(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
28 February 2002 (28.02.2002)

PCT

(10) International Publication Number
WO 02/17546 A2

(51) International Patent Classification[7]: H04L

(21) International Application Number: PCT/US01/26557

(22) International Filing Date: 23 August 2001 (23.08.2001)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
60/228,008 23 August 2000 (23.08.2000) US
60/272,793 1 March 2001 (01.03.2001) US

(71) Applicant: GEYSER NETWORKS, INC. [US/US]; 555 Del Rey Avenue, Sunnyvale, CA 94085 (US).

(71) Applicants and
(72) Inventors: LEE, Gordon [US/US]; 20567 Russell Lane, Saratoga, CA 95070 (US). HUANG, Kevin [US/US]; 4656 Clarendon Drive, San Jose, CA 95129 (US). CHEN, Wen-Lung [US/US]; 10455 Creston Drive, Los Altos, CA 94024 (US).

(74) Agents: GUZMAN, Leonard, T. et al.; McDermott, Will & Emery, 2700 Sand Hill Road, Menlo Park, CA 94025 (US).

(81) Designated States *(national)*: AU, BR, CN, DE, DK, ES, GB, IN, JP, KR, MX, PT, SG.

(84) Designated States *(regional)*: European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR).

Published:
— *without international search report and to be republished upon receipt of that report*

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: SYSTEM AND METHOD OF VIRTUALLY CONCATENATING VT1.5s ANS STS-1s OVER SONET AND SDH AND WDM

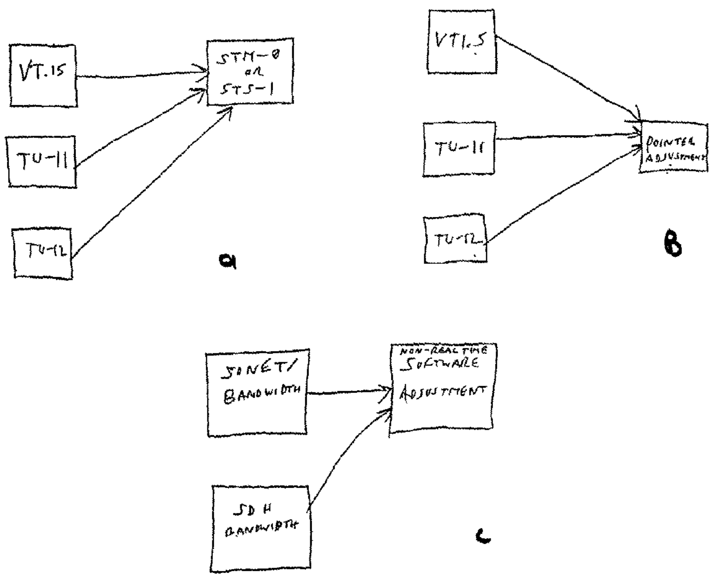

(57) Abstract: A system and method of virtually concatenating VT1.5s and STS-1s over SONET and SDH and WDM is provided. In an exemplary embodiment, the system and method includes (1) creating a nxSTS-1 virtually concatenated pipe out of STS-1s, (2) within the nxSTS-1 pipe, forming one or more nxVT pipes out of VTs, and (3) mapping an application to a nxVT pipe or a nxSTS-1 pipe based upon a service level agreement. In the system and method, the creating includes (a) terminating SONET framing, (b) processing pointers, and (c) compensating for delay, thereby aligning the STS-1s in the nxSTS-1 virtually concatenated pipe. In the system and method, the forming includes (a) processing VT overhead of the VTs and (b) mapping traffic to nxVT pipes or nxSTS-1 pipes according to a connection identification.

WO 02/17546 A2

SYSTEM AND METHOD OF VIRTUALLY CONCATENATING VT1.5s AND STS-1s OVER SONET AND SDH AND WDM

SPECIFICATION

RELATED APPLICATIONS

This application is related to U.S. Provisional Application No. 60/228,008, filed on August 23, 2000, to U.S. Provisional Application No. 60/272,793 , filed on March 1, 2001, and to co-pending and commonly assigned U.S. Patent Application No. (Number to be assigned) with Attorney Docket Number 55369-015, filed on August 23, 2001. The contents of U.S. Provisional Application No. 60/228,008, filed on August 23, 2000, of U.S. Provisional Application No. 60/272,793, filed on March 1, 2001, and of co-pending and commonly assigned U.S. Patent Application No. (Number to be assigned) with Attorney Docket Number 55369-015, filed on August 23, 2001, are hereby incorporated by reference. This application claims priority to U.S. Provisional Application No. 60/228,008, filed on August 23, 2000, and to U.S. Provisional Application No. 60/272,793, filed on March 1, 2001.

FIELD OF THE INVENTION

The present invention relates to optical networks. More particularly, the invention relates to a system and method of virtually concatenating VT1.5s and STS-1s over SONET and SDH and WDM.

BACKGROUND OF THE INVENTION

Both SONET (See Synchronous Optical Network (SONET) Transport Systems: Common Generic Criteria. GR-253-CORE, Issue 2, Revision 1. December, 1997.) and SDH (See International Telecommunication Union. Network Node Interface for the Synchronous Digital Hierarchy. Recommendation G.707. March, 1996.) enable the use of virtual concatenation to support both the dynamic resizing of transport trunks and the grooming of traffic. More recently, advances in the transport of routed datagram traffic leveraging the research and experience of ATM has resulted in the standardization of MPLS (See Internet Engineering Task Force. Multiprotocol Label Switching Architecture. IETF Draft Document. August, 1999 and http://www.ietf.org/internet-drafts/draft-ietf-mpls-arch-06.txt.). This work allows network devices to employ a standards-based method by which packet traffic can traverse a network, while receiving a previously agreed upon Quality of Service.

Referring to prior art Figure 1A, in standard SONET and SDH, a VT1.5 and a TU-11 or TU-12 can only be concatenated into a STS-1 and a STM-0 or a STM-1, respectively, with a minimum granularity of 51 megabits per second (Mbps).

