SYNCHRONIZATION IN INTELLIGENT DIGITAL NETWORKS

Gary J. Grimes, Mark S. Russo, Christopher Lanzafame, Christopher D. Near, and Bryan S. Moffitt

Gary J. Grimes is a distinguished member of technical staff in the Digital Terminal Development Department with AT&T Bell Laboratories in Denver, Colorado. Christopher Lanzafame is a member of technical staff in the Customer Systems Design Department with AT&T Bell Laboratories in Middletown, New Jersey. Bryan S. Moffitt is a consultant in the Systimax™ PDS (premises distribution system) Department with AT&T Bell Laboratories in Middletown, New Jersey. Christopher D. Near and Mark S. Russo are members of technical staff in the Network Planning Systems Engineering Department with AT&T Bell Laboratories in Holmdel, New Jersey. Mr. Grimes works on optical interconnection and synchronization issues for transmission terminals. He joined the company in 1978 and has a B.S. in physics from The
(continued on page 68)

Customer expectations regarding the performance of intelligent digital networks can be met only if careful consideration is given to synchronization at every level in the network. This includes synchronization of the public switched and private networks, and performance of the clocks that control customer premises equipment (CPE). Poor synchronization planning at the public or private network level, as well as poor CPE clock performance can degrade customer applications such as video, facsimile, digital data, encryption, voice-band data, and voice transmission. In this paper, we consider the synchronization philosophies and implementation of both AT&T's national telecommunications network and private customer-owned networks. We will also discuss the performance of typical CPE clocking systems and describe the synchronization systems used in AT&T's line of Definity® business communications systems.

Customer Expectations

Intelligent digital networks—which can consist of the public switched network's facilities and services, as well as customer-owned equipment—must provide nearly error-free performance under a variety of operating conditions. This includes automatic reconfiguration of the network (which may occur under software control, without operator intervention) during normal operations and failure modes.

The robustness of a network can be improved against observable degradation of services during failure modes, if a quality synchronization plan is in place. To attain such a plan, one must consider synchronization at each level in the network including:

- The synchronization performance of the public switched network
- The private network's synchronization architecture
- The performance of individual, CPE clocking systems.

AT&T has strived to improve the synchronization performance

Panel 1. Abbreviations, Acronyms, and Terms

ANSI — American National Standards Institute
BCM — bit-compression multiplexer
BITS — building integrated-timing supply
BSRF — basic synchronization reference frequency, formerly called the Bell System reference frequency
controlled slip — loss or repetition of an entire frame of DS1 data along frame boundaries
CCITT — International Telegraph and Telephone Consultative Committee
CCR — customer-controlled reconfiguration
CDU — clock distribution unit
CPE — customer premises equipment
DACS II — digital access and cross-connect system, model 2
DCS — digital cross-connect system
dropout — a brief loss of signal
DS1 — digital signal level of 1.544 Mb/s (24 channels)
EIA — Electronics Industries Association
GPS — global positioning system
MTIE — maximum time-interval error
PDS — premises distribution system
phase hit — a sudden change in the phase of a clocking signal
PRC — primary reference clock
PRINCE-S — private-network computer-based expert system for synchronization
PTTI — Precise Time and Time Interval
slip — loss or repetition of an entire frame of DS1 data
TIA — Telecommunications Industry Association
tip and ring — standard, analog, telephone-network loop interface
TMS — timing monitoring system
TR — AT&T Technical Reference
USITA — United States Independent Telephone Association; now called USTA, the United States Telephone Association

of both the public switched network and CPE. Thus, customers can expect and get the best possible error-free performance from integrated voice-and-data, intelligent digital networks. (Panel 1 defines acronyms and terms.)

Need for Synchronization in Digital Networks

Synchronization is the means used to keep all the equipment in a digital telecommunications network operating at the same average rate.[1] If the receiving end is to receive and properly interpret the digital signals transmitted over a communications link, it must stay in step, or be synchronized, with the transmitting end. When digital signals are transmitted over a network of digital communications links, switching nodes, and transmission interfaces, all entities must be synchronized to avoid data errors. This is known as *network synchronization*.

A major purpose of network synchronization is to ensure proper operation of DS1 equipment. DS1 data is transmitted in frames that consists of 192 data bits and 1 framing bit. We refer to the loss or repetition of an entire frame of DS1 data as a *slip*. (The loss or repetition of an entire frame of DS1 data along frame boundaries is more accurately called a *controlled slip*. For purposes of this paper, slips and controlled slips are synonymous.) Slips are caused by differences in clock rates within a network and by the effect of transmission-induced impairments on clocks.

