Synchronization status messaging ensures SONET interoperability

New network realities have not lessened the need for the synchronization of SONET. It's just that the challenges have become greater.

Phil Mann and Santiago Quijano
Symmetricom Inc.

Synchronous Optical Network (SONET) networks are built for survivability. Routing traffic around failures is fast and reliable, as is reconfiguration of circuits and services. SONET nodes communicate with each other to create this powerful environment. One of the communications protocols used by SONET, synchronization status messaging (SSM), assures the operator that all network elements are operating at a common frequency, within the boundaries set by standards.

Or does it?

Certainly within a closed ring, operated by a single service provider, properly configured and provisioned, SSM can provide this assurance. But today's networks don't always fit these criteria. New access providers have created topologies that challenge network architects to find reliable methods of synchronization.

The good news is that SSM remains a viable mechanism for assuring synchronization, even across network boundaries. However, the service provider must consider several interoperability requirements when designing SSM into a network that is a fragment of a larger SONET universe. From the perspective of a new network operator, SSM offers a number of key advantages in operating a node served by another carrier's SONET ring. SSM provides both information and control, and it also entails new responsibilities as the new operator's node interacts with the service provider's ring.

Key network resource

Every network element (NE) has an internal clock, which determines the rate at which data is shifted through its input and output buffers. As a basic requirement for time-division multiplexed equipment in telecommunications networks, every network element must be "traceable" to a primary reference source: the NE clock is "slaved" to a reference frequency that is ultimately derived from an upstream atomic clock such as cesium. If an NE is not traceable, it will operate at some frequency that is offset from that of other NEs with which it is linked. This offset creates "slips"—lost data—as the buffers at either side of the interface overflow or underflow.

Prior to SONET, carriers devised a method of stratifying clocks in the network. Table 1 describes the clock stratum levels used as the core of carrier network synchronization. This method set up a "master-slave" system of matching clock

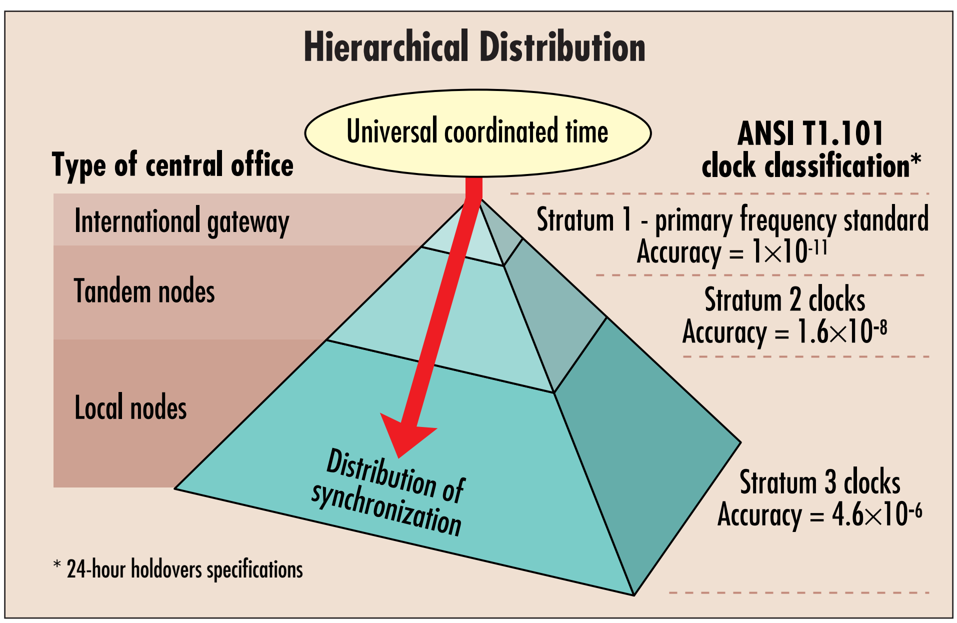

Fig. 1. The maintenance of proper synchronization requires that the timing hierarchy remain intact and that each clock output be traceable to the primary reference.

levels by class of node.

