FACULTAD DE INGENIERIA U.N.A.M.
DIVISION DE EDUCACION CONTINUA

FACULTAD DE INGENIERIA U.N.A.M.
DIVISION DE EDUCACION CONTINUA

CENTRO DE INFORMACION Y DOCUMENTACION "ING. BRUNO MASCANZONI"

El Centro de Información y Documentación Ing. Bruno Mascanzoni tiene por objetivo satisfacer las necesidades de actualización y proporcionar una adecuada información que permita a los ingenieros, profesores y alumnos estar al tanto del estado actual del conocimiento sobre temas específicos, enfatizando las investigaciones de vanguardia de los campos de la ingeniería, tanto nacionales como extranjeras.

Es por ello que se pone a disposición de los asistentes a los cursos de la DECFI, así como del público en general los siguientes servicios:

* **Préstamo interno.**
* **Préstamo externo.**
* **Préstamo interbibliotecario.**
* **Servicio de fotocopiado.**
* **Consulta a los bancos de datos: librunam, seriunam en cd-rom.**

Los materiales a disposición son:

* **Libros.**
* **Tesis de posgrado.**
* **Noticias técnicas.**
* **Publicaciones periódicas.**
* **Publicaciones de la Academia Mexicana de Ingeniería.**
* **Notas de los cursos que se han impartido de 1980 a la fecha.**

En las áreas de ingeniería industrial, civil, electrónica, ciencias de la tierra, computación y, mecánica y eléctrica.

El CID se encuentra ubicado en el mezzanine del Palacio de Minería, lado oriente.

El horario de servicio es de 10:00 a 19:30 horas de lunes a viernes.

Palacio de Mineria Calle de Tacuba 5 Primer piso Deleg. Cuauhtémoc 06000 México, D.F. APDO. Postal M-2285
Teléfonos: 512-8955 512-5121 521-7335 521-1987 Fax 510-0573 521-4020 AL 26

FACULTAD DE INGENIERIA U.N.A.M.
DIVISION DE EDUCACION CONTINUA

A LOS ASISTENTES A LOS CURSOS

Las autoridades de la Facultad de Ingeniería, por conducto del jefe de la División de Educación Continua, otorgan una constancia de asistencia a quienes cumplan con los requisitos establecidos para cada curso.

El control de asistencia se llevará a cabo a través de la persona que le entregó las notas. Las inasistencias serán computadas por las autoridades de la División, con el fin de entregarle constancia solamente a los alumnos que tengan un mínimo de 80% de asistencias.

Pedimos a los asistentes recoger su constancia el día de la clausura. Estas se retendrán por el periodo de un año, pasado este tiempo la DECFI no se hará responsable de este documento.

Se recomienda a los asistentes participar activamente con sus ideas y experiencias, pues los cursos que ofrece la División están planeados para que los profesores expongan una tesis, pero sobre todo, para que coordinen las opiniones de todos los interesados, constituyendo verdaderos seminarios.

Es muy importante que todos los asistentes llenen y entreguen su hoja de inscripción al inicio del curso, información que servirá para integrar un directorio de asistentes, que se entregará oportunamente.

Con el objeto de mejorar los servicios que la División de Educación Continua ofrece, al final del curso deberán entregar la evaluación a través de un cuestionario diseñado para emitir juicios anónimos.

Se recomienda llenar dicha evaluación conforme los profesores impartan sus clases, a efecto de no llenar en la última sesión las evaluaciones y con esto sean más fehacientes sus apreciaciones.

Atentamente
División de Educación Continua.

Palacio de Minería Calle de Tacuba 5 Primer piso Deleg. Cuauhtémoc 06000 México, D.F. APDO. Postal M-2285
Teléfonos: 512-8955 512-5121 521-7335 521-1987 Fax 510-0573 521-4020 AL 26

PALACIO DE MINERIA

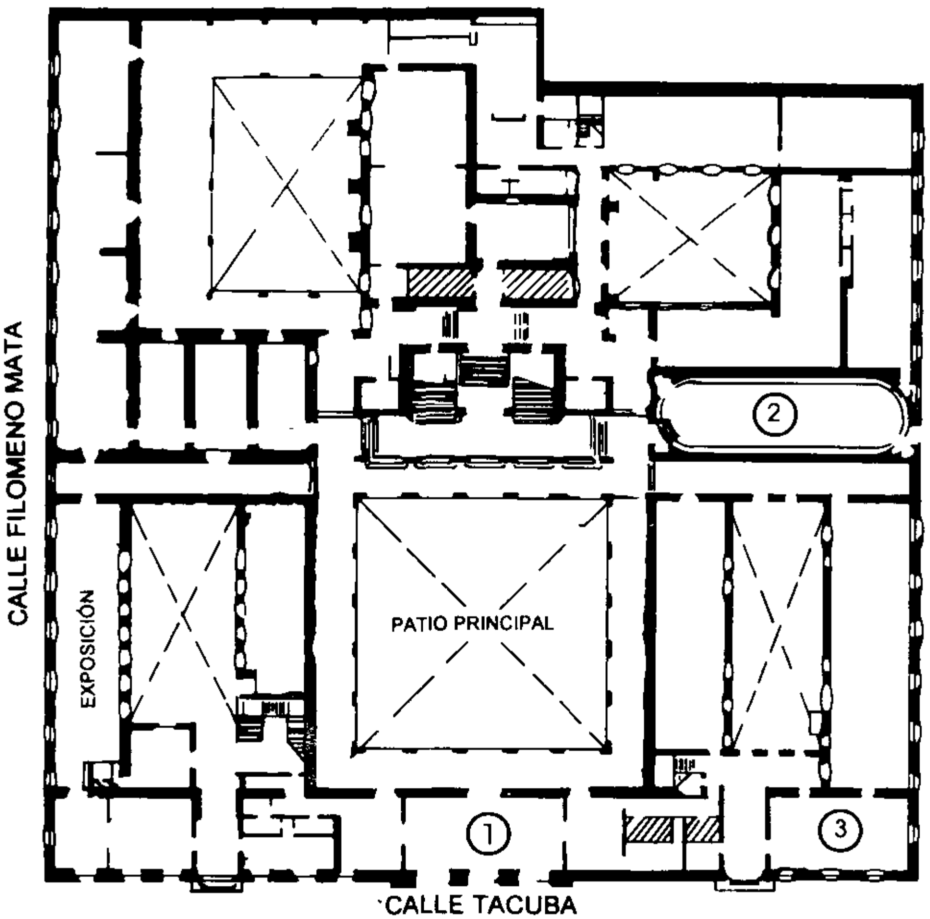

PLANTA BAJA

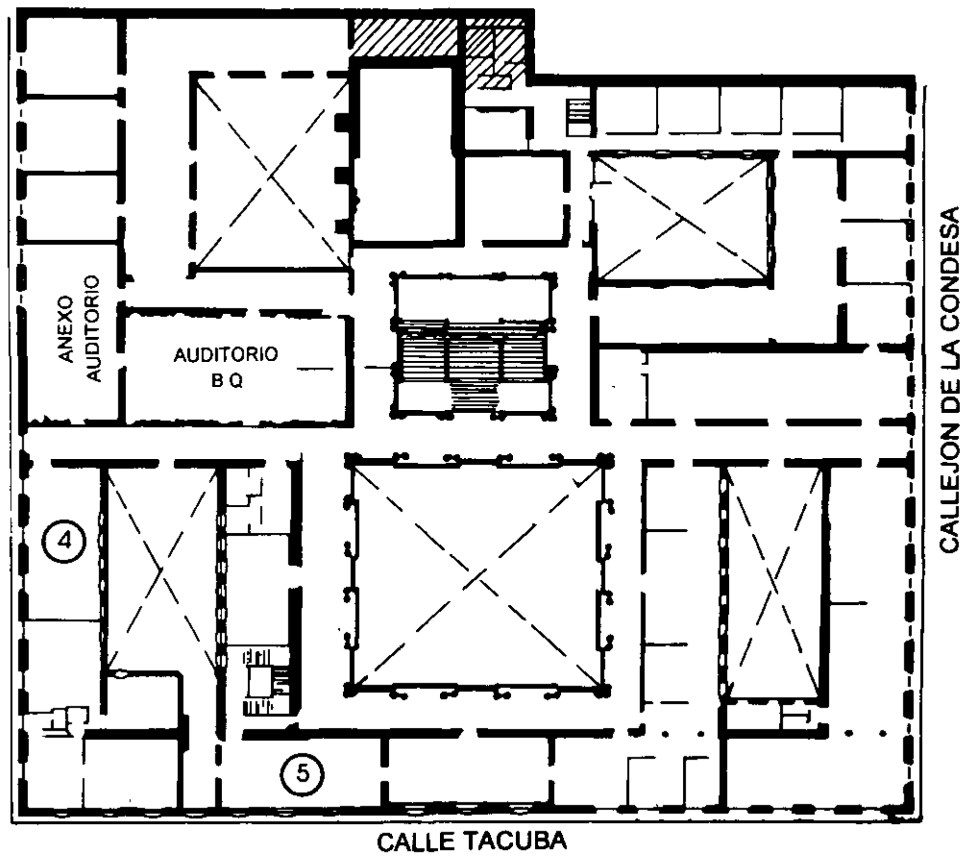

MEZZANINNE

PALACIO DE MINERIA

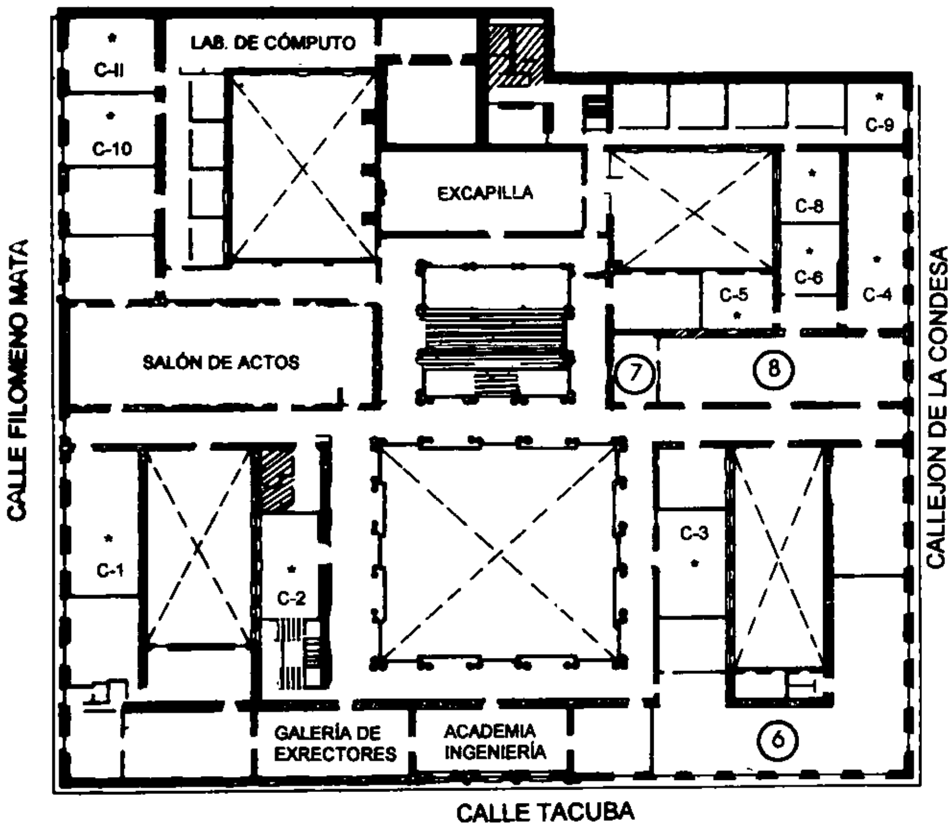

1er. PISO

GUÍA DE LOCALIZACIÓN

1. ACCESO
2. BIBLIOTECA HISTÓRICA
3. LIBRERÍA UNAM
4. CENTRO DE INFORMACIÓN Y DOCUMENTACIÓN "ING. BRUNO MASCANZONI"
5. PROGRAMA DE APOYO A LA TITULACIÓN
6. OFICINAS GENERALES
7. ENTREGA DE MATERIAL Y CONTROL DE ASISTENCIA
8. SALA DE DESCANSO

SANITARIOS

* AULAS

DIVISIÓN DE EDUCACIÓN CONTINUA
FACULTAD DE INGENIERÍA U.N.A.M.
CURSOS ABIERTOS

DIVISION DE EDUCACION CONTINUA
FACULTAD DE INGENIERIA, UNAM
CURSOS ABIERTOS

CURSO:*CC013 Redes de Alto Desempeño: Fast y Giga Ethernet, FDDI-II,...*
FECHA: *28 de febrero al 3 de marzo del 2000*

EVALUACIÓN DEL PERSONAL DOCENTE

(ESCALA DE EVALUACION: 1 A 10)

CONFERENCISTA	DOMINIO DEL TEMA	USO DE AYUDAS AUDIOVISUALES	COMUNICACIÓN CON EL ASISTENTE	PUNTUALIDAD
ING. SAUL MAGAÑA CISNEROS				

Promedio ________

EVALUACIÓN DE LA ENSEÑANZA

CONCEPTO	CALIF.
ORGANIZACIÓN Y DESARROLLO DEL CURSO	
GRADO DE PROFUNDIDAD DEL CURSO	
ACTUALIZACION DEL CURSO	
APLICACIÓN PRACTICA DEL CURSO	

Promedio ________

EVALUACIÓN DEL CURSO

CONCEPTO	CALIF.
CUMPLIMIENTO DE LOS OBJETIVOS DEL CURSO	
CONTINUIDAD EN LOS TEMAS	
CALIDAD DEL MATERIAL DIDÁCTICO UTILIZADO	

Promedio ________

Evaluación total del curso ________

Continúa...2

1. ¿Le agradó su estancia en la División de Educación Continua?

SI [ ] NO [ ]

Si indica que "NO" diga porqué: ____________________

2. Medio a través del cual se enteró del curso:

Periódico *La Jornada*	
Folleto anual	
Folleto del curso	
Gaceta UNAM	
Revistas técnicas	
Otro medio (Indique cuál)	

3. ¿Qué cambios sugeriría al curso para mejorarlo?

____

____

____

____

____

____

4 ¿Recomendaría el curso a otra(s) persona(s) ?

SI [ ] NO [ ]

5.¿Qué cursos sugiere que imparta la División de Educación Continua?

____

____

____

____

____

____

6. Otras sugerencias:

____

____

____

____

____

FACULTAD DE INGENIERIA U.N.A.M.
DIVISION DE EDUCACION CONTINUA

DIPLOMADO EN REDES DE COMPUTADORAS

MODULO IV

REDES DE ALTO DESEMPEÑO:
FAST Y GIGA ETHERNET, FDDI-II, "SWITCHING", ATM Y FRAME RELAY

FEBRERO - MARZO DEL 2000

Palacio de Minería Calle de Tacuba 5 Primer piso Deleg. Cuauhtémoc 06000 México, D.F. Tel.: 521-40-20 Apdo. Postal M-2285

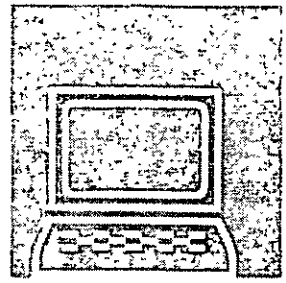

DIVISION DE EDUCACION CONTINUA DE LA FACULTAD DE INGENIERIA

DIPLOMADO EN REDES DE COMPUTADORAS (LAN, WAN y GAN)

Redes de Alto Desempeño : FAST y GIGA ETHERNET, FDDI-II, ATM y FRAMERELAY

Coord. Académico: Ing. Saúl S. Magaña Cisneros

MODULO 4

DIPLOMADO EN REDES DE COMPUTADORAS (LAN, WAN Y GAN)

REDES DE ALTO DESEMPEÑO: FAST y GIGA ETHERNET FDDI-II, SWITCHING, ATM y FRAMERELAY

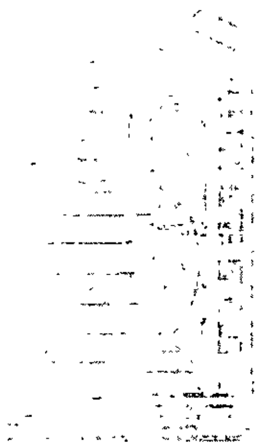

Presentación

Coord. Académico: Ing. Saúl S. Magaña Cisneros

Enero 1999

Redes de Alto Desempeño: FAST y GIGA ETHERNET, FDDI-II, "SWITCHING", ATM Y FRAME RELAY

Temas

1 FAST ETHERNET y GIGA ETHERNET

2 FDDI, FDDI II

3 SWITCHES

4 ATM

5 FRAME RELAY

6 APLICACIONES

TEMARIO

1.- FAST ETHERNET y GIGA ETHERNET

- Introducción
- Caracteristicas de 100BaseT
- Estándares y Normalización
- Tipos de cableado
- Características de los dispositivos Fast-Ethernet
- Características de los dispositivos GIGA-Ethernet
- Redes Conmutadas
- Alternativas de implementación

2.- FDDI, FDDI II

- Introducción
- Fibras ópticas
- Backbones
- Antecedentes de FDDI y FDDI-II
- Características
- Funcionamiento
- Dispositivos
- Normalización

3.- SWITCHES

- Introducción
- Caracteristicas
- Tecnologías Store and Forward y Cut-Through
- Switches ATM y Switches Ethernet

4.- ATM

- Introducción
- Componentes
- Servicios
- Estructura de la celda
- Modelo B-ISDN
- Niveles de adaptación, convergencia y físico
- Aplicaciones y casos de estudio

5.- FRAME RELAY

- Tecnologías antecesoras
- Terminologia y funcionamiento
- Estructura de frame
- Administración de la congestión
- Técnicas de reducción de tráfico
- Interfaces de administración local
- Estándares
- Aplicaciones y casos de estudio

6.- APLICACIONES

- Redes Virtuales
- Redes Multimedia
- Vídeo Conferencia
- Integración total de Redes; LAN=MAN=WAN=GAN

Enero 1999

Redes de Alto Desempeño: FAST y GIGA ETHERNET, FDDI-II, "SWITCHING", ATM Y FRAME RELAY

Temas

1 FAST ETHERNET y GIGA ETHERNET

2 FDDI, FDDI II

3 SWITCHES

4 ATM

5 FRAME RELAY

6 APLICACIONES

PRESENTACION

La constante evolución en las tecnologías de las redes de cómputo, las comunicaciones y las telecomunicaciones ha tenido como objetivo central incrementar su rendimiento, esto es crear **REDES DE ALTO DESEMPEÑO** para poder satisfacer las nuevas necesidades de los usuarios: Transmisión de grandes volúmenes de información, de datos, de voz, de vídeo a altas velocidades y cubriendo grandes distancias, permitiendo que el que hacer del hombre en este campo, cada día acorte el tiempo, mejore la seguridad en sus aplicaciones e incremente su productividad.

Las nuevas redes de alto rendimiento serán conformadas por enlaces locales basados en Fast y Giga Ethernet o FDDI-II, las comunicaciones entre redes estarán sustentadas en los servicios de **Cell Relay** que derivan en la tecnologia **ATM** y los enlaces remotos soportados por Freme-Relay, todo integrado en poderosos switches de niveles 2 y 3. Estas tecnologías de vanguardia, el día de hoy nos permiten alcanzar velocidades de transmisión de 622 mbps y "backplanes" de 4Gbps y su desempeño se seguirá incrementando

La marcada evolución en la tecnología de las Redes a ido acompañada de un alto desarrollo en los medios comunicación como hoy lo son los enlaces basados en fibras ópticas y cables telefónicos de altas velocidades como FTP y UTP niveles 6 y 7 además de los servicios ofrecidos por las compañías telefónicas como **ISDN** y **B-ISDN** los cuales nos ofrecen integración de múltiples servicios (**voz**, **datos**, **imagen** y **sonido**) gracias a sus amplios anchos de banda. Combinando ambas innovaciones, surgen fuertemente a partir de 1995, las **REDES DE ALTO DESEMPEÑO** que definen a las **Redes de cuarta generación.**

Las **REDES DE ALTO DESEMPEÑO** con sus elementos de comunicación, implican una serie de tecnologias y arquitectura modernas y avanzadas, que generan la necesidad del conocimiento y dominio de las mismas, y esto es imperante. Se requiere por lo tanto, de especialistas y ejecutivos bien capacitados y bien informados respectivamente, para un soporte técnico y toma de decisiones adecuados en este profundo y apasionante campo de las Redes. Conscientes esta necesidad, ofrecemos este curso como un módulo más del Diplomado, y/o como una oportunidad de actualización, tratando de lograr los siguientes

OBJETIVOS

Introducir a los participantes en las tecnologías de los Servicios Integrados de Redes Digitales de Banda Ancha (B-ISDN) y dar a conocer los nuevos estándares de las tecnologías de redes de alto desempeño.

Que el participante conozca los antecedentes y conceptos de las tecnologías Fast y Giga Ethernet, FDDI-II, "SWITCHING", CELL RELAY, ATM y FRAME RELAY para poder aplicarlo para la toma de decisiones o la implantación de estas tecnologías, según las condiciones del mercado mexicano.

A QUIEN VA DIRIGIDO

A todos aquellos profesionales y profesionistas que por sus necesidades laborales, estén involucrados con las Redes de Cómputo y requieran actualizarse en las Redes de Alto Desempeño, y a los Ejecutivos que necesiten bases técnicas en su responsabilidad de toma de decisiones.

REQUISITOS

Los participantes que estén sustentando el diplomado haber cursado al menos los módulos I y II

Para los participantes que tomen este modulo como un curso abierto, es necesario tener un buen nivel en microcomputacion y conocimientos avanzados en redes de computadoras y comunicaciones

Nuevas tecnologias de redes de computadoras

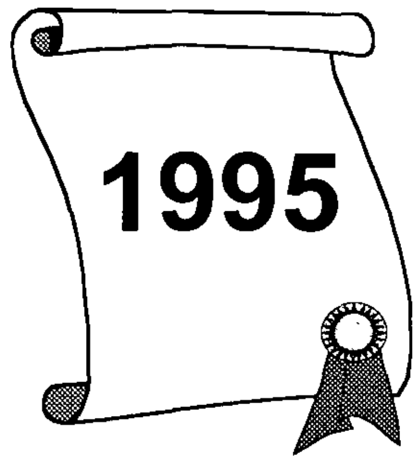

Notas:

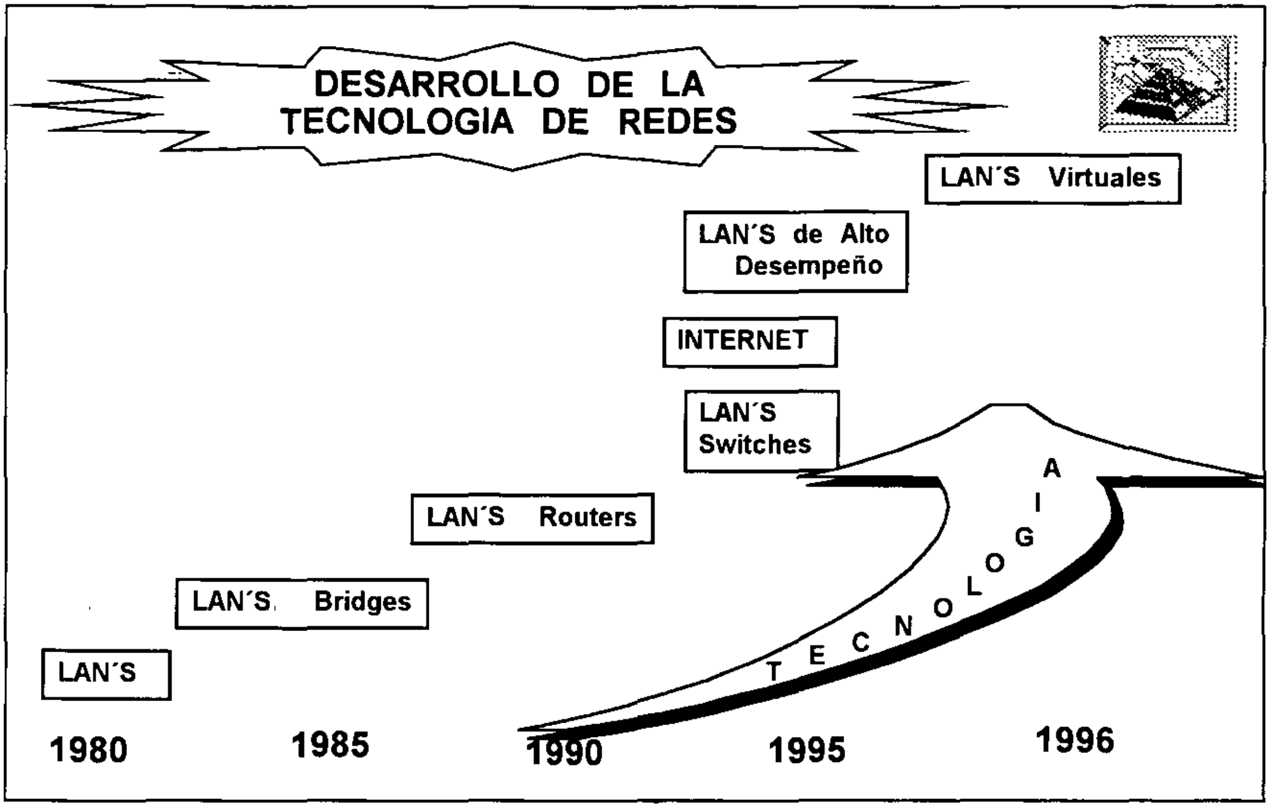
DESARROLLO DE LA TECNOLOGIA DE REDES
LAN´S Virtuales
LAN´S de Alto Desempeño
INTERNET
LAN´S Switches
LAN´S Routers
LAN´S Bridges
LAN´S
TECNOLOGIA
1980
1985
1990
1995
1996

Notas:

Redes de alto desempeño

- FDDI, FDDI - II
- FAST ETHERNET
- TECNOLOGIA SWITCHING
- ATM
- FRAME RELAY
- B - ISND

→

- REDES VIRTUALES
- REDES MULTIMEDIA
  VIDEOCONFERENCIAS

REDES

LAN = MAN = WAN = GAN

Notas:

COMUNICACION DIGITAL

- BANDA BASE
- BANDA ANCHA

Notas:

BANDA BASE

Características:

- Un solo canal
- Bajo costo
- Se modula y demodula la señal
- Utilizada por los estándares actuales de REDES locales

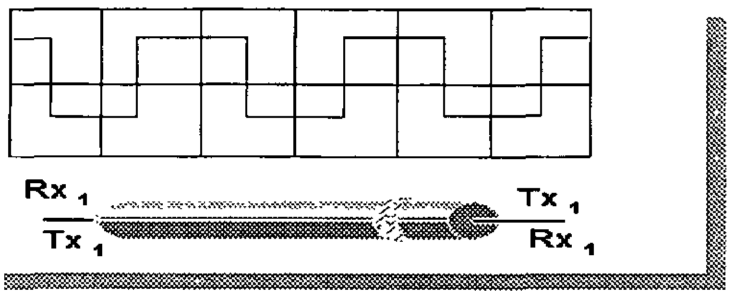

Notas:

TECNOLOGIAS EN SISTEMAS DE BANDA ANCHA

BANDA ANCHA

Características:

- **Varios Canales Paralelos**
- **Multiplexaje por Frecuencia**
- **—▷Un canal de Transmisión**
- **◁—Un Canal de Recepción**

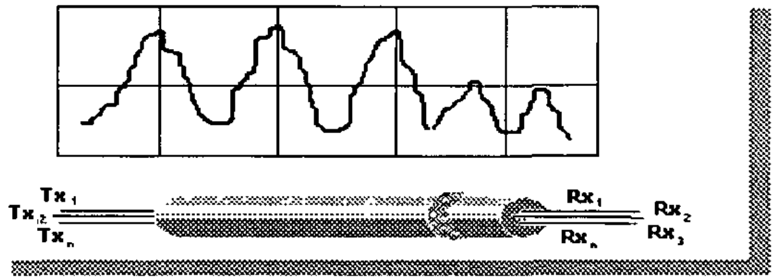

Notas:

SERVICIOS CONMUTADOS DE ALTA VELOCIDAD

Alta Velocidad:

- **ISDN** Integrated Service Digital Network
- **B-ISDN** Broadband-Integrated Service Digital Network

Notas:

ISDN

Acceso a los servicios de telecomunicaciones sin ISDN

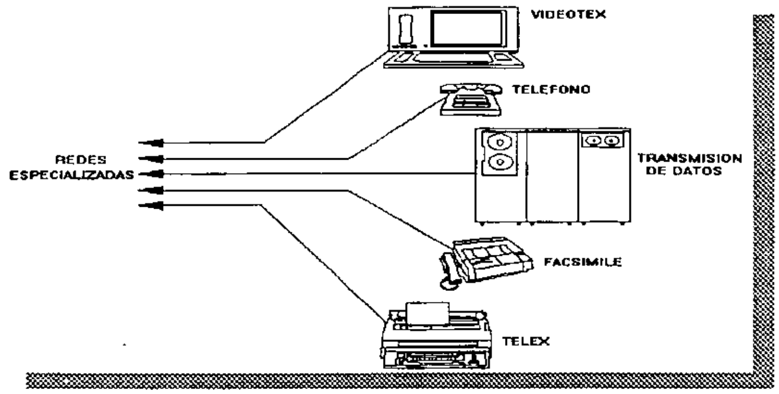

Notas:

ISDN

Acceso a los servicios de telecomunicaciones con ISDN

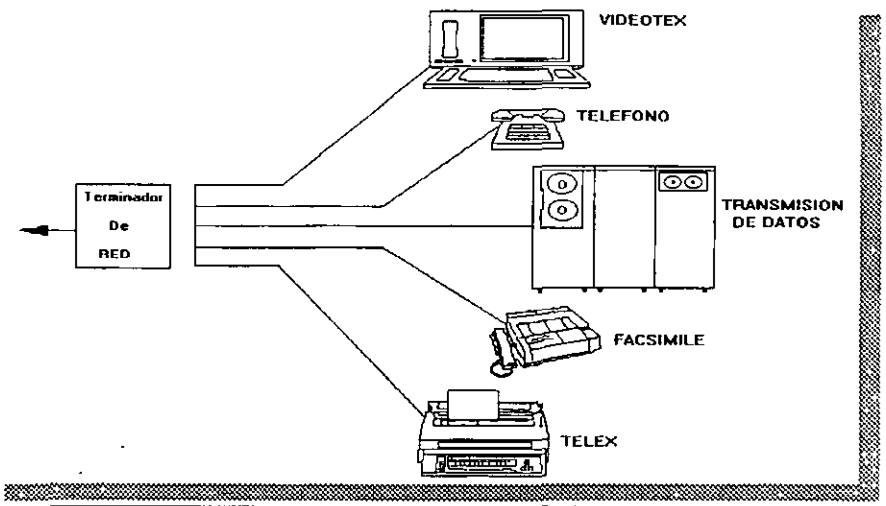

Notas:

ISDN

Acceso Básico

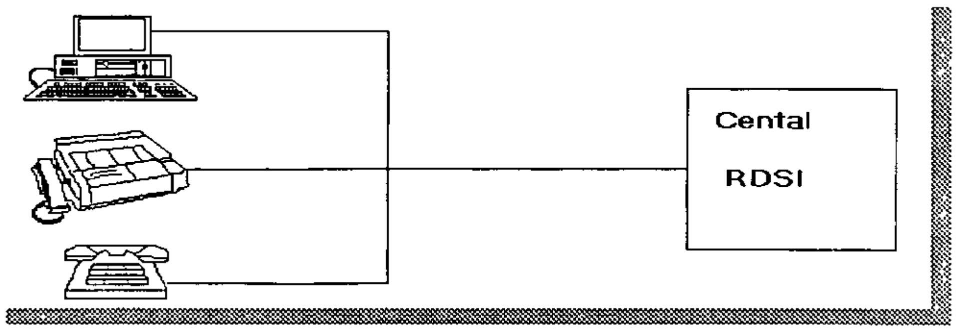

Notas:

ISDN

Acceso Primario

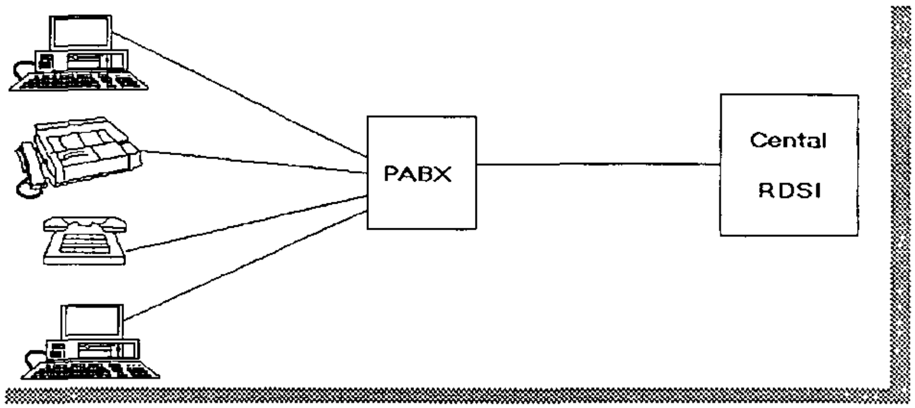

Notas:

ISDN Velocidades

Canal	Velocidad de Transmisión	Asociado A
B	64 Kbps	ISDN
D	16 Kbps y 64 Kbps	ISDN
E	64 Kbps	ISDN
H0	384 Kbps = 6B	BISDN
H11	1536 kbps = 24B	BISDN
H12	1920 Kbps = 30B	BISDN
.		
.		
.		
H4	120 a 140 Kbps	BISDN

Ejemplo: Canal 23B+D = 23X64 Kbps + 64 Kbps

Notas:

TECNOLOGIAS EN SISTEMAS DE BANDA ANCHA

INTRODUCCION

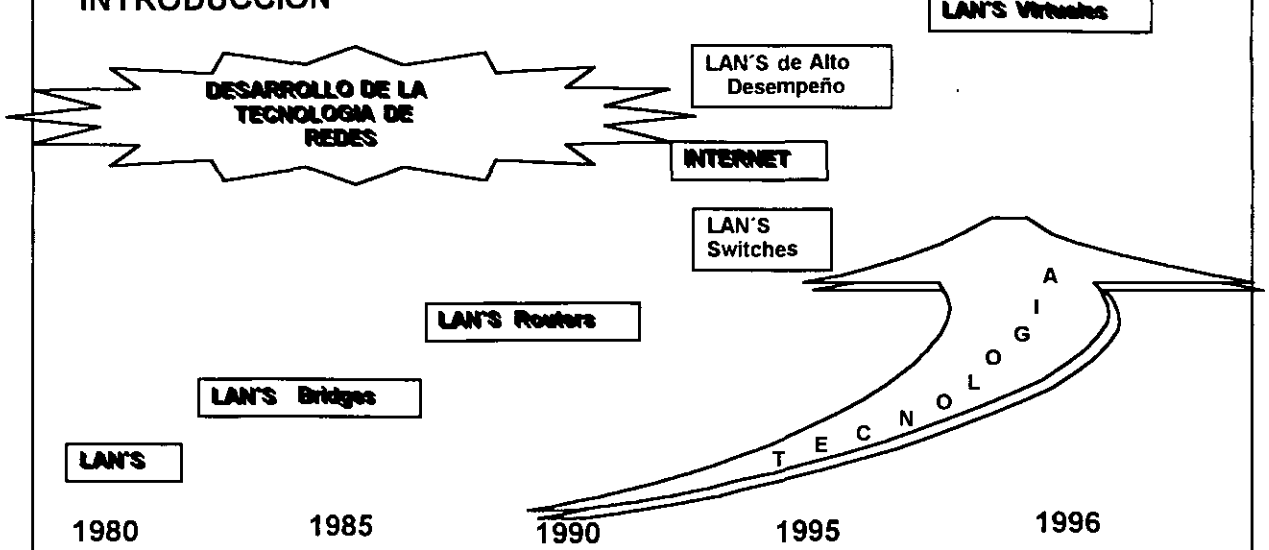

Notas:

Redes de alto desempeño

- FDDI, FDDI - II
- FAST ETHERNET
- TECNOLOGIA SWITCHING
- ATM
- FRAME RELAY
- B - ISDN

→

- REDES VIRTUALES
- REDES MULTIMEDIA
  VIDEOCONFERENCIAS

REDES
LAN = MAN = WAN = GAN

Notas:

B-ISDN. Estándares

- En 1988 se establece la recomendación I.121 del CCITT.
- En 1990 el grupo de estudio XVIII aprueba 13 recomendaciones básicas, entre ellas:
  - Aspectos generales de B-ISDN
  - Servicios específicos de Red
  - Características fundamentales de ATM
  - Aplicaciones ATM
  - Operación y mantenimiento de los accesos a B-ISDN
- A partir de 1992, se han generado nuevas recomendaciones y grupos de estudio, entre ellas la I.113 de vocabulario y términos.

Notas:

ORGANIZACIONES INVOLUCRADAS EN LA ESTANDARIZACION DE -ISDN

A nivel mundial

C C I T T	Comité Consultivo Internacional de Telegrafía y Telefonía
I S O	Internacional Standards Organizatión

En Europa

C E P T	European Conference of Posts and Telecommunications Administrations
E T S I	European Telecommunications Standards Institute

En Estados Unidos

A N S I	American National Standard Institute
EIA	Electronic Industries Association
BELLCORE	Bell Communications Research

Notas:

TECNOLOGIAS EN SISTEMAS DE BANDA ANCHA

B-ISDN.- INTRODUCCION

Diseñada para soportar conmutación de acuerdo a la demanda y conexiones en banda ancha tanto permanentes como semipermanentes para las aplicaciones punto-a-punto y punto-a-multipunto.

Soporta servicios de conmutación de circuitos y de conmutación de paquetes, aplicaciones "single media", "mixed-media" y "multimedia".

Notas:

BISDN .- CARACTERISTICAS

Conexiones conmutadas por demanda en Banda Ancha

- **Permanentes**
- **Semipermanentes**

Aplicaciones

- **Punto a punto**
- **Punto a multipunto**

Notas:

BISDN .- CARACTERISTICAS

Modos de Conmutación

- Paquetes
- Circuitos

Naturaleza de Servicios

- "Connection - oriented"
- "Connectionless"

Configuraciones

- Unidireccionales
- Bidireccionales

Notas:

BISDN. Características

Tráfico

- Velocidad constante CBR
  (Constant Bit Rate)
  - Sin negociación de velocidad
- Velocidad variable VBR
  (Variable Bit Rate)
  - Con negociación de velocidad

Notas:

TECNOLOGIAS EN SISTEMAS DE BANDA ANCHA

BISDN CARACTERISITCAS

- Conmutación por demanda
- Conexiones permanentes y semimermanentes
  - Punto a Punto
  - Punto a multipunto
- Conmutación de paquetes y conmutación de circuitos
  - Single media
  - Mexed media
  - Multimedia
  - "Conection less" y "Conection-oriented"
  - VBR y CBR

Notas:

ISDN.- TERMINOLOGIA:

- Grupos Funcionales.
- Puntos de referencia.

Notas:

ISDN.- TERMINOLOGIA:

Grupos Funcionales.

- **Terminadores de Red 1 (NT1).**
  Funciones equivalentes a las del nivel 1 del modelo de referencia OSI.
- **Teminadores de Red 2 (NT2)**
  Funciones equivalentes a las de los niveles 1, 2 y 3 del modelo OSI.
- **Equipo Terminal (TE)**
  Teléfonos digitales, Equipos terminales de datos y estaciones de trabajo que integran voz y datos.

Notas:

ISDN.- TERMINOLOGIA:

Grupos Funcionales.

- **Equipo terminal tipo 2 (TE2)**
  Equipo terminal con interfaces no-ISDN
- **Adaptador terminal (TA)**
  Grupo funcional que incluye las funciones para conectar equipo TE2 dentro de ISDN.

Notas:

ISDN.- TERMINOLOGIA:

Puntos de Referencia:

R: Interface funcional entre un grupo TE2 y un TA.

T: Interface entre el equipo NT2 y el NT1.

S: Interface entre equipos de usuario como pueden ser los TA o los TE1 y el equipo NT2.

U: Interface del lado de la red del equipo NT1.

Notas:

ISDN.- TERMINOLOGIA:

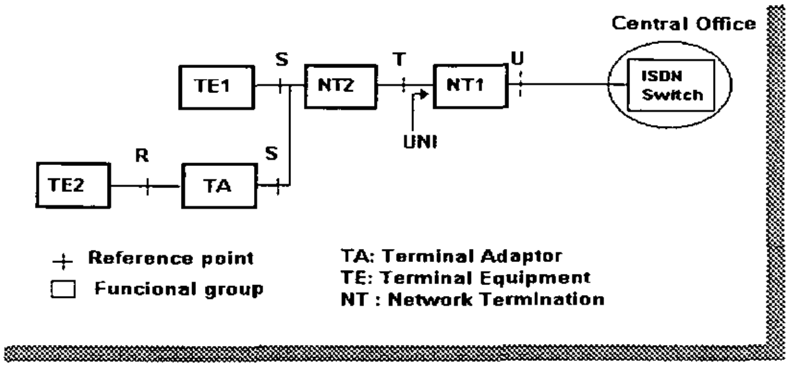

Notas:

ISDN.- EQUIPO

Canales de Acceso:

- **Canal B: 64Kbps para voz, datos en conmutación de circuitos o datos en conmutación de paquetes (B= bearer "portadora")**
- **Canal D: 16 ó 64Kbps para señalización, control o información del cliente en paquetes (D=delta).**
- **Canal H: 384Kbps (H0), 1,536Mbps (H11) ó 1,920 Mbps (H12) para teleconferencias, datos en alta velocidad o audio de alta calidad.**

Notas:

ISDN.- EQUIPO

UNI: User Network Interface

- **Basic Rate Access (o BRI basic rate interface).**
  Interface de usuario que provee 2 canales B y un canal D (2B+D).

