



(86) Date de dépôt PCT/PCT Filing Date: 2004/11/22  
 (87) Date publication PCT/PCT Publication Date: 2005/06/09  
 (45) Date de délivrance/Issue Date: 2013/11/12  
 (85) Entrée phase nationale/National Entry: 2006/05/17  
 (86) N° demande PCT/PCT Application No.: US 2004/039345  
 (87) N° publication PCT/PCT Publication No.: 2005/053208  
 (30) Priorité/Priority: 2003/11/24 (US10/720,742)

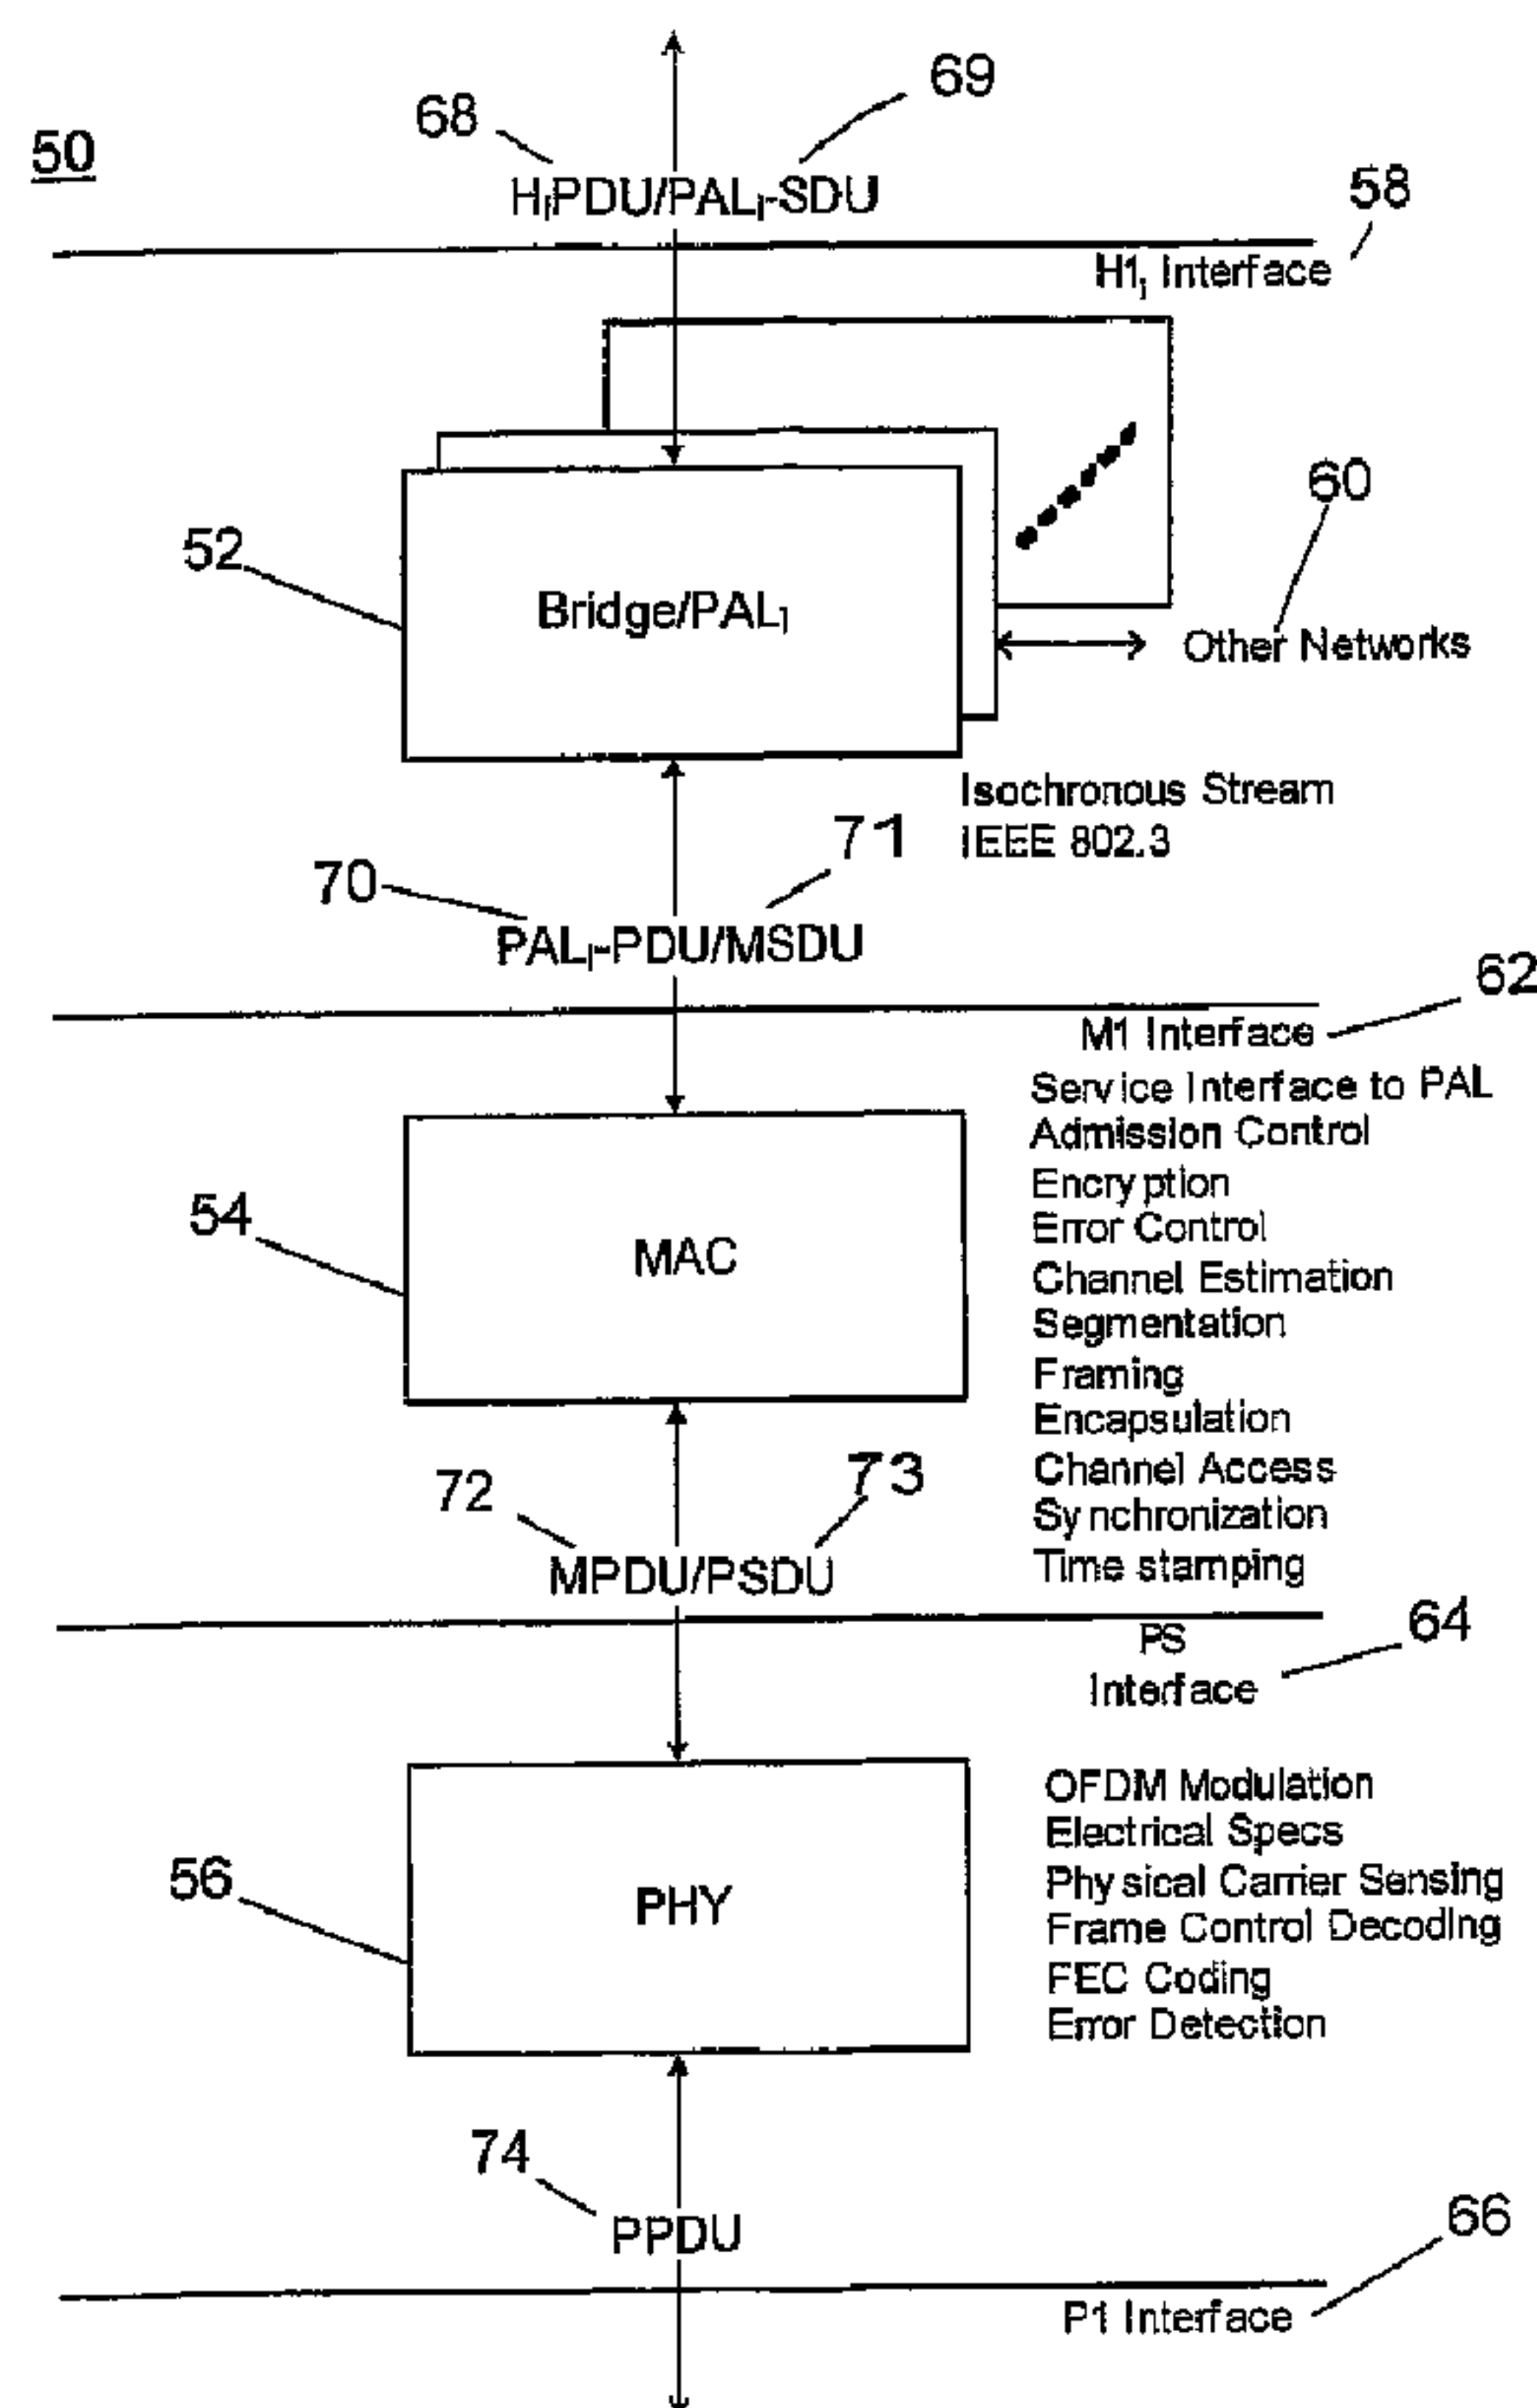
(51) Cl.Int./Int.Cl. *G06F 15/173* (2006.01)

(72) Inventeurs/Inventors:  
 YONGE, LAWRENCE W., III, US;  
 KATAR, SRINIVAS, US;  
 KOSTOFF, STANLEY J., II, US;  
 EARNSHAW, WILLIAM E., US;  
 BLANCHARD, BART W., US;  
 GARGRAVE, TIMOTHY R., US

(73) Propriétaire/Owner:  
 QUALCOMM INCORPORATED, US

(74) Agent: SMART & BIGGAR

(54) Titre : COUCHE DE CONTROLE D'ACCES DE MOYEN ENCAPSULANT LES DONNEES D'UNE PLURALITE D'UNITES DE DONNEES RECUES DANS UNE PLURALITE DE BLOCS TRANSMISSIBLES INDEPENDAMMENT  
 (54) Title: MEDIUM ACCESS CONTROL LAYER THAT ENCAPSULATES DATA FROM A PLURALITY OF RECEIVED DATA UNITS INTO A PLURALITY OF INDEPENDENTLY TRANSMITTABLE BLOCKS



Reference Network Architecture

(57) Abrégé/Abstract:

A method of operating in a network in which a plurality of stations communicate over a shared medium, comprising providing a physical layer (e.g., PHY) for handling physical communication over the shared medium; providing a high level layer (e.g., PAL) that



(57) **Abrégé(suite)/Abstract(continued):**

receives data from the station and supplies high level data units (e.g., MSDUs) for transmission over the medium; providing a MAC layer that receives the high level data units from the high level layer and supplies low level data units (e.g., MPDUs) to the physical layer; at the MAC layer, encapsulating content from a plurality of the high level data units; dividing the encapsulated content into a plurality of pieces (e.g., segments) with each piece capable of being independently retransmitted; and supplying low level data units containing one or more of the plurality of pieces.

