Virtual Concatenation + LCAS;

Providing Scalable SONET/SDH Bandwidth

White Paper

By Fredrik Olsson, Jim Shupenis, Alisdair Gunn

Agilent Technologies
Semiconductor Products Group

Agilent Technologies

Abstract

SONET and SDH are the primary low-level protocols for all types of metro and long-haul traffic. They were originally designed to provide static bandwidth connections between points with a high degree of resiliency, which is suitable for voice and ATM traffic.

As these networks are used more and more to carry data, one major inefficiency is becoming apparent – the lack of scalability in connection sizes. Modern Ethernet links can provide bandwidth in very small increments, using flow control, and can be changed in an instant. SONET/SDH links at higher speeds are limited to factors of four, however they are traditionally plagued by long provisioning times. So while the SONET/SDH incumbents have superior service guarantees and offer a complete service, including the profitable transport of voice traffic, Ethernet based Metro Carriers are able to offer a better service; buy exactly the bandwidth you need now and change it at any time.

This situation is now about to change with new Metro and Long Haul products soon to hit the market. These products will support new features like Virtual Concatenation and LCAS.

The remainder of this paper uses SONET terminology for readability, but the cases are the same as for SDH networks.

Benefits of Virtual Concatenation

The use of virtual concatenation provides four key advantages; Scalability, Efficiency, Compatibility and Resiliency.

Scalability – SONET pipes can be sized to match the desired data rate and thereby avoid unnecessary waste. While traditional contiguous concatenation comes in coarse steps, virtual concatenation allows the bandwidth to be tuned in small increments on demand

For example, if a contiguous concatenated STS-12c (599 Mb/s payload) was not quite big enough for a specific link, a full STS-48c (2,400 Mb/s payload) needs to be allocated. However, with virtual concatenation, an STS-12v (599 Mb/s) can be slightly up-sized to an STS-13v (649 Mb/s) to meet the needs and thereby leave the additional bandwidth for other links.

Figure 1 shows a typical case where Gigabit Ethernet Data is transported over a SONET network. With legacy contiguous concatenation, the utilization is poor. With virtual concatenation, an OC-48 link can actually carry two full Gigabit Ethernet links and still have 6 STS-1s available to carry other traffic.

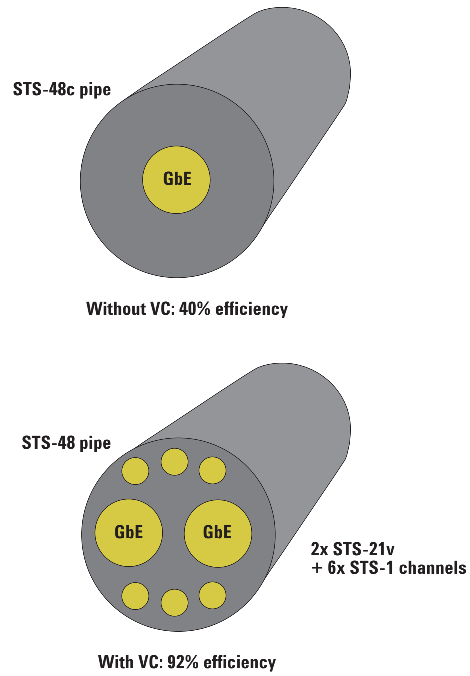

Figure 1. Scalability by using Virtual Concatenation.

Efficiency – Virtually Concatenated channels are more easily routed through a network and eliminate stranded bandwidth.

As an analogy, imagine a team of 100 people that needs to travel from Denver, Colorado to Rome, Italy for a meeting. The legacy "contiguous concatenation" travel agency (Network Management System) has a challenging task. The complete group must not only take the same flights but also sit in adjacent seats on each plane. If no such itinerary is possible, the group's order (connection request) is rejected. Here half-empty planes are common as the travel agency is forced to reject orders from customers whose requirements cannot be met. The capacity utilization is poor.

