

Virtual concatenation tutorial: enhancing SONET/SDH networks for data transport

Larry Choy

*Cisco Systems, 170 West Tasman Drive, San Jose, California 95134-1706
lchoy@cisco.com*

Received 5 November 2001; revised manuscript received 10 December 2001

Synchronous optical network (SONET) and synchronous digital hierarchy (SDH) contiguous concatenation have enabled high-speed networking in both metro and core networks for many narrow-band and broadband services including digital subscriber loop (DSL), cable, and Ethernet. Ironically, the tremendous success and rise of data traffic have, in part, put the focus on some of the inefficiencies and inflexibilities of SDH contiguous concatenation. Here I examine these shortcomings, survey available bandwidth efficiency techniques for data transport, examine two relatively new mechanisms known as virtual concatenation and link capacity adjustment scheme (LCAS) that address these issues, and finally present several potential applications for these technologies. © 2001 Optical Society of America

OCIS codes: 060.4510, 060.4250.

1. Introduction

The synchronous optical network (SONET)^{1,2} and the synchronous digital hierarchy (SDH)³ networks have indeed proven their value to the telecommunication and networking industry. The ingredients for this success include bandwidth scalability, infrastructure scalability both nationally and internationally, proven operation administration management and provisioning (OAMP) functionality, and robust protection and restoration. For the rest of this paper the term SDH will be used to represent both SONET and SDH unless indicated otherwise, and examples will be provided in either SONET or SDH format.

The rise of the Internet and the widespread adoption of Internet technologies deployed within corporate Intranets have fueled the explosive growth of bursty Internet protocol (IP) traffic and have put to question some of the inefficiencies of SDH's handling of data traffic that result from its 64-Kbit/s voice-transport origins.

Fortunately, the International Telecommunication Union Telecommunication Standardization Sector (ITU-T) and the American National Standards Institute (ANSI) T1X1.5 working group (T1X1.5) have made SDH virtual concatenation available. Virtual concatenation would allow use of best-size SDH path bandwidth containers in the carriage of higher-layer protocol without additional overhead or delay. It comes in two flavors, high-order (HO) virtual concatenation for higher-speed applications and low-order (LO) virtual concatenation for modest-speed applications. In addition, these groups have also developed a method known as the link capacity adjustment scheme (LCAS^{2,4,5}), which is a low-level bandwidth capacity control protocol for virtual concatenation.

Before the examination of virtual concatenation and the LCAS, a review of the inefficiencies with contiguous concatenation is provided in Section 2.

2. Contiguous Concatenation and Its Inefficiencies

Contiguous concatenation has been a part of the SDH specifications from the early days of SDH networks. It was conceived to accommodate high-speed data applications that use protocols such as asynchronous transfer mode (ATM) and high-level data-link control (HDLC) over SDH, and others. These protocols required high-bandwidth mappings over SDH. To accommodate these payload mappings, multiple SDH payload

containers are transported and switched across the SDH network as a single module with the first SDH container payload pointer used in normal mode and the subsequent payload pointers set to concatenation mode thus linking all the units together. The payload containers would arrive in phase alignment at destination. This method of transport is known as contiguous concatenation. SDH international examples of HO contiguous concatenation include but are not limited to VC-4-4c and VC-4-16c. SONET examples of HO contiguous concatenation include but are not limited to STS-3c and STS-12c (STS is synchronous transport signal). Within the SDH multiplex section, H1 and H2 pointer overhead bytes are responsible for locating the SDH payload and provide contiguous concatenation indications to the receiver.

Fixed and contiguous rates could cause waste of bandwidth resulting from large jumps and mismatches to native application data rate. Examples are illustrated in Table 1 identifying the amount of potential waste in specific cases. In the case of Fast Ethernet applications, if an STS-3c circuit is provisioned, 33% wastage of provisioned bandwidth is not used. In the 67% of the bandwidth intended for the Fast Ethernet service there is also potential waste resulting from the statistical and bursty nature of data traffic where the actual use could be much lower. It is also possible that the service provider may want to offer different incremental levels within the 100-Mbit/s service, for example, charging by 1-Mbit/s increments. But with the contiguous fixed-rate method of SDH, there is no accommodation to share, reallocate, and reuse that bandwidth.

Beyond the mismatched inefficiencies, some of the legacy SDH equipment may not support contiguous concatenation transport switching at higher rates such as VC-4-4c and VC-4-16c in some regions. In addition, when the equipment supported such functionality, locating a contiguous-concatenation-capable set of virtual container slots across a provider's SDH network is a time-consuming task that adds to the cost of the service.

Virtual container fragmentation is another problem, similar to the memory and disk fragmentation problem seen in computer systems. This virtual container fragmentation problem is exhibited when physically available SDH paths exist but are not usable because of the noncontiguous concatenation nature of the available containers. Figure 1, steps 1–4, below illustrates the virtual container fragmentation problem.