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

(19) World Intellectual Property  
Organization  
International Bureau



(43) International Publication Date  
9 June 2005 (09.06.2005)

PCT

(10) International Publication Number  
WO 2005/053208 A3

(51) International Patent Classification<sup>7</sup>: G06F 15/173

(21) International Application Number:

PCT/US2004/039345

(22) International Filing Date:

22 November 2004 (22.11.2004)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

10/720,742 24 November 2003 (24.11.2003) US

(63) Related by continuation (CON) or continuation-in-part (CIP) to earlier application:

US 10/720,742 (CON)

Filed on 24 November 2003 (24.11.2003)

(71) Applicant (for all designated States except US): INTELLELLON CORPORATION [US/US]; 5100 West Silver Springs Boulevard, Ocala, FL 34482 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): YONGE, Lawrence, W., III [US/US]; 8380 Juniper Road, Ocala, FL 34480 (US). KATAR, Srinivas [IN/US]; 5001 SW 20th Street, Apartment 2607, Ocala, FL 34474 (US). KOSTOFF, Stanley, J., II [US/US]; 13 Needles Drive, Ocala, FL 34482 (US). EARNSHAW, William, E. [US/US]; 48 NE 56th Terrace, Ocala, FL 34470 (US). BLANCHARD, Bart, W. [US/US]; 22278 112th Terrace, Fort McCoy, FL 32134 (US). GARGRAVE, Timothy, R. [US/US]; 5021 North East 7 Place, Ocala, FL 34470 (US).

(74) Agent: LEE, Roger, G.; Fish & Richardson P.C., 225 Franklin Street, Boston, MA 02111-2804 (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 (patent), 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, IS, 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).

**Declarations under Rule 4.17:**

— as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii)) for the following designations 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, UZ, VC, VN, YU, ZA, ZM, ZW, ARIPO patent (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LU, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG)

— as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii)) for the following designations 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, UZ, VC, VN, YU, ZA, ZM, ZW, ARIPO patent (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LU, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG)

**Published:**

— with international search report

[Continued on next page]

(54) Title: METHOD OF OPERATING IN A NETWORK IN WHICH A PLURALITY OF STATIONS COMMUNICATE OVER A SHARED MEDIUM

(57) Abstract: A method of operating in a network in which a plurality of stations communicate over a shared medium, comprising providing a physical layer (e.g., PHY) for handling physical communication over the shared medium; providing a high level layer (e.g., PAL) that receives data from the station and supplies high level data units (e.g., MSDUs) for transmission over the medium; providing a MAC layer that receives the high level data units from the high level layer and supplies low level data units (e.g., MPDUs) to the physical layer; at the MAC layer, encapsulating content from a plurality of the high level data units; dividing the encapsulated content into a plurality of pieces (e.g., segments) with each piece capable of being independently retransmitted; and supplying low level data units containing one or more of the plurality of pieces.

WO 2005/053208 A3

**WO 2005/053208 A3**



---

**(88) Date of publication of the international search report:**  
10 November 2005

*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.*

**MEDIUM ACCESS CONTROL LAYER THAT ENCAPSULATES DATA  
FROM A PLURALITY OF RECEIVED DATA UNITS INTO A PLURALITY  
OF INDEPENDENTLY TRANSMITTABLE BLOCKS**

**TECHNICAL FIELD**

5           This invention relates to network protocols, and more particularly to medium access control layers that encapsulate data from a plurality of received data units.

**BACKGROUND**

Networking protocols are normally developed in layers, with each layer responsible for a different facet for the communication. Layers exchange structured  
10 information. Each layer receives Service Data Units (SDUs) from higher layers, which are processed to generate Protocol Data Units (PDUs). Protocol Data Units are handed over to the lower layers for service. Similarly, the PDUs received from the lower layers are processed to generate SDUs, which are handed over to the higher layers. PDUs not only carry the SDUs but also carry management information that is relevant for  
15 managing the layer functionality. Defining the structure of SDUs and PDUs for a given protocol layer is critical to enable proper layer functionality. Some examples of network protocol layers include the well-known Transmission Control Protocol (TCP) and Internet Protocol (IP). The structure of TCP data units has provisions to enable end-to-end delivery. The structure of IP data units enables efficient routing.

20           Networks use medium access control layer (MAC) to enable coordinated access to the medium. Medium access layer uses the functionality of the physical layer (PHY) to provide services to the higher layer. MAC service to the higher layers can include guarantees on Quality of Service (QoS). QoS provides guarantees on bandwidth, latency, jitter and packet loss probability for traffic streams. Jitter refers to deviation in  
25 the time of delivery of data over the network.

**SUMMARY**

In general, the invention features a method of operating in a network in which a plurality of stations communicate over a shared medium, comprising providing a physical layer (e.g., PHY) for handling physical communication over the shared  
30 medium; providing a high level layer (e.g., PAL) that receives data from the station and supplies high level data units (e.g., MSDUs) for transmission over the medium;

54830-2

providing a MAC layer that receives the high level data units from the high level layer and supplies low level data units (e.g., MPDUs) to the physical layer; at the MAC layer, encapsulating content from a plurality of the high level data units; dividing the encapsulated content into a plurality of pieces (e.g., segments) with each piece capable of being  
5 independently retransmitted; and supplying low level data units containing one or more of the plurality of pieces.

According to one aspect of the present invention, there is provided a method of operating in a network in which a plurality of stations communicate over a shared medium, the method comprising providing a physical layer for handling physical communication over  
10 the shared medium; providing a high level layer that receives data from the station and supplies high level data units for transmission over the medium; providing a medium access control (MAC) layer that receives the high level data units from the high level layer and supplies low level data units to the physical layer; at the MAC layer, encapsulating content from a plurality of the high level data units into a stream of sub-frames; dividing the  
15 encapsulated stream into a plurality of pieces with each piece capable of being independently retransmitted; and supplying low level data units, at least some of the low level data units each containing a plurality of pieces into which the encapsulated stream was divided, and at least some of the low level data units containing boundary demarcation information indicating boundaries between the sub-frames in the stream; wherein at least some information common  
20 to the high level data units is not repeated for each high level data unit encapsulated in the stream; wherein the high level data units each comprise a payload, and encapsulating comprises forming the stream comprising the payloads from a succession of high level data units; wherein the stream comprises a succession of sub-frames, each sub-frame comprising a header and a plurality of the payloads; wherein each sub-frame is divided into the plurality of  
25 pieces capable of being independently retransmitted; and wherein division of a sub-frame into the plurality of pieces comprises dividing the sub-frame into a plurality of sub-blocks, and forming at least some pieces from a plurality of sub-blocks.

54830-2

According to another aspect of the present invention, there is provided a method of operating in a network in which a plurality of stations communicate over a shared medium, comprising providing a physical layer for handling physical communication over the shared medium; providing a high level layer that receives data from a station and supplies  
5 high level data units for transmission over the medium; providing a medium access control (MAC) layer that receives the high level data units from the high level layer and supplies low level data units to the physical layer; at the MAC layer, encapsulating content from a plurality of the high level data units into a stream of sub-frames; dividing the encapsulated stream into a plurality of pieces with each piece capable of being independently retransmitted; and  
10 supplying low level data units, at least some of the low level data units each containing a plurality of the pieces into which the encapsulated stream was divided, and at least some of the low level data units containing boundary demarcation information indicating boundaries between the sub-frames in the stream; wherein at least some information common to the high level data units is not repeated for each high level data unit encapsulated in the stream;  
15 wherein the high level data units each comprise a payload, and encapsulating comprises forming the stream comprising the payloads from a succession of high level data units; wherein the stream comprises a succession of sub-frames, each sub-frame comprising a header and a plurality of the payloads; and wherein the stream is divided into a plurality of sub-blocks, wherein a plurality of sub-blocks are grouped to form a segment, with a segment  
20 crossing sub-frame boundaries in the stream, wherein a segment constitutes one of the pieces.

According to still another aspect of the present invention, there is provided a method of operating in a network in which a plurality of stations communicate over a shared medium, comprising providing a physical layer for handling physical communication over the shared medium; providing a high level layer that receives data from a station and supplies  
25 high level data units; and providing a layer that receives the high level data units from the high level layer and supplies low level data units to the physical layer; wherein supplying the low level data units comprises encapsulating content from a plurality of the high level data units into a stream of sub-frames, dividing the encapsulated stream into a plurality of sub-blocks, forming a plurality of pieces, with each piece including one or more sub-blocks,

54830-2

to provide pieces capable of being independently retransmitted, and supplying low level data units, at least some of the low level data units each containing a plurality of the pieces that include the sub-blocks into which the encapsulated content was divided, and at least some of the low level data units containing boundary demarcation information indicating boundaries  
5 between the sub-frames in the stream.

According to yet another aspect of the present invention, there is provided a method of operating in a network in which a plurality of stations communicate over a powerline communications medium, comprising providing a physical layer for handling physical communication over the powerline communications medium; providing a high level  
10 layer that receives data from the station and supplies high level data units for transmission over the powerline communications medium; providing a medium access control (MAC) layer that receives the high level data units from the high level layer and supplies low level data units to the physical layer; at the MAC layer, encapsulating content from a plurality of the high level data units into a stream of sub-frames; dividing the encapsulated content into a  
15 plurality of pieces with each piece capable of being independently retransmitted; and supplying low level data units containing the plurality of the pieces.

According to a further aspect of the present invention, there is provided a system in which a plurality of stations communicate over a powerline communications medium, comprising: one station of the plurality of stations configured to operate according to  
20 a protocol that includes: a physical layer configured to handle physical communication over the powerline communications medium, a high level layer configured to receive data from the one station and further configured to provide high level data units for transmission over the powerline communications medium, and a medium access control (MAC) layer configured to receive the high level data units from the high level layer and further configured to provide  
25 low level data units to the physical layer; wherein the one station is further configured to: encapsulate content from a plurality of the high level data units into a stream of sub-frames at the MAC layer of the one station; divide the encapsulated content into a plurality of pieces

54830-2

with each piece being capable of independent retransmission; and supply low level data units containing the plurality of the pieces.

