

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

(19) World Intellectual Property Organization International Bureau



(43) International Publication Date
31 March 2005 (31.03.2005)

PCT

(10) International Publication Number
WO 2005/029906 A1

(51) International Patent Classification⁷: H04Q 11/00

(21) International Application Number:
PCT/US2004/030365

(22) International Filing Date:
14 September 2004 (14.09.2004)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
60/502,856 15 September 2003 (15.09.2003) US
60/504,936 23 September 2003 (23.09.2003) US

(71) Applicant (for all designated States except US): TEKNOVUS, INC. [US/US]; 1363 Redwood Way, Petaluma, CA 94954 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): KRAMER, Glen [US/US]; 821 Rolling Hill Drive, Petaluma, CA 94954 (US). BOYD, Edward, W. [US/US]; 1806 Pembridge Street, Petaluma, CA 94954 (US).

(74) Agent: PARK, A., Richard; 508 Second Street, Suite 201, Davis, CA 95616 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NA, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.

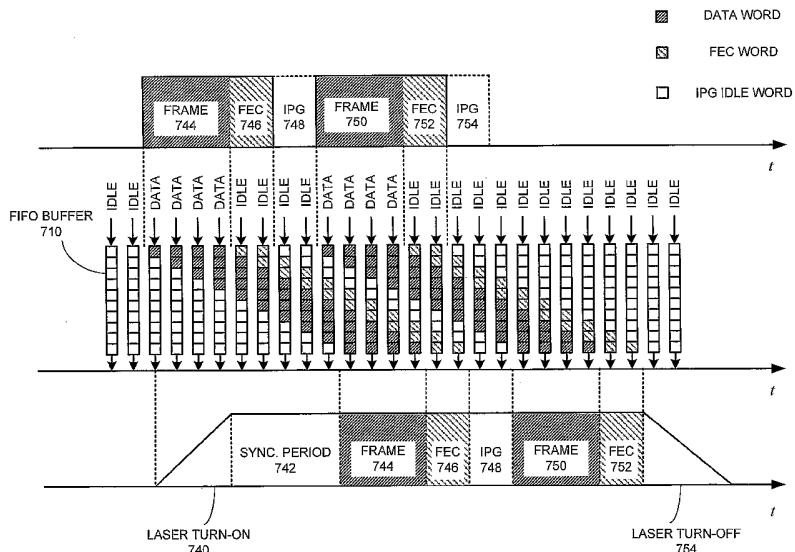
(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IT, LU, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

- with international search report
- before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments

[Continued on next page]

(54) Title: METHOD AND APPARATUS FOR TRANSMISSION CONTROL IN AN ETHERNET PASSIVE OPTICAL NETWORK



WO 2005/029906 A1

(57) Abstract: One embodiment of the present invention provides a system that facilitates transmission control in an Ethernet passive optical network, which includes a central node and at least one remote node, and wherein a remote node implements a data-link layer and a physical layer. During operation, the system starts by receiving, at the physical layer of a remote node, a word which is communicated from the data-link layer of the remote node, wherein the word may be a data word or an idle word. The system then delays the word for a pre-determined amount of time before allowing the word to be transmitted by a transmitter, thereby providing time for turning the transmitter on or off. The system also turns the transmitter on or off based on the content of the received words.



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.

METHOD AND APPARATUS FOR TRANSMISSION CONTROL IN AN ETHERNET PASSIVE OPTICAL NETWORK

5

Inventors: Glen Kramer and Edward W. Boyd

BACKGROUND

10

Field of the Invention

[0001] The present invention relates to the design of Ethernet passive optical networks. More specifically, the present invention relates to a method and an apparatus for controlling transmission to reduce interference caused by laser noise from different nodes in an Ethernet passive optical network .

Related Art

[0002] In order to keep pace with increasing Internet traffic, optical fibers and associated optical transmission equipment have been widely deployed to substantially increase the capacity of backbone networks. However, this increase in the capacity of backbone networks has not been matched by a corresponding increase in the capacity of access networks. Even with broadband solutions, such as digital subscriber line (DSL) and cable modem (CM), the limited bandwidth offered by current access networks creates a severe bottleneck in delivering high bandwidth to end users.

[0003] Among the different technologies that are presently being developed, Ethernet passive optical networks (EPONs) are one of the best candidates for next-generation access networks. EPONs combine ubiquitous Ethernet technology with inexpensive passive optics. Hence, they offer the simplicity and scalability of Ethernet with the cost-efficiency and high capacity of passive optics. In particular, due to the high bandwidth of optical fibers, EPONs are capable of accommodating

broadband voice, data, and video traffic simultaneously. Such integrated service is difficult to provide with DSL or CM technology. Furthermore, EPONs are more suitable for Internet Protocol (IP) traffic, because Ethernet frames can directly encapsulate native IP packets with different sizes, whereas ATM passive optical networks (APONs) use fixed-size ATM cells and, consequently, require packet fragmentation and reassembly.

[0004] Typically, EPONs are used in the “first mile” of the network, which provides connectivity between the service provider’s central offices and business or residential subscribers. Logically, the first mile is a point-to-multipoint network, with a central office servicing a number of subscribers. A tree topology can be used in an EPON, wherein one fiber couples the central office to a passive optical splitter, which divides and distributes downstream optical signals to subscribers and combines upstream optical signals from subscribers (see FIG. 1).

[0005] Transmissions within an EPON are typically performed between an optical line terminal (OLT) and optical networks units (ONUs) (see FIG. 2). The OLT generally resides in the central office and couples the optical access network to a metro backbone, which is typically an external network belonging to an Internet Service Provider (ISP) or a local exchange carrier. An ONU can be located either at the curb or at an end-user location, and can provide broadband voice, data, and video services. ONUs are typically coupled to a one-by-N (1xN) passive optical coupler, where N is the number of ONUs, and the passive optical coupler is typically coupled to the OLT through a single optical link. (Note that one may cascade a number of optical splitters/couplers to accommodate more ONUs.) This configuration can significantly save the number of fibers and the amount of hardware required by EPONs.

[0006] Communications within an EPON can be divided into downstream traffic (from OLT to ONUs) and upstream traffic (from ONUs to OLT). In the downstream direction, because of the broadcast nature of the 1xN passive optical coupler, downstream data frames are broadcast by the OLT to all ONUs and are subsequently extracted by their destination ONUs. In the upstream direction, the

ONUs need to share channel capacity and resources, because there is only one link coupling the passive optical coupler with the OLT.

[0007] Correspondingly, an EPON typically employs some arbitration mechanism to avoid data collision and to provide fair sharing of the upstream fiber-5 channel capacity. This is achieved by allocating a transmission timeslot to each ONU.

An ONU typically buffers data it receives from a subscriber until the ONU's local time reaches the start time of its transmission timeslot. When its turn arrives, the ONU "bursts" all stored frames to the OLT at full channel speed.

[0008] One issue in designing an EPON is the interference with data10 transmission caused by laser noise from different ONUs. A laser transmitter within an ONU generates spontaneous emission noise, even in the absence of data transmission. Hence, if an ONU's laser remains on between transmission timeslots, its spontaneous emission noise may impair the signal quality of data transmitted by another ONU. This impairment worsens when the noise from multiple closely-located ONUs15 interferes with the data signal from a distant ONU.

[0009] One can avoid the laser-noise problem by turning off a laser when it is not transmitting data. This can be accomplished by using a laser-control signal, which is generated in the data link layer and received in the physical layer.

[0010] However, sending a control signal across several sublayers violates a20 basic layering principle, which is, a layer or a sublayer may only signal its immediate neighboring layers or sublayers. Note that in an ONU, the data link layer includes a medium access control (MAC) client sublayer, a multi-point MAC control sublayer, and a MAC sublayer. Furthermore, the physical layer includes a reconciliation sublayer, a Gigabit media independent interface (GMII), a physical coding sublayer25 (PCS), a physical medium attachment (PMA) sublayer, and a physical medium dependent (PMD) sublayer. In this environment, a laser-control signal is typically generated by the MAC control sublayer and is received by the PMD sublayer, thereby bypassing the sublayers between the MAC control sublayer and the PMD sublayer.

[0011] Hence, what is needed is a method and an apparatus for efficiently controlling laser transmission for an ONU without sending a control signal across multiple sublayers.

SUMMARY

5 **[0012]** One embodiment of the present invention provides a system that facilitates transmission control in an Ethernet passive optical network, which includes a central node and at least one remote node, and wherein a remote node implements a data-link layer and a physical layer. During operation, the system starts by receiving, at the physical layer of a remote node, a word which is communicated from the data-link layer of the remote node, wherein the word may be a data word or an idle word.
10 The system then delays the word for a pre-determined amount of time before allowing the word to be transmitted by a transmitter, thereby providing time for turning the transmitter on or off. The system also turns the transmitter on or off based on the content of the received words.

15 **[0013]** In a variation of this embodiment, turning the transmitter on or off based on the content of the received words involves turning on the transmitter if the transmitter is currently powered off and if the received word is a data word.

20 **[0014]** In a variation of this embodiment, turning the transmitter on or off based on the content of the received words involves turning off the transmitter if the transmitter is currently powered on and if the number of consecutive received idle words is equal to or greater than a pre-determined number.

25 **[0015]** In a variation of this embodiment, delaying the word for a pre-determined amount of time involves: buffering the word in a first-in-first-out (FIFO) queue; shifting each word buffered in the FIFO queue towards the head of the FIFO queue at regular time intervals; and removing a word at the head of the FIFO queue at regular time intervals.

[0016] In a variation of this embodiment, the system determines the amount of time for which the word is delayed based on a laser turn-on time, an automatic gain control (AGC) time, and a data and clock recovery (CDR) time.