Over time and for larger networks, the number of unused bandwidth containers will

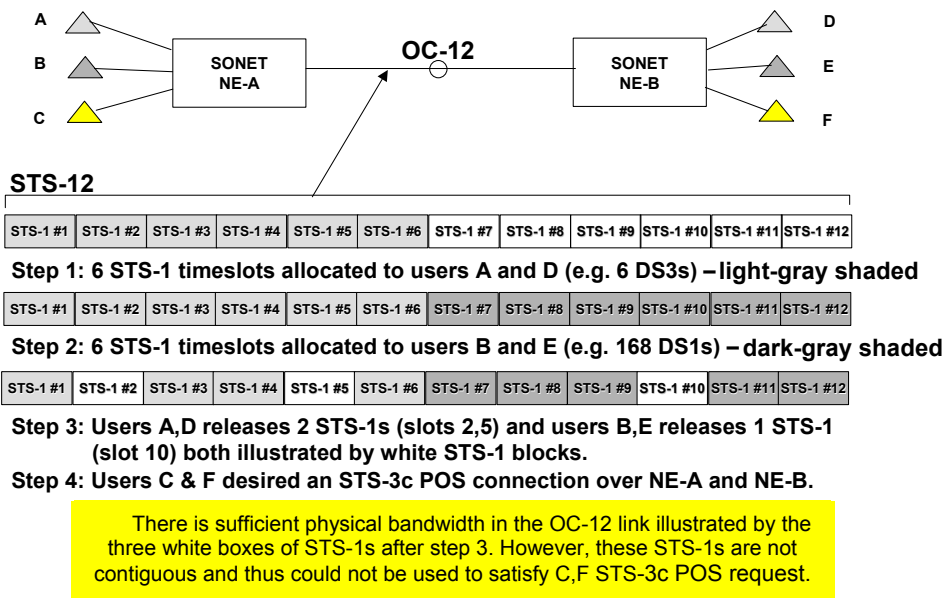


Fig. 1. Virtual container fragmentation problem.

build up resulting in suboptimal use of the metro and core infrastructure. To counter this problem, service providers and vendors at times employ regrooming to recover this

bandwidth. In some situations, regrooming may not be an option. Regardless, the recovery of bandwidth will be time consuming in both planning and execution.

A number of techniques are available to improve data-transport efficiency. A listing and a brief review of these techniques are provided in Section 3.

3. Available Efficient Data-Transport Techniques

A number of techniques are available to provide more-efficient data transport. These techniques are listed below with a brief discussion of how they relate to SDH networks and their position relative to addressing contiguous concatenation issues at the SDH layer:

- Time division multiplex (TDM) inverse multiplexing.
- Inverse multiplexing over ATM (IMA⁶).
- Multilink point-to-point protocol (ML-PPP⁷).
- Resilient packet ring (RPR).
- Equal cost multipath forwarding (ECMF⁸).
- Ethernet link aggregation [Institute of Electrical and Electronics Engineers (IEEE) 802.3ad standard]

TDM inverse multiplexing can operate end to end over TDM networks but imposes a nonstandard solution requiring same vendor equipment at both ends of the link. In addition, it is applicable only to lower-speed DS1/E1 links. IMA operates at a layer above SDH but typically only with DS1/E1 ATM links. Since IMA is an ATM technology, the ATM and ATM adaptation layer 5 (AAL5) overhead inefficiencies in packet data transport apply. Although each solution enables load spreading of traffic across multiple DS1/E1 links and thus provides improved transport in its respective part or layer of the network, TDM inverse multiplexing does not address the contiguous concatenation issues at the SDH layer at either the HO or the LO bandwidth levels.

ML-PPP enables fragmentation and reassembly of upper layer packets, for example, IP packets inside point-to-point protocol (PPP), based on multilink headers. This method typically operates over DS/E1 links and is a good solution for increasing bandwidth incrementally for lower-speed data transport. However, quality of service (QOS), latency, and jitter must be considered when this technology is applied at much higher speeds if stringent service levels are required. In addition, ML-PPP operates at a layer above SDH and thus does not address the contiguous concatenation issue directly.

RPR is an emerging ring-based media access control (MAC) layer protocol currently under standards development in the IEEE 802.17 working group. Among other things, it offers improved efficiency as well as protection and restoration of data transport over fiber rings. Since it is a MAC layer technology, it can, in principle, operate over many different kinds of links including SDH and Ethernet. Since it is a layer above SDH, not only is it unable to resolve the contiguous concatenation at the SDH layer; it in fact uses contiguous concatenation. Indeed, RPR can operate and be complementary when operating over virtual concatenation-enabled networks.

ECMF is packet or flow spreading technique for IP packet networks. Packet or flow spreading of traffic across multiple paths is performed to increase packet network efficiency and transfer resiliency. Therefore it is a technology that operates above the SDH or optical networks and does not address the contiguous concatenation issue directly.

Ethernet link aggregation is a technology that enables the spreading of Ethernet framed traffic over multiple Ethernet links. As such, it is not applicable to SDH networks.