The primary methods used to control the slip rate and limit it to an acceptable level are:

- Run all clocks in the network at the same average rate.
- Place buffers to absorb small phase variations.

Table I identifies the possible effect of a single slip on customer applications.

In the United States, equipment clocks are categorized into four strata according to accuracy and performance level.[2] Stratum 1 clocks have the highest accuracy requirements; stratum 2, 3, and 4 clocks have progressively lower accuracy and performance requirements. Timing references are passed hierarchically from the high-performance clocking systems to equivalent or lower performance clocking systems within a network to synchronize the many network-equipment clocks to one, nominal reference frequency. Most of the clocking systems used in the AT&T switched network are stratum 1 and stratum 2. Most clocks found in CPE, such as PBXs and multiplexers, are stratum 4.

However, even when all the equipment clocks are synchronized to a single master clock, it still is impossible to eliminate slips in a digital network. Therefore, synchronization plans are developed to minimize and control the rate at which slips occur. A synchronization plan specifies the timing flow through the network and takes into account equipment performance, facility performance, and network topology.

The Evolving Synchronization Architecture

The synchronization architecture of AT&T's national telecommunications network has been going through a transition. Before 1990, analog signals derived from a cesium clock ensemble were used to synchronize the AT&T switched network.[1] This clocking system, known as the *basic synchronization reference frequency* (BSRF), was located in Hillsboro, Missouri, and had a long-term accuracy of a few parts in 10^{12}. The BSRF output consisted of a 2.048-MHz (megahertz) reference frequency that was distributed on analog radio and coaxial systems to key locations within the AT&T network.

When the BSRF was initially installed, the AT&T analog network used frequency-division multiplexing, and even the most critical equipment could tolerate a 10^{-9} frequency offset. The BSRF distribution system could easily provide this frequency accuracy even in the

Table I. Performance Impact of One Slip

Application	Potential impact
Encrypted text or voice	Encryption key is retransmitted.
Video	Freeze frame occurs for several seconds, with a loud pop on the audio.
Digital data	Data is deleted or repeated; possible misframe.
Facsimile	Delete 4 to 8 scan lines. Drop call.
Voice-band data	Transmission errors occur for 0.01 to 2 sec. Drop call.
Voice	Click may occur.

presence of transmission impairments such as dropouts, protection switching, and phase hits.

But with AT&T's move to an all-digital network, the performance objective for network timing of 1×10^{-11} became 100 times greater than the frequency accuracy required for analog transmission.[2,3]

The BSRF initially provided timing to digital offices by using the existing analog distribution system. The equipment clocks in the digital offices were designed to accept this signal and meet the 1×10^{-11} recommendation.[4-6] With this architecture, typical performance was within several parts in 10^{12}; but under certain conditions, it was possible for a node to degrade to parts in 10^{10} for a short period.[7]

Today, AT&T has nearly finished deploying a new network of distributed clocking systems that will be used to synchronize its digital network. These systems, called *primary reference clocks* (PRCs), are being installed within the continental United States, Puerto Rico, and Hawaii. The PRCs use the global positioning system (GPS) to provide long-term timing accuracy to the digital network.[8] (The GPS is a system of satellites used primarily for navigation and positioning and is owned and operated by the U.S. Government.) The primary reference clocks track multiple GPS satellites and periodically receive accurate timing data. Each day, a PRC time and frequency error is estimated and used to steer the duplicated