Using this hierarchical master-slave distribution system, the operator could be sure that all clocks were traceable to a primary frequency source. Timing was distributed from a Stratum 1 primary reference source at the top of the network to all other offices within that network (see Fig. 1). Following the hierarchy, DS-1 span lines distribute timing from higher-level offices to lower-level offices. Locking to the frequency of the incoming span line, a Stratum level 2 or 3 clock serves as the central timing source at each central office. A specific class of clocks—timing signal generators (TSGs)—was developed to meet the Stratum 2 and 3 requirements. All NEs in an office are slaved to this office timing supply.

The hierarchical method of distribution features strong constraints against the formation of "timing loops" and loss of traceability. When the output of a clock is connected back to its input, a timing loop is created where the clock drifts off frequency over time. Clearly, the clock is no longer traceable to the primary reference, and synchronization of all traffic through the office it serves is seriously degraded. While this problem can arise in hierarchical networks, careful design and provisioning minimize the possibility.

Another loss of traceability occurs when a higher-level clock (e.g., Stratum 2) receives its reference from a lower-level clock (e.g., Stratum 3). The problem occurs when the lower-level clock loses its connection to the primary reference because of a facility interruption or other network anomaly. The downstream Stratum 2 clock is then driven off frequency by the lower accuracy of the upstream Stratum 3 clock.

SONET, of course, has replaced hierarchical networks with a flat landscape of rings and chains. Given the use of rings, it is easy to see how timing could become circular in nature and not readily traceable to a primary reference source. Guarding against timing loops and loss of traceability takes on a new complexity in this "network of peers."

To help solve this problem, SONET networks generate synchronization status messages to tell the NE whether the incoming DS-1 is traceable to a primary reference source or if its timing is of a lower quality. The NE can then decide, based on these messages, which of several timing inputs will serve as its clock reference (see Table 2).

In some cases, older systems require upgrades for full SSM operation, and standards are evolving for the interface between SONET NEs and the TSG. But many SONET rings fully rely on internal SSM messages for traceability during rearrangements, failures, and fault recovery, counting on SSM-capable SONET NEs and TSGs to exchange SSM messages and switch to the highest-quality timing reference available.

The design objectives for synchronization status messaging within SONET networks are very straightforward:

- Maintain traceability to a primary reference source
- Prohibit higher-stratum clocks from being slaved to lower-stratum clocks
- Assure that each NE has a primary and secondary sync reference
- Avoid timing loops in the event of a failure.

The reliability of SONET networks in service today is sound testimony to the effectiveness of SSM as a means of keeping these networks on frequency.

SSM remains a viable mechanism for assuring synchronization.

But today's networks are becoming more fragmented as new operators connect new nodes into existing infrastructures. There are more network boundaries and interface points to manage. If synchronization traceability is not considered, the operator risks service degradation within the newly deployed node.

Using access or transport not directly in control of the network operator poses the greatest difficulty. The simplest example is where one operator's access node is served by another carrier's SONET ring. Assuring traceable synchronization is made easier if synchronization status messages can be received from the serving carrier's ring.

Take as an example a large-scale Asynchronous Transfer Mode (ATM) service, where the operator is one of several providers in a SONET ring. Proper operation of ATM requires the application of "network clock" at ingress and egress switch points. Network clock ensures that the transmission is traceable to a primary reference clock (PRS in SSM-speak, the message that indicates a quality level of PRS Stratum 1).

The timing output of the SONET add/drop multiplexer (ADM) can deliver SSM to the node, informing the operator that the ADM clock is indeed traceable to PRS Stratum 1. Although the SSM quality message is a reflection of the last node where the SONET payload signal originated and not the actual quality delivered by the ADM, reading the SSM at the far end is a good indicator of traceability. Once the operator has access to this information, it can be incorporated into a nodal timing solution that assures traceability to a primary reference source.

