Implementation and Performance Evaluation of a New ADSL Link Aggregation Model

Daniel J. Edeen and David H.C. Du
Department of Computer Science and Engineering
University of Minnesota
edeen@cs.umn.edu, du@cs.umn.edu

***Abstract* - ADSL provides up to 8 Megabits (Mbps) downstream and 1 Mbps upstream, insufficient for many applications. The transmission of High Definition Television (HDTV) requires nearly 20 Mbps with current compression techniques, far exceeding the capabilities of a single ADSL circuit. Numerous link aggregation models have been developed to combine multiple transmission channels into a single, high-speed connection. Some of these models will operate with ADSL, but require constraints on link rates or introduce significant additional overhead. In this paper, we propose a new link aggregation model for ADSL that operates with any combination of circuit rates, supports dynamic link rate changes, and requires minimal overhead. We implemented this model using standard ADSL equipment and measured performance under a variety of conditions. Our implementation provided a data rate of 20 Mbps at 13,500 feet using four ADSL circuits, with very low latency and jitter. We also tested this model with ADSL2 and ADSL2+ transceivers, demonstrating aggregate data rates of 35 Mbps with two circuits and up to 70 Mbps with four circuits.**

I. INTRODUCTION

In recent years, tens of millions of ADSL [1] (Asymmetric Digital Subscriber Line) circuits have been deployed throughout the world, delivering high-speed access to the Internet and private data networks. ADSL operates on existing copper telephone circuits, providing higher bandwidth towards the subscriber than towards the network. The maximum bit rate varies with circuit length, ranging from 8 Mbps downstream at shorter distance to approximately 2 Mbps at 18,000 feet. Two enhancements to ADSL have recently been standardized: ADSL2 [2] and ADSL2+ [3]. The maximum downstream rate for ADSL2 is 12 Mbps, and 24 Mbps for ADSL2+. This additional bandwidth diminishes with circuit length, converging to standard ADSL link rates at 10,000 to 12,000 feet. Even with these enhancements, many applications require more bandwidth than a single ADSL circuit can deliver. The transmission of High Definition Television (HDTV) with MPEG-2 compression requires nearly 20 Mbps [4]-[6], more than twice the maximum rate of an ADSL circuit.

Link aggregation models have been developed for a variety of circuit types. These models provide higher bandwidth by combining multiple transmission channels into a single, logical connection. In this paper, we propose a new link aggregation model for ADSL that provides unlimited bandwidth, low latency and jitter, and requires minimal overhead. We introduce a novel symbol-based synchronization method that supports any combination of link rates in the group. We develop protocols to manage data partitioning and reassembly, adjust to dynamic link rate changes, and support link addition and removal.

We implemented this model using standard ADSL transceivers and measured performance under a variety of conditions. HDTV encoded at 20 Mbps was successfully transmitted over four ADSL circuits at a distance of 13,500 feet. We also tested this model with ADSL2 and ADSL2+ transceivers, demonstrating data rates as high as 70 Mbps. In addition to bandwidth requirements, some applications are sensitive to the latency and jitter of a network connection. ITU-T Recommendation G.114 [7] specifies a maximum one-way delay of 150 milliseconds for the transmission of high-quality voice service. The maximum end-to-end latency of our system was less than 37 milliseconds, and jitter ranged from 11.9 to 25.3 milliseconds.

A. Existing Link Aggregation Models

Link aggregation models have been developed for a variety of transmission channels, including HDSL, ISDN, ATM, SHDSL, Frame Relay, PPP, and Ethernet [8]-[13]. The models range from physical layer coupling with rigid link constraints, to data flow multiplexing across independent, heterogeneous connections. Each of the models share a common set of properties: data is received and partitioned, transmitted across multiple links, and reassembled at the opposite endpoint. These models also differ in a number of ways. Some require identical links, others support different link types or bit rates. Data can be partitioned at multiple levels within the network protocol hierarchy, and some models support link addition and removal. Current link aggregation models can be classified into three categories based on the layer they occupy in the ISO/OSI Network Model [14]: bonding models, frame-based models, and packet-based models.