The work at the "virtual concatenation" travel agency is much more relaxed. The individuals in the 100 people group can take different flights and different routes and thereby arrive in Rome at different times. Once all members of the group have landed, the group is reassembled for the meeting. Now planes are mostly full, as every seat can be utilized and no orders are rejected. Hence the revenue is maximized for the airline (carrier) and expensive capacity upgrades can be delayed.

Virtual concatenation allows for more efficient usage of an existing network's available bandwidth.

Compatibility – Virtual concatenation works across legacy networks. Only the end nodes of the network are aware of the containers being virtually concatenated, as this is fully transparent to the network. Hence, with virtual concatenation, large data channels can be routed over older networks that do not support large contiguous channels.

Resiliency – Best effort type traffic is often carried over unprotected links or in the protection channel of high-priority traffic. In the event of a link failure, the high priority instantly reclaims the protection bandwidth and the best effort traffic across that link is halted.

For a contiguous channel, this means that the best effort traffic service is interrupted completely i.e. all data is lost.

Individual members of a virtually concatenated channel are recommended to be routed as diversely as possible across a network. So if one link goes down, the others are likely still to be operational. The virtually concatenated channel thereby loses only one tributary in the event of a link failure and the link can still continue to provide the best effort service, albeit with a reduced bandwidth.

LCAS – Bandwidth on Demand

While it is beneficial to use virtual concatenation alone, powerful advantages are gained in coupling it with the Link Capacity Adjustment Scheme, LCAS.

Virtual concatenation provides the means for creating right-sized pipes. But in many applications the size of a right-sized pipe changes with time. What was right-sized a week ago may very well be a seriously under-sized channel today. Bandwidth requirements change.

When a virtual channel is resized, traffic is disrupted and lost. Strict Service Level Agreements, SLAs, often limit the amount of acceptable traffic disruptions and thereby effectively limit channel resizing.

LCAS is the solution to this problem. With LCAS, channels can be resized at any time without disturbing the traffic on the link. Also, connectivity checks are continuously performed and failed links automatically removed and added back as the link is repaired, without intervention of the (slow) network management system.

Now, carriers can dynamically change the bandwidth allocated to a connection. As an example, bandwidth demand may increase during the nights and at weekends. A customer may buy a 100 Mb/s connection that increases to 1000 Mb/s between 2:00 am and 3:00 am every night as their computer system creates data back-ups. During the daytime this bandwidth is not needed and can be re-allocated to other customers.

SONET Refresher

Data that is to be transported by SONET is first wrapped with two layers of overhead. The first layer is the path overhead, POH, which is attached to the original data when entering the network and stripped off at the final destination. The data with the POH attached is called the synchronous payload envelope, SPE. A contiguous stream of SPEs are injected into the network by the source node and are then transported to the destination node.

POH allows for end-to-end data integrity checks. All information for Virtual Concatenation and LCAS is carried in the H4 POH byte.

The second layer of overhead is the transport overhead, TOH. This is examined and recreated at each node along the path to measure performance of each link as well as to carry network management commands. The TOH is not relevant to this Virtual Concatenation and LCAS discussion and will not be shown in the figures shown below.

Figure 2 shows an example of how a six-byte packet (A to F) is carried in the payload surrounded by pad bytes.

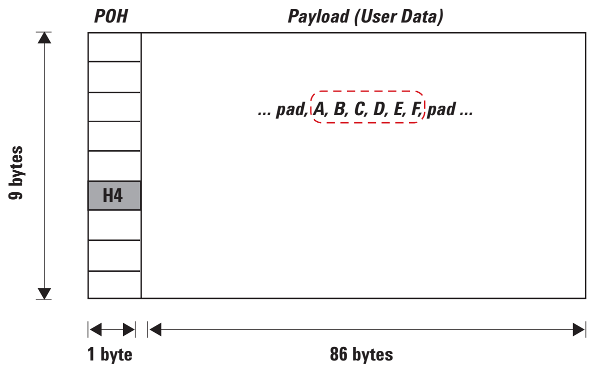

Figure 2. An STS-1 SPE.