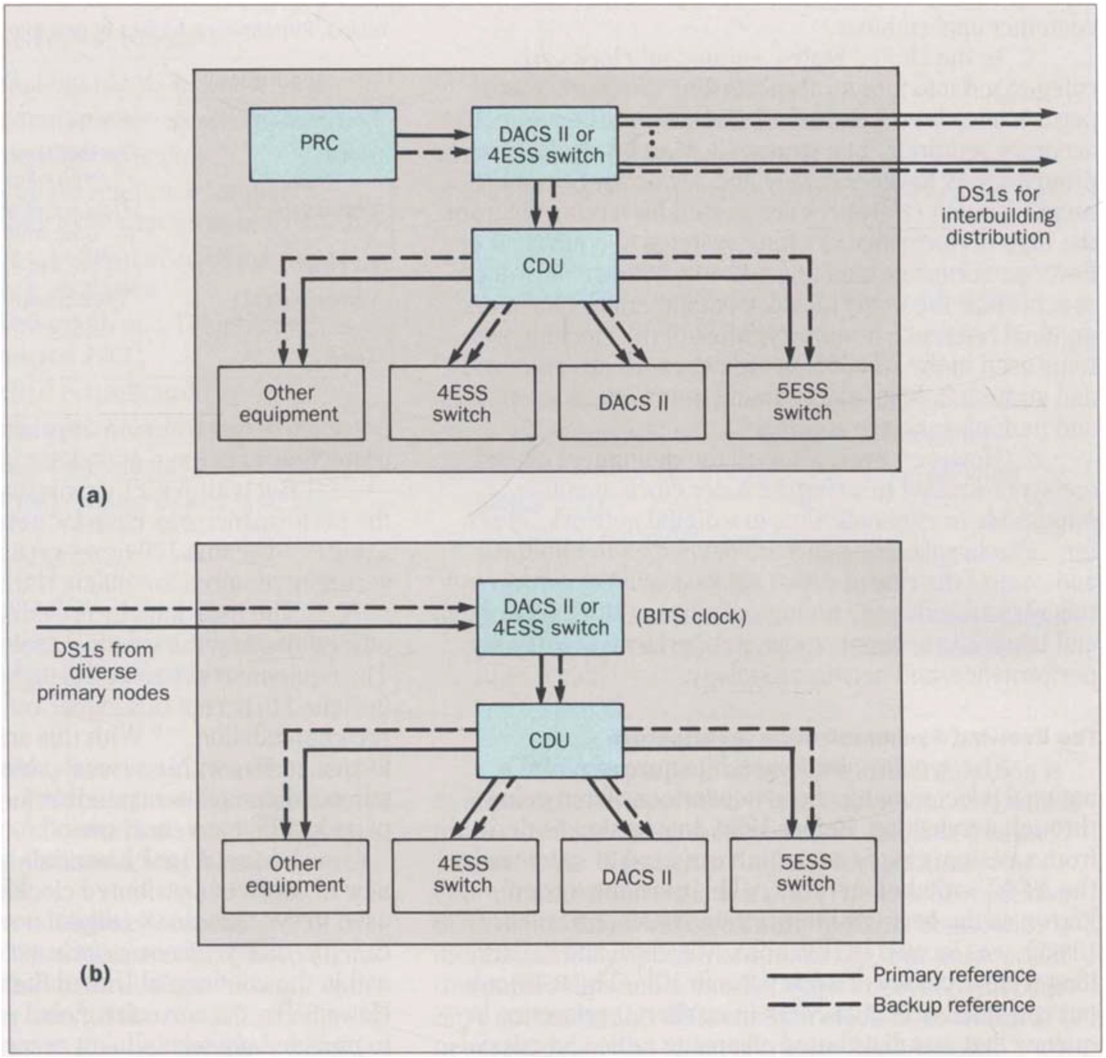

Figure 1. Representative node architectures. (a) Within a primary node, PRC outputs provide timing to the stratum 2 clock of a DACS II transmission system or a 4ESS™ switch. Timing for other clocks at the node is then provided via the CDU, which uses the primary timing reference to generate multiple references for the other clocks. (b) At secondary nodes, the best performing clock in the building (i.e., the BITS clock) receives timing references from two remote primary nodes and then disseminates intrabuilding timing, via the CDU, to all the clocks that need such a reference.

rubidium oscillators that provide the PRC system's timing output. As a result, this system achieves a long-term frequency stability of about 1×10^{-13}.

In addition to improved accuracy compared to the BSRF-based synchronization plan, the new synchronization network will offer verification capabilities. These capabilities will make it possible to monitor the timing performance of each office, using AT&T's timing monitoring system (TMS), which is collocated with the PRCs. With this system, AT&T can detect and correct timing degradations before they begin to affect service.

PRC timing references are transmitted from each primary node (i.e., locations that contain a PRC) to other nodes, called secondary nodes, that do not contain a PRC. The TMS simultaneously monitors the timing references it receives from at least two other PRC nodes, as well as from the secondary nodes. The monitoring of neighboring PRC locations provides long-term stability data. By using the TMS, AT&T will track and verify the long-term timing stability of each PRC location. The architecture is designed so that secondary nodes are monitored independently by different TMSs.[7]

Within a primary node, PRC outputs are used to provide timing to the stratum 2 clock of either a DACS II transmission system or a 4ESS™ switching system. This clock, in turn, provides timing to all the other clocks in that node, using a device known as a clock distribution unit (CDU). The CDU's function is to take a designated timing reference and generate multiple references that will be used by the other clocks in that node. Figure 1a illustrates this architecture.