- **Primary Rate Access (o PRI primary rate interface)**
  Interface de usuario que provee 23 canales B y un canal D (23B+D).

- **Para canales H se prevee que en el futuro se utlice una interface de red tipo H+D.**

Notas:

CAPITULO 1

DIPLOMADO EN REDES DE COMPUTADORAS (LAN, WAN Y GAN)

REDES DE ALTO DESEMPEÑO: FAST y GIGA ETHERNET FDDI-II, SWITCHING, ATM y FRAMERELAY

Fast Ethernet y Giga Ethernet

Fast Ethernet

REQUERIMENTOS DE ALTA VELOCIDAD Y SOLUCIONES PROPUESTAS.

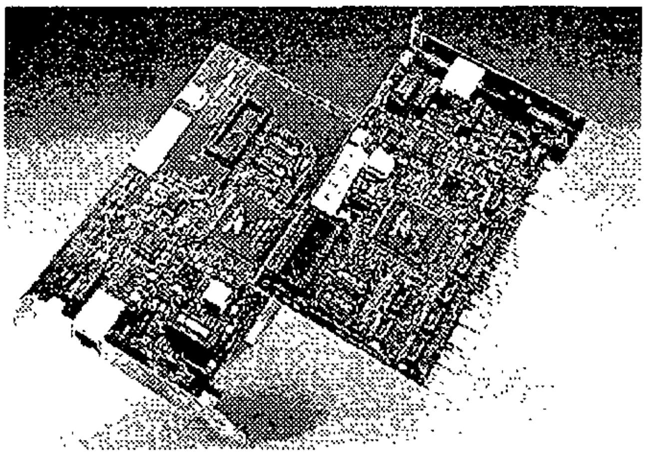

Día con día, cada vez más usuarios de PC´s se agregan a las redes. Al final de 1994 solo el 40% de las PC ´s en el mundo estaban conectadas en redes. Al mismo tiempo, la tecnología estaba logrando avances significativos como el lanzamiento comercial de el INTEL PENTIUM y tecnologías como POWER PC, tecnologías de sistemas de almacenamiento en disco duro avanzadas que decrementaban los costos, con el objeto de dar potencia a aplicaciones de redes basadas en PC´s de propósito crítico, aplicaciones que hasta recientemente han sido posibles solo en un mainframe.

La capacidad de las PC´s ha crecido en forma exponencial, al igual que las aplicaciones que corren en éstas, por lo que las tecnologías para conectar las PC´s entre si, empiezan a ser un factor determinante en la funcionalidad de las redes locales.

Aunque no todos los usuarios requieren una red con capacidad de 100 mbps. muchas aplicaciones "lan-intensive" ya empujan los 10 mbps existentes y pueden beneficiarse con la tecnología actual de 100 mbps

Surgieron aplicaciones de datos intensivos como multimedia, trabajo en grupo y bases de datos cliente-servidor, que pronto harán de los 100mbps parte crítica de la mayoría de las Lan´s.

Así mismo, como los servidores de red son ahora mas poderosos, han sido reubicados de conexiones locales a centrales de datos, donde necesitan conexiones de alta velocidad a 100 mbps al "backbone" para proporcionar capacidad centralizada al costo óptimo.

¿Que tecnología está mejor situada dentro del crecimiento de los requerimientos de alta velocidad de las redes de hoy?

La respuesta depende del usuario y de las necesidades de la red. **FAST ETHERNET** es una excelente alternativa por las siguientes razones:

ventajas de Fast Ethernet

- ◊ Alto rendimiento.
- ◊ Tecnología basada en estándares.
- ◊ Migración a costo aceptable con máximo aprovechamiento del equipo ya existente (infraestructura de cableado, sistemas de administración de red etc...)
- ◊ Soporte de los principales vendedores en todas las áreas de productos de red.
- ◊ Costo óptimo.

↳ Alto rendimiento.

Una de las mejores razones para cambiar a fast ethernet para grupos de trabajo, es la disponibilidad de manejo de ambas demandas agregadas, de una red multiusuario y el excesivo tráfico ocasionado por el alto desempeño de las PC´s y las sofisticadas aplicaciones empleadas. Fast Ethernet es la solución óptima para grupos de trabajo.

↳ Tecnología basada en estándares.

Fast Ethernet está diseñada para ser la evolución más directa y simple de ethernet 10 base-T, la clave de su simplicidad es que fast ethernet usa csma/cd definido en el media access control.

El 100 base-T es una versión escalada del (M.A.C.), usado en ethernet convencional, sólo que más rápido, es la misma tecnología robusta, confiable y económica usada por 40 millones de usuarios hasta hoy, lo que es más, la misma compatibilidad entre 10 base-T y 100 base-T permite la fácil migración a conexiones de alta velocidad sin cambiar el cableado, depurando técnicas de administración de red y más.

Adicionalmente, ambas tecnologías ofrecen ambientes compartidos con conexiones ethernet compartidas o conmutadas permitiendo 10 0 100 mbps a todas las estaciones conectadas al hub, esto es ideal para grupos de trabajo de tamaño mediano con incrementos de demanda de ancho de banda ocasionales, ethernet compartido delibera el ancho de banda a un costo muy bajo.

Ambientes conmutados proveen el máximo ancho de banda para cada puerto conmutado del hub. Para grupos de trabajo grandes con demanda agregada que excede los 100 mbps, ethernet conmutado es la mejor solución.

↳ Costo efectivo de migración.

Como el protocolo natural de 10 base-T, virtualmente no cambia en fast ethernet, éste puede ser introducido fácilmente en ambientes de ethernet estandar. la migración es simple y económica en muchos aspectos importantes.

- ◊ Las especificaciones de el cableado para red 100 base-T permiten a fast ethernet correr en la mayoría de cableados comunes en ethernet, incluso categorías 3,4 y 5 de utp, stp y fibra óptica.

- ◊ Experiencia administrativa. los administradores pueden relevar en ambientes 100 base-T con herramientas de análisis de red familiares.

- ◊ La administración informática se traduce fácilmente de ethernet a 10MBPS a redes fast ethernet lo que significa recapacitación mínima del personal de administración y mantenimiento de la red.

Software de administración. Las redes fast ethernet pueden ser administradas con un protocolo simple como smnp.

Soporte de software. El software de aplicación y manejo de redes no cambia en redes 100 base-T.

Migración flexible. Adaptadores autosensibles de velocidad dual pueden correr a 10 ó 100 mbps en el medio existente, al igual que los concentradores con 10 100 mbps permiten el cambio dependiendo de la transmisión que se esté realizando

↳ Soporte de los principales fabricantes.

Fast ethernet es soportado por más de 60 fabricantes importantes, incluyendo empresas líder en adaptadores, conmutadores, estaciones de trabajo y empresas de semiconductores como 3Com, SMC, Intel, Sun Microsystems y Synoptics que empezaron a comercializar productos interoperables a fines de 1994.

Estas empresas son miembros de la Fast Ethernet Alliance (FEA), un consorcio cuyo objetivo es acelerar la tecnología fast ethernet a través de la Norma 802.3 del IEEE. Además la FEA estableció procedimientos de prueba y estándares para asegurar la interoperabilidad para los fabricantes de productos 100 Base-T.

↳ Valor óptimo.

Como la estandarización progresa rápidamente y los productos estarán disponibles por una gran variedad de fabricantes, el precio/desempeño de fast ethernet estará regido por la competitividad de las tecnologías de alta velocidad.

Al principio, los precios de fast ethernet superaban 10 veces el desempeño por menos de la mitad del costo por conexión. Ahora los precios están casi a la par de la tecnología de 10 Base-T y aún tienen las ventajas sobre otras tecnologías no ethernet.

↳ La tecnología tras fast ethernet.

Fast ethernet es una extensión del estandar existente 802.3 del IEEE, la nueva tecnología usa el mismo control (Media Access Control), de 802.3 conectado a través de otro control (Media Independient-Interface), a otros tres controles de nivel físico, la especificación de M.I.I., es similar a la AUI de 10 mbps y proporciona una sola interface que puede soportar transceivers externos con alguna de las especificaciones 100 Base-T.

100 base-T soporta tres especificaciones: 100 baseTx, 100 base T4 y 100 base Fx, el estándar 100 base-T, también define una interface para concentrador universal y una interface de manejo.

En el diseño del MAC para 100 base-T, el IEEE reduce el tiempo de transmisión de cada bit, del MAC de 10 mbps de csma/cd multiplicado por un factor de 10 proporcionando turbo velocidad al paquete. Desde que el MAC está especificado de manera independiente de la velocidad, la funcionalidad en el formato del paquete no cambia, la longitud, el control de errores y la información de manejo son idénticos a 10 Base-T.

Alternativas de cableado.

- ◊ 100 base-T soporta 3 especificaciones físicas.
- ◊ 100 Base Tx: Cable UTP o STP de un par trenzado eia 568 o categoría 5 para datos.
- ◊ 100 Base T4: Cable UTP de 4 pares trenzados para voz y datos categoría 3, 4 ó 5.
- ◊ 100 Base Fx: sistema estándar de 2 fibras ópticas.

La flexibilidad de estas especificaciones permite a 100 base-T, implementar un ambiente de cable 10 Base-T virtual, permitiendo a los usuarios conservar la infraestructura de cableado mientras emigran a fast ethernet.

Las especificaciones 100 base Tx y 100 Base T4, juntas cubren todas las especificaciones de cableado que existen para redes 10 Base-T, las especificaciones fast ethernet pueden ser mezcladas e interconectadas a un hub como lo hacen las especificaciones 10 Base-T.

100 Base Tx está basado en la especificación PMD (Physical Media Dependent), desarrollada por el ansi x3t9.5, éste combina el MAC escalado con los mismos chips del transceiver y el PHY desarrollados para FDDI y CDDI. Como estos chips están disponibles y el estándar de señalización está completo, 100 Base-T ofrece una solución de tecnología aprobada y basada en estándares y soporta ambientes de cableado 10 Base-T.

100 Base-T permite transmisión a través de cable UTP 5 instalado virtualmente en las redes nuevas.

100 Base T4 es una tecnología de señal desarrollada por 3Com y otros miembros de Fast Ethernet Alliance para manejar las necesidades de cableado UTP 3 instalado en la mayoría de las antigüas redes basadas en 10 Base-T, esta tecnología permite a 100 Base-T correr sobre cableados UTP 3, 4 ó 5 permitiendo a las redes con cableado UTP 5 moverse a la tecnología de 100 Base-T sin tener que recablear.

100 Base FX es una especificación para fibra, ideal para grandes distancias o BackBones o ambientes sujetos a interferencia eléctrica.

✎ Auto-Negociación 10 / 100 MBPS

Para facilitar la migración de 10 a 100 MBPS el estándar 100 Base-T, incluye un sensor automático de velocidad, esta función opcional permite transmitir a 10 o 100 MBPS con comunicación automática disponible en ambos casos.

Auto-Negociación es usado en adaptadores 10 / 100 MBPS este proceso se da fuera de banda sin interposición de señal, para comenzar, una estación 100 Base-T advierte sus capacidades enviando un barrido de pulsos de prueba para verificar la integridad del enlace llamados FAST LINK PULSE, generados automáticamente al encender el equipo.

Si la estación receptora es un hub con capacidad 10 Base-T únicamente, el segmento operará a 10 MBPS, pero si el hub soporta 100 Base-T, este será censado por el FLP y usara el algoritmo de auto-negociación para determinar la mayor velocidad posible en el segmento, y enviar FLP´s al adaptador para poner ambos dispositivos en modo 100 Base-T.

El cambio ocurre automáticamente sin intervención manual o de software, (una RED o un segmento de RED puede ser forzado a operar a 10 MBPS a través de un manejo de mayor jerarquía, aunque éste sea capaz de trabajar a 100 MBPS, si asi se desea.)

✎ REGLAS DE TOPOLOGIA.

Fast Ethernet preserva la longitud crítica de 100 metros para cable UTP, como resultado del MAC escalado de la interface Ethernet.

Otras reglas topológicas de 100 MBPS son diferentes de las reglas Ethernet.

La figura 3 ilustra la clave de las reglas topológicas 10 Base-T y muestra ejemplos de como éstas permiten la interconexión en gran escala.

La máxima distancia en cable UTP es 100 metros igual que en 10 Base-T.

- ◊ En UTP se permiten máximo 2 concentradores y una distancia total de 205 mts.
- ◊ En topologías con un solo repetidor un segmento de fibra óptica de hasta 225 metros, puede conectarse a un backbone colapsado.
- ◊ Conexiones MAC to MAC, Switch to Switch, o End Station to Switch, se usan segmentos de hasta 450 mts., de fibra óptica bajo 100 Base FX.

◊ Para distancias muy largas una versión completamente duplex de100 Base FX puede ser usada para conectar dos dispositivos a más de 2 KM de distancia.

Al principio, estas reglas topológicas pudieron parecer restrictivas , pero ahora en las redes con backbone, que usan fibra óptica , concentradores y/o ruteadores o puentes, Fast Ethernet puede ser fácilmente implementado en redes de gran escala o corporativas.

↳ ETAPAS DE MIGRACION.

La migración hacia fast ethernet está determinada en etapas, permitiendo al Administrador de la RED emigrar fast ethernet cuando y donde lo necesite.

Aqui tenemos una secuencia típica.

◊ Determine el tipo de cableado instalado, si este es categoría 5, se usan adaptadores100 Base TX, las categorías 3 ó 4 requieren adaptadores 100 Base-T4.

◊ Instale adaptadores de velocidad dual 10 /100 MBPS en PC´s nuevas; para prepararse a la migración de la nueva tecnología, las PC´s deben estar configuradas con adaptadores de velocidad dual, entonces podrán soportar ethernet compartido, ethernet conmutado, fast ethernet y aún fast ethernet conmutado.

◊ Instale concentradores 100 Base-T conforme el número de PC¨s se incremente, o conforme el tráfico de la RED empiece a crecer, comience la migración con hubs de velocidad dual, use un puente 10 / 100 MBPS para nodos que trabajen aún con 10 Base-T.

◊ Instale hubs conmutados 10 / 100 MBPS para las PC´s que ya existen en la RED, para usarse con las PC´s que no requieren tanta velocidad de comunicación, que además, necesitan conectarse a backbones o servidores a alta velocidad, el único cambio requerido en las conexiones ethernet 10 Base-T compartido a los puertos conmutados 10 /100 MBPS.

◊ Extienda 100 Base-T a los backbones. Conecte los grupos de trabajo y servidores a un backbone de alta velocidad, un puente o un ruteador con capacidad fast ethernet.

CAPITULO 2

DIPLOMADO EN REDES DE COMPUTADORAS (LAN, WAN Y GAN)

REDES DE ALTO DESEMPEÑO: FAST y GIGA ETHERNET FDDI-II, SWITCHING, ATM y FRAMERELAY

FDDI, FDDI - II

FDDI

Fiber Distributed Data Interface

Red anillo Token-Passing 100 Mb/s con redundancia. (ANSI-X3T9)

Anillo principal = Conexión Punto a Punto entre nodos para transmisión de datos

Anillo Secundario = Transmisión de datos/respaldo del anillo principal en caso de falla

FDDI proveé comunicaciones par conmutación de paquetes y transmisión de datos en tiempo real.

Notas:

FDDI

* FDDI emplea una codificación 4B/. tasa de transmisión a 100 Mb/-125 Mhz 80% de eficiencia en el ancho de banda
* ETHERNET Y TOKEN-RING emplea una codificación Manchester
* Tasa de transmisión - ETHERNET: 10Mb/s-20 Mhz
  - TOKEN-RING 16Mb/s-32 Mhz

50% DE EFICIENCIA EN EL ANCHO DE BANDA

Notas:

FDDI

FDDI: VS TOKEN - RING 16 MB/S

* Reloj distribuido recuperación de errores	{ Monitor Activo
* Doble anillo	{ Anillo Sencillo
* Rotación del " TOKEN"	{ Sistema de reservación por prioridad
* Uso de Fibra Optica	{ Uso de Par Trenzado/Fibra Optica

Notas:

FDDI

TOKEN-PASSING ofrece una transmisión de datos más eficiente, ya que conforme aumenta el tráfico se requiere un mayor ancho de banda. TRT 85 %.

CSMA/CD Resulta más eficiente cuando se utiliza un menor ancho de banda.

Notas:

FDDI

* FDDI Ofrece hasta 1000 conexiones físicas (500 Estaciones) y una distancia total de 200 Km. de extremo a extremo.

* La distancia máxima entre nodos activos es la de 2 Km

* Fibras Opticas empleadas:

  A) Fibra tipo unimodo, con gran ancho de banda (GHz) y largas distancias (20-30 Km)

  B) Fibra tipo multimodo. Fibras con nucleo 50-62.5 Micras y Medianas distancias (10-20 Km.) a 1300 nanometros.

Notas:

FDDI

CONSIDERACIONES

Manejo

SMT (Interface SNMP)
Estadística de las estaciones reset. Soporte para deshabilitar.

300KM-180 Miles

El control es crítico para las Redes de gran tamaño y capacidad.

Notas:

FDDI: BACKBONES TOPOLOGIAS

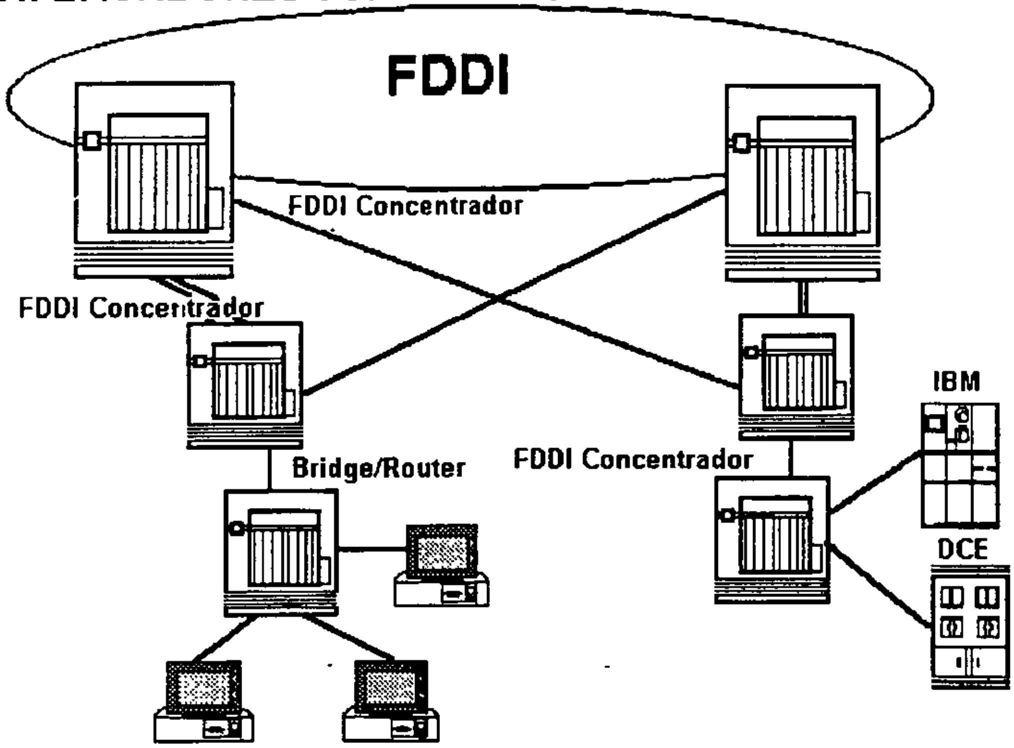

Notas:

FDDI

ESTACIONES

Tipo Clase A: Se conecta directamente al anillo doble

Tipo Clase B: Se conectan al concentrador puertos mültiples en Red estrella o Estaciones ɔn posibilidad de conexión sencilla. Los concentradores pueden ser conectados en cascada.

Notas:

USANDO HOSTS CON FDDI

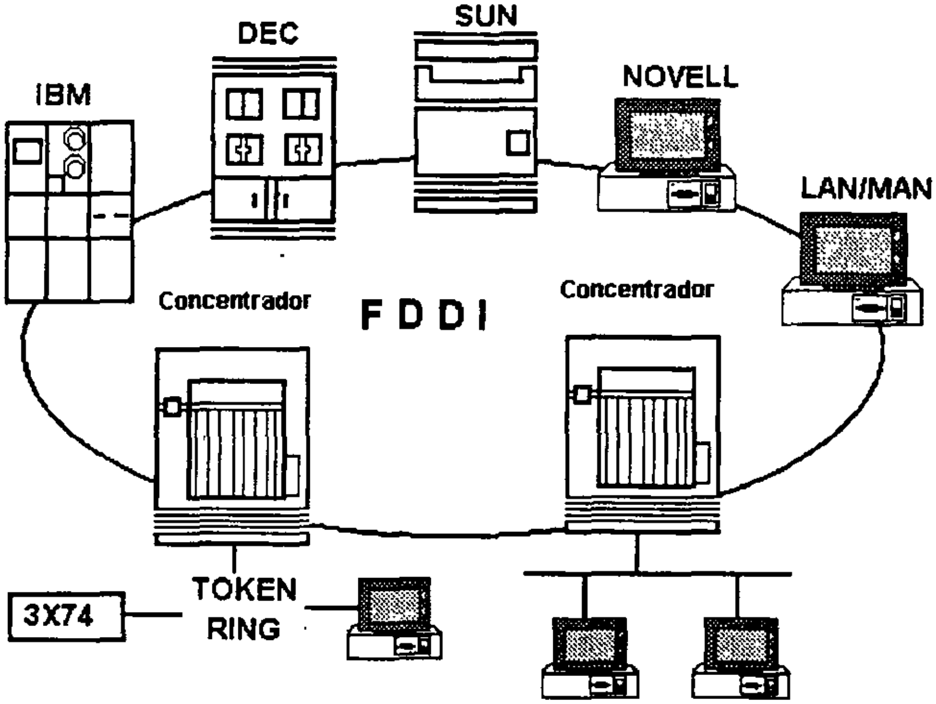

Notas:

FIBRAS OPTICAS

Hasta hace cerca de una década, las comunicaciones fueron realizadas a través de medios como cable coaxial o cable telefónico, Desde hace algunos años y ahora más fuerte que nunca se introduce un nuevo medio de comunicación: **las fibras ópticas.**

El uso de la luz como un medio de comunicación no es nuevo. El fuego fué usado como señal de comunicación en los amaneceres de la historia humana. La clave Morse fue utilizada particularmente en comunicaciones de una embarcación a otra usando espejos para reflejar la luz y transmitir señales.

En 1860 Alejandro Graham Bell demostró la transmisión de voz usando espejos.

Estos vibraban debido a las ondas sonoras generadas por la voz, de manera que la luz reflejada por los espejos era modulada por el sonido. La luz modulada en el receptor era enfocada en una lámina de Selenio, la resistencia de la lámina y su respectiva corriente variaba con los cambios de intensidad de la luz incidente. Esta corriente se aplicaba a un dispositivo parecido a un altavoz moderno
Todos estos métodos dependían del medio ambiente y solo cubrían distancias pequeñas y para aplicaciones visuales en línea directa, en 1960 con la invención del láser, el interés por la comunicación luminosa tomo fuerza, aunque, contando con el láser, los métodos de comunicación por luz al aire libre seguían dependiendo del ambiente y limitados en alcance.

El primer intento para transmitir a larga distancia a través de fibra de vidrio fue realizado en 1966, pero las excesivas impurezas de la fibra de vidrio generaban grandes pérdidas de energía de la luz que viajaba a través de ésta. La transmisión seguía limitada en distancia, además de que el tamaño de los lasers con que se contaba en aquel tiempo hacían muy difícil el acoplamiento de la energía luminosa en las fibras de manera eficiente.

Con el desarrollo del diodo láser, del diodo LED, y más tarde la introducción de alta pureza, llegó la era de la comunicación por fibra: transmisión a largas distancias sin la necesidad de reamplificar la señal.

La historia del desarrollo de la tecnología de fibra óptica se centra en aplicaciones de comunicación y desarrollo e investigación gubernamental, los avances mas significativos se lograron recientemente en la década de los 70's y los 80's, aunque la teoría general de la propagación de la luz se desarrolló a lo largo de muchos años de investigaciones intentos y fracasos.

Una fibra óptica es una delgada varilla transparente hecha de vidrio o plástico puro, a través del cual la luz puede propagarse con una pérdida de señal muy baja, la estructura de una fibra óptica moderna consiste en el tubo de vidrio delgado recubierto por otro material con distintas características ópticas, éste evita que la señal que viaja a través de la fibra óptica se refracte fuera de la misma ocasionando pérdidas en la señal.

El uso de fibra óptica para transmitir señales de comunicación tiene muchas ventajas importantes sobre los medios de comunicación convencionales:

- ◊ La baja pérdida en la energía de la señal.
- ◊ La baja tasa de distorsión en los pulsos de la señal transmitida.
- ◊ El ancho de banda es mucho mayor que en UTP o coaxial.
- ◊ No es susceptible de ruido o interferencia eléctrica o electromagnética.
- ◊ Es muy segura, no es posible "robarse" la señal de la fibra óptica.
- ◊ Soporta ambientes hostiles, contaminación, salinidad, humedad o radiación. Es inmune.
- ◊ No existe una conexión eléctrica entre receptor y transmisor.
- ◊ El costo de la fibra óptica es casi el mismo que el del cable coaxial.
- ◊ Las velocidades de transmisión son muy altas.

Recientes desarrollos han permitido fibras ópticas con 0.2 dB de atenuación por kilómetro, además de los desarrollos de equipos para trabajar con fibra óptica con capacidad de operación de hasta 1 Ghz y mas de 3000 canales de comunicación individuales.

Las fibras ópticas se clasifican en dos tipos: unimodo y multimodo.
Llamadas así por el número de modos de propagación de la longitud de onda de operación

Fibra multimodo

Es un tipo de fibra en la cual hay más de un modo de propagación de señal. Van desde las que tienen dos modos hasta cientos de modos de propagación. Las aplicaciones típicas de estas fibras son la telecomunicación con anchos de banda de 1 a 2 Ghz, cableado de inmuebles, con anchos de banda de 500 a 1000 Mhz y enlaces donde la potencia y el ancho de banda son necesarios, generalmente 50 a 100 Mhz son suficientes.

Fibra unimodo

La fibra unimodo es fabricada con los mismos materiales y bajo los mismos procesos que las fibras multimodo, la diferencia es el tamaño del centro de la fibra que es mas pequeño y la cantidad de impurezas que es diferente a la fibra multimodo, hace la diferencia de características de operación.

Las siguientes tablas ofrecen un panorama general de características

✎ Dimensiones

Fibra óptica Tipo	diámetro del núcleo (micras)	diámetro del revestimiento (micras)	longitud de onda (Nanometros)
unimodo	8.10	125	1300,1500
multimodo	50	125	850,1300

✎ cuadro comparativo de atenuación.

Medio de comunicación Tipo	Longitud de onda o Frecuencia	Atenuación (dB / Km.)
COAXIAL	100 Mhz	61
Fibra Óptica Multimodo	850 Nm	2.4 - 3.2
Fibra Óptica Multimodo	300 Nm	1.0 - 1.5
Fibra Óptica Unimodo	1300 Nm	menor a 0.5
Fibra Óptica Unimodo	1300 Nm	menor a 0.25

✎ Distancias máximas cubiertas por un segmento de línea de comunicación

Medio de comunicación Tipo	Distancia máxima sin repetidor (Mts) (Rango dinámico típico 35 dB)
COAXIAL	570
Fibra óptica Multimodo a 850 Nm	10, 000
Fibra óptica Multimodo a 1300 Nm	20, 000
Fibra óptica Unimodo a 1300 Nm	60, 000
Fibra óptica Unimodo a 1550 Nm	120, 000

ASPECTO DE LA FIBRA OPTICA

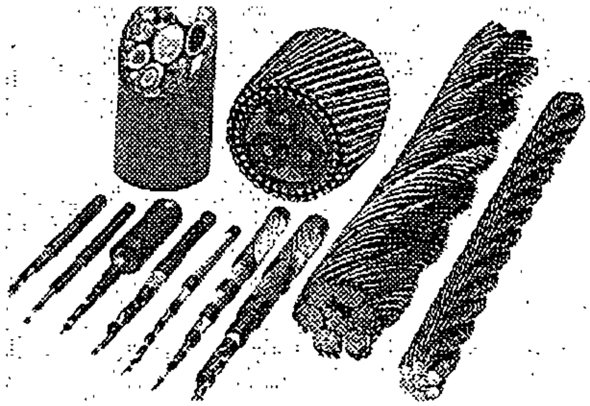

existe una gran variedad de presentaciones para fibras ópticas dependiendo de las aplicaciones.

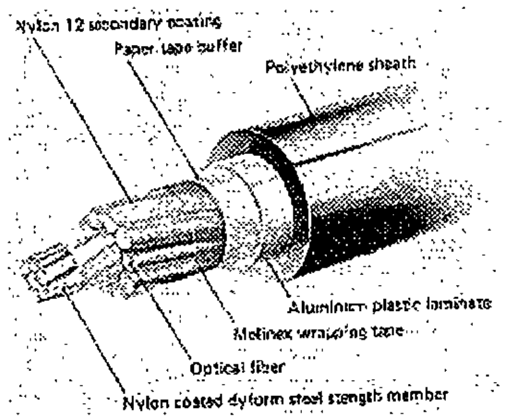

CABLE DE FIBRA OPTICA PARA ESTRUCTURA

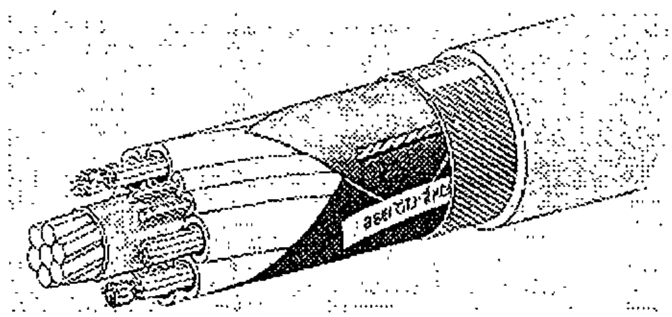

TUBO DE FIBRA OPTICA DE USO INDUSTRIAL

↳ CONECTORES DE FIBRA OPTICA.

Son dispositivos de unión, que realizan la función de acoplamiento entre dos fibras ópticas o en los extremos de éstas, permitiendo un fácil manejo, instalación y mantenimiento de la fibra óptica.

Los parámetros que definen la calidad de un conector para un sistema de transmisión dado son los siguientes:

- ◊ Pérdida por inserción.
- ◊ Facilidad para su ensamble y montaje.
- ◊ Estabilidad al ambiente.
- ◊ Confiabilidad.
- ◊ Inserción de perturbaciones al sistema.
- ◊ Costo.

Aunque normalmente es imposible optimizar todos los parámetros, la elección de un conector es el resultado de un balance de necesidades específicas, debe tenerse el cuidado no solo de seleccionar el conector adecuado, sino que también debe ponerse especial atención en el momento del manejo y ensamble de los conectores.

FDDI

La nuevas tecnologías de interconexión de redes tienden al uso de la fibra óptica, como medio de comunicación, tiene una capacidad de transmisión de datos y de seguridad muy altas. Las fibras ópticas pueden soportar transmisiones de varios cientos de Mbps. Los cableados por medio de fibra óptica pueden soportar grandes distancias sin necesidad de repetidores, además de ser un medio inmune a la interferencia electromagnética.

Los costos de conexión con fibra óptica son típicamente altos, pero podemos esperar que estos precios bajen significativamente en los próximos años.

Ya existen en el mercado, proveedores que cuentan con las tarjetas necesarias para poder realizar conexiones con fibra óptica para las topologías Ethernet y Token Ring.

Muchas compañías están optando por la fibra óptica por diversas razones, entre ellas está la velocidad de transmisión de la que es capaz. Por ejemplo, **FDDI**[1] soporta velocidades de transmisión de hasta 100 Mbits por segundo. En comparación con Ethernet que transmite a 10 Mbits por segundo o Token Ring que transmite a 4 ó 16 Mbits por segundo.

El comité 802.6 de la IEEE ha adoptado estándares para redes de área metropolitana, y el American National Standars Institute ha desarrollado los estándares FDDI y FDDI-II .

Además, la fibra óptica tiende a ser más segura que el cableado de cobre. Una red interconectada por medio de fibra óptica puede trabajar cerca de equipo eléctrico altamente sensible sin interferir uno con el otro. Un cable de fibra óptica entre dos edificios no atraerá rayos como el cable de cobre.

[1] Fiber Distributed Data Interface

Al hablar de redes interconectadas por medio de fibra óptica, generalmente se está hablando de FDDI, diversos productos capaces de soportar FDDI han estado saliendo lentamente al mercado y se han dejado ver en diversas exposiciones de computadoras.

Como Token Ring, FDDI usa una topología con forma de anillo y un Token eléctrico para pasar el control de la red de una estación a otra, más no es compatible con Token Ring.

La mayor parte de las redes actuales con FDDI usan un doble anillo en donde cada nodo se une a los dos anillos independientes, transmitiendo los datos en sentidos opuestos. Esta configuración mejora la velocidad de transmisión así como la confiabilidad de la red, pero es muy caro.

Hasta ahora, FDDI se ha usado para interconectar PC's de alta velocidad o estaciones de trabajo con redes, o bien como *backbone* para interconectar estaciones más lentas, de igual manera que una carretera une los diferentes pueblos. Conectarse a FDDI es caro, dado el alto costo de los componentes ópticos, así como el costo del transreceptor y los integrados necesarios para FDDI.
Debido a sus características de ancho de banda, la fibra óptica se usa principalmente para backbones (que es un segmento que une varias redes locales) .

Existe también FDDI-II que es una segunda versión de FDDI que nos permite transmitir voz y video además de datos. De manera distinta a FDDI que tiene un reloj corriendo de manera independiente, FDDI-II tendrá un marco de 125 microsegundos, permitiendo ser sincronizado con la red de comunicaciones.

FAQ´S sobre FDDI y FDDI-II

Q. What does FDDI stand for?

Fiber Distributed Data Interface

Q. What is the difference between FDDI and FDDI-II?

Both FDDI and FDDI-II runs at 100 M bits/sec on the fiber.
FDDI can transport both async and sync types of frames.
FDDI-II has a new mode of operation called Hybrid Mode.
Hybrid mode uses a 125usec cycle structure to transport isochronus traffic, in addition to sync/async frames.
FDDI and FDDI-II stations can be operated in the same ring only in Basic mode.

Q. What is the name of the standards and where can I get them?

ANSI X3T9.5 standards
American National Standards Institute
1430 Broadway, New York, NY 10018, USA
Attention: Sales Dept.

- IEEE Standards
IEEE Service Center
445 Hoes Lane, Piscataway, NJ 08855, USA
- X3T9.5 Documents
Global Engineering Documents
(USA) 1-800-854-7179

Q. What are other good sources of printed information?

- FDDI Technology and Applications: Edited Mirchandanı and Khanna
- Handbook of Computer Communications Standards Vol 2: By Stallings
- Call up DEC to ask for the free FDDI tutorial book
- Dig up 1986-1992 issue of IEE Local Computer Network Conference

Q. I've heard that FDDI uses a token passing scheme for access arbitration, how does this work?

A token is a normal FDDI frame with a fixed format.
The station waits until a token comes by, grabs the token, transmits the the frames and release the token. The amount of frames that can be transmitted is determined by timers in the MAC protocol chips.

[You really need a diagram for the station and/or topology.]

Q. I've heard that FDDI is a counter-rotating ring, what does this mean?

FDDI is a dual ring technology. And each ring is running in the opposite direction to improve fault recovery.

Q. What is a dual ring of trees?

See the diagram.

Q. What is dual homing?

When a DAS is connected to two concentrator ports, it is called dual-homing. One port is the active link, where data is transmitted and the other port is a hot standby. The hot standby will constantly testing the link and will kick in if the active link failed or disconnected. The B-port in a DAS is the active port and the A-port is the hot-standby.

Q. What is a DAS?

DAS (Dual Attach Station) is a station with two peer ports (A-Port and B-Port). The A-port is going to the B-Port of another DAS, and the B-port is going to connect to the A-Port the yet another DAS. ie:

```
      ---------        ---------       ---------
 +--->|A     B|------>|A     B|----->|A     B|----+
 | +--|---------|<------|---------|<-----|---------|<-+ |
 | |                                      | |
 | +------------------------------------------+ |
 +----------------------------------------------+
```

Q. What is a SAS?

SAS (Single Attach Station) is a station with one peer port (S-Port). It is usually connected to the M-Port of a concentrator.

Q. What is a wrapped ring?

When a link in the dual-ring is broken or not connected, the two adj ports connecting to the borken link will be disconnected and the both stations enter the wrap state.

```
          Wrap   Wrap
      ---------        ---------       ---------
 +--->|A     B|-X  X->|A     B|----->|A     B|----+
 | +--|---------|<------|---------|<-----|---------|<-+ |
 | |                                      | |
 | +------------------------------------------+ |
 +----------------------------------------------+
```

Q. Do I need a concentrator port for each workstation, or can workstations be chained together?

Usually you will need a concentrator port (M-Port) to connect each SAS. DAS can be hooked up to the main rings or concentrator port(s).

Q. If I use a concentrator, what are the advantages/disadvantages?

Advantages: Fault tolerance. When a link breaks, the ring can be segmented. A concentrator can just bypass the problem port and avoid most segmentations. It also gives you better physical planning. Usually people prefer tree physical topology. Generally star configuration of a concentrator system is easier to troubleshoot.

Disadvatages: A concentrator represents a single point of failure.
There may also be more costly.

you can build a tree as deep as you want. We have

a Q. Can I cascade concentrators? Are there limitations as to how many?
Yes. And dual-rings of concentrator here connecting machine rooms and wiring rooms. And from the there we connect to other concentrators to different offices. Then we have a concentrator in the lab to different machines. There is a maximum of 500 stations on an FDDI LAN.

Q. What is a bypass and what are the issues in having or not having one?
Bypass is a ($600-$1200) device that is used to skip a station on the ring if it is turned off. Therefore, you don't need to use concentrator to avoid the segmentation problems. One problem with them is that they increase the db loss of the fiber, so you can't have too many of them (3 activated in a row maximum, I believe).

Q. What are the minimum/maximum distances on fiber runs?
no min, 2 km max for multimode fiber. 20 km max (may be as high as 60km, we're not sure) for single mode fiber.
500 m for the new Low Cost Fiber.

Q. What are the types of fiber that are supported?
Multimode (62.5/125 micron graded index multimode fiber)
and other fiber like 50/125. 85/125. 100/140 allowed
Single mode (8-10 micron)
The new Low Cost Fiber (plastics?) standard.

Q. I've hear of FDDI over Copper, what type of cable does this scheme use?
Type 1 STP - distance between connections must be less than 100 m
Category 5 UTP - distance between connections must be less than 100 m
(The ANSI standard for STP and UTP is incomplete, but a number of companies are already shipping proprietary twisted-pair solutions until the standard is completed, which is expected later this year.)

? Q. Is there any advantage to seperating the fiber pairs (will the ring work better if only one strand is broken on a DAS connection?)

Q. I have ethernet, can I bridge/route between the 2 topologies?
Yes. But from what we are hearing some protocols are having problems.
Only TCP/IP is handling frame fragmentation correctly. (See below).
It should also be noted that frame fragmentation will not work for DECNET, IPX, LAT, Appletalk, NETBEUI etc.
IP is the only protocol that has a standard method of fragmenting.
Other protocols destined for Ethernet Lans must stay below the 1500 MTU.

Q. I've heard that there is a frame length difference, what are the issues and problems here?

FDDI frames has a max size of 4500 bytes and Enet only 1500 bytes. Therefore your bridge or router needs to be smart enough to fragment the packets (eg into smaller IP fragments). Or you need to reduce your frame size to 1500 bytes (of data).

Q. What does an FDDI frame look like?

PA	Preamble (II) (8 or more Idle symbol pairs)
SD	Starting Delimiter (JK) (J followed by K control symbol)
FC	Frame Control (nn) (Tell you if it is a token, MAC frame, LLC frame, SMT frame, frame priority, sync or async)
DA	Destination Address (nn) (6 bytes of MAC Address in MSb first format)
SA	Source Addrewss (nn) (6 bytes of MAC Address of this station)
INFO	Information field (nn) (Varibale Length. Usually starts with LLC header, then SNAP field, then the payload eg IP packet)
ED	Ending Delimiter (T) (one T control symbol)
FS	Frame Status (EAC) (Three symbols of status of Error, Address_match, and Copied. Each symbol is either SET or RESET. eg If EAC == RSS, then then frame has no error, some station on the ring matched the DA, and some station on the ring copied the frame into its buffer.

Q. So FDDI is 100 Megbits per second, what is the practical maximum bps?

Depends. You can get aggregate usage up to 95Mbit/s with no problem. But 75Mbps is pretty good. Actually, this question depends so much on how you construct your test, what equipment you use, etc, that the best idea is to let the user decide.

Q. What happens when I bridge between a 100 Mbps FDDI and a 10Mbps ethernet if the FDDI traffic destined for the ethernet gets above 8 Mbps? 10 Mbps?

After the buffer fills Frames start dropping. This is not a problem unique to FDDI however. Consider ethernet to T1, or multiple ethernets to a single ethernet.

Q. What is the latency across a bridge/router?. (Yes I know that different vendors are different, but what is a the window?)

No idea.

Q. Are there FDDI repeaters?

Yes. But it is not a standard yet. A group in the ANSI committee is looking into making FDDI repeater a standard. Other companies like ODS has something like simgle mode to multimode converter.

Q. What type of test and trouble shooting equipment is available for FDDI?

Digital Technology Inc (DTI), W&D, HP, and Tekelec all sell FDDI analyzers. The Sniffer from Network General also has a module that works with the NPI FDDI Cards. SGI has a nice looking ringmap program. IBM has a product called DatagLANce. Most Ethernet tools will also work with FDDI in the protocol level. Also a optical time domain reflectometer (TDR) is recommended for db loss checking and distance measurements, though it has been reported that an FDDI link tester is less expensive and will do the job.