Referring to prior art Figure 1B, in standard SONET and SDH, VT1.5s and TU-11s or TU-12s require pointer adjustment, which is costly mechanically complex, and difficult to implement.

Referring to prior art Figure 1C, in standard SONET and SDH, bandwidth of connections can only be adjusted via software, in a non-real time manner.

In addition, standard SONET's GR-253 allows for 4 frames per multiple-frame. In standard SONET, the GR-253 VT1.5 virtual tributary structure uses bytes interleaved with 3 separate columns for one T1 mapping.

SUMMARY OF THE INVENTION

The present invention provides a system and method of virtually concatenating VT1.5s and STS-1s over SONET and SDH and WDM. In an exemplary embodiment, the system and method includes (1) creating a nxSTS-1 virtually concatenated pipe out of STS-1s, (2) within the nxSTS-1 pipe, forming one or more nxVT pipes out of VTs, and (3) mapping an application to a nxVT pipe or a nxSTS-1 pipe based upon a service level agreement. In the system and method, the creating includes (a) terminating SONET framing, (b) processing pointers, and (c) compensating for delay, thereby aligning the STS-1s in the nxSTS-1 virtually concatenated pipe. In the system and method, the forming includes (a) processing VT overhead of the VTs and (b) mapping traffic to nxVT pipes or nxSTS-1 pipes according to a connection identification. In the system and method, the mapping includes (a) processing traffic and (b) mapping traffic to a nxVT pipe or a nxSTS-1 pipe based upon the service level agreement.

The present invention provides for the virtual concatenation of a VT1.5 and a TU-11 or TU-12 a granularity of a VT1.5 and a TU-11 or TU-12, respectively, does not require pointer adjustment, and supports real-time dynamic bandwidth allocation.

BRIEF DESCRIPTION OF THE DRAWINGS

Figures 1A, 1B, and 1C illustrate prior art details of standard SONET and SDH.

Figure 2 illustrates nxSTS-1 and nxVT virtual concatenation in accordance with an exemplary embodiment of the present invention.

Figure 3 illustrates a SONET STS-1 path ring in accordance with an exemplary embodiment of the present invention.

Figure 4 illustrates a nxSTS-1 system with STS-1s going through different routes in a public network in accordance with an exemplary embodiment of the present invention.

Figure 5 illustrates a STS-1 superframe or framing structure in accordance with an exemplary embodiment of the present invention.

DETAILED DESCRIPTION OF THE INVENTION

Introduction

The present invention provides a system and method of virtually concatenating VT1.5s and STS-1s over SONET and SDH and WDM. The present invention allows users to setup connections or pipes with configurable bandwidth over either nxSTS-1/nxAU-3/nxAU-4 or nxVT1.5/nxTU-11/nxTU-12 within a nxSTS-1/nxAU-3/nxAU-4 pipe on an existing SONET/SDH network. This provides a connection or pipe of adjustable bandwidth with a granularity of close to 1.5 Mbps to fit the needs of applications. The resulting connection can be treated as a TDM like connection.

By replacing "STS-1" with "AU-3" or "AU-4" and "VT" or "VT1.5" with "TU-11" or "TU-12", the present invention applies to nxAU-3/nxAU-4 and nxTU-11/nxTU-12 for SDH networks. For simplicity, these connections are called "nxVT" for both SONET and SDH networks. By replacing "STS-1" with "VT" or "VT1.5", the present invention applies to nxSTS-1 and nxAU-3/nxAU-4.

On top of the present invention, a dynamic bandwidth allocation (DBA) protocol, which is described in co-pending and commonly assigned U.S. Patent Application No. (Number to be assigned) with Attorney Docket Number 55369-015, allows for dynamically changing the throughput of all nxVT connections, based on the real-time traffic loads of applications using the nxVT connections. The DBA protocol allows for the efficient use of the SONET/SDH bandwidth through statistical multiplexing. The same dynamic bandwidth allocation protocol applies to nxSTS-1 and nxAU-3/nxAU-4.

The present invention provides for virtual concatenation, which includes creating a logical connection or pipe by combining multiple, n (where n is a positive integer), STS-1 or VT connections or pipes, which may be contiguous or non-contiguous, into a single connection or pipe, nxSTS-1 or nxVT, respectively, in order to support a connection or pipe with a higher throughput than the throughput of the original STS-1 or VT pipes.

Geyser nxSTS-1 and nxVT virtual concatenation

Referring to Figure 2, the present invention provides for the following two kinds of "virtual concatenated" connections: (1) a nxSTS-1 210, 220, 230; and (2) a nxVT 214, 216, 218.

In an exemplary embodiment of the present invention, nxSTS-1 210, 220, 230 is formed from (a) multiple STS-1s residing in a single or multiple OC-n (OC-3, 12, 48, 192, etc.) interface(s) or (b) multiple λ in a DWDM system or (c) multiple fiber cable. The STS-1s may be scattered over multiple wavelengths or interfaces or may reside in the same fiber or different fibers. The STS-1s may be randomly picked from the OC-n interfaces or λ or fiber to form nxSTS-1 pipe 210, 220, 230. The STS-1s may be contiguous or non-contiguous. Multiple nxSTS-1s 210, 220, 230 can be setup in the same SONET or SDH system.

In an exemplary embodiment of the present invention, nxVT 214, 216, 218 is formed with multiple VTs 212 in a single nxSTS-1 pipe 210. Again, the VTs 212 do not need to be contiguous. Any sets of VTs inside the nxSTS-1 pipe 210 may be used to form an nxVT connection 214, 216, 218. nxVT 214, 216, 218 may be formed across STS-1s within the same nxSTS-1 210. Multiple nxVTs 214, 216, 218 may be setup in one nxSTS-1 connection 210. In the present invention, multiple nxSTS-1s 210, 220, 230 may be formed, where each nxSTS-1 210, 220, 230 may support one or more nxVT connections 214, 216, 218 within the nxSTS-1. In an exemplary embodiment of the present invention, each nxVT 214, 216, 218 cannot be formed across multiple nxSTS-1s 210, 220, 230. In an exemplary embodiment of the present invention, a STS-1 is logically broken up with a granularity of one VT. In an exemplary embodiment of the present invention, a connection may be made with, at a minimum, one VT. In an exemplary embodiment of the present invention, a connection may be made with one STS-1.