In secondary nodes (i.e., those that do not contain a PRC), one clock will be selected as the building integrated-timing supply or BITS clock. The BITS concept is described in detail in AT&T's Digital Synchronization Network Plan.[9] According to this plan, the best performing clock in the building (i.e., the most accurate clock within the highest stratum at the site) is to receive two interbuilding timing references and then disseminate intrabuilding timing to all other clocks that need such a reference. Hence, the secondary node's BITS clock receives its timing references from two diverse, remote primary nodes. This timing signal is then distributed to other clocks in the building, using the CDU (as we described for the primary nodes). Figure 1b illustrates the secondary-node architecture.

While the BSRF synchronization architecture, on average, met the 1×10^{-11} performance requirement, the performance of the PRC network offers significant improvement over this requirement. Performance gains occur in two areas:

- The PRC's inherent performance advantage will not only provide timing accuracies of 1×10^{-13}, but will also make the verification capabilities possible.
- The ability to verify network timing will allow AT&T to detect synchronization problems before they affect service, and effectively provide an even higher level of performance for the AT&T switched network.

Private-Network Synchronization Planning

Private, customer-owned networks present additional complexities for synchronization planning. Typically, these networks have diverse architectures; i.e., the equipment and facilities they contain come from multiple vendors and carriers and are interconnected in complex arrangements.

In addition, the requirements for CPE clock have been minimum specifications. Most CPE currently use stratum 4 clocks, which were not designed to transfer timing references between equipment. Consequently, the private-network arena is the most difficult one in which to obtain high-quality synchronization performance and is the one most susceptible to overall performance problems.

Synchronization Architecture. Ideally, the timing flow in digital private networks follows the hierarchical source-to-receiver method. That is, synchronization references are passed from the high-performance clocks to the low-performance clocks. While most private-network equipment are usually equipped with stratum 4 clocking systems, some CPE have stratum 3 systems.

Most public-switched-network services contain stratum 1, traceable timing (i.e., timing that originates from a stratum 1 source). Hence, if the master-timing source in a private network originates from a stratum 1, traceable, public-switched-network link, then such timing should be distributed to all CPE throughout the network. Ideally, intrabuilding timing will be distributed via a BITS-type arrangement, as described for the AT&T public network. But this is unlikely to happen because most CPE do