[0017] In a variation of this embodiment, if the received word is a data word, the system temporarily prohibits the data link layer from sending another data frame immediately after the current data frame which is being received at the physical layer, thereby allowing a sufficient number of idle words to be inserted between data frames.

5 [0018] In a variation of this embodiment, if there are a sufficient number of consecutive idle words received at the physical layer, the system allows the data link layer to send another data frame to the physical layer.

BRIEF DESCRIPTION OF THE FIGURES

10 [0019] FIG. 1 illustrates an Ethernet passive optical network wherein a central office and a number of subscribers are coupled through optical fibers and a passive optical splitter (prior art).

[0020] FIG. 2 illustrates an EPON in normal operation mode (prior art).

[0021] FIG. 3 illustrates bridged Ethernet segments (prior art).

15 [0022] FIG. 4A illustrates transmission of downstream traffic with point-to-point emulation in an EPON (prior art).

[0023] FIG. 4B illustrates transmission of upstream traffic with point-to-point emulation in an EPON (prior art).

20 [0024] FIG. 5A illustrates the structure of a transmission timeslot in an EPON (prior art).

[0025] FIG. 5B illustrates the transmission of upstream data bursts from multiple ONUs (prior art).

[0026] FIG. 6 illustrates generating a laser control signal from the data link layer to control a laser which belongs to the physical layer (prior art).

25 [0027] FIG. 7 illustrates the process of controlling a laser transmitter based on buffered words received from the data link layer in accordance with one embodiment of the present invention.

30 [0028] FIG. 8A presents a flow chart illustrating the process of buffering words received from the data link layer without signaling the data link layer for frame-gap control in accordance with one embodiment of the present invention.

[0029] FIG. 8B presents a flow chart illustrating the process of buffering words received from the data link layer with signaling the data link layer for frame-gap control in accordance with one embodiment of the present invention.

5

DETAILED DESCRIPTION

[0030] The following description is presented to enable any person skilled in the art to make and use the invention, and is provided in the context of a particular application and its requirements. Various modifications to the disclosed embodiments will be readily apparent to those skilled in the art, and the general principles defined herein may be applied to other embodiments and applications without departing from the spirit and scope of the present invention (e.g., general passive optical network (PON) architectures). Thus, the present invention is not intended to be limited to the embodiments shown, but is to be accorded the widest scope consistent with the principles and features disclosed herein.

15 [0031] The data structures and procedures described in this detailed description are typically stored on a computer readable storage medium, which may be any device or medium that can store code and/or data for use by a computer system. This includes, but is not limited to, application specific integrated circuits (ASICs), field-programmable gate arrays (FPGAs), semiconductor memories, magnetic and 20 optical storage devices such as disk drives, magnetic tape, CDs (compact discs) and DVDs (digital versatile discs or digital video discs), and computer instruction signals embodied in a transmission medium (with or without a carrier wave upon which the signals are modulated).

25 **Passive Optical Network Topology**

[0032] FIG. 1 illustrates a passive optical network, wherein a central office and a number of subscribers are coupled together through optical fibers and a passive optical splitter (prior art). As shown in FIG. 1, a number of subscribers are coupled to a central office 101 through optical fibers and a passive optical splitter 102. Passive 30 optical splitter 102 can be placed in the vicinity of end-user locations, so that the

initial fiber-deployment cost is minimized. Central office 101 can be coupled to an external network 103, such as a metropolitan area network operated by an Internet service provider (ISP). Note that although FIG. 1 illustrates a tree topology, a passive optical network can also be based on other topologies, such as a ring or a bus.

5

Normal Operation Mode in EPON

[0033] FIG.2 illustrates an EPON in normal operation mode (prior art). To allow ONUs to join an EPON at arbitrary times, an EPON typically has two modes of operation: a normal operation mode and a discovery (initialization) mode. Normal operation mode involves regular upstream data transmissions, and allows an OLT to assign transmission opportunities to all initialized ONUs.

[0034] As shown in FIG. 2, in the downstream direction, OLT 201 broadcasts downstream data to ONU 211, ONU 212, and ONU 213. While all ONUs may receive the same copy of data, each ONU selectively forwards only the data destined to itself to its corresponding users, which are user 221, user 222, and user 223, respectively.

[0035] In the upstream direction, OLT 201 first schedules and assigns transmission timeslots to each ONU according to the ONU's service-level agreement. When not in its transmission timeslot, an ONU typically buffers the data received from its user. When its scheduled transmission timeslot arrives, an ONU transmits the buffered user data within the assigned transmission timeslot.

[0036] Since every ONU takes turns in transmitting upstream data according to the OLT's scheduling, the upstream link's capacity can be efficiently utilized. However, for the scheduling to function properly, the OLT needs to discover and initialize a newly joined ONU. During discovery, the OLT may collect information critical to transmission scheduling, such as the ONU's round-trip time (RTT), its media access control (MAC) address, its service-level agreement, etc. (Note that in some cases service-level agreement may already be known to the OLT),

30 General Ethernet Requirement

[003*] FIG. 3 illustrates bridged Ethernet segments (prior art). The IEEE 802 standards allow an Ethernet segment to operate in a point-to-point mode. In a point-to-point Ethernet segment, a link couples two hosts, or a host and an Ethernet bridge. Point-to-point mode is a common form of operation in a switched Ethernet, such as 5 Gigabit Ethernet.

[0038] When multiple Ethernet hosts need to communicate with each other, an Ethernet bridge typically couples and switches between multiple point-to-point Ethernet segments to allow inter-segment communications. As shown in FIG. 3, Ethernet bridge 310 has multiple ports. Point-to-point segments 321, 322, and 323 are 10 coupled to ports 311, 312, and 313, respectively. If the host on segment 322 sends a data frame to the host on segment 321, the data frame will be forwarded by Ethernet bridge 310 from port 312 to port 311 according to its destination Ethernet (MAC) address. Generally, a bridge does not forward a frame back to the port on which it arrives.

15

Point-to-Point Emulation (PtPE) in EPON

[0039] In an EPON, because the upstream transmission from an ONU to the OLT is point-to-point communication, the operation of EPON ideally conforms to the point-to-point Ethernet operation as defined by the IEEE 802 standards. However, the 20 EPON architecture does not automatically satisfy the requirement of bridged point-to-point Ethernet: if the EPON upstream link is coupled to one Ethernet bridge port, and all the upstream traffic is received at that port, users connected to different ONUs on the same EPON will be unable to communicate with each other. The Ethernet bridge located within the OLT will not switch among the upstream data, because they are 25 received on the same port. Such a configuration forces data traffic among ONUs within the same EPON to be processed on layer 3 (network layer) and switched by equipment that resides outside the EPON (e.g., an IP router to which the OLT is connected). This is a very inefficient way of delivering intra-EPON traffic.

[0040] To resolve this problem, and to ensure seamless integration of an 30 EPON with other Ethernet networks, devices attached to the EPON medium ideally

have an additional sublayer that can emulate a point-to-point medium. This sublayer is referred to as the reconciliation sublayer (RS), and resides below the data link layer to preserve existing Ethernet MAC operation defined in the IEEE 802.3 standards.

Operation of RS relies on tagging Ethernet frames with tags unique for each ONU.

5 These tags are called logic link IDs (LLIDs) and are placed in the preamble before each frame.

[0041] FIG. 4A illustrates transmission of downstream traffic with point-to-point emulation in an EPON (prior art). In PtPE mode, OLT 400 has multiple MAC ports (interfaces), each of which corresponds to an ONU. When sending an Ethernet 10 frame downstream from MAC port 431, reconciliation sublayer 440 in OLT 400 inserts LLID 461 which is associated with MAC port 431. Although the frame is broadcast through the passive optical coupler to every ONU, only the reconciliation sublayer module located within an ONU with a matching LLID (ONU 451 with LLID 461 in this example) will accept the frame and pass it to its data link layer for further 15 verification. MAC sublayers in other ONUs (ONU 452 with LLID 462, and ONU 453 with LLID 463) generally do not receive that frame. Accordingly, it appears as if the frame is sent on a point-to-point link to the destination ONU.

[0042] FIG. 4B illustrates transmission of upstream traffic with point-to-point emulation in an EPON (prior art). In the upstream direction, ONU 451 inserts 20 its assigned LLID 461 in the preamble of each transmitted frame. Accordingly, reconciliation sublayer 440 of OLT 400 disseminates the frame to MAC port 431.

Sharing of the Upstream Link Among Multiple ONUs

[0043] In an EPON, multiple ONUs typically share one uplink for transmitting 25 upstream data, because there is only one link between the optical splitter and the OLT. To coordinate this time-shared transmission, the OLT assigns transmission timeslots to each ONU. One issue in implementing such a time-shared transmission scheme is the interference with data transmitted from one ONU caused by laser noise from other ONUs.

10

[0044] A laser transmitter, when powered on, generates spontaneous emission noise even in the absence of data transmission. If the laser transmitter of an ONU is not turned off when it is not transmitting data, the accumulated noise from multiple ONUs may interfere with the data transmission of another ONU. This interference 5 may worsen if the noise-generating ONUs are closer to the optical splitter than the data-transmitting ONU is.

[0045] One approach to solve this interference problem is to turn off a laser between transmission slots. FIG. 5A illustrates the structure of a transmission timeslot in an EPON (prior art). An upstream data burst contained in a transmission 10 timeslot comprises several parts beside its data payload. As shown in FIG. 5A, a timeslot may contain a laser turn-on period 501, an automatic gain control (AGC) period 502, a clock and data recovery (CDR) period 503, a data/idle payload 504, and a laser turn-off period 505. Note that, because a laser generally does not turn on or off 15 instantaneously, the laser response time should be accounted for when calculating the right moment to turn on or turn off a laser with regard to the assigned timeslot.