In an exemplary embodiment of the present invention, nxSTS-1 210, 220, 230 (a) may be used as a single connection or pipe for an application 222, 232 or (b) may consist of many nxVTs 214, 216, 218 where each nxVT 214, 216, 218 may be used as a connection or pipe for an application 215, 217, 219, respectively. The resulting connection may be a TDM, Packet based, or ATM based connection. The number, n, of VTs 212 or STS-1s used for a connection is selected to fit the bandwidth required by application using the connection. For example, a 10 Mbps Ethernet application may need a 6xVT connection to accommodate its throughput, while a MPEG application may need a17xVT connection. The present invention places no limit on the number, n, of STS-1s that may be used to form an nxSTS-1 pipe 210, 220, 230. The number, n, of

VTs 212 used to form nxVT connection 214, 216, 218 is limited by nxSTS-1 210.

In a packet or ATM based connection, the bandwidth usage depends on the actual traffic pattern. Usually the data packet traffic can be very bursty. Assigning a fixed value of n for nxVT connection 214, 216, 218 may not be optimal. A more optimal solution is to dynamically adjust the value of n, and therefore the bandwidth of the connection, based on the real traffic load running in the pipe.

The Timing and Framing Path Ring Closing

Referring next to Figure 3, a SONET STS-1 path ring 300 in accordance with an exemplary embodiment of the present invention is shown. In the present invention, nxSTS-1 210, 220, 230 and nxVT 214, 216, 218 are based on the SONET path layer running as a ring 300. Figure 3 shows the nxSTS-1 ring implementation 300. Path ring 300 includes nodes 310, 320, 330, and 340. In Figure 3, all nodes 310, 320, 330, and 340 use a common synchronized timing, either bits clock or recovered clock, in order to avoid delay compensation buffer overflow or underflow.

Nodes 310, 320, 330, and 340 include pointer processing (PP) buffers 314, 315, 324, 325, 334, 335, and 344, 345, respectively. Nodes 310, 320, 330, and 340 also include delay compensators (DC) 312, 322, 332, and 342, respectively. In an exemplary embodiment, node 310 is a SONET framing termination node.

In an exemplary embodiment, node 320 is a delay compensation node performing J1 alignment. Node 320 makes J1 aligned to compensate the J1 delay among STS-1s caused by PP XMT and RCV buffer (i.e. pointer adjustment). Node 320 can freely generate SJ1/J1/H4.

In an exemplary embodiment, node 330 is a delay compensation node performing SJ and H4 termination. Node 330 writes from PP buffer 334 to a SJ1 byte of DC 332 according to a received SPE. Node 330 also includes (1) a free running SJ1/J1/SPE module 338, (2) a SJ node internal system clock 339, which inputs to free running SJ1/J1/SPE module 338, and (3) a H4 coding module 336 . Node 330 reads out from the SJ1 byte of DC 332 using a freely generated SPE from free running SJ1/J1/SPE module 338.

In an exemplary embodiment, node 340 is a delay compensation node performing multiple frame alignment. DC 342 includes a J1 byte with a maximum delay. In an exemplary embodiment, node 340 receives STS-1s from a first network input 346 and STS-1s from a second network input 348.

In an exemplary embodiment, in path ring 300, only one node includes system clock 339. In an exemplary embodiment, in path ring 300, only one node includes free running SJ1/J1/SPE module 338, thereby minimizing delay. In an alternative embodiment, in path ring 300, at least one node may include a free running SJ1/J1/SPE module. In an exemplary embodiment, in path ring 300, only one node includes H4 coding module 336. In an alternative embodiment, in path ring 300, at least one nodes may include a H4 coding module.

In an exemplary embodiment, every node 310, 320, 330, 340 maps to or de-maps from path layer framing structure 300. In order for the SONET network to work as a ring topology 300, at least one of the nodes, e.g. node 330, acts as a master node to provide the timing information for everyone else. If the master timing were not available, every node would try to follow the timing received from its previous node. Since it is a ring topology, this would lead to an infinite loop, and the timing would become unstable. The present invention avoids this kind of dead loop of framing and timing by having a master node 330. In ring topology 300, master node 330 generates the following timings: (a) super frame pulse SJ1/J1, where J1 marks the 1st byte of SONET-SPE, via free running SJ1/J1/SPE module 338; and (b) the Frame sequence number H4, via H4 coding module 336.

Super Frame Pulse – SJ1/J1

In an exemplary embodiment, nxSTS-1 210, 220, 230 has an 8 SONET frames based super frame structure. In present invention, one super frame consists of 8 SONET frames. The start of a super frame is represented by a SJ1 pulse. SJ1 is the position of the first J1 in one super frame. In an exemplary embodiment, nxSTS-1 210, 220, 230 could have a n x SONET frames based super frame structure, where n is positive integer (n=1,2,3,……).

In an exemplary embodiment, master node 330 generates a SJ1 pulse and the J1 pulse freely via free running SJ1/J1/SPE module 338 and with input from free running clock 339. Once this timing is generated, all other nodes 310, 320, and 340 receive SJ1/J1 from its previous node and perform pointer adjustment based on the incoming SJ1/J1. Eventually this SJ1 pulse is propagated across the whole ring 300 and sent back to master node 330.