Q. What about network station management? Does FDDI support SNMP?

Yes. There is a FDDI-SNMP MIB translation from the SNMP working group.

Q. What is a beaconing ring? Does FDDI beacon?

Beacon is a special frame that FDDI MAC sends when something is very wrong. When Beaconing for a while, SMT will kick in trying to detect and solve the problem.

Q. How about interoperability, does one manufacture's equipment work with others?

Just like any networking products, Ethernet, Token, FDDI, ATM, there is a possibility that one vendor does not work with another. But most of the equipment shipping today is tested at InterOp, UNH or ANTC, are this is the equipment that will meet the minimum interoperability requirements. Ask the vendor what type of testing they did and ask them to ship you a system for field trial before you pay big bucks for it.

Q. Can I interface FDDI to a PC (ISA Bus), PC (EISA Bus), PC (Micro channel Bus), Macintosh, Sun workstation, DECstation 5000, NEXT computer, Silicon Graphics, Cisco router, WellFleet router, SNA gateway (McData), other?

Yes. I am not sure if NeXT has any FDDI adaptor software, but there are ~5 different NuBus FDDI cards in the market. But FDDI adaptors are available for all other buses or vendors.

Q. What is the maximum time a station has to wait for media access. What type of applications care?

MaxTime = ~(#of stations * T_neg)
(T_neg ist the negotiated target token rotation time)
Usually this won't happened. It is only a very very heavily loaded ring but the station be waiting for that long. If this is the case, then change the T_request of the station to some lower value (eg 8 msec).

Q. Can I bridge/route TCPIP, SNA, Novell, Sun protocols, DecNet, Banyan Vines, Appletalk, X windows, LAT?

Yes for IP, Novell, DecNet, X windows.
Don't know about the others.

Q. What are the applications that would use FDDI's bandwidth?

Basically anything will be at least a bit faster. From NFS to images transmission. Even if a single station cannot take advantage of the 100M bit/sec, the aggregate bandwidth will help a lot if your Ethernet is saturated. However, note that though FDDI has higher bandwidth than ethernet, the signals travel at the same speed. The propogation of a signal on the transmission line is the same for ethernet, token ring, and FDDI.

Q.What are the effects of powering off a workstation on a DAS or SAS connection?

Depends. Let's do SAS first, it is easier. If a SAS is connected to a concentrator, then the concentrator will bypass the SAS connection using an internal data path. If the DAS is connected to a concentractor, then the concentrator will also bypass the DAS. If the DAS is connected to the trunk rings without using an optical bypass switch, then the trunk ring will wrap. If multiple stations power off on the trunk rings, then the rign will be badly segmented. Now if the DAS is using an optical bypass switch, the switch will kick in and prevent the ring from wrapping.

Q. What are the effects of disconnecting the fiber on a DAS or SAS connection?

SAS connecting to concentrator:

Same as above.

DAS dual-home to a concentrator:

If A-port fiber breaks, no effect on B port since A port is a backup port. (And SMT will NOT send out alert msg.)
If B-port fiber breaks, A-port will kick in, complete PCM and be used as the primary connection.

DAS on trunk rings, with no optical bypass:

If one fiber breaks, then the ring will wrap.
If both fibers break, ring will wrap, station won't be communicate.

DAS on trunk rings using optical bypass:

If one fiber between bypass and the next station breaks, then the ring will wrap.
If both fibers between bypass and the next station break, ring will wrap, station won't be able to communicate.
If one fiber between bypass and the host station breaks, then the ring will wrap.
If two fiber between bypass and the host station breaks, then the ring will wrap.

Q. What is one recommended topolpgy?

Connect backbone concentrators and ring monitors to the trunk rings, and connect all the workgroup concentrators and users stations to the backbone concentrators. Connect bridges and routers to backbone concentrators using dual-homing.

Q. What is Graceful Insertion? Should I demand it from my vendors?

Graceful Insertion is a method to insert a station (or a tree) in a concentrator without losing any data frames (and not going into Ring_Non_Op mode). The theory goes as Graceful Insertion can minimize ring non_op and losing frame, therefore it saves you transmission timeout of lost frame in upper layer protocol (eg TCP) and retransmission effort. The following is the counter argument: Graceful Insertion can hold up the ring for more time that the FDDI ring non-op recovery time. And Upper layer protocol is designed to perform frame recovery and retransmission anyway. And no vendor can gaurantee 100% Graceful Insertion anyway. Should I get Graceful Insertion in my concentrators? If it is free, take it. You are going to get ring_op no matter what (eg insertion in the trunk ring and station power down).

Q. Is there a Graceful De-insertion?

No.

Q. Can you name a few FDDI Concentrator vendors?

IBM, Optical Data System, SynOptics, Cabletron, DEC, Chipcom, NPI, Synernetics, 3Com, Interphase, Ungermann-Bass, Timeplex, Crescendo/Cisco, Sumitomo etc ... (vendors feel free to email me to be included here)

Q. Can I run FDDI on electrical cable?

DEC is already sell a FDDI link that runs on coax.

ANSI is currently finishing up the TP-FDDI Standard for running FDDI on twisted-pair media (Category 5 Cable).

ANSI is also working on a standard (long term TP working group) to run FDDI on telephone cable. [Please comment.]

IBM and a group of vendors (SynOptics, National Semiconductor ...) promote SDDI that runs FDDI on Shielded Twisted-Pair cable.

(this is incomplete), there is much work being done on FDDI over various types of electrical cable, most notably twisted pair.

Q. What does SMT stand for? What does it do? Do I need it?

Station ManagemenT (SMT). It is part of the ANSI FDDI Standards that provides link-level management for FDDI. SMT is a low-level protocol that addresses the management of FDDI functions provided by the MAC, PHY, and PMD. It performs functions like ring recovery, frame level management, link control, etc. Every stations on FDDI need to have SMT. The latest version of the SMT standard is version 7.3, but most vendors ship products with SMT version 6.2.

Q. Who supports FDDI-II?

National Semiconductor Corp, IBM, Apple Computer, XDI, Alpha Inc, etc

Q. Who is working on Synchronous frame type utilitization?

Alpha, IBM, and many more companies. Try to contact scoop4@aol.com and warren@lgevm2.vnet.ibm.com. They are working with a group of companies to define the usage of SYNC frame in FDDI-I rings.

Q. Can I connect two Single attach stations together and form a two stations ring without a concentrator?

yes. You can do that if both stations support the S-S port connection. Most vendors support the S-S connections.

Q. What are ports? What are the different type of ports?

A port is the basically the fiber optic connector on the card. FDDI SMT defines 4 types of ports (A, B, M, S). A dual-attach station has two ports, one A-port and one B-port. A single attach station has only one port (S-port). A concentrator will have many M-port for connecting to other stations' A, B or S-ports.

Q. What are the port connection rules?

When connecting DASs, one should connect the A-port of one station to the B-port of another. S-port on the SAS is to connect to the M=port on the concentrators. A and B-port on DASs can also connect to the M-port of concentrator. But M-ports of the concentrator will not connect to each other.

In more detail, SMT suggested tthe following rules:

	A	B	M	S
A	-	+	+	-
B	+	-	+	-
M	+	+	X	+
S	-	-	+	-

==> '+' is the preferred connection

==> '-' connection has possible problems, and a vendor can choose to disable that connection in the default configuration

==> 'X' indicates a legal connection and will be rejected

CAPITULO 4

DIPLOMADO EN REDES DE COMPUTADORAS (LAN, WAN Y GAN)

REDES DE ALTO DESEMPEÑO: FAST y GIGA ETHERNET FDDI-II, SWITCHING, ATM y FRAMERELAY

ATM

ATM and Cell Relay Service

1.1 Introduction

1.1.1 Background

Asynchronous transfer mode (ATM), as the term is used in current parlance, refers to a high-bandwidth, low-delay switching and multiplexing technology that is now becoming available for both public and private networks. ATM principles and ATM-based platforms form the foundation for the delivery of a variety of high-speed digital communication services aimed at corporate users of high-speed data, LANs interconnection, imaging, and multimedia applications. Residential applications, such as video distribution, videotelephony, and other information-based services, are also planned. ATM is the technology of choice for evolving broadband integrated services digital network (B-ISDN) public networks, for next-generation LANs, and for high-speed seamless interconnection of LANs and WANs. ATM supports transmission speeds of 155 Mbits/s and 622 Mbits/s, and will be able to support speeds as high as 10 Gbits/s in the future. Networks operating at these speeds have been called gigabit networks. As an option, ATM will operate at the DS3 (45 Mbits/s) rate; some proponents are also looking at operating at the DS1 (1.544 Mbits/s) rate. While ATM in the strict sense is simply a Data Link Layer protocol, ATM and its many supporting standards, specifications, and agreements constitute a platform supporting the integrated delivery of a variety of switched high-speed digital services.

Cell relay service (CRS) is one of the key new services enabled by ATM. CRS can be utilized for enterprise networks that use completely private communication facilities, use completely public communication facilities, or use a hybrid arrangement. It can support a variety of evolving corporate applications, such as desk-to-desk videoconferencing of remote parties, access to remote multimedia video servers (for example, for network-based client/server video systems), multimedia conferencing, multimedia massaging, distance learning, business imaging (including CAD/CAM), animation, and cooperative work (for example, joint document editing). CRS is one of three "fastpacket" technologies, that have entered the scene in the 1990s [the other two are frame relay service and Switched Multimegabit Data Service (SMDS)]. A generic ATM platform supports all of these fastpacket services (namely, it can support cell relay service, frame relay service, and SMDS), as well as circuit emulation service.

1993 saw the culmination of nine years of ATM standards-making efforts. Work started in 1984 and experienced an acceleration in the late 1980s and early 1990s. With the ITU-TS (International Telecommunications Union Telecommunication Standardization) standards and the ATM Forum implementers' agreements, both of which were finalized in 1993, the technology is ready for introduction in the corporate environment. In particular, a user-network interface (UNI) specification that supports switched cell relay service as well as the critical point-to-multipoint connectivity, important for new applications, has been finalized (multiservice UNIs are also contemplated). In 1993, the ATM Forum also published a broadband intercarrier interface (B-ICI) specification; this specification is equally critical for wide-area network (WAN) inter-LATA service. At press time, a variety of vendors were readying end-user products for 1994 market introduction; some prototype products have been on the market since the early 1990s. A number of carriers either already provide services or are poised to do so in the immediate future.

A key aspect of B-ISDN in general and ATM in particular is the support of a wide range of data, video, and voice applications in the same public network. An important element of service integration is the provision of a range of services using a limited number of connection types and multipurpose user-network interfaces. ATM supports both nonswitched permanent virtual connections (PVCs) and switched virtual connections (SVCs). In a PVC service, virtual connections between endpoints in a customer's network are established at service subscription time through a provisioning process; these connections or paths can be changed via a subsequent provisioning process or via a customer network management (CNM) application. In SVC, the virtual connections are established as needed (that is, in real time) through a signaling capability. ATM supports services requiring both circuit-mode and packet-mode information transfer capabilities. ATM can be used to support both connection-oriented (e.g., frame relay service) and connectionless services (e.g., SMDS).

1.1.2 Course of Investigation: applying ATM to enterprise networks

This book is aimed at corporate practitioners who may be interested in determining how they can deploy ATM and cell relay technology in their networks at an early time and reap the benefits. The purpose of this first chapter is to provide an overview of key ATM/cell relay service concepts. These concepts will be revisited in more depth in the chapters that follow.

The book has four major segments: (1) platform technology applicable to all B-ISDN services, (2) cell relay service, (3) interworking and support of basic multimedia, and (4) use ofATM in corporate enterprise networks. Table 1.1 provides a roadmap of this investigation.

The text is not a research monograph on open technical issues related to ATM, such as traffic deseriptors, ingress/egress traffic policing, object-oriented signaling, etc. A literature search undertaken in the spring of 1993 showed that about 5000 papers and trade articles have been written on ATM in the previous nine years, including Refs. 7 through 15. The purpose of this book, therefore, is to stick to the facts and avoid unnecessary hype. There are a few books already available, but these tend to focus on protocol issues. This text aims at a balance between standards, platforms, interworking, and, most important, deployment issues.

In summary, a network supporting cell relay service accepts user data units (called cells) formatted according to a certain layout and sends these data units in a connection-oriented manner (i.e., via a fixed established path), with sequentiality of delivery, to a remote recipient (or recipients). Every so often a cell may be dropped by the network to deal with network congestion; however, this is a very rare event. The user needs a signaling mechanism in order to tell the network what he or she needs. The signaling mechanism consists of a Data Link Layer capability (where the Data Link Layer has been partitioned into four sublayers) and an application-level call-control layer. ATM switches and other network elements supporting cell relay service can also support other fastpacket services. If the user wishes to use ATM to achieve a circuit-emulated service, certain adaptation protocols in the user equipment will be required. Other adaptation protocols in the user equipment are also needed to obtain fastpacket services over an ATM platform. ATM supports certain operations and maintenance procedures that enable both the user and the provider to monitor the "health" of the network. Figure 1.1 is a physical view of an ATM network.

A glossary of some of the key ATM and related concepts, based on a variety of ATM standards and documents, is given in Table 1.2

1.1.3 Early corporate applications of ATM

Table 1.3 depicts some of the proposed applications for ATM/cell relay service.

TABLE 1.1 Areas of Investigation In This Text

1.	ATM and cell relay service: an overview
2.	ATM platform aspects and ATM proper
3.	ATM Adaptation Layer
4.	Signaling
5.	Cell relay service-a formal definition
6.	Cell relay service-traffic and performance issues
7.	Support of fastpacket services and CPE
8.	ATM interworking: support of basic multimedia
9.	Third-generation LANs
10.	Network management
11.	Typical user equipment and public carrier service availability
12.	How to migrate a pre-ATM enterprise network to CRS

1.2 Basic ATM Concepts

1.2.1 ATM protocol model: an overview

ATM's functionality corresponds to the Physical Layer and *part* of the Data Link Layer of the Open Systems Interconnection Reference Model (OSIRM). This protocol functionality must be implemented in appropriate user equipment (for example, routers, hubs, and multiplexers) and in appropriate network elements (for example, switches and service multiplexers). A *cell* is a block of information of short fixed length (53 octets) that is composed of an "overhead" section and a payload section (5 of the 53 octets are for overhead and 48 are for user information), as shown in Fig. 1.2. Effectively, the cell corresponds to the Data Link Layer frame that is taken as the atomic building block of the cell relay service. The term *cell relay* is used because ATM transports user cells reliably and expeditiously across the network to their destination. ATM is a transfer mode in which the information is organized into cells; it is asynchronous in the sense that the recurrence of cells containing information from an individual user is not necessarily periodic.

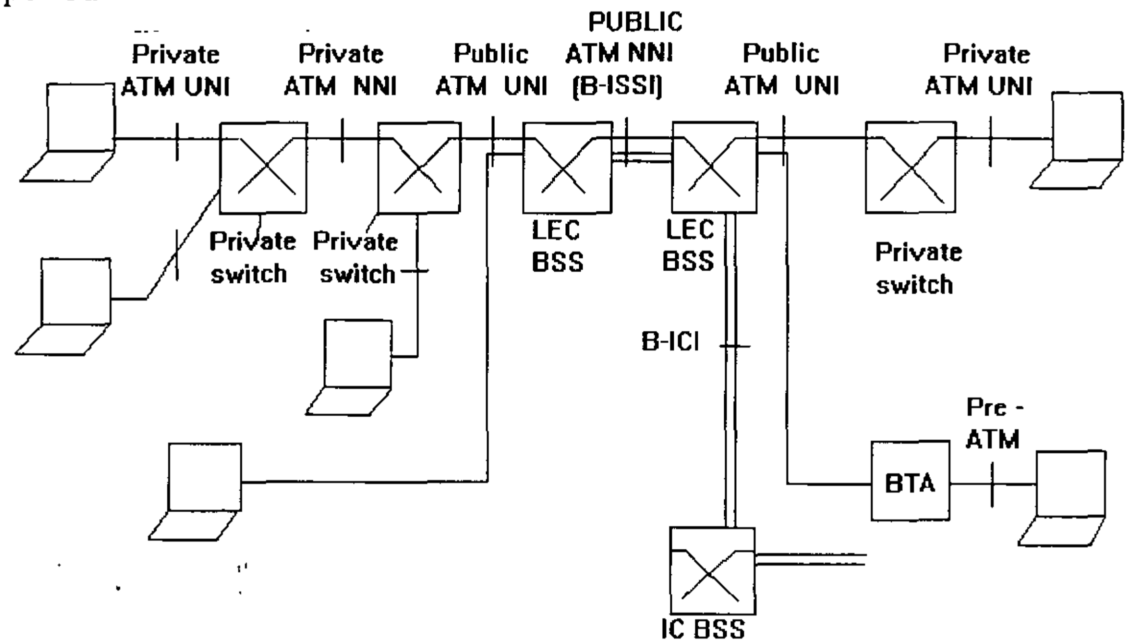

Figure 1.1 A physical view of an ATM/CRS private/public network. BSS = broadband switching system (B-ISDN switch); BTA = broadband terminal adapter; B-ISSI = broadband interswitching system interface; BICI = broadband intercarrier interface; LEC = local exchange carrier; IC = interexchange carrier.

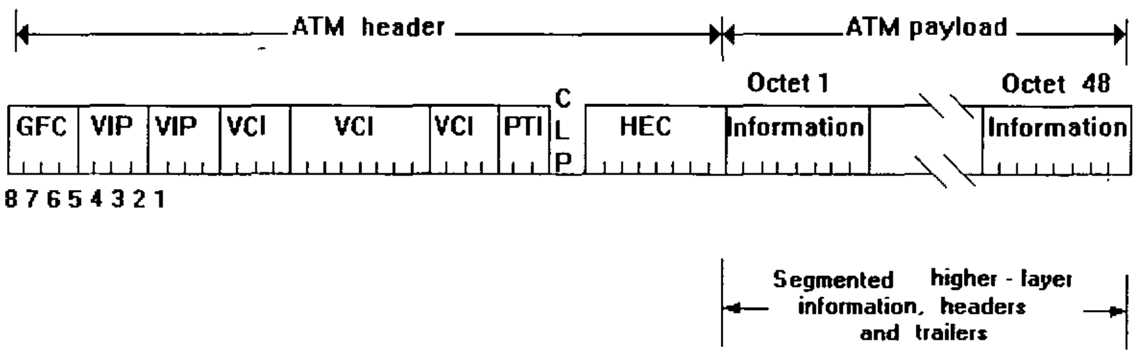

Figure 1.2 ATM cell layout

The ATM architecture utilizes a logical protocol model to describe the functionality it supports. The ATM logical model is composed of a User Plane, a Control Plane, and a Management Plane. The *User Plane* with its layered structure, supports user information transfer. Above the Physical Layer, the ATM Layer provides information transfer for all applications the user may contemplate; the ATM Adaptation Layer (AAL), along with associated services and protocols, provides service-dependent functions to the layer above the AAL.

TABLE 1.2 Glossary of Key ATM Terms

AAL	A layer that adapts higher-layer user protocols (e.g., TC/IP, APPN) to the ATM protocol (layer)
AAL connection	An association established by the AAL between two or more next higher layer entities
Asynchronous time-division multiplexing	A multiplexing technique in which a transmission capability is organized in a priori unassigned time slots The time slots are assigned to cells upon request of each application's instantaneous real need.
Asynchronous transfer mode	A transfer mode in which the information is organized into cells. It is asynchronous in the sense that the recur rence of cells containing information from an individual user is not necessarily periodic.
ATM Layer connection	An association established by the ATM Layer to support communication between two or more ATM service users (i.e., between two or more next higher layer entities or between two or more ATM management entities). The communication over an ATM Layer connection may be either it is bidirectional or unidirectional. When it is bidirectional, two VCCs are used. When it is unidirectional, only one VCC is used.
ATM Layer link	A section of an ATM Layer connection between two adjacent active ATM Layer entities (ATM entities).
ATM link	A virtual path link (VPL) or a virtual channel link (VCL).
ATM peer-to-peer connection	A virtual channel connection (VCC) or a virtual path connection (VPC).
ATM traffic descriptor	A generic list of traffic parameters that can be used to capture the intrinsic traffic characteristics of a requested ATM connection.
ATM user-user connection	An association established by the ATM Layer to support communication between two or more ATM service users [i.e., between two or more next-higher-layer entities or between two or more ATM management (ATMM) entities]. The communication over an ATM Layer connection may be either bidirectional or unidirectional. When it is bidirectional, two VCCs are used. When it is unidirec tional, only one VCC is used.
Broadband	A service or system requiring transmission channels capable of supporting rates greater than the Integrated Service Digital Network (ISDN) primary rate.
Call	An association between two or more users or between a user and a network entity that is established by the use of network capabilities. This association may have zero or more connections.
Cell	ATM Layer protocol data unit.
Cell delay variation	A quantification of variability in cell delay for an ATM Layer connection.

TABLE 1.2 Glossary of Key ATM Terms *(continued)*

Cell header	ATM Layer protocol control information.
Cell loss ratio	The ratio ofthe nutnber of cells "lost" by the network (i.e., cells transmitted into the network but not received at the egress ofthe network) to the number ofcells transmitted to the network.
Cell transfer delay	The transit delay of an ATM cell successfully passed between two designated boundaries.
Connection	The concatenation ofATM Layer links in order to provide an end-to-end information transfer capability to access points.
Connection admission control (CAC)	The procedure used to decide if a request for an ATM connection can be accepted based on the attributes of both the requested connection and the existing connections.
Connection endpoint (CE)	A terminator at one end of a layer connection within a SAP.
Connection endpoint identifier (CEI)	Identifier of a CE that can be used to identify the connection at a SAP.
Corresponding entities	Peer entities with a lower-layer connection among them
Header	Protocol control information located at the beginning of a protocol data unit.
Layer connection	A capability that enables two remote peers at the same layer to exchange information.
Layer entity	An active element within a layer.
Layer function	A part of the activity of the layer entities.
Layer service	A capability of a layer and the layers beneath it that is provided to the upper-layer entities at the boundary between the layer and the next higher layer.
Layer user data	Data transferred between corresponding entities on behalf of the upper-layer or layer management entities for which they are providing services.
Multipoint access	User access in which more than one terminal equipment (TE) is supported by a single network termination.
Multipoint-to-multipoint connection	A collection of associated ATM VC or VP links and their associated endpoint nodes, with the following properties: (1) All N nodes in the connection, called endpoints, serve as root nodes in a point-to-multipoint connection to all of the (N - 1) remaining endpoints. (2) Each of the endpoints on the connection can send information directly to any other endpoint .[the receiving endpoint cannot distinguish which of the endpoints is sending information without additional (e.g., higher-layer) information].
Multipoint-to-point connection	A multipoint-to-point connection where the bandwidth from the root node to the leaf nodes is zero, and the return bandwidth from the leaf node to the root node is nonzero.
Network node interface (NNI)	The interface between two network nodes.

TABLE 1.2 Glossary of Key ATM Terms *(continued)*

Operation and maintenance (OAM) cell	A cell that contains ATM Layer Management (LM) information. It does not form part of the upper-layer information transfer.
Peer entities	Entities within the same layer.
Physical Layer (PHY) connection	An association established by the PHY between two or more ATM, entities. A PHY connection consists of the concatenation of PHY links in order to provide an end-to-end transfer capability to PHY SAPs.
Point-to-multipoint connection	A collection of associated ATM VC or VP links, with associated endpoint nodes, with the following properties: (1) One ATM link, called the root link, serves as the root in a simple tree topology. When the root node sends information, all of the remaining nodes on the connection, called Leaf Nodes, receive copies ofthe information. (2) Each of the leaf nodes on the connection can send information directly to the root node. The root node cannot distinguish which leaf is sending information without additional (higher-layer) information. (3) The leaf nodes cannot communicate with one another directly with this connection type.
Point-to-point connection	A connection with only two endpoints.
Primitive	An abstract, implementation-independent interaction between a layer service user and a layer service provider or between a layer and the Management Plane.
Protocol	A set of rules and formats (semantic and syntactic) that determines the communication behavior of layer entities in the performance of the layer functions.
Protocol control informa tion (PCI)	Information exchanged between corresponding entities, using a lower-layer connection, to coordinate their joint operation.
Protocol data unit (PDU)	A unit of data specified in a layer protocol and consisting of protocol control information and layer user data.
Relaying	A function of a layer by means of which a layer entity receives data from a corresponding entity and transmits them to another corresponding entity.
Service access point (SAP)	The point at which an entity of a layer provides services to its layer management entity or to an entity of the next higher layer.
Service data unit (SDU)	A unit of interface information whose identity is preserved from one end of a layer connection to the other.
Source traffic descriptor	A set of trafric parameters belonging to the ATM traffic descriptor used during the connection setup to capture the intrinsic traffic characteristics of the connection requested by the source.

TABLE 1.2 Glossary of Key ATM Terms *(continued)*

Structured data transfer	The transfer of AAL user information supported by the CBR AAL when the AAL user data transferred by theAAL are organized into data blocks with a fixed length corresponding to an integral number of octets.
Sublayer	A logical subdivision of a layer.
Switched connection	A connection established via signaling.
Symmetric connection	A connection with the same bandwidth value specified for both directions.
Traffic parameter	A parameter for specifying a particular traffic aspect of a connection.
Trailer	Protocol control information located at the end of a PDU.
Transit delay	The time difference between the instant at which the first bit of a PDU crosses one designated boundary and the instant at which the last bit of the same PDU crosses a second designated boundary.
Unstructured data transfer	The transfer of AAL user information supported by the CBR AAL when the AAL user data transferred by the AAL are not organized into data blocks.
Virtual channel (VC)	A communication channel that provides for the sequential unidirectional transport of ATM cells.
Virtual channel connection (VCC)	A concatenation of VCLs that extends between the points where the ATM service users access the ATM Layer. The points at which the ATM cell payload is passed to or received from the user of the ATM Layer (i.e., a higher layer or ATM management entity) for processing signify the endpoints of a VCC. VCCs are unidirectional
Virtual channel link (VCL)	A means of unidirectional transport of ATM cells between the point where a VCI value is assigned and the point where that value is translated or removed.
Virtual path (VP)	A unidirectional logical association or bundle of VCs.
Virtual path connection (VPC)	A concatenation of VPLs between virtual path terminators (VPTs). VPCs are unidirectional.
Virtual path link (VPL)	A means of unidirectional transport of ATM cells between the point where a VPI value is assigned and the point where that value is translated or removed.

In approximate terms, the AAL supplies the balance of the Data Link Layer not included in the ATM Layer. The AAL supports error checking, multiplexing, segmentation, and reassembly. It is generally implemented in user equipment but may occasionally be implemented in the network at an interworking (i.e., protocol conversion) point. The *Control Plane* also has a layered architecture and supports the call control and connection functions. The *Control Plane* uses AAL capabilities as seen in Fig. 1.3; the layer above the AAL in the Control Plane provides call control and connection control.

TABLE 1.3 Possible early applications of ATM in real enviroments (partial list)

Application	Advantages of ATM use	Associated true-to-life business issues
WAN interconnection of existing enterprise network	High bandwidth; switched service	Unknown cost; geographic availability, equipment availability
WAN interconnection of existing LAN, especially FDDI (fiber distributed data interface) LANs	High bandwidth; switched service	Unknown cost; geographic availability
WAN interconnection of mainframe and supercomputer channel	High bandwidth; only service that supports required throughput (200 Mbits/s); switched service	Unknown cost; geographic availability; equipment availability
WAN interconnection of ATM-based LANs	High bandwidth; switched service; multipoint connectivity	New application, not widely deployed; unproven business need; unknown cost; geographic availability
Support of distributed multimedia	High bandwidth; switched service; multipoint connectivity	New application, not widely deployed; unproven business need; unknown cost; geographic availability
Support of statewide distance learning with two way video	High bandwidth; switched service; multipoint connectivity	New application, not widely deployed; unproven market; other solutions exist; unknown cost; geographic availability
Support of videoconferencing (including desktop video)	High bandwidth; switched service; multipoint connectivity	Not widely deployed; unproven market; other solutions exist, particularly at lower end (e.g., 384 Kbits/s H.200 video); unknown cost; geographic availability
Residential distribution of video (video dial tone)	High bandwidth; switched service; multipoint con-nectivity	Unproven market., other solutions exist, particularly CATV, expensive for this market; needs MPEG II (Motion Picture Expert Group) hardware; geographic availability

It deals with the signaling necessary to set up, supervise, and release connections. The *Management Plane* provides network supervision functions. It provides two types of functions: Layer Management and Plane Management. Plane Management performs management functions related to the system as a whole and provides coordination among all planes, Layer Management performs management functions relating to resources and parameters residing in. its protocol entities. See Fig. 1.3. (The various protocols identified in this figure will be discussed at length later.)

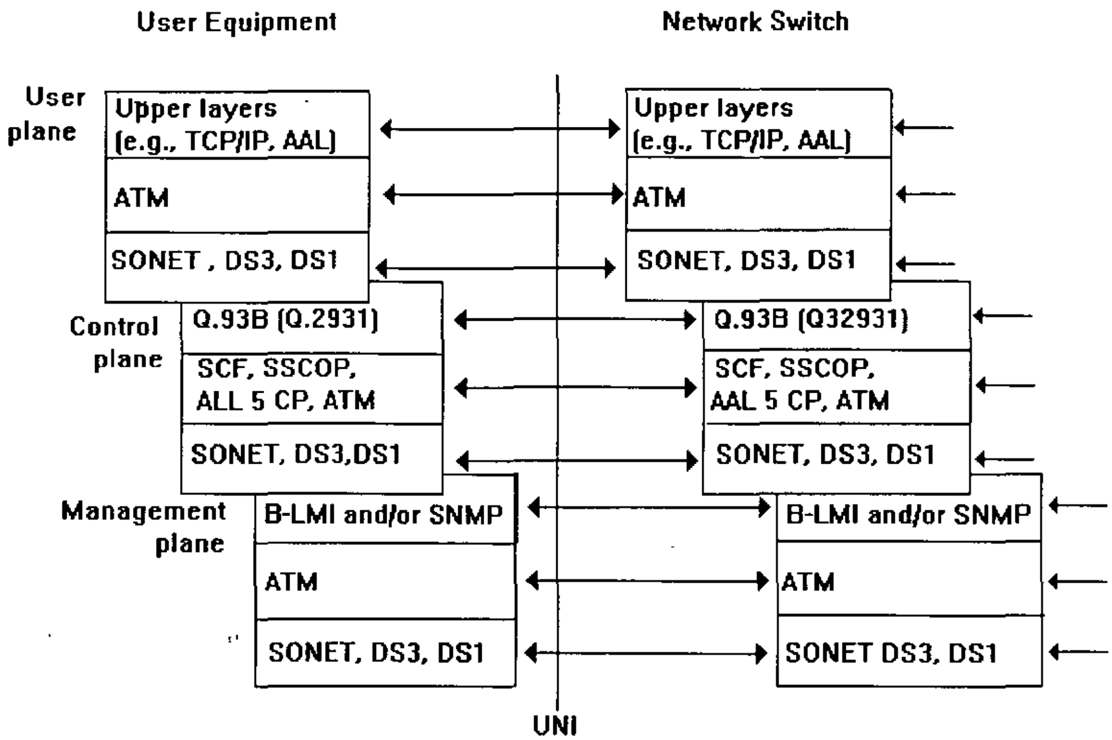

Figure 1.3 Planes constituting the ATM protocol model.

As noted in this description, four User Plane protocol layers are needed to undertake communication in an ATM-based environment:

1. A layer below the ATM Layer, corresponding to the Physical Layer. The function of the Physical Layer is to manage the actual medium-dependent transmission. Synchronous Optical Network (SONET) is the technology of choice for speeds greater than 45 Mbits/s.

2. The ATM Layer (equating approximately, for comparison, to the upper part of a LAN's medium access control layer), which has been found to meet specified objectives of throughput, scalability, interworking, and consistency with international standards. The function of the ATM layer is to provide efficient multiplexing and switching, using cell relay mechanisms.

3. The layer above the ATM Layer, that is, the AAL. The function of the AAL is to insulate the upper layers of the user's application protocols [e.g., TCP/IP (Transmission Control Protocol/Internet Protocol)] from the details of the ATM mechanism.

4. Upper layers, as needed. These include TCP/IP, IBM APPN, OSI TP, etc.

Several layers are needed in the Control Plane. Early PVC service users do not need the signaling stack in the Control Plane (this situation is analogous to the early PVC frame relay environment). SVC service needs both an information transfer protocol stack and a companion signaling protocol stack.

ATM is intended to support a variety of user needs, including highspeed data, video, and multimedia applications. These applications have varying quality of service (QOS) requirements. For example, video-based services have stringent delay, delay variation, and cell loss goals, while other applications have different QOS requirements. Carriers are proposing to support a number of service classes in order to tailor cell relay to a variety of business applications. In particular, there have been proposals to support a "guaranteed" and a "best efforts" class.

1.2.2 Classes of ATM applications.

Two main service categories of ATM have been identified (from the network point of view): (1) interactive broadband service and (2) distributive broadband service. See table 1.4.

1.2.3 Virtual connections

Just as in traditional packet switching or frame relay, information in ATM is sent between two points not over a dedicated, physically owned facility, but over a shared facility composed of virtual channels. Each user is assured that, although other users or other channels belonging to the same user may be present, the user's data can be reliably, rapidly, and securely transmitted over the network in a manner consistent with the subscribed quality of service. The user's data is associated with a specified virtual channel. ATM's " sharing" is not the same as a random access technique used in LANs, where there are no guarantees as to how long it can take for a data block to be transmitted: in ATM, cells coming from the user at a stipulated (subscription) rate are, with a very high probability and with low delay, " guaranteed" delivery at the other end, almost as if the user had a dedicated line between the two points. Of course, the user does not, in fact, have such a dedicated (and expensive) end-to-end facility, but it will seem that way to users and applications on the network. Cell relay service allows for a dynamic transfer rate, specified on a per-call basis. Transfer capacity is assigned by negotiation and is based on the source requirements and the available network capacity. Cell sequence integrity on a virtual channel connection is preserved by ATM.

Cells are identified and switched by means of the label in the header, as seen in Fig. 1.2. In ATM, a *virtual channel* (VC) is used to describe unidirectional transport of ATM cells associated by a common unique identifier value, called the *virtual channel identifier* (VCI). Even though a channel is unidirectional, the channel identifiers are assigned bidirectionally. The bandwidth in the return direction may be assigned symmetrically, or asymmetrically, or it could be zero. A *virtual path* (VP) is used to describe unidirectional transport of ATM cells belonging to virtual channels that are associated by a common identifier value, called the *virtual path identifier* (VPI). See Fig. 1.4.

VPIs are viewed by some as a mechanism for hierarchical addressing. In theory, the VPI/VCI address space allows up to 16 million virtual connections over a single interface; however, most vendors are building equipment supporting (a minimum of) 4096 channels on the user's interface. Note that these labels are only locally significant (at a given interface). They may undergo remapping in the network; however, there is an end-to-end identification of the user's stream so that data can flow reliably. Also note that on the network trunk side more than 4096 channels per interface are supported.

Figure 1.5 illustrates how the VPI/VCI field is used in an ATM WAN. Figure 1.6 depicts the relationship of VPs and VCs as they might be utilized in an enterprise network.

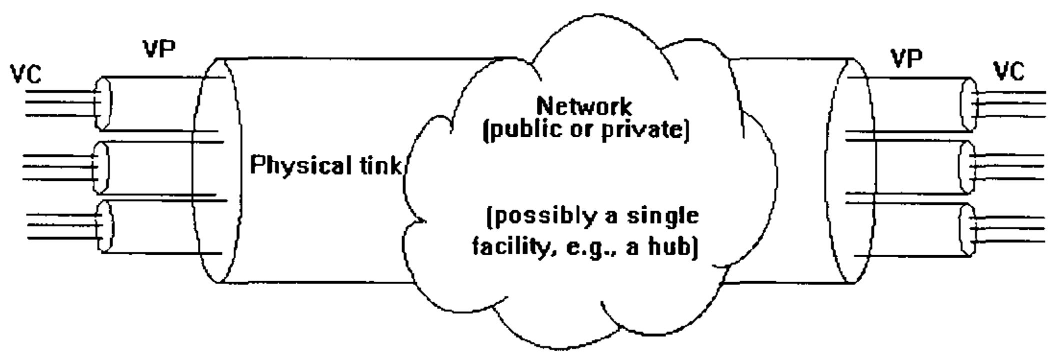

Figure 1.4 Relationship of VCs VPs

TABLE 1.4 Broadband Service Supported by ATM/Cell Relay

Interactive services	*Conversational services* provide the means for bidirectional communication with real-time, end-to-end information transfer between users or between users and servers. Information flow may be bidirectional simmetric or bidirectional asymmetric. Examples: High speed data transmission, image transmission, videotelephony, and videoconferencing. *Messaging services* provide user-to-user communication between individual users via storage units with store-and-forward, mailbox, and/or message handling (e.g., information editing, processing, and conversion) functions. Examples: Message handling services and mail services for moving pictures (films), store-and-forward image and audio information. *Retrieval services* allow users to retrieve information stored in information repositories (information is sent to the user on demand only). The time at which an information sequence is to start is under the control of the use. Examples: Film, high-resolution images, information on CD-ROMs,. and audio information
Distributive services	Distribution services without user individual presentation control provide a continuous flow of information that is distributed from a central source to an unlimited number of authorized receivers connected to the network. The user can access this flow of information without having to determine at which instant the distribution of a string of information will be started. The user cannot control the start and order of the presentation of the broadcast information, so that depending on the point in time of the user's access, the information will not be presented from its beginning. Examples: broadcast of television and audio programs. *Distribution services with user individual presentation control* provide information distribution from a central source to a large number of users. Information is rendered as a sequence of information entities with cyclical repetition. The user has individual access to the cyclically distributed information, and can control the start and order of presentation. Example: broadcast videography.

1.3 ATM Protocols: An Introductory Overview

Figure 1.7 depicts the cell relay protocol environment, which is a particularization of the more general B-ISDN protocol model described earlier. The user's equipment must implement these protocols, as must the network elements to which the user connects. Some of the key functions of each layer are described next.

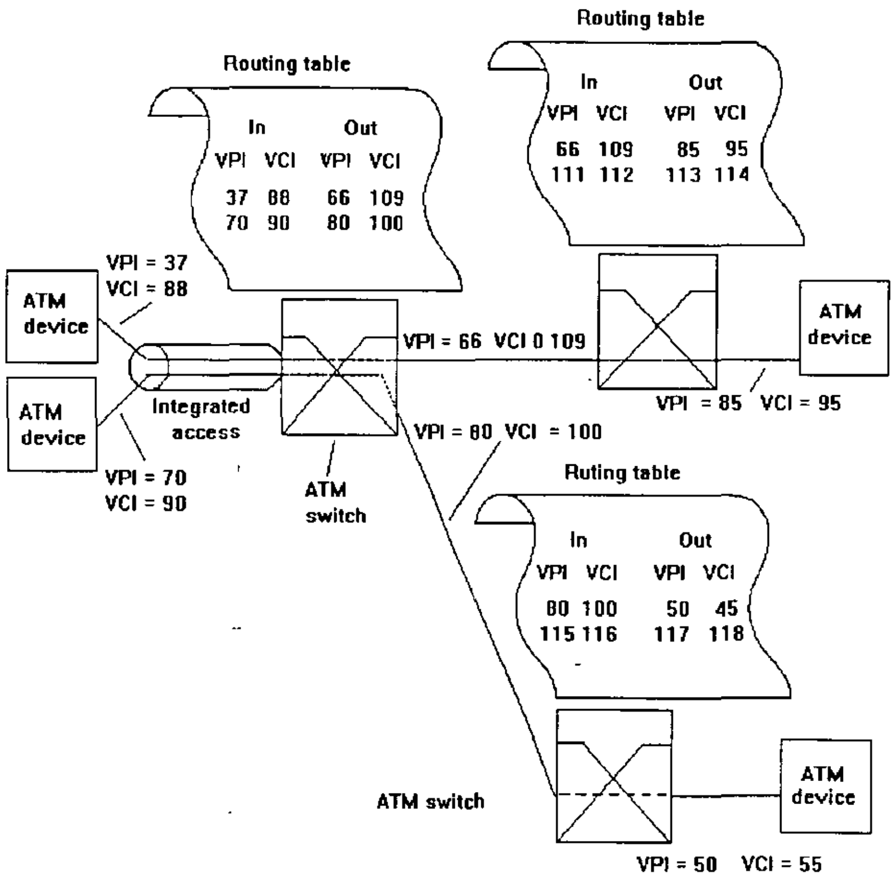

Figure 1.5 Ilustrative use of VPIs and VCIs.

1.3.1 Physical Layer functions

The Physical Layer consists of two *logical* sublayers: the Physical Medium-Dependent (PMD) Sublayer and the Transmission Convergence (TC) Sublayer. The PMD includes only physical medium-dependent functions. It provides bit transmission capability, including bit transfer, bit alignment, line coding, and electrical-optical conversion. The Transmission Convergence Sublayer performs the functions required to transform a flow of cells into a flow of information (i.e., bits) that can be transmitted and received over a physical medium. Transmission Convergence functions include (1) transmission frame generation and recovery, (2) transmission frame adaptation, (3) cell delineation, (4) header error control (HEC) sequence generation and cell header verification, and (5) cell rate decoupling.