Although none of these techniques address contiguous concatenation efficiency issues at the SDH layer directly, each of these mechanisms offers improved data-transport efficiency at its respective layer or part of the data-transport networks. In addition, those techniques that operate above the SDH layer, for example, ECMF, are indeed complementary and can bring additional packet-transfer efficiency and resiliency

when combined with techniques such as virtual concatenation and SDH protection and restoration methods. Given that none of these techniques resolves the issues of SDH contiguous concatenation, virtual concatenation has been developed in the standards to address these issues directly. In Section 4 we explore virtual concatenation, what it is, how it works, what its benefits are, and some other issues.

4. Virtual Concatenation

Contiguous concatenation requires SDH contiguous concatenation bandwidth mapping end to end throughout SDH transport. Virtual concatenation breaks the bandwidth into individual payload containers, SDH HO or LO paths, at the source transmitter end but logically represents them in a virtual concatenation group (VCG). The members of the VCG are routed and transported individually across the SDH network and recombined into a continuous bandwidth at the far end destination VCG receiver. The routing of the individual members can be in a physically diverse manner. Because of differing propagation delay associated with individual members of the VCG, a multiframing and sequencing approach has been developed to help identify the differential delay and the associated realignment process. Intermediate network elements (NEs) do not need to be path terminating and therefore do not need to be virtual concatenation aware.

Since individual paths can be transmitted, received, and recombined in different ways and in different increments, when compared with traditional contiguous concatenation, the strict sequential nature of members' virtual container and traditional large-bandwidth steps of contiguous concatenation bandwidth (for example, in SONET, the next step up from STS-3c/~155 Mbit/s is STS-12c/~622 Mbit/s) constraints can be alleviated with this method. Hence virtual concatenation requires concatenation functionality only at the source and destination SDH NE ends, whereas contiguous concatenation requires concatenation functionality at source, destination, and every intermediate SDH NE.

Internationally, HO VCG members can operate at VC-3 and VC-4 containers denoted by VC-3-Xv and VC-4-Xv where the range of X is from 1 to 256 and the associated lowercase v denotes virtual concatenation. For SONET, members operate at STS-1 or STS-3c containers denoted by STS-1-Xv and STS-3c-Xv. LO virtual concatenation has similar nomenclature.^{2,3}

The net result is significant bandwidth savings resulting from the flexible increments of bandwidth and the better size matching to native application data rates. Table 1 below illustrates the waste of bandwidth in contiguous concatenation when compared with virtual concatenation.

Table 1. Comparison of Contiguous versus Virtual Concatenation Stranded Bandwidth^a

Application	Application Native Rate	SONET Contiguous Concatenation/ Bandwidth Wasted	SONET Virtual Concatenation/ Bandwidth Wasted
Ethernet	10 Mbit/s	STS-1/80%	VT-1.5-7v/11%
Fast Ethernet	100 Mbit/s	STS-3c/33%	STS-1-2v/0%
Gigabit Ethernet	1 Gbit/s	STS-48c/58%	STS-3c-7v/5%
ESCON ^b	200 Mbit/s	STS-12c/67%	STS-1-4v/0%
Fibre Channel	1 Gbit/s	STS-48c/58%	STS-3c-7v/5%

^aRef. 10.

^bEnterprise Systems Connection.

A number of benefits are offered by virtual concatenation and can be summarized as follows:

- It enables much closer alignment of native application data rates to SDH bandwidth containers, even if the physically available SDH bandwidth is not SDH contiguous. This enables recovery of stranded bandwidth.

- It is a method to solve the virtual container fragmentation problem outlined above. This in turn could ease planning and traffic reengineering in the network yet still allows optimal use of physically available bandwidth resources.
- In applications requiring stringent QOS and guaranteed bandwidth transport, virtual concatenation offers strict bandwidth guarantee, separation, and flexible incremental bandwidth allocations to upper layer protocols.
- It imposes no new requirements on intermediate SDH NEs. Virtual concatenation support is required only at the source and destination points of the required service.
- It is applicable to many emerging broadband service environments, since the basic mechanism is protocol agnostic. It is specifically useful in networks with large SDH infrastructure, which is common in many carrier networks, thus allowing the service provider to leverage further its investment.
- The flexibility for best size of fit data- and voice-application traffic makes this a potentially useful network evolution tool to migrate from a TDM/voice-centric network toward a more convergence IP/Ethernet-oriented packet-based multiservice network.
- Interworking functions have been defined⁹ that enable NEs to send contiguous concatenation signals to and from virtual concatenation signals. This capability allows current contiguous concatenation equipment to take advantage of virtual concatenation capabilities.

Figure 2 and the associated bulleted description below provide an overview of the HO virtual concatenation mechanism.