[0046] FIG. 5B illustrates the transmission of upstream data bursts from multiple ONUs (prior art). As shown in FIG. 5B, five ONUs 521 – 525 communicate with OLT 500. Upstream data bursts 541 and 542, transmitted by ONUs 521 and 522 respectively, share the same uplink. The lasers within ONUs 522 – 525 are turned off 20 when ONU 521 is transmitting data burst 541. After completing data transmission, ONU 521 turns off its laser to allow ONU 522 to transmit data burst 542.

ONU Transmission Control

[0047] To turn on a laser when the assigned timeslot arrives and to turn off the 25 laser when the transmission is complete, one may simply generate a laser on/off control signal from the data link layer to control the laser which belongs to the physical layer. This is because an ONU's local clock typically resides in the data link layer, based on the IEEE 802.3ah standards, and hence it is suitable for the data link layer to detect the beginning and the end of a transmission timeslot. However, this

approach violates a basic layering principle, which is, a layer or a sublayer may only signal its immediate neighboring layers or sublayers.

[0048] FIG. 6 illustrates generating a laser control signal from the data link layer to control a laser which belongs to the physical layer (prior art). As shown in 5 FIG. 6, the OSI model layers comprise a physical layer 612, a data link layer 614, a network layer 616, a transport layer 618, a session layer 620, a presentation layer 622, and an application layer 624. EPON equipment mainly operates within physical layer 612 and data link layer 614. Based on the IEEE 802.3 standards, data link layer 614 in an EPON includes an LLC/MAC client sublayer 622, a multi-point MAC 10 control sublayer 628, and a MAC sublayer 630; and the physical layer includes a reconciliation sublayer 636, a GMII interface 638, a PCS sublayer 640, a PMA sublayer 642, a PMD sublayer 644, a medium dependent interface (MDI) 646, and the physical medium 648 (optical fiber). In particular, a physical-layer device (PHY) generally implements the PCS, PMA, and PMD sublayers.

15 [0049] A laser on/off control signal 650 is typically generated at multi-point MAC control sublayer 628 and sent to PMD sublayer 644 to turn on or turn off the laser. Such a signaling mechanism, although simple, violates the layering principle that in general a sublayer may not signal across multiple layers/sublayers. One way to avoid this layering violation is to eliminate laser control signal 650 which travels 20 across multiple sublayers.

[0050] One embodiment of the present invention introduces a pre-determined amount of delay for every word received from the data link layer at the physical layer. In this way, the physical layer can determine, by observing the words which are delayed, when there is data to be transmitted, and can have enough time to turn on the 25 laser and to prepare for the data transmission. Similarly, the physical layer can determine when there is no more data to be transmitted, and can turn off the laser correspondingly. To implement the delay mechanism described above, one may use a first-in-first-out (FIFO) queue to buffer each word.

[0051] FIG. 7 illustrates the process of controlling a laser transmitter based on 30 buffered words received from the data link layer in accordance with one embodiment

of the present invention. As shown on the top of FIG. 7, the data received from the data link layer include data frames 744 and 750. Following frame 744 is the space for optional forward error correction (FEC) parity bits 746 and inter-packet gap (IPG) 748. Similarly, following frame 750 is the optional FEC space 752 and IPG space 754. Note that an FEC encoder may be implemented before or after the FIFO queue. Hence, FEC spaces 746 and 752 may contain FEC parity data for the preceding frame (if FEC encoding is implemented before the FIFO queue), or idles (if FEC encoding is implemented after the FIFO queue).

[0052] Typically, the length of an FEC space is a function of the preceding frame's length. In the illustrated example, FEC space 746 depends on the length of frame 744, and FEC space 752 depends on the length of frame 750. In addition, there is a minimum required size for IPG. For example, the IEEE 802.3ah standards define IPG to be 12-word long. FEC, on the other hand, is optional.

[0053] When words arrive at the physical layer, they are stored in a FIFO buffer 710 which has a pre-determined delay. Each word, after being stored at the tail of FIFO buffer 710, shifts towards the head of FIFO buffer 710 at regular time intervals.

[0054] To determine when to turn on or turn off the laser, the system ideally monitors the content of the received words. When a received word is a data word and the laser is currently powered off, the system turns on the laser. The received word is then stored in FIFO buffer 710. The total delay introduced by FIFO buffer 710, when combined with other delays introduced in the physical layer, is ideally sufficient to accommodate laser turn-on period 740 and synchronization period 742. Note that synchronization period 742 may include an automatic gain control (AGC) period and a clock and data recovery (CDR) period, which allow the OLT to prepare for reception of the upstream data.

[0055] FEC space 746 following frame 744 accommodates FEC parity bits for frame 744. Because FEC encoding may be implemented before or after the FIFO queue, FEC space 746 as received from the data link layer may contain encoded FEC parity data or may contain only idles. After FEC space 746, IPG space 748 ensures

that there is a minimum gap between frame 744 (including the corresponding FEC space 768) and the subsequent frame 752. In the example shown in FIG. 7, IPG space 748 is shorter than the sum of laser turn-on period 740 and synchronization period 742. Therefore, the system does not turn off the laser during the period corresponding to IPG space 748.

[0056] Similarly, frame 750, its corresponding FEC space 752, and IPG space 754 are stored in FIFO buffer 710 when they arrive at the physical layer. The system only turns off the laser when the physical layer has received a sufficiently large number of consecutive IPG idle words. The minimum number of consecutive IPG 10 idle words required for turning off the laser is ideally sufficient to accommodate the laser turn-on period 740 and/or synchronization period 742, in order to allow enough time for the laser to be turned on when a data word follows these idle words.

[0057] In one embodiment of the present invention, the system may require FIFO buffer 710 to contain only IPG idle words to turn off the laser. In this way, one 15 can ensure that there is sufficient time to turn on the laser and to provide synchronization time for the next data word. As shown in FIG. 7, FIFO buffer 710 contains only IPG idle words after the last word within FEC space 752 leaves FIFO buffer 710. At this moment, the system turns off the laser, which completely shuts off after laser turn-off period 754. Note that, if FEC is not used, the system may turn off 20 laser as soon as FIFO buffer 710 contains only IPG idle words after the last data word of frame 750 leaves FIFO buffer 710.

[0058] In some cases where an ONU resides close to the OLT, the OLT may need less than the default AGC and CDR time for synchronization. Correspondingly, the system in such an ONU may shorten the delay introduced by FIFO buffer 710 to 25 reflect the reduced synchronization period and to reduce transmission overhead. In general, the system may re-configure the delay of FIFO buffer 710 for more flexibility.

[0059] It is important that the correct amount of FEC space and IPG space be reserved between two consecutive data frames. The reservation of FEC space and IPG space may be accomplished in the data link layer or in the physical layer. In the 30 first scenario, the data link layer (e.g., the multi-point MAC control sublayer) is

responsible for detecting data frame boundaries and inserting necessary space for FEC parity data and IPG. Correspondingly, the transmission-control mechanism in the physical layer does not need to signal the data link layer as to when to transmit a data frame and when to pause transmission.

5 [0060] In the second scenario, the data link layer is not responsible for inserting space for FEC and IPG between data frames. Instead, the physical layer is responsible for detecting data frame boundaries and inserting necessary space for FEC and IPG. Correspondingly, the physical layer ideally signals the data link layer as to when to transmit a data frame and when to pause transmission.

10 [0061] FIG. 8A presents a flow chart illustrating the process of buffering words received from the data link layer without signaling the data link layer for frame-gap control in accordance with one embodiment of the present invention. The system starts by receiving a word from the data link layer (step 810). The system then determines whether the received word is a data word or an idle word (step 812).

15 [0062] If the received word is a data word, the system sets an idle length counter to 0 (step 813). Idle length represents the number of consecutive idles and is used in determining whether the laser should be turned off. The system then determines whether the laser is on (step 814). If the laser is off, the system turns on the laser (step 816). Otherwise, the system proceeds directly to remove the first word 20 in the FIFO buffer (step 840). Next, the system shifts the words in the FIFO buffer towards the head of the buffer (step 841), and appends the received data word to the tail of the FIFO buffer (step 842). The system then prepares to receive the next word from the data link layer (step 810).

25 [0063] If the received word is an idle word, the system first increments the idle length counter by 1 (step 820). Next, the system determines whether the laser is currently on, and whether the idle length is greater than a delay bound, which is defined as the sum of laser turn-on period and the synchronization time (step 852). If both conditions are true, the system turns off the laser (step 854), and removes the first word in the FIFO buffer (step 840). Otherwise, the system does not change the

current state of the laser, and directly proceeds to remove the first word in the FIFO buffer (step 840).

[0064] FIG. 8B presents a flow chart illustrating the process of buffering words received from the data link layer and signaling the data link layer for frame-gap control in accordance with one embodiment of the present invention. The system starts by receiving a word from the data link layer (step 860). The system then determines whether the received word is a data word or an idle word (step 862).

[0065] If the received word is a data word, the system sets an idle length counter to 0 (step 864), and temporarily prohibits the data link layer from sending a back-to-back frame immediately after the current frame which is being received (step 866). (This may be achieved by setting a carrier-sense signal to “true,” based on the IEEE 802.3 standards. This carrier-sense signal, when set to “true,” allows the current data frame to finish transmission from the data link layer, and prevents the data link layer from sending the next data frame.) The system then determines whether the laser is on (step 870). If the laser is off, the system turns on the laser (step 872). Otherwise, the system removes the first word at the head of the FIFO buffer (step 874). Next, the system shifts the words in the FIFO buffer towards the head of the buffer (step 876), and appends the received data word to the tail of the FIFO buffer (step 878). The system then prepares to receive the next word from the data link layer (step 860).