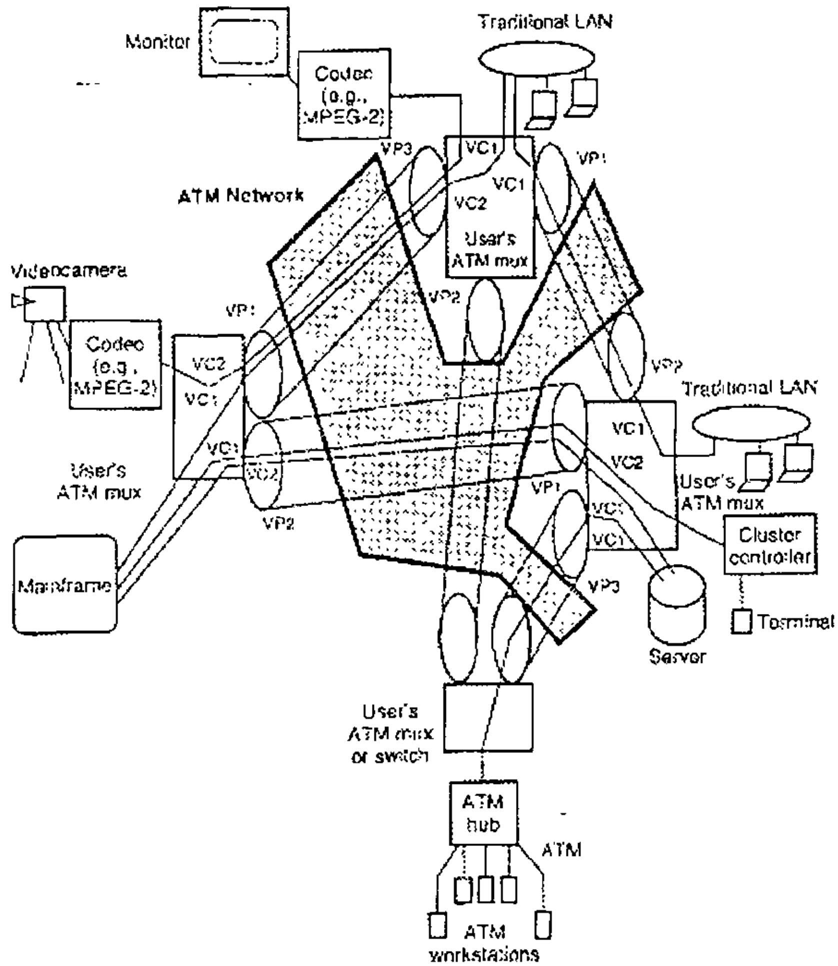

Figure 1.6 Example of use of VPs and VCs in an enterprise network (broadband switches not shown for simplicity). *Note*: VPs and VCs can be preprovisioned (PVCs) or on-demand (SVC whit signaling)

The transmission frame adaptation function performs the actions that are necessary to structure the cell flow according to the payload structure of the transmission frame (transmit direction) and to extract this cell flow out of the transmission frame (receive direction). In the United States, the transmission frame requires SONET envelopes above 45 Mbits/s. Cell delineation prepares the cell flow in order to enable the receiving side to recover cell boundaries. In the transmit direction, the payload of the ATM cell is scrambled. In the receive direction, cell boundaries are identified and confirmed, and the cell flow is descrambled. The HEC mechanism covers the entire cell header, which is available to this layer by the time the cell is passed down to it. The code used for this function is capable of either single-bit correction or multiple-bit error detection. The transmitting side computes the HEC field value. Cell rate decoupling includes insertion and suppression of idle cells, in order to adapt the rate of valid ATM cells to the payload capacity of the transmission system.

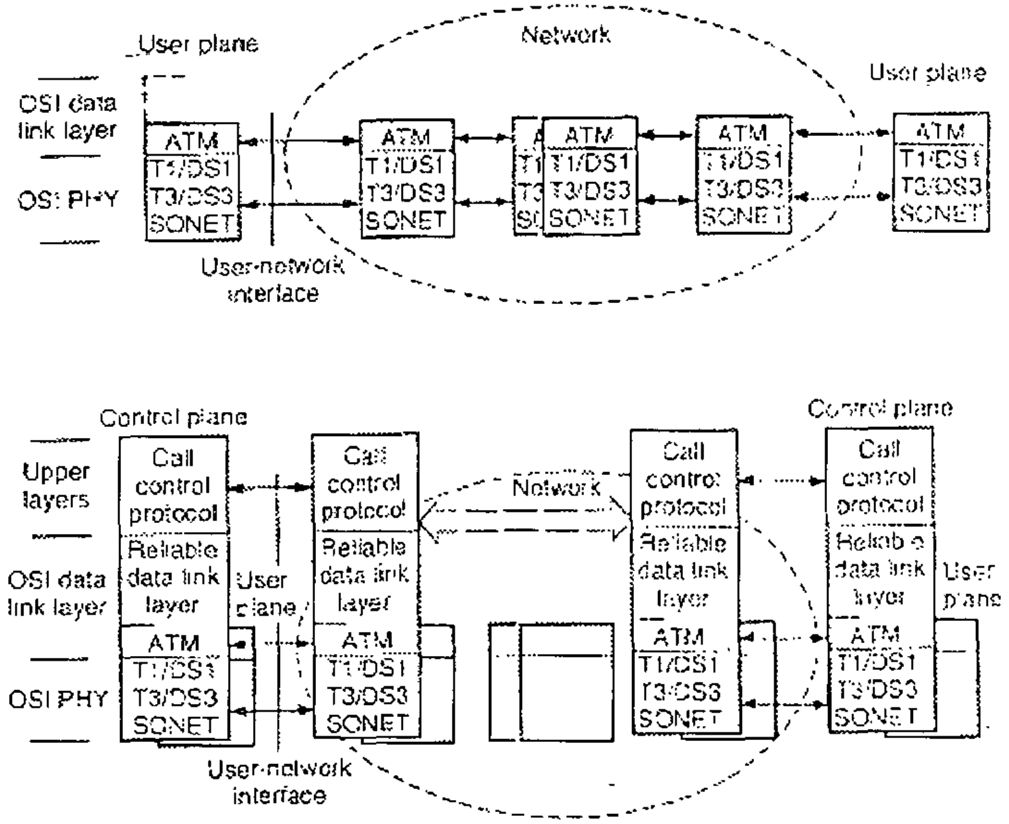

Figure 1.7 CRS environment, protocol view. Top: User Plane (information flow). Bottom· Control Plane (signaling).

The service data units crossing the boundary between the ATM Layer and the Physical Layer constitute a flow of valid cells. The ATM Layer is unique, that is, independent of 'the underlying Physical Layer. The data flow inserted in the transmission system payload is physical medium-independent; the Physical Layer merges the ATM cell flow with the appropriate information for cell delineation, according to the cell delineation mechanism.

The transfer capacity at the UNI is 155.52 Mbits/s, with a cell-fill capacity of 149.76 Mbits/s because of Physical Layer framing overhead. Since the ATM cell has 5 octets of overhead, the 48-octet information field quates to a maximum of 135.631 Mbits/s of actual user information. A second UNI interface is defined at 622.08 Mbits/s, with a service bit rate of approximately 600 Mbits/s. Access at these rates requires a fiber-based loop. Other UNIs at the DS3 rate and perhaps at the DS1 rate are also being contemplated in the United Stlates. The DS1 UNI is discussed in the context of an electrical interface (T1); so is the DS3 UNI.

1.3.2 ATM Layer functions

ATM supports a flexible transfer capability common to all services, including connectionless services (if these are provided). The transport functions of the ATM Layer are independent of the Physical Layer implementation. As noted, connection identifiers are assigned to each link of a connection when required and are released when no longer needed. The label in each ATM cell is used to explicitly identify the VC to which the cells belong. The label consists of two parts: the VCI and the VPI. A VCI identifies a particular VC link for a given virtual path connection (refer to Fig. 1.6). A specific value of VCI is assigned each time a VC is switched in the network. With this in mind, a VC can be defined as a unidirectional capability for the transport of ATM cells between two consecutive ATM entities where the VCI value is translated. A VC link is originated or terminated by the assignment or removal of the VCI value.
The functions of ATM include the following

Cell multiplexing and demultiplexing. In the transmit direction, the cell multiplexing function combines cells from individual VPs and VCs into a noncontinuous composite cell flow. In the receive direction, the cell demultiplexing function directs individual cells from a noncontinuous composite cell flow to the appropriate VP or VC.

Virtual path identifier and virtual channel identifier translation. This function occurs at ATM switching points and/or cross-connect nodes. The value of the VPI and/or VCI field of each incoming ATM cell is mapped into a new VPI and/or VCI value (this mapping function could be null).

Cell header generation / extraction. These functions apply at points where the ATM Layer is terminated (e.g., user's equipment). The header error control field is used for error management of the header. In the transmit direction, the cell header generation function receives cell payload information from a higher layer and generates an appropriate ATM cell header except for the HEC sequence (which is considered a Physical Layer function). In the receive direction, the cell header extraction function removes the ATM cell header and passes the cell information field to a higher layer.

For the UNI, as can be seen in Fig. 1.2, 24 bits are available for cell routing: 8 bits for the VPI and 16 bits for the VCI. Three bits are available for payload type identification; this is used to provide an indication of whether the cell payload contains user information or network information. In user information cells, the payload consists of user information and, optionally, service adaptation function information. In network information cells, the payload does not form part of the user's information transfer. The header error control field consists of 8 bits.

The initial thinking was that if the cell loss priority (CLP) is set by the user (CLP value is 1), the cell is subject to discard, depending on the network (congestion) conditions. If the CLP is not set (CLP value is 0), the cell has higher priority. More recent thinking proposes not making use of this bit on the part of the user (i.e., it must always be set to 0 by the user).

ATM is discussed further in Chap. 2.

1.3.3 ATM Adaptation Layer

Additional functionality on top of the ATM Layer (i.e., in the ATM Adaptation Layer) may have to be provided by the user (or interworking) equipment to accommodate various services. The ATM Adaptation Layer enhances the services provided by the ATM Layer to support the functions required by the next higher layer. The AAL function is typically implemented in the user's equipment, and the protocol fields it requires are nested within the cells' payload.

The AAL performs functions required by the User, Control, and Management Planes and supports the mapping between the ATM Layer and the next higher layer. Note that a different instance of the AAL functionality is required in each plane. The AAL supports multiple protocols to fit the needs of the different users; hence, it is service-dependent (namely, the functions performed in the AAL depend upon the higher-layer requirements). The AAL isolates the higher layers from the specific characteristics of the ATM Layer by mapping the higher-layer protocol data units into the information field of the ATM cell and viceversa. The AAL entities exchange information with the peer AAL entities to support the AAL functions.

The AAL functions are organized in two logical sublayers, the Convergence Sublayer (CS) and the Segmentation and Reassembly Sublayer (SAR). The function of the CS is to provide the AAL service to the layer above it; this sublayer is service-dependent. The functions of the SAR are (1) segmentation of higher-layer information into a size suitable for the information field of an ATM cell and (2) reassembly of the contents of ATM cell information fields into higher layer information.

Connections in an ATM network support both circuit-mode and packet-mode (connection-oriented and connectionless) services of a single medium and/or mixed media and multimedia. ATM supports two types of traffic: constant bit rate (CBR) and variable bit rate (VBR). CBR transfer rate parameters for on-demand services are negotiated at call setup time. (Changes to traffic rates during the call may eventually be negotiated through the signaling mechanism; however, initial deployments will not support renegotiation of bit rates.) CBR transfer rate parameters for permanent services are agreed upon with the carrier from which the user obtains service. This service would be used, for example, to transmit real-time video. VBR services are described by a number of traffic-related parameters (minimum capacity, maximum capacity, burst length, etc.). VBR supports packet like traffic (e.g., variable-rate video, LAN interconnection, etc.). The AAL protocols are used to support these different connection types.

In order to minimize the number of AAL protocols, however, a service classification is defined based on the following three parameters: (1) the timing relation between source and destination (required or not required), (2) the bit rate (constant or variable, already discussed), and (3) the connection mode (connection-oriented or connectionless). Other parameters, such as assurance of the communication, are treated as quality of service parameters, and therefore do not lead to different service classes for the AAL. The five classes of application are:

Class A service is an on-demand, connection oriented, constant-bit rate ATM transport service. It has end-to-end timing requirements. This service requires stringent cell loss, cell delay, and cell delay, variation performance. The user chooses the desired bandwidth and the appropriate QOS during the signaling phase of an SVC call to establish a Class A connection (in the PVC case, this is prenegotiated). This service can provide the equivalent of a traditional dedicated line and may be used for videoconferencing, multimedia, etc.

Class B service is not currently defined by formal agreements. Eventually it may be used for (unbuffered) compressed video.

Class C service is an on-demand, connection-oriented, variable-bitrate ATM transport service. It has no end-to-end timing requirements. The user chooses the desired bandwidth and QOS during the signaling phase of an SVC call to establish the connection.

Class D service is a connectionless service. It has no end-to-end timing requirements. The user supplies independent data units that are delivered by the network to the destination specified in the data unit. SMDS is an example of a Class D service.

Class X service is an on-demand, connection-oriented ATM transport service where the AAL, traffic type (VBR or CBR), and timing requirements are user-defined (i.e., transparent to the network). The user chooses only the desired bandwidth and QOS during the signaling phase of an SVC call to establish a Class X connection (in the PVC case, this is prenegotiated).

Three AAL protocols have been defined in support of these User Plane applications: AAL Type 1, AAL Type 3/4, and AAL Type 5. Type 1 supports Class A, Type 3/4 supports Class D, and Type 5 supports Class X. It appears that the computer communication community (e.g., LAN and multiplexing equipment) will use AAL Type 5. Additionally, the ATM service likely to be available first (and the one supported by evolving computer equipment vendors) is Class X (that is, cell relay service).

Note that two stacks must be implemented in the user's equipment in order to obtain VCs on demand (i.e., SVC service) from the network. With this capability, the user can set up and take down multiple connections at will. The Control Plane needs its own AAL; there has been agreement to use AAL 5 in the Control Plane. Initially only PVC service will be available in the United States. In this mode, the Control Plane stack is not required, and the desired connections are established at service initiation time and remain active for the duration of the service contract. Also note that AAL functions (SAR and CS) must be provided by the user equipment (except in the case where the network provides interworking functions). Additionally, the user equipment must be able to assemble and disassemble cells (i.e., run the ATM protocol).

AAL is discussed further in Chap. 3. Signaling is discussed in Chap. 4.

1.4 Multiservice ATM Platforms

SMDS and frame relay PVC are currently available fastpacket services. SMDS is a high-performance, packet-switched public data service being deployed by the Regional Bell Operating Companies (RBOCs), GTE, and SNET in the United States. SMDS is also being deployed in Europe. Frame relay PVC is a public data service that is widely available today and is expected to be deployed by all RBOCs and most interexchange carriers by the end of 1994.

Frame relay SVC should be available in the 1994 - 1995 time frame. ATM is a switching and multiplexing technology that is being embraced worldwide by a wide spectrum of carriers and suppliers. This new technology can switch and transport voice, data, and video at very high speeds in a local or wide area. What is the relationship of SMDS and frame relay to ATM?

SMDS and frame relay are carrier services, whereas ATM is a technology, as indicated at the beginning of this chapter. ATM will be used by carriers to provide SMDS, frame relay, and other services, including cell relay service (a fastpacket service based on the native ATM bearer service capabilities). Customers who deploy SMDS or frame relay now will be able to take advantage of the benefits of ATM technology without changing the services they use as carriers upgrade their networks to ATM. The customer's investment in SMDS or frame relay equipment and applications is thus preserved.[17]

SMDS is based on well-defined specifications and provides switched, LAN-like transport across a wide area.[1] SMDS service features include a large maximum packet size, an addressing structure that enables data transfer among all SMDS customers, the ability to send the same SMDS packet to several destinations by specifying one address (group addressing), address screening, and strict quality of service values. As ATM technology is deployed within public carrier networks, SMDS service features will not change. The current SMDS interface between the customer and the network uses an access protocol based on the IEEE 802.6 standard. As ATM technology is deployed, this existing SMDS interface will be maintained. The published requirements for ATM switching and transmission technology specify that the existing well-defined SMDS communications interface with the customer must be supported by ATM. When a carrier introduces ATM-based switching systems, customers need not see any effect on their SMDS service. Any technology conversion will be made within the carrier networks. Thus, customers reap the benefits of the latest technology development, while maintaining a consistency and continuity in the service they already employ. Because ATM and IEEE 802.6 technology are both cell based and have the same size cells, such conversion will be facilitated.[17]

In addition, with the introduction of ATM, SMDS can be combined with other services over a new ATM multiservice communications interface. In this case, the communications interface between the customer and the network is based on ATM protocols for all the services on the multiservice interface, including SMDS and frame relay service. This combination was foreseen in the development of ATM standards. In fact, AAL 3/4 (the ATM Adaptation Layer for SMDS) was specifically designed by ITU–T to carry connectionless services like SMDS. Figure 1.8 depicts the typical platform configuration for carrier-provided ATM-based services.

With its large capacity and multiservice capability, ATM provides SMDS with a faster and more scalable technology platform whose cost can be shared among multiple services. SMDS, along with frame relay PVC, is encouraging the use of high-speed, wide-area public networking in the United States. SMDS and frame relay provide ATM

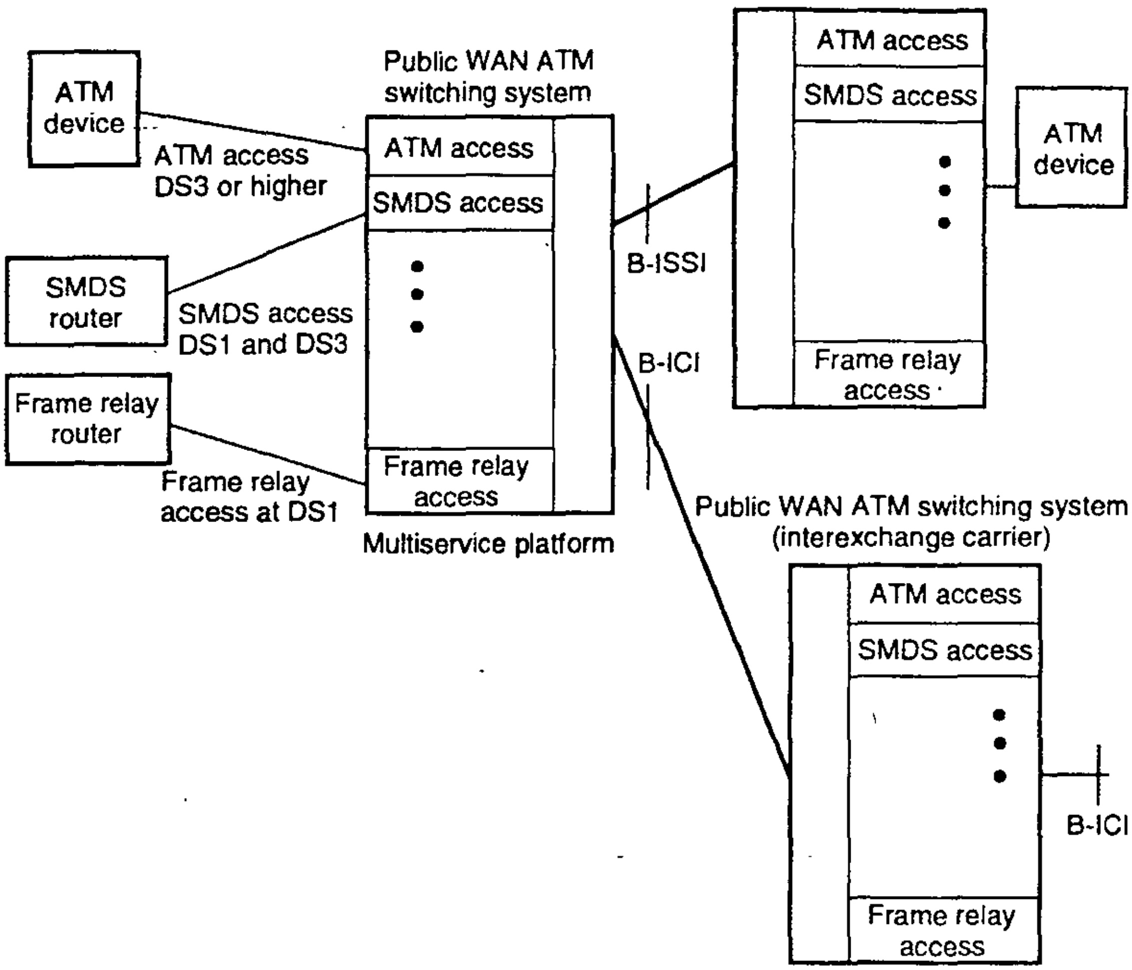

Figure 1.8 Multiservice broadband switching system. B-ISSI = broadband interswitching system interface; B-ICI = broadband interexchange carrier interface.

with significant revenue-producing services that will justify its deployment and allow users and carriers to benefit from the multiservice technology platform.

Frame relay PVC will be a key low-cost, low-overhead broadband data service available in public networks for at least the rest of this decade. The service is currently provided via both frame and ATM switching platforms; ATM simply provides a faster, more scalable platform, as discussed, for SMDS.[18] It appears that frame relay PVC access rates will probably not be extended beyond DS3 (currently, the standards and the deployed services only cover speeds up to 2.048 Mbits/s). This presents the PVC-oriented customer with the possibility of needing to interwork emerging cell relay PVC service with frame relay PVC service. If a user requires PVC service at access speeds of DS3 and above (for example, to aggregate traffic), it is likely that the user will use the ATM cell relay PVC service. This is because customer premises equipment with high-speed wide-area interfaces (e.g., routers) will use ATM technology, thus making cell relay PVC a good choice. As new applications are developed that require these speeds, it is likely that cell relay

PVC service will need to interwork with the users' large installed base of lower-speed wide-area networks for years to come. To meet this need to interwork, the Frame Relay Forum, the ATM Forum, and standards bodies are working on specifications to assure the smooth interworking of these services (ITU-T I.555, in particular).

For the same reasons that carriers are choosing ATM technology (i.e., speed and flexibility), workstation, computer, hub, and LAN manufacturers are turning to ATM for their next-generation networking needs. This is happening because current networks based on Ethernet, FDDI, etc., have limitations when handling the multimedia communications (video, voice, and data) that will flow among future workstations in a network. These manufacturers see global multimedia communications among devices as essential. To meet these networking needs, future workstations and computers will transport user information in ATM cells. Public carriers will offer cell relay service that will transport ATM cells across metropolitan area networks (MANs), across WANs, and internationally as networks evolve. Cell relay service is targeted initially toward high-end users with multimedia needs to transport video, voice, and data across their WANs. When ATM technology extends from the desktop and throughout the network, cell relay service will join SMDS and frame relay as another service that data communications managers can use to support evolving high-bandwidth corporate applications.

Cell relay service is described in Chaps. 5 and 6. Additional aspects of fastpacket are covered in Chap. 7.

1.5 Commercial Availability of ATM Equipment and Network Services

As with any other service, at least three parties are needed to make this technology a commercial reality (if any of these three parties fails to support the service, the service will not see any measurable commercial deployment): (1) carriers must deploy the service, (2) equipment manufacturers must bring user products to the market,* and (3) users must be willing to incorporate the service in their networks. (Some observers add two more forces: agencies supporting R&D and standardization, and the trade press to "educate" the end users.) The early phases of ATM research, including all of the work already accomplished in standards organizations (that is, the topics treated in Chaps. 2 through 10 of this book), cover the first item. The industry activity discussed briefly below and in Chap. 11 covers the second item. The user analysis that will follow (not covered

*In order for item 1 to occur, some vendors must bring out network products; this point refers to user products (see ISDN switches versus availability of cost-effective terminal adapters).

in this book), where users assess applicability, cost, support of embedded base, and manageability, all of it in situ, in their own environment (rather than in a multicolor brochure), covers the third item.

The paragraphs to follow describe industry activities that show encouraging signs of the acceptance of cell relay as a commercially viable networking technology. However, as with all new technologies, there are a number of potential hurdles and roadblocks that can delay or deter its success. History has shown that in spite of industry standards, interoperability problems can exist if different manufacturers implement subsets (or supersets) of the required networking features. Networking hardware may preceed the availability of software applications designed to exploit the networking power of ATM, and this may slow user acceptance of cell relay. In addition, advances in existing technologies (e.g., the emergence of "fast" Ethernet) may extend the life cycle of existing products and slow the acceptance of new technologies. These challenges must be met to make ATM cell relay a long-term commercial success.

Vendors are in the process of bringing products to the market. By 1994 there already were several vendors of ATM hubs and a dozen vendors of ATM workstation plug-ins. Some equipment vendors are building stand-alone premises switches; others are adding switching capabilities to their hubs and at the same time are developing ATM adapter cards for workstations to allow them to connect to the hub. Some are also working on bridge-router cards for ATM hubs that enable Ethernet LANs to connect to ATM. About three dozen vendors had announced firm equipment plans by publication time. Over 320 companies have joined the ATM Forum, which is an organization whose goal is to expedite and facilitate the introduction of ATM-based services. PC/workstation cards are expected to become available for about $1000 per port, although the initial cost was in the $2800–5000 range.

Carriers are deploying broadband switching systems (BSSs) based on ATM technology to support a variety of services. As noted earlier, ATM is designed to be a multi-service platform. For example, frame relay and SMDS will be early services supported on these platforms; another early service is cell relay service, which allows users to connect their ATM equipment using the native ATM bearer service.

Early entrants, including Adaptive, AT&T Network Systems, Cabletron, Digital Equipment Corporation, Fore Systems, Fujitsu, GDC, Hughes, Newbridge, Stratacom, Sun, SynOptics, and Wellfleet, were demonstrating ready or near-ready products for a variety of user networking needs in 1994. The first products were targeted to the local connectivity environment, but WAN products are also expected soon. Additionally, about a dozen vendors have working carrier-grade switching products.

Hubs and switches to support the bandwidth-intensive applications listed earlier, such as video, are becoming available. Typical premises switches now support 8 to 16 155-Mbits/s ports over shielded twisted pair or multimode fibers [lower speeds (45 or 100 Mbits/s) are also supported]. Some systems can grow to 100 ports. Typical backplane throughput ranges from 1 or 2 Gbits/s, up to 10 Gbits/s. A number of these products support not only PVC but also SVC; some also support multipoint SVC service. Products already on the market (e.g., from Hughes LAN, Synoptics, Newbridge, Adaptive, Fore Systems, etc.) are priced as low as $1500 per port. Some of the hubs also act as multiprotocol routers, either (1) accepting ATM devices internally for WAN interconnection over SMDS and frame relay networks, (2) accepting ATM devices internally for WAN interconnection over a cell relay network, or (3) accepting traditional devices internally for WAN interconnection over a cell relay network (these are stand-alone ATM multiprotocol routers).

One major push now is in the network management arena. Users need the capability to integrate the support of ATM products into the overall enterprise network, specifically the corporate management system. Some typical features recently introduced include automatic reconfiguration of virtual connections in case of failure, loopback support, performance and configuration management, and Simple Network Management Protocol (SNMP) functionality [with private management information base (MIB) extensions].

Interface cards for high-end workstations (e.g., SPARCstation) are also appearing (e.g., Synoptics, Adaptive, etc.). These typically support 45 Mbits/s (DS3) on twisted-pair cable and 100 or 155 Mbits/s on multimode fiber, consistent with the ATM Forum specification. Some even support prototype 155-Mbits/s connectivity on shielded twisted pair. These boards are already available for as little as $1250.

Specifically for WAN cell relay service, Sprint has already demonstrated a prototype service operating at the DS3 rate. A three-phase approach has been announced publicly by the company. Phase 1 (1993) entails frame relay interconnectivity with local exchange carriers, Phase 2 (1993–1994) supports PVC cell relay service at the DS3 rate, and Phase 3 (1994–1995) enhances the Cell Relay Service to 155 Mbits/s. AT&T, Wiltel, BellSouth, NYNEX, and Pacific Bell have also announced deployment plans for ATM platforms and for cell relay service. There is strong support for the introduction of cell relay service at the local level. Now users can expect public cell relay service in a number of key metropolitan areas.

In addition to the international and domestic standards, additional details and clarifications are needed to enable the deployment of the technology. To this end, in 1992, Bellcore completed generic requirements that suppliers need in order to start building ATM equipment that will enable the BOCs to offer PVC cell relay services. Work on generic requirements for ATM equipment that provides SVC cell relay was completed at Bellcore in 1994. In particular, Bellcore has already published (preliminary) requirements to define nationally consistent cell relay PVC exchange and cell relay PVC exchange access services, including

"Cell Relay PVC Exchange Service," 1993 [CR PVC exchange service is a public cell relay intra-LATA service offering from local exchange carriers (LECs)]

"Cell Relay PVC Exchange Access CRS (XA-CRS)," 1993 [a PVC XA-CRS is provided by an LEC to an interexchange carrier (IC) in support of the IC's inter-LATA cell relay PVC offering]

"Cell Relay SVC Exchange Service," 1993

The Framework Advisories, Technical Advisories, and Technical Requirements can be used by (1) LECs interested in providing nationally consistent cell relay PVC exchange service to their customers, (2) suppliers of ATM equipment in the local customer environment (e.g., ATM LANs, ATM routers, ATM DSUs, ATM switches), and (3) suppliers of ATM equipment in LEC networks.

The development of nationally consistent LEC cell relay (as well as an exchange access cell relay) service is critical to provide a consistent set of service features and service operations for customers who will want to use the service on a national basis. The following phases of nationally consistent service have been advanced. It is possible that LECs may be offering "pre-nationally consistent" cell relay PVC to meet customers' near-term demand for the service in the late 1993–early 1994 period. These carriers are expected to support a nationally consistent cell relay PVC exchange service at some point thereafter.

- Phase 1.0: Nationally consistent cell relay PVC exchange service based on a core set of service features by the fourth quarter of 1994. The core set is proposed to be a subset of the preliminary generic requirements published by Bellcore in 1993.
- Phase 2.0: Nationally consistent cell relay PVC exchange service based on generic requirements published by Bellcore in 1994 by the second quarter of 1995. Phase 2.0 builds on the capabilities of Phase 1.0 and supports expanded capabilities in some areas, such as traffic management, congestion management, and customer network management.

- Phase 3.0: This will see the initial support of a cell relay SVC exchange service in mid to late 1995 based on generic requirements expected to be published in 1994.

Figure 1.9 depicts the set of Bellcore generic requirements in support of ATM, SMDS, cell relay, and frame relay.[17] These are just some of the key documents that form the foundation for ATM. Standards bodies such as the ITU-T and ANSI (American National Standards Institute) T1S1, and industry bodies such as the ATM Forum and the Frame Relay Forum also publish related documents.

1.6 Typical Examples of Cell Relay Usage in an Enterprise Context

1.6.1 Front-end and back-end usages

Cell relay/ATM is being contemplated at the local-area network level as well as the wide-area network level. Several approaches have been followed by vendors:

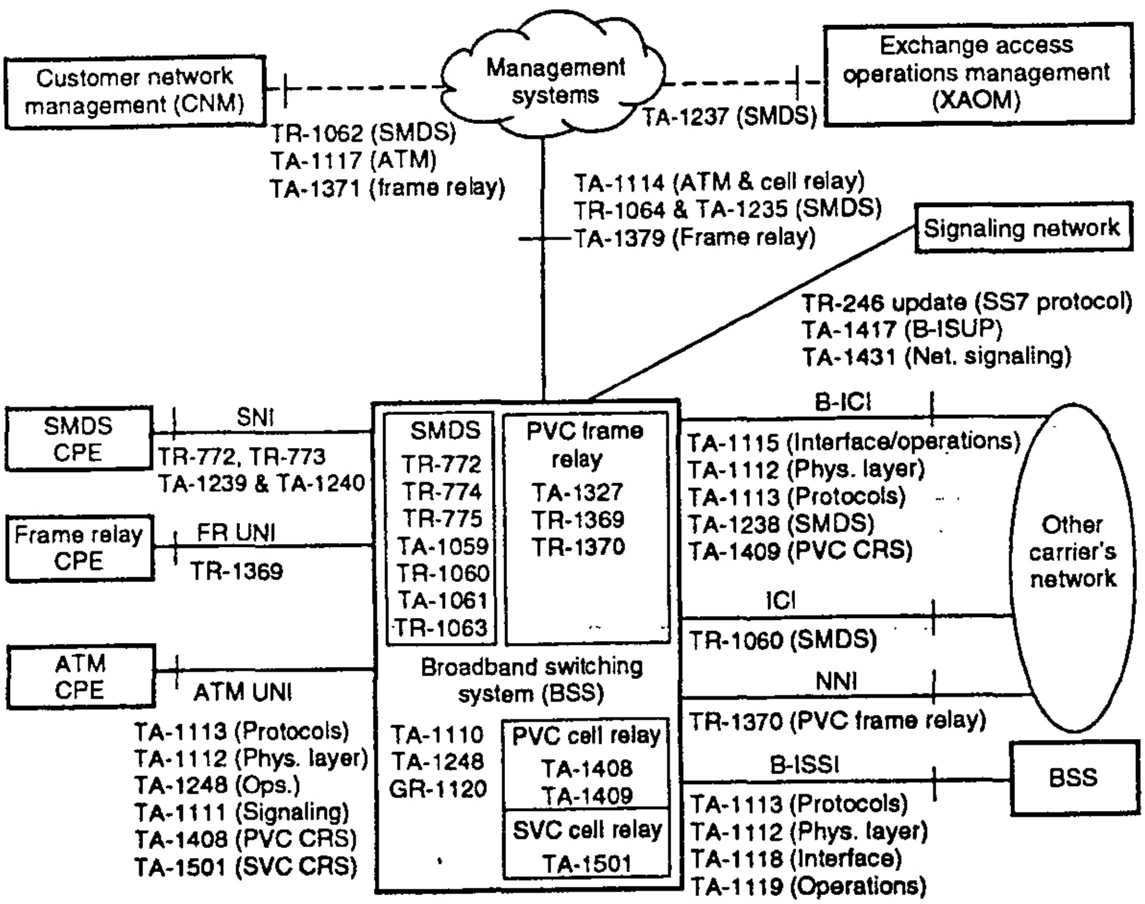

Figure 1.9 ATM, SMDS, cell relay, and PVC frame relay generic requirements.

1. TR-NWT-00246, Bellcore Specification of Signaling System 7, B-ISUP, Issue 2, December 1993.
2. TR-TSV-000772, Generic System Requirements in Support of SMDS, May 1991.
3. TR-TSV-000773, SMDS Requirements, Objectives, and Interfaces, Revision 1, December 1993.
4. TR-TSV-000774, SMDS Operations Technology Network Element Generic Requirements, Issue 1, March 1992, Supp. 1, March 1993.
5. TR-TSV-000775, Usage Measurement Generic Requirements In Support of Billing for Switched Multi-Megabit Data Service, Issue 1, June 1991.
6. TA-NWT-001248, Generic Operations Requirements for Broadband Switching Systems, Issue 2, October 1993.
7. TA-TSV-001059, Generic Requirements for SMDS Networking, Bellcore, Issue 2, August 1992.
8. TR-TSV-001060, Switched Multi-Megabit Data Service Generic Requirements for Exchange Access and Intercompany Serving Arrangements, Issue 1, December 1991, and Revision 1, August 1992; Revision 2, March 1993.
9. TA-TSV-001061, Operations Technology Network Element Generic Requirements in Support of Inter-Switch and Exchange Access SMDS, Issue 1, May 1991.
10. TR-TSV-001062, Generic Requirements for SMDS Customer Network Management Services, Bellcore, Issue 1, March 1992.
11. TR-TSV-001063, Operations Technology Network Element Generic Requirements in Support of Exchange Access SMDS and Intercompany Serving Arrangements, Issue 1, March 1992; Revision 1, March 1993.
12. TR-TSV-001064, SMDS Phase 1 Operations Information Model, December 1993.
13. TA-NWT-001110, Broadband ISDN Switching System Generic Requirements, Issue 2, July 1993.
14. TA-NWT-001111, User to Network Access Signaling Requirements, July 1993.
15. TR-NWT-001112, Broadband ISDN User to Network Interface and Network Node Interface Physical Layer Generic Criteria, July 1993.
16. TA-NWT-001113, Asynchronous Transfer Mode (ATM) and ATM Adaptation Layer (AAL) Protocols Generic Requirements, Issue 2, July 1993.
17. TA-NWT-001114, Generic Requirements for Operations Interfaces Using OSI Tools: Broadband ISDN Operations, Issue 2, October 1993.
18. TA-NWT-001115, Broadband InterCarrier Interface (B-ICI) Requirements, September 1993.
19. TA-NWT-001117, ATM Customer Network Management (CNM), September 1993.
20. TA-TSV-001118, Broadband InterSwitching System Interface (B-ISSI) and Network Generic Requirements, July 1993.
21. TA-NWT-001119, B-ISSI Operations, December 1993.
22. GR-1120-CORE, Guide to Generic Requirements for Usage Information to Support Billing for ATM Broadband Networking, Issue 1, December 1993.
23. TA-NWT-001235, Exchange Access SMDS Operations Interface Model, April 1993.
24. TA-TSV-001237, A Framework for High Level Generic Requirements for SMDS Exchange Access Operations Management Services, July 1993.
25. TA-TSV-001238, SMDS 155 Mbps ATM B-ICI, December 1992.
26. TA-TSV-001239, Low Speed SMDS Access via Data Exchange Interface (DXI), June 1993.
27. TA-TSV-001240, Frame-Based Access to SMDS via SRI, June 1993.
28. TA-NWT-001248, B-ISDN Network Operations Criteria, Issue 2, October 1993.
29. FA-NWT-001327, Frame Relay NE Operations Functional Requirements, Bellcore, Issue 1, December 1992.
30. TR-TSV-001369, Frame Relay (PVC) Exchange Service Definition, May 1993.
31. TR-TSV-001370, Exchange Access Frame Relay (PVC) Service Definition, May 1993.
32. TA-TSV-001371, Frame Relay (PVC) Customer Network Management Service, September 1993.
33. TA-NWT-001379, Frame Relay Network Operations Using OSI, July 1993.
34. TA-TSV-001408, Generic Requirements for Exchange PVC Cell Relay Service, Issue 1, August 1993.
35. TA-TSV-001409, Generic Requirements, Issue 1, November 1993.
36. TA-NWT-001417, B-ISUP Generic Requirements, Issue 1, February 1994.
37. TA-NWT-001431, CCS Network Signaling Specification Supporting B-ISDN Generic Requirements, Issue 1, May 1994.
38. TA-NWT-001501, Generic Requirements for Exchange SVC Cell Relay Service, December 1993.

Figure 1.9 *(Continued)*

1. Use of ATM technology between traditional local or remote LAN hubs; Fig. 1.10 shows a case of interconnection of remote hubs. (The LAN hubs are implicit in the figure.)
2. Introduction of ATM cards on traditional routers for access to a public cell relay service (see Fig. 1.11).
3. Introduction of ATM-based LAN hubs, extending ATM all the way to the desktop, for front-end applications (see Fig. 1.12).
4. Development of private-enterprise ATM switches to support generic corporate networking.
5. Development of carrier-grade multiservice ATM switches (also known as broadband switching systems) to support services such as cell relay service, frame relay service, and SMDS.
6. Development of related equipment (for example, Fig. 1.13 depicts usage in a channel extension environment).

Some industry proponents expect to see Fortune 1000 users passing the majority of their LAN-to-WAN traffic through premises-based ATM switches by 1997. Approximately 50 percent of the ATM traffic in these companies is expected to be in support of LAN interconnection, for LANs serving traditional business applications, and for traditional enterprise data applications, such as mainframe channel extension; the other 50 percent of the traffic is expected to be split fairly evenly among application supporting real-time video, imaging, real-time voice, and multimedia.

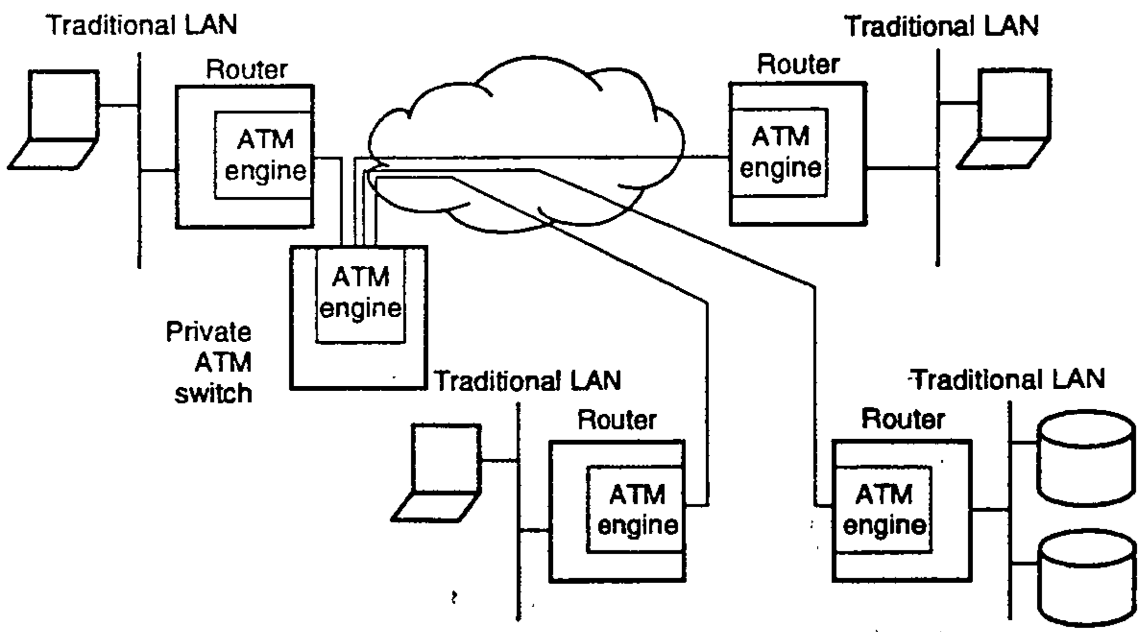

Figure 1.10 Private ATM technology to interconnect dispersed LAN hubs. ATM engine = the logic implementing ATM, control, and, optionally, user plane protocols.