Preferred implementations of the invention may include one or more of the following. At least some information common to the encapsulated high level data units may  
5 not be repeated for each high level data unit encapsulated in a low level data unit. The information common to the encapsulated high level data units may comprise destination and source addresses. The high level data units may each comprise a payload, and encapsulating may comprise forming a queue comprising the payloads from a succession of high level data. The queue may comprise a succession of sub-frames, each sub-frame comprising a header and  
10 a plurality of payloads. Each sub-frame may be divided into the plurality of pieces capable of being independently retransmitted. Division of a sub-frame into the plurality of pieces may comprise dividing the sub-frame into a plurality of sub-blocks, and forming at least some pieces from a plurality of sub-blocks. Each piece may constitute a segment that is transmitted as a physical layer block. The invention may further comprise parity pieces derived from  
15 other pieces and capable of being used at a destination to recover one or more lost pieces at the destination without having to retransmit the lost pieces. Each piece may be transmitted as a physical layer block, and the parity pieces may also be transmitted as parity physical layer blocks. The physical layer blocks may be encoded using forward error correction. Some of the pieces making up a low level data unit may constitute retransmitted pieces that failed to be  
20 correctly transmitted in an earlier attempt. At least some retransmitted pieces may be transmitted with greater forward error correction. Each sub-frame may further comprise a delivery time stamp associated with at least some payloads. Clock information characterizing the time setting of a clock in a transmitting station may be transmitted to a receiving station within a header of the low level data units, and the clock information may be used by the  
25 receiving station along with the delivery time stamps to establish the time at which payloads are delivered. The time at which a payload is delivered may be set to be substantially the time specified by the time stamp. The invention may further

comprise an integrity check value associated with each sub-frame or with a plurality of sub-frames. Each of the plurality of payloads in a sub-frame may have identical length. Each sub-frame may further comprise MAC management information. The MAC layer may have the capability of transmitting data in a plurality of sessions within a

5 regularly-repeated contention free interval, wherein a station to which data is transmitted may be identified by a destination address and a station from which data is transmitted may be identified by a source address, and wherein the queue may contain payloads for the same session, same source address, and same destination address.

The MAC layer may have the capability of transmitting data in a plurality of sessions

10 within a regularly-repeated contention free interval, wherein a station to which data is transmitted may be identified by a destination address and a station from which data is transmitted may be identified by a source address, and wherein the queue may contain sub-frames for the same session, same source address, and same destination address.

The sessions may be transmitted in a substantially contention-free manner. The

15 sessions may be transmitted within time slots of a regularly-repeated contention-free interval. A stream identifier (e.g., MSID) may be used to associate content of a queue with a particular session. The stream identifier may also be used to associate content of a queue with a priority level for contention-based transmission over the medium. There may be a plurality of queues, each containing payloads having a unique combination of

20 stream identifier, source address, and destination address. Each queue may contain a payload having a unique combination of stream identifier, source address, destination address, and type of high level layer. The queue may be divided into a plurality of sub-blocks, wherein a plurality of sub-blocks may be grouped to form a segment, with a segment crossing sub-frame boundaries in the queue, wherein a segment may constitute

25 one of the pieces. Each sub-block may be shorter than a sub-frame. At least some segments may contain a number of sub-blocks corresponding to other than an integral number of sub-frames. The sub-blocks may be of equal length. The sub-blocks may have an associated sequential numbering adapted for use at the receiving station for re-establishing the correct sequential order of the sub-blocks. The sub-blocks may have a

30 predetermined size, which combined with the associated sequential numbering, may eliminate the need for buffer reordering when out of order segments are received. The sub-blocks may be of equal size. The invention may further comprise, for at least some of the low level data units, forming the low level data unit from a plurality of

segments. Each segment in the low level data unit may form the body of a separate block transmitted by the physical layer. Individual segments may be individually encrypted. Encryption information common to a plurality of segments may be carried in a header. Some encryption information may be carried in a header and frame control  
5 of the low level data unit and in a header of the block. Some encryption information may be carried in frame control of the low level data unit and in a header of the block. Each block may separately undergo forward error correction, and forward error correction bits for each block may be transmitted in the low level data unit. The level of forward error correction used may be different for different blocks. The level of  
10 forward error correction used may provide greater error correction capability for selected blocks that are being retransmitted after failing to be correctly transmitted in an earlier attempt. Most of the blocks may be identical in length. The initial and final block of a low level data unit may be of a different length than the remaining blocks. Information common to the plurality of segments forming the low level data unit may  
15 be transmitted in a header for the low level data unit. The information common to the plurality of segments may be transmitted only in the header. The low level data unit may further comprise a frame control field.

In another aspect, the invention features a method of operating in a network in which a plurality of stations communicate over a shared medium, comprising providing  
20 a physical layer (e.g., PHY) for handling physical communication over the shared medium; providing a high level layer (e.g., PAL) that receives data from the station and supplies high level data units (e.g., MSDUs) for transmission over the medium; providing a MAC layer that receives the high level data units from the high level layer and supplies low level data units (e.g., MPDUs) to the physical layer; at the MAC layer,  
25 forming low level data units by encapsulating content from a plurality of the high level data units; and adaptively escalating the robustness of transmission of the low level data units depending on the frequency of transmission errors.

Preferred implementations of the invention may include one or more of the following. The invention may further comprise incorporating forward-error correction  
30 information into the transmitted stream of low level data units, and the step of adaptively escalating may comprise adaptively varying the forward-error correction information depending on the frequency of transmission errors. Varying the forward-error correction information may comprise varying one or both of the amount and type

of forward-error correction information. Decisions on adaptively escalating may be made at a transmitting station. The low level data units may comprise a plurality of pieces (e.g., segments). The forward error correction information may comprise information associated with provided with the pieces for use at a destination for recovering a piece that is received with errors. The forward error correction information may comprise parity pieces derived from other pieces and capable of being used at a destination to recover one or more lost pieces at the destination without having to retransmit the lost pieces. Each piece may be transmitted as a physical layer block, and the parity pieces may also be transmitted as parity physical layer blocks.

10           These and other embodiments may have one or more of the following advantages.

          The invention provides mechanisms to generate MAC protocol data units (MPDU) from the MAC Service data units (MSDU) in such a manner that enables efficient end-to-end delivery of packets. These mechanisms provide support to enhance Quality of Service (QoS) support and efficient delivery of management information. The format of the MPDU enables efficient retransmission of corrupted data and seamless integration with the underlying physical layer.

          Multiple higher layers of the networking protocols can be seamlessly interfaced with the MAC.

20           The MAC layer provides various Classes of service for application payloads. At the MAC layer, each Class encompasses a coherent set of Quality of Service (QoS) guarantees and can be translated naturally to such behavior in the MAC as channel access, number of retries, etc. This enables scalability and improved QoS guarantees. Supports both connection based and connection less service.

25           Mechanisms are provided to exchange MAC Management information between MAC layer and higher layers in a manner that would simplify implementation. Several types of MAC Management entities can be defined.

          Processing on the MSDUs reduces redundant information while maintaining functionality.

30           Transmission of management information is enabled in an in-band manner along with application data.

          Transmission of urgent MAC management information is enabled in an out-of band manner.

Efficient encryption of information is enabled to provide data privacy.

Testing of end-to-end delivery of MSDUs is enabled by means of a Integrity check vector (ICV).

A segmentation process enables maximum possible MPDUs to generated, thus  
5 increasing the MPDU efficiency.

There is a mapping of MPDU on to FEC Blocks at the PHY and the choice of FEC Block sizes enable efficient retransmission.

A MPDU header carries information common to all PBs, thus increasing MPDU efficiency

10 Transmission of MPDUs is enabled with low end-to-end jitter.

Bridging and forwarding of MSDUs are supported.

PHY error detection and correction by means of ARQ process is enabled.

An ARQ process is augmented by an Escalation mechanism and an outer erasure code, which enables improved guarantees on QoS parameters.

15 There is a simplified reassembly process with duplicate rejection capability.

These advantages are illustrated in the Detailed Description of the preferred embodiment that follows.

The details of one or more implementations of the invention are set forth in the accompanying drawings and the description below. Other features, objects, and  
20 advantages of the invention will be apparent from the description and drawings, and from the claims.

## DESCRIPTION OF DRAWINGS

FIG. 1 is a network configuration.

25 FIG. 2 is a reference network architecture.

FIG. 3 is a format for a MSDU.

FIG. 4 is a format for a Sub-Frame.

FIG. 5 is a format for a Sub Frame header.

FIG. 6 is a block of Sub-Frames protected by a single ICV.

30 FIG. 7 is a Sub-Frame generated from a MSDU Payload.

FIG. 8 is a Sub-Frame generated from multiple MSDU Payloads.

FIG. 9 is a MAC Encapsulation.

FIG. 10 is a MPDU generated from a Sub-Frame Stream.

FIG. 11 is a format of a MPDU Header.

FIG. 12 is a format for a PHY Block.

### DETAILED DESCRIPTION

5           There are a great many possible implementations of the invention, too many to describe herein. Some possible implementations that are presently preferred are described below. It cannot be emphasized too strongly, however, that these are descriptions of implementations of the invention, and not descriptions of the invention, which is not limited to the detailed implementations described in this section but is  
10 described in broader terms in the claims.

As shown in FIG. 1, network configuration 2 includes communications medium 3 and network 4 in which electronic devices 6, 8, and 10 (e.g., audiovisual equipment) communicate over medium 3. Electronic devices 6, 8, and 10 include media access controllers (MAC) 12, 14, and 16 that manage communication access to the network 4  
15 for electronic devices 6, 8, and 10, respectively. MACs 12, 14, and 16 implement the data link layer and connect to the physical layer (PHY) of the Open Systems Interconnection (OSI) network architecture standard. In a general sense, MACs 12, 14, and 16 represent stations on network 4 that send messages to one another over medium 3. Communications medium 3 is a physical communication link between electronic  
20 devices 6, 8, and 10 and may include optical fiber, coaxial cable, unshielded twisted pair, in addition to other media such as power lines. Electronic devices 6, 8, and 10 communicate with one another based on requirements of software applications running on electronic devices 6, 8, and 10. This communication creates traffic of messages on network 4.

25           FIG. 2 shows the major system interfaces and their associated data units for a portion of a reference network architecture 50 used by the network configuration 2. This portion may be implemented at each station. The abstract objects that make up the layers of a network system are sometimes called protocols. That is, a protocol provides a communication service that higher-level objects (such as application processes, or  
30 higher-level layers) use to exchange messages. Three layers of the network architecture are shown: Bridge/PAL<sub>i</sub> 52, MAC 54, and Physical layer (PHY) 56, separated by M1 Interface 62 and PS interface 64, respectively.

H1<sub>i</sub> **58** denotes the  $i^{\text{th}}$  Host Interface, with one interface for each protocol supported. The H1 interface **58** defines the point of demarcation for the  $i^{\text{th}}$  Host Protocol Data Units (H<sub>i</sub>PDU) **68** and the  $i^{\text{th}}$  Protocol Adaptation Layer Service Data Unit (PAL<sub>i</sub>SDU) **69** to higher layers of the network architecture **50**.

5 For each protocol supported, the corresponding Protocol Adaptation Layer (PAL) **52** may be implemented partially in host software and partially in firmware and/or hardware. Examples of architecture **50** support IEEE 802.3 and Isochronous Stream protocols as well as provide access to the proprietary protocols through interface **60**. The PAL **52** provides support for Higher Layer Adaptation (HLA) functionality and/or Bridging functionality. Both HLA and Bridging operations support translation of host data packets including PAL Protocol Data Units (PAL<sub>i</sub>PDU) **70** to MAC Service Data Units (MSDUs) **71** and vice versa, translation of host address from the H1 interface **58** to MAC **12, 14, 16** addresses. HLA and bridging operations also support determination of traffic classes and QoS parameters in addition to  
10  
15 Establishment of streams in coordination with the MAC **12, 14, 16**.

The PALs **52** also support address discovery and routing functions for bridging operations. Each PAL **52** provides binding and mapping from the stream identifiers provided by the MAC layer **54** at session setup time with the higher layer entities as necessary.

20 Each PAL **52** has an associated PAL Type (PLT) at the MAC layer **54**, to enable routing of the associated MAC Service Data Units (MSDUs) **71** at the receiver MAC (e.g., **12, 14, 16**). In addition, information about available overall channel bandwidth as well as available bandwidth for a specific class of traffic is provided by the MAC layer **54** to the PAL **52** to support rate adaptation.

25 The M1 interface **62** is common to all Protocol Adaptation Layers and defines the demarcation between the PAL **52** and the MAC layer **54**, with PAL Protocol Data Units (PAL<sub>i</sub>PDU) **70** being passed down from the PAL **52** to the MAC layer **54** as MAC Service Data Units (MSDUs) **72** and vice versa.

The Medium Access Control (MAC) layer **54** processes MAC Service Data  
30 Units (MSDUs) **71** from the PAL **52** and generates PHY Service Data Units (PSDU) **73** for delivery to the Physical Layer **56**. MAC layer **54** processing includes Service interface to PAL **52**, Network Management, Admission Control, Encryption, Error Control (ARQ), Retransmission, Escalation, Channel Estimation - Modulations, FEC,

etc., Tone Map as a function of time, Framing, Segmentation & Reassembly, Packet Encapsulation and De-encapsulation, Channel Access (Contention Free Bursting, managed sessions, CSMA/CA, etc.), Time Stamping, Synchronization - With Multimedia Clocks, and Contention Free Sessions.