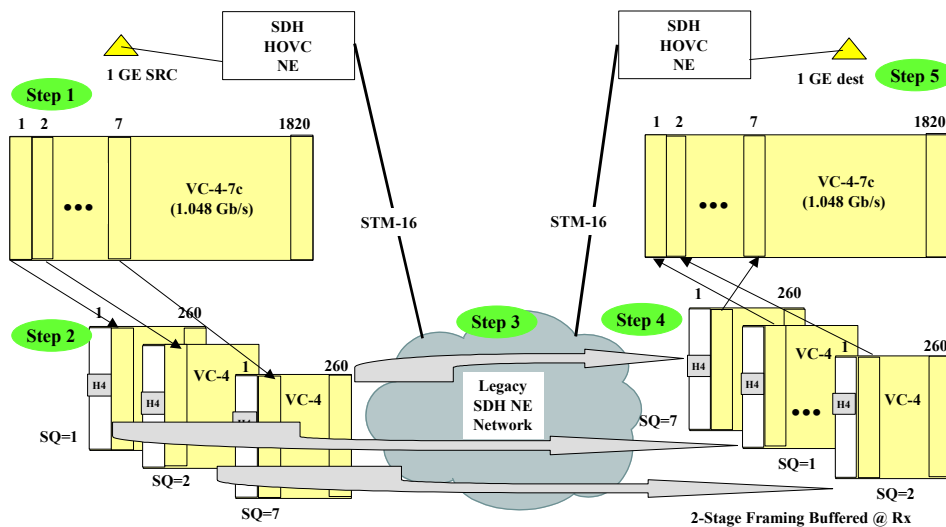


Fig. 2. High-order virtual concatenation Gigabit Ethernet (GE) transport example.

- SDH NEs at source and destination path-termination equipment (PTE) nodes perform HO virtual concatenation functions. Intermediate SDH NEs are transparent to the virtual concatenation operation.
- HO virtual concatenation is fundamentally a unidirectional operation. Source PTE maps the traffic to be transported into local memory that forms a continuous SDH signal (step 1). This is then allocated into SDH path (members of VCG) containers before the individual SDH paths are sent out the network side interface (step 2). These paths are then individually transported across the SDH network, possibly diversely routed (step 3). At the destination PTE, the individual SDH paths are received in buffered memory (step 4); validation and

differential delay realignment are performed and are then recombined back in local memory into a HO continuous payload (step 5). This process is repeated in the reverse direction for bidirectional transport.

- HO virtual concatenation transport is asymmetrical. That is, forward bandwidth can be different than reverse bandwidth. This asymmetry of bandwidth offers additional opportunity for increased efficiency.
- HO virtual concatenation employs two-stage multiframing. The first stage uses part of the H4 POH pointer bits 5–8 Multi-Frame Indicator 1 (MFI1) and cycles through 0–15 (16 basic frames). The second-stage multiframing Multi-Frame Indicator 2 (MFI2) uses bits 1–4 H4 byte frame 0 and frame 1 and cycles through 0–255. The result is a two-stage framing process that yields a total of 4096 frames or a 512-ms two-stage multiframe cycle.
- A sequence indicator (SQ), bits 1–4 of H4 byte in basic frame 14 and 15, identifies the order in which the members of VCG are split up. The sequence numbers are assigned by source node and interpreted by destination node. As many as 256 members are allowed in a single HO VCG.

Low-order virtual concatenation^{2,3} employs a similar mechanism to route lower-order virtual containers over the SDH infrastructure and reassemble the lower-order paths at the far end. These mechanisms use LO path overhead bytes V1 and V2 pointers and its associated overhead byte K4 (Z7 in SONET) bit 2 in a multiframe manner to derive a 32-bit control sequence for LO virtual concatenation.

Virtual concatenation is a relatively new method (proprietary variants have existed); therefore there are a couple of possible issues to consider that fall into the category of data protocols mapping into virtual concatenation and virtual concatenation signal protection. Since virtual concatenation has been devised primarily for data traffic, it is important that efficient mappings be available. Besides the traditional popular PPP over SDH (POS¹¹) and ATM mapping that is required for virtual concatenation, emerging protocol mappings include Ethernet (from 10 Mbit/s to 10 Gbit/s) and Fibre Channel protocols for various applications including campus and metro internetworking, content networking, and storage networking.

At present there is an Internet Engineering Task Force (IETF) draft¹⁰ that specifies the changes required for allowing POS mapping over virtual concatenation, and the ATM mapping over virtual concatenation is covered by the availability of ITU-T G.707³ and ANSI T1.PP.105.02.¹² However, there is no specific mapping currently outlined for Ethernet and Fibre Channel protocols, which are becoming popular in metro deployments and possibly will be used beyond the metro in the future. There is some relief in sight based on work in the ANSI T1X1.5 and ITU-T standards working groups looking at the Generic Framing Procedure (GFP¹³), which when available will provide mapping procedures for a host of protocols, including Ethernet and Fibre Channel. These mappings will be available over SDH and optical transport network (OTN) networks. However, proprietary mappings in these situations will likely prevail until standards are fully available, implemented, and deployed.

Protection and restoration of VCG are other potential areas for consideration. Current virtual concatenation standards do not specify protection at the VCG grouping. Virtual concatenation could rely on standard SDH protection and restoration at path and line layers of the network. This may be sufficient. However, given that the members of the VCG could be taking various paths through the SDH network, ensuring full protection coverage on all paths and links traversed could become complex if there is no coordination. LCAS, to be discussed in Section 5, could provide some remedy but at the expense of reduced VCG total bandwidth. This then leaves the option of vendor-specific solutions to provide coverage if full protection and restoration of this service is desired.