The operator's new opportunity

It might be instructive to consider how a new operator with a single node served by another carrier's SONET ring might use SSM. Assume that the operator is providing both voice and packet data services over some loop transport carrier directly to end users (see Fig. 2).

SSM gives the operator access to information about the quality of synchronization being received from the SONET ring. The question becomes: What action can be taken if the synchronization quality falls below the minimum needed for proper operation of services switched through this node? With Stratum 3 synchronization quality, the node's traffic at interfaces to customer premises equipment and other carriers' networks begins slipping within a few minutes, as all NE clocks drift toward a 1×10^{-7} offset from the primary reference.

This synchronization level can have a devastating effect on end-user services.

Table 1: Clock Stratum Levels Before SONET

	Clock level	Technology	Cost
Primary reference sources	Stratum 1	Cesium atomic clock	Moderate to high
		GPS radio receiver Loran radio receiver	Low to moderate
Timing signal generators	Stratum 2	Rubidium atomic clock Quartz crystal clock	Low to moderate
	Statum 3	Quartz crystal clock	Low to very low

Ingress and egress of constant bit-rate services over ATM, for example, are difficult if not impossible to sustain at this level.

Of course, the operator for this node can conceivably contact the carrier responsible for the SONET ring and request that Stratum 1-traceable timing be restored. But that leaves two problems: How does the operator detect the change in SSM quality level and can the carrier react in real time to correct the problem? It is easy to see how significant delays could occur, particularly if the reason for the SSM change is related to a complex failure that is preoccupying the serving carrier.

The solution depicted in Figure 3 allows the operator to control the quality of synchronization in the node, and to do so in a cooperative way with the SONET ring carrier. Because the economics of synchronization have changed dramatically in the operator's favor, it is now possible to deploy an affordable primary reference source directly into the node, not just at the top of the network. An installed PRS is now less than $10,000, down from the $25,000-to-$70,000 range of recent years.

Using the ADM timing output as one reference and the PRS as another, the operator can choose the best source of timing for the node. TSGs can be provisioned to accept both inputs, implement the SSM transaction tables that govern reference switching, and deliver output messages according to the inputs received. In the Figure 3 example, the TSG receives an ST3 message from the ADM, switches to the PRS reference, and continues to deliver PRS-quality timing to the node. If the ADM is collocated with the operator's node, then the operator can deliver a PRS-quality timing output back to the ADM. This strategy ensures that the ADM is locked to the