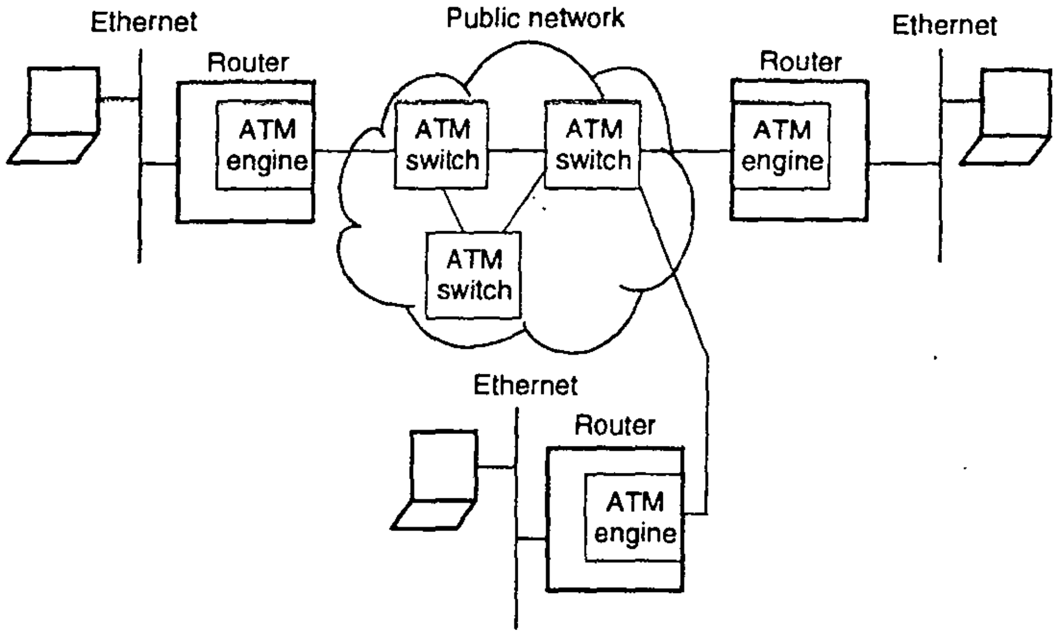

Figure 1.11 Routers used in conjunction with a public cell relay service. ATM engine = the logic implementing ATM, control, and, optionally, user plane protocols.

Figure 1.14 depicts a typical "full-blown" ATM/cell relay arrangement for both WAN and LAN applications. This supports ATM to the desktop for such applications as desk-to-desk videoconferencing and multimedia. Figure 1.15 depicts an example of the protocol machinery across a router/public switch arrangement that is expected to be a common deployment scenario in client/server environments. Figure 1.16 depicts an example in

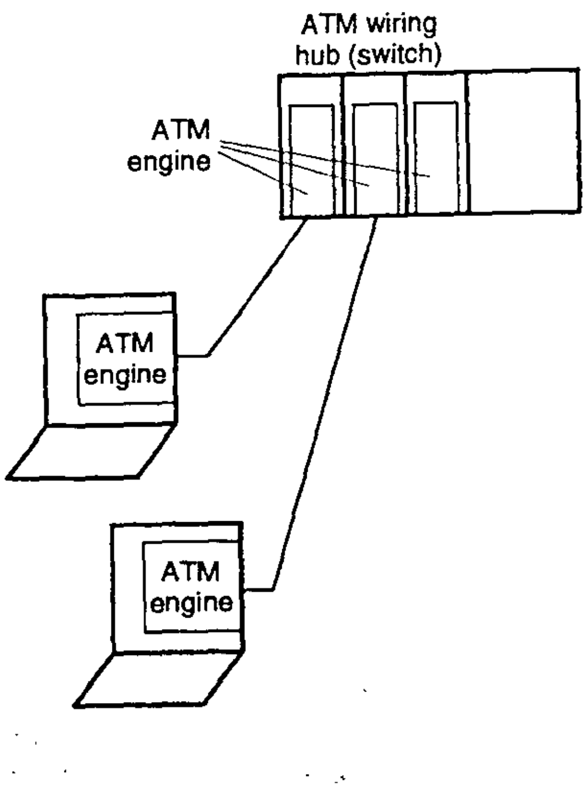

Figure 1.12 ATM to the desktop.

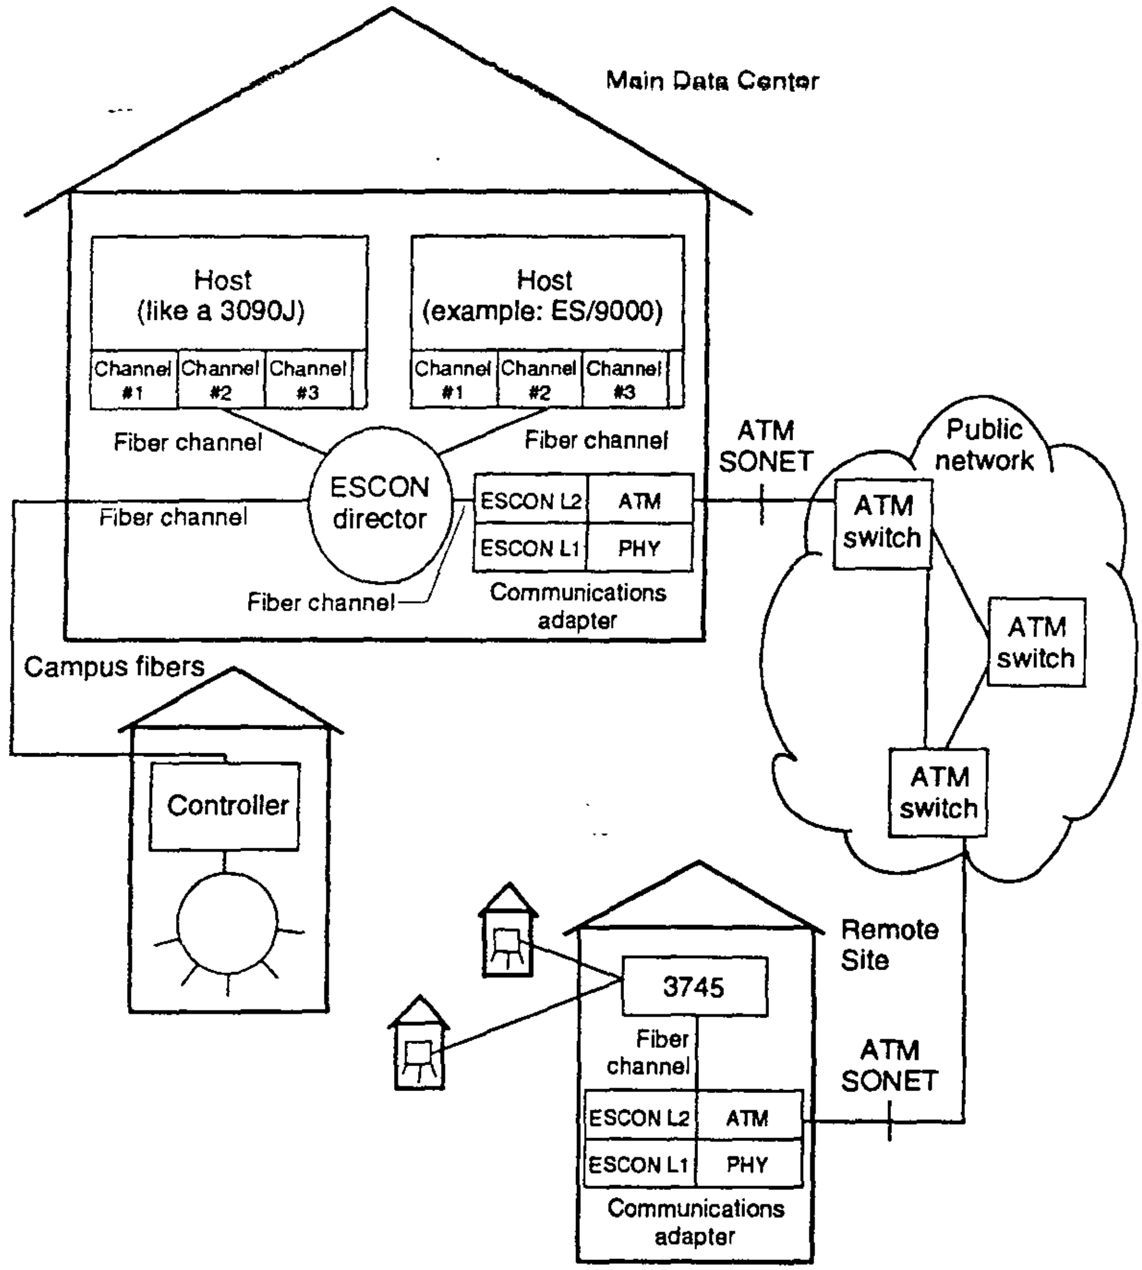

Figure 1.13 Channel extension via ATM services.

a videoconferencing application, also from a protocol point of view. Figure 1.15 shows an example in a corporate network supporting business imaging.

Figure 1.18 depicts a more complete enterprisewise use of cell relay service, while employing a public WAN CRS network. For this example, ATM-ready workstations and devices connected to an ATM-based hub with ATM WAN router capabilities (the router could also be a separate device) can get direct access to the ATM WAN. Some of the hub and router vendors are taking this path to the market. The figure also shows that traditional LAN users can employ an ATM-ready router to obtain the benefit of cell relay WAN services without having to replace their

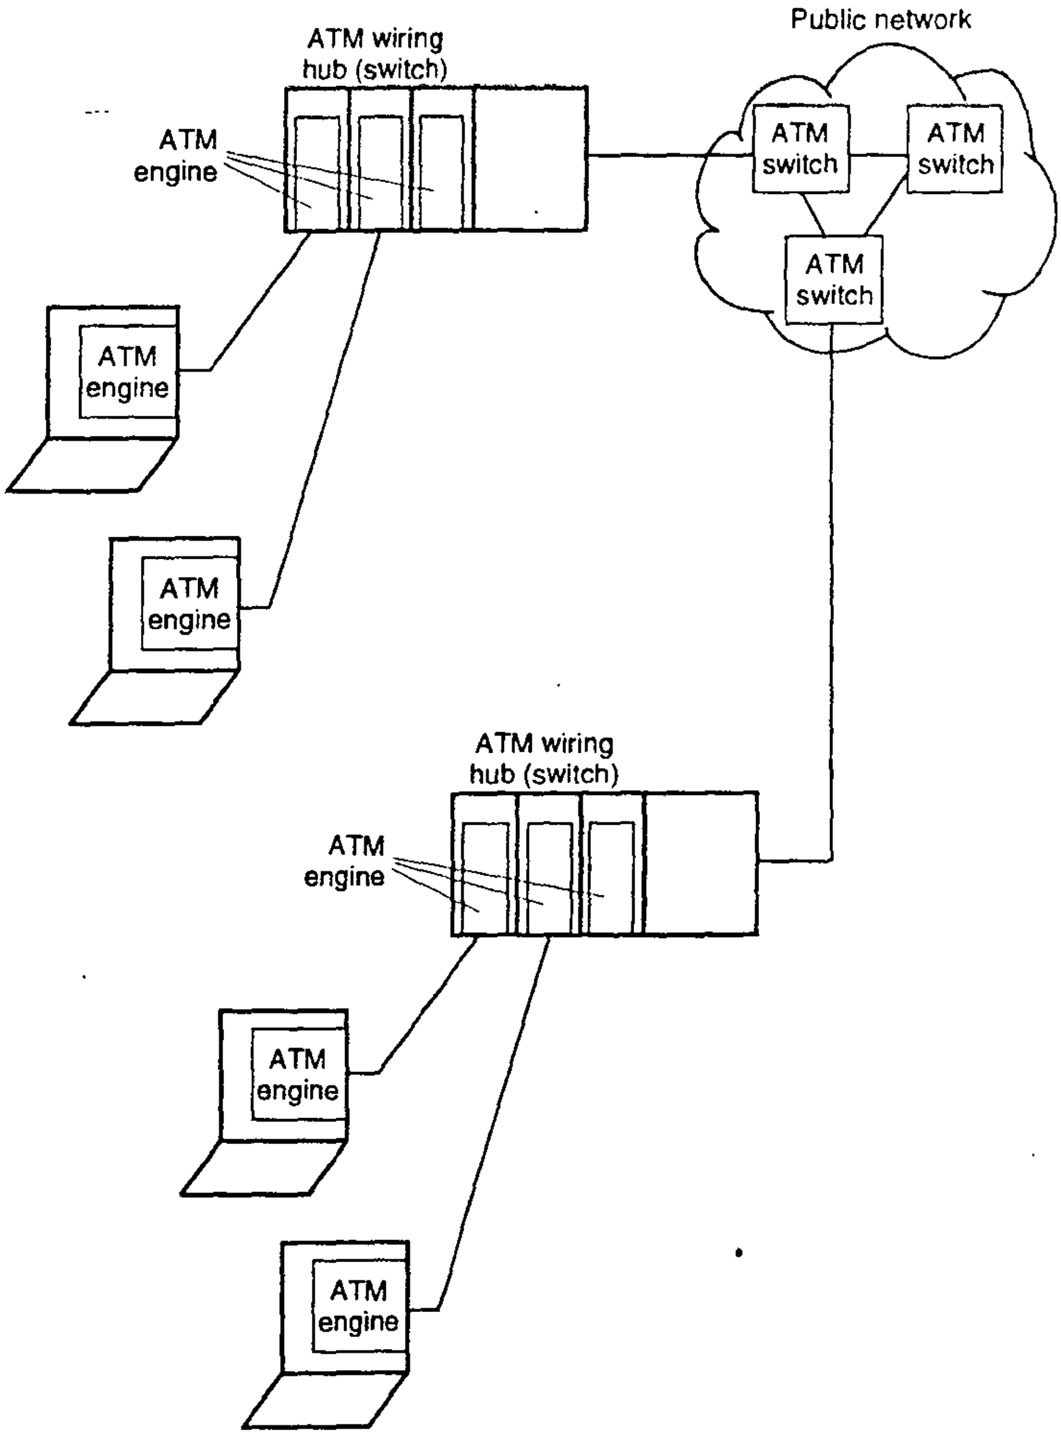

Figure 1.14 Example of usage of ATM in an enterprise network.

desktops or in-house wiring. It also depicts another route to the market, followed by some of the more sophisticated multiplexer manufacturers: The multiplexer can connect traditional data devices, mainframe channels, and video to a cell relay WAN network by supporting ATM on the trunk side. Some of these multiplexers also support traditional LANs on the house side over a frame relay interface. (*Note:* Carrier-deployed ATM "service nodes" in close proximity of the user location but on the

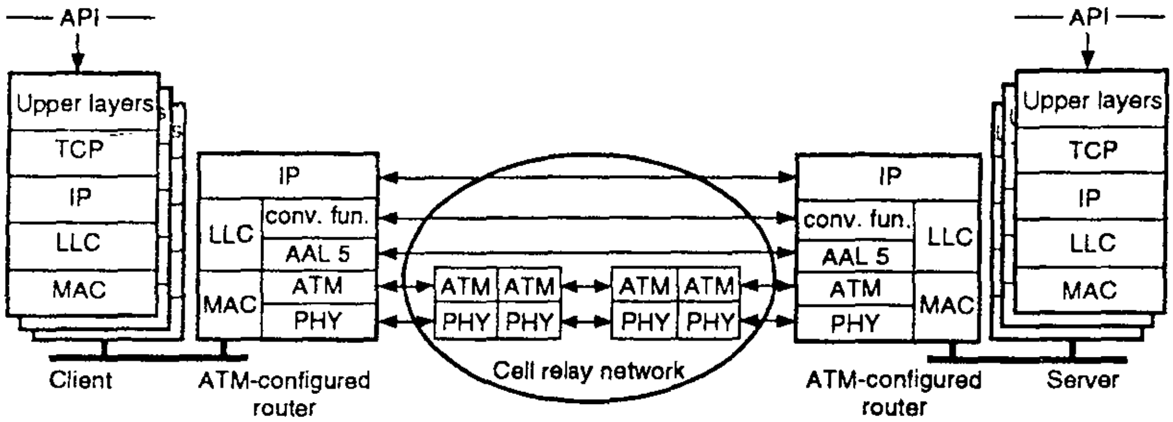

Figure 1.15 Typical corporate application from a protocol-stack point of view. conv.fun. = convergence function.

network side of the interface support these same services plus LAN emulation service.)

Figure 1.19 depicts some user applications of cell relay service in the case where the user wants to develop a private ATM/cell relay service WAN. Note the need to (1)install privately managed switches, (2) use dedicated high-speed WAN lines, and (3) backhaul remote locations to a remote switching site. Public cell relay service may prove less demanding in terms of users' responsibility. Hybrid arrangements are also possible.

1.6.2 Client/server issues

The client/server architecture being put in place in many organizations is truly distributed in the sense that the corporate user has access to data regardless of where the data are located, be they on a system in another campus, another city, another state, or another continent. Client/server applications require extensive interchange of data blocks, often entailing multiple transactions. Low end-to-end delay is critical in making client/server computing possible.[19]

Applications requiring large transfers (e.g., 50–100 kbits) are not unusual in these environments, particularly for imaging video, and

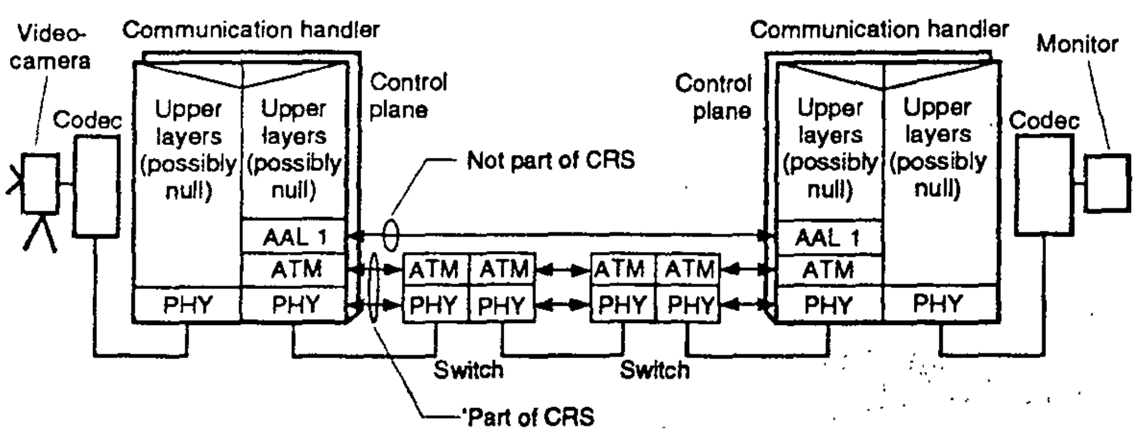

Figure 1.16 Example of video application over ATM/cell relay arrangement.

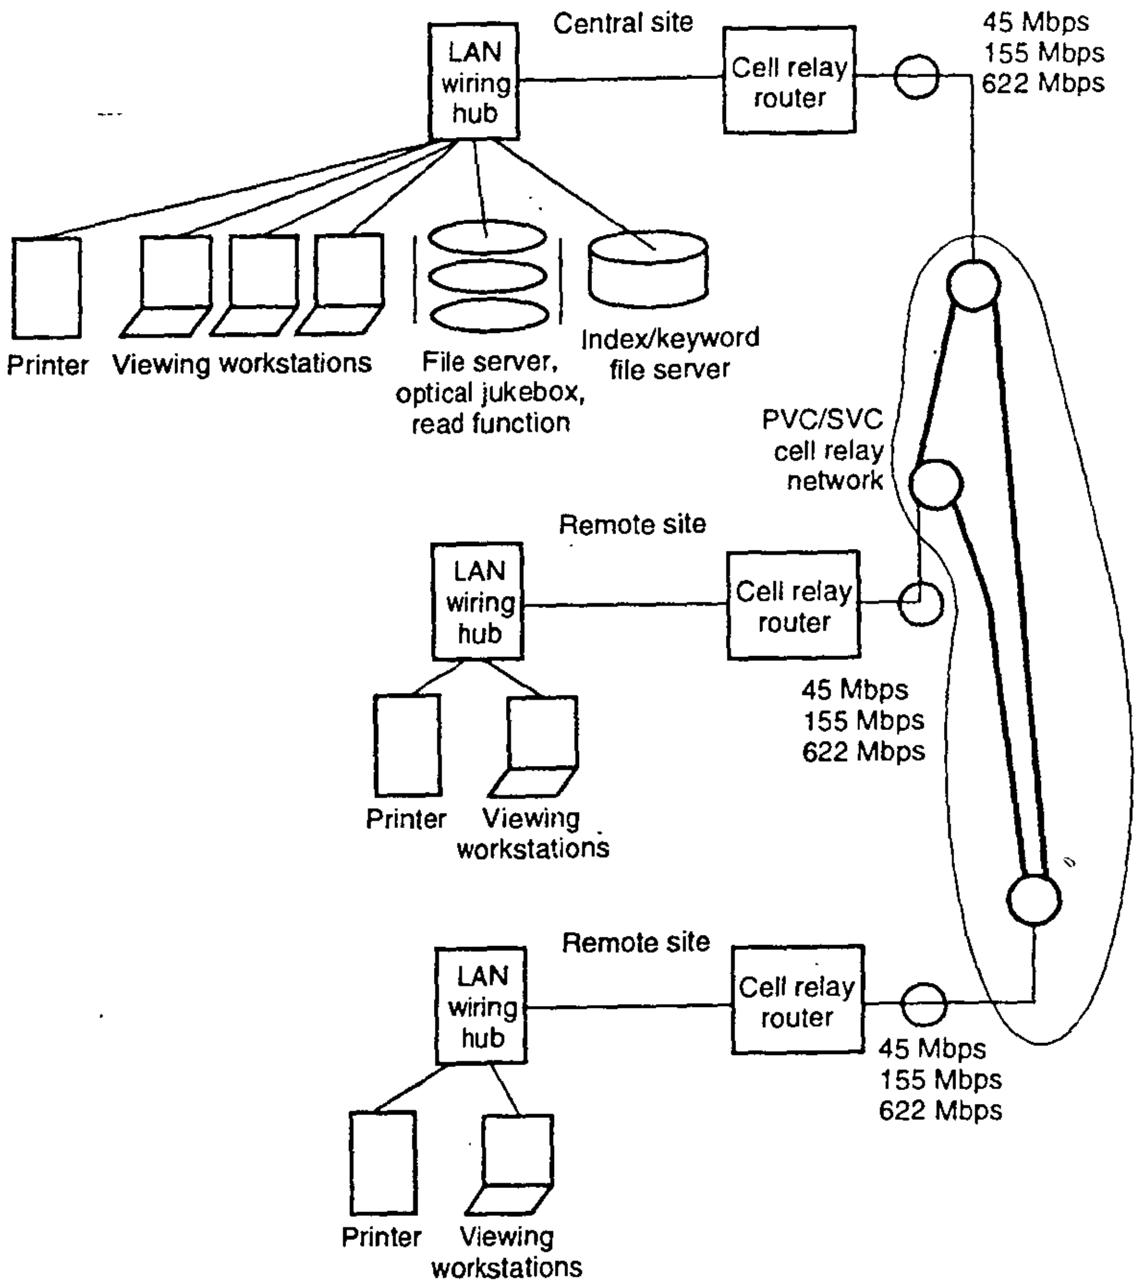

Figure 1.17 Use of ATM/cell relay to support imaging.

multimedia applications (the last two applications also have stringent delay sensitivities). A 100-MByte data unit across the application programming interface (API) running on a remotely located LAN-resident server is segmented into approximately 60 Ethernet frames. Each Ethernet frame is then segmented into approximately 30 cells by an ATM-configured router for delivery over a public cell relay network.

Some wish to clarify the implications of the interplay between the network(or private ATM switch) performance in terms of cell loss/mutilation, response time, latency, and the end-to-end error correction protocols (e.g., included in TCP). For example, if one of the 29 cells that

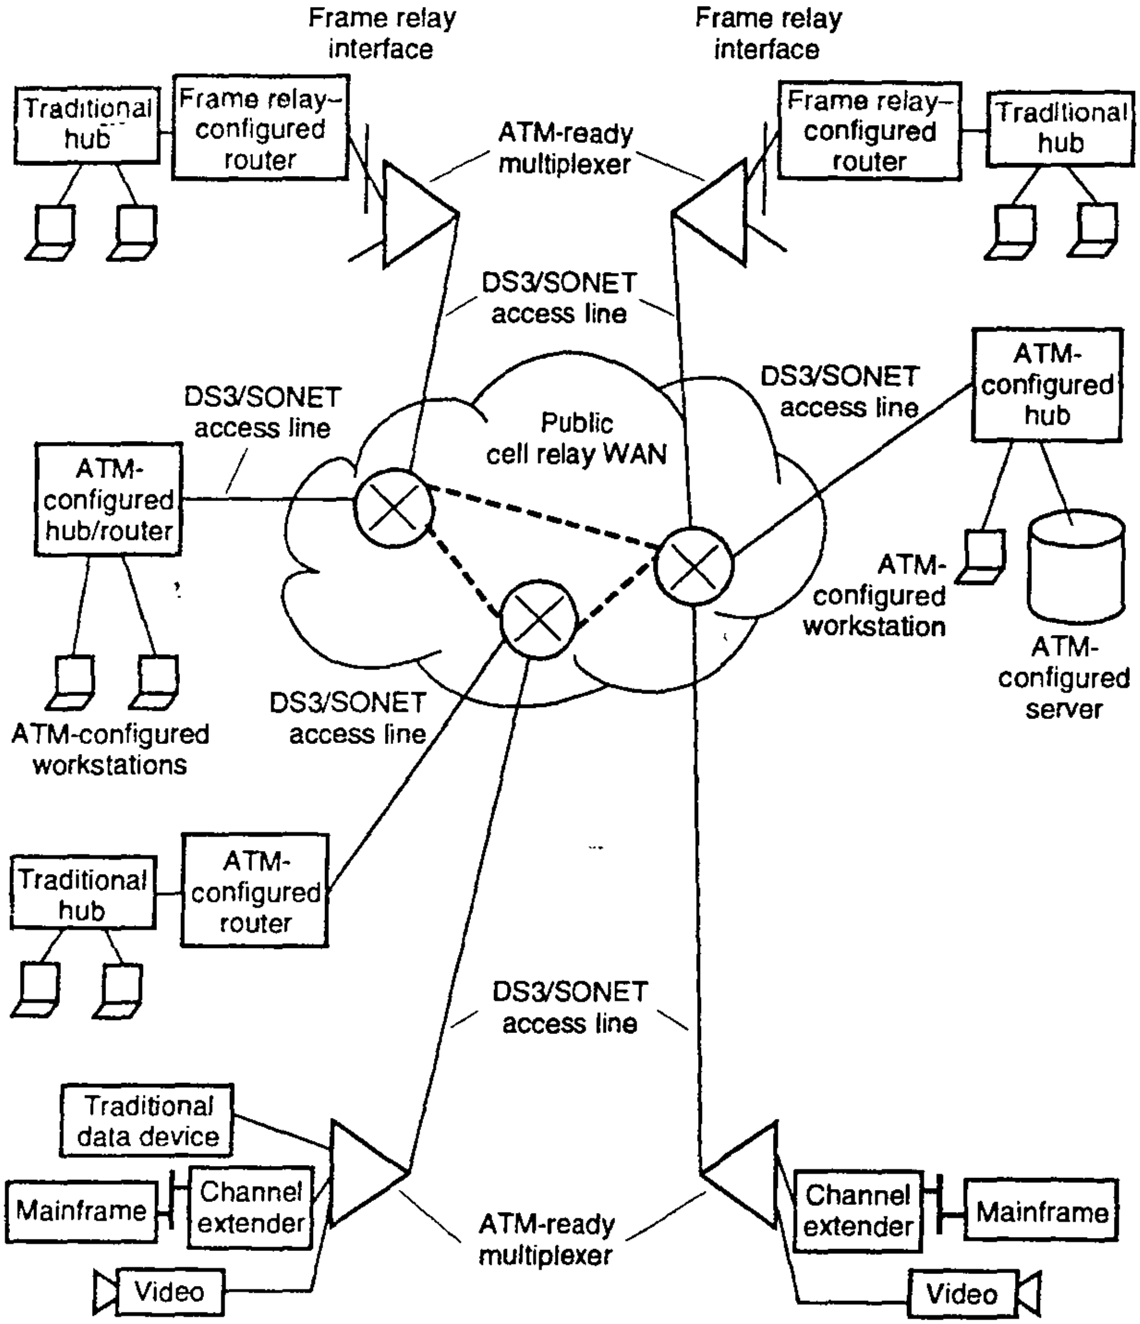

Figure 1.18 CRS to support enterprise networking in WAN applications.

made up a frame is lost, the entire frame (30 cells) needs to be retransmitted by TCP. Under heavy user load as well as coterminous ATM switch overload (whether public or private), the combination of client/server architecture and ATM communication could result in degradation, saturation, or instability. A number of simulation-based studies have shown that, when properly engineered, the network should behave as expected.

Chapter 9 covers ATM-based LANs, while Chaps. 11 and 12 cover other details pertaining to the deployment of ATM in users' environments.

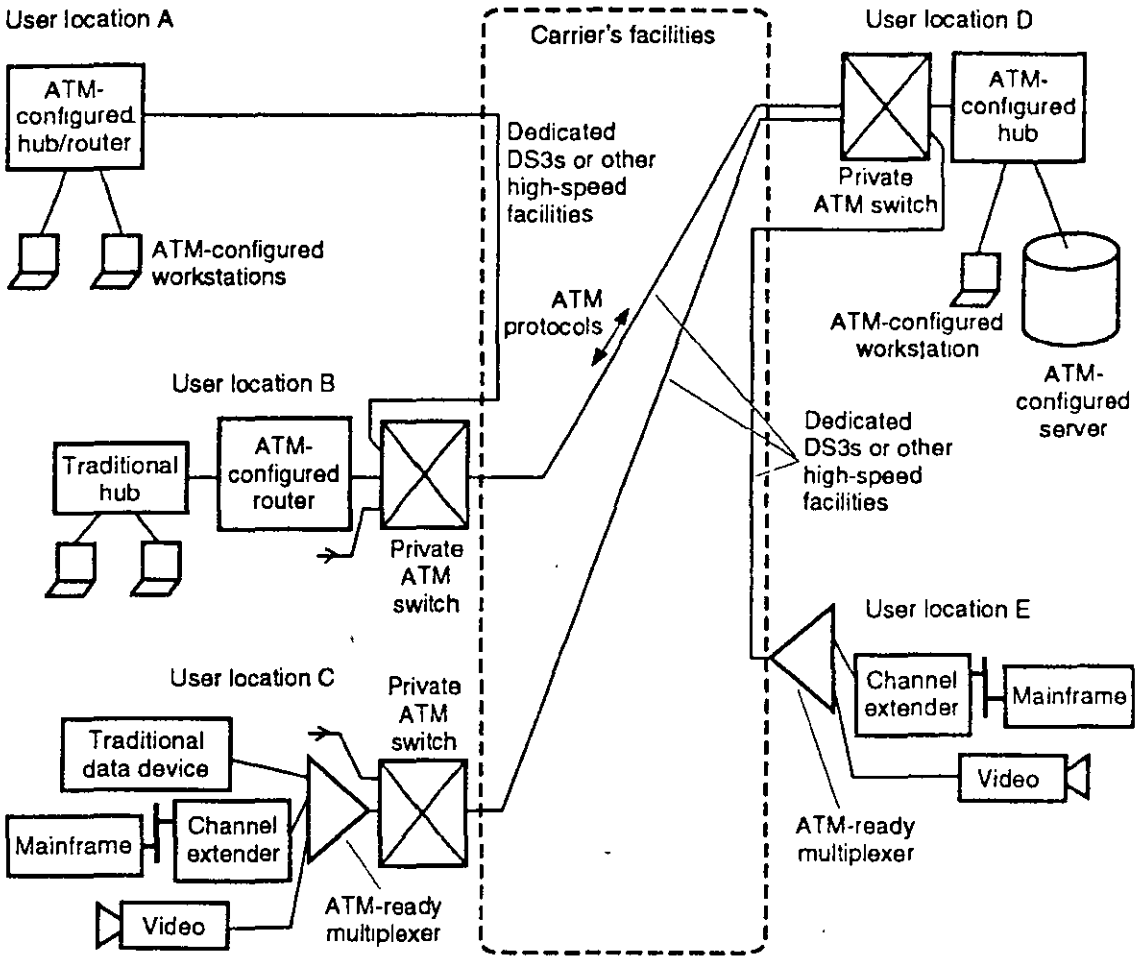

Figure 1.19 CRS to support enterprise networking in WAN applications (private network).

1.7 The Value of Standards

It is a well-known fact that standards benefit not only an industry but an entire economy. Many industries would not have arisen (e.g., the VCR industry, the CD audio industry, television, radio, etc.) if it were not for standards. Standards make a level playing field, fostering competition; this is in contrast to vendor proprietary approaches, where only those vendors have access to a market or have disproportionate control of it. However, for a standard to be effective, it must be widely available, without restrictions on promulgation, discussion, commentary, proliferation, distribution, and duplication. In our opinion, a standard is not an open standard if it is restricted, copyrighted, or patented, if it represents someone's intellectual property, or if it is "owned" by someone (sounds mighty close to a proprietary system to us!) because all of these factors frustrate the exact purpose for which the standard aims to exist (or has a reason to exist). There is much discussion at large about "free trade," "free movement of information," and "lack of censorship."

Standards are developed by industry consensus. This means that representatives from many companies, typically several dozen, have input into the standard. These proceedings can go on for years, and the representatives of these companies travel to many meetings and invest company resources back home to work on technical issues, prepare contributions, review contributions, and act as editors, chairs, etc. In the end, no one individual or institution should be able to claim ownership. There must be a free flow of specification information. Developers must be able to obtain copies. Programmers must be able to use the material. Documentaries must be able to write down the standard and comment on how they implemented various aspects. Educators must be able to discuss the standard and promulgate it to users. Otherwise, such a standard may go nowhere, as many examples of voluminous standards from the (late) 1980s illustrate.

Given this philosophical imperative, and in spite of the less than eloquent case made in these terse paragraphs, we have taken the approach of discussing here, in this text, the dozens of standards that support cell relay service and ATM, regardless of their source. In the end, all stand to benefit from such oper and uninhibited discussion at the birth of this new technology. Since this book is only a brief synopsis of the estimated 15 cubic feet of standards material that forms the basis for ATM (ITU-T, ANSI T1S1, ATM Forum, Frame Relay Forum, Bellcore, and other documents), the reader is constantly referred to the original documents for the full-scale detail. In particular, developers, who stand to benefit commercially from their efforts, should definitely refer to the original documentation for the necessary level of detail. The purpose of this book is strictly pedagogical and for the end user. Each of the more than 100 documents alluded to earlier can be obtained from the original source for $100 or less.

References

1. D. Minoli, *Enterprise Networking—Fractional T1 to SONET, Frame Relay to BISDN*, Artech House, Norwood, Mass., 1993.
2. D. Minoli, *1st, 2nd, and Next Generation LANs*, McGraw-Hill, New York, 1994.
3. D. Minoli, *Imaging in Corporate Environments: Technology and Communication*, McGraw-Hill, New York, 1994.
4. D. Minoli and B. Keinath, *Distributed Multimedia: Through Broadband Communication Services*, Artech House, Norwood, Mass., 1994.
5. D. Minoli et al., *ATM Layer Bearer Service/Cell Relay Service Extended Stage 1 Description for Public Service Offerings*, T1S1.5/93-021, February 1993.

6. T1S1.5/93-52, *Broadband Aspects of ISDN Baseline Document, T1S1 Technical Subcommittee*, August 1990, Chief Editor: Erwin Fandrich.
7. D. Minoli, "The New Wide Area Technologies: SMDS and B-ISDN," *Network Computing*, pp. 88ff., August 1991.
8. D. Minoli, "Understanding ATM—Part 1," *Network Computing*, pp. 128ff., Oct. 15, 1992.
9. D. Minoli, "Understanding ATM—Part 2," *Network Computing*, pp. 156ff., Nov. 15, 1992.
10. D. Minoli, "Third-Generation LANs," *UNIX Expo 92 Proceedings*, Bruno Blemheim Inc., Fort Lee, N.J., 1992.
11. D. Minoli, "Third Generation LANs," *Proceedings of Texpro 1993*, Pacific Bell, San Francisco, April 1993.
12. D. Minoli, "Cell Relay and ATM," WAN Insert to *Network Computing* and *Communications Week*, pp. 22 ff., August 1993.
13. D. Minoli, "Wide Area Networking for Multimedia?", WAN Insert to *Network Computing* and *Communications Week*, pp. 60ff., August 1993.
14. D. Minoli, "Broadband Integrated Services Digital Network," Datapro Communications Series: *Broadband Networking*, Report #2890, April 1992.
15. D. Minoli, "ATM and Cell Relay Concepts," Datapro Communications Series: *Broadband Networking*, Report #2880, April 1992.
16. J. T. Johnson, "Applications Catch Up to ATM," *Data Communications*, pp. 41–42, July 1993.
17. F. Gratzer and S. Walters, "ATM and Fast Packet Services—Perfect Together," *Bellcore DIGEST*, vol. 10 (6), pp. 3ff., 1993.
18. D. Minoli, "Designing Scalable Networks," *Network World Collaboration*, pp. 17ff., January 10, 1994.
19. D. Minoli, *Analyzing Outsourcing: Reengineering Information and Communication Systems*, McGraw-Hill, New York, 1995.

Asynchronous Transfer Mode

As noted in Chap. 1, ATM is a new transport and switching technology that can be used in a variety of telecommunications and computing environments. ATM is a cell-based technology, designed to support user applications requiring high-bandwidth, high-performance transport and switching. This chapter provides a summary description of the peer-to-peer ATM protocol at the user-network interface in support of cell relay service and other ATM capabilities. It describes functionality in the User Plane, thereby enabling a PVC service. The addition of Control Plane support enables the user to obtain an SVC service; the operation of the ATM Layer in the Control Plane is nearly identical to that of the User Plane (the Control Plane functionality is discussed in Chap. 4). Some aspects of the underlying transport mechanism are also briefly covered at the end of the chapter.

A description of general aspects of the access interface(s) between the user and the network is followed by a description of the protocol across such an interface. The protocols and related requirements are associated with two functional OSIRM layers: the Data Link Layer and the Physical Layer. Figure 2.1 depicts this peer-to-peer protocol view of the service. Figure 2.2 depicts communication through a set of network peers. As described in ITU-T Recommendation X.210, *Open Systems Interconnection, Layer Service Definition Conventions*,[1] the service defined at the Data Link Layer also relies on the capabilities of the Physical Layer. This view of cell relay service in general and of the ATM protocol in particular establishes requirements on what an entity in the ATM Layer (whether the entity is in the network or in the user's equipment),

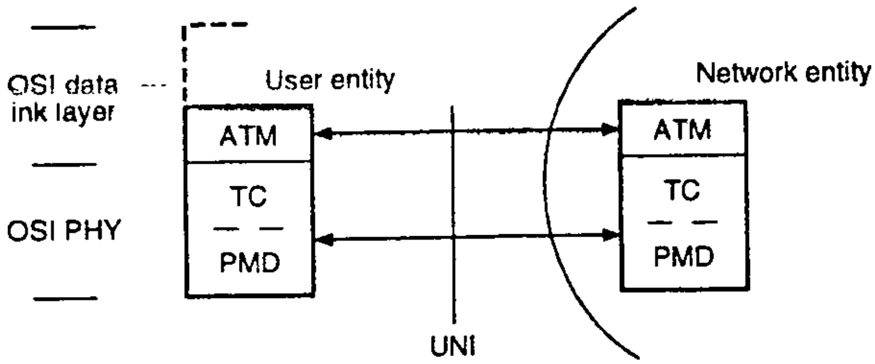

Figure 2.1 Peer entities across the user-network interface. TC = Transmission Convergence Sublayer; PMD = Physical Medium–Dependent Sublayer.

defined in ITU-T Recommendation I.361, *B-ISDN ATM Layer Specification*,[2] and in T1S1.5/92-410, *Broadband ISDN—ATM Layer Functionality and Specification*,[3] expects the remote peer entity to support. The physical aspects of the UNI supporting cell relay service are based on the B-ISDN UNI defined in ITU-T Recommendation I.432, *B-ISDN User-Network Interface—Physical Layer Specification*[4] and on the ATM Forum's *UNI Specification*[5] for public UNIs. This discussion only provides an overview; the reader interested in additional details should consult Refs. 6 and 7.

This chapter only covers the interface between user equipment and a public network; intra-CPE interfaces (for example, for ATM-based LANs), although similar in many respects to the interface between the CPE and the network, are not addressed. Table 2.1 depicts some of the key ITU-T standards in support of ATM in general and the peer-to-peer cell relay protocol in particular.

2.1 Access Interface

This section defines the concept of access interface. This is accomplished by defining an access reference configuration, functional entities (groups), and logical reference points.

An *access reference configuration* for B-ISDN is defined in ITU-T Recommendation I.413, *B-ISDN User-Network Interface*.[8] This configu-

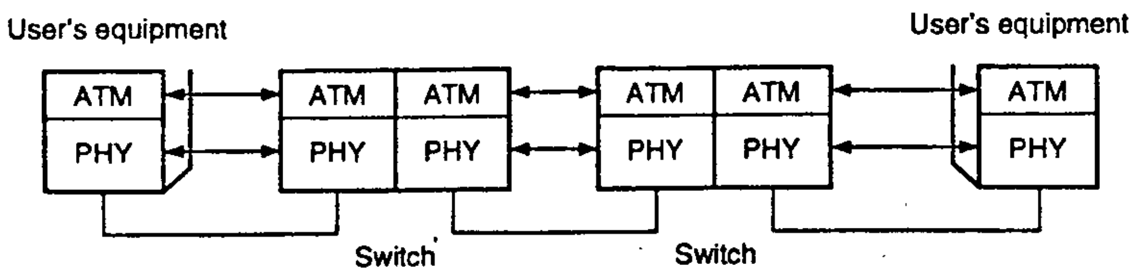

Figure 2.2 Cascaded ATM entities.