Virtual concatenation flexibility could be enhanced with LCAS, a mechanism recently coming off the standards track that allows the dynamic addition and removal of

SDH path capacity in a hitless fashion. In Section 5 we explore LCAS, what it is, how it works, what its benefits are, and some other issues.

5. Link Capacity Adjustment Scheme

The LCAS^{2,4,5} is an optional low-level capacity control mechanism that enables addition and removal of HO and LO SDH path(s) in a hitless fashion into and out of a VCG. This includes the removal of failed members within the VCG. The following is a summary of LCAS's basic capabilities:

- It is able to increase and decrease VCG capacity on increments of member (SDH path) container bandwidths.
- It is able to do this without affecting service.
- The capacity control is unidirectional. The forward LCAS VCG capacity can be different than that of the reverse direction, and both can change without coordination.
- Automatic removal of failed member links without taking down the VCG is supported and can dynamically add the member back when it recovers. Failed member removal is likely not hitless.
- LCAS allows interworking of LCAS VCG to non-LCAS VCG. That is, a LCAS transmitter can transfer to a non-LCAS receiver, and a non-LCAS transmitter can transfer with a LCAS receiver.
- Prior to the LCAS operation of incrementing or decrementing a member into or out of a VCG, it is assumed that the member to be operated on has been pre-provisioned prior with a management system or signaling system or a combination of both.

LCAS offers a number of benefits, which and can be summarized as follows:

- It offers flexibility to add and remove bandwidth capacity incrementally within a VCG without affecting service or taking down the entire VCG service. Therefore in-service resizing of VCG service capacity is possible.
- It offers new service granularities and run-time control of this bandwidth when virtual concatenation is combined with LCAS. This could be employed as an internal service as well.
- The load-sharing operation of LCAS can provide reduced-bandwidth service when a member fails without affecting the total service. Therefore one can choose to explore this type of restoration scheme in addition to the more traditional robust SDH protection and restoration methods.
- The above inherent load-sharing-type restoration scheme can potentially be a component of a service that when combined with packet-level prioritization and congestion-avoidance schemes produce new types of enhanced service offerings.
- The basic capability can be used in combination with differing trigger mechanisms including management invocations, signaling protocol invocations, and on-board NE application schemes with variety of possible trigger algorithms and methods. This is assuming that member paths are set up across the SDH network prior to invocation of LCAS operation.
- It should be noted that if the mechanism is to be deployed across service provider to customer boundaries, prior agreements on behavior, billing, security, and other policies may require further consideration.

LCAS operation requires the VCG transmitter end and the VCG receiver end to maintain states on each member of the group and to exchange control messaging to manage the addition and removal of member links within the group. A set of parameters is necessary for enabling this control messaging protocol. In the case of HO virtual concatenation, the

LCAS control messaging channel is via the H4 path overhead byte (the same byte is used for basic virtual concatenation operation), and for LO virtual concatenation, which uses the K4 overhead byte bit 2, a 32-bit multiframe control sequence is employed. The following provides a summary of the parameters used in virtual concatenation and LCAS control messaging:

- *Frame indicator*. The mechanism employed between the VCG transmitter and VCG receiver to determine the differential delay and use for realignment between members in the same VCG. The value should be the same for all members of the VCG. For HO virtual concatenation, the range is 0–4095, and for LO virtual concatenation the range is 0–31.
- *Sequence indicator (SQ)*. The number assigned to members within the VCG where the number is unique within one VCG and sequential (corresponding to reassembly order of the virtual concatenation). The SQ range for HO virtual concatenation is 0–255 (256 maximum members within one VCG) and 0–63 for LO virtual concatenation (64 maximum members within one VCG).
- *Control (CTRL)*. LCAS control command from transmitter to receiver VCG:
 - FIXED Indication of fixed-bandwidth operation, non-LCAS mode.
 - ADD Identified member is about to be added to the group.
 - NORM Normal transmission, steady-state no-change mode.
 - EOS End of sequence and normal transmission mode.
 - IDLE Member is not part of or about to be removed from this group.
 - DNU Do not use this member of VCG. The receiver has detected a failure.
- *Group identifier (GID)*. The group identifier for the VCG. All members of the same VCG have the same GID value. This is to ensure that all member signals come from same intended transmitter.
- *Member status (MST)*. A summary status report of all members of a VCG (OK or failed) from the receiver back to the transmitter. A multiframe approach is used to send the full report of all members in the VCG.
- *Resequencing acknowledge (RS-Ack)*. Indication from receiver to transmitter that the changes initiated from the transmitter have been accepted and that the transmitter can begin accepting the new member status information.
- *Cyclic redundancy check (CRC)*. Validation check to protect the integrity of each virtual concatenation control message. If the CRC check fails, the virtual concatenation overhead contents are not used.