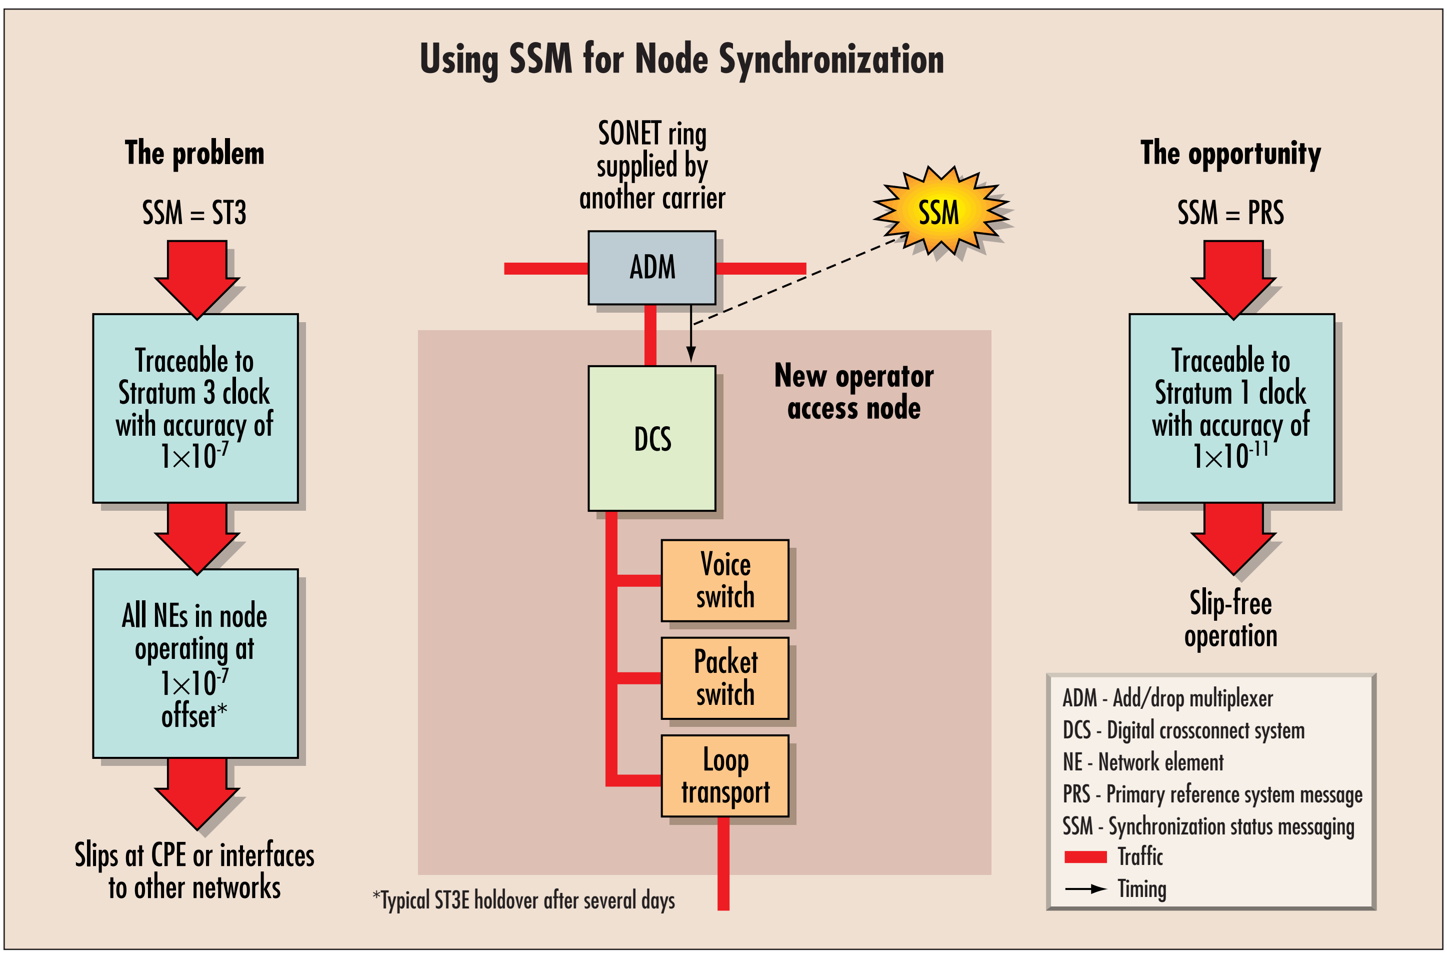

Fig. 2. SSM can prove useful in applications where a new operator's single node is operated by another carrier's SONET ring. With typical synchronization schemes, timing slips can easily occur.

