR exhibits major limitations, primarily performance degradation under heavy network utilization, and it also may not meet carrier requirements for phase and time, per ITU Standards G.823/G.824's Mean Time Interval Error (MTIE) and Time Deviation (TDEV) masks. ACR also does not provide proper phase alignment, or Time of Day information. In addition, ACR coding is proprietary (vendor specific), point-to-point only, non-interoperable, and subject to timing loops. Most carriers shy away from proprietary solutions whenever possible, making it less than ideal for backhaul applications.

With *Differential Mode*, timing is not carried through the PWE pipe at all, instead there is a common stratum level reference clock source installed at each end of the path. This approach enables timing to essentially be "re-inserted" at the receiving end, and both ends are now synchronized to a common reference. One way to accomplish this is to put expensive GPS Stratum 1 clocks at both ends, but this approach is not cost effective, nor completely desirable, as GPS antennas at remote sites are subject to lightning damage and vandalism. Another way that *is* cost effective is to employ a Stratum 1 PTP Grandmaster at the MSC, and have it port timing information to isolated PTP Slave clocks at cell sites, which then in turn output traditional TDM T1 timing to the installed equipment. In this case, a forklift upgrade is not required, and frequency, time and phase are all accounted for.

Network Time Protocol, (NTP), is one of the oldest ethernet protocols still in use, and was originally designed by Dave Mills of the University of Delaware some 25 years ago. NTP is available in two levels, the "standard" version, and Simple Network Time Protocol, (SNTP), a de-featured subset of NTP. The latest version of NTP, Version 4 (NTPv4) can usually maintain time to within 10-20 milliseconds using traditional software interrupt based solutions over the public Internet, and can achieve accuracies of microseconds or better in local area networks under ideal conditions and the latest generation of timing solutions. That said, accuracy can degrade substantially if networks become congested, for example if links go down, or if IP routers send packets through more hops that originally engineered. Even the best engineered network will still be subject to variability, latency errors, with lost or misdirected packets due to how NTP is implemented.

NTP is indeed quite simple in operation; the NTP client polls the NTP server at regular intervals with "What time is it?" and the server responds with a time stamp. The problem with using NTP for precise timing applications is that there is no allowance to account for network delays other than through multiple poll time averaging techniques and buffering. NTP, even in the latest implementation does not meet the higher precision requirements for telecom network synchronization, especially those required for 3G/4G wireless backhaul.

At the turn of the 20^{th} Century the most viable mode of transportation for the masses was either driven by horse or steam. Then along came this little innovation in technology called the "automobile" and its state-of-the art petrol based engine, which revolutionized transportation as we knew it. **Precision Time Protocol, (PTP)**, popularly known as IEEE standard 1588, is the latest revolution in IP timing technology. Originally designed to provide precise timing for critical industrial automation applications it is now providing the highest level of accurate frequency and time to wireless backhaul networks. Currently standardized in 2008 as Version 2 (IEEE-1588v2), PTP overcomes the ethernet NTP latency and jitter issues, providing an unprecedented accuracy in the *nanosecond* range. Network delays and latency are greatly reduced by measuring the roundtrip delay between the master and slave clock (client), using a technique where the slave and master communicate with short messages to each other in order to measure and cancel out delay and latency inaccuracies. Previously, expensive GPS based clocks at each cell site were required to obtain the order of magnitude required for 3G and new 4G services when using IP backhaul. Now, a central Grandmaster Clock at the Mobile Switching Center (MSC) and low cost PTP Slave clocks at the cell sites are all that are required, greatly lowering both capital and operating costs for the carrier.

Synchronous Ethernet, commonly abbreviated Sync-E is yet another technology for delivering precise frequency, (but not time) through an IP network. Sync-E is standardized through ITU Standards G.8261, G.8262 and G.8263, but is slow to be fully adopted, and for a very good reason. In Sync-E, the timing information is not carried within the packet structure, as in NTP or PTP, but is carried in Layer 1, the Physical Layer. Typically in ethernet, the transmit clock is derived from a simple inexpensive +/- 100ppm crystal, and the receiver is able to lock onto it, no problem. The good news is that there is certainly nothing keeping one from making the transmit clock Primary Reference Source (PRS) traceable with a Stratum 1 GPS, Rubidium or Cesium standard. The bad news comes from the fact that now the "PRS traceable" clock has to be passed through the various routers and gateways from the MSC or central office to the cell towers while not being altered or degraded with additional wander or jitter. This means that all the ethernet network equipment in the middle has to be fully compatible with the Sync-E standards for passing timing, and this often requires a forklift upgrade to accomplish. This is generally not a problem with newly engineered infrastructure build-outs from scratch, but for backhaul carriers trying to utilize existing equipment already deployed and paid for it becomes more of a challenge, and of course greatly increases capital expenditures. However, there is an advantage to deployment of Sync-E also; since timing is no longer carried within the ethernet packet structure, it is not affected by timing impairments caused by packet loss or packet delay variation (latency). Also, the data is not instilled with bandwidth robbing overhead or continuous polling requests and replies. Still the requirement for complete network compatibility with Sync-E is an obstacle, very expensive for the typical carrier and service provider to overcome, as is the lack of any form of time stamping capability. Remember that Sync-E can be utilized only for *frequency* synchronization, not *time* synchronization.

