



Knowledge Network

<http://kn.theiet.org/magazine/issues/0909/doing-time-0909.cfm>

Doing time

Published on 19 May 2009

By Bryan Betts



Mobile operators are having to find new ways of synchronising their base station networks as they shift to packet-based backhaul connections. *E&T* explores their options.

Mobile phone basestations need a highly accurate and shared time signal. Without it, error rates rise as basestation clocks get out of sync, and adjacent cells are unable to synchronise their transmission frequencies. As a result, calls get dropped and users suffer.

In the world of GSM, this has not been a major issue. Most basestations have a leased line for backhaul, typically a copper cable carrying a T1 (1.5Mbit/s) or E1 (2Mbit/s) circuit, or fibre carrying a SONET/SDH link. The basestation recovers a primary reference clock signal from a reliable end-to-end synchronisation chain within the backhaul network, and uses that to calibrate its own embedded quartz oscillator.

Problems have arisen with the move from basic mobile phones to video and data-capable 3G devices, with higher data rates and bandwidth needs. Renting a T1/E1 backhaul already represents a large part of the expense of operating a basestation – between 30 and 50 per cent in some cases – so the cost of switching to an even more expensive T3 (45Mbit/s) or E3 (34Mbit/s) service is prohibitive.

“From 2006 to 2010, the number of mobile phone users is expected to grow by only 30 per cent, while backhaul expenses will skyrocket due to an exponential increase in bandwidth required for video and multimedia applications,” says Eitan Schwartz, vice president for pseudowire and Ethernet access at Israeli network equipment developer RAD Data Communications.

“Moreover, the average revenue per user is not likely to grow much, even for new data services, with competition from fixed-line offerings keeping a lid on pricing,” he adds. “On the other hand, a quick search on the Internet reveals that complaints are mounting about 3G data speeds not meeting the promised performance and being inconsistent, with inadequate coverage. If mobile operators are to reach profitability targets while providing untarnished performance, they must improve the efficiency of their networks by dropping the cost per Mbit/s of bandwidth.”

Cost drivers search for new sync sources

The cost of these backhaul connections is pushing operators into using cheaper packet-based networks. Unfortunately this breaks the end-to-end clock synchronisation chain that enabled GSM networks to keep in sync, says Emir Halilovic, programme manager for carrier network equipment in central and eastern Europe at research company IDC.

“Lately we’ve seen the importance of the issue increasing because people are thinking of using IP for backhaul,” he says. “Before, we had leased lines for backhaul, but now IP and Ethernet are coming up as the only solutions people are going to use. Since Ethernet does not have timing in the normal IP protocol, they have to think about synchronisation and that opens the way for all sorts of solutions.”

The problem is that packet-based networks such as Ethernet are not deterministic. Packets arrive when they arrive – they may be buffered along the way, consecutive packets might take different routes across the network, they may even be lost and have to be retransmitted. It is up to the receiving station to reassemble them and make sense of it all. That is fine for an email message or Web download, and it might even keep your PC’s clock accurate to within milliseconds, but it is not going to deliver the sub-microsecond levels of synchronous timing accuracy needed by cellular basestations.

Time and frequency alignment

Basestations need more than just accurate time – they also need accurate frequencies, says Tomer Carmeli, a product manager at Tel Aviv-based mobile backhaul specialist Ceragon Networks.

“There are two technologies for the air interface of a cell site, FDD [frequency division duplex] and TDD [time division duplex],” he says. GSM (2G and 2.5G) and UMTS (3G) networks are primarily FDD – although some UMTS implementations also use TDD – while CDMA cellphone networks and WiMax data networks use TDD.

“FDD uses different frequency subsets for each call or session, so you need very accurate frequencies, because if a subscriber roams from cell to cell, in order to switch the call the two cells need to use the same frequencies. So two adjacent cells need to have the same clock source, which can be done by distributing the same frequency to cells.

“GSM uses FDD so just needs frequency alignment – typically that means supplying a traceable clock through the transmission network. TDD uses timeslots as well as frequencies. It not only needs the same frequency but also the same clock tick – phase alignment – which is a harder requirement.”

On the plus side, TDD is more efficient than FDD in its use of spectrum and is better suited to asymmetric traffic such as data. That’s because two TDD end-points share the same frequency, transmitting in different timeslots, whereas FDD needs two different frequency bands, one for upstream transmission and another for downstream.

However, the additional cost that TDD brings is producing industry pressure towards FDD, some in the industry suggesting that WiMax should merge with the FDD-based 4G mobile phone specification LTE (Long Term Evolution).

GPS timing for basestations

In many cases the need for network synchronisation is not yet pressing. There are a lot of cell sites that host both 2G and 3G basestations, and if there is still an E1 line in place for the 2G site, operators can take the clock from that and use it to drive Synchronous Ethernet to the 3G equipment.

But, if there is no longer a circuit-based connection from which to source a timing signal because of a switch to an all-IP or circuit-emulated service, and if it's impossible to send a reliable timing signal in packet form, how are these networks going to synchronise your basestations?