[0066] If the received word is an idle word, the system first increments the idle length counter by 1 (step 880). The system then determines whether the number of consecutive idle words (idle length) is equal to or greater than a minimum frame gap (step 882). (Note that the minimum frame gap ideally includes the space for FEC parity bits (if FEC is used) and the default IPG space.) If so, the system allows the data link layer to send the next frame (step 884). (This may be achieved by setting a carrier-sense signal to “false,” based on the IEEE 802.3 standards.) Otherwise, the system determines whether the laser is currently on, and whether the idle length is greater than a delay bound, which is defined as the sum of laser turn-on period and the synchronization time (step 886). If both conditions are true, the system turns off the

laser (step 88), and removes the first word in the FIFO buffer (step 874). Otherwise, the system does not change the current state of the laser, and directly proceeds to remove the first word in the FIFO buffer (step 874).

The foregoing descriptions of embodiments of the present invention have been 5 presented for purposes of illustration and description only. They are not intended to be exhaustive or to limit the present invention to the forms disclosed. Accordingly, many modifications and variations will be apparent to practitioners skilled in the art. Additionally, the above disclosure is not intended to limit the present invention. The scope of the present invention is defined by the appended claims.

What Is Claimed Is:

1 1. A method for controlling transmission in an Ethernet passive optical
2 network (EPON) which includes a central node and at least one remote node, wherein
3 a remote node implements a data-link layer and a physical layer, the method
4 comprising:

5 receiving, at the physical layer of a remote node, a word which is
6 communicated from the data-link layer of the remote node, wherein the word may be a
7 data word or an idle word;

8 delaying the word for a pre-determined amount of time before allowing the
9 word to be transmitted by a transmitter, thereby providing time for turning the
0 transmitter on or off; and

11 turning the transmitter on or off based on the content of the received words.

1 2. The method of claim 0, wherein turning the transmitter on or
2 off based on the content of the received words involves turning on the
3 transmitter if the transmitter is currently powered off and if the received word
4 is a data word.

1 3. The method of claim 0, wherein turning the transmitter on or
2 off based on the content of the received words involves turning off the
3 transmitter if the transmitter is currently powered on and if the number of
4 consecutive received idle words is equal to or greater than a pre-determined
5 number.

1 4. The method of claim 0, wherein delaying the word for a pre-
2 determined amount of time involves:
3 buffering the word in a first-in-first-out (FIFO) queue;
4 shifting each word buffered in the FIFO queue towards the head of the
5 FIFO queue at regular time intervals; and

removing a word at the head of the FIFO queue at regular time intervals.

1 5. The method of claim 0, further comprising determining the
2 amount of time for which the word is delayed based on a laser turn-on time, an
3 automatic gain control (AGC) time, and a data and clock recovery (CDR) time.

1 6. The method of claim 0, wherein if the received word is a data
2 word, the method further comprises temporarily prohibiting the data link layer
3 from sending another data frame immediately after the current data frame
4 which is being received at the physical layer, thereby allowing a sufficient
5 number of idle words to be inserted between data frames.

1 7. The method of claim 0, wherein if there are a sufficient number
2 of consecutive idle words received at the physical layer, the method further
3 comprises allowing the data link layer to send another data frame to the
4 physical layer.

1 8. An apparatus for controlling transmission in an Ethernet
2 passive optical network (EPON) which includes a central node and at least one
3 remote node, wherein a remote node implements a data-link layer and a
4 physical layer, the apparatus comprising:

5 a receiving mechanism configured to receive, at the physical layer of a
6 remote node, a word which is communicated from the data-link layer of the
7 remote node, wherein the word may be a data word or an idle word;

8 a delay mechanism configured to delay the word for a pre-determined
9 amount of time before allowing the word to be transmitted by a transmitter,
10 thereby providing time for turning the transmitter on or off; and

11 a control mechanism configured to turn the transmitter on or off based
12 on the content of the received words.

1 9. The apparatus of claim 8, wherein while turning the transmitter
2 on or off based on the content of the received words, the control mechanism is
3 configured to turn on the transmitter if the transmitter is currently powered off
4 and if the received word is a data word.

1 10. The apparatus of claim 8, wherein while turning the transmitter
2 on or off based on the content of the received words, the control mechanism is
3 configured to turn off the transmitter if the transmitter is currently powered on
4 and if the number of consecutive received idle words is equal to or greater than
5 a pre-determined number.

1 11. The apparatus of claim 8, wherein while delaying the word for a
2 pre-determined amount of time, the delay mechanism is configured to:
3 buffer the word in a first-in-first-out (FIFO) queue;
4 shift each word buffered in the FIFO queue towards the head of the
5 FIFO queue at regular time intervals; and
6 to remove a word at the head of the FIFO queue at regular time
7 intervals.

1 12. The apparatus of claim 8, wherein the delay mechanism is
2 further configured to determine the amount of time for which the word is
3 delayed based on a laser turn-on time, an automatic gain control (AGC) time,
4 and a data and clock recovery (CDR) time.

1 13. The apparatus of claim 8, wherein if the received word is a data
2 word, the delay mechanism is further configured to temporarily prohibit the
3 data link layer from sending another data frame immediately after the current
4 data frame which is being received at the physical layer, thereby allowing a
5 sufficient number of idle words to be inserted between data frames.

1 14. The apparatus of claim 8, wherein if there are a sufficient
2 number of consecutive idle words received at the physical layer, the delay
3 mechanism is further configured to allow the data link layer to send another
4 data frame to the physical layer.

1 15. A computer-readable storage medium storing instructions that
2 when executed by a computer cause the computer to perform a method for
3 controlling transmission in an Ethernet passive optical network (EPON) which
4 includes a central node and at least one remote node, wherein a remote node
5 implements a data-link layer and a physical layer, the method comprising:

6 receiving, at the physical layer of a remote node, a word which is
7 communicated from the data-link layer of the remote node, wherein the word
8 may be a data word or an idle word;

9 delaying the word for a pre-determined amount of time before allowing
10 the word to be transmitted by a transmitter, thereby providing time for turning
11 the transmitter on or off; and

12 turning the transmitter on or off based on the content of the received
13 words.

1 16. The computer-readable storage medium of claim 15, wherein
2 turning the transmitter on or off based on the content of the received words
3 involves turning on the transmitter if the transmitter is currently powered off
4 and if the received word is a data word.

1 17. The computer-readable storage medium of claim 15, wherein
2 turning the transmitter on or off based on the content of the received words
3 involves turning off the transmitter if the transmitter is currently powered on
4 and if the number of consecutive received idle words is equal to or greater than
5 a pre-determined number.

1 18. The computer-readable storage medium of claim 15, wherein
2 delaying the word for a pre-determined amount of time involves:
3 buffering the word in a first-in-first-out (FIFO) queue;
4 shifting each word buffered in the FIFO queue towards the head of the
5 FIFO queue at regular time intervals; and
6 removing a word at the head of the FIFO queue at regular time
7 intervals.

1 19. The computer-readable storage medium of claim 15, wherein
2 the method further comprises determining the amount of time for which the
3 word is delayed based on a laser turn-on time, an automatic gain control
4 (AGC) time, and a data and clock recovery (CDR) time.

1 20. The computer-readable storage medium of claim 15, wherein if
2 the received word is a data word, the method further comprises temporarily
3 prohibiting the data link layer from sending another data frame immediately
4 after the current data frame which is being received at the physical layer,
5 thereby allowing a sufficient number of idle words to be inserted between data
6 frames.

1 21. The computer-readable storage medium of claim 15, wherein if
2 there are a sufficient number of consecutive idle words received at the physical
3 layer, the method further comprises allowing the data link layer to send
4 another data frame to the physical layer.

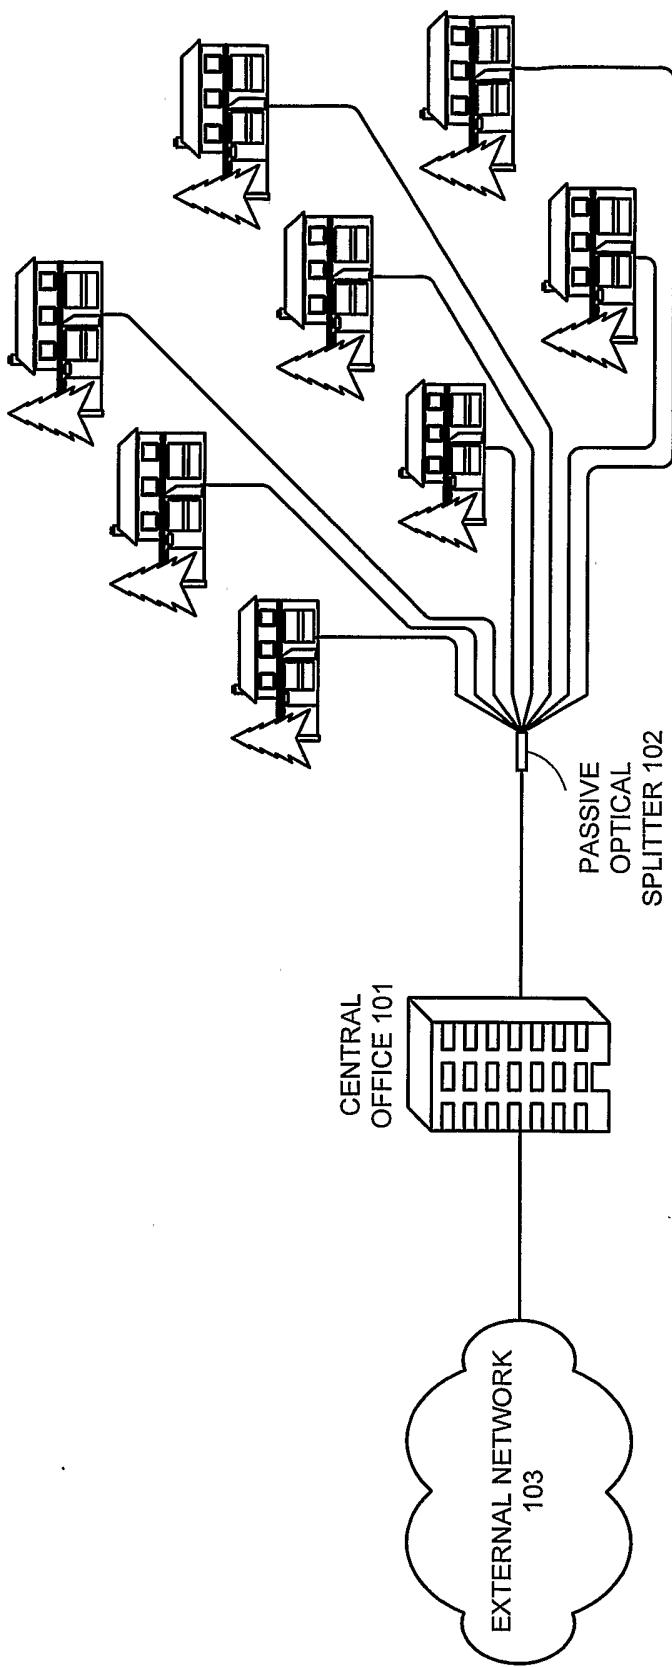


FIG. 1 (PRIOR ART)