5 The Physical Layer Signaling (PS) Interface 64 separates the MAC layer 54 and the PHY 56 with MAC Protocol Data Units (MPDUs) 72 being passed to the PHY 56 from the MAC layer 54 as PHY Service Data Units (PSDUs) 73 across the PS Interface 64 and vice versa.

The Physical Layer (PHY) 56 Protocol provides the following operations.  
10 Service interface to MAC layer 54, OFDM Modulation, Forward Error Correction Coding, Physical Carrier Sensing, Frame Control Decoding, Error detection, and information needed for channel estimation and tone map selection.

MSDUs 71 are received by the MAC (e.g., 12, 14, or 16) at the MAC layer 54 from higher layers of the network architecture 50. Details of the format of the MSDUs  
15 71 are described in more detail below. MSDUs 71 arrive either by themselves or in association with a connection. One or more MSDUs 71 are processed by the MAC (e.g., 12, 14, or 16) to produce a Sub-Frame. The term Sub-Frame is used to refer to the data element composed of Sub-Frame Header, optional MAC Management Information, optional Delivery Time Stamp, the Payload from one or more MSDUs 71,  
20 and an optional Integrity Check Value (ICV). When a Sub-Frame is generated from multiple MSDUs 71, all MSDU 71 payloads have the same length and have identical SA 104, DA 102, MSID 118, and PLT 112. Grouping of MSDUs 71 into a Sub-Frame is done for efficiency when small fixed length MSDU 71 payloads (such as MPEG Transport Stream packets) are sent in the same stream. The format of the Sub-frame is  
25 described in more detail below. Sub-Frames are grouped into Sub-Frame streams. Each sub-frame stream is delivered independently by the MAC (e.g., 12, 14, or 16).

Each MAC 12, 14, 16 supports eight different Classes of services. Each Class encompasses a coherent set of Quality of Service (QoS) characteristics for an application and can be translated naturally to such behavior in the MAC (e.g., 12, 14,  
30 16) as channel access, number of retries, etc. Classes 0 to 3 are used by non-connection oriented MSDUs while Classes 4 to 7 are used by connection oriented services. Each MSDU 71 and hence the corresponding sub-frame stream is associated with a Class. The Sub-Frame can also carry delivery time stamp, which enable support for jitter free

delivery of the MSDU 71. Reliable end to end delivery of packets can be confirmed by means of integrity check sequence that can span on or more sub-frames.

Sub-Frames that belong to the same stream are partitioned into Segments and are transmitted as part of a MAC protocol Data Unit (MPDU) 72. Segment and MPDU 5 72 contents are described in detail below. Segments can be encrypted to provide data privacy. Details of encryption and decryption process are presented in more detail below. Each MPDU 72 contains Frame control information, MPDU header and one or more PHY Blocks (PBs). The Frame Control carries information that is relevant to all stations in the network and is broadcast. MPDU header carries information relevant to 10 all PHY Blocks. The PHY Blocks carry Segments as their payload. Details of the MPDU header and PHY Block are described below. At the physical layer level, each PB is mapped onto a FEC Block except the first PB. The first FEC Block contains MPDU header and the first PB. This mapping of segments onto the FEC blocks at the PHY level enable efficient retransmission as errors at the physical layer occur on 15 granularity of FEC blocks. PHY Blocks contains PB Header and PB integrity check sequence (PBCS). PBCS is used to test the integrity of PB. PB header is used along with the MPDU header for proper reassembly of segments and generation of Sub-Frames.

MPDUs 72 are acknowledged by a receiver layer (e.g., MAC 54) to indicate 20 reception of MPDUs. Segments that cannot be delivered reliable can be retransmitted. Segments in an MPDU 72 can be transmitted in an escalated mode. Escalated Segments are transmitted by the PHY 56 using more robust encoding, thus enabling higher probability of error free delivery. More details on Escalation are provided below. There is interactive use of PHY level 56 escalation and MAC level 54 retransmissions 25 to enable reliable end to end delivery of packets along with QoS enhancements.

### **MAC Service Data Unit (MSDU)**

MAC Service Data Unit (MSDU) 71 is the information payload that the MAC layer 54 has been asked to transport by the higher layer of the network architecture. As shown in FIG. 3, a MSDU format 100 includes a Source Address (SA) 102, a 30 Destination Addresses (DA) 104, a Traffic Information 106, a MAC Management Information 108, and a MSDU Payload 110. The Traffic information field 106 includes a Protocol Adaptation Layer (PAL) Type (PLT) 112, a Delivery Time Stamp

Flag (DTSF) 114, a MAC Management Flag (MMF) 116, and a MAC Stream Identifier (MSID) 118.

The salient features of the MSDU format 100 include support for multiple higher layers of the network architecture to interface with the MAC layer 54. Each higher layer of the network architecture 50 is provided with a unique PAL Type 112, which is carried in each MSDU 71 that is generated by the higher layer of the network architecture 50. This enables proper routing of the MSDUs 71 at the receiving MAC layer 54.

The MSDU format 100 also includes support for identifying streams of MSDUs 71 that belong to the same session or require a specific Class of service. This is achieved by means of MAC Stream identifiers (MSID) 118. Sessions can be established by negotiation between the higher layer of the network architecture and the MAC 12. During this process, each session is provided with a unique MSID 118. MSDUs 71 that belong to a session carry the MSID 118 to which each MSDU 71 is associated. In this example, MSIDs 118 enable MAC 12 to use resources allocated for that session, thus providing guarantees on various QoS parameters. A set of MSIDs 118 can be reserved for use by MSDUs 71 that do not belong to any session. In this example, MSID 118 indicates the traffic Class to which the MSDUs 71 belong. Internal to the MAC layer 54, each Class of traffic is provided with a coherent set of access parameters and allocations thus providing differentiated services. In general, established sessions can also be divided into various classes, with each class providing guarantees in a specific range of QoS parameters. In this case, MSID 118 can be used to explicitly determine the traffic Class, which is provided during connection setup.

The format of the MSDU 71 also enables an exchange of MAC Management information between the higher layers of the network architecture 50 and the MAC layer 54 by means of the optional MAC Management field 108. This feature simplifies the interface between the MAC layer 54 and the higher layers of the network architecture. Furthermore, this feature can also be used to exchange management information between higher layers of the network architecture 50.

The MSDU format 100 also provides support for the layer of the network architecture 50 that is higher than the MAC layer 54 to control when a delivery time stamp has to be inserted.

The Destination Address (DA) field **102** and Source Address (SA) field **104** are 6 octets each and carry addressing information between transmitting MAC **12** and receiving MAC **14**. An octet is a sequence of eight bits. An octet is thus an eight-bit byte. These fields **102** and **104** are identical to a 48-bit MAC address format described in the Institute of Electrical and Electronics Engineers (IEEE) Standard 802.3.

The 2-octet Traffic Information field **106** contains a 2-bit PAL Type (PLT) field, a 1-bit MAC Management Flag (MMF), a 1-bit DTS Flag, and a 12-bit MAC Stream ID (MSID) field as shown by Table 1.

Table 1. MSDU Traffic Information

Field	Length (bits)	Definition
PLT	2	PAL Type
MMF	1	MAC Management Information Flag
DTSF	1	Delivery Time Stamp Flag
MSID	12	MAC Stream Identifier

10

The PAL Type (PLT) **112** enables the MAC layer **54** to distinguish between various types of higher layers. This is used for proper routing of the MSDU **71** at the receiver layer. MAC layer **54** supports IEEE 802.3 and Isochronous Streams (IS).

Table 2 shows the interpretation of the PLT fields.

15

Table 2. PAL Type

PLT Value	Interpretation
0b00	Ethernet PAL
0b01	Isochronous Stream
0b10	Reserved
0b11	Reserved

20

The MAC Management Flag (MMF) **114** is set to 0b1 to indicate that the corresponding MSDU **71** is associated with an embedded MAC Management Information (MMI) field **108**.

The Delivery Time Stamp Flag (DTSF) **116** is set to 0b1 by the PAL **52** to indicate that this MSDU payload **110** should be associated with a Delivery Time Stamp in a Sub-Frame that may contain other MSDU payloads **110** that do not have a DTS (as indicated by a DTSF value of 0b0).

25

The MAC Stream ID (MSID) **118** is a 12-bit field that is associated with the payload being carried by the MSDU **71**. MSIDs **118** with values from 0 to 3 are used by MSDUs **71** that do not belong to an established connection and map on to MAC

Service Classes 0 to 3. The remaining MSIDs **118** may be used by connection-based services and are assigned by the MAC layer **54** during the connection setup process.

Table 3. MAC Stream Identifier

MSID Value	Interpretation
0x000	Class 0
0x001	Class 1
0x002	Class 2
0x003	Class 3
0x004 - 0xff	Negotiated Stream IDs

5

The MSDU format **100** can contain MAC Management Information **108**. The presence of this field **108** is indicated by the MMF flag **114** in the Traffic Information field **106**. If MAC Management Information **108** is present in the Sub-Frame, its format and content shall be as described in the Jitter Control Section below.

10 The MSDU Payload field **110** depends on the higher layer (e.g., PAL **52**) that generated the MSDU **71**. The MSDU Payload **110** is not interpreted by the MAC layer **54**.

The Sub-Frame may contain MAC Management Information **108** and no MSDU Payload **110**, or a MSDU Payload **110** and no MAC Management Information  
15 **108**, or it may contain both.

### Sub-Frame

The MAC layer **54** processes one or more MSDUs **71** to generate a Sub-Frame. As shown in FIG. 4, a Sub-Frame **150** includes a Sub-Frame Header **152**, Optional MAC Management information **154**, Optional Delivery time stamp **156**, payload **110**  
20 from one MSDU and an optional integrity check sequence (ICV) **158**. Sub-Frame header **152** contains MAC Management Flag **182**, Integrity Check Sequence Flag (ICVF) **184**, and Sub-Frame Payload length **186**. The format of Sub-Frame **150** is also specified in Table 4.

Table 4. Sub-Frame Format

Field	Length	Definition
SFH	2 octets	Sub-Frame Header
MAC Management Information	0-M octets	Optional MAC Management Information
DTS	3 octets	Optional Delivery Time Stamp
MSDU Payload	variable octets	Optional MSDU Payload
ICV	4 octets	Optional Integrity Check Value

As shown in FIG. 5, the Sub-Frame Header **152** is a 2-octet field that carries information about the presence of MAC Management Information and Integrity Check Value (ICV) in the Sub-Frame as well as the length of the Sub-Frame. This information includes MAC Management Flag **182**, Integrity Check Value flag **184**, and length field **186**. The Sub-Frame header is also specified in Table 5.

Table 5. Sub-Frame Header

Field	Length	Definition
MMF	1 bit	MAC Management Flag
ICVF	1 bit	ICV Flag
LEN	14 bits	Sub-Frame Length

The MAC Management Flag **182** is set to 0b1 to indicate the presence of MAC Management information **154**. MAC Management information **154**, if present, shall follow the sub-frame header **152**.

The Integrity Check Value Flag **184** is set to 0b1 to indicate the presence of an ICV field **158** in the corresponding Sub-Frame **150**. The ICV field **158**, if present, follows the Sub-Frame payload **110**.

The Length field **186** is a 14-bit used to specify the length of Sub-Frame **150**, excluding the 2-octet Sub-Frame Header **152** and the 4-octet ICV (if present) **158**.

The Sub-Frame **150** can contain MAC Management Information **154** as indicated by the MMF flag **182** in the Sub-Frame Header **152**. If the MAC Management Information **154** is present in the Sub-Frame **150**, its format and content is as described in the Jitter Control Mechanism section below.

The optional Delivery Time Stamp (DTS) **156** is the 24-bit value of the sender's local 25 MHz multimedia clock at the time at which the MSDU **71** arrived from the sender's PAL **52**, plus the delivery latency associated with this MSDU **71**. This value indicates the time at which the MSDU **71** should be presented to the destination's PAL

52. The DTS field 156 shall be included in a Sub-Frame 150 only when required for jitter control as negotiated at stream set-up. At that time, the option of one DTS 156 per Sub-Frame 150 or one DTS 156 per MSDU payload 110 shall be selected for the stream. The DTS 156 will precede the MSDU payload(s) 110 to which it applies, and these payloads 110 will be grouped according to the DTS Flag 116 in the MSDU traffic information 106. All the MSDUs 100 with DTSF=0b0 will be grouped into a single Sub-Frame 150 with the next MSDU 100 whose DTSF=0b1.