LCAS is an optional capability that can be deployed with virtual concatenation. Its specifications have recently^{2,5} become available, and clarification of the original specification has also just been updated.⁴ As such, deployment experience of this real-time capacity-control protocol technology, to the extent that there is any, is limited. Multivendor interoperability would certainly take time. Beyond the issue of potential deployment experience is how to ensure that member links are indeed diversely physically routed over a SDH network. If service definition requires more graceful reduction of bandwidth when member links fail, then it would be necessary to ensure that such member links are traversed over physically diverse paths. Clarification of how this could be accomplished within the framework of the standard would seem relevant to the successful deployment of this technology.

We have just completed a brief technology review of virtual concatenation and LCAS. To complement this, a few situations are offered in Section 6 to explore the potential applications of this technology.

6. Application Situations

6.A. Metro Provider Leveraging Virtual Concatenation in OC-48 Leased Line

A metro service provider, operating in multiple metro areas denoted in Fig. 3 below as Metro SP-A, leases OC-48 (up to 48 STS-1s) bandwidth from a national backbone transport carrier, denoted by Carrier-B, to interconnect its metros across the country. Metro SP-A provides voice, video, and data services (OC is optical carrier). Since the Metro SP-A leases bandwidth from a backbone carrier, it would want to maximize the bandwidth used in the intercity routes and where possible gain better control of its usage. With virtual concatenation and LCAS, the Metro Service provider is able to mix voice, video, and IP/Ethernet data traffic over the same OC-48 lease facility, via one or more VCGs transported over the intercity routes in noncontiguous nonstandard size traffic containers. This optimizes use of the leased facilities and in some instances defers the need for additional facilities.

As Metro SP-A grows its subscriber base, especially in the data/IP arena, it will require more bandwidth. If there is excess bandwidth within the voice/TDM component, the metro service provider can move the virtual concatenation partition, either through management provisioning or in combination with LCAS capability to satisfy the demand. Based on preprovisioned SDH paths, LCAS can help Metro SP-A perform the necessary increase of data services bandwidth without service disruption and maintain transparency in the backbone provider. Metro SP-A can now more flexibly modify its capacity partitioning between TDM and data services and more easily adjust to traffic pattern changes between inter city routes. Virtual concatenation and LCAS enables this flexible implementation without lengthy intercarrier provisioning, procedures, and possibly truck roll. Deployment of virtual concatenation equipment is required only at points of interface between the metro provider and the backbone carrier provider. It should be noted that this does not necessarily reduce the number of OC-48 leased facilities but rather increases use of these facilities, enabling greater flexibility and control of this capacity.

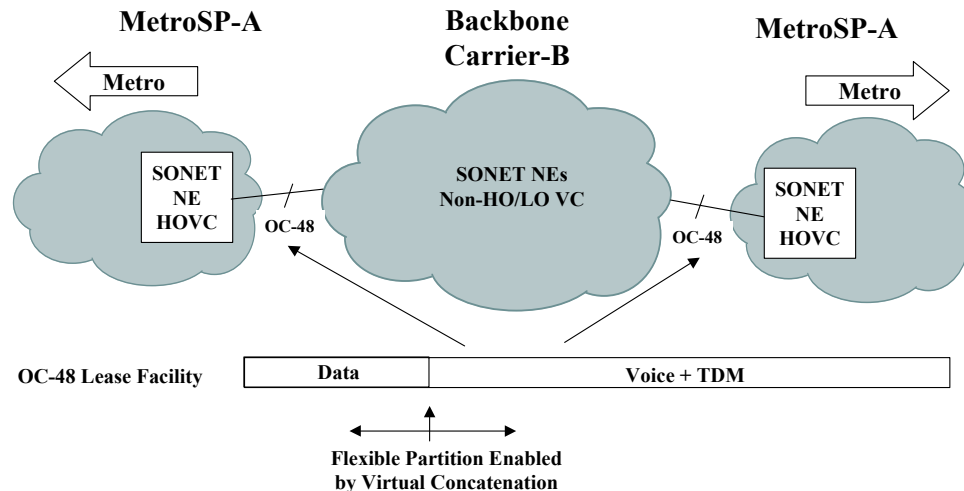


Fig. 3. Metro service provider subscribing OC-48 service.

6.B. New Ethernet Service for an Integrated End-to-End Metro and Core Provider

An end-to-end metro and core service provider, shown in Fig. 4 below as Service Provider A, desires to improve use of its fragmented (VC-4/~150 Mbit/s) SDH network to deliver end-to-end high-speed optical Ethernet services in addition to its present ones

based on voice, TDM, and leased line service offerings. Because of the fiscal and capital constraint environment that the provider is now experiencing, a low initial investment approach is desired.

Figure 4 illustrates that this service provider has used up much of its bandwidth for voice and TDM services denoted by the shaded VC-4 blocks. However, this service provider also has physically available but fragmented bandwidths, denoted by white VC-4 blocks, that it has not currently leveraged. With virtual concatenation, the service provider can group these VC-4 blocks together into VCGs and deploy them as bandwidth pools for a new high-speed Ethernet service. This is illustrated in the figure with the end-to-end GE arrow consuming the white block VC-4 bandwidth, which without virtual concatenation would be left stranded and unused. The service provider would have to deploy virtual-concatenation-capable SDH equipment or other data-networking equipment capable of virtual concatenation adaptation. However, the overall investment is reduced, since virtual concatenation equipment would need to be deployed only incrementally at the edges of the metro networks in initial GE service areas, and, on the basis of revenue, return on investment, and other parameters, additional investment can then be made from a business perspective.