In an exemplary embodiment, due to the delay/pointer processing in each node, master node 330 receives an incoming SJ1 different from the freely generated SJ1 from free running SJ1/J1/SPE module 338 at the transmitting side. Master node 330 compensates for the difference between incoming SJ1 and the freely generated SJ1 by buffering all the data between the two timings. In this way, master node 330 compensates and ensures that the payload delay across ring 300 is a multiple of super-frame delay, thereby resulting in stable timing. In an exemplary embodiment, for a shared nxSTS-1 pipe 210, 220, 230, ring 300 is stable only when the total ring delay is equal to the time of n x super frame, where n is a positive integer.

In an exemplary embodiment, the generation of J1 and SJ1 may be performed in different nodes. In a preferred embodiment, in order to avoid increasing the complexity of protection, the generation of J1 and SJ1 is performed in the same node 330. In an exemplary embodiment, SJ1 generation may cover J1 generation, and therefore, only SJ1 generation may needed in STS-1 Path ring 300.

In order for master node 330 to compensate for the delay between the received SJ1 and the freely generated SJ1, the timing, or clock, of all nodes needs to be synchronized. Non-synchronized timing would eventually cause buffer overflow or underflow since the SJ1 framing processed by all the nodes using different timing could cause the pointer adjustment to further deviate from the original SJ1. Eventually, this would cause buffer overflow underflow. Therefore, in an exemplary embodiment, in path ring 300, only one node includes system clock 339, which provides the synchronized timing or clock for all nodes 310, 320, 330, and 340 in ring 300.

Frame Sequence Number H4

In an exemplary embodiment, a sequence number of 8 bits is put at the STS-1 H4 byte overhead, which serves two purposes. First, the SJ1 may be embedded in the H4 byte. The SJ1 is the first J1 of the super-frame. The super-frame will carry 8 SONET frames with H4 sequence number 0, 1, 2, 3, …, 255. Th SJ1 pulse may be specified to be 1 whenever H4 equals 8n, where n = 0, 1, 2, 3,.., etc. In order to minimize the ring delay in DBA implementation, master node 330 re-generating J1/SJ1 will generate the H4 code also, via H4 coding module 336.

Secondly, the H4 sequence number is used to specify the delay of each STS-1 when they go through different routes. The H4 byte is used to specify an 8-bit sequence number, which can differentiate up to 256 SONET frames, or equivalently 32 msec. This allows for a delay compensation of up to 16 msec.

In an exemplary embodiment, other H4 coding method may be used instead of a flat coding.

Delay Compensation

Figure 4 shows a nxSTS-1 system 400 with STS-1s going through different routes in a public network in accordance with an exemplary embodiment of the present invention. System 400 includes nodes 420 and 430, routers 440 and 442, and public networks 410 and 450, logically interconnected as shown in Figure 4. Nodes 420 and 430 are examples of node 340 from Figure 3. Public networks 410 and 450 may be a SONET or SDH or DWDM network. In an exemplary embodiment, node 430 combines network inputs 346 and 348.

In an exemplary embodiment, the virtual concatenated STS-1s' SJ1/J1s are aligned at every node 310, 320, 330, 340, 420, 430 in order for packet to map or de-map correctly. Referring to Figure 4, the present invention provides for the following three kinds of delay compensation: (1) pointer processor buffer delay compensation; (2) super frame termination delay compensation; and (3) configured maximum SONET framing delay compensation.

Pointer Processor Buffer Delay Compensation

In an exemplary embodiment, referring to Figure 3, in path ring 300, a downstream node (e.g. node 320) only needs to compensate for the J1 delay variation caused by the transmitter pointer processor (e.g. PP 314) from an upstream node (e.g. node 310) and the receiver pointer processor (e.g. PP 324) at the downstream node (e.g, node 320). Each of the nodes 310, 320, 330, 340 only introduces minimal additional latency to compensate the pointer adjustment delay, which is in the worst of 2 times of the Pointer_Processor_buffer_size delay. In an exemplary embodiment, only one master node 330 in ring 300 exists and compensates the whole ring delay to be a multiple of super-frame delay.

Super-Frame Termination Delay Compensation

In an exemplary embodiment, referring to Figure 3, master node 330 needs to keep a stable and synchronized SJ1 pulse. Master node 330 attempts to buffer all of the data in a nxSTS-1 payload such that the relative difference between the received SJ1 pulse and the locally generated SJ1 pulse from module 338 remains fixed. This is equivalent to doing a delay compensation such that the whole ring delay is exactly an integer multiple of the super-frame delay. In this way, the received data in the nxSTS-1 payload can be safely passed to the next node without buffer overflow or under flow.

Configured Maximum SONET Framing Delay Compensation

Referring to Figure 4, in an exemplary embodiment, some STS-1s may go through a different route or fiber or λ in a network from other STS-1s. The delay of each STS-1 route can be different and each STS-1 can be floating in OC-n (where n is an integer) signal or WDM networks. This is especially true when the STS-1s go through different public networks 410 and 450 as shown in Figure 4.

The delay variations among all the STS-1s are much higher when the traffic is going through public networks 410 and 450. The H4 byte is used to carry an 8-bit sequence number. The sequence number will cover up to 256 SONET frames. It allows the node to compensate for a delay of plus or minus 16 msec. In order to cover even larger delay variations, more bytes of overhead or hierarchy H4 coding recommended in the ITU-T standard may be used.

Geyser Framing Structure

Referring next to Figure 5, a STS-1 superframe or framing structure 500 in accordance with an exemplary embodiment of the present invention is shown. The present invention's STS-1 framing structure supports the nxSTS-1 and nxVT virtual concatenation shown in Figure 2.

STS-1 framing structure 500 is compatible with GR-253 in STS-1 frames. STS-1 framing structure 500 may pass safely through the standard SONET networking environment. On top of this standard compliant STS-1 frame, the present invention provides the following two additions: (1) a super-frame (SFRM) 500; and (2) a virtual tributary structure.