The Sub-Frame Payload field 160 contains the payload 110 from one or more MSDUs 71 depending on how the Sub-Frame 150 was formed.

10 The Integrity Check Value (ICV) 158 is a Cyclic Redundancy Code (CRC)-32 error checking code computed over one or more Sub-Frames 150. The ICV Flag (ICVF) 158 in the Sub-Frame header 152 is used to determine the Sub-Frames 150 over which the ICV 158 is computed. The ICV 158 does not cover the Sub-Frame headers 152. FIG. 6 shows a block of Sub-Frames 150 protected by a single ICV 158.

15 Sub-Frames 150 that are generated from MSDUs 71 belonging to the same {SA 104, DA 102, PLT 112 and MSID 118} tuple are grouped together to form a sub-frame stream. When a MPDU 72 is generated by the MAC layer 54, its payload contains Sub-Frame(s) 150 from only one sub-frame stream at a time.

The salient features of Sub-Frame 150, and Sub-Frame Stream generation process include removing information that is common to all the MSDUs 71 that belong to a single stream while a sub-frame 150 is generated. This information is only transmitted once per MPDU 72, thus increasing protocol efficiency.

Multiple MSDU payloads 110 can be transmitted in a single Sub-Frame 150. This improves the protocol efficiency when small fixed length MSDU payloads 110 are sent in the same stream.

The structure of the Sub-Frame 150 provides a mechanism for carrying management information along with MSDU payload 110.

Sub-Frames 150 also provide a mechanism for transmitting delivery time stamps 156. These delivery time stamps 156 provide the time at which the Sub-Frame 150 has to be delivered to the higher layer of the architecture 50 at the receiver MAC (e.g., 12, 14, 16).

The structure of the Sub-Frame **150** allows for inserting an ICV **158** on each Sub-Frame **150** or a group of Sub-Frames **150** at a time. The ICV **158** enables end-to-end check for proper reception of Sub-Frames **150**.

The Sub-Frame **150** is generated by processing one or more MSDUs **71**. The generation of a Sub-Frame **150** from an MSDU **71** is shown in FIG. 7 for the case of a Sub-Frame **150** formed from a single MSDU **71**. When a Sub-Frame **150** is generated from multiple MSDUs **71**, all MSDU payloads **110** have the same length and belong to an established session. This is done for efficiency when small fixed length MSDU payloads **110** are sent in the same stream. FIG. 8 shows the generation of a Sub-Frame **150** for the case when the Sub-Frame **150** is formed from multiple MSDUs **71**.

### **Sub-Frames Streams, Sub-Blocks and Segments**

As shown in FIG. 9, a Sub-Frame Stream **200** includes Sub-Frames **150** generated from MSDUs **71** that belong to the same {SA, DA, MSID, PLT} tuple. A group of Sub-Frames **150** that are protected by a single Integrity Check Value (ICV) **158** forms an ICV Block, which is the basic entity that is subjected to end-to-end MAC delivery services. This process of generating a Sub-Frame Stream **200** from MSDUs **71** is called encapsulation.

As shown in FIG. 10, the Sub-Frame Stream **200** is divided into fixed size Sub-Blocks **250**. One or more such Sub-Blocks **250** are then grouped into a Segment **252** to form the basic entity processed by the MAC layer **54** to ensure reliable delivery services. Sub-blocks **250** are numbered entities used for reassembly at the receiver. The Sub-Frame **150** boundary demarcation information is transmitted to the receiver in the MPDU Header. Each segment is padded as necessary, optionally encrypted, and then inserted into a PHY Block (PB) Body. In some examples, padding zeros and a length field are added to a Segment **252** if the buffer is depleted when the Segments **252** are being formed.

### **MAC Protocol Data Unit (MPDU) and FEC Blocks**

The term *MAC Protocol Data Unit (MPDU)* **254** is the information that the PHY **56** has been asked to transport by the MAC layer **54**. The MPDU **72** is composed of a Frame Control field **256**, MPDU Header **258** and one or more PHY Blocks **266**. Frame Control carries broadcast information. The MPDU header **258** and the first

PHY Block **266** are transmitted using a single FEC Block **268**. The subsequent PHY Blocks **266** are transmitted in separate FEC Blocks **266**. The first FEC Block **268** in an MPDU **72** is of a larger size to accommodate the fixed length MPDU header **258** along with the PHY Block **266**. All the PHY Blocks **266** have a fixed size except for the last one in the MPDU **72**.

The salient features of the MPDU format include that all the information that is common to all Segments **252** in an MPDU **72** is transmitted as part of the MPDU header **258**, thus improving the efficiency of communication. Furthermore, segmentation across Sub-Frame boundaries provides high MPDU transmission efficiency under a very large range of MSDU, Sub-Frame sizes. The MPDU header **258** is protected by a special integrity check, which provides better performance on marginal channels. The MPDU header **258** carries local clock time stamp information. This time stamp can be used by the receiver MAC (e.g., **14**) to synchronize with the transmitter MAC **12**, thus enabling jitter free service. The mapping of MPDU header **258** and first PHY Blocks **266** on to the first FEC Block **268** that has a larger size to enable MPDU header **258** overhead enables efficient retransmission of lost PHY Blocks **266**. Support for Escalating the PHY Block **266** encoding is provided. This mechanism can be used in conjunction with retransmissions to enhance QoS guarantees. There is also support of Multicast with partial ARQ, bridging and forwarding.

The format of MPDU Header **258** is shown in FIG. 11. The receiver MAC **14** uses information contained in the MPDU header **258** along with the information in the PB header **260** to decrypt and to reassemble the Sub-Frames **150**. The MPDU header **258** includes MPDU Control **300**, DA **302**, SA **304**, ODA **306**, OSA **308**, and HCS **310**. The fields that comprise the 12 octets of the MPDU Control **300** are shown in Table 6.

Table 6. MPDU Control Format

Field	Length (bits)	Definition
NEPB	2	Number of Empty PBs
MSID	12	MAC Stream ID
PLT	2	PAL Type
TS	24	Time Stamp
EKS	12	Encryption Key Select
SFPBN	6	Sub-Frame boundary PHY block number
SFO	10	Sub-Frame boundary offset in PB

Number of Empty PHY blocks (NEPB) is two bits of the MPDU header which are used to indicate the number of empty PBs **266** at the end of the PPDU Payload. The restrictions on the frame length at high data rates cause increments of as many as 3 FEC blocks between successive valid frame sizes. The sender MAC (e.g., **12**) may  
5 only require one of these FEC blocks **268** to hold data, and so there may be zero, one, or two empty PBs at the end of the PHY PDU Payload, as indicated by NEPB.

The MAC Stream ID (MSID) field carries the MAC Stream ID that is associated with the payload being carried by this MPDU. MSIDs 0 to 3 are used by MPDUs that carry connectionless Class 0 to 3 traffic respectively. The remaining  
10 MSIDs may be used by connection-based services and are assigned by the MAC during the connection setup process.

The PAL Type (PLT) field defines the PAL Type (PLT) that is being carried by the MPDU. The MAC receiver uses this to reassemble and to route the MSDUs to the correct PAL.

15 The Time Stamp (TS) field is a 24-bit Time Stamp representing the value of the local transmitter's Multimedia clock with reference to the start of the preamble when the MPDU was transmitted. The TS field is used for jitter-free delivery (in conjunction with the Delivery Time Stamp (DTS) in the Sub-Frame Header), Tone Map (TM) timing and in managing the Periodic Contention Free Channel Access.

20 The Encryption Key Select (EKS) field is an Index of the Encryption Key used for encrypting the Segments. In some examples, EKS is 12 bits long, providing additional keys for access networks. A value of 0x000 indicates that the segments are encrypted using the stations default encryption key. A value of 0xffff indicates that the Segments in the MPDU **72** are not encrypted. Preferred implementations can also  
25 obtain the EKS by processing the frame control header fields.

The Sub-Frame Boundary PHY Block Sequence Number (SFPBN) field carries a number representing the relative position within the MPDU of the PHY Block that contains a Sub-Frame boundary. A value of 0b000000 indicates the first PB, 0b000001 indicates the second PB, etc. A value of 0b111111 indicates that no Sub-Frame  
30 boundary exists in the current MPDU **72**.

The Sub-Frame boundary offset (SFO) field carries the offset in bytes of the Sub-Frame boundary (i.e., the first octet of the first new Sub-Frame) within the PHY Block indicated by SFPBN. A value of 0x000 indicates the first byte.

The Destination Address (DA) 302, Source Address (SA) 304, Original Destination Address (ODA) 306, and Original Source Address (OSA) 308 fields carry the addressing associated with the MPDU 72.

The Destination Address (DA) 302 is a 48-bit address for the receiver to which this MPDU 72 is being sent in the current transmission. The address format follows the IEEE 802.3 Ethernet Standard.

The Source Address (SA) 304 is a 48-bit address for the station (e.g., MAC 12) that is sending this MPDU 72 in the current transmission. The address format follows the IEEE 802.3 Ethernet Standard.

The Original Destination Address (ODA) 306 is a 48-bit address for the receiver that is the ultimate destination of this MPDU 254. The address format follows the IEEE 802.3 Ethernet Standard.

The Original Source Address (OSA) 308 is a 48-bit address for the station (e.g., MAC 12) from which this MPDU 72 originated. The address format follows the IEEE 802.3 Ethernet Standard.

The contents of the DA 302, SA 304, ODA 306 and OSA 308 fields in the MPDU header 258 are used to indicate whether the MPDU 72 being transmitted is a Regular MPDU or a Multicast MPDU with Response. Table 7 summarizes the interpretation of these addresses.

Table 7. ODA, OSA, DA, and SA fields interpretation

DA	SA	ODA	OSA	Interpretation
ODA	OSA	Unicast	Unicast	Regular MPDU
not ODA, Unicast	OSA	Unicast	Unicast	Bridged/Forwarded MPDU from the Original Source
ODA	not OSA, Unicast	Unicast	Unicast	Bridged/Forwarded MPDU designated to the Original Destination
not ODA, Unicast	not OSA, Unicast	Unicast	Unicast	Bridged/Forwarded MPDU between two intermediate stations
not ODA, Unicast	Unicast	M/B	Unicast	Multicast or Broadcast MPDU with DA indicating the address of the responder (for partial ARQ)
not ODA, unicast	not OSA, Broadcast	Unicast	Unicast	Bridged/Forwarded MPDU with DA indicating the address of the responder (for partial ARQ) and SA indicating the set of station to which the MPDU is intended

M/B ≡ Multicast/Broadcast

The Header Check Sequence (HCS) is a 32-bit CRC computed over all the MPDU Header fields. After receiving the MPDU, stations shall compute the 32-bit

CRC based on the above process to detect transmission errors. If any transmission error is detected, the entire MPDU is discarded. To reduce the probability of errors in the MPDU header, the first FEC Block may be more robustly encoded than the standard FEC block.

5 Each PHY Block (PB) or PB with MPDU header is mapped onto a single Forward Error Correction (FEC) block at the physical layer. A Long MPDU can carry one or more PHY blocks. Each PB contains a PB Header (PBH), PB Body (PBB) and PB Check Sequence (PBCS). The MPDU Header is always carried as an addition field pre-pended to the first PB in the MPDU.

10 The salient features of the PHY Block format include that the PHY Block Check Sequence (PBCS) provides a very highly reliable error detection mechanism. Further mapping of PHY Blocks on to the FEC Blocks enable efficient retransmission.

The PHY Block format also enables the Sub-Block Sequence number to simplify reassembly and provides duplicate rejection at the receiver.

15 The PHY Block header format also provides a mechanism to transmit MAC Management frame in an out of band manner. This mechanism enables fast exchange of important MAC Management information.

The PHY Block body size is chosen to enable zero encryption overheads in the PHY Block Body. The overall encryption mechanism simplifies implementation.

20 Three sizes, 263, 519, and 775 octets (with 256, 512, or 768 octets of PBB for the segment it contains, respectively) are supported for PHY blocks **266**. However, there are six FEC block information field sizes, namely 263, 519, and 775 octets for FEC Blocks containing only a PHY Block and 303, 559, and 815 octets for FEC Blocks containing a PHY Block and MPDU Header or SMPDU header and VFs field (in  
25 SACK long MPDUs). The larger size accommodates an additional 40 Octets for the header and the extra data. The first FEC block in a PPDU contains an MPDU header and a PB, while the rest contain only one PB each. When the PHY Body is filled with FEC blocks that form the PHY Payload, maximum size PBs shall be used for all but the  
30 constraints, the sender (e.g., MAC 12) shall fill as much of the PHY Body as possible with PHY Payload.