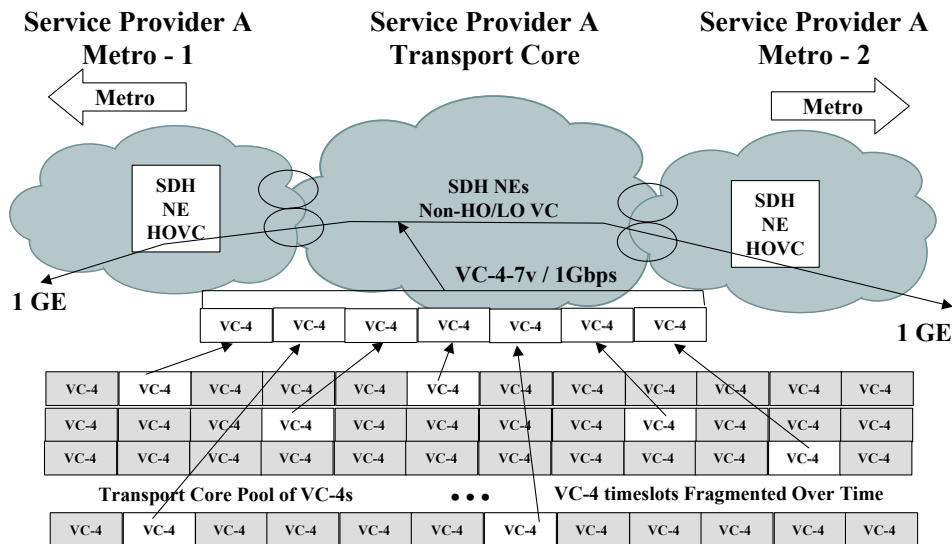


Fig. 4. Integrated end-to-end metro and core transport provider.

6.C. Networks Enabled with Virtual Concatenation, Link Capacity Adjustment Schemes, and Generalized Multi-Protocol Label Switching

Use of virtual concatenation to combine multiple noncontiguous individual transport containers to higher aggregate rates at much more gradual and granular increments for transport may have an interesting application when combined with generalized multiprotocol label switching (GMPLS^{14,15}) over SDH and OTNs. Specifically, Fig. 5 and the description below provides an illustration with the Optical Internetworking Forum (OIF) Optical User Network Interface (O-UNI) in combination with virtual concatenation and LCAS in an overlay deployment model. O-UNI is an implementation agreement currently being developed by the OIF to enable overlay network implementation within the GMPLS architectural framework.

Client edge devices, whether they be SDH NEs, routers, or ATM switches, can use O-UNI as a mechanism for dynamic signaling of bandwidth requests into the optical network, one SDH path at a time as necessary. When the bandwidth is confirmed (with a signaling protocol response), that SDH path bandwidth could be incorporated into a VCG as defined by the source and destination client devices. This VCG assignment is

transparent to the optical network servicing the O-UNI bandwidth request. The source and destination client devices have the role of virtual concatenation application adaptation for various traffic types: data, voice, and video. The following figure illustrates the concept where the blue functional block is an example of where virtual concatenation adaptation or interworking conversion can occur. Other points of virtual concatenation adaptation and conversion are possible.

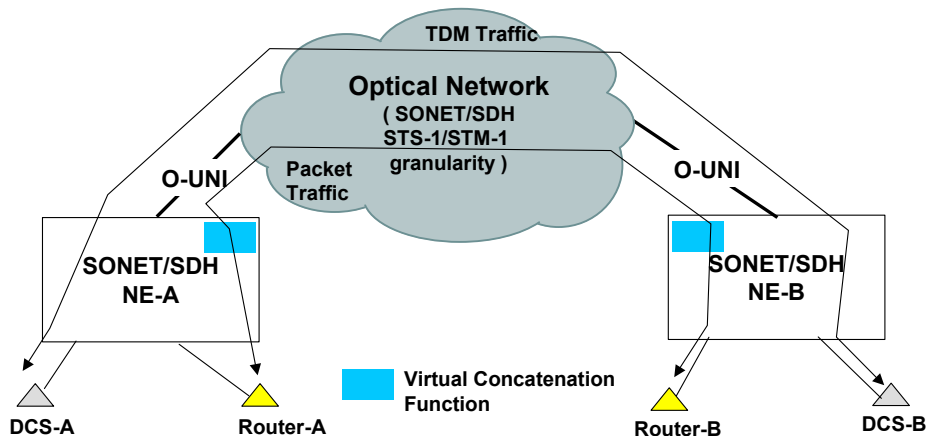


Fig. 5. Applications enabled by virtual concatenation, LCAS, and O-UNI.