TABLE 2.1 Key ITU-T Standards in Support of ATM

F.811	B-ISDN Connection-Oriented Bearer Service
F.812	B-ISDN Connectionless Bearer Service
I.113	B-ISDN Vocabulary of Terms
I.121R	Broadband Aspects of ISDN [Basic Principles and Evolution]
I.150	B-ISDN ATM Functional Characteristics
I.211	B-ISDN Service Aspects
I.311	B-ISDN General Network Aspects
I.321	B-ISDN Protocol Reference Model and Its Applications
I.327	B-ISDN Functional Architecture Aspects
I.356	Quality of Service Configuration and Principles
I.361	B-ISDN ATM Layer Specification
I.362	B-ISDN AAL Functional Description
I.363	B-ISDN AAL Specification
I.371	Traffic Control and Resource Management
I.374	Network Capabilities to Support Multimedia
I.413	B-ISDN UNI
I.432	B-ISDN UNI Physical
I.555	Interworking with Frame Relay
I.555	Interworking with ISDN
I.610	B-ISDN OAM Principles
I.cls	Support for Connectionless Data Service on B-ISDN
Q.93B (now Q.2931)	B-ISDN Call Control
Q.SAAL 1 and 2 (now Q.2110 and Q.2130)	Signaling AALs [Q.2110, Service-Specific Connection-Oriented Protocol (SSCOP); Q.2130, Service-Specific Co-ordination Function (SSCF)]

ration forms the basis for the definition of access interfaces supporting cell relay service.

Functional entities are logical abstractions of functions typically found in network equipment and in users' equipment, also known as customer premises equipment (CPE). Public network switch-termination functions are modeled by the broadband line terminator/exchange terminator (B-LT/ET) functional group. The CPE is modeled by the broadband network termination 2 (B-NT2) functional group; NT2 functions include concentration, switching, and resource management. Broadband network termination 1 (B-NT1) functions support line termination, line maintenance, and performance monitoring. The broadband terminal

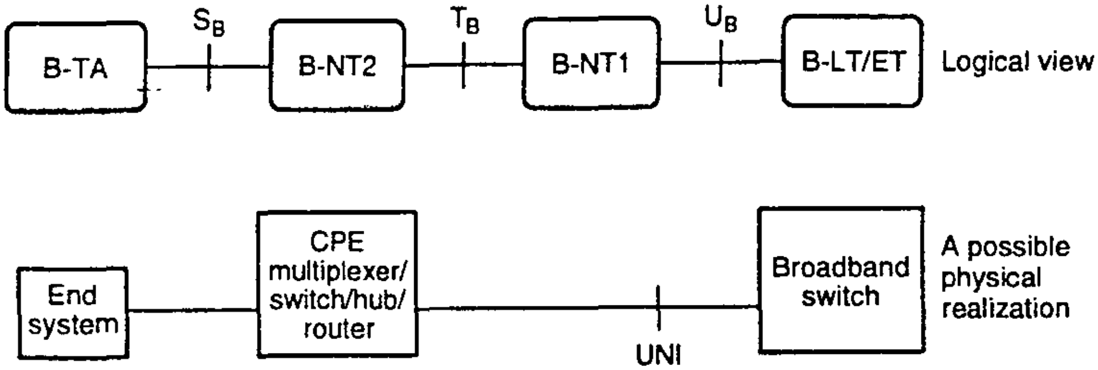

Figure 2.3 B-ISDN access reference configurations. B-TA = broadband terminal adapter.

equipment, such as a workstation, is modeled by the broadband terminal equipment (B-TE) functional group.

Logical reference points are defined between B-ISDN functional entities. *TB* is the logical reference point between a B-NT2 and a B-NT1. UB is the logical reference point between a B-NT1 and a B-LT/ET. In this description, the UNI is associated with the UB reference point. See Fig. 2.3.

Note: This description only covers the case where there is a single B-NT2 (however, several B-TEs may be connected to the B-NT2). The case where the B-NT2 is null and there are several B-TEs connected to a single UNI is not addressed in the initial view of ATM services in the United States.

2.2 ATM-Level Protocol

2.2.1 Overview

UNI protocols define the way in which users communicate with the public network for the purpose of accessing the service provided by the network. Figure 2.4 illustrates the B-ISDN Protocol Reference Model, which is the basis for the protocols that operate across the UNI (this is another common way to represent the protocol model of Fig. 1.3). The B-ISDN Protocol Reference Model is described in ITU-T Recommendation I.121. This model is made up of three planes, already discussed in Chap. 1: the User Plane, the Control Plane, and the Management Plane. Table 2.2 provides a summary of the functions supported by each plane.

The UNI specified at this level includes the functions associated with the User Plane at the Physical Layer and the ATM Layer. The Physical Layer provides access to the physical medium for the transport of ATM cells. It includes methods for mapping cells to the physical medium (i.e., the Transport Convergence Sublayer) and methods dependent on the physical medium (i.e., the Physical Medium–Dependent Sublayer). The ATM layer provides for the transport of cells between end-user locations. An ATM cell contains a header that contains control information, iden-

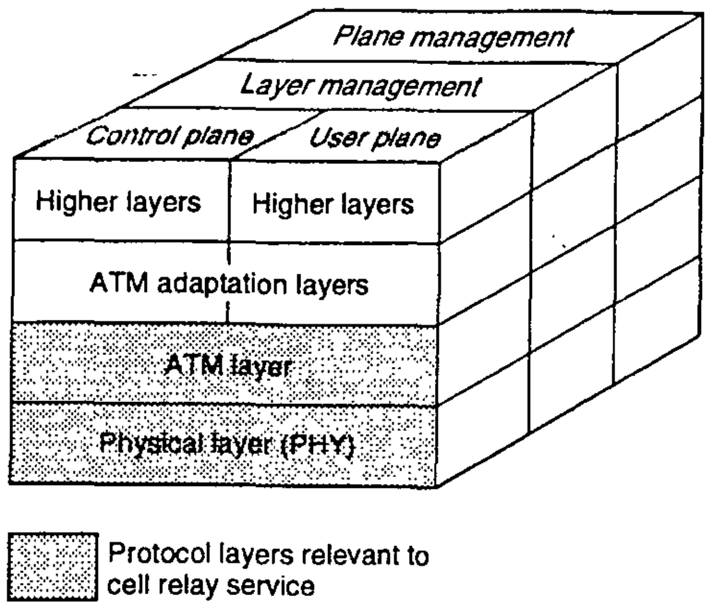

Figure 2.4 B-ISDN protocol reference model.

tifies the type of cell, and contains routing information that identifies a logical channel (i.e., a VPC or a VCC) over which the cell is to be forwarded.

The interactions of each protocol layer with other layers and with its own layer management are described in terms of primitives. Primitives describe abstractly the logical exchange of information and control

TABLE 2.2 Functions of Various Planes of the Protocol Model

User Plane	Provides for the transfer of end-user information. It consists of the Physical Layer and the ATM Layer. The model also includes ATM Adaptation Layers and higher layers necessary for each end-user application. (Because these layers are specific to each application, they are not part of the cell relay service described here and in Chap. 5.)
Control Plane	Provides for the transfer of information to support connection establishment and control functions necessary for providing switched services. The Control Plane shares the ATM and Physical Layer with the User Plane. Also, it contains AAL procedures and higher-layer signaling protocols. The Control Plane is discussed in Chap. 5.
Management Plane	Provides for operations and management functions and the capability to exchange information between the User and the Control Planes. The Management Plane is made up of the *Layer Management* (for layer-specific management functions such as detection of failures and protocol abnormalities) and the *Plane Management* (for management and coordination functions related to the complete system). The Management Plane is discussed in Chap. 10.

through a service access point, while not imposing any constraint on the implementation. Figures 2.5, 2.6, and 2.7 depict some aspects of this protocol machinery.

2.2.2 ATM Layer

The ATM Layer provides for the transport of fixed-size cells between end-user locations. It is implemented in users' equipment (workstations, routers, private switches, etc.) and in network equipment. ATM cells from end users are forwarded across virtual connections through the public network. These connections are provided at subscription time or in real time via signaling (as described in Chap. 4). The ATM Layer also provides multiplexing functions to allow the establishment of multiple connections across a single UNI.

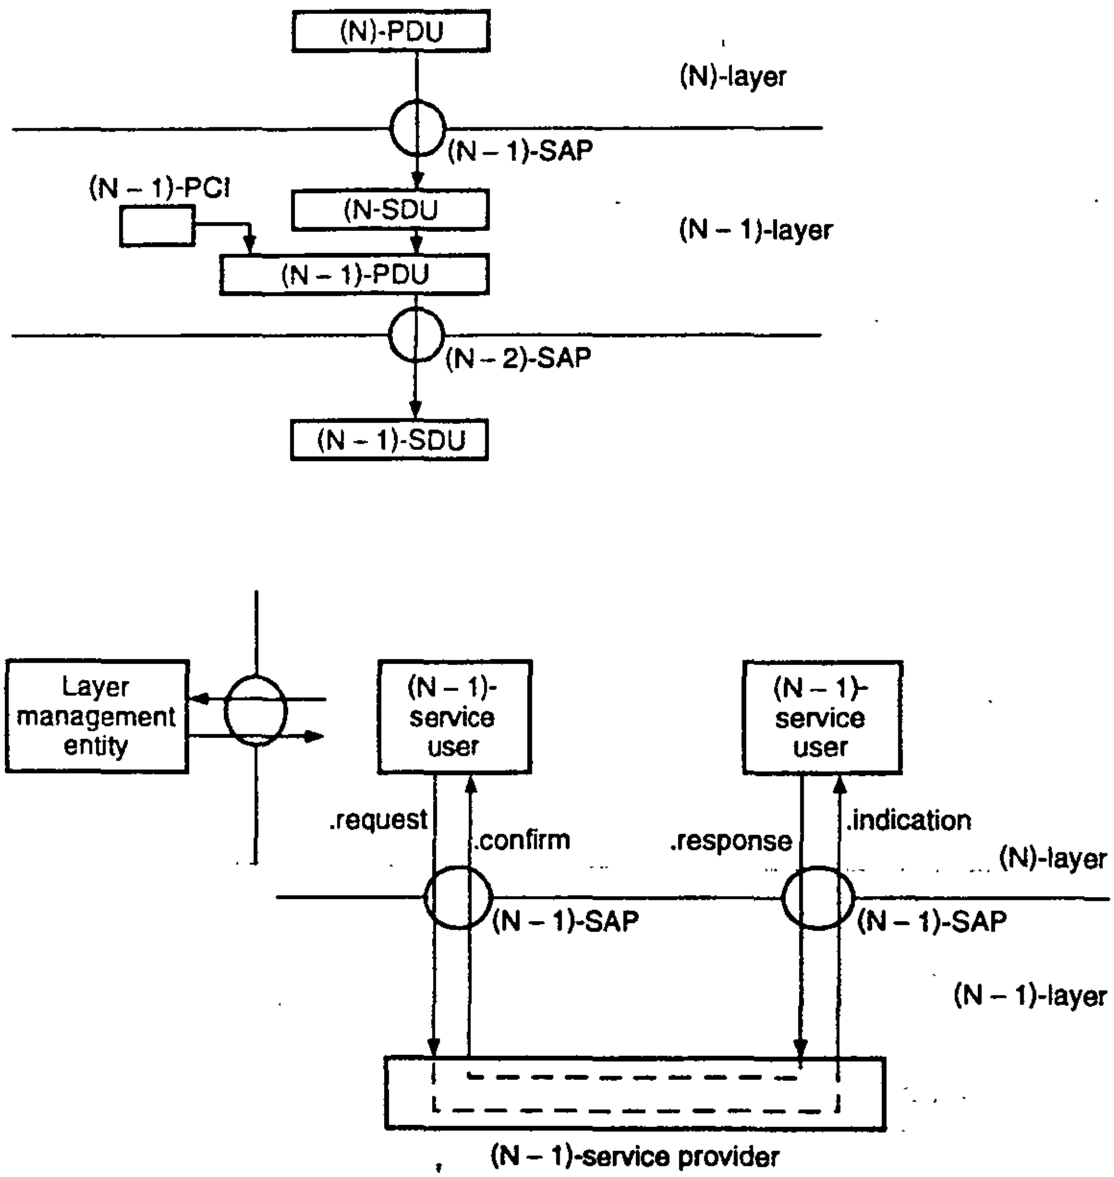

Figure 2.5 SAPs (top) and primitives (bottom). SAP = service access point; PDU = protocol data unit; SDU = service data unit; PCI = protocol control information.

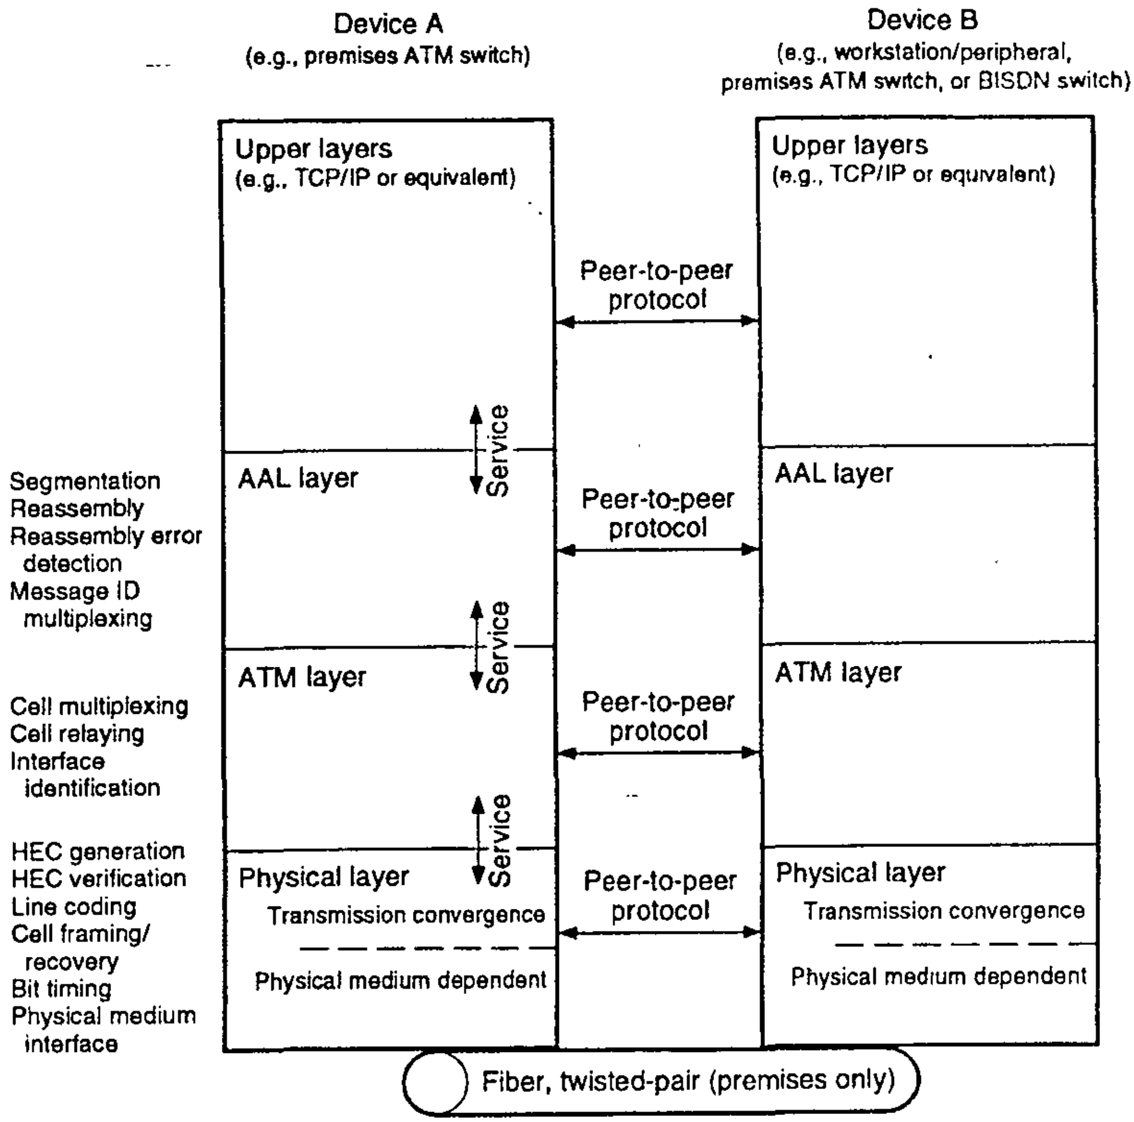

Figure 2.6 ATM protocols.

Service provided to the upper layer. The ATM-Layer service is based on fixed-size ATM service data units which consist of 48 octets. It provides for the transparent transfer of ATM SDUs between communicating peer upper-layer entities. To accomplish this, the ATM Layer generates a 53-octet ATM cell by prepending a 5-octet header to the ATM SDU. The header contains routing and protocol control information. The interaction between the ATM Layer and its service users is implemented by the primitives shown in Table 2.3.

Service expected from the lower layer. The ATM Layer expects the Physical Layer to support the transparent transport of ATM cells between peer ATM entities. The exchange of information between the ATM Layer and the Physical Layer is implemented by the primitives shown in Table 2.4. The PHY-SDU parameter in these primitives contains the 53-octet cell to be transmitted between peer ATM entities.

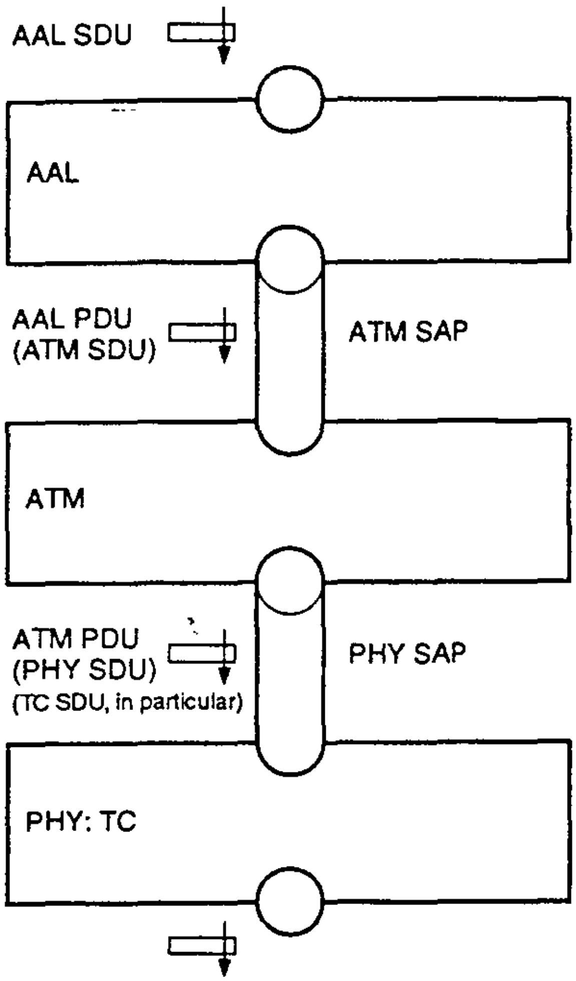

Figure 2.7 Pertinent ATM SAPs.

ATM cell format. The ATM cell format used across the UNI is shown in Fig. 2.8 (which is another way of looking at Fig. 1.2). Table 2.5 describes the meaning of the fields.

End-to-end operations administration and maintenance capabilities need to be supported. For VPs, operation functions are supported via specially marked ATM cells, which are transmitted over VCs with specific VCI values* (these are known as F4 flows). For VCs, operation functions are supported via cells marked with an appropriate codepoint in the Payload Type Indicator field (these are known as F5 flows). The functions supported are shown in Table 2.6. Figure 2.9 illustrates the difference between these two OAM flows.

Table 2.7 provides the encoding for the PTI field. Code point 100_B (B = binary) indicates a segment OAM F4 cell flow used to monitor the status of a segment within the virtual connection. Code point 101_B indicates an end-to-end OAM F5 cell flow used to monitor the status of a connection end to end. Code point 110_B is reserved for future traffic control and resource management procedures.

*VCI is 4 for end-to-end operations and 3 for segment information.

TABLE 2.3 ATM Layer Primitives

ATM-DATA.request (ATM_SDU, Submitted_Loss_Priority, Congestion_Indication, SDU_Type)	Used to request transmission of an ATM SDU across a VPC or VCC to a peer entity
ATM-DATA.indication (ATM_SDU, Received_Loss_Priority, Congestion_Indication, SDU_Type)	Used by the ATM Layer to indicate to the service user the arrival of an ATM cell

Description of parameters:

ATM_SDU: The 48 octets of information to be transferred by the ATM Layer between peer communicating upper-layer entities.

Submitted_Loss_Priority: The relative importance of the ATM_SDU contained in this primitive. Two values are possible. A value of "high" indicates that the resulting ATM cell has higher (or equivalent) loss priority than a cell with a value of "low." A high value may be translated to a cell loss priority value of 0 in the cell header. Similarly, a low value may be translated to a CLP value of 1 in the cell header.

Congestion_Indication: This parameter indicates whether this cell has passed through one or more network nodes experiencing congestion. It has two values: True or False.

SDU_Type: This parameter indicates the type of SDU to be transferred between peer upper layer entities. It can take only two values, 0 and 1, and its use is as determined by the higher layer. For example, AAL Type 5 sets SDU_Type to 1 to indicate the last cell of a frame. In other words, this field is currently used by the AAL Type 5 Common Part protocol to distinguish between cells that contain the last segment of an AAL Type 5 Common Part PDU and those that do not. AAL Type 1 and AAL Type 3/4 always set the bit to 0.

Received_Loss_Priority: This parameter indicates the CLP field marking of the received ATM_PDU. Two values are possible. A value of "high" indicates that the received ATM cell has higher (or equivalent) loss priority than a cell with a value of "low." A high value may be translated to a cell loss priority value of 0 in the cell header. Similarly, a low value may be translated to a CLP value of 1 in the cell header.

ATM Layer procedures. This section summarizes the functions performed by ATM layer entities.

ATM sending procedures. These procedures are performed by an ATM entity to send ATM cells to a peer ATM entity. The procedures are organized according to the categories of functions performed by the ATM Layer.

ATM layer connections. As described earlier, the ATM service is provided by means of virtual connections. For the PVC cell relay service, connections are established at subscription time. For SVC service,

TABLE 2.4 Physical Layer Primitives

PHY-DATA.request (PHY_SDU)	Requests the Physical Layer to transport an ATM cell between peer ATM entities over an existing connection.
PHY-DATA.indication (PHY_SDU)	Indicates to the ATM Layer that an ATM cell has been received over an existing connection.

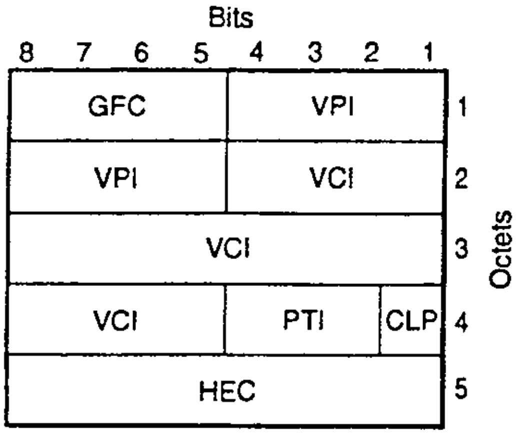

Figure 2.8 ATM cell format.

connections are established by a signaling mechanism. As will be seen in Chap. 4, about one dozen parameters need to be specified to describe a connection (for example, called party, bandwidth, quality of service, etc.).

Cell rate decoupling. A sending ATM entity must add unassigned cells to the assigned cell stream to be transmitted, so that a continuous cell stream matching the line rate of the UNI is provided to the Physical Layer. This is necessary in order for the Physical Layer to perform adequate cell delineation functions. Unassigned cells are empty cells which have the first 4 octets of the cell header encoded as depicted in Fig. 2.10. Unassigned cells do not carry information. Therefore, they must be extracted at the receiving ATM entity and not passed to the upper layer.

Loss priority indication. Traffic management functions may use tagging as a way to control traffic entering the network across the UNI. The network may choose to tag cells that violate a traffic descriptor for the connection by setting the CLP bit to 1. If cell discarding is necessary, these cells would be discarded first. Some traffic management procedures are discussed in Chap. 6.

ATM receiving procedures. This section describes the procedures an ATM entity executes when receiving an ATM cell to ensure its proper processing. These procedures include the provision for sequenced processing of ATM cells which arrive across a virtual connection.

Sequenced ATM processing. ATM cells received across a virtual connection must be processed in sequence to ensure adequate service to the higher layers.

Cell validation procedures. The cell validation procedures determine whether a received cell is an unassigned cell and detect invalid header patterns. These procedures also detect cells received with inactive VPI/VCI values (e.g., VPI/VCI values which identify inactive connections). Unassigned cells and cells found to be in error are discarded.

TABLE 2.5 ATM Cell Fields

Generic Flow Control (GFC)	The 4-bit GFC field has only local significance and may be used to provide standardized local functions at the customer site (e.g., passive bus support); the field is ignored and may be overwritten by the public network.
Virtual Path Identifier/Virtual Channel Identifier	The 24-bit VPI/VCI field indicates the virtual connection over which a cell is to be forwarded. The number of connections needed across the UNI is less than 2^{24}, therefore, only some bits of the VPI and VCI subfields are used. Those bits are called *allocated bits*, and all other bits in the VPI/VCI field are set to 0. A VPI value of 0 is not available for user-to-user virtual path identification. Similarly, a VCI value of 0 is not available for user-to-user virtual channel identification.
Payload Type Indicator (PTI)	The 3-bit PTI field indicates whether the cell contains user information or layer management information. Code points 000 to 011 indicate user information; these PTI values identify two types of end-user information and whether the cell has experienced congestion (the two types of information are used by the end-user application). For user data, the public network does not change the SDU_Type indicated by the PTI field. The public network can, however, change the PTI value from Congestion_Experienced = False to Congestion_Experienced = True. Code points 100 to 111 identify different types of operations flows. See Table 2.7.
Cell Loss Priority	This 1-bit field allows the user to indicate the relative cell loss priority of the cell. The network may attempt to provide a higher cell loss priority (or equivalent) for cells marked with high priority than for cells marked with low priority. The current view is to only let the user set CLP to the value 0.
Header Error Control	The 8-bit HEC field is used by the Physical Layer to detect transmission errors in the cell header and in some cases for cell delineation.

Cell discrimination based on PTI value. A receiving ATM Layer entity processes cells according to the type of payload they contain as indicated by the value in the PTI field. User cells (PTI values 000–100) are forwarded across the appropriate virtual channel. If neces-

TABLE 2.6 Layer Management Functions Included in Cell Relay Service

Fault management functions	Alarm surveillance: AIS (alarm indication signal)
	Alarm surveillance: FERF (far-end receive failure; now known as remote defect indicator)
	Connectivity verification: cell loopback continuity check
Performance management functions	Forward monitoring Backward reporting Monitoring/reporting
Activation/deactivation	Performance monitoring Continuity check

sary, PTI values may be modified to indicate whether the cell experienced congestion.

Layer Management cells (PTI values of 101–111) are used to provide various operations flows to support functions like performance monitoring and trouble sectionalization. CPE supporting the UNI is not required to support these operations flows. However, network equipment must support them so that it can interface with end-user equipment supporting these functions. (This topic is revisited in Chap. 10.)

2.2.3 Layer Management

There are two types of interactions between the ATM entity and the ATM Management entity. One interaction is for the exchange of local information between these two entities. The primitives are shown in Table 2.8 (the parameters are not shown for simplicity). The other interaction is for peer-to-peer communication between ATM Management entities. The primitives for this interaction are shown in Table 2.9. For more details, refer to Ref. 2, 5, or 6. (This topic is revisited in Chap. 10.)

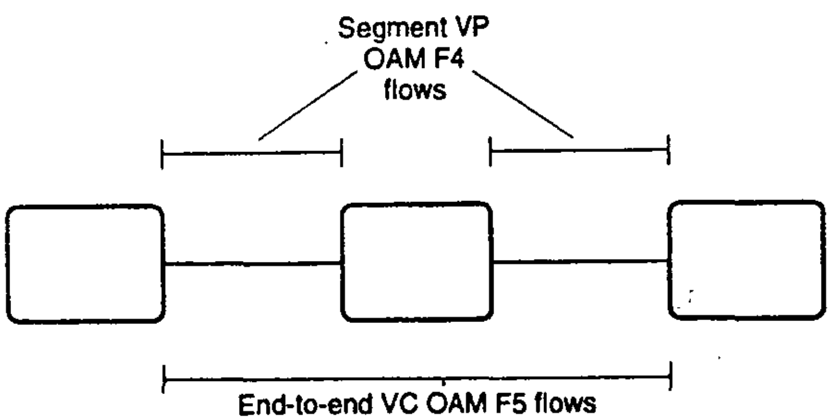

Figure 2.9 OAM F4 and F5 flows.

TABLE 2.7 PTI Code Points

PTI code-point	Meaning
000	User data—SDU_Type 0, no congestion experienced
001	User data—SDU_Type 1, no congestion experienced
010	User data—SDU_Type 0, congestion experienced
011	User data—SDU_Type 1, congestion experienced
100	Segment OAM F5 flow cell
101	End-to-end OAM F5 flow cell
110	Reserved for future traffic control and resource management functions
111	Reserved for future use

2.2.4 Physical Layer

Although the emphasis of this chapter is on the ATM Layer, a brief discussion of the underlying Physical Layer is also provided. Figure 2.11 depicts some of the key Physical Layer protocols supported.

As noted, the Physical Layer is made up of two sublayers: the Transmission Convergence Sublayer and the Physical Medium–Dependent Sublayer. The TC Sublayer "maps" the cell stream to the underlying framing mechanism of the physical transmission facility and generates the required protocol control information for the Physical Layer (e.g., SONET overhead octets). It also generates the HEC. The PMD Sublayer deals with the electrical or optical aspects of the physical interface (e.g., timing, power, jitter).

The UNI providing the service's access interface includes the physical characteristics of facilities that provide actual realizations of the U_B reference point. In practical terms, this access interface specifies the means and characteristics of the connection mechanism between CPE supporting cell relay service and a LEC's switch providing the same service. UNIs are specified by characteristics such as physical and electromagnetic/optical characteristics, channel structures and access

← 4 octets →

GFC	VPI	VCI	PTI	CLP
AAAA	0s	0s	XXX	0

A: This bit is available for use by appropriate ATM layer function.
X: This bit is a don't care bit.

Figure 2.10 First four octets of cell header for unassigned cells.

TABLE 2.8 ATM Management Primitives for Local Communication

ATMM-MONITOR.indication	Issued by an ATM Layer Management entity to deliver the content of an ATM_PDU received by the ATM entity, to facilitate an OAM function
ATMM-ASSIGN.request	Issued by an ATM Layer Management entity to request the establishment of an ATM link
ATMM-ASSIGN.confirm	Issued by an ATM Layer Management entity to confirm the establishment of an ATM link
ATMM-REMOVE.request	Issued by an ATM Layer Management entity to request the release of an ATM link
ATMM-REMOVE.confirm	Issued by an ATM Layer Management entity to confirm the release of an ATM link
ATMM-ERROR.indication	Issued by an ATM Layer Management entity to indicate an error and invoke appropriate management actions
ATMM-PARAMETER-CHANGE.request	Issued by an ATM Layer Management entity to request a change in a parameter of the ATM link

capabilities, user-network protocols, maintenance and operations characteristics, performance characteristics, and service characteristics.

The physical access channel for ATM-based fastpacket services such as cell relay service supports one of the following access rates: 622.080 Mbits/s (future); 155.520 Mbits/s; 44.736 Mbits/s; 1.544 Mbits/s (perhaps in the future). The corresponding channel signal formats are STS-12c (Synchronous Transport Signal Level 12, concatenated), STS-3c, DS3 (Digital Signal Level 3), and DS1.

Physical-Layer mappings. The mapping of cells onto the DS1, DS3, and SONET STS-3c has also been defined.[9] Some key aspects of how cells are inserted over the underlying framing mechanism are discussed below.

TABLE 2.9 ATM Management Peer-to-Peer Primitives

ATMM-DATA.request (ATM_SDU, Submitted_Loss_Priority, PHY_CEI(s))	Issued by an ATM Layer Management entity to request transfer of a management ATM_SDU
ATMM-DATA.indication (ATM_SDU, Received_Loss_Priority, PHY_CEI, Congestion_Indication)	Issued to an ATM Layer Management entity to indicate the arrival of a management ATM_SDU

Note: CEI is the connection endpoint identifier.

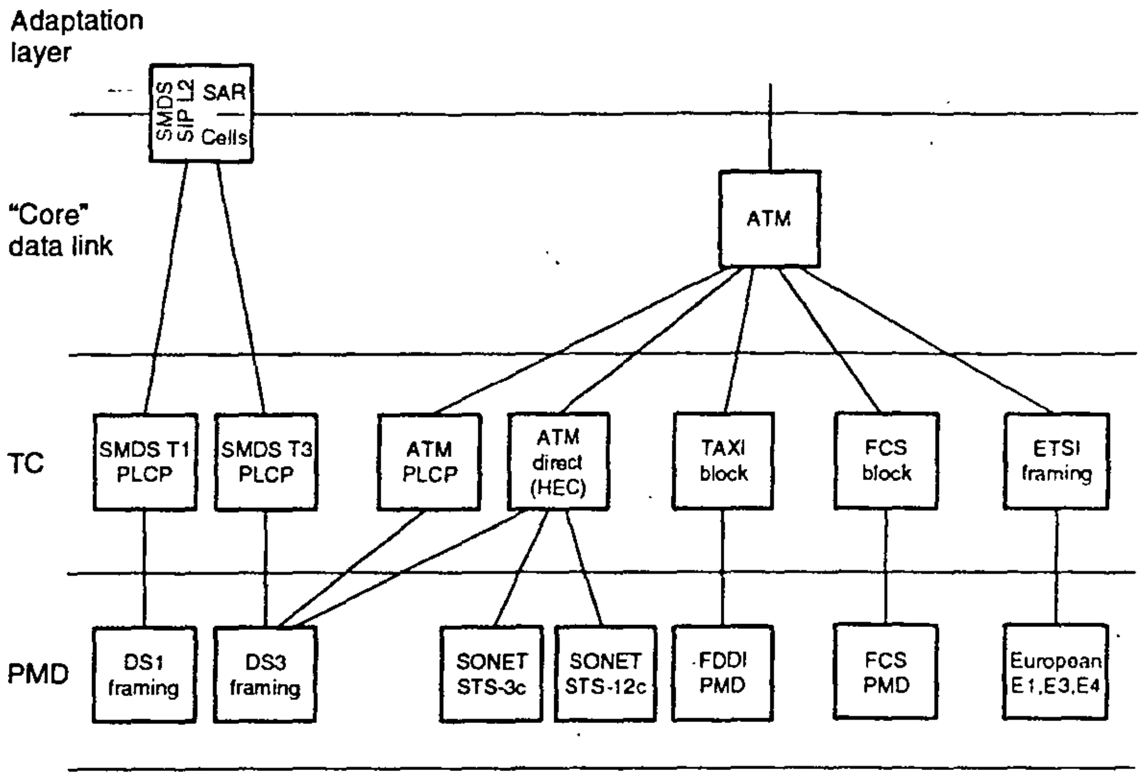

Figure 2.11 Key Physical Layer protocols supported. PLCP = Physical Layer convergence procedure; FCS = fiber channel standard; FDDI = fiber distributed data interface; ETSI = European Telecommunications Standards Institute.

The challenge at the receiving end is to extract the cell from the underlying frame, that is, to establish cell boundaries.

Mapping of ATM cells into 1544-kbit/s DS1 frame. Frame format. The multiframe structure for the 24-frame multiframe as described in ITU-T Recommendation G.704 is used. The ATM cell is mapped into bits 2 to 193 (i.e., time slots 1 to 24 described in Recommendation G.704) of the 1544-kbit/s* frame, with the octet structure of the cell aligned with the octet structure of the frame (however, the start of the cell can be at any octet in the DS1 payload; (see Fig. 2.12).

Cell rate adaption. The cell rate adaption to the payload capacity of the frames is performed by the insertion of idle cells, as described in ITU-T Recommendation I.432, when valid cells are not available from the ATM Layer.

Header error control generation. The Header Error Control value is generated and inserted in the specific field in compliance with ITU-T Recommendation I.432.

Scrambling of the ATM cell payload (optional). As an option, the ATM cell payload (48 bytes) can be scrambled before it is mapped into the 1544-kbit/s signal. In the reverse operation, following termination

*As of press time, however, standards for the delivery of ATM over a DS1 access were still being investigated.

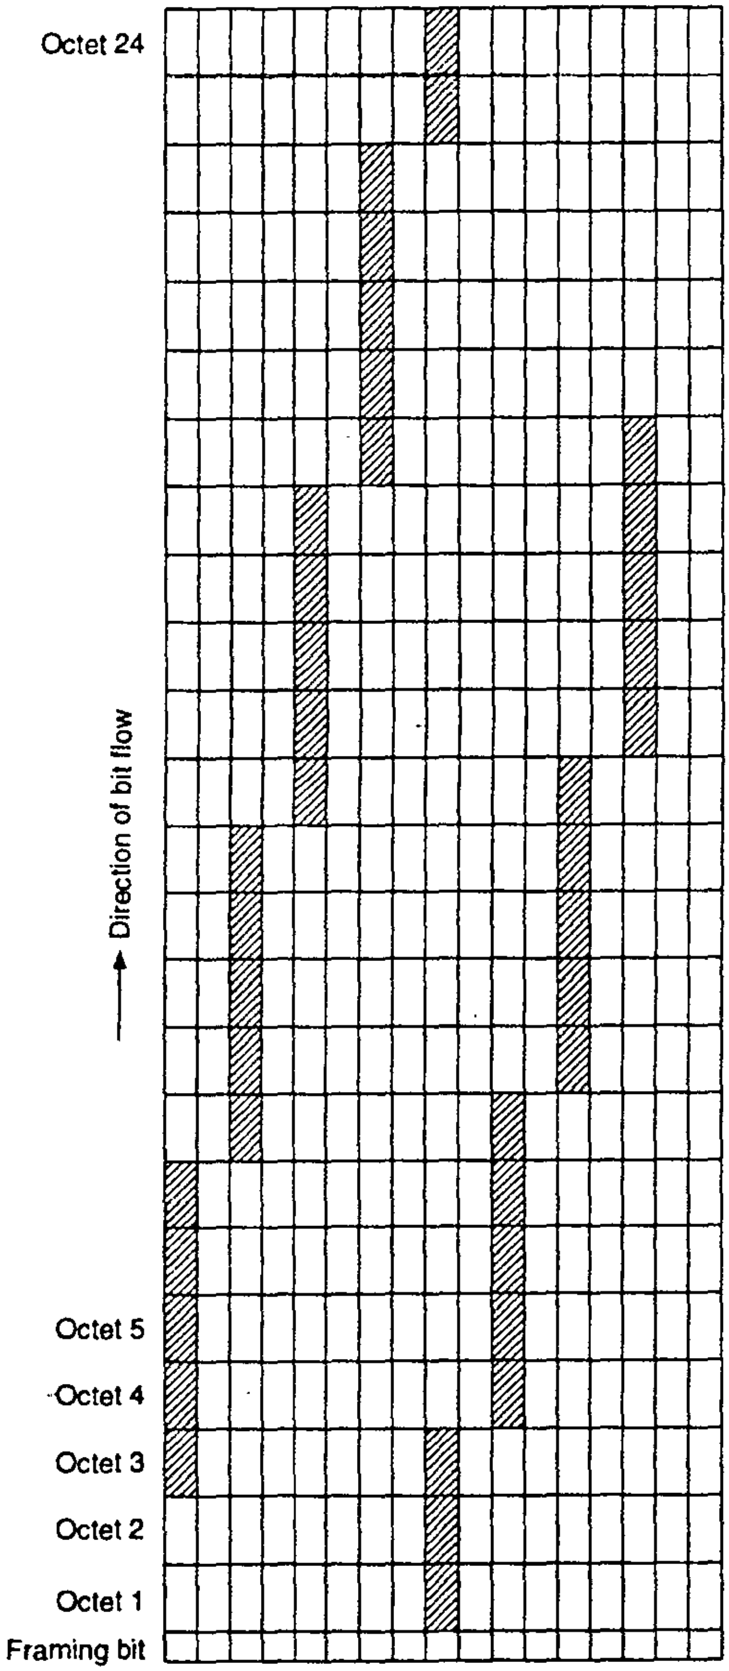

Figure 2.12 Direct mapping of cells onto DS1 frame (example).

of the 1544-kbit/s signal, the ATM cell payload is descrambled before being passed to the ATM Layer. The self-synchronizing scrambler with the generator polynomial $x^{43} + 1$ is used.

Cell delineation. Cell delineation is performed using the header error control mechanism as defined in ITU-T Recommendation I.432. This direct mapping approach means that the algorithm parses 5 octets on the fly until a 5-octet boundary is found through the HEC procedure. Once the header boundary is found, the rest of the cell boundary is established by counting 48 additional octets.

Cell header verification and extraction. The cell header verification is performed in compliance with ITU-T Recommendation I.432. Only valid cells are passed to the ATM Layer.

Mapping of ATM cells into 44,736-kbit/s DS3 frame

Frame format. The multiframe format at 44,736 kbits/s, as described in ITU-T Recommendation G.704, is used.

Two mappings are available:

1. Physical Layer Convergence Protocol (PLCP)–based mapping of ATM cells, derived from SMDS principles
2. A direct (HEC-based) mapping, established in 1993

This discussion focuses on PLCP, since the direct mapping is similar to the DS1 mapping.