Bonding models require identical link rates, and perform link aggregation at the physical layer, operating on bits or bytes of data. Dual duplex High-Bit-Rate Digital Subscriber Line (HDSL) combines two 784 kbps circuits into a single 1.544 Mbps DS1 circuit [15]. Digital Channel Aggregation (DCA) bonds multiple 56 kbps or 64 kbps circuits into a single logical connection [16]. Four-wire mode Single-pair High-speed Digital Subscriber Line (SHDSL) bonds two identical rate circuits into a single channel, with a total bandwidth ranging from 384 kbps to 4.624 Mbps [10].

Frame-based models operate at the data link layer and use an underlying transport protocol to perform aggregation. Inverse Multiplexing For Asynchronous Transfer Mode (IMA) [9] and Multilink Frame Relay [13] are examples. Both models require identical link rates, and are subject to all the protocol overhead associated with the transport protocols.

Packet-based models perform aggregation at the network layer. Packets from a single data flow may be transported over a single link (load balancing), or may be demultiplexed across multiple links. The links may have different transmission characteristics in some cases, and individual packets may be fragmented to improve performance. PPP Multilink (ML-PPP) [11] and 802.3ad Ethernet link aggregation [12] are two examples of packet-based models.

B. Application of Existing Models to ADSL

Several existing link aggregation models can be used with ADSL circuits, but each one introduces some limitations. Bonding is usually implemented in hardware and requires identical link rates. Bonding with ADSL would limit the link rate of every circuit in the group to the worst-case minimum of any circuit, and would require disabling features that allow dynamic link rate changes.

Frame-based models rely on an underlying transport protocol such as ATM (Asynchronous Transfer Mode) to perform data framing and synchronization, introducing significant overhead. ATM virtual circuits require more than 10% overhead for cell headers alone. Most frame-based models also require identical link rates.

Packet-based models operate at the network layer, using an end-to-end protocol such as PPP (Point-to-Point Protocol) [17] to encapsulate each packet. With an end-to-end connection, the aggregation group must be extended beyond the parallel circuits to include multiple nodes and links. The physical layer characteristics of each ADSL circuit are decoupled from the aggregation model, making it difficult to control performance of the group. If the data traffic contains mixed packet sizes, an additional fragmentation and reassembly protocol may be required to improve performance. Packet-based aggregation models cannot guarantee good performance when used with ADSL circuits.

II. PROPOSED SOLUTION

In this section we describe a new link aggregation model for ADSL that operates just above the ADSL physical layer, supports any combination of link rates, and provides nearly 100% utilization of each circuit in the group. This model requires minimal overhead, and the endpoints use local link state information to accommodate dynamic rate changes.

A. Symbol-Based Synchronization and Framing

Packets received at each endpoint are partitioned, striped across multiple symbol payloads, and transmitted over multiple ADSL circuits to the opposite endpoint. Reassembly of the original data stream requires synchronization and framing across links to identify the original byte order. We propose a new symbol-based synchronization method for this purpose.

ADSL transceivers use DMT (Discrete Multitone) to transmit data and control information. The transmission spectrum is partitioned into multiple orthogonal subcarriers, each with a fixed bandwidth of 4.3125 kHz. Data is encoded into a single DMT symbol transmitted over all subcarriers once every 250 microseconds. The symbol rate (4 kHz) is fixed regardless of link rate, but the payload capacity of each symbol varies as the link rate changes. Our synchronization model takes advantage of this property to perform aggregation with any combination of ADSL link rates. A synchronization marker is inserted periodically on each ADSL circuit, consuming the payload of a single DMT symbol per link. Markers are transmitted at approximately the same symbol period on each circuit. After each marker, a fixed number of DMT symbols containing application data blocks are transmitted, as shown in Fig. 1. Each ADSL transceiver at an endpoint may derive timing independently, leading to phase or frequency variance across links. Phase variance is compensated for in the receiver buffering, but significant frequency variance requires some compensation method.