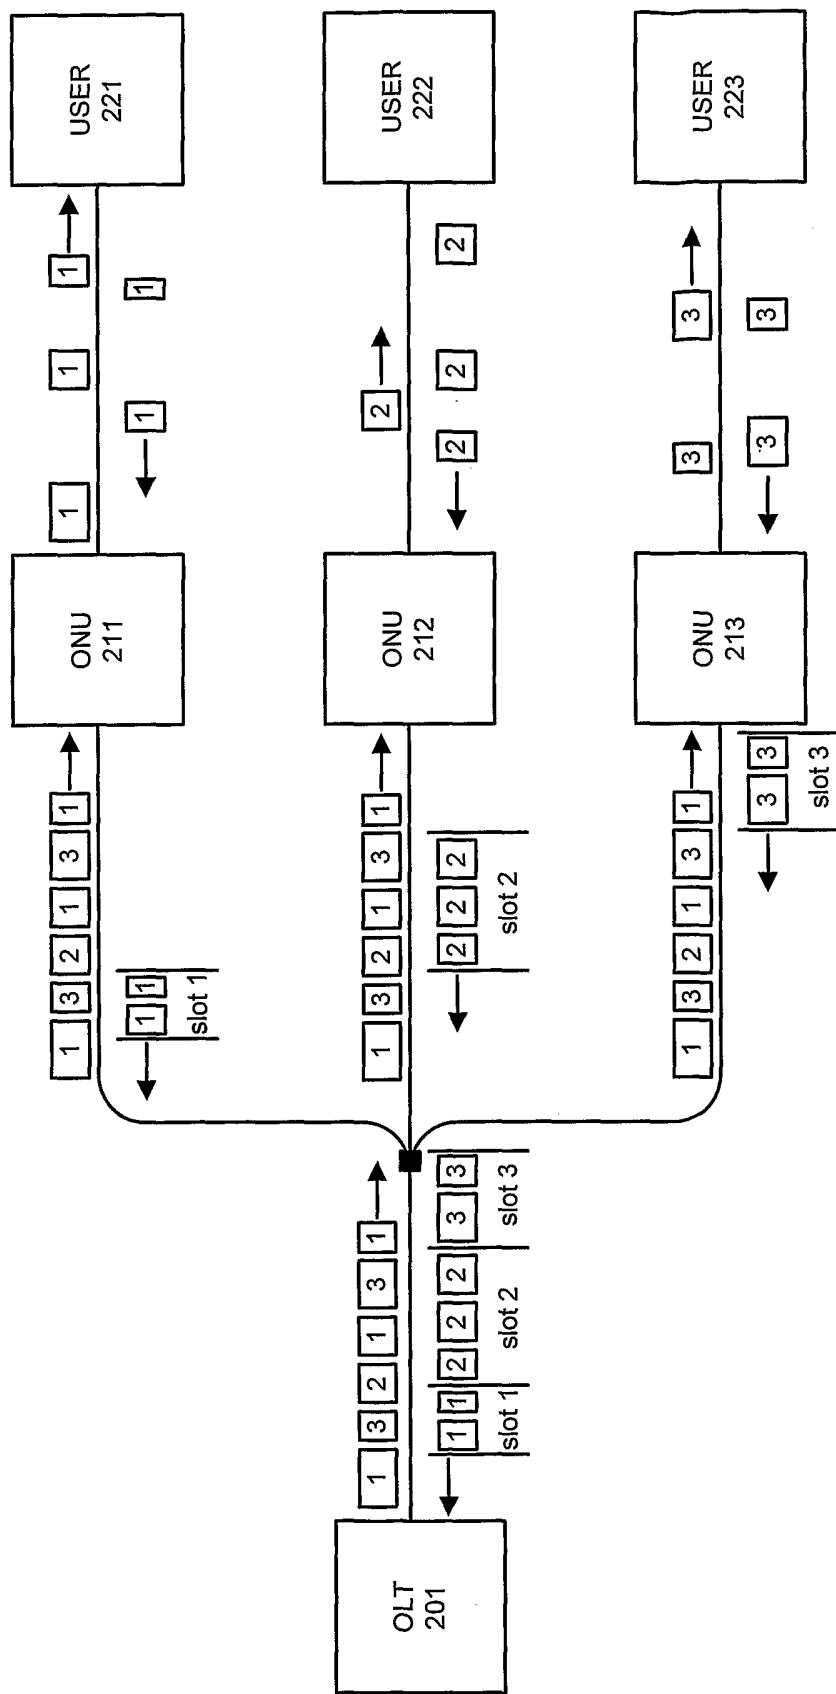


FIG. 2 (PRIOR ART)

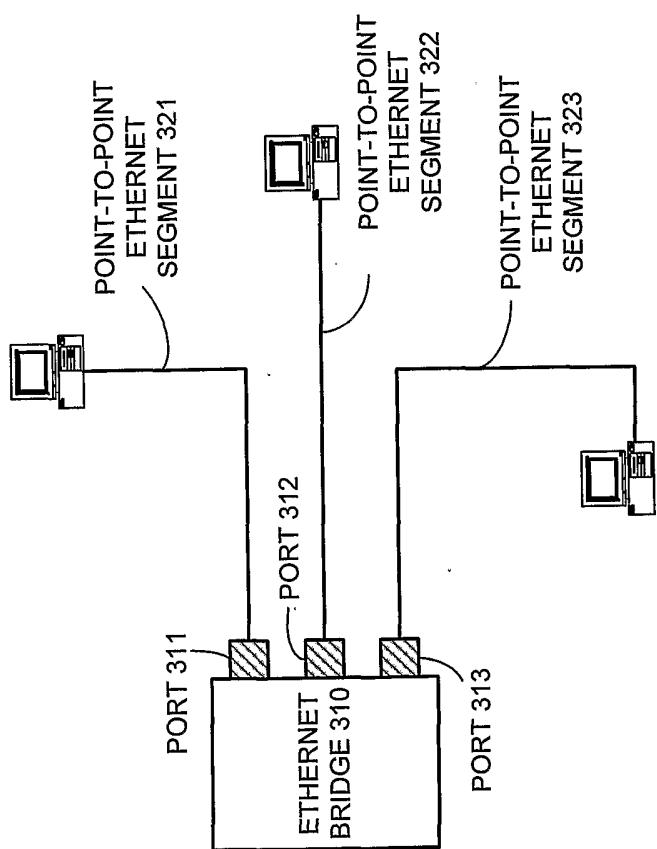
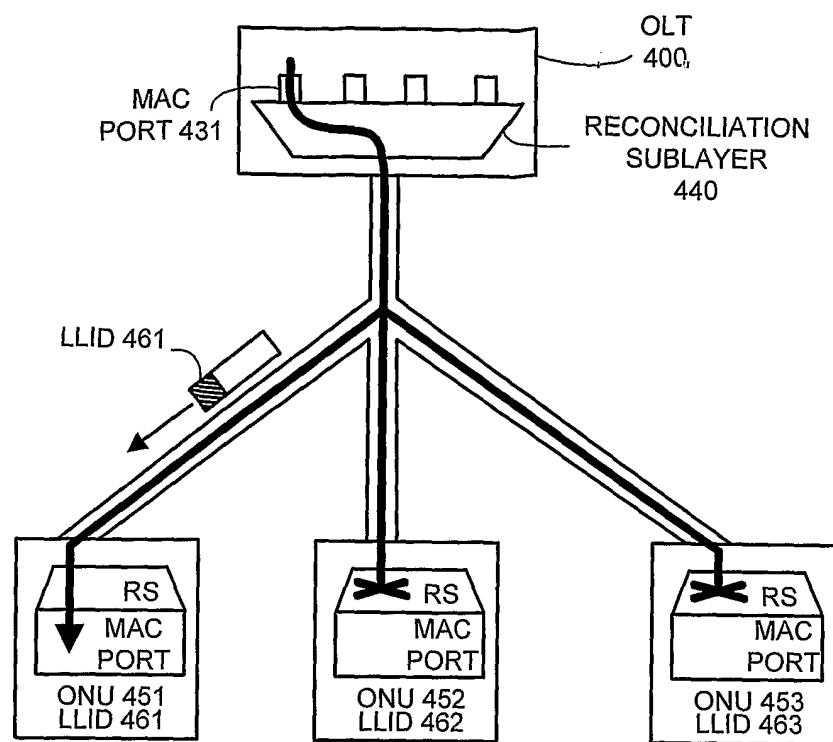
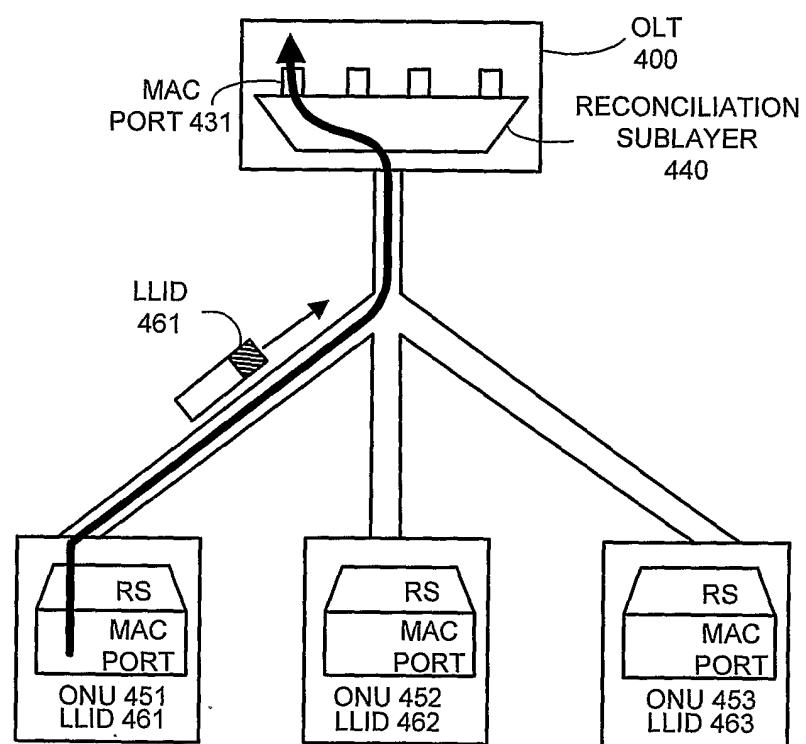