The fields in the 3-octet PB Header are shown in Table 8.

Table 8. PB Header Format

Field	Length	Definition
SBSN	14 bits	Sub-Block Sequence Number
PBLT	2 bits	PB Length Type
ECV	1 bit	Erasur e Code Version
EGL	5 bits	Erasur e Group Length
PBN	2 bit	Parity Block Number

The PB Header consists of a 14-bit Sub-Block Sequence Number and a 2-bit Length Type (PBLT) field, 1-bit Erasure Code Version, 5-bit Erasure Group Length  
5 and a 2-bit Parity Block Number.

The Sub-Block Sequence Number (SBSN) field indicates the sequence number of the first Sub-Block in the segment. The SBSN can be used by the receiver to properly insert the received Segments in the reassembly buffer. The process of numbering Sub-Blocks combined with fixed Sub-Block sizes eliminates the need for  
10 buffer reordering when out of order segments are received. Dividing the queue into sub-blocks of equal size and sending the sequence number in the PHY Block header simplifies reassembly while reducing the overhead required to carry the sequence number. The overhead is reduced because numbering is done one sub-block at a time rather than one byte (or one bit) at a time. For example, using 256 byte blocks  
15 compared to byte number saves 8-bits of space in the PHY block header. Reassembly is simplified because the receiver exactly knows where to put each sub-block.

SBSN numbers shall be initialized to 0 when a CF session is set up, and wrap around as long as the CFID is in use. For non-CF traffic (MSIDs 0-3), it is initialized to 0, wraps around as needed. For CSMA/CA traffic, the last SBSN shall be stored  
20 until twice the maximum Sub-Frame lifetime after which the SSBN shall be reset to 0. The first segment with a reset SBSN should have SFPBN=0 and SFO = 0 also. When EGL is non-zero (i.e., Parity PB), this field carries the sequence number of the first sub-block in the last segment of the erasure group.

The PHY Block Length Type (PBLT) is a 2-bit field that indicates whether the PHY Block Body (PBB) is full, short 1 octet, or short more than 1 octet. The PBLT values and meanings are given in Table 9.

Table 9. PBLT Values and Meaning

PBLT Value	Meaning
0b00	The PBB is full, all octets are valid
0b01	The last octet of the PBB is not valid, the segment length is (PBB length – 1) octets (i.e., 767 octets)
0b10	The segment contained in the PBB is more than 1 octet shorter than the PBB. In this case the last two octets of the PBB form a length field that explicitly gives the segment length in octets..
0b11	The segment contained in the PBB is destined for the MAC Management Queue for this {SA, DA} pair. The last two octets of the PBB form a length field that explicitly gives the segment length in octets.

5

In the case of PBLT = 0b10 or 0b11, the implied 2-octet length field contains the valid data length of the Segment carried by the PBB. The rest of the Segment is zero padded. The PHY Payload length may be large enough to hold more FEC blocks than are required by the MAC, which means that the last FEC block will not hold a PB. In this case, the transmitter inserts an empty PB with the PBLT=0b10 and a length field of 0x00 so that the receiver will discard this PB. The NEPB field of the MPDU Header indicates the number of these PBs so the receiver can discard them without having to decrypt them. When PBLT=0b11, then the receiver reassembles the segment contained in the PBB into the MAC Management Sub-Frame queue associated with this {SA, DA} pair. The MSB of the length field in the PBB of PBs with PBLT=0b11 shall be interpreted as the Sub-Frame Boundary Flag (SFBF). This bit allows the sender to indicate to the receiver that the first octet of the PBB is a sub-frame boundary (when SFBF=0b1).

An Erasure Group Length field when set to 0b00000, indicates a normal PB. A non-zero value of the EGL indicates parity PB. In this case, the value in the EGL field is the number of normal PBs (or the length of erasure group) covered by this parity PB. A value of 0b00001 indicates erasure group of length one and so on. A value of 0b11111 indicates an erasure group of size 31.

A Parity Block Number field is valid only when the EGL is set to a non-zero value. PBN indicates the sequence number of the parity block and is used by the receiver to recover lost segments. This field shall be set to 0b00 for this version

The PHY Block (PB) body carries the encrypted Segment as the payload. Note that a Segment may have to be zero-padded before encryption to ensure that it fits exactly into the PB Body. The PB Header and the PBCS are not encrypted.

The PHY Body Check Sequence (PBCS) is a CRC-32 and is computed over the PB Header and the encrypted PB Body. The PBCS of the first PB in an MPDU is not computed over the MPDU header.

### MAC Management Information Fields

MAC Management Information (MMI) can be transmitted as part of an MSDU or a Sub-Frame. When MMI is transmitted as part of an MSDU, the presence of this field is indicated by setting the MAC Management flag in the Traffic Information to 0b1 (refer to Section 1). When the MMF flag is set, the MMI field immediately follows the end of the Traffic Information.

When MMI is transmitted as part of a Sub-Frame, the presence of this field is indicated by setting the MAC Management flag in the Sub-Frame header to 0b1 (Refer to Section 2). When the MMF flag is set, the MAC Management Information field immediately follows the end of the Sub-Frame header. Table 10 shows the structure of the MMI field. Note that the MMI field has variable structure and that the sub-fields are so defined as to specify the particular structure of the MMI field.

Table 10. MAC Management Information Field Format

Field	Length	Definition
NE	1 octet	Number Of MAC Data Entries (L)
MEHDR <sub>1</sub>	1 octet	First MAC Management Entry Header
MELEN <sub>1</sub>	2 octet	First MAC Management Entry Length (= N <sub>1</sub> )
MMENTRY <sub>1</sub>	N <sub>1</sub> octets	First MAC Management Entry Data
• • •		
MEHDR <sub>i</sub>	1 octet	i <sup>th</sup> MAC Management Entry Header
MELEN <sub>i</sub>	2 octet	i <sup>th</sup> MAC Management Entry Length (= N <sub>i</sub> )
MMENTRY <sub>i</sub>	N <sub>i</sub> octets	i <sup>th</sup> MAC Management Entry Data
• • •		
MEHDR <sub>L</sub>	1 octet	Last MAC Management Entry Header
MELEN <sub>L</sub>	2 octet	Last MAC Management Entry Length (= N <sub>L</sub> )
MMENTRY <sub>L</sub>	N <sub>L</sub> octets	Last MAC Management Entry Data

20

The 1-octet Number of Entries (NE) field specifies the number of separate MAC Management Entries that are contained in the MMI field. Supposing that NE is L, then the MMI field contains L structures, one for each MAC Management Entry. Each such structure includes a MAC Management Entry Header (MEHDR), a MAC

Management Entry Length (MELEN), and the associated MAC Management Entry data (MMENTRY).

For the  $i^{\text{th}}$  MMENTRY, the  $i^{\text{th}}$  MAC Management Entry Header (MEHDR <sub>$i$</sub> ) field specifies a 1 octet header. The MAC Management Entry Header structure is as shown in Table 11.

Table 11. MAC Management Entry Header Field

Field	Bit Number	Bits	Definition
MEV	7-6	2	MAC Entry Version
METYPE	5-0	6	MAC Entry Type

The 2-bit MAC Management Entry Version (MEV) field indicates the version in use for interpretation of MAC Entries. If the received MEV is not equal to 0b00, the receiver discards the MAC Management Entry and uses the MAC entry length field to determine the number of octets to ignore before continuing to process the remainder of the Sub-Frame.

The 6-bit MAC Management Entry Type (METYPE) field defines the MAC entry command or request that follows. Several METYPES are defined that enables such functions as layer management, Session set up etc.

The MAC Entry Length field (MELEN <sub>$i$</sub> ) contains the length in octets of the MMENTRY field to follow. If MMENTRY does not exist, MELEN is set to zero. This field provides for transparent extension of MAC management, without rendering older equipment obsolete. If an MSDU or a Sub-Frame is received with an METYPE value that is not understood, the receiver can still properly parse the MSDU or Sub-Frame and process its contents, ignoring what it does not understand. The format of MMENTRY depends on the MEHDR with which it is associated.

### Jitter Control Mechanism

A Jitter Control mechanism enables station to deliver MSDUs with a very low jitter in the order of a few nano seconds. This mechanism uses the Delivery time stamp in the Sub-Frames to determine when the corresponding MSDU has to be delivered to the higher layer at the receiver. Synchronization of the clocks of the transmitters (e.g., MAC 12) and receivers (e.g., MAC 14) is obtained by transmitters inserting its local clock time stamp in MPDU header and receiver using this to synchronize with the transmitter.

The salient features of jitter control mechanism include support for very low end-to-end jitter. The jitter control mechanism also includes support for higher layers of the network architecture to control the insertion of Delivery time stamps. This support for higher layers reduces overhead while providing the needed functionality.

5 The jitter control mechanism can also use tracking algorithms to obtain close synchronization with the transmitters clock, thus enabling low end-to-end jitter guarantees in the order of nano-seconds. Furthermore, multi-streaming applications can use jitter control mechanism to provide synchronization between multiple receiver MACs.

10 Each MAC maintains a 25MHz System Clock. Any MSDU that belongs to a jitter-free session is associated with a 24-bit Delivery Time Stamp (DTS) when the MSDU arrives at the MAC. This timestamp is inserted into the Sub-Frame that is generated from the MSDU (and possibly other MSDUs). When multiple MSDUs are combined into a single Sub-Frame with a single timestamp, the DTS Flag (DTSF) in  
15 the MSDU header indicates which MSDUs are to generate the timestamp. When an MSDU with the DTSF=0b1 arrives, its timestamp is generated and inserted into the Sub-Frame along with the MSDU payload and all other MSDU payloads that arrived since the last MSDU with DTSF=0b1. At the receiver, all of these MSDU payloads are delivered by the time indicated by the DTS in the Sub-Frame, with the last MSDU  
20 payload delivered at the indicated time. The PAL sending the MSDUs **71** to the source MAC (e.g., **12**) takes care not to exceed the maximum Sub-Frame size before a time stamped MSDU **71** is sent.

The DTS is the sum of the system clock value when the MSDU **71** is received plus the end to end latency associated with the traffic (this is determined during the call  
25 admission process and the QoS for this traffic type). Every MPDU **72** carries the transmitter's System Clock time stamp (with respect to the start of the preamble) in the MPDU header **258**. The receiver may use jitter control algorithm to provide very low jitter guarantees.

The receiving MAC (e.g., **14**) delivers jitter-free traffic to the destination PAL  
30 at the time indicated in the delivery time stamp (DTS) based on the information derived from the System Clock timestamps in the MPDU headers **258**.

## ARQ, Escalation, and Erasure Codes

MPDUs 72 are acknowledged by the receiver to indicate reception station. Segments that cannot be delivered reliable can be retransmitted. A retransmitted segment is packaged in a new PB in the front of the next available MPDU 72 and is retransmitted. The retransmitted PBs will normally be escalated to improve their chances of correct reception. The number of escalated PHY Blocks in the MPDU 72 can be indicated in the frame control header. MAC layer can also use parity PBs to ensure reliable delivery of regular PBs. Parity PBs are generated by from a group of regular PBs and can be used to recover one or more lost PBs at the destination without having to retransmit them. These mechanisms enable latency sensitive packets to be delivered more effectively with a limited number of retries. Escalation and Erasure codes tradeoff data rate of the channel with the number of retries required to get a certain packet loss rate.

## Encryption

Some implementations allow MACs to transmit segments in an encrypted form, thus providing privacy of data. Encryption information may include an Network Encryption Key (NEK) that indicates the key to be used to decrypt a block and an Initialization Vector (IV) that is used to initialize the decryption algorithm. Both NEK and IV should be correctly known to the receiver to properly decrypt the PB. The Encryption Key Select (EKS) field in the MPDU Header is used to refer to the index of the Network Encryption Key (NEK) used for encryption. The NEK to be used for encrypting any Segment and the corresponding EKS are exchanged between station prior to the transmission of MPDU. The Initialization Vector (IV) used for encrypting the first PHY Block is obtained by concatenating fields from Frame Control, MPDU header and PHY block header. Other preferred implementations may obtain the EKS by processing the fields of the Frame Control. For example, the EKS can be derived from a substantially unique session identifier carried in the Frame Control. The Initialization vector can be generated from the fields of the frame control and the PHY Block header. Once the MPDU is delivered to the destination, the PBCS of each PB is checked and then the good PBs are decrypted and delivered to the receiver buffer. PB failures are reported to the transmitting station by a SACK and are re-encrypted and retransmitted, using a current Network Encryption Key (NEK) and a new Initialization

Vector (IV). This process reduces the overhead for transmission of initialization vector. Further, proper choice of PHY Block body length can be used to reduce the encryption pad that might be needed.