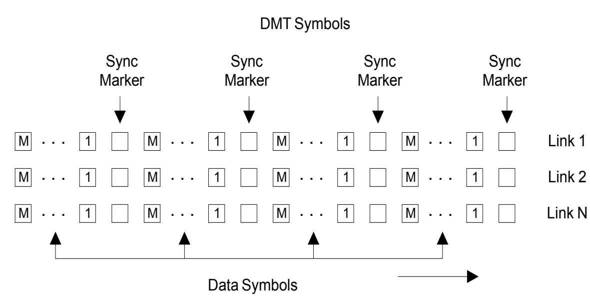

Fig. 1. Symbol-based synchronization model.

B. The Group Control Protocol

In addition to symbol-based synchronization, we introduce a framing structure to provide data sequence information to the receiver. Frames begin with the transmission of a marker on each link, followed by a fixed number of data symbols. The next frame begins with a new set of markers. The framing structure is shown Fig. 2.

To support framing and group operations, we developed a Group Control Protocol (GCP) that relies on the exchange of messages transmitted in-band between endpoints. The format of a GCP message is shown in Fig. 3. The Information Channel (IC) is used to communicate state, configuration, and control messages between endpoints. The IC uses a four byte multi-message data structure, as shown in Fig. 4.

The inbound data stream is partitioned into blocks of bytes that match the symbol payload size of each link in the group. Partitioning and link allocation occurs in round-robin link order, and a single packet may be striped across multiple links and symbols. The receiver stores each block received in individual link buffers, and reconstructs the original byte stream after a complete GCP frame has been received. This operation is shown in Fig. 5. The buffering mechanism compensates for constant variance in symbol arrival time across links. If timing across links varies in frequency, frame arrivals will drift over time. We propose an idle symbol stuffing mechanism to compensate for timing frequency variance, similar to cell stuffing in the IMA model. Compensation is requested by the receiving endpoint using the Information Channel in the GCP message structure.

A GCP group may be in one of the four different states: Down (DN), Starting (ST), Active-1 (A-1), or Active-N (A-N). The state transition diagram for a GCP group is shown in Fig. 6. Each ADSL link may be in one of five different states: Not in Group, No Sync (NGNS), Not in Group, Sync (NGS), In Group, No Sync (IGNS), In Group, Sync (IGS), or Active (ACT). The state transition diagram for a link is shown in Fig 7.

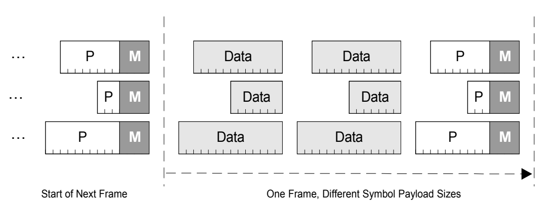

Fig. 2. GCP framing structure, with markers (M), padding (P), and data blocks.

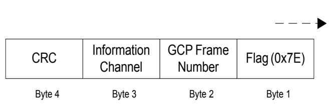

Fig. 3. GCP message format

ADSL circuits in a group are configured to operate at the maximum link rate available. When a circuit is initialized, the transceivers analyze circuit conditions and select highest link rate the circuit will support. Multiple circuits between two endpoints may operate at the same or different rates. ADSL transceivers may also change the link rate dynamically through Dynamic Rate Adaption (DRA). Each time a link rate changes, the new payload size must be reflected in the data partitioning and reassembly operations. The endpoints make these adjustments based on local link state available from the transceivers.

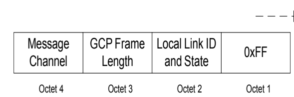

Fig. 4. Information channel message format.

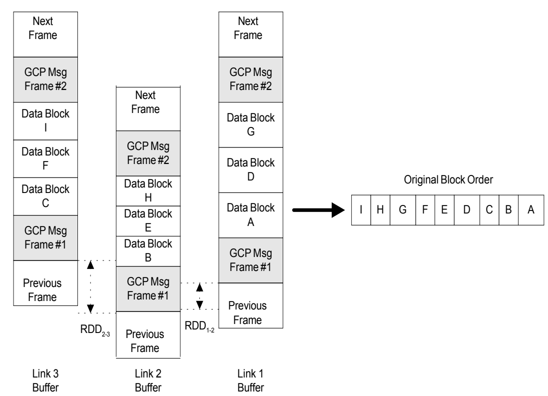

Fig. 5. Reassembly of original data stream.