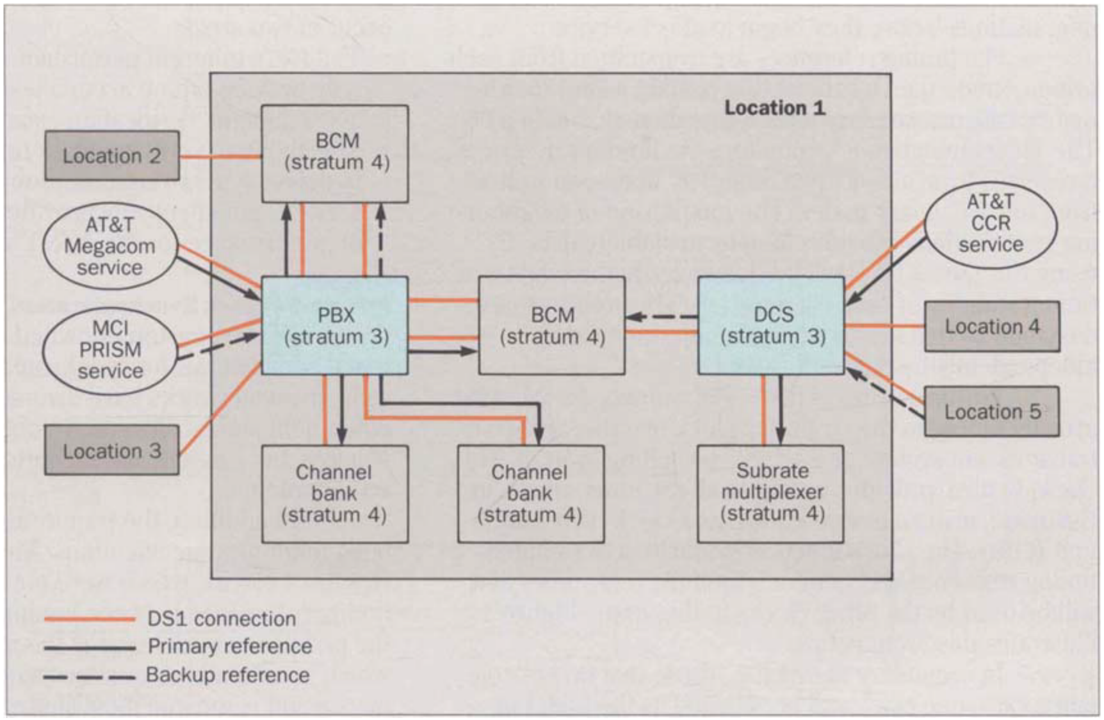

Figure 2. Synchronization plan for all equipment at a representative location in a private network. This location is connected to four other locations in the network and to three public-switched-network services (AT&T's Megacom® and CCR services, and MCI's PRISM℠ service). The timing for each network service can be traced to a stratum 1 clock.

not have dedicated timing interfaces and can only derive timing from a traffic-carrying signal.

Figure 2 illustrates a typical private-network location. This location is connected to four other locations in the private network, as well as to three public-switched-network services: AT&T's Megacom® service and CCR (customer-controlled reconfiguration) service, and MCI's PRISM℠ service. (PRISM is a service mark of MCI Corporation.) Each public-switched-network service contains timing that can be traced to a stratum 1 clock. All equipment at the location requires synchronization, and the diagram shows a reasonable synchronization plan for the site. This plan adheres to the hierarchical timing-flow requirements and permits diverse timing sources into the location.

However, the location does not contain a true BITS architecture. Instead, the location illustrated in Figure 2 has two "BITS" clocks: the PBX, and the digital cross-connect system (DCS). A true BITS architecture offers several advantages:

- Maximum use of the highest performance clock at the location (without intralocation cascading of timing references)
- The BITS clock serves as a central point for maintaining synchronization within the location.

Because of implementation problems, we are unlikely to see private-network BITS designs. Customer premises lack bridging repeaters for routing timing signals within the site. In addition, some stratum 3 and most stratum 4 CPE systems lack dedicated timing interfaces.

A major focus for any synchronization plan is to overcome topological complexity and constraints. Private networks can suffer from excessive cascading of CPE, lack of diversity of timing references, and hierarchical conflicts. There may be limited connectivity within the private network and to public switched networks. Locations within the private network may lack provisioning for BITS architectures.

Before a good synchronization plan can be developed for a given network, one must consider these issues in detail. But one must understand that these problems may be compounded because equipment and facility performance will also vary and affect network performance. The original stratum 4 clock specifications lead to consistently poor performance in complicated private networks.

CPE Clocks. Originally, stratum 4 clocks acted as the CPE's receive clock, which terminates timing; they were not designed to transfer timing references between systems. But as the private, customer-owned networks have grown more complex, the clocks in CPE (such as PBXs and multiplexers) have been used to transfer timing. Private networks often have long chains of stratum 4 systems through which timing is cascaded.

The main problem with stratum 4 clocks is not their low oscillator accuracy (i.e., 3.2×10^{-5}), but rather their inability to switch cleanly between timing references. If the timing reference for a stratum 4 clock is sufficiently impaired, the clock will switch to its backup timing source, which may be either another timing reference or the clock's internal oscillator. The original stratum 4 clock specification does not have a phase build-out mechanism. Phase build-out would slowly achieve and maintain a small phase difference at the output when switching between timing references. But because stratum 4 clocks do not provide for phase build-out, the reference switch would cause a phase hit at the CPE's digital interfaces. If this CPE is providing a timing reference to another piece of equipment, the receiving CPE will react to this degradation (which looks like an error burst) by inducing a reference switch of its own. Thus, these impairments will propagate through the chain of cascaded stratum 4 clocks, seriously affecting the network's performance.

Stratum 3 and stratum 2 clocks incorporate phase build-out and, therefore, do not suffer from this phenomenon. If a reference switch occurs, these clock systems correct for phase differences between references.

Today, digital network standards define a new level of stratum 4 clock that incorporates phase build-out when references are switched. For example, AT&T's TR 62411 (the interface specification for Accunet® T1.5 service),[10] which was last published in December 1988, defined an enhanced stratum 4 clock that can be used for timing transfer. This clock, which TR 62411 calls *stratum 4, type I*, meets the requirements for maximum time-interval error (MTIE) that regulates the total phase movement and the rate of change of phase movement of the timing signal under reference-switching conditions.