To meet requirements for real-time applications, wireless service providers must deliver highly accurate **Frequency**, **Phase**, and **Time of Day** on an end-to-end network basis (Reference Appendix 1):

Frequency is required at the physical layer and reduces the effects of call overlap and T1 slips. GSM, 3G and 4G are especially susceptible to frequency offset. Additionally Femtocells, which are indoor basestations providing last mile backhaul on the Internet Service Provider's (ISP's), rather than the mobile operator's network, require a network based frequency synchronization methodology.

Phase synchronization is required by CDMA systems and the most stringent of these are WCDMA, WiMAX and new technology, Long Term Evaluation (LTE). Frequency requirements vary between 1-3 µs. Absolute phase and frequency synchronization is also required for spectral efficiency. Inefficient frequency can reduce the spectrum utilization to 85% or lower. After all, service providers pay billions of dollars for spectrum, so it is in their best interest to maximize its efficiency.

***Time of Day* (TOD)** is required by carriers to adhere to operational needs such as billing, root-cause analysis and alarms, as well as one-way measurement of latency and jitter.

Conclusion

OK, so what IS the ideal timing solution for wireless backhaul? Depending upon their personal goals, and of course their resources, the settlers that came in droves to the lures of the Old West certainly had many viable alternatives as to where to plant their stake in the ground. So it is the same with the backhaul carrier. As we have learned, Sync-E does not provide time synchronization and requires a major investment in new infrastructure, so it may not be viable for all budgets. PWE has at its core proprietary vendor specific coding, and other issues such as poor phase and frequency stability that make it unattractive. Consensus has formed among Tier 1 mobile operators and equipment manufacturers including Vodafone, BT, FT Orange, T-Mobile, ATT, Cisco, Ericsson, Alcatel-Lucent and others that frequency delivery as specified in ITU-G.1862 for Sync-E is the preferred solution for frequency and IEEE 1588v2 Precision Time Protocol (PTP) is the preferred solution for delivering both phase and time, with Network Time Protocol (NTP) delivering accurate time at Layer 3. So, the answer is really a *combination* of technologies. The best *single* and most *cost effective* solution for the backhaul provider centers around Precision Time Protocol, IEEE-1588v2. However, even that option requires a carefully engineered network to fully gain the performance advantages offered, such as controlling the number of router hops between the MSC and cell site, and ensuring bandwidth utilization constraints are met.

About CXR Larus

CXR Larus has been in the timing business since the beginning, and offers a comprehensive suite of standards-compliant grandmaster and slave clocks, network time servers and clients that provide Edge-to-Core network timing and synchronization. With support for NTP, PTP/IEEE 1588v2, Synchronous Ethernet (future), 802.1ag, 802.3ah and legacy synchronous interfaces such as T1, E1, CC, 1PPS, 1/5/10MHz, etc., we offer the wireless carrier a seamless migration path from the past into the future.

CXR Larus products are fully standards compliant, hardware based timestamp solutions designed around a uniform architecture across all Next Generation timing platforms. CXR Larus is currently providing solutions to multiple tier-1 wireless service and telecommunications equipment providers worldwide. Our Next Generation timing portfolio includes:

- The **StarSync 6840 Stratum 1 GPS Integrated Zero Footprint NTP server** is designed for access and metro edge networks.
- The **StarSync 6850 Stratum 1 PRS/PRC and NTP/PTP Server/Grandmaster** for core and aggregation networks.
- The **StarSync 6800 Edge NTP Client/PTP Slave** recovers GPS timing and sync from the network.
- The **StarClock TiemPo 6400** is a Convergence Building Integrated Timing Supply (BITS), providing advanced and configurable timing solutions for both Next Generation and legacy TDM timing applications.

For further information on the full range of CXR Larus timing products please visit our website, www.cxrlarus.com, or call one of our timing experts at 1-408-573-2700. We can also be reached by e-mail at sales@cxrlarus.com.

Appendix 1- Wireless System Timing Requirements

Wireless System Technology	Basestation Timing Requirements		FemtoCell Timing Requirements	
	Frequency Synchronization	Phase (Time) Synchronization	Frequency Synchronization	Phase (Time) Synchronization
GSM	<50 ppb	Not required	<100 ppb	Not required
CDMA	<50 ppb	± 3 µs	<100 ppb	± 1 µs ± 10 µs during holdover over 8 hr
3GPP2 CDMA2000	<50 ppb	± 3 µs ± 10 µs over 8hr	<100 ppb	± 3 µs ± 10 µs worst case over 8 hr
WCDMA/UMTS - FDD	<50 ppb - Wide Area <100 ppb - Local Area	Not required	<250 ppb	Not required
WCDMA/UMTS - TDD	<50 ppb	± 2.5 µs	<250 ppb	± 2.5 µs between basestations
WiMAX /FDD	<50 ppb	Not required	<8 ppm	Not required
WiMAX /TDD	<50 ppb	± 1 µs	<8 ppm	± 1 µs with GPS ± 5 µs without GPS ± 25 µs during holdover
TD-SCDMA	<50 ppb	± 3 µs	<100 ppb	± 2.5 µs between basestations
LTE	<50 ppb	Under development Some use ± 3 µs / ± 10 µs	Under development	Under development