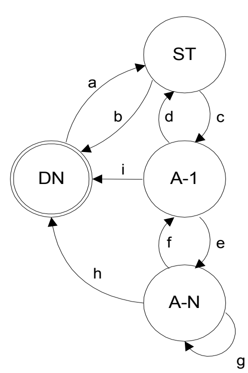

(Start) Groups start in the DOWN state.
(a) Group is started locally.
(b) Group is stopped locally.
(c) First link in group becomes active.
(d) Only link in group becomes inactive.
(e) Second link in group becomes active.
(f) One of two links in group becomes inactive.
(g) Additional links added or removed from group, total links exceeds one.
(h) Group is stopped locally.
(i) Group is stopped locally.

Fig. 6. Group state transition diagram.

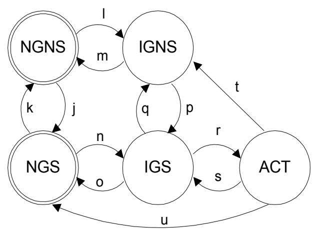

(Start) Links start in the NGNS or NGS state.
(j) (p) Link establishes ADSL sync.
(k) (q) (t) Link loses ADSL sync.
(l) (n) Link added to group locally.
(m) (o) (u) Link removed from group locally.
(r) GCP messages received on link.
(s) Far end does not put link in Active state. | GCP messages no longer received. | Far end signals error condition with IC message.

Fig. 7. Link state transition diagram.

C. *Strengths of This Model*

This link aggregation model provides higher performance and more flexibility than existing models when used with ADSL. The model supports any combination of link rates in the group, allows link addition and removal, and adapts to dynamic link rate changes. Overhead is minimal, and each link is fully utilized, regardless of link rate. The model operates with standard ADSL, ADSL2, and ADSL2+ transceivers, and does not require additional protocols such as PPP or ATM.

III. Implementation and Performance Measurements

A. *Implementation Architecture*

We implemented a four circuit version of this model with standard ADSL transmission equipment. We used two computers to provide the aggregation endpoints, each equipped with five Ethernet Network Interface Cards (NIC). Up to four NICs were connected to ADSL circuits, and the fifth was connected to a traffic generator and analyzer. Both endpoint computers ran software we developed to perform data partitioning, reassembly, and GCP protocol processing. The transceivers only supported ATM bearer channels, introducing 10.4 % overhead on each circuit for ATM cell headers. We measured throughput in this mode, and estimated performance of packet-mode transceivers by subtracting the cell header overhead. The ADSL circuits were configured as multi-protocol bridges using ATM AAL-5 encapsulation [18]. An overview of this system is shown in Fig. 8.

B. *Performance Measurements*

Network performance can be characterized by the latency, jitter, throughput, and loss rate of a connection [19]. We measured end-to-end latency and jitter with different packet sizes and link rates. Latency ranged from 23.8 to 36.3 milliseconds. Jitter ranged from 11.9 to 25.3 milliseconds. We measured throughput with ADSL, ADSL2, and ADSL2+ with different link rates, packet sizes, and circuit lengths. All tests were conducted with white noise injected on each circuit at a power level of -140 dBm/Hz. The results are shown in Figs. 9 through 16 and Table I.

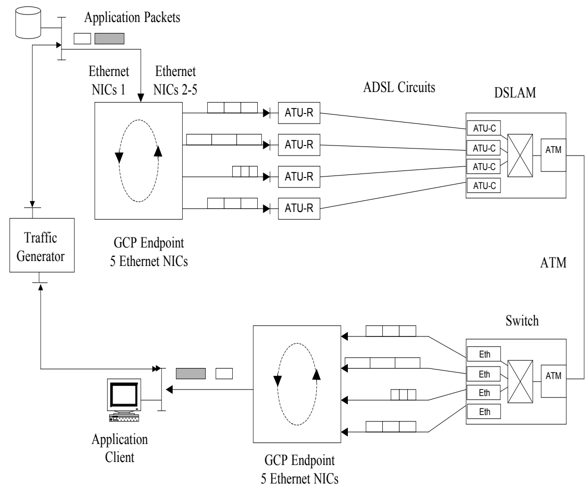

Fig. 8. Implementation architecture.