Virtual Concatenation Explained

The basic principle of virtual concatenation is really quite simple. A number of smaller containers are concatenated and assembled, to create a bigger container that carries more data per second.

Virtual concatenation is possible for all container sizes from VC-11/VT-1.5 up to and including VC-4/STS-3c. Smaller containers allow for finer granularity but less maximum channel sizes and also require the network to be able to switch down to that level. Table 1 shows the containers for which virtual concatenation is defined and their bandwidth ranges.

SONET Name	SDH Name	Min Size (Mb/s)	Max Size (Mb/s)	Granularity (Mb/s)
VT1.5	VC-11	1.6	102	1.6
VT2	VC-12	2.2	139	2.2
VT3	—	3.4	217	3.4
VT6	VC-2	6.8	434	6.8
STS-1 SPE	VC-3	48	12,000	48
STS-3c SPE	VC-4	150	38,000	150

Table 1. Virtual concatenation base containers and approximate bandwidths.

The remainder of this paper uses examples with virtual concatenation of STS-1/VC-3 containers that are often the base containers of choice in a SONET network.

Figure 3 shows how the packet from the previous section is now carried by three virtually concatenated STS-1s. The packet is byte-interleaved between the STS-1s. These three STS-1s are then transported independently by the network to the destination. When traveling through the intermediate SONET network, these STS-1s will incur different delays. At the destination, realignment of the STS-1s is necessary before extracting the packet.

In general, a virtually concatenated channel made up of Nx STS-1 is transported as individual STS-1s across the network and at the receiver the individual STS-1s are re-aligned and sorted to recreate the original payload.

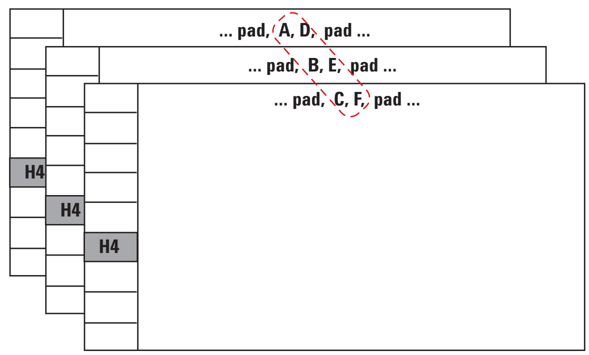

Figure 3. Three virtually concatenated STS-1s.

The H4 POH byte carries information on how to reassemble the SPEs. A sequence of 16 H4 bytes makes a complete message. This message contains two important numbers; the multi-frame indicator, MFI, and the sequence number, SQ.

Virtual Concatenation—MFI

The MFI indicator is a running frame number, which is incremented with each new frame. In Figure 4, the three STS-1s shown all have the same MFI. The next three STS-1s transmitted will have their MFIs incremented by one.

At the destination, due to different delays, the MFIs will no longer necessarily be the same for the three STS-1s. For example, the first STS-1 may have a trace that is 1.25 ms (ten frames) faster than the second and third STS-1. Hence at any given instance, the MFI number of the first STS-1 at the destination will be ten higher than the other STS-1s. In order to extract the packet (ABCDEF), the destination node needs to compensate for the different network delays by delaying the first STS-1 by 1.25 ms.

The 12-bit MFI number allows end nodes to compensate for up to 256 ms of differential delay.

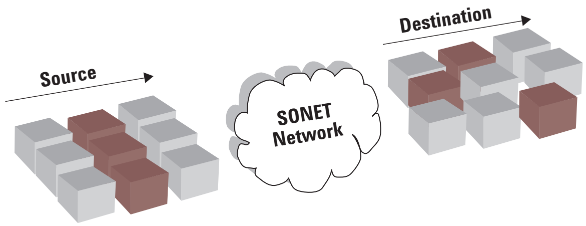

Figure 4. STS-1s experience different delays.