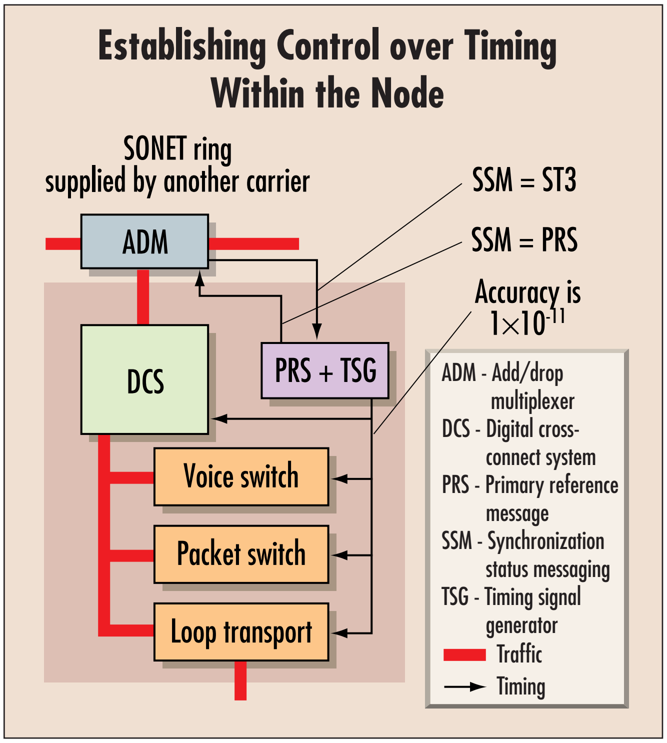

Fig. 3. A combination of PRS and TSG can make timing within the node possible and economical.

same frequency as all other NEs in the node and is not a source of slips. Accordingly, the TSG can deliver an SSM message of PRS to the ADM to indicate the quality of the timing input.

As a side benefit, installing the TSG can improve sync distribution within the node. Following the guidelines in Bellcore's GR-378 specification, the operator can directly time all NEs in the node from the output of the TSG. All NEs are always slaved to a PRS-traceable clock, not through an intervening NE such as a digital crossconnect system.

By installing a local PRS, the operator gains control over synchronization within the node. But with control comes responsibility. With the TSG providing SSM information to a collocated ADM, the operator becomes an active participant in the health of the SONET ring.

The operator must comply with the SONET carrier's synchronization plan for the overall SONET ring, so that SSM messages emanating from this node drive appropriate behavior within the ring itself. The two scenarios depicted in Figure 4 illustrate this concept.

In the first scenario, there is no downstream effect on the ring when the ADM's incoming message changes from PRS to ST3. The node's PRS becomes the reference chosen by the TSG, whose output message remains at PRS. If provisioned to do so, the ADM will switch to the TSG as its primary timing input and will continue to output PRS.

In the second scenario, a downstream effect clearly does occur when the ADM's incoming SSM changes from PRS to ST3. The difference is that the operator in this scenario has not elected to deploy a PRS, relying on a Stratum 2 level TSG and the ADM's timing output for synchronizing the node. With the loss of PRS on the ADM timing feed, the TSG sends out an ST2 message, which is its own stratum classification. Depending on the status of its other inputs, the ADM could switch to the TSG as its reference, and an ST2 message would be forwarded downstream. Receipt of this message by downstream NEs could cause rearrangements of the ring's synchronization trails. Proper network design and provisioning are needed to make sure that insertion of this "external" TSG reference will not create timing