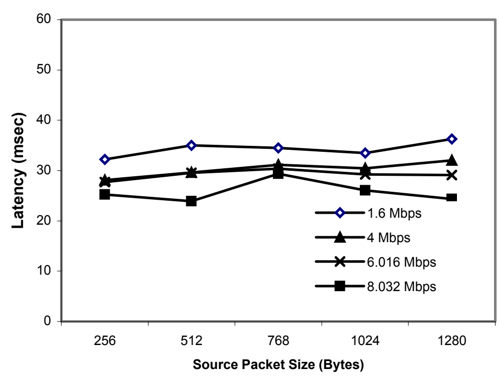

Fig. 9. Average end-to-end latency, with four ADSL circuits and various link rates.

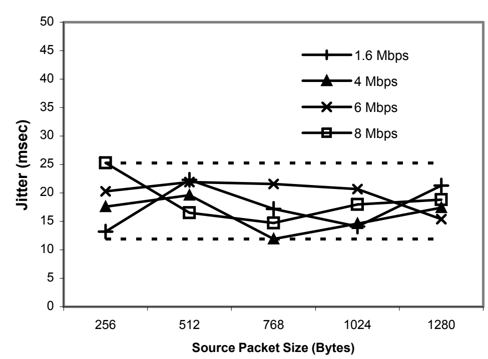

Fig. 10. End-to-end jitter, with four ADSL circuits and various link rates.

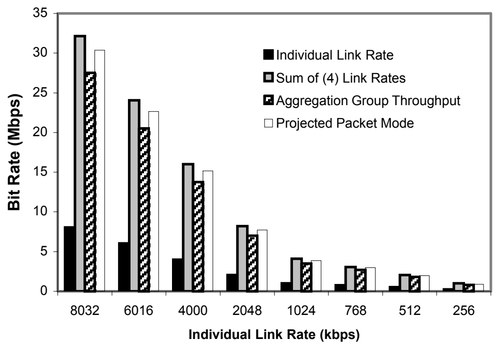

Fig. 11. Downstream throughput vs. link rate, four ADSL circuits with identical link rates.

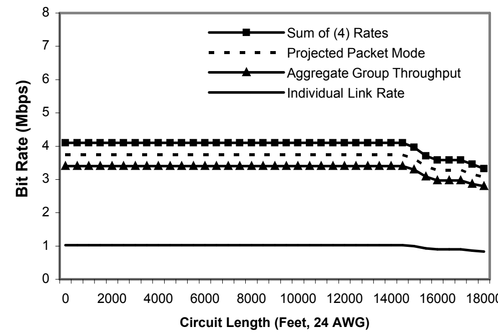

Fig. 14. Upstream throughput vs. circuit length, four ADSL circuits with identical link rates.

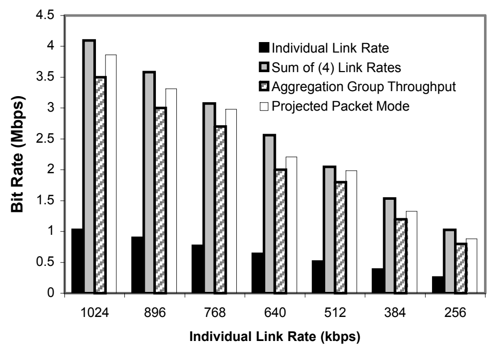

Fig. 12. Upstream throughput vs. link rate, four ADSL circuits with identical link rates.

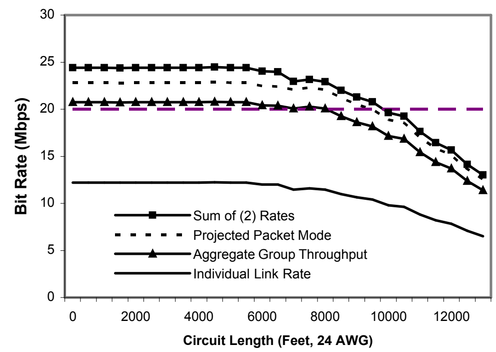

Fig. 15. Downstream throughput vs. circuit length, two ADSL2 circuits with identical link rates.