In this deployment the service provider could leverage O-UNI to signal data bandwidth as necessary at individual SDH path increments while sharing the same access pipe with TDM traffic from the same office. The virtual concatenation functions in NE-A and NE-B would enable best-size virtual concatenation bandwidth across the optical network. Once the O-UNI bandwidth is set up, if LCAS is available on NE-A and NE-B, the bandwidth could be added in a hitless fashion within the control of the subscriber (NE-A and NE-B) independent of the optical network provider. This offers greater flexibility to the subscriber organization to allocate, consolidate, or partition its own traffic and traffic mix as necessary, at any time, without bilateral coordination with the optical network provider. When O-UNI is deployed in combination with LCAS and virtual concatenation, hitless modification of SDH HO virtual concatenation bandwidth is possible and is transparent to the serving optical network with control under the administration of the subscriber organization.

Additional applications are possible by means of embedding the virtual concatenation, interworking, and LCAS capabilities inside the optical network. This is particularly interesting when these capabilities are combined with the efforts in progress with GMPLS signaling extensions¹⁵ to include virtual concatenation and dynamic bandwidth adjustments, thus enabling a variety of deployment models ranging from overlay to peer-to-peer and other combinations in between.

7. Conclusion

SDH networks are mature and have contributed much to the telecommunications and networking industry. Given the rise of Internet and IP/data-oriented traffic, there is room for improvement. Virtual concatenation and LCAS are some of the mechanisms emerging from the standards track that offer some relief. They are also tools to help manage the migration toward a more convergent packet-oriented multiservice network. In this paper a number of contiguous concatenation issues have been reviewed along with other techniques that improve data transport. It was illustrated that these techniques, although useful, do not fundamentally address the issues directly at the SDH layer. This is then

complemented with a discussion on the definitions, mechanisms, benefits, and other issues involving virtual concatenation and LCAS. Finally, a number of situations for virtual concatenation were reviewed to illustrate potential application deployments.

References and Links

1. Telecordia Technologies, "SONET transport systems: common criteria," GR-253-CORE Issue 3, September 2000, <http://telecom-info.telcordia.com/site-cgi/ido/index.html>.
2. S. Gorshe, "Revised draft T105 SONET base standard," ANSI T1X1.5/2001-062, January 2001, <http://www.t1.org/t1x1/x15-hm.htm>.
3. ITU-T, "Network node interface for synchronous digital hierarchy (SDH)," G.707/Y.1322 (prepublication document) (International Telecommunication Union, October 2000), <http://www.itu.int/publications/online/index.html>.
4. H. V. Helvoort, S. Gorshe, and T. Wilson, "Removal of LCAS ambiguities from T1.105," ANSI T1X1.5/2001-218, September 2001, <http://www.t1.org/t1x1/x15-hm.htm>.
5. ITU-T, "New recommendation G.7042/Y.1305 (ex. G.lcas) (for consent)," G.7042/Y.1305 SG 15, Question 11/15 (International Telecommunication Union, October 2001), <http://www.itu.int/publications/online/index.html>.
6. ATM Forum, Inverse Multiplexing for ATM (IMA) version 1.1 af-phy-0086.001, March 1999, <http://www.atmforum.com/techspecfs1.html>.
7. K. Sklower, B. Lloyd, G. McGregor, D. Carr, and T. Coradetti, "The PPP MultiLink Protocol (MP)," IETF RFC1990, August 1996, <http://www.faqs.org/rfcs/rfc-titles.html>.
8. C. Metz, "IP protection and restoration," IEEE Internet Comput. (March–April 2000), pp. 97–102, <http://computer.org/Internet/>.
9. ITU-T, "Characteristics of synchronous digital hierarchy (SDH) equipment functional blocks," G.783 (prepublication document) (International Telecommunication Union, October 2000), <http://www.itu.int/publications/online/index.html>.
10. N. Jones and C. Murton, "Extending PPP over SONET/SDH with virtual concatenation, high order and low order payloads," IETF draft draft-ietf-pppext-posvcholo-04.txt (Internet Engineering Task Force, June 2001), <http://search.ietf.org/search/brokers/internet-drafts/query.html>.
11. A. Malis and W. Simpson, "PPP over SONET/SDH," IETF RFC 2615 (Internet Engineering Task Force, June 1999), <http://www.faqs.org/rfcs/rfc-titles.html>.
12. ANSI, "Synchronous optical network (SONET) payload mappings," T1.PP.105.02-2001 (prepublished document), September 2000, <https://www.atis.org/atis/docstore/searchform.asp>.
13. S. Gorshe, ed., "Current draft of G.fgp correspondence document," ANSI T1X1.5/2001-186, September 2001, <http://www.t1.org/t1x1/x15-hm.htm>.
14. E. Mannie, ed., "Generalized multi-protocol label switching (GMPLS) architecture," IETF Draft draft-ietf-ccamp-gmpls-architecture-00.txt (Internet Engineering Task Force, June 2001), <http://search.ietf.org/search/brokers/internet-drafts/query.html>.
15. E. Mannie, ed., "GMPLS extensions for SONET and SDH control," IETF Draft draft-ietf-ccamp-gmpls-sonet-sdh-02.txt (Internet Engineering Task Force, October 2001), <http://search.ietf.org/search/brokers/internet-drafts/query.html>.