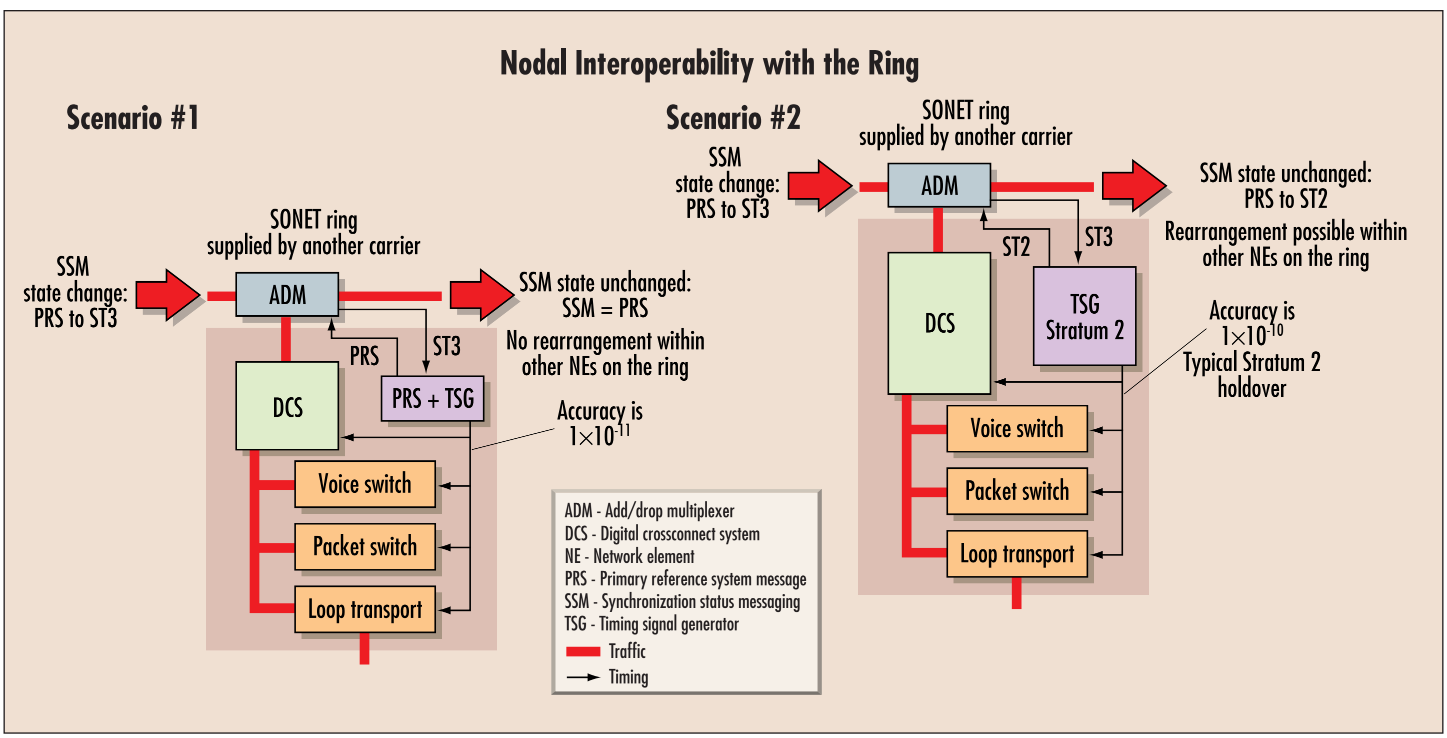

Fig. 4. In scenario #1, a change of the ADM's incoming message from PRS to ST3 has no downstream effect, thanks to proper implementation of the TSG. In scenario #2, the lack of a PRS has a detrimental effect on network synchronization.

loops within the ring.

Of course, the operator can choose not to provide a timing input to the SONET carrier's ADM, avoiding any impact on the ring whatsoever. Interesting trade-offs abound in this decision:

- The operator runs the risk of slips occurring at the interface between the ADM and the NEs in the node if there is no TSG timing interface to the ADM.
- The SONET carrier could benefit from the addition of a PRS or a TSG into this ring, but could have a more complex SSM environment to manage.
- Standards are not final for an SSM interface between the ADM and the TSG, so both the operator and the SONET carrier would need to collaborate on an approach that allows seamless upgrade of SSM functions on both sides of the interface.

Experience with the interface between SONET rings and external TSGs is still somewhat limited, and experimentation continues, both in the lab and in actual networks. Standards are still being debated and may evolve well beyond the current drafts being considered.

But operators have networks to run and services to deliver. SSM is viable today, given proper planning and management.