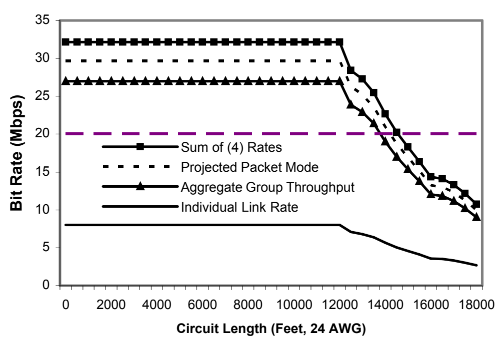

Fig. 13. Downstream throughput vs. circuit length, four ADSL circuits with identical link rates.

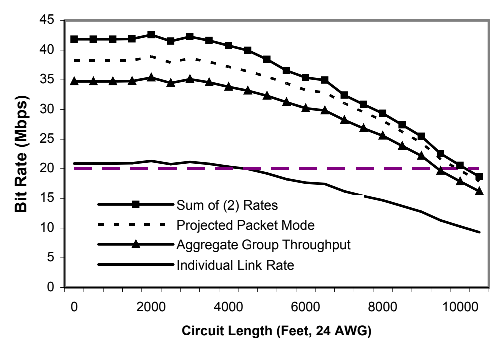

Fig. 16. Downstream throughput vs. circuit length, two ADSL2+ circuits with identical link rates.

We tested this model with a link rate variance ratio of up to 12:1, as shown in Table I. HDTV transmission at 20 Mbps was supported at a distance of 13,500 using four ADSL circuits, 8000 feet using two ADSL2 circuits, and 9500 feet using two ADSL2+ circuits (all 24 AWG wire). The dashed lines in Figs. 13, 15, and 16 indicate HDTV requirements. ADSL2 and ADSL2+ link rates converged to ADSL at 10,000 to 12,000 feet.

Our four circuit implementation demonstrated maximum aggregate bandwidth of 26.9 Mbps with ADSL, 41.4 Mbps with ADSL2, and 70.6 Mbps with ADSL2+. Based on experimental results, we calculated throughput for six and eight circuits, and estimated performance for packet mode transceivers. The results are shown in Table II.

TABLE I

DOWNSTREAM THROUGHPUT (KBPS) WITH DIFFERENT LINK RATES

Link 1&2 Rate	Link 3&4 Rate	Link Rate Ratio	Link Rate Sum - ATM Over-head	Through-put 256 Byte Packets	Through-put 1280 Byte Packets
8032	6016	1.33:1	25445	25100	23600
8032	4000	2:1	21793	20800	20300
1024	512	2:1	2782	2700	2600
8032	1984	4:1	18142	17500	16700
2048	512	4:1	4637	4200	4000
8032	992	8:1	16345	15400	14500
2560	320	8:1	5217	5000	4700
3840	320	12:1	7535	6300	6200

TABLE II

MAXIMUM DISTANCE (FEET) VS. AGGREGATION GROUP BANDWIDTH, 24 AWG, -140 DBM/HZ WHITE NOISE

Group Band-width	Num-ber of Circuits	ADSL ATM Mode	ADSL Packet Mode	ADSL 2 ATM Mode	ADSL 2 Packet Mode	ADSL 2+ ATM Mode	ADSL 2+ Packet Mode
20 Mbps	2			8000	9500	9000	9500
	4	13500	13500	13500	13500	13500	13500
	6	15500	15500	15500	15500	15500	15500
	8	17500	17500	17500	17500	17500	17500
30 Mbps	2					6000	7000
	4			11000	12000	11000	12000
	6	13500	14000	13500	14000	13500	14000
	8	15000	15500	15000	15500	15000	15500
40 Mbps	4			8000	9500	9000	9500
	6	12000	12000	12000	12500	12000	12500
	8	13500	14000	13500	14000	13500	14000
50 Mbps	4					8000	8500
	6			10500	11000	10500	11000
	8	12000	13000	12000	13000	12000	13000
100 Mbps	6					4000	5500
	8					8000	8500