STS-1 Super-Frame

STS-1 super-frame (SFRM) 500 consists of 8 STS-1 frames, as compared with GR-253's 4 frames per multiple-frame. In STS-1 super-frame (SFRM) 500, one STS-1 SPE has a total of 87(column) x 9 (row) x 8 (frame) bytes. The STS-1 overhead is defined to be the same as GR-253 to maintain compatibility. Similarly, the two fix stuffed columns 550 and 552 (column #29 and column #58) remain the same but may be optionally used for data payload.

Super-frame structure 500 may also apply to super-frames that consist of n frames, where n is a positive integer.

Virtual Tributary Structure

Inside STS-1 super-frame 500, a virtual tributary (VT) structure 510 is defined differently from the standard SONET GR-253 VT1.5 structure. In standard SONET, the GR-253 VT1.5 virtual tributary structure uses bytes interleaved with 3 separate columns for one T1 mapping. In the present invention, virtual tributary structure 510 is provided which uses a word interleaved multiplexing scheme. The format of the present invention's SFRM structure for VT is shown in Figure 5.

In Figure 5, in an exemplary embodiment, one word of VT is mapped to one slot within a STS-1 super-frame. The first row and the second 1/3 row within one super-frame carry the overhead byte for 28 VT1.5s. Each VT1.5 in one super frame has an overhead of one word (4 bytes) and a payload of 212 bytes. The payload slots start from the second 1/3 row of second row. The total VT capacity per Super-Frame is 3x28x9x8 = 6048 bytes. For convenience, the present invention's word interleaved virtual tributary is abbreviated as VT.

Actually VT1.5 overhead can be distributed in any location of the super-frame.

STS-1 superframe 500 includes 28 slots for 28 VTs. Each VT includes 53 words. In Figure 5, a sloti[j] refers to the ith VT and the jth word in the VT, where i is between 1 and 28 and j is between 1 and 53. STS-1 superframe 500 is 1000 μs in duration. The path overhead bytes of a SONET STS-1 are stored in path overhead 520. P, Q, U, and V bytes, as described in co-pending and commonly assigned U.S. Patent Application No. (Number to be assigned) with Attorney Docket Number 55369-015, of a VT.15 are stored in bytes 530. The payload of the 1st VT1.5 is stored in slot1[1] (slot 1, word 1) 540 to slot1[53] (slot 1, word 53) 542, as shown in Figure 5. As can be followed, the ith VT1.5 is stored in sloti[1] (slot i, word 1) to sloti[53] (slot i, word 53).

STS-1 Superframe and Virtual Tributary Structure

STS-1 superframe 500 and virtual tributary structure 510 has several important features.

First, with a super-frame of 8 frames instead of 1 or 4 in standard SONET, the virtual tributary payload overhead 520 is largely decreased. For example, the overhead of the present invention's VT is 4/216 = 1.85%, as compared with the standard SONET GR-253 overhead of 16/216 = 7.4%. The present invention's payload per VT is 1.696 Mbps as compared with the standard SONET GR-253's payload per VT of 1.6 Mbps.

Second, super-frame 500 of 8 frames allows for each virtual tributary payload 510 to carry exactly 4 ATM cells. Whenever an ATM cell is mapped onto the present invention's VT tributary 510, the ATM cell boundary is aligned with the super-frame boundary. In this case, the present invention allows for the easy transmission or reception of ATM cells without the need of either implementing ATM cell delineation or additional coding such as HDLC.

Third, the byte interleaving based standard SONET stream was defined for 8-bit wide data path. When the throughput gets higher, it's mandatory to use a wider data path in the design to support the higher data rate. By defining a word interleaved virtual tributary, the present invention makes it very easy to match between the VT data stream and the internal data path. No additional memory or conversion is needed.

Conclusion

The present invention relates to optical networks. More particularly, the invention relates to a system and method of virtually concatenating VT1.5s and STS-1s over SONET and SDH and WDM.

Having fully described a preferred embodiment of the invention and various alternatives, those skilled in the art will recognize, given the teachings herein, that numerous alternatives and equivalents exist which do not depart from the invention. It is therefore intended that the invention not be limited by the foregoing description, but only by the appended claims.

CLAIMS

We claim:

1. A method for virtually concatenating VT1.5s and STS-1s in a SONET pipe comprising:

creating a nxSTS-1 virtually concatenated pipe out of STS-1s;

within the nxSTS-1 pipe, forming one or more nxVT pipes out of VTs; and

mapping an application to a nxVT pipe or a nxSTS-1 pipe based upon a service level agreement.

2. The method of claim 1 wherein the creating comprises:

terminating SONET framing;

processing pointers; and

compensating for delay, thereby aligning the STS-1s in the nxSTS-1 virtually concatenated pipe.

3. The method of claim 1 wherein the forming comprises:

processing VT overhead of the VTs; and

mapping traffic to nxVT pipes or nxSTS-1 pipes according to a connection identification.

4. The method of claim 1 wherein the mapping comprises:

processing traffic; and

mapping traffic to a nxVT pipe or a nxSTS-1 pipe based upon the service level agreement.

5. The method of claim 3 wherein the processing comprises:

capturing a superframe;

retrieving a VT overhead message;

based upon the VT overhead message, performing a connection table lookup;

putting a traffic word from a traffic first-in, first-out buffer in the superframe based upon the result of the connection table lookup; and

modifying a overhead of the superframe, where the superframe is an outgoing superframe.

6. The method of claim 3 wherein the processing comprises:

capturing a superframe;

retrieving a VT overhead message;

based upon the VT overhead message, performing a connection table lookup;

putting a traffic word from a traffic first-in, first-out buffer in the superframe based upon the result of the connection table lookup; and

modifying a overhead of each VT.

7. A system for virtually concatenating VT1.5s and STS-1s in a SONET pipe comprising:

a creating module configured to create a nxSTS-1 virtually concatenated pipe out of STS-1s;

a forming module configured to form, within the nxSTS-1 pipe, one or more nxVT pipes out of VTs; and

a mapping module configured to map an application to a nxVT pipe or a nxSTS-1 pipe based upon a service level agreement .