The next level of interoperability is monitoring and management of the sync resource for the services within the node itself and for the node as an integrated part of the ring's synchronization messaging plan. Two things seem clear: There exists a mutual need for interoperability of SSM, and both the operator and SONET ring carrier need to separately manage their own parts of the network.

Is there an intersection of these two requirements? Perhaps. Consider the high-level solution proposed in Figure 5. With modern management systems, it is possible to share information between network operators, given proper protection and message filtering. Achieving this level of interoperability is more a business issue for the operator and the SONET ring carrier than a technical matter.

For the operation of the node itself and the services within it, a management solution can usually be justified without even considering benefits arising from interoperability. But once deployed, the operator's management system can be the vehicle for any future arrangement with the SONET ring carrier.

The ability to manage synchronization within the node so that service level agreements (SLAs) with the operator's customers can be fulfilled can be an overriding justification and involves monitoring

- incoming SSM messages from all available sources
- reference source selected on the basis of SSM messages received

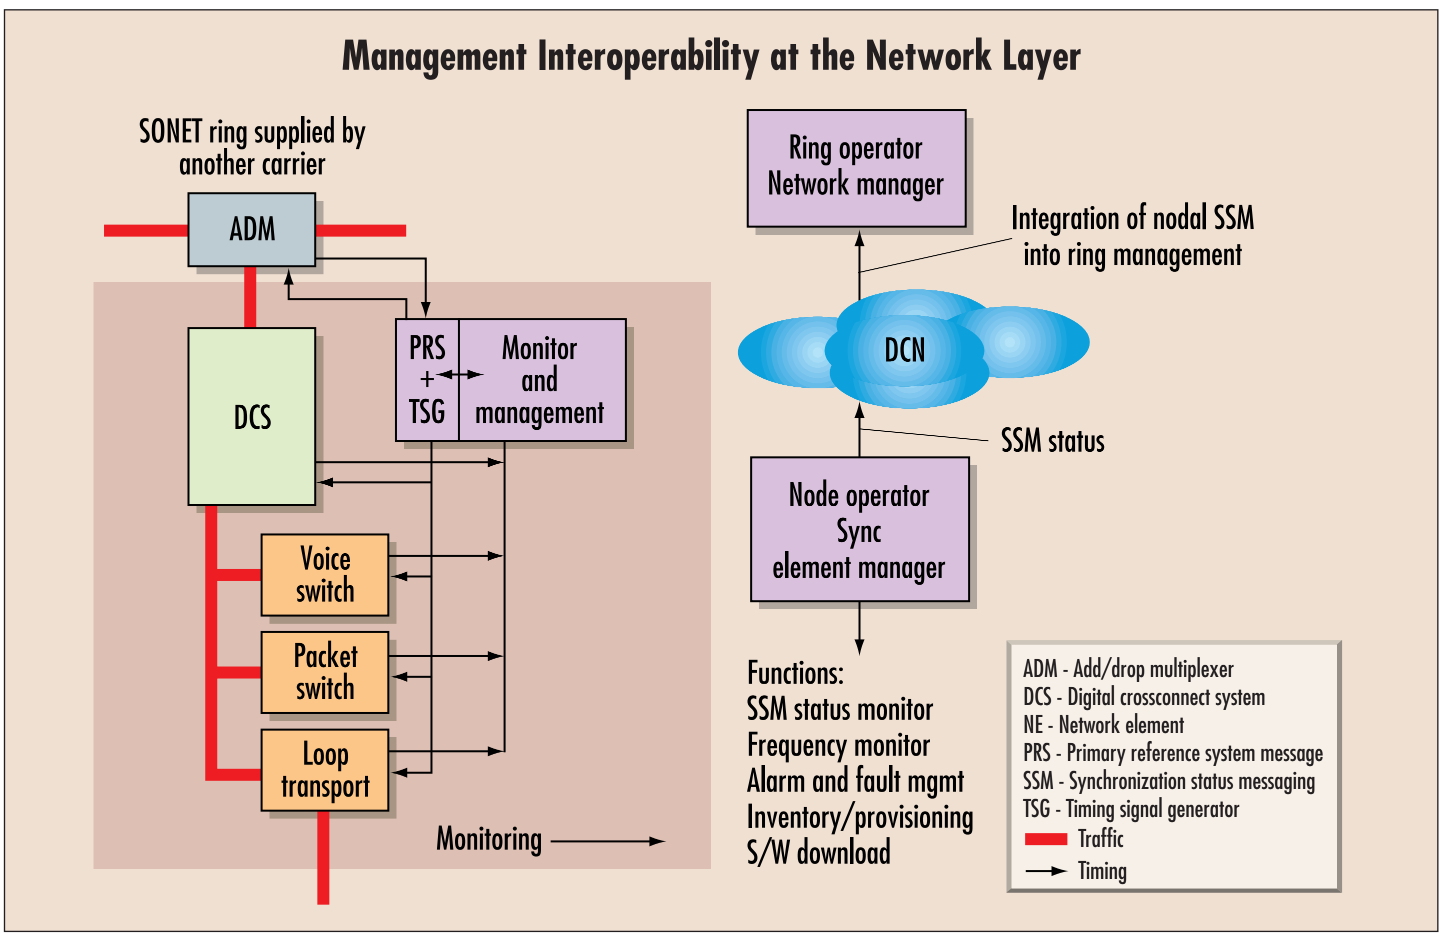

Fig. 5. Both the operator and the SONET ring carrier can manage the synchronization of their own parts of the network separately using modern management systems.

- outgoing SSM messages and comparison to the SSM state table prescribed in standards
- to determine if non-SSM inputs to the TSG are in either a "locked" or "holdover" state so they may also be used as references
- frequency characteristics of all NE clocks in the node to determine whether they are properly slaved and traceable
- frequency alarm thresholds on selected NEs to resolve problems before they become SLA violations.