The ATM PLCP defines a mapping of ATM cells onto existing 44,736-kbit/s facilities. The DS3 PLCP consists of a 125-μs frame within a standard 44,736-kbit/s payload. Note that there is no fixed relationship between the PLCP frame and the 44,736-kbit/s frame; i.e., the PLCP can begin anywhere inside the 44,736-kbit/s payload. The PLCP frame, Fig. 2.13, consists of 12 rows of ATM cells, each preceded by 4 octets of overhead. Nibble stuffing is required after the twelfth cell to fill the 125-μs PLCP frame. Although the PLCP is not aligned with the 44,736-kbit/s framing bits, the octets in the PLCP frame are nibble-aligned with the 44,736-kbit/s payload envelope. Nibbles begin after the control bits (F, X, P, C, or M) of the 44,736-kbit/s frame. The stuff bits are never used in the 44,736-kbits/s, i.e., the payload is always inserted. The reader interested in a detailed explanation of the DS3 framing format may refer to Ref. 10 or other material. Octets in the PLCP frame are described in the following sections.

Cell rate adaption. The cell rate adaption to the payload capacity of the PLCP frame is performed by the insertion of idle cells, as described in ITU-T Recommendation I.432, when no valid cells are available from the ATM Layer.

PLCP (1 octet)	Framing (1 octet)	POI (1 octet)	POH (1 octet)	PLCP payload (53 octets)	
A1	A2	P11	Z6	First ATM cell	
A1	A2	P10	Z5	Second ATM cell	
A1	A2	P09	Z4	Third ATM cell	
A1	A2	P08	Z3		
A1	A2	P07	Z2		
A1	A2	P06	Z1		
A1	A2	P05	X		
A1	A2	P04	B1		
A1	A2	P03	G1		
A1	A2	P02	X		(13 or 14
A1	A2	P01	X	Eleventh ATM cell	nibbles)
A1	A2	P00	C1	Twelfth ATM cell	Trailer

Figure 2.13 PLCP frame. POI = path overhead indicator; POH = path overhead; BIP-8 = bit interleaved parity-8; X = unassigned (receiver to ignore). [*Note:* Order and transmission of all PLCP bits and octets are from left to right and top to bottom. This figure shows the most significant bit (MSB) on the left and the least significant bit (LSB) on the right.]

Header error control generation. The HEC generation is based on the algorithm described in ITU-T Recommendation I.432.

Cell delineation. Since the cells are in predetermined locations within the PLCP, framing on the 44,736-kbit/s signal and then on the PLCP is sufficient to delineate cells.

Cell header verification and extraction. The cell header verification is consistent with ITU-T Recommendation I.432. Only valid cells are passed to the ATM Layer.

PLCP overhead utilization. The following PLCP overhead bytes/nibbles are activated across the UNI:

- A1: Frame alignment
- A2: Frame alignment
- B1: PLCP path error monitoring
- C1: Cycle/stuff counter
- G1: PLCP path status
- Px: Path overhead identifier
- Zx: Growth octets
- Trailer nibbles

Frame alignment (A1, A2). The PLCP framing octets use the same framing pattern: A1 = 11110110, A2 = 00101000.

PLCP path error monitoring (B1). The BIP-8 field supports path error monitoring, and is calculated over a 12 × 54 octet structure consisting of the POH field and the associated ATM cells (648 octets) of the *previous* PLCP frame.

Cycle / stuff counter (C1). The cycle/stuff counter provides a nibble-stuffing opportunity cycle and length indicator for the PLCP frame. A stuffing opportunity occurs every third frame of a three-frame (375-μs) stuffing cycle. The value of the C1 code is used as an indication of the phase of the 375 μs stuffing opportunity cycle, as follows:

C1 code	Frame phase of cycle	Trailer length
11111111	1	13
00000000	2	14
01100110	3 (no stuff)	13
10011001	3 (stuff)	14

Notice that a trailer containing 13 nibbles is used in the first frame of the 375 ms stuffing opportunity cycle. A trailer of 14 nibbles is used in the second frame. The third frame provides a nibble-stuffing opportunity. A trailer containing 14 nibbles is used in the third frame if a stuff occurs. If it does not, the trailer will contain 13 nibbles.

PLCP path status (G1). The PLCP path status is allocated to convey the received PLCP status and performance to the transmitting far end. This octet permits the status of the full receive/transmit PLCP path to be monitored at either end of the path.

Path overhead identifier (P00–P11). The path overhead identifier (POI) indexes the adjacent path overhead (POH) octet of the PLCP.

Growth octets. These are reserved for future use. The receiver ignores the values contained in these fields.

Trailer nibbles. The content of each of the 13 or 14 trailer nibbles is 1100.

Other Mappings. Other mappings have been defined. Direct mappings for E1, DS2, and STS-3c are available.[4]

References

1. ITU-T Recommendation X.210, *Open Systems Interconnection, Layer Service Definition Convention,* Geneva, Switzerland, 1989.
2. ITU-T Recommendation I.361, *B-ISDN ATM Layer Specification*, Geneva, Switzerland, June 1992.
3. T1S1.5/92-410, *Broadband ISDN—ATM Layer Functionality and Specification,*" August 1992.
4. ITU-T Recommendation I.432, *B-ISDN User-Network Interface—Physical Layer Specification*, Geneva, Switzerland, June 1992.

5. ATM Forum, *ATM User-Network Interface Specification*, Version 3.0, August, 1993.
6. Bellcore, *Asynchronous Transfer Mode (ATM) and ATM Adaptation Layer (AAL) Protocols Generic Requirements*, TA-NWT-001113, Issue 2, July 1993.
7. ANSI T1.ATM-1993, *Broadband ISDN—ATM Layer Functionality and Specification*, New York.
8. ITU-T Recommendation I.413, *B-ISDN User-Network Interface*, Geneva, Switzerland, 1991.
9. ITU Draft Recommendation G.804, *ATM Cell Mapping into Plesiochronous Digital Hierarchy*, Geneva, Switzerland, February 1993.
10. D. Minoli, *Enterprise Networking—Fractional T1 to SONET, Frame Relay to BISDN*, Artech House, Norwood, Mass., 1993.

ATM Adaptation Layer

3.1 Introduction

As discussed in the previous two chapters, the Protocol Reference Model applicable to both the User Plane and the Control Plane (see Fig. 3.1) is divided into three protocol layers: the *Physical Layer*, the *ATM Layer*, and the *AAL and Service-Specific Layers*.

- The Physical Layer provides the ATM Layer with access to the physical transmission medium. Its functions include transmission of bits across the physical medium, timing recovery, line coding, cell delineation, cell scrambling and descrambling, and generation and checking of the header error control.
- The ATM Layer provides for the transport of ATM cells between the endpoints of a virtual connection. It is the basis for native cell relay service as well as other services. ATM cells are delivered across the network in the same sequence they are received from the CPE.
- The AAL maps the upper-layer data into cells for transport across the network. The Service-Specific Layers perform application-dependent processing and functions.

This chapter focuses on AAL protocols. As noted, the AAL performs the functions necessary to adapt the capabilities provided by the ATM Layer to the needs of higher-layer applications using CRS or other ATM-based services.[1–4] AALs are typically implemented in end user equipment, as shown, for example, in Fig. 1.16, but can also (occasionally) be found in the network, as seen later. The functions of the AAL include segmentation and reassembly of the higher-layer data units and mapping them into the

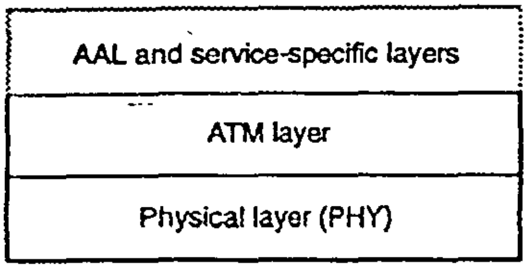

Figure 3.1 Protocol reference model.

fixed-length payload of the ATM cells. Effectively, AAL protocols allow a user with some preexisting application, say using TCP/IP, to get the benefits of ATM. To date, three AAL protocol types have been standardized: AAL Type 1 for circuit emulation (or CBR) services, and AAL Type 3/4 and AAL Type 5 for VBR services. A number of service-specific parts have also been standardized. For many years "AAL" meant segmentation/reassembly and error detection only. With the recent inclusion of service-specific functions into the AAL, the functionality has been significantly increased. Two examples of service-specific parts are briefly discussed at the end of this chapter. In AAL Type 1, 1 octet of the cell payload is reserved for control; the remaining 47 octets are utilized for user information. AAL Type 3/4 reserves 4 octets of each cell payload for control use. AAL Type 5 provides all 48 octets of each cell (except for the last cell of a higher-layer packet; see Sec. 3.5.2) for user information.

Note: In this discussion, the term *user* is employed consistent with protocol parlance, unless noted otherwise. Namely, it represents the (protocol) entity just above the AAL Layer; it does not refer to the ultimate user of the (corporate) network. Such a corporate user would access ATM through the top of the protocol stack, e.g., via an application such as E-mail over TCP/IP over ATM.

Recall, for positioning, as we proceed, that AAL provides the balance of capabilities to "fill out" part, but not all, of the Data Link Layer in the OSIRM. Typically the stack {AAL, ATM, PHY} runs just under the Logical Link Control of a traditional LAN, or directly under TCP/IP in an ATM-based LAN or ATM-based WAN.

The novice reader may choose to skip this chapter on first reading; alternatively, the reader may read the first few sections to understand what the AAL aims at doing, without concentrating on how it does it.

3.2 AAL Model

Architecturally, the AAL is a layer between the ATM Layer and the "service layer" (the service layer is shown in Fig. 3.5). The purpose of the ATM Adaptation Layer is to provide the necessary functions to support the service layer that are not provided by the ATM Layer. The functions provided by the AAL depend upon the service. VBR users may require such functions as PDU delimitation, bit error detection and correction, and cell loss detection. CBR users typically require source clock frequency recovery and detection and possible replacement of lost cells.

Figure 3.2 depicts the positioning of the AAL in the context of the corporate user equipment. AAL capabilities can also be used at an interworking point in the carrier's network, as shown in Fig. 3.3 (this topic is reexamined in Chap. 7). Figure 3.4 shows a classification of services that has been used for specifying ATM Adaptation Layers for different services.

Five AAL protocol types to support the following services are covered in this chapter:

- CBR service using the AAL 1 protocol
- VBR service using the AAL 3/4 Common Part protocol
- VBR service using the AAL 5 Common Part protocol
- Frame relay service (the Frame Relay Service-Specific AAL protocol, which utilizes the AAL 5 Common Part protocol)
- UNI signaling service (the UNI Signaling AAL protocol, which utilizes the AAL 5 Common Part protocol)

The AAL for VBR services consists of two parts: a Common Part (CP) and a Service-Specific Part (SSP). The SSP is used to provide those additional capabilities, beyond those provided by the CP, that are necessary to support the user of the AAL. For some applications the SSP may be "null"; in these cases, the user of the AAL utilizes the AAL Common Part (AALCP) directly. For all AAL types, the AAL receives information from the ATM Layer in the form of 48-octet ATM service data units (ATM_SDU). The AAL passes information to the ATM Layer in the form of a 48-octet ATM_SDU. Figure 3.5 depicts some of the more common protocol arrangements.

Section 3.3 discusses the AAL description for Class 1 (e.g., circuit emulation services), and Sec. 3.4 discusses the AAL description for Class

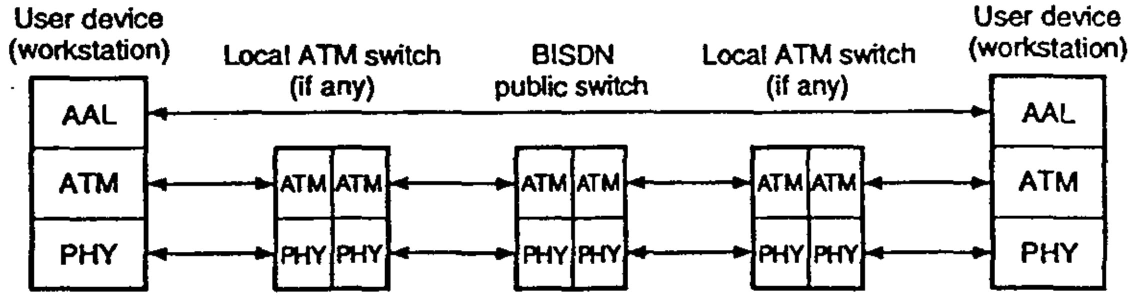

Figure 3.2 The positioning of AAL in CPE.

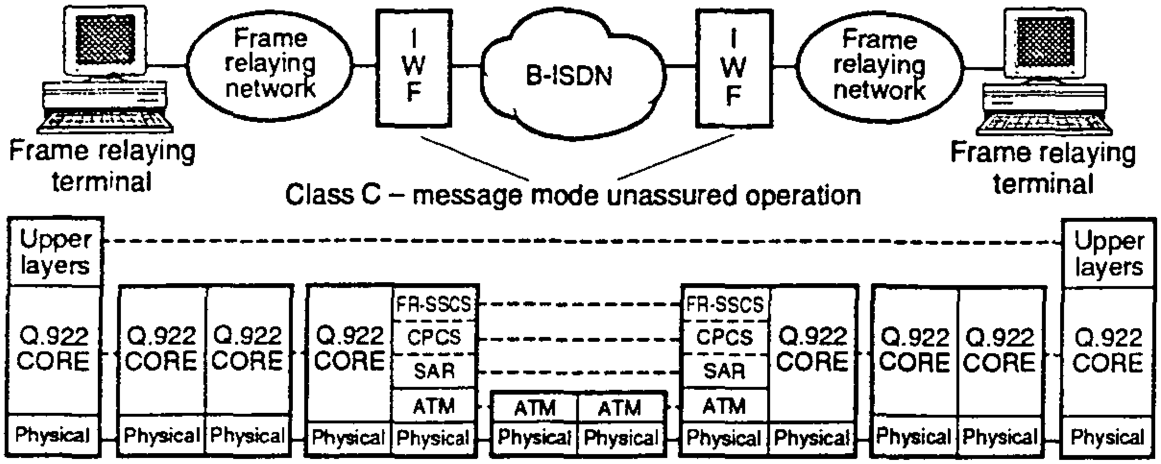

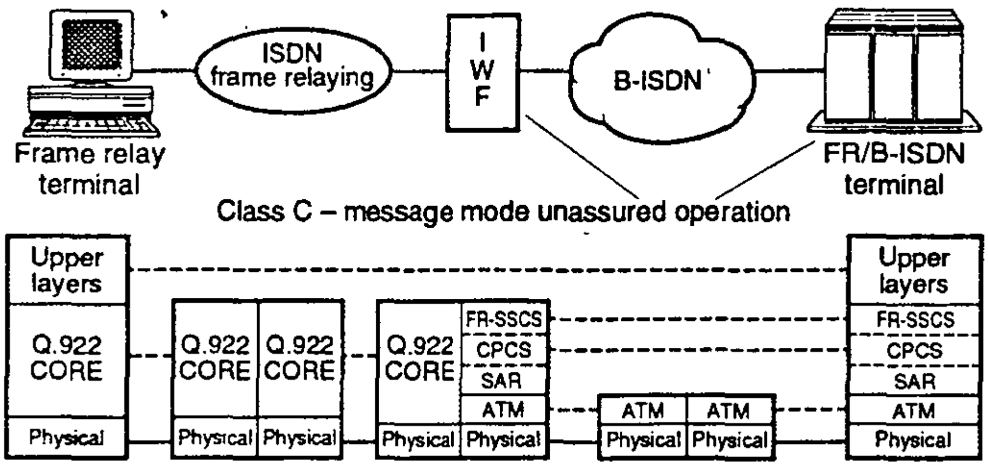

Figure 3.3 Use of AAL protocols at interworking points.

3/4 (e.g., connectionless data services, such as SMDS). Maximum commonality between Class 4 and Class 3 (e.g., connection-oriented data services) AALs has been sought, and people now refer to this AAL as AAL 3/4. The AAL specification for Class 2 services (e.g., variable-bit-rate video services) may occur at a future date. Section 3.5 describes AAL 5, Sec. 3.6 covers the Frame Relay Service-Specific AAL, and Sec. 3.7 briefly covers the signaling AAL.

3.3 AAL Type 1

3.3.1 Overview

One of the services possible with an ATM platform is emulation of a dedicated line (typically at 1.544 or 45 Mbits/s). This type of service is also known as Class A or CBR service. To support CBR services, an adaptation layer is required in the user's equipment for the necessary

<table>
<tr><th>Attributes</th><th>Class 1</th><th>Class 2</th><th>Class 3</th><th>Class 4</th></tr>
<tr><td>Timing between source and destination</td><td colspan="2">Related</td><td colspan="2">Nonrelated</td></tr>
<tr><td>Bit rate</td><td>Constant</td><td colspan="3">Variable</td></tr>
<tr><td>Connection mode</td><td colspan="3">Connection-oriented</td><td>Connection-less</td></tr>
</table>

Figure 3.4 Classification of services for AAL specification. Examples of services: Class 1, circuit emulation; Class 2, variable bit rate video; Class 3, connection-oriented data; Class 4, support of connectionless data transfer; Class X, unrestricted.

functions that cannot be provided by the ATM cell header. Some characteristics and functions that may be needed for an efficient and reliable transport of CBR services are identified below.

Ideally, CBR services carried over an ATM-based network should appear to the corporate user as equivalent to CBR services provided by the circuit switched or dedicated network. Some characteristics of these CBR services are

1. Maintenance of timing information
2. Reliable transmission with negligible reframes
3. Path performance monitoring capability

CBR services with the above characteristics can be provided by assigning the following functions for the CBR Adaptation Layer:

1. Lost cell detection
2. Synchronization
3. Performance monitoring

(These functions may not be required by all the CBR services.)

Therefore, the CBR AAL performs the functions necessary to match the service provided by the ATM Layer to the CBR services required by its service user. It provides for the transfer of AAL_SDUs carrying information of an AAL user supporting constant-bit-rate services. This layer is service-specific, with the main goal of supporting services that

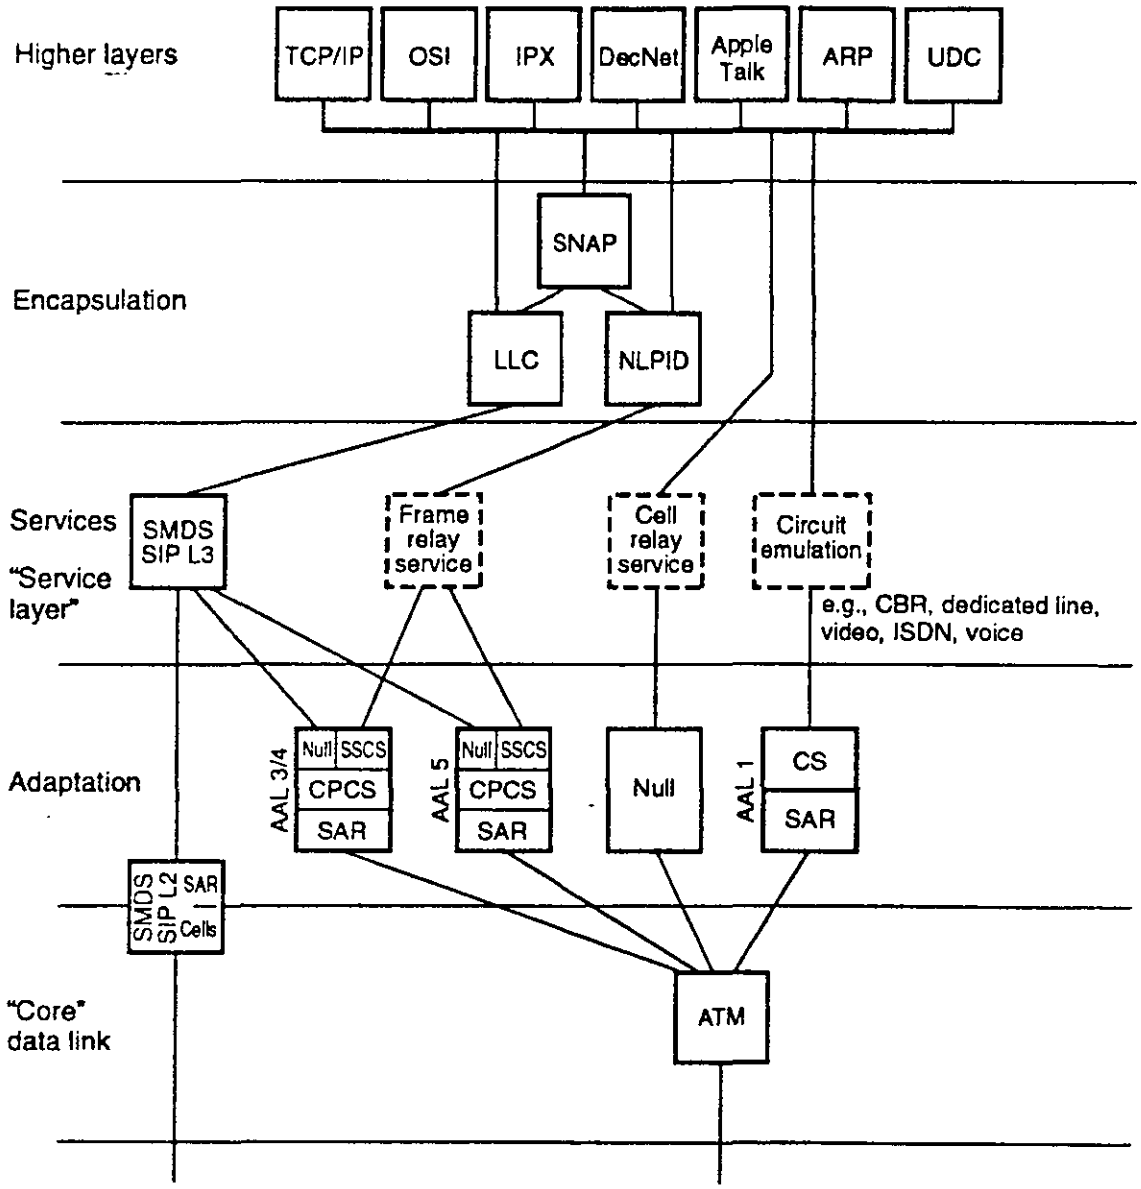

Figure 3.5 Support of user applications. CPCS = common part CS; SSCS = service-specific CS; LLC = logical link control; SNAP = Subnetwork Access Protocol; NLPID = Network Layer Protocol ID.

have specific delay, jitter, and timing requirements, such as circuit emulation. It provides timing recovery, synchronization, and indication of lost information.

The AAL 1 functions are grouped into Segmentation and Reassembly Sublayer functions and Convergence Sublayer functions. The existing agreements in ITU-T Recommendation I.363 and the ANSI CBR AAL Standard[3] provide two basic modes of operation for the CBR AAL:

- Unstructured data transfer (UDT)
- Structured data transfer (SDT)

When the UDT mode is operational, the AAL protocol assumes that the incoming data from the AAL user are a *bit stream* with an associated bit clock. When the SDT mode is operational, the AAL protocol assumes that the incoming information is *octet blocks* of a fixed length (such as an $n \times 64$ kbit/s channel with 8-kHz integrity) with an associated clock. While the SDT mode of operation has not been completely specified in the standards, a substantial enough body of agreements exists to assume that by the end of 1994 a complete SDT mechanism will be defined.

3.3.2 CBR AAL services

AAL Type 1 services and functions. The CBR AAL functions are grouped into two sublayers, the SAR Sublayer and the Convergence Sublayer. The SAR is responsible for the transport and bit error detection (and possibly correction) of CS protocol control information. The CS performs a set of service-related functions. It blocks and deblocks AAL_SDUs, counting the blocks, modulo 8, as it generates or receives them. Also, it maintains bit count integrity, generates timing information (if required), recovers timing, generates and recovers data structure information (if required), and detects and generates indications to the AAL management (AALM) entity of error conditions or signal loss. The CS may receive reference clock information from the AALM entity which is responsible for managing the AAL resources and parameters used by the AAL entity. The services provided by AAL Type 1 to the AAL user are

- Transfer of service data units with a constant source bit rate and the delivery of them with the same bit rate
- Transfer of timing information between the source and the destination
- Transfer of structure information between the source and the destination
- Indication of lost or errored information that is not recovered by AAL Type 1, if needed

Specifically, the functions are:

1. Segmentation and reassembly of user information
2. Handling of cell delay variation
3. Handling of cell payload assembly delay
4. Handling of lost and misinserted cells
5. Source clock recovery at the receiver
6. Recovery of the source data structure at the receiver

7. Monitoring of AAL-PCI for bit errors
8. Handling of AAL-PCI bit errors
9. Monitoring of the user information field for bit errors and possible corrective actions

SAR functions. The SAR functions are

- Mapping between the CS_PDU and the SAR_PDU (the SAR Sublayer at the transmitting end accepts a 47-octet block of data from the CS and then prepends a 1-octet SAR_PDU header to each block to form the SAR_PDU).
- Indicating the existence of a CS function (the SAR can indicate the existence of a CS function; the use of the indication mechanism is optional).
- Sequence numbering (for each SAR_PDU payload, the SAR sublayer receives a sequence number value from the CS).
- Error protection (the sequence number and the CSI bits are protected).

A buffer is used to handle cell delay variation. When cells are lost, it may be necessary to insert an appropriate number of dummy SAR_PDUs. Figure 3.6 depicts the AAL Type 1 frame layout.

Convergence Sublayer functions. The functions of the CS are

- Handling of cell delay variation for delivery of AAL_SDUs to the AAL user at a constant bit rate (the CS layer may need a clock derived at the S_B or T_B interface to support this function).
- Processing the sequence count to detect cell loss and misinsertion.
- Providing the mechanism for timing information transfer for AAL users requiring recovery of source clock frequency at the destination end.

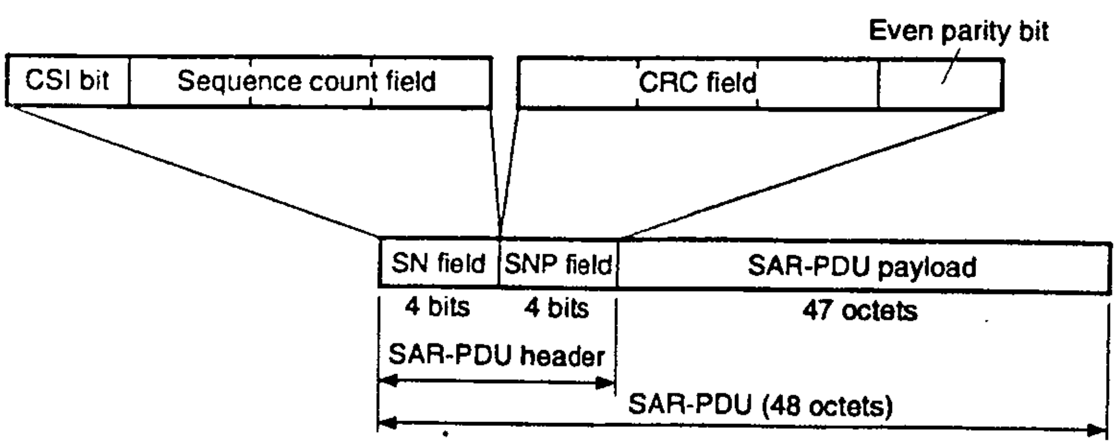

Figure 3.6 AAL Type 1 frame layout. SN = sequence number; SNP = sequence number protection; CSI = Convergence Sublayer indication.

- Providing the transfer of Structure information between source and destination for some AAL users.
- Supporting forward error correction (particularly for video)

For those AAL users that require transfer of structured data [e.g., 8-kHz structured data for circuit-mode bearer services for 64-kbit/s-based ISDN (see Chap. 8)], the Structure parameter is used. This parameter can be used when the user data stream to be transferred to the peer AAL entity is organized into groups of bits. The length of the structured block is fixed for each instance of the AAL service. The length is an integer multiple of 8 bits. An example of the use of this parameter is to support circuit-mode services of the 64-kbit/s-based ISDN. The two values of the Structure parameter are

Start. This value is used when the DATA is the first part of a structured block, which can be composed of consecutive data segments.

Continuation. This value is used when the value Start is not applicable.

The use of the Structure parameter depends on the type of AAL service provided; its use is agreed upon prior to or at the connection establishment between the AAL user and the AAL.

I.363 notes that "for certain applications such as speech, some SAR functions may not be needed." For example, I.363 provides the following guidance for CS for voice-band signal transport [which is a specific example of CBR service (see Chap. 8)]:

- *Handling of AAL user information.* The length of the AAL_SDU (i.e., the information provided to the AAL by the upper-layer protocols) is 1 octet (for comparison, the SAR_PDU is 47 octets).
- *Handling of cell delay variation.* A buffer of appropriate size is used to support this function.
- *Handling of lost and misinserted cells.* The detection of lost and inserted cells, if needed, may be provided by processing the sequence count values. The monitoring of the buffer fill level can also provide an indication of lost and misinserted cells. Detected misinserted cells are discarded.

P and non-P formats. The 47-octet SAR_PDU payload used by CS has two formats called non-P and P formats, as seen in Fig. 3.7. These are used to support transfer of information with Structure.

Note that in the non-P format, the entire CS_PDU is filled with user information.

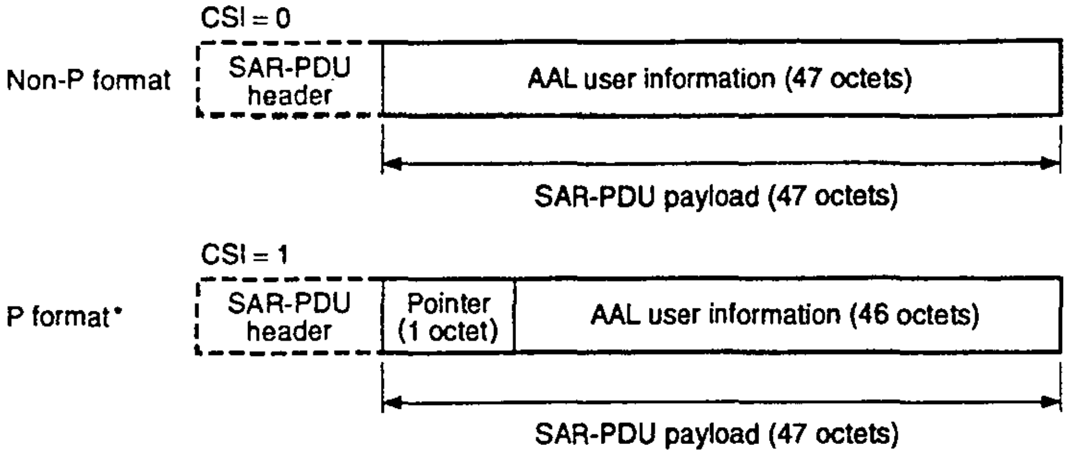

* Used when the SAR-PDU SN = 0, 2, 4, or 6

Figure 3.7 Non-P and P formats.

Partially filled cells. I.363 notes that SAR_PDU payload may be filled only partially with user data in order to reduce the cell payload assembly delay. In this case, the number of leading octets utilized for user information in each SAR_PDU payload is a constant that is determined by the allowable cell payload assembly delay. The remainder of the SAR_PDU payload consists of dummy octets.

Clocking issues. Besides the UDT/SDT issues discussed earlier, the other basic CBR service attribute that determines the AAL functionality required to support a service is the status of the CBR service clock:[5]

- Synchronous
- Asynchronous

Since the service clock is assumed to be frequency-locked to a network clock in the synchronous case, its recovery is done directly with a clock available from the network. For an asynchronous service clock, the AAL provides a method for recovering the source clock at the receiver. Two methods are available, the synchronous residual time stamp (SRTS) method and the adaptive clock method. The SRTS method is used to recover clocks with tight tolerance and jitter requirements, such as DS1 or DS3 clocks. The adaptive clock recovery method has not been described in enough detail to determine what types of service clocks are supported [presumably less accurate clocks with looser low-frequency jitter (i.e., wander) specifications] or what, if any, added agreements are needed. However, since adaptive clock recovery is common in user equipment, this method is assumed to be available.

The support of DS1 and DS3 CBR service

- Uses the entire 47-octet information payload available with the basic CBR AAL protocol.
- Uses the UDT mode of operation.
- Uses the SRTS method of timing recovery, if the service clock is asynchronous.
- Maintains bit count integrity by inserting the appropriate alarm indication signal for the service supported as a DS1 and DS3 error control measure.

3.3.3 CBR AAL mechanism

The CBR AAL provides its service over preestablished AAL connections. The establishment and initialization of an AAL connection is performed through the AALM. The transfer capacity of each connection and other connection characteristics are negotiated prior to or at connection establishment (the CBR AAL is not directly involved in the negotiation process, which may be performed by management or signaling). The AAL receives from its service user a constant-rate bit stream with a clock. It provides to its service user this constant-rate bit stream with the same clock. The CBR service clock can be either synchronous or asynchronous relative to the network clock. The CBR service is called synchronous if its service clock is frequency-locked to the network clock. Otherwise, the CBR service is called asynchronous.

The service provided by the AAL consists of its own capability plus the capability of the ATM Layer and the Physical Layer. This service is provided to the AAL user (e.g., an entity in an upper layer or in the Management Plane). The service definition is based on a set of service primitives that describe in an abstract manner the logical exchange of information and control. Functions performed by the CBR AAL entities are shown in Table 3.1.

The logical exchange of information between the AAL and the AAL user is represented by two primitives, as shown in Table 3.2.

Service expected from the ATM Layer. The AAL expects the ATM Layer to provide for the transparent and sequential transfer of AAL data units, each of length 48 octets, between communicating AAL entities over an ATM Layer connection, at a negotiated bandwidth and QOS. The ATM Layer transfers the information in the order in which it was delivered to the ATM Layer and provides no retransmission of lost or corrupted information.

TABLE 3.1 Functions Performed by CBR AAL

Detection and reporting of lost SAR_PDUs	Detects discontinuity in the sequence count values of the SAR_PDUs and senses buffer underflow and overflow conditions.
Detection and correction of SAR_PDU header error	Detects bit errors in the SAR_PDU header and possibly corrects a 1-bit error.
Bit count integrity	Generates dummy information units to replace lost AAL_SDUs to be passed to the AAL user in an AAL-DATA.indication.
Residual time stamp (RTS) generation	Encodes source service clock timing information for transport to the receiving AAL entity.*
Source clock recovery	Recovers the CBR service source clock.
Blocking	Maps AAL_SDUs into the payload of a CS_PDU.
Deblocking	Reconstructs the AAL_SDU from the received SAR_PDUs and generates the AAL-DATA. indication primitive.
Structure pointer generation and extraction	Encodes in a 1-octet structure pointer field at the sending AAL entity the information about periodic octet-based block structures present in AAL-DATA.request primitives. The receiving AAL entity extracts the structure pointer received in the CS_PDU header field to verify locally generated block structure.

*Refer to Ref. 3 for a description of the time stamp mechanism.

Interactions between the SAR and the Convergence Sublayer. The logical exchange of information between the SAR and the Convergence Sublayer is represented by the primitives of Table 3.3.

Interacting with the Management Plane. The AALM entities in the Management Plane perform the management functions specific to the AAL. Also, the AALM entities, in conjunction with the Plane Management, provide coordination of the local interactions between the User Plane and the Control Plane across the layers.

The AAL entities provide the AALM entities with the information required for error processing or abnormal condition handling, such as indication of lost or misdelivered SAR_PDUs and indication of errored SAR_PDU headers.

TABLE 3.2 Primitives for CBR AAL

AAL-DATA.request (AAL_SDU, Structure)	This primitive is issued by an AAL user entity to request the transfer of an AAL_SDU to its peer entity over an existing AAL connection. The time interval between two consecutive AAL-DATA.request primitives is constant and a function of the specific AAL service provided to the AAL user.
AAL-DATA.indication (AAL_SDU, Structure, Status)	This primitive is issued to an AAL user entity to notify the arrival of an AAL_SDU over an existing AAL connection. In the absence of error, the AAL_SDU is the same as the AAL_SDU sent by the peer AAL user entity in the corresponding AAL-DATA.request. The time interval between two consecutive AAL-DATA.indication primitives is constant and a function of the specific AAL service provided to the AAL user.

Description of parameters:

AAL_SDU: This parameter contains 1 bit of AAL user data to be transferred by the AAL between two communication AAL user peer entities.

Structure: This parameter is used to indicate the beginning or continuation of a block of AAL_SDUs when providing for the transfer of a structured bit stream between communicating AAL user peer entities (structured data transfer service). The length of the blocks is constant for each instance of the AAL service and is a multiple of 8 bits. This parameter takes one of the following two values: Start and Continuation. It is set to Start whenever the AAL_SDU being passed in the same primitive is the first bit of a block of a structured bit stream. Otherwise, it is set to Continuation. This parameter is used only when SDT service is supported.

Status: This parameter indicates whether the AAL_SDU being passed in the same indication primitive is judged to be nonerrored or errored. It takes one of the following two values: Valid or Invalid. The Invalid value may also indicate that the AAL_SDU being passed is a dummy value. The use of this parameter and the choice of the dummy value depend on the specific service provided.

TABLE 3.3 SAR Primitives

SAR-DATA.invoke (CSDATA, SCVAL, CSIVAL)	This primitive is issued by the sending CS entity to the sending SAR entity to request the transfer of a CSDATA to its peer entity.
SAR-DATA.signal (CSDATA, SNCK, SCVAL, CSIVAL)	This primitive is issued by the receiving SAR entity to the receiving CS entity to notify it of the arrival of a CSDATA from its peer CS entity.

Description of parameters:

CSDATA: This parameter represents the interface data unit exchanged between the SAR entity and the CS entity. It contains the 47-octet CS_PDU.

SCVAL: This 3-bit parameter contains the value of the sequence count associated with the CS_PDU contained in the CSDATA parameter.

CSIVAL: This 1-bit parameter contains the value of the CSI bit.

SNCK: This parameter is generated by the receiving SAR entity. It represents the results of the sequence number protection error check over the SAR_PDU header. It can assume the values of SN-Valid and SN-Invalid.

3.4 ATM Adaptation Layer Functions for VBR (or Bursty Data) Services

As seen in Fig. 3.5, AAL functions for VBR services such as SMDS and frame relay consist of a set of core functions and a set of optional functions. This AAL is now commonly referred to as AAL Type 3/4. As an example, SMDS over ATM uses AAL Type 3/4. The purpose of the ATM Adaptation Layer Type 4/3 Common Part (CPAAL3/4) protocol is to support the upper-layer data transfer needs while using the service of the ATM Layer. This protocol provides for the transport of variable-length frames (up to 65,535 octets in length) with error detection. The CPAAL3/4 provides service over preestablished connections. Termination of a CPAAL3/4 connection also coincides with termination of an ATM Layer service. The establishment and initialization of a CPAAL3/4 connection is performed by interaction with CPAAL3/4 Layer Management entities. There is a dual view of the AAL3/4 Layer.

1. View in terms of Service-Specific Parts and Common Part, as shown in the left-hand side of Fig. 3.8. Core functions are required by all bursty data applications; these functions are known as CP. Optional SSPs are selected as needed. For some applications the SSP is null, implying that the user of the AAL3/4 Layer utilizes the Common Part directly.

2. View in terms of a combination of SAR, the Common Part of the Convergence Sublayer, and SSP, as shown in the right-hand side of Fig. 3.8. SAR and the Common Part of the Convergence Sublayer taken together make up the CP; the Common Part of CS and SSP together form the CS. In other words, the Convergence Sublayer has been

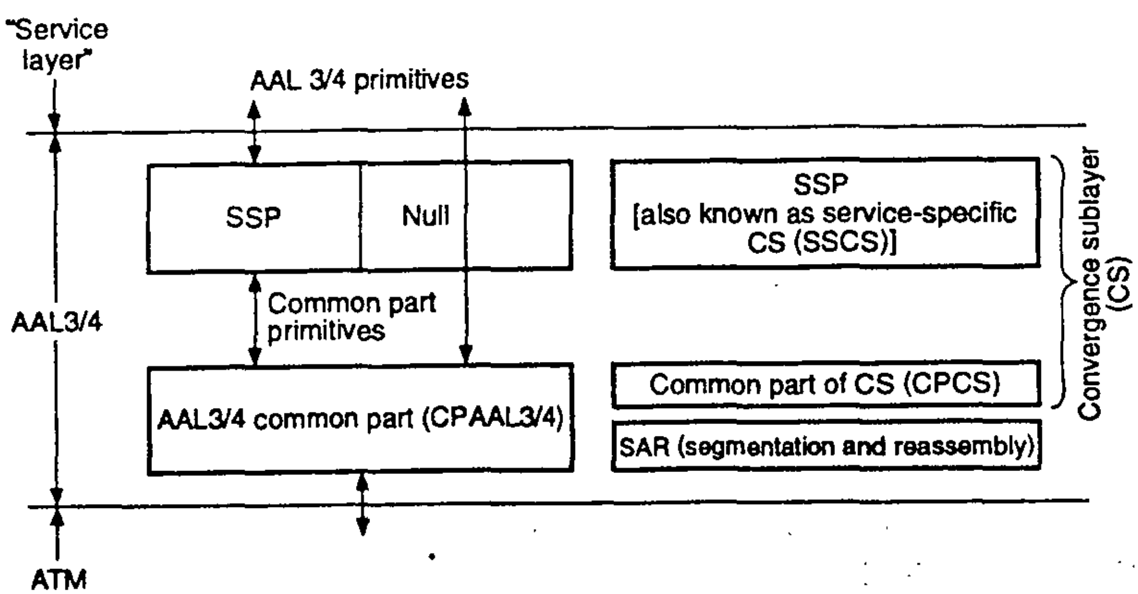

Figure 3.8 Model of AAL3/4. Left: CP/SSP view; right: CS/SAR view.