In addition, a new EIA/TIA standard[11] is being prepared that defines the same enhanced stratum 4 system (currently referred to as *stratum 4E*). [The standard is a joint effort of the Electronics Industries Association (EIA) and the Telecommunications Industry Association (TIA).] We expect this standard to be published before the end of 1991.

Private-Network Performance. Overall, the typical synchronization performance of a digital private network can be up to 100,000 times worse than the performance of the AT&T network (in terms of frequency accuracy).[7] When new CPE synchronization standards for the telecommunications industry are adopted and implemented in private networks, this performance should improve by several orders of magnitude.

Generally, a high-quality synchronization plan is needed to minimize the occurrence of slips, error bursts, and phase hits in private networks.

Synchronization-Plan Design Issues. Some basic concepts need to be considered during the design of a synchronization plan:

- Whenever possible, it is important to adhere to the hierarchical rules of timing transfer. That is, timing should always be passed from low numbered stratum devices (e.g., stratum 1, stratum 2) to equivalent or higher numbered stratum devices (e.g., stratum 3, stratum 4).

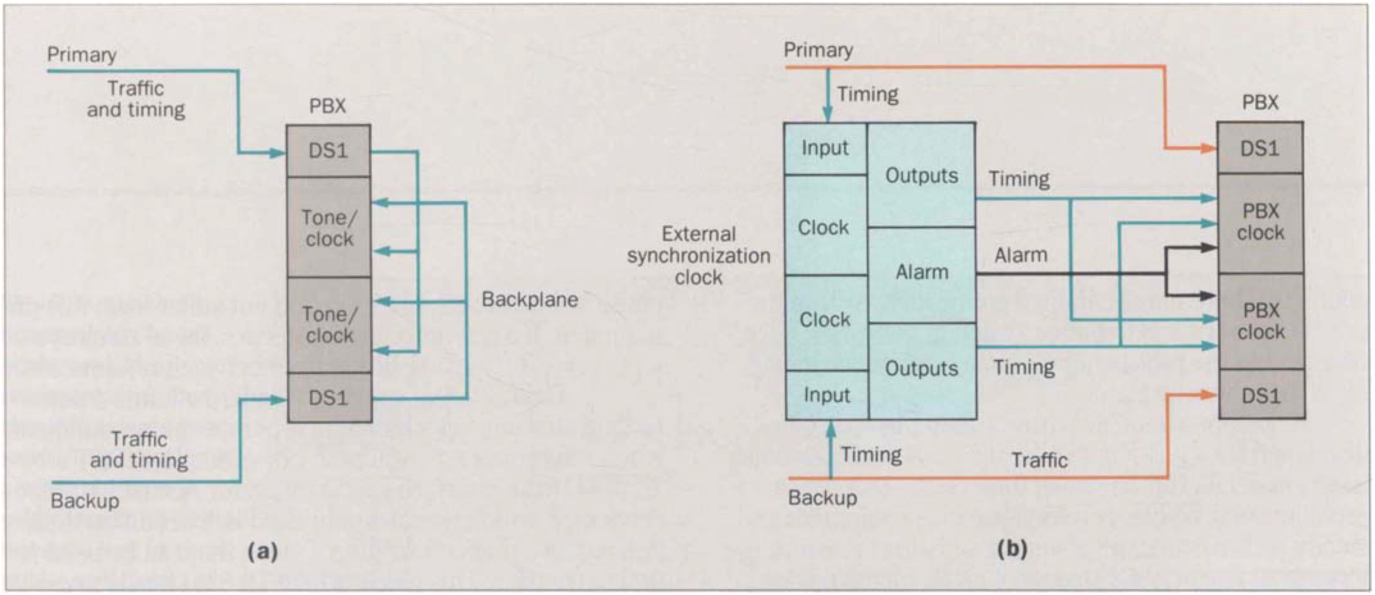

Figure 3. Reference clock scheme for Generic 1 of the Definity communications system, a PBX with dual input references. (a) Standard stratum 4, type I clock. The primary and backup inputs carry both timing and traffic. DS1-interface cards pass timing to the PBX's internal clock circuit packs on dedicated backplane paths. (b) Proposed stratum 3 clock. Here, the timing and alarm information goes to an external clock, which passes it to the PBX's internal clock circuit packs.

- Stratum 1, traceable timing sources need to be used whenever possible, and timing sources should be diverse whenever possible.
- Timing loops must be avoided. These loops occur when two or more CPE transfer timing to each other and form a loop without a designated master source.
- The cascading of timing references through CPE should be minimized.