Networks can behave in unexpected ways. When they do, traceable synchronization can be interrupted or SSM states may become flawed. The operator needs tools to create new synchronization trails by manually setting SSM parameters or by intervening so service is not impaired or interrupted. Controls available as part of a synchronization management system should include the following:

- SSM override input and output
- Settings for incoming message persistence before reference switching
- Provision to hold off switching
- Settings for wait-to-restore time
- Generate squelch or AIS
- Assign SSM to locked or holdover state on a non-SSM facility.

Beyond the day-to-day operational advantages of a managed SSM environment, the operator can also use the management system to evolve each TSG to reflect changes in standards or to accommodate a new type of NE. From a central element manager, the operator can install new message types and new SSM transaction definitions as dictated by changing standards. Also, SSM operation throughout the network can be optimized by integrating transaction parameters specific to certain types of NEs, deviating from standards as a conscious decision. All these capabilities are part of a complete SSM management system, where commands can be sent to configure TSG parameters or download new software to implement a change in standards.

SSM is a valuable tool for assuring quality of synchronization in today's rings-and-chains SONET environment. New operators can derive substantial benefits from the use of SSM in new nodes served by SONET ring carriers. At the same time, the operator has a responsibility to make sure the node is an integral part of the SONET ring carrier's overall SSM design for the ring.

TSGs and associated management systems are now available that enable operators to take full advantage of SSM capabilities today and as standards and network topologies evolve. ◆

Table 2: Quality Level Messages

DS-1 SSM message	Description
PRS	PRS Stratum 1
STU	Sync traceability unkown
ST-2	Stratum 2
TNC*	Transit node clock
ST-3E*	Stratum 3E
ST3	Stratum 3
SMC	SONET minimum clock
ST-4	Stratum 4
DUS	Don't use for sync

Per Bellcore GR-378-CORE Issue 2, February 1999
*Standardization of these messages is still in progress.

Phil Mann is marketing director and Santiago Quijano is senior product manager at Symmetricom Inc. (San Jose, CA).