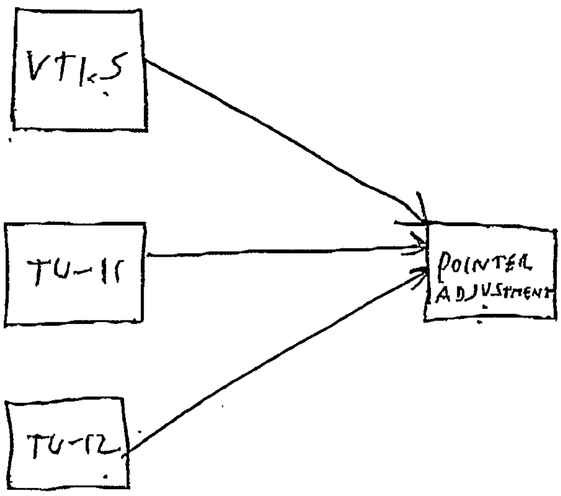

(PRIOR ART)

FIG. 1B

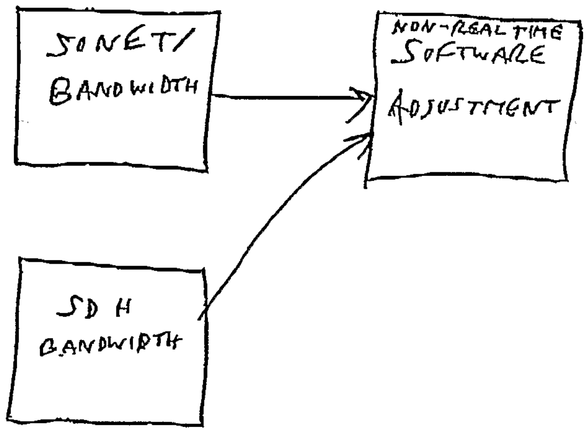

(PRIOR ART)

FIG. 1C

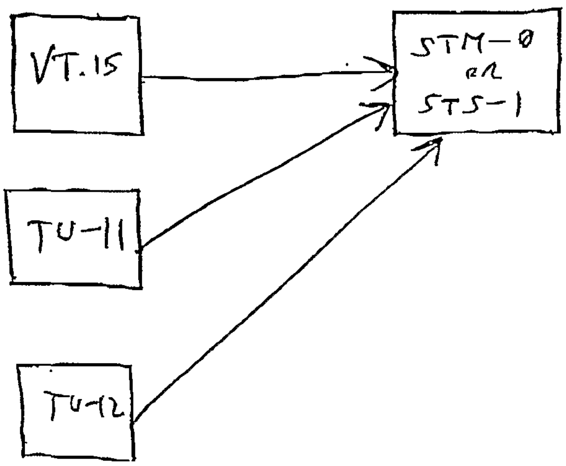

(PRIOR ART)