These are some of the fundamental concepts that are the basis of more than 3000 rules that AT&T has developed for synchronization planning of digital private networks. The rules have been incorporated into a proprietary expert system, called *PRINCE-S*, that was developed by AT&T Bell Laboratories. PRINCE-S, which stands for private-network computer-based expert system for synchronization, is currently deployed to support AT&T's Network Design Organization.

Expert System for Synchronization Planning. AT&T developed the PRINCE-S expert system to mechanize the task of synchronization planning for a private network in a consistent, accurate, and reliable way. AT&T synchronization engineers have used the system to help them design synchronization plans.

The system generates network maps that specify primary and backup timing-reference distribution for each CPE, and estimates the synchronization performance of the network end to end (and for each CPE). The system's recommendation represents the best synchronization performance attainable, based on information used to model the network topology, equipment, facilities, and public-switched-network connectivity.

Synchronization Planning Service. AT&T now offers a synchronization planning service for private networks. This service gives customers the combined expertise of the PRINCE-S expert system and of certified synchronization engineers to help them develop synchronization plans for their digital private networks.

This offering, the Accunet synchronization planning service, analyzes the customer's existing synchronization plan (if one exists), and provides an *economic* synchronization design and an *optimized* synchronization design. The economic design outlines the best synchronization performance that can be achieved without the purchase of additional equipment. The optimized design is intended for more demanding applications or customer requirements and typically will recommend the placement of additional equipment to achieve optimized synchronization performance.

The planning service provides consultation and on-going planning support for network growth and changes. This planning is critical for maintaining the integrity of synchronization performance as the network evolves.

Definity Communications System

The Definity PBX product line illustrates AT&T's commitment to provide quality, CPE clocking systems. The desire to conform to the TR 62411 interface requirements for Accunet digital services began the evolution of clock performance for AT&T's System 75 (which was stratum 4) to the standard stratum 4, type I clock in Generic 1 of the Definity communications system. Generic 2 of the Definity communications system is now available and has a stratum 3 option; a similar option has been proposed for Generic 1.

The upgrade to the stratum 4, type I clock did not require architectural changes to Generic 1 of the Definity communications system (Figure 3a). Traffic and timing inputs enter the PBX system via the DS1-interface circuit pack. Each DS1 pack derives a timing signal from the incoming data stream, and provides it to the tone/clock circuit pack. The tone/clock circuit pack makes the decision to switch from the designated, primary timing reference to either the backup timing reference or an on-board, local oscillator based on the "health" of each clock. It determines health by detecting a loss of signal or clock-frequency drift.

Previously, switch-level processing controlled reference switching. Because intelligence was moved to the circuit pack, the Generic 1 PBX was able to meet the maximum-time-interval-error requirements in TR 62411.

Generic 2 of the Definity communications system works the same way (i.e., as a stratum 4 system). Its system-clock synchronizer circuit pack is located in the module control for single-module systems, or in the time-multiplexed switch for multiple-module systems.

To upgrade Generic 1 of the Definity communications system to accommodate the proposed stratum 3 clock, we can alter the architecture as shown in Figure 3b. As illustrated, an external clock taps the primary and backup inputs and passes the signals to the DS1-interface circuit packs. Now, the DS1 circuit packs deal only with traffic. Timing is provided to the PBX-clock circuit pack directly from the external clock via the tip and ring input of the clock circuit pack. Alarm information is passed to the clock circuit pack, which passes it to switch-level processing.

In the external clock, separate circuit packs control output-timing signals, alarms, clock accuracy and holdover, and the DS1-input interface. These packs are duplicated to provide failure-mode reliability, as well as permit field replacement without affecting the customer. The clock circuit pack on the external clock monitors the two incoming DS1 clocks, and (as required by TR 62411) shifts between them or to a holdover clock if both fail. Alarm information that the external clock provides includes failure of either input reference, failure of either circuit pack in the external clock, or failure of either power source. The external clock can be powered by either ac (alternating current) or 48 volt dc (direct current).

When the stratum 3 option is provided for Generic 1 of the Definity communications system, then Generic 1 and 2 will use a common, external, stratum 3 clock. This will minimize costs and allow smooth migration between PBX products. The PBX-clock circuit pack in Generic 1 of the Definity system currently operates as a stratum 4, type I clock.