FIG. 3 (PRIOR ART)

4/9

**FIG. 4A (PRIOR ART)**

6/8

**FIG. 4B (PRIOR ART)**

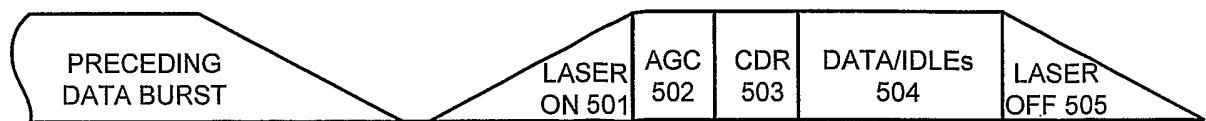


FIG. 5A (PRIOR ART)

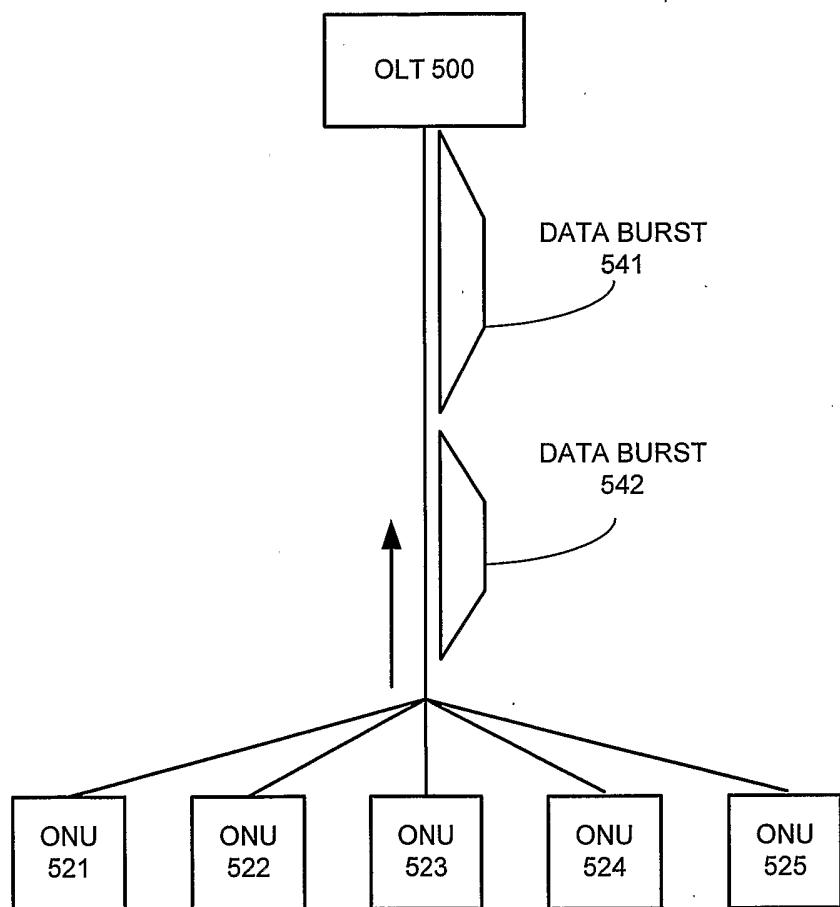


FIG. 5B (PRIOR ART)

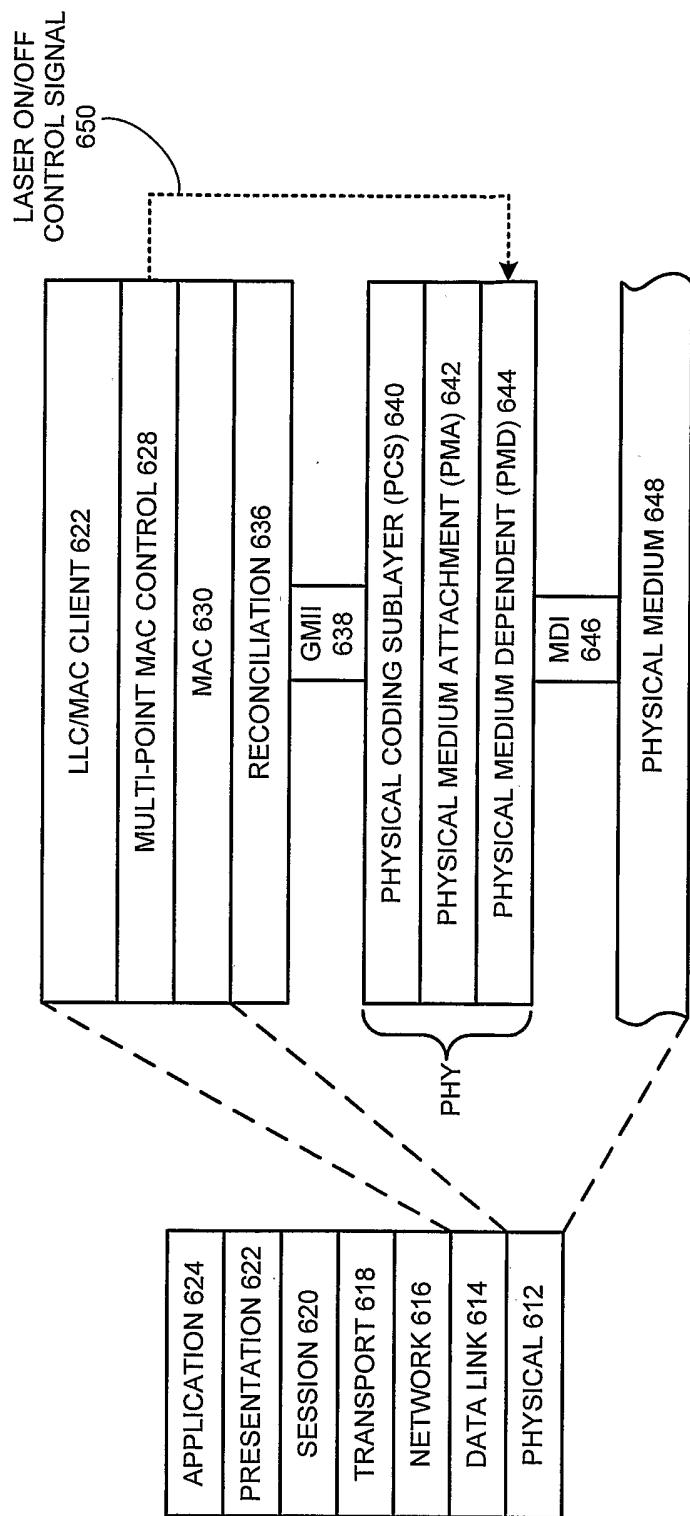


FIG. 6 (PRIOR ART)