Other implementations of the invention are within the following claims.

5

54830-2

CLAIMS:

1. A method of operating in a network in which a plurality of stations communicate over a shared medium, the method comprising
- 5 providing a physical layer for handling physical communication over the shared medium;
- providing a high level layer that receives data from the station and supplies high level data units for transmission over the medium;
- providing a medium access control (MAC) layer that receives the high level data units from the high level layer and supplies low level data units to the physical layer;
- 10 at the MAC layer, encapsulating content from a plurality of the high level data units into a stream of sub-frames;
- dividing the encapsulated stream into a plurality of pieces with each piece capable of being independently retransmitted; and
- supplying low level data units, at least some of the low level data units each
- 15 containing a plurality of pieces into which the encapsulated stream was divided, and at least some of the low level data units containing boundary demarcation information indicating boundaries between the sub-frames in the stream;
- wherein at least some information common to the high level data units is not repeated for each high level data unit encapsulated in the stream;
- 20 wherein the high level data units each comprise a payload, and encapsulating comprises forming the stream comprising the payloads from a succession of high level data units;
- wherein the stream comprises a succession of sub-frames, each sub-frame comprising a header and a plurality of the payloads;

54830-2

wherein each sub-frame is divided into the plurality of pieces capable of being independently retransmitted; and

wherein division of a sub-frame into the plurality of pieces comprises dividing the sub-frame into a plurality of sub-blocks, and forming at least some pieces from a plurality  
5 of sub-blocks.

2. The method of claim 1 wherein the information common to the high level data units comprises destination and source addresses.
3. The method of claim 1 wherein each piece constitutes a segment that is transmitted as a physical layer block.
- 10 4. The method of claim 1 further comprising parity pieces derived from other pieces and capable of being used at a destination to recover one or more lost pieces at the destination without having to retransmit the lost pieces.
5. The method of claim 4 wherein each piece is transmitted as a physical layer block, and the parity pieces are also transmitted as parity physical layer blocks.
- 15 6. The method of claim 5 wherein the physical layer blocks are encoded using forward error correction.
7. The method of claim 1 wherein some of the pieces making up a low level data unit constitute retransmitted pieces that failed to be correctly transmitted in an earlier attempt.
8. The method of claim 7 wherein at least some retransmitted pieces are  
20 transmitted with greater forward error correction than forward error correction used in the earlier attempt.
9. The method of claim 1 wherein each sub-frame further comprises a delivery time stamp associated with at least some payloads.

54830-2

10. The method of claim 1 wherein clock information characterizing the time setting of a clock in a transmitting station is transmitted to a receiving station within a header of the low level data units, and the clock information is used by the receiving station along with the delivery time stamps to establish the time at which payloads are delivered.
- 5 11. The method of claim 10 wherein the time at which a payload is delivered is substantially the time specified by the time stamp based on information derived from the clock information.
12. The method of claim 1 further comprising an integrity check value associated with each sub-frame or with a plurality of sub-frames.
- 10 13. The method of claim 1 wherein each of the plurality of payloads in a sub-frame have identical length.
14. The method of claim 1 wherein each sub-frame further comprises MAC management information.
- 15 15. The method of claim 1 wherein the MAC layer has the capability of transmitting data in a plurality of sessions within a regularly-repeated contention free interval, wherein a station to which data is transmitted is identified by a destination address and a station from which data is transmitted is identified by a source address, and wherein the stream contains a queue of payloads for the same session, same source address, and same destination address.
- 20 16. The method of claim 1 wherein the MAC layer has the capability of transmitting data in a plurality of sessions within a regularly-repeated contention free interval, wherein a station to which data is transmitted is identified by a destination address and a station from which data is transmitted is identified by a source address, and wherein the queue contains sub-frames for the same session, same source address, and same destination address.

54830-2

17. The method of claim 15 wherein the MAC layer processes data transmitted in the sessions according to contention-free channel access processing.
18. The method of claim 17 wherein the sessions are transmitted within time slots of a regularly-repeated contention-free interval.
- 5 19. The method of claim 15 wherein a stream identifier is used to associate content of a queue with a particular session.
20. The method of claim 19 wherein the stream identifier is also used to associate content of a queue with a priority level for contention-based transmission over the medium.
21. The method of claim 19 wherein there are a plurality of queues of payloads in  
10 the stream, each containing payloads having a unique combination of stream identifier, source address, and destination address.
22. The method of claim 21 wherein each queue contains payloads having a unique combination of stream identifier, source address, destination address, and type of high level layer.
- 15 23. The method of any one of claims 19 to 22, wherein the stream identifier is a MAC Stream identifier (MSID).
24. A method of operating in a network in which a plurality of stations communicate over a shared medium, comprising
- 20 providing a physical layer for handling physical communication over the shared medium;
- providing a high level layer that receives data from a station and supplies high level data units for transmission over the medium;
- providing a medium access control (MAC) layer that receives the high level data units from the high level layer and supplies low level data units to the physical layer;

54830-2

at the MAC layer, encapsulating content from a plurality of the high level data units into a stream of sub-frames;

dividing the encapsulated stream into a plurality of pieces with each piece capable of being independently retransmitted; and

5 supplying low level data units, at least some of the low level data units each containing a plurality of the pieces into which the encapsulated stream was divided, and at least some of the low level data units containing boundary demarcation information indicating boundaries between the sub-frames in the stream;

10 wherein at least some information common to the high level data units is not repeated for each high level data unit encapsulated in the stream;

wherein the high level data units each comprise a payload, and encapsulating comprises forming the stream comprising the payloads from a succession of high level data units;

15 wherein the stream comprises a succession of sub-frames, each sub-frame comprising a header and a plurality of the payloads; and

wherein the stream is divided into a plurality of sub-blocks, wherein a plurality of sub-blocks are grouped to form a segment, with a segment crossing sub-frame boundaries in the stream, wherein a segment constitutes one of the pieces.

25. The method of claim 24 wherein each sub-block is shorter than a sub-frame.

20 26. The method of claim 24 wherein at least some segments contain a number of sub-blocks corresponding to one or more sub-frames including at least one sub-frame whose associated sub-blocks comprise less than the whole sub-frame.

27. The method of claim 24 wherein the sub-blocks are of equal length.

54830-2

28. The method of claim 24 wherein the sub-blocks have an associated sequential numbering adapted for use at the receiving station for re-establishing the correct sequential order of the sub-blocks.
29. The method of claim 28 wherein the sub-blocks have a predetermined size,  
5 which combined with the associated sequential numbering, eliminates the need for buffer reordering when out of order segments are received.
30. The method of claim 29 wherein the sub-blocks are of equal size.
31. The method of claim 24 further comprising, for at least some of the low level data units, forming the low level data unit from a plurality of segments.
- 10 32. The method of claim 31 wherein each segment in the low level data unit forms the body of a separate block transmitted by the physical layer.
33. The method of claim 31 wherein individual segments are individually encrypted.
34. The method of claim 33 wherein encryption information common to a plurality  
15 of segments is carried in a header.
35. The method of claim 34 wherein some encryption information is carried in a header and frame control of the low level data unit and in a header of the block.
36. The method of claim 33 wherein some encryption information is carried in frame control of the low level data unit and in a header of the block.
- 20 37. The method of claim 32 wherein each block separately undergoes forward error correction, and forward error correction bits for each block are transmitted in the low level data unit.
38. The method of claim 37 wherein the level of forward error correction used is different for different blocks.

54830-2

39. The method of claim 38 wherein the level of forward error correction used provides greater error correction capability for selected blocks that are being retransmitted after failing to be correctly transmitted in an earlier attempt.
40. The method of claim 32 wherein most of the blocks are identical in length.
- 5 41. The method of claim 40 wherein the initial and final block of a low level data unit can be of a different length than the remaining blocks.
42. The method of claim 31 wherein information common to the plurality of segments forming the low level data unit is transmitted in a header for the low level data unit.
- 10 43. The method of claim 37 wherein the information common to the plurality of segments is transmitted only in the header.
44. The method of claim 37 wherein the low level data unit further comprises a frame control field.
45. The method of claim 1, further comprising:  
 adaptively escalating the robustness of transmission of the low level data units  
 15 depending on the frequency of transmission errors.
46. The method of claim 45 wherein  
 the method further comprises incorporating forward-error correction  
 information into the transmitted stream of low level data units, and  
 wherein the step of adaptively escalating comprises adaptively varying the  
 20 forward-error correction information depending on the frequency of transmission errors.
47. The method of claim 46 wherein varying the forward-error correction information comprises varying one or both of the amount and type of forward-error correction information.

54830-2

48. The method of claim 45 wherein decisions on adaptively escalating are made at a transmitting station.

49. The method of claim 45 wherein each of the low level data units contains a plurality of the pieces.

5 50. The method of claim 48 wherein the forward error correction information comprises information associated with provided with the pieces for use at a destination for recovering a piece that is received with errors.

51. The method of claim 48 wherein the forward error correction information comprises parity pieces derived from other pieces and capable of being used at a destination to  
10 recover one or more lost pieces at the destination without having to retransmit the lost pieces.

52. The method of claim 51 wherein each piece is transmitted as a physical layer block, and the parity pieces are also transmitted as parity physical layer blocks.

53. A method of operating in a network in which a plurality of stations communicate over a shared medium, comprising

15 providing a physical layer for handling physical communication over the shared medium;

providing a high level layer that receives data from a station and supplies high level data units; and

20 providing a layer that receives the high level data units from the high level layer and supplies low level data units to the physical layer;

wherein supplying the low level data units comprises

encapsulating content from a plurality of the high level data units into a stream of sub-frames,

54830-2

dividing the encapsulated stream into a plurality of sub-blocks,

forming a plurality of pieces, with each piece including one or more sub-blocks, to provide pieces capable of being independently retransmitted, and

supplying low level data units, at least some of the low level data units each  
 5 containing a plurality of the pieces that include the sub-blocks into which the encapsulated content was divided, and at least some of the low level data units containing boundary demarcation information indicating boundaries between the sub-frames in the stream.

54. The method of claim 53 wherein a plurality of the plurality of pieces each include a same number of the sub-blocks.

10 55. The method of claim 53 wherein at least one of the plurality of pieces includes one or more sub-blocks and padding.

56. The method of claim 55 wherein an amount of padding in a piece is selected based on the number of sub-blocks in the piece and the size of the pieces.

15 57. The method of claim 53 wherein the high level data units each comprise a payload, and encapsulating comprises forming the stream comprising the payloads from a succession of high level data units.

58. The method of claim 57 wherein the stream comprises a succession of sub-frames, each sub-frame comprising a header and a plurality of the payloads.

20 59. The method of claim 58 wherein each sub-frame is divided into the plurality of pieces capable of being independently retransmitted.

60. The method of claim 59 wherein division of a sub-frame into the plurality of pieces comprises dividing the sub-frame into the plurality of sub-blocks of equal size, and forming at least some pieces from a plurality of sub-blocks.

54830-2

61. The method of claim 1 wherein the boundary demarcation information for a given low level data unit comprises information that indicates whether a boundary between subframes exists within the low level data unit.

62. The method of claim 61 wherein, if such boundary does exist within the  
5 low level data unit, the boundary demarcation information further comprises information that indicates where the boundary occurs within the low level data unit.

63. The method of claim 62 wherein the information that indicates whether a boundary between high level data units exists within the low level data unit comprises a field having a value that indicates which piece in the low level data unit includes the boundary, or  
10 having a value that indicates that no boundary exists within the low level data unit.

64. The method of claim 63 wherein the information that indicates where the boundary occurs within the low level data unit comprises an offset indicating a relative position of the boundary within the piece including the boundary.

65. The method of claim 62 wherein the field is a header associated with the  
15 low-level data unit.

66. The method of claim 1 wherein at least one of the low-level data units contains a portion of at least one subframe, and at least one of the low-level data units contains a different portion of the at least one sub-frame.