"The first option is a GPS receiver on each site. This is a viable option as it provides both frequency synchronisation and highly accurate timing – it is the *de facto* method for WiMax and CDMA," says Carmeli.

GPS has its drawbacks. GPS signals can be jammed or spoofed, which could open GPS-equipped basestations to denial-of-service attacks, and there is the political dimension – the GPS network is operated and controlled by the US military.

More important is its cost – the GPS receivers used for network synchronisation have a much higher specification than those in the average portable satellite navigation system, plus they need all the right interfaces and cabling to talk to telecoms equipment. These receivers also need a good view of the sky, which can be a challenge if the basestation is in the basement of a tower block, or a network operator is introducing femtocells to provide 3G coverage in the home, with a backhaul connection over home broadband.

Femtocell access points still need synchronisation, with the relevant specs requiring a local oscillator frequency accuracy of +/- 0.1 parts per million over one time slot, says Sanjay Bhatia, director of product line management at Genband, which develops gateway devices for fixed-mobile convergence.

He adds that, where there is 3G coverage from macrocells, devices such as femtocells can make use of that, sniffing frequency offsets from nearby basestations and using them to recalibrate their own oscillators. If not, they must find other methods.

Synchronous Ethernet

Service providers are requesting alternatives to GPS, he says, most notably timing-over-packet schemes such as Synchronous Ethernet, which is defined in the ITU standards G.8261, G.8262 and G.8264, and the IEEE 1588 version 2 standard, which defines the precision time protocol, PTP.

Synchronous Ethernet uses the timing information inherent in the underlying medium, now that Ethernet is a switched transport rather than a collision-based one.

"Synchronous Ethernet is saying we have been using a traceable network, now it's IP but it is still a bitstream," says Ceragon's Carmeli. "For example, Gigabit Ethernet is 125 million symbols per second at the physical layer. So it took the same concept and generalised it to other links."

He adds that microwave-based wireless systems such as Ceragon's are synchronous, so you can recover a clock signal from them much as you can from a T1/E1 link. However, Synchronous Ethernet does not supply the phase synchronisation or alignment needed by TDD networks.

Where the connection is not synchronous or where phase alignment is needed, IEEE 1588v2 can be used. This involves a master clock broadcasting sync messages and its current time to its client basestations; these respond and the master then tells each client when it received its response. Each client then uses the message and response times to calculate its offset from the master using its own clock recovery algorithm – a form of phase-locked loop (PLL) that is not specified in the standard.

There are a few issues here. The scheme assumes that the exchange of messages is quick enough for the offset to be constant, and that the network latency is the same in each direction. Carmeli adds that it can also be affected by factors such as network loading.

Jitter and wander

“In a perfect network we would have a fixed latency so we can estimate when packets arrive,” he says. “Timing-over-packet works via educated guesses then a series of fixes using a PLL. The problem is that networks are not perfect, and there is delay and variation. The algorithm of the PLL needs to be very sophisticated to filter this noise, which generates jitter and wander.

“You do need to be much more aware of traffic patterns and network loading for timing-over-packet, because as you add more traffic, you can potentially add more noise, so it is more tricky to engineer. A lot of operators have a timing-over-packet strategy, which is okay for fibre but can have issues with microwave [links].

“It is nice because you can have master and client clocks over any packet network, in theory,” he continues. “For good performance it needs control over network performance, and it needs a well-engineered network. It is definitely not going to work over the Internet, for example. But you can use it to emulate phase alignment.”

He adds: “The trend in the industry now when you need phase alignment is to use Synchronous Ethernet with IEEE 1588v2 running on top for the phase alignment. A good-quality, cheap 1588 client is not so easy to do because it needs a good local oscillator, which is expensive. So it is better to use them together, because then 1588 only needs to emulate the phase, it doesn’t need to guess the frequency.”

Despite these issues, IEEE 1588v2 – or PTPv2 – is one of the most popular schemes around. It is now supported by most major equipment vendors, and many multi-vendor interoperability tests have been successfully completed.

It will be even more important as Carrier Ethernet – a telco-focused evolution of Ethernet, designed to run on wide-area networks – is introduced, says Dirk Lindemeier, who is responsible for mobile backhaul solutions development at Nokia Siemens Networks.

“Carrier Ethernet economically meets exploding bandwidth requirements currently constrained by the prohibitive costs of legacy networks,” he told the Carrier Ethernet World Congress, late last year. He added that, with synchronisation based on IEEE 1588v2, “We have demonstrated a complete solution for packet-based synchronisation working live over a multi-vendor Carrier Ethernet network.”

Most of the network synchronisation specifications and concepts are not new – IEEE 1588 has existed for several years, although the focus is now on version 2, which was approved in March 2008. What’s new is that the shift to packet-based networks is now under way, so the need for network synchronisation is much more pressing.

Mobile phone networks may be the highest profile application for network synchronisation. Accurate timing is also used in production-line networks, and in home multimedia systems. However, it is mobile networks that need tight synchronisation, and it is their growth that has driven the shift to packet-based backhaul and the consequent need to develop new ways of keeping everything on time and in line.