FIG. 1A

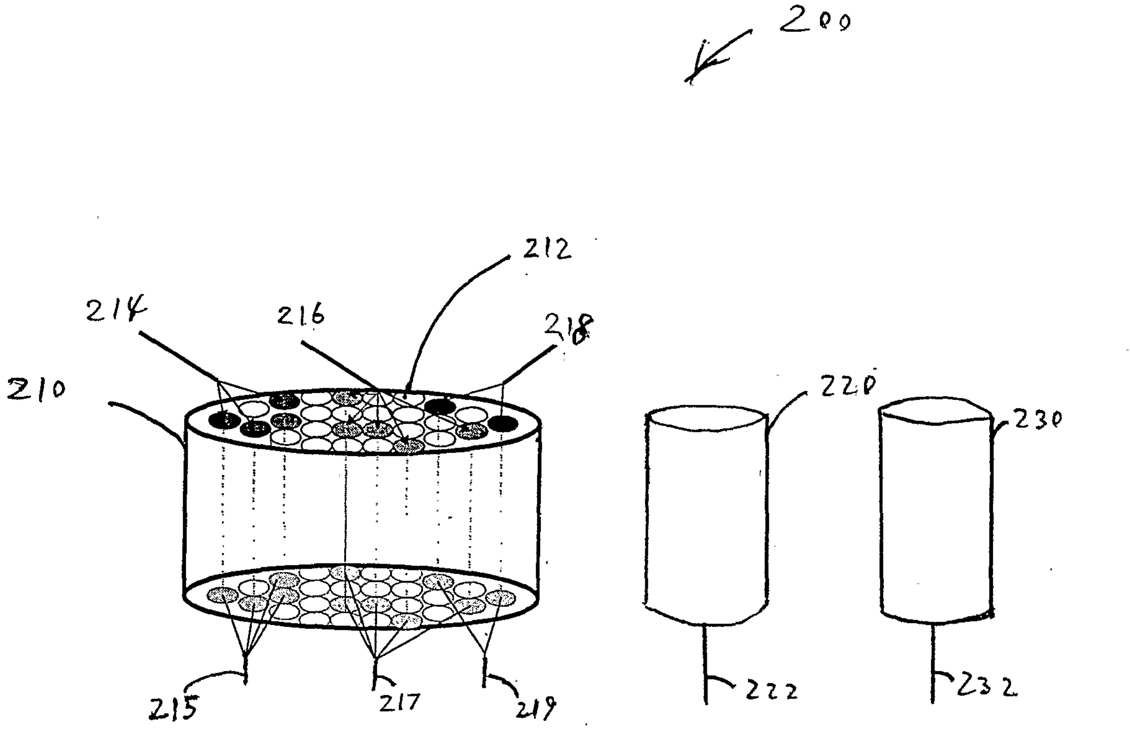

FIG. 2

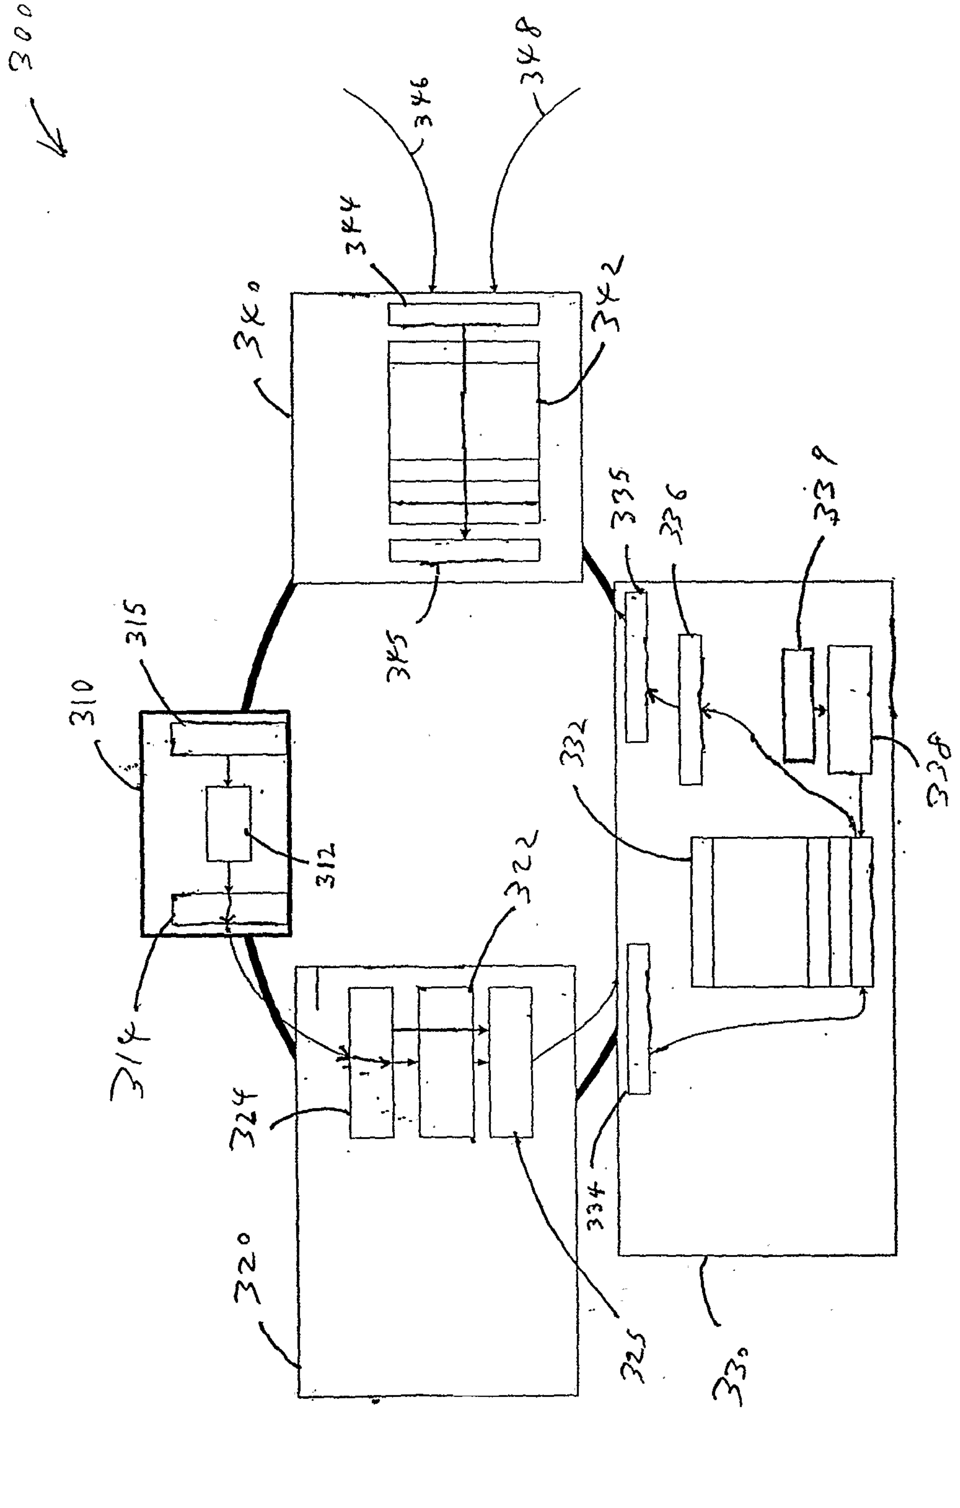
300
310
312
314
315
320
322
324
325
330
332
334
335
336
338
339
340
342
344
345
346
348