Virtual Concatenation—SQ

The source node labels each STS-1s, in a virtually concatenated channel, with a sequence number indicating its relative position. An STS-Xv channel will have SQ number zero to (X-1). In Figure 4, the upper STS-1 is assigned SQ #0, the middle SQ #1 and the lower SQ #2. At the destination, the STS-1s are reordered according to the sequence numbers to guarantee that the packet is extracted correctly as “ABCDEF”, not as “BCAEFD” or “CBAFED”.

This SQ number relieves the network management of having to keep track of the order of each individual trace through the network. As long as the intended STS-1s are routed to the destination node, the order within a channel is sorted out at the destination.

LCAS Explained

Virtual concatenation allows for any-sized bandwidth. LCAS is a protocol to synchronize the re-sizing of a pipe in use, so it can be changed without corrupting packets in the process. It also provides automatic recovery of a link after tributaries failures.

Virtual concatenation can be used without LCAS, but LCAS requires Virtual Concatenation. LCAS is resident in the H4 POH byte of the SONET overhead, the same byte as virtual concatenation. The H4 bytes from a 16-frame sequence make up a message for both virtual concatenation and LCAS. Virtual concatenation uses 4 of the 16 bytes for its MFI and SQ numbers. LCAS uses 7 others for its purposes, leaving 5 reserved for future development.

While virtual concatenation is a simple labeling of individual STS-1s within a channel, LCAS is a two-way handshake protocol. Status messages are continuously exchanged and consequent actions taken.

Each STS-1 carries one of six LCAS control commands.

- Fixed—LCAS not supported on this STS-1.
- Add—Request to add this STS-1 to a channel, thereby increasing the bandwidth of an existing channel or creating a new channel.
- Norm—This STS-1 is in use.
- EOS—This STS-1 is in use and the last STS-1 of this channel, i.e. the STS-1 with the highest SQ number.
- Idle—This STS-1 is not part of a channel.
- Do not use—This STS-1 is supposed to be part of a channel, but is removed due to a broken link reported by the destination.

A typical sequence when up-sizing a link is as follows:

1. The network management system adds a new trace through the network between the source and destination node.
2. The network management system orders the source to add this new link to the existing channel.
3. The source node starts sending “Add” control commands in this STS-1.
4. The destination notices the add command and returns an OK in the link status for this STS-1.
5. The source sees the OK, assigns this STS-1 an SQ number one higher than currently in use by this channel.
6. At a frame boundary, the source includes this STS-1 in the byte interleaving and sets the control command to “EOS”, indicating that this STS-1 is in use and the last in the sequence.
7. The STS-1 that previously was “EOS” now becomes “Norm” as it is no longer the one with the highest SQ number.

Multiple STS-1s can be added to or removed from a link concurrently to allow for fast resizing.

Summary

Virtual Concatenation and LCAS will change how carriers use their Metro and Long Haul Networks and will allow them to provide services more efficiently. SONET and SDH networks are overcoming old deficiencies like long provisioning times and inflexible services and will thereby remain highly competitive into the foreseeable future.

www.agilent.com/semiconductors

For product information and a complete list of distributors, please go to our web site.

For technical assistance call:
Americas/Canada: +1 (800) 235-0312 or (408) 654-8675

Europe: +49 (0) 6441 92460

China: 10800 650 0017

Hong Kong: (+65) 271 2451

India, Australia, New Zealand: (+65) 271 2394

Japan: (+81 3) 3335-8152(Domestic/International), or 0120-61-1280(Domestic Only)

Korea: (+65) 271 2194

Malaysia, Singapore: (+65) 271 2054

Taiwan: (+65) 271 2654

December 20, 2001
5988-5107EN

Agilent Technologies