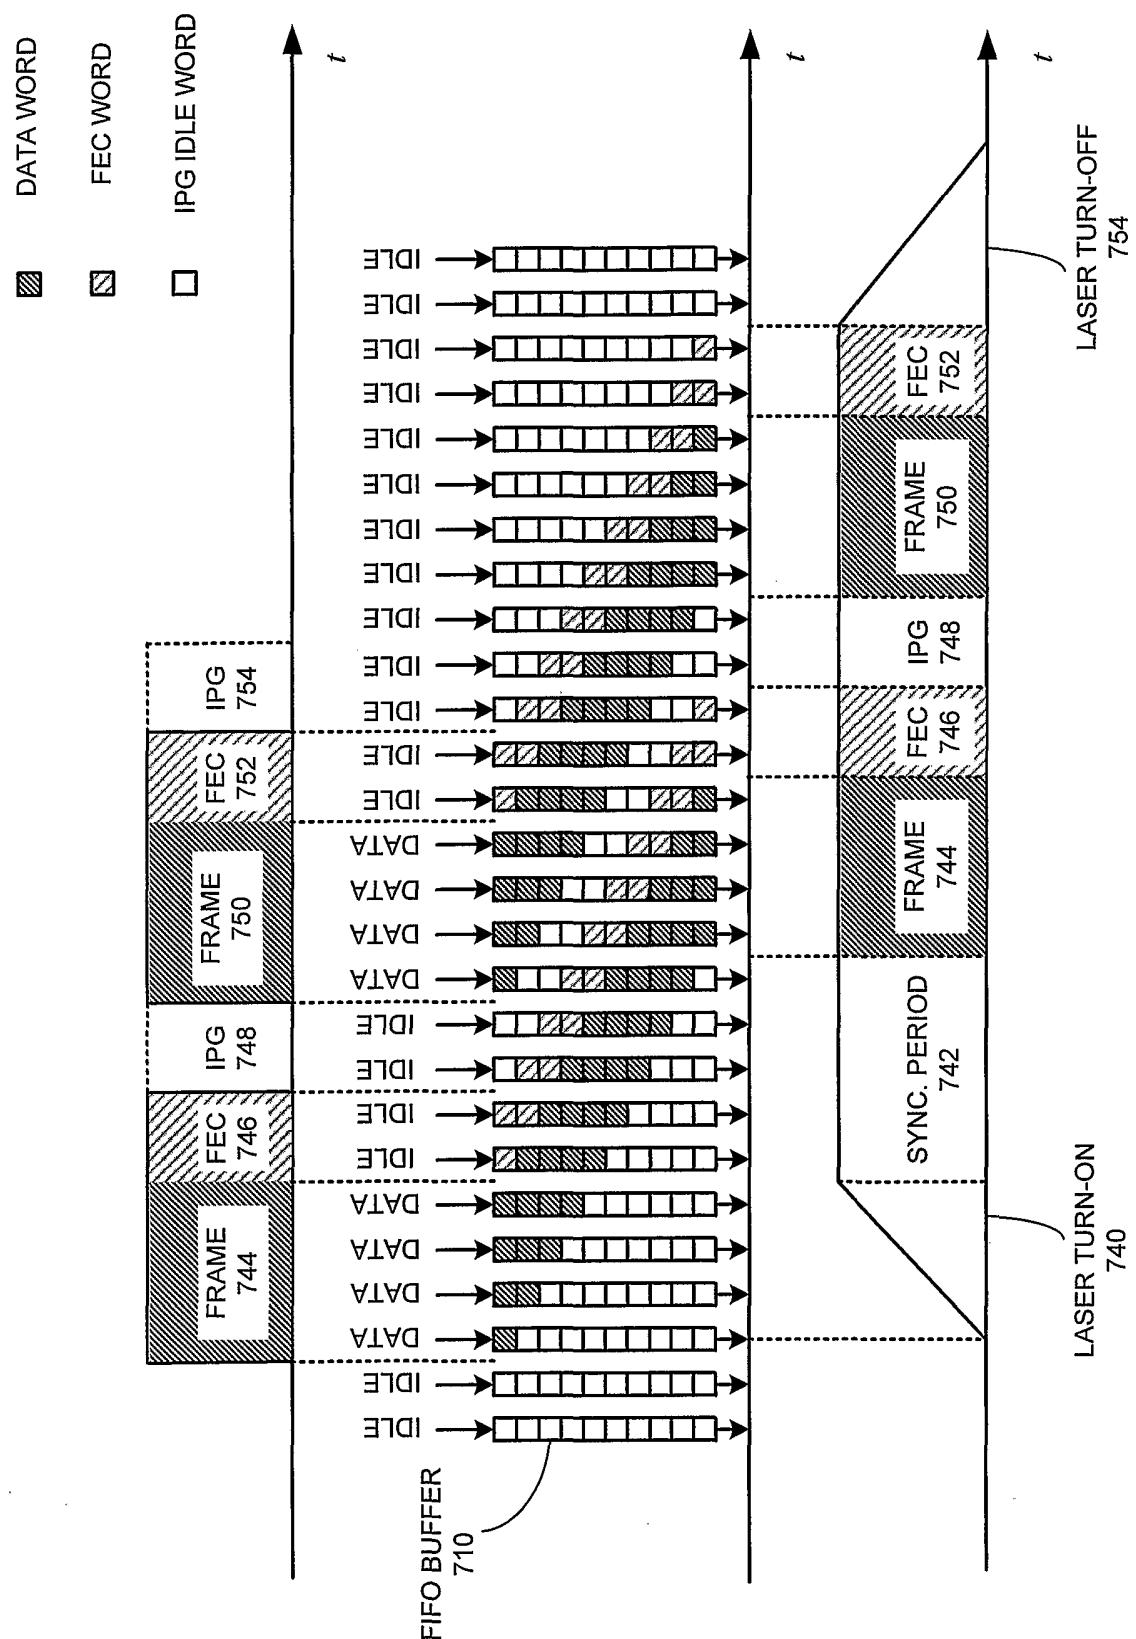


FIG. 7

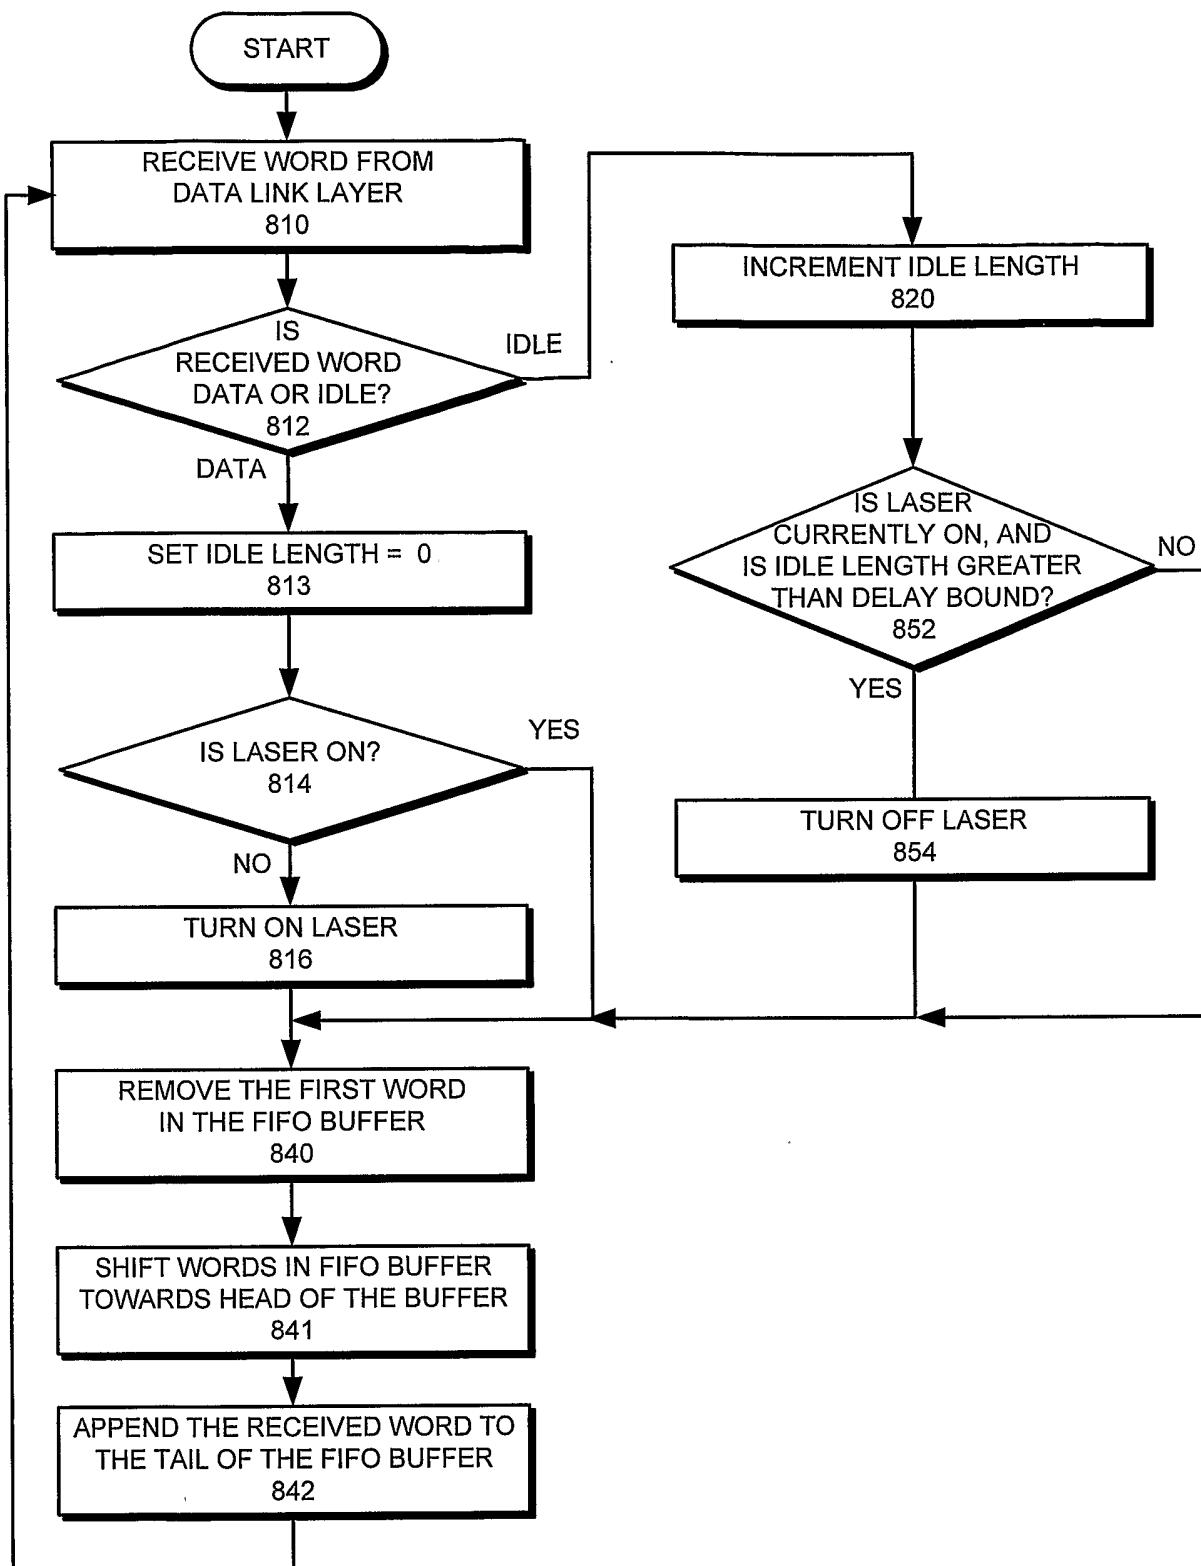


FIG. 8A

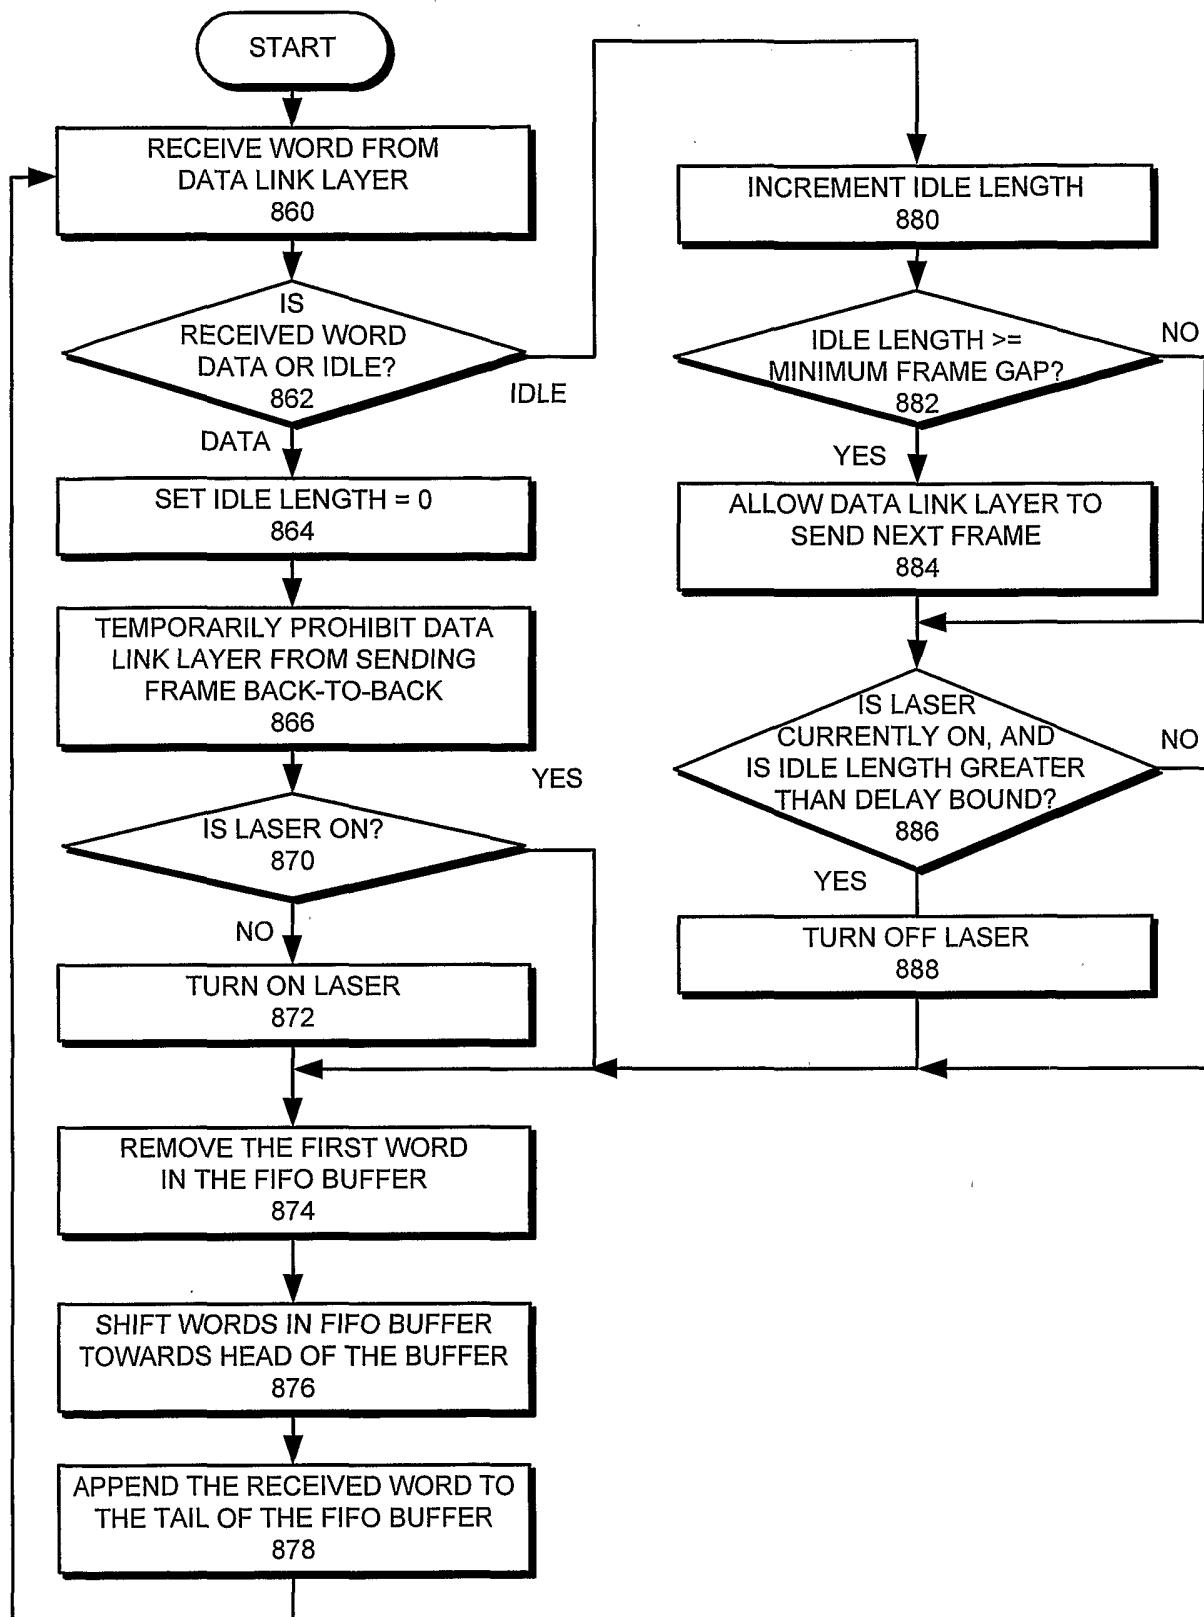


FIG. 8B

INTERNATIONAL SEARCH REPORT

International Application No

PCT/US2004/030365

A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 H04Q11/00

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 7 H04Q H04B H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category °	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	EP 0 765 045 A (ALCATEL BELL NV) 26 March 1997 (1997-03-26) column 3, line 51 - column 4, line 31	1, 8, 15
A		2-7, 9-14, 16-21
A	KRAMER G ET AL: "ETHERNET PASSIVE OPTICAL NETWORK (EPON): BUILDING A NEXT-GENERATION OPTICAL ACCESS NETWORK" IEEE COMMUNICATIONS MAGAZINE, IEEE SERVICE CENTER. PISCATAWAY, N.J, US, vol. 40, no. 2, February 2002 (2002-02), pages 66-73, XP001125420 ISSN: 0163-6804 page 71, left-hand column, line 40 - right-hand column, line 8	1-21

Further documents are listed in the continuation of box C.

Patent family members are listed in annex.

° Special categories of cited documents:

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier document but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