IV. CONCLUSIONS AND CONTRIBUTIONS

Link aggregation models have been developed for numerous types of circuits. Some models will operate with ADSL, but impose link rate constraints, require additional overhead, or provide limited performance. We developed a new model for ADSL link aggregation that operates with any combination of link rates, and supports the addition and removal of links from the group. We introduced a novel symbol-based synchronization mechanism, and developed a simple and efficient protocol to manage this model. We constructed a four circuit implementation using standard ADSL, ADSL2, and ADSL2+ transceivers, demonstrating aggregate data rates of 26.9 Mbps, 41.4 Mbps, and 70.6 Mbps, respectively. The highest end-to-end latency measured was 36.3 milliseconds, and jitter ranged from 11.9 to 25.3 milliseconds, well within the requirements for interactive applications such as digitized voice. The aggregation model we have developed operates with existing transceivers and line codes, and extends the use of ADSL, ADSL2, and ADSL2+ to new types of applications.

REFERENCES

[1] ITU-T Recommendation G.992.1, "Asymmetric Digital Subscriber Line (ADSL) Transceivers, " July 1999.

[2] ITU-T Recommendation G.992.3, "Asymmetric Digital Subscriber Line (ADSL) Transceivers - 2 (ADSL2), " July 2002.

[3] ITU-T Recommendation G.992.5, "Asymmetric Digital Subscriber Line (ADSL) Transceivers - Extended Bandwidth ADSL2 (ADSL2plus), " May 2003.

[4] Advanced Television Standards Committee (ATSC), "Digital Television Standard (A/53), Revision C Including Amendment No. 1," May, 2004.

[5] ISO/IEC 13818-1, "Generic Coding of Moving Pictures and Associated Audio Information: Systems, " 1996.

[6] ISO/IEC 13818-2, "Generic Coding of Moving Pictures and Associated Audio Information: Video, " 1996.

[7] ITU-T Recommendation G.114, "One-way transmission time, " February 1996.

[8] ANSI T1.601-1992, "ISDN - Basic Access Interface for Use on Metallic Loops for Application at the Network Side of NT, Layer 1 Specification, " 1992.

[9] ATM Forum Technical Committee, "Inverse Multiplexing for ATM (IMA) Specification, Version 1.1, " AF-PHY-0086.001, March 1999.

[10] ITU-T Recommendation G.991.2, "Single-Pair High-Speed Digital Subscriber Line (SHDSL) Transceivers, " February 2001.

[11] Sklower, K., Lloyd, B., McGregor, G., Carr, D. and Coradetti, T., "The PPP Multilink Protocol (MP)," Network Working Group, RFC1990, August 1996.

[12] IEEE 802.3ad Standard, "Amendment to Carrier Sense Multiple Access With Collision Detection (CSMA/CD)Access Method and Physical Layer Specifications - Aggregation of Multiple Link Segments", 2000.

[13] Frame Relay Forum, "End-to-End Multilink Frame Relay Implementation Agreement," FRF.15, August 1999.

[14] Stallings, R. *Handbook of Computer-Communications Standards, Volume I: The Open Systems Interconnection (OSI) Model and OSI-related Standards*. New York: Macmillan Publishing Co., 1987.

[15] ANSI Committee T1 - Technical Report No. 28, "A Technical Report on High-Bit-Rate Digital Subscriber Lines (HDSL) ", February 1994.

[16] ISO/IEC 13871, "Information Technology – Telecommunications and Information Exchange Between Systems – Private Telecommunications Networks – Digital Channel Aggregation, " December 1995.

[17] Simpson, W., Editor, "The Point-to-Point Protocol (PPP)," Network Working Group, RFC1633, July 1994.

[18] Grossman, D. and J. Heinanen, "Multiprotocol Encapsulation over ATM Adaptation Layer 5," Network Working Group, RFC2684, September 1999.

[19] H. Zhang, "Service Disciplines For Guaranteed Performance Service in Packet-Switching Networks, *Proceedings of the IEEE*, vol. 83, no. 10, October 1995.