FIG. 3

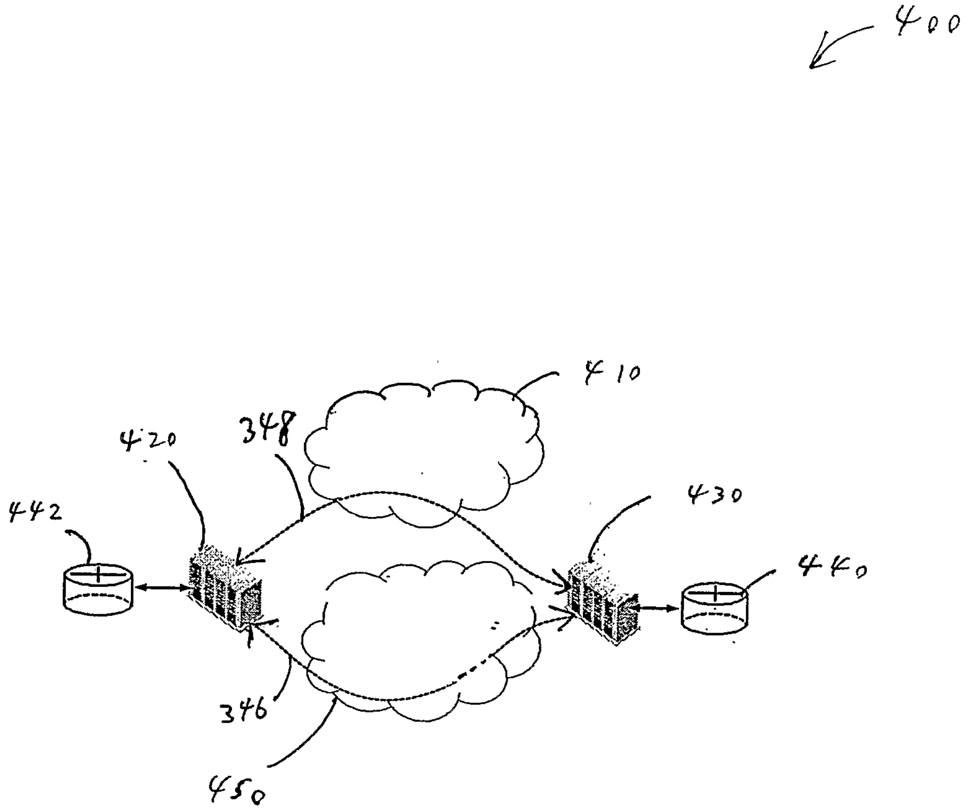

FIG. 4

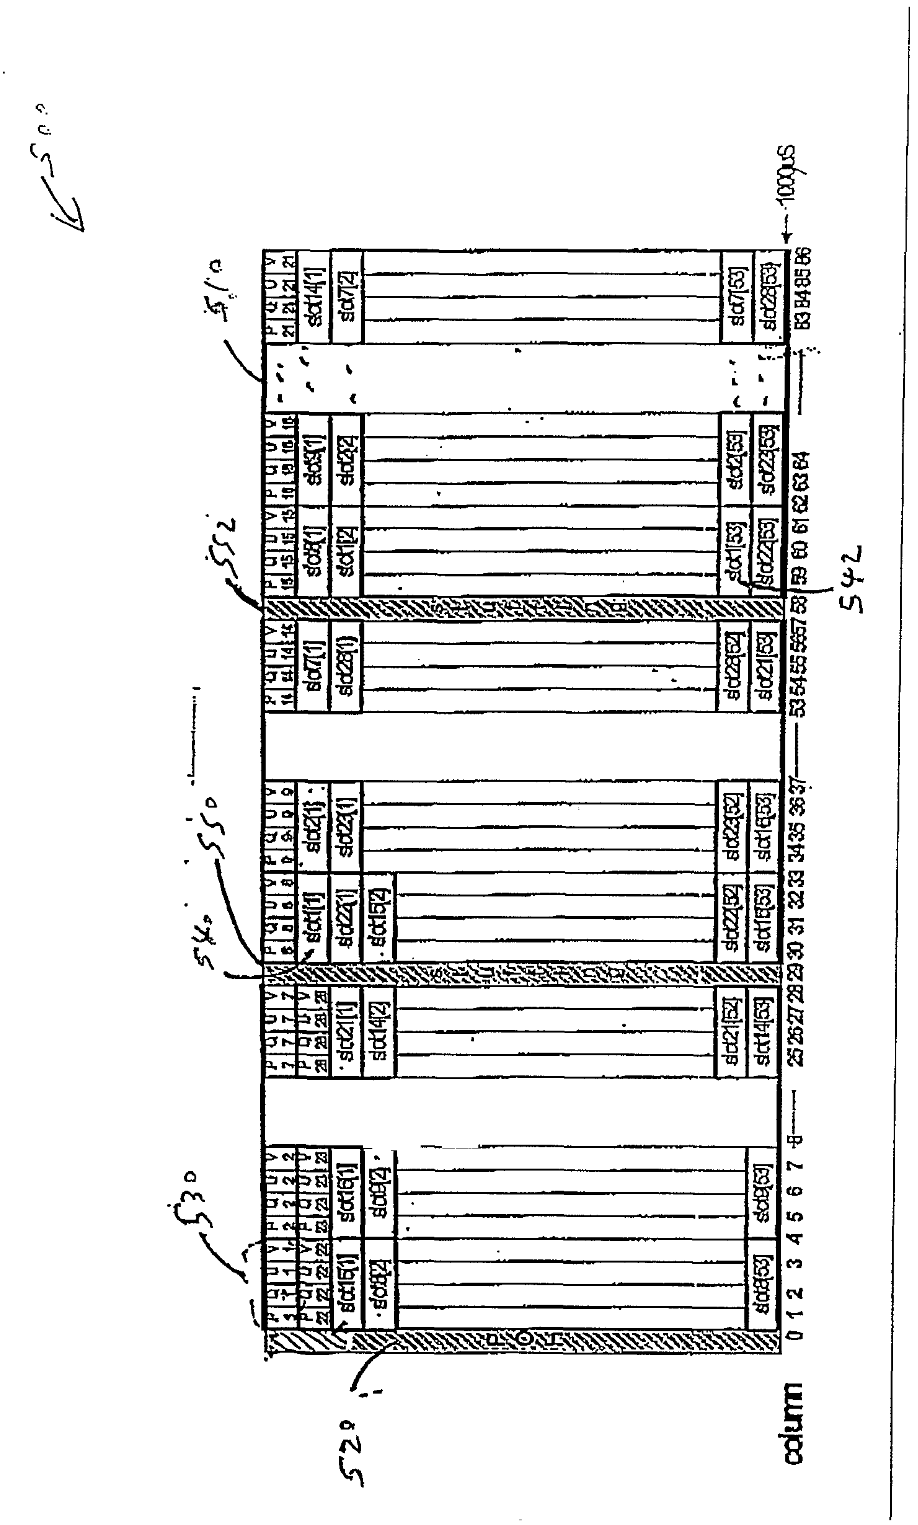

FIG. 5