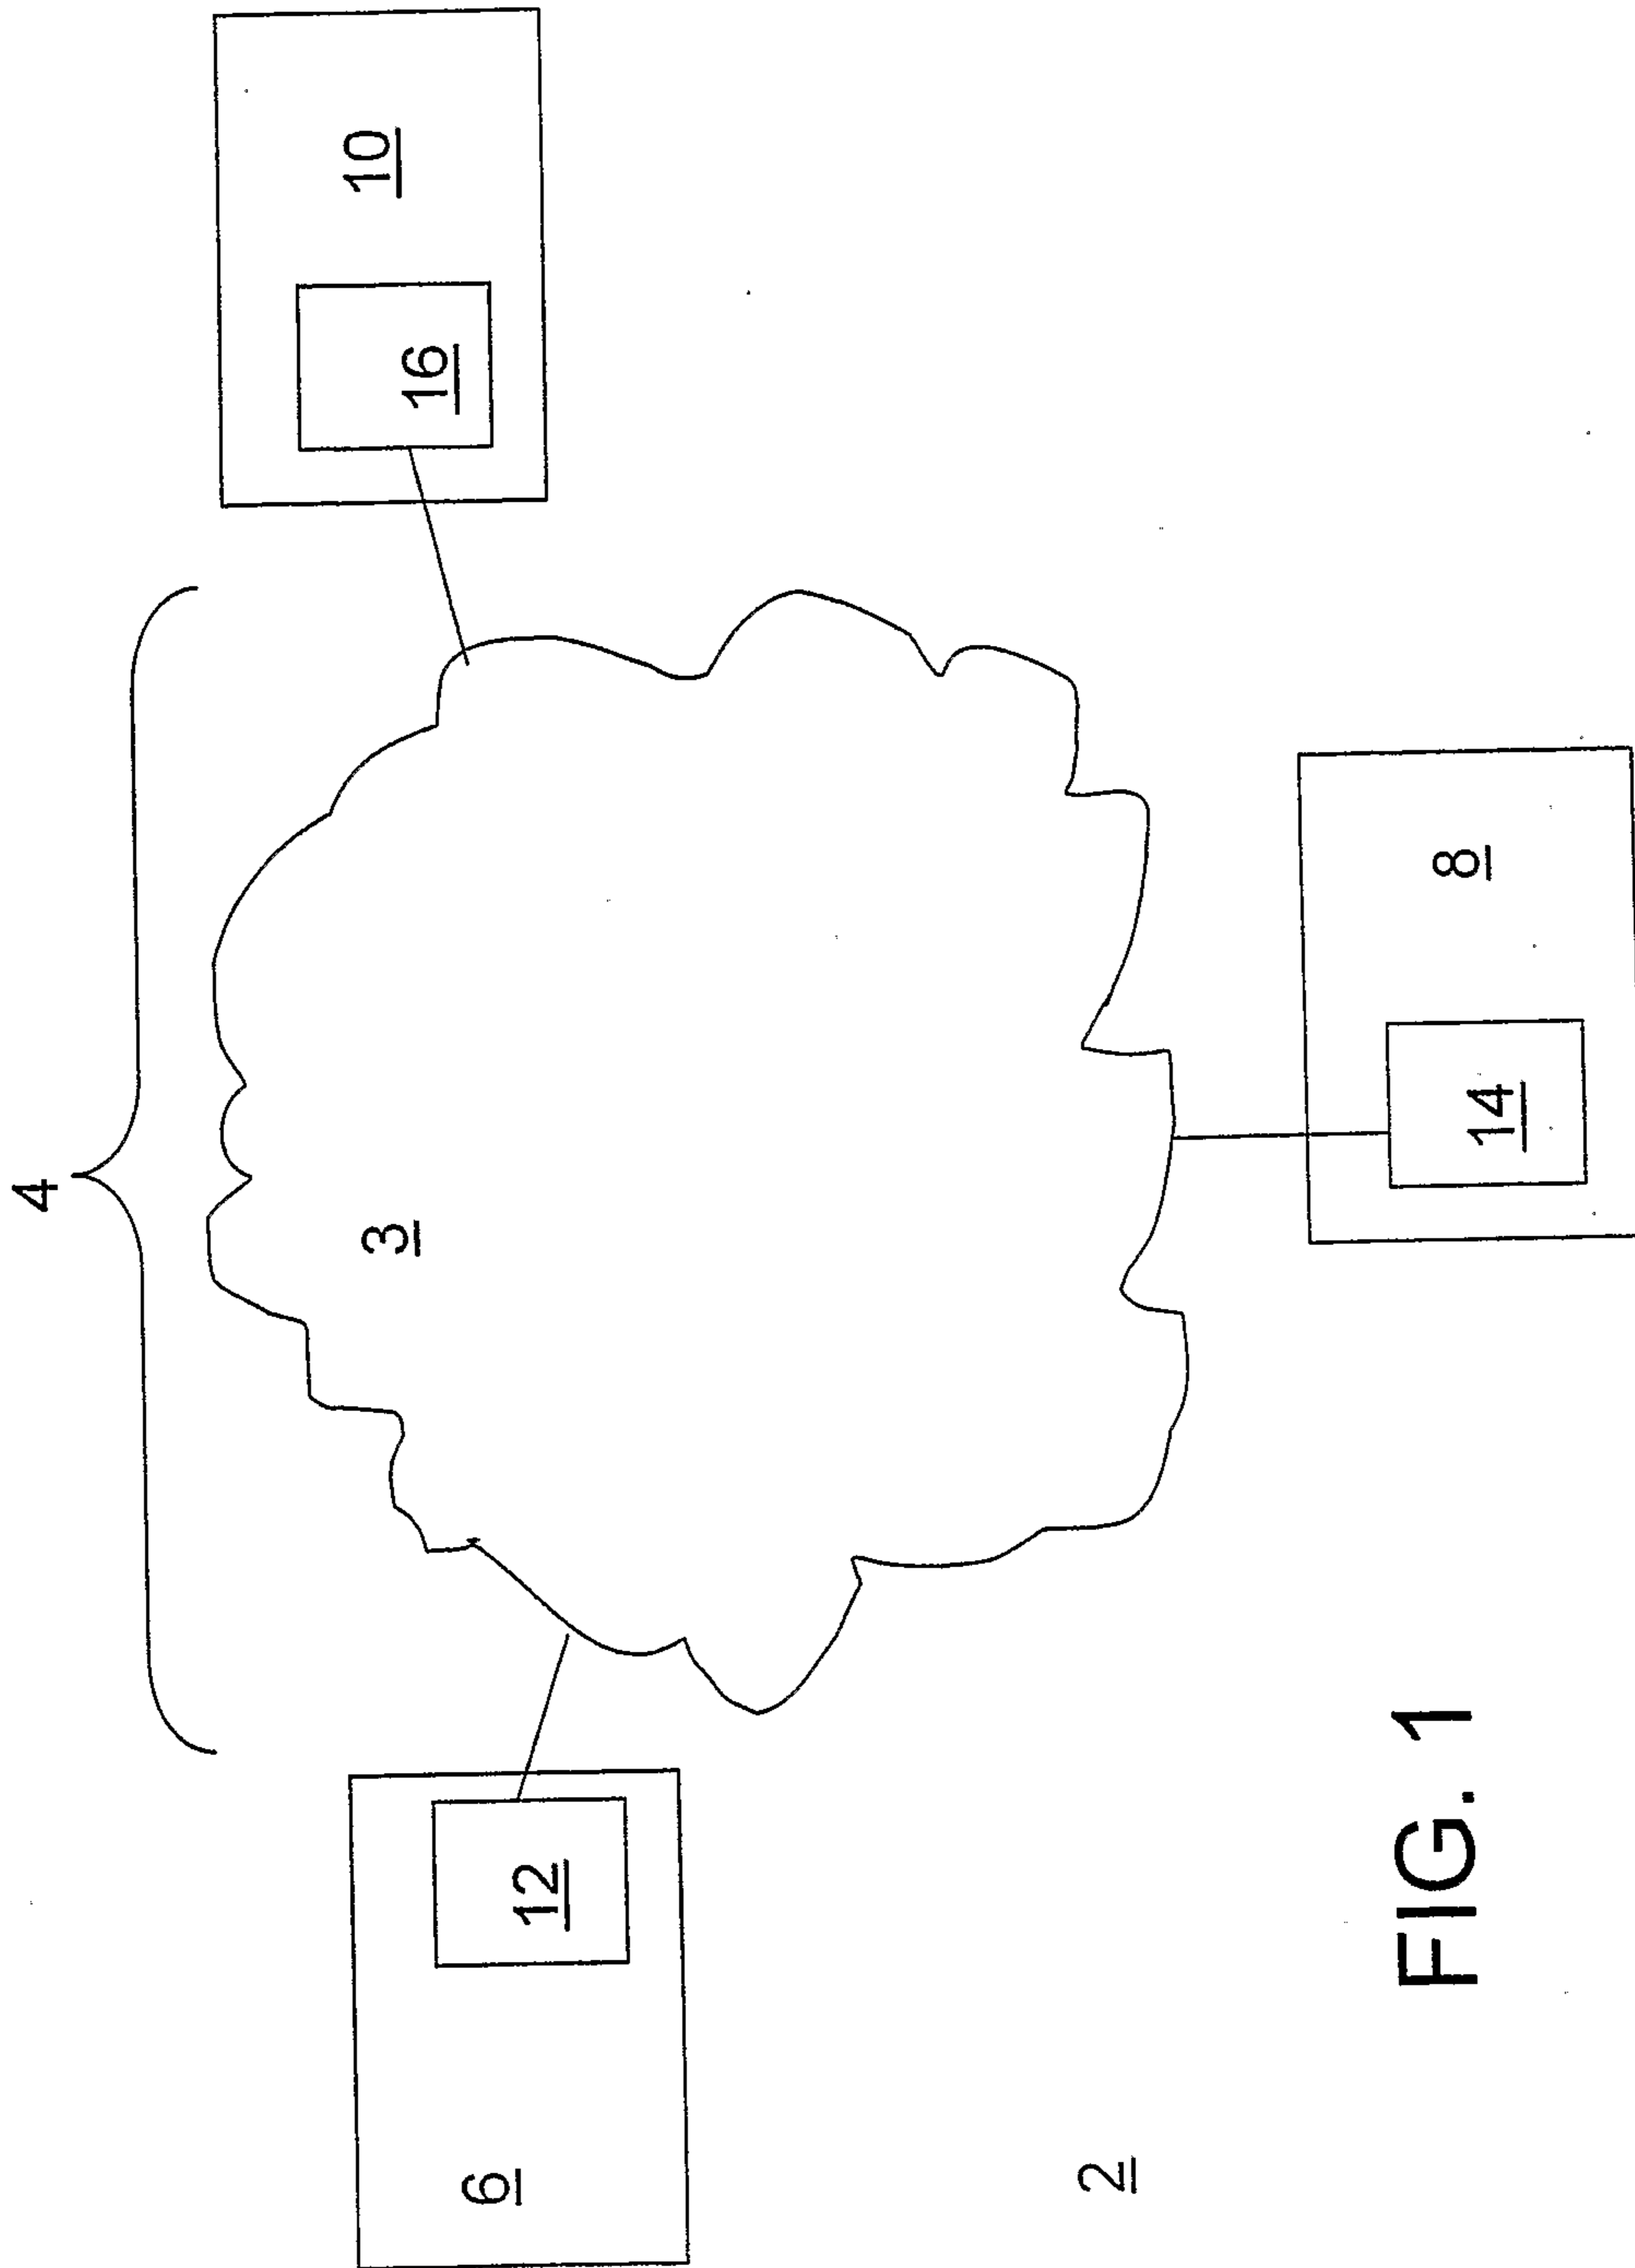


FIG. 1

2/11




FIG. 2 Reference Network Architecture




FIG. 3. MSDU Format




FIG. 4 Sub-Frame




FIG. 5 Sub-Frame Header




FIG. 6 Block of Sub-Frames protected by a single ICV




FIG. 7 Sub-Frame generation from a MSDU Payload




FIG. 8 Sub-Frame generation from multiple MSDU Payloads




FIG. 9  
MAC ENCAPSULATION

Sub-Frame Stream  
belong to the same  
(SA, DA, MSID, PLT)

200






FIG. 10 Generation of MPDU from Sub-Frame Stream




MPDU Header  
**FIG. 11**



PHY Block Format

**FIG. 12**



Reference Network Architecture