Date of the actual completion of the international search

Date of mailing of the international search report

12 January 2005

18/01/2005

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Meurisse, W

INTERNATIONAL SEARCH REPORT

International Application No
PCT/US2004/030365

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	EP 1 333 612 A (SAMSUNG ELECTRONICS CO LTD) 6 August 2003 (2003-08-06) claim 1 -----	1-21
A	EP 1 091 525 A (OKI ELECTRIC IND CO LTD) 11 April 2001 (2001-04-11) figure 3 -----	1-21

INTERNATIONAL SEARCH REPORT
Information on patent family members

Int	national Application No
PCT/US2004/030365	

Patent document cited in search report	Publication date		Patent family member(s)		Publication date
EP 0765045	A 26-03-1997	EP	0765045 A1		26-03-1997
		AU	714781 B2		13-01-2000
		AU	6554396 A		27-03-1997
		CA	2186149 A1		22-03-1997
		DE	69524751 D1		31-01-2002
		DE	69524751 T2		22-08-2002
		ES	2167399 T3		16-05-2002
		JP	9181686 A		11-07-1997
		US	5754319 A		19-05-1998
EP 1333612	A 06-08-2003	KR	2003065842 A		09-08-2003
		EP	1333612 A2		06-08-2003
		JP	2003283522 A		03-10-2003
		US	2003147654 A1		07-08-2003
EP 1091525	A 11-04-2001	JP	2001111554 A		20-04-2001
		EP	1091525 A2		11-04-2001
		US	6735211 B1		11-05-2004