Summary

We have explained how AT&T provides the means to give the best possible synchronization performance in intelligent digital networks. This encompasses the synchronization architecture of AT&T's national telecommunications network, as well as synchronization planning for customer-owned networks and CPE clock-performance issues.

We also briefly described AT&T's PRINCE-S expert system, Accunet synchronization planning service, and the synchronization systems used in AT&T's Definity product line.

Acknowledgments

We would like to thank John Abate, George Zampetti, Karen Swayze, Ed Butterline, and Pablo Alcivar for their contributions to the body of knowledge described in this paper. Thanks also go to Peter Ting, Elie Khawand, Rick Bournique, Prem Mehrotra, Jack Amoroso, and Kevin Hurley for the development and support of the PRINCE-S expert system.

References

1. J. F. Oberst, "Keeping Bell System Frequencies on the Beam," *Bell Laboratories Record*, Vol. 52, No. 3, March 1984, pp. 84–89.
2. "Synchronization Interface Standards for Digital Networks," *American National Standard for Telecommunications*, ANSI T1.101–1987, American National Standards Institute, New York, New York, 1987.
3. "Timing Requirements at the Output of Primary Reference Clocks Suitable for Plesiochronous Operation of International Digital Links," CCITT Recommendation G.811, CCITT IXth Plenary Assembly, Melbourne, Australia, 1988.
4. "Administration Plan for Synchronization of the Switched Digital Network," USITA Technical Advisory No. 58, Washington, D.C., June 29, 1979; reissued in 1981 and 1982.
5. J. E. Abate, L. H. Brandenburg, J. C. Lawson, and W. L. Ross, "The Switched Digital Network Plan," *The Bell System Technical Journal*, Vol. 56, No. 7, September 1977, pp. 1297–1320.
6. J. E. Abate, J. R. Rosenberger, and M. Yin, "Keeping the Integrated Services Digital Network in sync," *Bell Laboratories Record*, Vol. 59, No. 7, September 1981, pp. 217–220.
7. J. E. Abate et al., "AT&T's New Approach to the Synchronization of Telecommunications Networks," *IEEE Communications Magazine*, Vol. 27, No. 4, April 1989, pp. 35–45.
8. J. E. Abate, G. P. Zampetti, and E. Butterline, "Use of GPS to Synchronize the AT&T National Telecommunications Network," *Proceedings of the 20th Annual Precise Time and Time Interval (PTTI) Applications and Planning Meeting*, Washington, D.C., November 29, 1988, pp. 65–76.
9. AT&T, "Digital Synchronization Network Plan," Technical Reference PUB 60110, AT&T Information Release Service, Route 202/206 Bedminster, New Jersey, December 1983.
10. AT&T, "ACCUNET® T1.5 Service Description and Interface Specifications," Technical Reference (TR) 62411, AT&T Information Release Service, Route 202/206, Bedminster, New Jersey, December 1988.
11. "Private Digital Network Synchronization, DRAFT," EIA/TIA SP–2198, Electronics Industries Association, Washington, D.C., December 1990.

Biographies (continued)

Colorado College (Colorado Springs), an M.S. in physics from the University of Wisconsin at Madison, and a Ph.D. in electrical engineering from the University of Colorado (Boulder). Mr. Lanzafame is responsible for circuit-pack development and support, services support, and factory support for the Definity business communications system. He joined the company in 1979 and has an A.A.S. in electronics technology from the State University of New York (Farmingdale) and has a B.S.E.E. from Monmouth College (West Long Branch, New Jersey). Mr. Moffitt is responsible for architectural planning for the Systimax premises distribution system. He joined the company in 1978 and has an A.S.E.E. from Hartford State Technical College (Connecticut), a B.S.E.E from the University of Hartford (West Hartford, Connecticut), and an M.S.E.E. from Carnegie Melon University (Pittsburgh, Pennsylvania). Mr. Near is the prime-knowledge engineer of the PRINCE-S expert system, and is currently developing formal methodologies for testing expert systems. He joined the company in 1987 and has B.S. degrees in electrical engineering and applied mathematics from Northwestern University (Evanston, Illinois) and both an M.S. and Ph.D. in electrical engineering from Cornell University (Ithaca, New York). Mr. Russo, a member of the Network Synchronization Group, is responsible for the development of tools and methodologies to aid in the synchronization planning of domestic and international digital networks. He joined the company in 1987 and has a B.S. from Stony Brook University (New York) and an M.S. from Columbia University (New York, New York), both in electrical engineering.

(Manuscript received April 3, 1991)