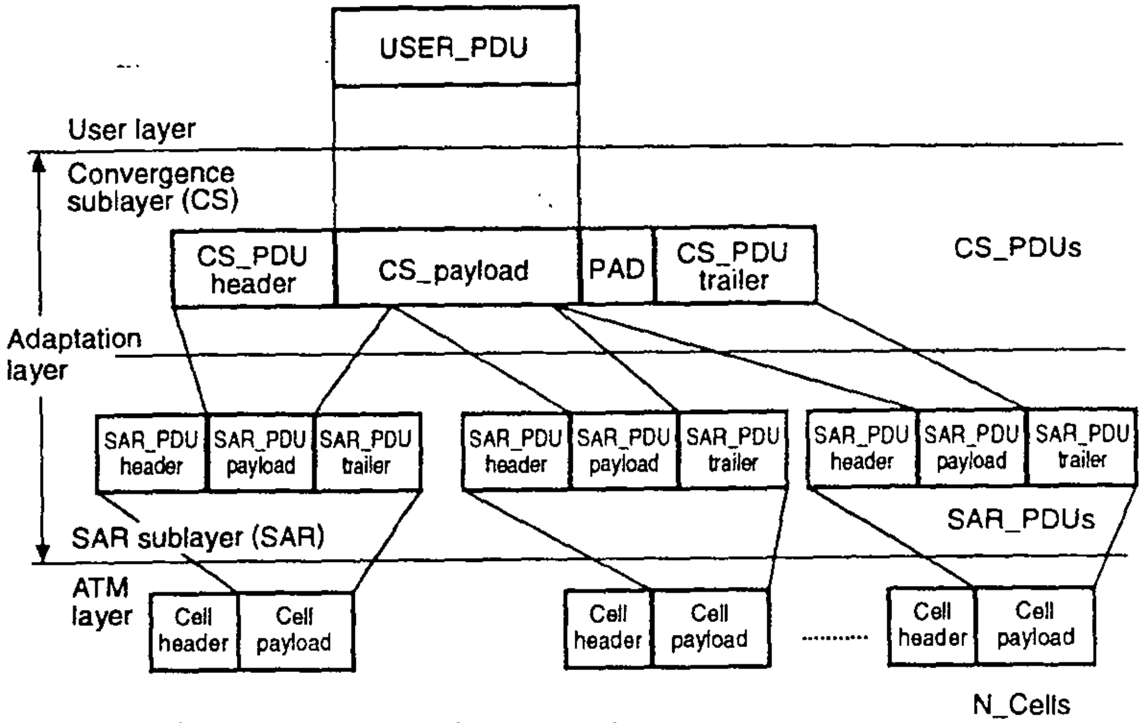

Figure 3.9 Adaptation Layer model for bursty data services.

subdivided into the Common Part CS (CPCS) and the Service-Specific CS (SSCS). In this view, functions are provided by the operation of two logical sublayers, the CS and the SAR. Figure 3.9 shows the operation of AAL3/4 in terms of the PDUs.

The SAR Sublayer is common to all VBR services using AAL3/4, whereas the Convergence Sublayer provides additional, service-specific functions (note that some VBR services may use AAL5). The functions of the Common Part are clearly common by definition. In addition to this, achieving the maximum commonality in the Convergence Sublayer protocol for bursty data services has also been an objective, as implied in Fig. 3.5. For these services, the user presents a variable-size PDU for transmission across the ATM network. The transmission is accomplished by using fixed-length cells to transport data in ATM, as discussed in Chap. 2. At the receiving end of the ATM connection, the user layer receives the PDU that has been reassembled by the SAR and CS protocols.

The discussion that follows looks at AAL3/4 first from a CP point of view (the left-hand model in Fig. 3.8), then from the SAR point of view (the right-hand side of Fig. 3.8). As noted, the functions of the CPAAL3/4 in this view have been grouped into two sublayers: CPAAL3/4 Segmentation and Reassembly (CPAAL3/4_SAR) and CPAAL3/4 Convergence Sublayer (CPAAL3/4_CS). The CPAAL3/4_SAR deals principally with the segmentation and reassembly of data units so that they can be mapped into fixed-length payloads of the ATM cells, while the CPAAL3/4_CS deals mainly with checking missassembled CPAAL3/4_CS_PDUs.

CPAAL3/4 Layer Management is responsible for the following capabilities: assignment of the CPAAL3/4 association necessary for the establishment of CPAAL3/4 connections between peer CPAAL3/4 entities, resetting the parameters and state variables associated with a CPAAL3/4 connection between peer CPAAL3/4 entities, and monitoring performance for the quality of the ATM connection service provided through notification of errors.

3.4.1 Services provided to the upper layer

The CPAAL3/4 provides, on behalf of its user, for the sequential and transparent transfer of variable-length, octet-aligned CPAAL3/4_SDUs from one corresponding CPAAL3/4 peer to one or more CPAAL3/4 peers. The service is unassured: CPAAL3/4_SDUs may be lost or corrupted. Lost or corrupted CPAAL3/4_SDUs are not recovered by the CPAAL3/4. As an option, corrupted CPAAL3/4_SDUs may be delivered to the remote peer with an indication of the error (this option is known as corrupted data delivery option).

Specifically, the functions performed by the CPAAL3/4 are[6]

- Data transfer between CPAAL3/4 peers
- Preservation of CPAAL3/4_SDUs (delineation and transparency of CPAAL3/4_SDUs)
- CPAAL3/4_SDU segmentation
- CPAAL3/4_SDU reassembly
- Error detection and handling (detects and handles bit errors, lost or gained information, and incorrectly assembled CPAAL3/4_SDUs)
- Multiplexing and demultiplexing (optional multiplexing of multiple CPAAL3/4 connections or interleaving of CPAAL3/4_CS_PDUs)
- Abort (termination of task in case of partially transmitted/received CPAAL3/4_SDUs)
- Pipelining (forwarding PDUs before the entire PDU is received)

This layer provides its user two services:

1. Message-mode service: In this service mode, the CPAAL3/4_SDU passed across the CPAAL3/4 interface is exactly equal to one CPAAL3/4 interface data unit (CPAAL3/4_IDU), as seen in Fig. 3.10.

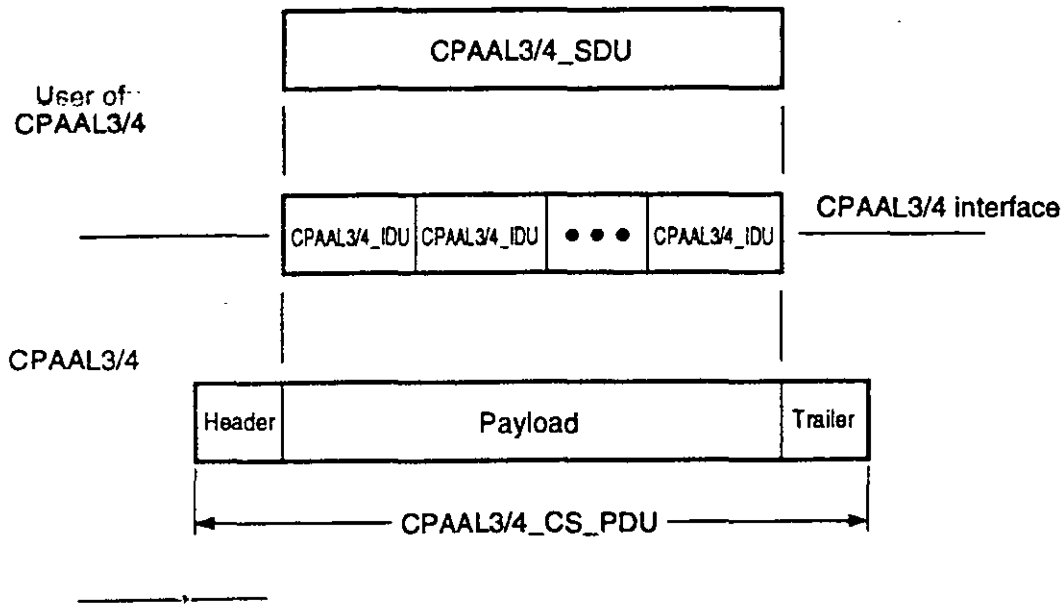

Figure 3.11 Streaming-mode service.

the .signal to an .indication. Table 3.4 provides additional information on these primitives.

Services from the ATM Layer. The CPAAL3/4 expects the ATM Layer (discussed in Chap. 2) to provide for the transparent and sequential transport of 48-octet CPAAL3/4 data units (that is, CPAAL3/4_SAR_PDUs) between communicating CPAAL3/4 peers over preestablished connections at a negotiated QOS. The information is transferred to the ATM Layer in the order in which it is to be sent, with no retransmission of lost or corrupted information.

Interaction with CPAAL3/4 Management entities. Management information is exchanged using five management primitives. See Ref. 4 for details.

3.4.2 SAR Sublayer functions

There is a single SAR function for all bursty data services. Hence, the SAR control fields that appear in each cell payload must be the same, regardless of the service and whether or not the fields are used by a particular application. A single SAR for these services leads to lower overall costs for equipment providers and network providers, and hence for end users (e.g., diagnostic generation, testing, and maintenance are simpler when only a single SAR function is used for all services).

The SAR control fields include the following:[6]

Segment_Type field to identify the cell payload as being beginning of message (BOM), continuation of message (COM), end of message (EOM), or only a single-segment message (SSM).

TABLE 3.4 CPAAL3/4 Primitives

CPAAL3/4-UNITDATA.invoke	Issued by a CPAAL3/4 entity to request the transfer of a CPAAL3/4_IDU over an existing CPAAL3/4 connection. This IDU is not subject to any flow control and is always transmitted. The transfer of the IDU is subject to the service mode being used (message versus streaming).
CPAAL3/4-UNITDATA.signal	Issued to a CPAAL3/4 entity to indicate the arrival of a CPAAL3/4_IDU over an existing CPAAL3/4 connection.
CPAAL3/4-U-ABORT.invoke	Issued by a CPAAL3/4 entity using streaming-mode service to request the termination of a CPAAL3/4_SDU that has been partially transferred. The issue of this primitive also causes the generation of an abort message by the CPAAL3/4 to its peer entity if the transmission of the message has already started. (This primitive is not used in message mode.)
CPAAL3/4-U-ABORT.signal	Issued by a CPAAL3/4 entity using streaming-mode service to indicate the termination of a partially delivered CPAAL3/4_SDU by instruction from its peer entity. (This primitive is not used in message mode.)
CPAAL3/4-P-ABORT.signal	Issued by a CPAAL3/4 entity using streaming-mode service to indicate to its user that a partially delivered CPAAL-3/4_SDU is to be discarded because of the occurrence of some error; it has local significance. (This primitive is not used in message mode.)

Description of parameters:

ID (Interface data): This parameter contains the interface data unit (CPAAL3/4_IDU) exchanged between CPAAL3/4 entities [it may be the entire CPAAL3/4_SDU (message mode) or segments (streaming mode)].

M (more): Used only in streaming mode to indicate whether the CPAAL3/4_IDU communicated in the ID parameter contains the ending segment of the CPAAL3/4_PDU (=0) or does not (=1).

ML (maximum length): Used only in streaming mode to indicate the maximum length of the CPAAL3/4_SDU; it has values from 0 to 65,535.

RS (reception status): Indicates that the CPAAL3/4_IDU delivered may be corrupted.

LP (loss priority): Indicates the loss priority assigned to the CPAAL3/4_SDU. Two levels of priority are supported, but how to map this parameter to and from the ATM_Submitted_Loss_Priority (discussed in Chap. 2) has not yet been worked out.

CI (congestion indication): Indicates the detection of congestion experienced by the received CPAAL3/4_SDU.

Sequence_Number field to improve the reassembly error detection process.

Message_ID (M_ID) field, which, for connectionless services, allows for the collection of the cell payloads that make up a CS PDU.

Cell Fill field that allows the identification of the fill within a cell payload. It can be used to locate the last octet in the end of message cell. The last octet in the EOM cell could also be identified from the length field associated with the PDU; additionally, data pipelining could be provided by a series of partially filled single-segment message cells. However, in the latter case, significant additional processing is required to reconstruct the original data unit compared with the case where partial fills are indicated by a cell-associated length field.

Error Control field which provides error detection capabilities across the adaptation header and the information payload. The error check is made across all 48 octets irrespective of whether the cell is fully or partially filled.

On transmission, the process is used by the sending CPAAL3/4 entity. The SAR Sublayer accepts variable-length CPAAL3/4_CS_PDUs from the Convergence Sublayer and maps each CPAAL3/4_CS_PDU into a sequence of CPAAL3/4_SAR_PDUs, by placing at most 44 octets of the CPAAL3/4_CS_PDU into a CPAAL3/4_SAR_PDU payload, along with additional control information, described below, used to verify the integrity of the CPAAL3/4_SAR_PDU payload on reception and to control the reassembly process. The sending CPAAL3/4 entity transfers the CPAAL3/4_SAR_PDUs to the ATM Layer for delivery across the network.

On reception, CPAAL3/4_SAR_PDUs are validated, and the user data in the CPAAL3/4_SAR_PDU (note that a CPAAL3/4_SAR_PDU can be partially filled) are passed to the Convergence Sublayer.

3.4.3 Convergence Sublayer functions

On transmission, the Convergence Sublayer accepts variable-length user protocol data units (USER_PDUs) from the service layer. The Convergence Sublayer prepends a 32-bit header to the USER_PDU, then appends from 0 to 3 pad octets to the USER_PDU to build it out to an integral multiple of 32 bits. Next, it appends a 32-bit trailer to the concatenated header, USER_PDU, and pad structure. This collection (the header, USER_PDU, pad, and trailer) is referred to as a CPAAL3/4_CS_PDU. The header and trailer fields are used to detect loss of data and to perform additional functions as required by the service user. After appending the trailer, the Convergence Sublayer passes the CPAAL3/4_CS_PDU to the SAR Sublayer for segmentation and then transmission.[6]

On reception, the Convergence Sublayer validates the collection of CPAAL3/4_SAR_PDU payloads received from the SAR Sublayer by using the information contained in the Convergence Sublayer header and trailer. It removes the pad octets, if any, and presents the validated CPAAL3/4_CS_PDU payload to the user (i.e., the service layer).

3.4.4 SAR Sublayer fields and format

The SAR Sublayer functions are implemented using a 2-octet adaptation header and a 2-octet adaptation trailer. The header and trailer, together with 44 octets of user information, make up the payload of the ATM cell. The sizes and positions of the fields are given in Fig. 3.12. The use of the error control field for error detection is mandatory. The 10-bit CRC has the capability of single-bit error correction over the 48 octets. If the underlying transmission system produces single-bit errors, error correction may be applied at the receiver.

Figure 3.12 shows the CPAAL3/4_SAR_PDU components of the Adaptation Layer, which include a SAR_PDU_Header and an SAR_PDU_Trailer. These two fields encapsulate the SAR_PDU_Payload, which contains a portion of the CPAAL3/4_CS_PDU.

The SAR_PDU_Header is subdivided into three fields: a Segment_Type field, a Sequence_Number field, and a Message Identification (MID) field. The SAR_PDU_Trailer is subdivided into two fields: a Payload_Length field and a Payload CRC field. Details of the purpose and encoding of each subfield follow.[6]

Segment_Type subfield. The 2-bit Segment_Type subfield is used to indicate whether a CPAAL3/4_SAR_PDU is a BOM, COM, EOM, or SSM. Table 3.5 shows the encodings for the Segment_Type subfield.

Sequence_Number subfield. Four-bits are allocated to the SAR_PDU Sequence_Number (SAR_SN) subfield, allowing the streams of

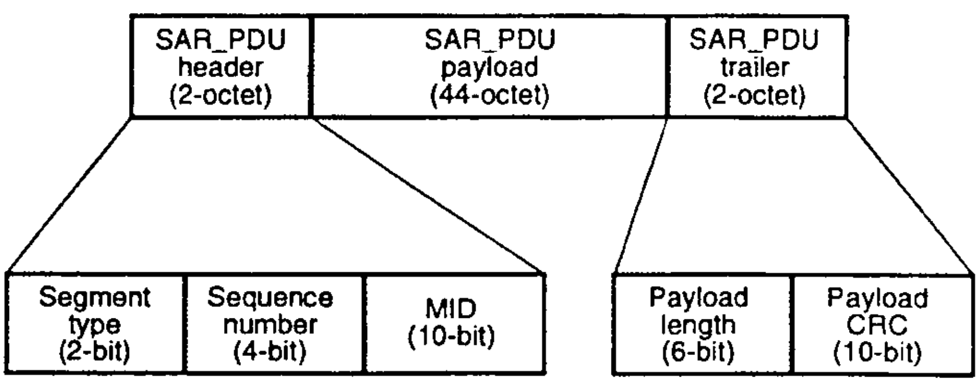

Figure 3.12 CPAAL3/4_SAR_PDU Sublayer format of AAL. MID = message identifier, or multiplexing identifier.

TABLE 3.5 Encoding of the Segment _Type Subfield

Segment_Type	Encoding
BOM	10
COM	00
EOM	01
SSM	11

CPAAL3/4_SAR_PDUs and CPAAL3/4_CS_PDUs to be numbered modulo 16. The SAR_SN is set to all 0s for the first CPAAL3/4_SAR_PDU associated with a given CPAAL3/4_CS_PDU (i.e., the BOM). For each succeeding CPAAL3/4_SAR_PDU of that CPAAL3/4_CS_PDU, the SAR_SN is incremented by 1 relative to the SAR_SN of the previous CPAAL3/4_SAR_PDU of the CPAAL3/4_CS_PDU. When reassembling a CPAAL3/4_CS_PDU, a state variable is maintained that indicates the value of the next expected SAR_SN for the CPAAL3/4_CS_PDU. If the value of the received SAR_SN differs from the expected value, the CPAAL3/4_SAR_PDU is dropped, the partially reassembled errored CPAAL3/4_CS_PDU is discarded, and any following CPAAL3/4_SAR_PDUs associated with this corrupted CPAAL3/4_CS_PDU are dropped.

The use of this function allows the detection of most consecutive losses of COM cells as soon as the following COM or EOM cell of the CPAAL3/4_CS_PDU is received. If the number of COMs of a given CPAAL3/4_CS_PDU that is lost is an integer multiple of 16, the SAR_SN cannot detect them. Therefore, the use of the length field at the CS Sublayer is still required to detect any modulo 16 consecutive losses of CPAAL3/4_SAR_PDUs that may occur during situations like network congestion or protection switching events.

In addition, the use of this function will allow for immediate detection of most cases of cell insertion.

The use of Sequence_Number to detect situations in which two CPAAL3/4_CS_PDUs are inadvertently merged into one and the resulting length matches the length field in the CPAAL3/4_CS_PDU trailer is weak. This is due to the fact that this error event requires that the lengths of the original CPAAL3/4_CS_PDUs be the same. This implies that the same number of CPAAL3/4_SAR_PDUs will probably be required to transport two CPAAL3/4_CS_PDUs. Therefore, the SAR_SNs of the received CPAAL3/4_SAR_PDUs will probably be consecutive, and so the SAR Sublayer will not detect this error event. As a result, the use of the Etag at the CS Sublayer is still required.

Message Identification (MID) subfield. The 10-bit MID subfield is used to reassemble CPAAL3/4_SAR_PDUs into CPAAL3/4_CS_PDUs. All CPAAL3/4_SAR_PDUs of a given CPAAL3/4_CS_PDU will have the same MID. Note that this provides the basis for reassembly of discrete connectionless packets. Use of this subfield as the basis for a multiplexing or reassembly capability for connection-oriented services is for further study.

Payload_Length subfield. The 6-bit Payload_Length subfield is coded with the number of octets from the CPAAL3/4_CS_PDU that are included in the current CPAAL3/4_SAR_PDU. This number has a value between 0 and 44 inclusive. This subfield is binary coded with the most significant bit left-justified. BOM and COM cells take the value 44; EOM cells take the values 4, 8, ..., 44; SSM cells take the values 8, 12, ..., 44.

SAR_PDU_Payload. The CPAAL3/4_CS_PDU is left-justified in the SAR_PDU_Payload of the CPAAL3/4_SAR_PDU. Any part of the SAR_PDU_Payload that is not filled with CS information shall be coded as zeros.

Payload_CRC subfield. The 10-bit Payload_CRC subfield is filled with the value of a CRC calculation that is performed over the entire contents of the CPAAL3/4_SAR_PDU payload, including the SAR_PDU_Header, the SAR_PDU_Payload, and the SAR_PDU_Trailer. The CRC-10 generating polynomial has the capability of single-bit error correction over the CPAAL3/4_SAR_PDU. The following generator polynomial is used to calculate the Payload_CRC:

$$G(x) = x^{10} + x^3 + 1$$

The CRC remainder is placed in the CRC subfield with the most significant bit left-justified in the CRC subfield.

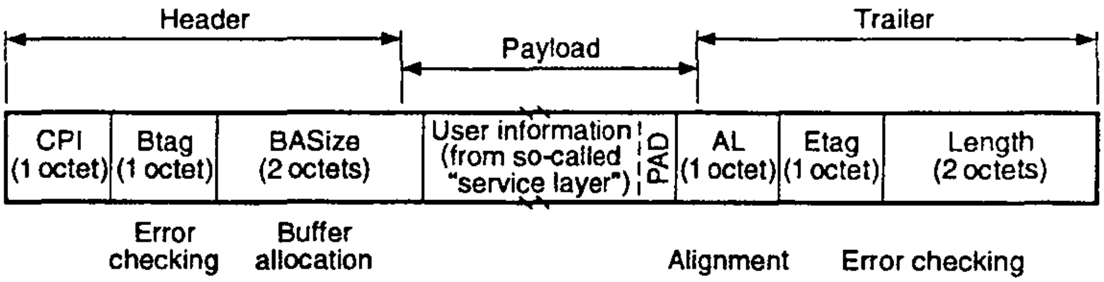

Figure 3.13 CPAAL3/4_CS_PDU Sublayer format of AAL.

3.4.5 Convergence Sublayer fields and format

Figure 3.13 depicts the Convergence Sublayer format of the AAL3/4.

There are two Adaptation Layer control fields: the CS_PDU_Header and the CS_PDU_Trailer, both of which are 4 octets long. The CS_PDU_Header and CS_PDU_Trailer encapsulate the user's protocol data units (USER_PDU). In addition, there may be from 0 to 3 pad octets added to align the CPAAL3/4_CS_PDU with a 32-bit boundary.

The CS_PDU_Header is subdivided into three fields: an 8-bit Common Part Indicator field, an 8-bit Beginning Tag (Btag) field, and a 16-bit Buffer Allocation size (BAsize) field. Likewise, the CS_PDU_Trailer is also subdivided into three fields: an 8-bit filler field, an 8-bit End Tag (Etag) field, and a 16-bit Length field.[6]

Common Part Indicator subfield. The 8-bit Common Part Indicator (CPI) subfield is used to identify the message type, i.e., to interpret subsequent fields for the CPAAL3/4-CS functions in the CPAAL3/4_CS_PDU header and trailer. It also indicates the counting unit for the values specified in the BAsize and Length fields.

CS_PDU Header—Btag subfield. For a given CPAAL3/4_CS_PDU, the same value appears in the 8-bit Btag field of the CS_PDU_Header and in the Etag field in the CS_PDU_Trailer. This allows the identification of a BOM segment and an EOM segment, and hence all intervening COM segments, as belonging to the same CPAAL3/4_CS_PDU. This correlation is required to implement segment loss detection over a CPAAL3/4_CS_PDU. As each CPAAL3/4_CS_PDU is transmitted, the Etag value is changed so that the entire range of Etag field values (0 to 255) is cycled through before reuse to aid in this segment loss protection.

BAsize subfield. The 16-bit Buffer Allocation size (BAsize) subfield is used to predict the buffer requirements for the CPAAL3/4_CS_PDU. Therefore, it must be greater than or equal to the true CPAAL3/4_CS_PDU length. This field is binary coded with the most significant bit left-justified in the subfield. If message-mode service is being provided, the BAsize value is encoded to be equal to the length of the USER_PDU field contained in the CPAAL3/4_CS_PDU Payload field. If streaming-mode service is being provided, the BAsize value is encoded to be equal to the maximum length of the CPAAL3/4_SDU.

USER_PDU field. The variable-length USER_PDU field contains user information. It contains the CPAAL3/4_SDU. It is octet aligned, as it is limited in length to the value of the BAsize field multiplied by the value of the counting unit (as identified in the CPI field).

Pad Field. The Pad field consists of 0, 1, 2, or 3 octets set to zero, so that the CPAAL3/4_CS_PDU is padded out to a 32-bit boundary.

AL. This 8-bit subfield is used to achieve 32-bit alignment in the CPAAL3/4_CS_PDU trailer. This is strictly a filler octet and does not contain any additional information.

Etag subfield. The 8-bit Etag subfield in the CPAAL3/4_CS_PDU trailer has the same value as the Btag subfield in the corresponding CPAAL3/4_CS_PDU header. As was mentioned earlier, the Btag and Etag subfields in the CS_PDU_Header and CS_PDU_Trailer are correlated in order to detect segment loss and misassembly. This field is binary coded with the most significant bit left-justified.

Length subfield. The 16-bit Length subfield specifies the length, in octets, of the USER_PDU (that is, the length of the user information contained in the CPAAL3/4_CS_PDU Payload field). This field is binary coded with the most significant bit left-justified in the subfield. It is used in conjunction with the Btag and Etag fields for the purpose of detecting misassembled CPAAL3/4_CS_PDUs.

3.5 AAL Type 5

The goal of the AAL Type 5 is to support, in the most streamlined fashion, those capabilities that are required to meet upper-layer data transfer over an ATM platform. The AAL Type 5 Common Part (CPAAL5) protocol provides for the transport of variable-length frames (1 to 65,535 octets) with error detection (the frame is padded to align the resulting PDU with an integral number of ATM cells). A length field is used to extract the frame and detect additional errors not detected with the CRC-32 mechanism. ANSI had a Letter Ballot for AAL Type 5 Common Part at press time, and ITU-TS had a draft version of I.363 (Section 6); approval was expected.

The Convergence Sublayer has been subdivided into the Common Part CS (CPCS) and the Service-Specific CS (SSCS), as shown in Fig. 3.14. Different SSCS protocols, to support specific AAL user services or groups of services, may be defined. The SSCS may also be null, in the sense that it provides only for the mapping of the equivalent primitives of the AAL to CPCS and vice versa. SSCS protocols are specified in separate Recommendations, not in, say, ITU-T I.363. This discussion

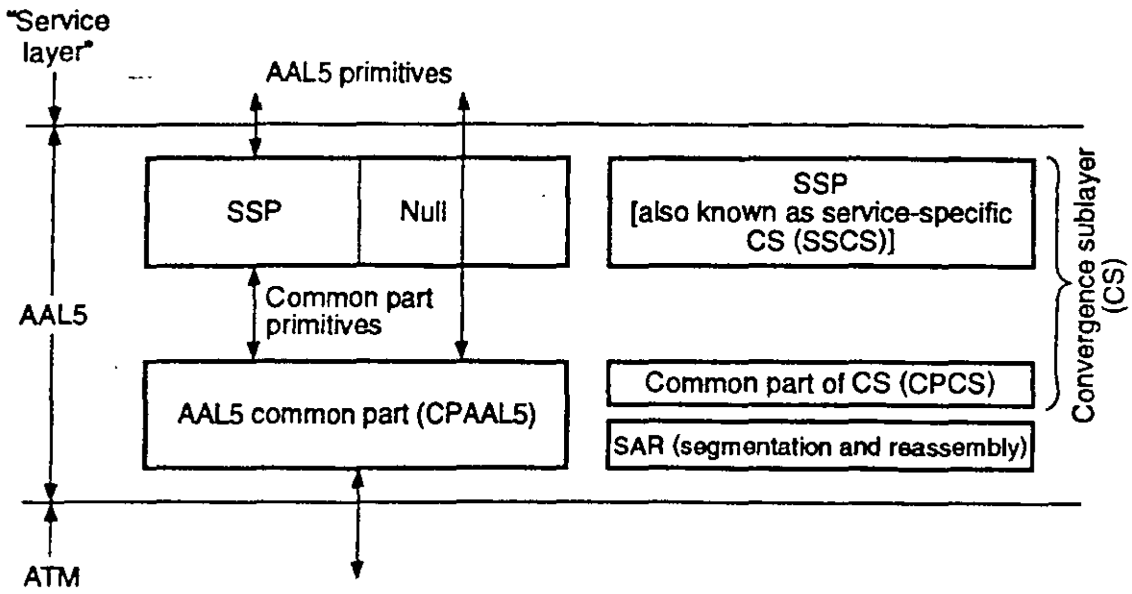

Figure 3.14 Structure of AAL Type 5.

therefore focuses on CPCS and SAR. Notice that CPAAL5 = SAR + CPCS. Also see Fig. 3.15.

3.5.1 Service provided by CPAAL5

The Common Part of AAL Type 5 provides the capability to transfer the CPAAL5_SDU from one CPAAL5 user to another CPAAL5 user through the ATM network. During this process, CPAAL5_SDUs may be corrupted or lost (in this case, an indication of the error is provided). Corrupted or

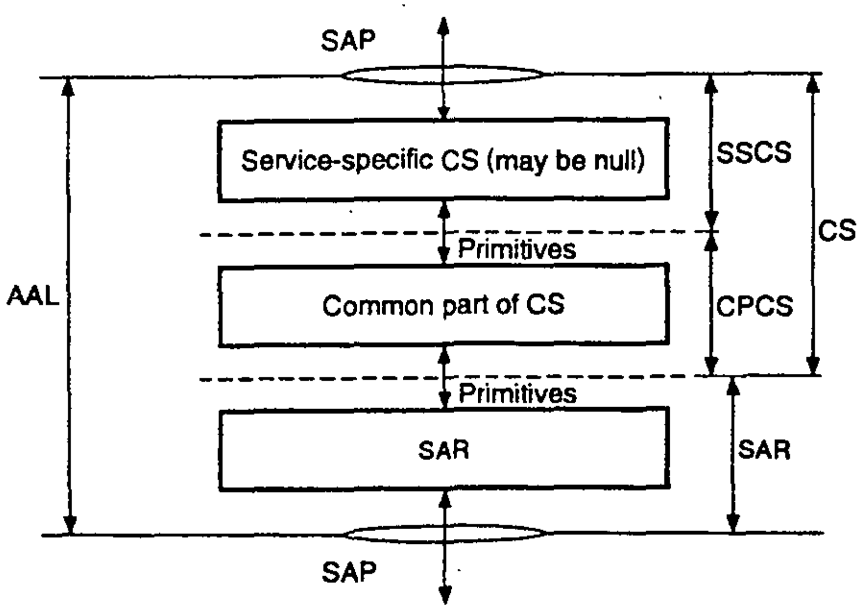

Figure 3.15 Another view of the structure of AAL Type 5.

lost CPAAL5_SDUs are not recovered by CPAAL5. CPAAL5 supports a message mode and a streaming mode. The message-mode service, streaming-mode service, and assured and nonassured operations as defined below for CPAAL5 are identical to those defined for AAL Type 3/4.

1. *Message-mode service.* The CPAAL5_SDU is passed across the CPAAL5 interface in exactly one Common Part AAL interface data unit (CPAAL5_IDU). This service provides the transport of fixed-size or variable-length CPAAL5_SDUs.
   a. In the case of small fixed-size CPAAL5_SDUs, an internal blocking/deblocking function in the SSCS may be applied; it provides the transport of one or more fixed-size CPAAL5_SDUs in one SSCS_PDU.
   b. In the case of variable-length CPAAL5_SDUs, an internal CPAAL5_SDU message segmentation/reassembling function in the SSCS may be applied. In this case, a single CPAAL5_SDU is transferred in one or more SSCS_PDUs.
   c. Where the above options are not used, a single CPAAL5_SDU is transferred in one SSCS_PDU. When the SSCS is null, the CPAAL5_SDU is mapped to one CPCS_SDU.
2. *Streaming-mode service.* The CPAAL5_SDU is passed across the CPAAL5 interface in one or more CPAAL5_IDUs. The transfer of these CPAAL5_IDUs across the CPAAL5 interface may occur separated in time. This service provides the transport of variable-length CPAAL5_SDUs. Streaming-mode service includes an abort service by which the discarding of an CPAAL5_SDU that has been partially transferred across the AAL interface can be requested.
   a. An internal CPAAL5_SDU message segmentation/reassembling function in the SSCS may be applied. In this case, all the CPAAL5_IDUs belonging to a single CPAAL5_SDU are transferred in one or more SSCS_PDUs.
   b. An internal pipelining function may be applied. It provides the means by which the sending CPAAL5 entity initiates the transfer to the receiving CPAAL5 entity before it has the complete CPAAL5_SDU available.
   c. Where option *a* is not used, all the CPAAL5_IDUs belonging to a single CPAAL5_SDU are transferred in one SSCS_PDU. When the SSCS is null, the CPAAL5_IDUs belonging to a single CPAAL5_SDU are mapped to one CPCS_SDU.

Both modes of service may offer the following peer-to-peer operational procedures:

- *Assured operations.* Every assured CPAAL5_SDU is delivered with exactly the data content that the user sent. The assured service is provided by retransmission of missing or corrupted SSCS_PDUs. Flow control is provided as a mandatory feature. The assured operation may be restricted to point-to-point AAL connections.
- *Nonassured operations.* Integral CPAAL5_SDUs may be lost or corrupted. Lost and corrupted CPAAL5_SDUs will not be corrected by retransmission. An optional feature may be provided to allow corrupted CPAAL5_SDUs to be delivered to the user (i.e., optional error discard). Flow control may be provided as an option.

Description of AAL connections. The CPAAL5 provides the capability to transfer the CPAAL5_SDU from one AAL5-SAP to another AAL5-SAP through the ATM network. CPAAL5 users have the ability to select a given AAL5-SAP associated with the QOS required to transport that CPAAL5_SDU (for example, delay- and loss-sensitive QOS).

The CPAAL5 in nonassured operation also provides the capability to transfer the CPAAL5_SDUs from one AAL5-SAP to more than one AAL5-SAP through the ATM network.

CPAAL5 makes use of the service provided by the underlying ATM Layer. Multiple AAL connections may be associated with a single ATM-Layer connection, allowing multiplexing at the AAL; however, if multiplexing is used in the AAL, it occurs in the SSCS. The AAL user selects the QOS provided by the AAL through the choice of the AAL5-SAP used for data transfer.

Primitives for the AAL. These primitives are service-specific and are contained in separate Recommendations on SSCS protocols.

The SSCS may be null, in the sense that it provides only for the mapping of the equivalent primitives of the AAL to CPCS and vice versa. In this case, the primitives for the AAL are equivalent to those for the CPCS but are identified as CPAAL5-UNITDATA.request, CPAAL5-UNITDATA.indication, CPAAL5-U-Abort.request, CPAAL5-U-Abort.indication, and CPAAL5-P-Abort.indication, consistent with the primitive naming convention at an SAP.

Primitives for the CPCS of the AAL. As there is no SAP between the sublayers of the AAL5, the primitives are called .invoke and .signal instead of the conventional .request and .indication to highlight the absence of the SAP.

CPCS-UNITDATA.invoke and CPCS-UNITDATA.signal. These primitives are used for data transfer. The following parameters are defined:

- Interface data (ID). This parameter specifies the interface data unit exchanged between the CPCS and the SSCS entity. The ID is an integral multiple of 1 octet. If the CPCS entity is operating in message-mode service, the ID represents a complete CPCS_SDU; when operating in streaming-mode service, the ID does not necessarily represent a complete CPCS_SDU.
- More (M). In message-mode service, this parameter is not used. In streaming-mode service, this parameter specifies whether the interface data communicated contains a beginning/continuation of a CPCS_SDU or the end of a complete CPCS_SDU.
- CPCS loss priority (CPCS-LP). This parameter indicates the loss priority for the associated CPCS_SDU. It can take only two values, one for high priority and the other for low priority. The use of this parameter in streaming mode is for further study. This parameter is mapped to and from the SAR-LP parameter.
- CPCS congestion indication (CPCS-CI). This parameter indicates that the associated CPCS_SDU has experienced congestion. The use of this parameter in streaming mode is for further study. This parameter is mapped to and from the SAR-CI parameter.
- CPCS user-to-user indication (CPCS-UU). This parameter is transparently transported by the CPCS between peer CPCS users.
- Reception status (RS). This parameter indicates that the associated CPCS_SDU delivered may be corrupted. This parameter is utilized only if the corrupted data delivery option is used.

Depending on the service mode (message- or streaming-mode service, discarding or delivery of errored information), not all parameters are required.

CPCS-U-Abort.invoke and CPCS-U-Abort.signal. These primitives are used by the CPCS user to invoke the abort service. They are also used to signal to the CPCS user that a partially delivered CPCS_SDU is to be discarded by instruction from its peer entity. No parameters are defined. These primitives are not used in message mode.

CPCS-P-Abort.signal. This primitive is used by the CPCS entity to signal to its user that a partially delivered CPCS_SDU is to be discarded because of the occurrence of some error in the CPCS or below. No parameters are defined. This primitive is not used in message mode.

Primitives for the SAR sublayer of the AAL. These primitives model the exchange of information between the SAR sublayer and the CPCS.

As there is no SAP between the sublayers of the AAL5, the primitives are called .invoke and .signal instead of the conventional .request and .indication to highlight the absence of the SAP.

SAR-UNITDATA.invoke and SAR-UNITDATA.signal. These primitives are used for data transfer. The following parameters are defined:

- Interface data (ID). This parameter specifies the interface data unit exchanged between the SAR and the CPCS entity. The ID is an integral multiple of 48 octets. It does not necessarily represent a complete SAR_SDU.
- More (M). This parameter specifies whether the interface data communicated contains the end of the SAR_SDU.
- SAR loss priority (SAR-LP). This parameter indicates the loss priority for the associated SAR interface data. It can take on two values, one for high priority and the other for low priority. This parameter is mapped to the ATM Layer's submitted loss priority parameter and from the ATM Layer's received loss priority parameter.
- SAR congestion indication (SAR-CI). This parameter indicates whether the associated SAR interface data has experienced congestion. This parameter is mapped to and from the ATM Layer's congestion indication parameter.

3.5.2 Functions, structure, and coding of AAL5

Functions of the SAR Sublayer. The SAR Sublayer functions are performed on an SAR_PDU basis. The SAR Sublayer accepts variable-length SAR_SDUs which are integral multiples of 48 octets from the CPCS and generates SAR_PDUs containing 48 octets of SAR_SDU data. It supports the preservation of SAR_SDUs by providing for an "end of SAR_SDU" indication.

SAR_PDU structure and coding. The SAR Sublayer function utilizes the ATM-Layer-user-to-ATM-Layer-user (AUU) parameter of the ATM Layer primitives to indicate that a SAR_PDU contains the end of a

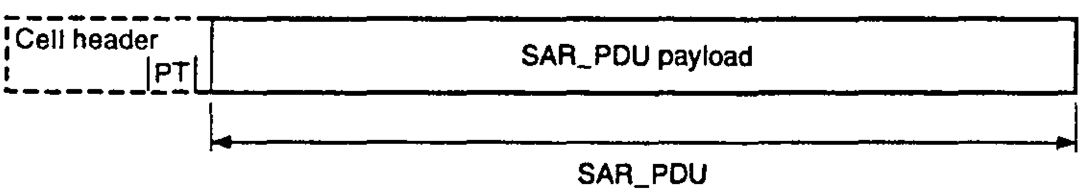

Figure 3.16 SAR_PDU format for AAL5. [*Note:* The payload type (PT) field belongs to the ATM header. It conveys the value of the AUU parameter end-to-end.]

SAR_SDU. A SAR_PDU where the value of the AUU parameter is 1 indicates the end of a SAR_SDU; a value of 0 indicates the beginning or continuation of a SAR_SDU. The structure of the SAR_PDU is shown in Fig. 3.16.

Convergence Sublayer. The CPCS has the following service characteristics.

- Nonassured data transfer of user data frames with any length measured in octets from 1 to 65,535 octets.
- The CPCS connection will be established by management or by the Control Plane.
- Error detection and indication (bit error and cell loss or gain).
- CPCS_SDU sequence integrity on each CPCS connection.

Functions of the CPCS. The CPCS functions are performed per CPCS_PDU. The CPCS provides several functions in support of the CPCS service user. The functions provided depend on whether the CPCS service user is operating in message or streaming mode.

1. *Message mode service.* The CPCS_SDU is passed across the CPCS interface in exactly one CPCS-IDU. This service provides the transport of a single CPCS_SDU in one CPCS_PDU.
2. *Streaming mode service.* The CPCS_SDU is passed across the CPCS interface in one or more CPCS-IDUs. The transfer of these CPCS-IDUs across the CPCS interface may occur separated in time. This service provides the transport of all the CPCS-IDUs belonging to a single CPCS_SDU into one CPCS_PDU. An internal pipelining function in the CPCS may be applied which provides the means by which the sending CPCS entity initiates the transfer to the receiving CPCS entity before it has the complete CPCS_SDU available. Streaming-mode service includes an abort service by which the discarding of a CPCS_SDU partially transferred across the interface can be requested.

Note: At the sending side, parts of the CPCS_PDU may have to be buffered if the restriction "interface data are a multiple of 48 octets" cannot be satisfied.

The functions implemented by the CPCS include:

1. *Preservation of CPCS_SDU.* This function provides for the delineation and transparency of CPCS_SDUs.
2. *Preservation of CPCS user-to-user information.* This function provides for the transparent transfer of CPCS user-to-user information.

3. *Error detection and handling.* This function provides for the detection and handling of CPCS_PDU corruption. Corrupted CPCS_SDUs are either discarded or optionally delivered to the SSCS. The procedures for delivery of corrupted CPCS_SDUs are for further study. When delivering errored information to the CPCS user, an error indication is associated with the delivery. Examples of detected errors would include received length and CPCS_PDU Length field mismatch including buffer overflow, an improperly formatted CPCS_PDU, and CPCS CRC errors.
4. *Abort.* This function provides for the means to abort a partially transmitted CPCS_SDU. This function is indicated in the Length field.
5. *Padding.* A padding function provides for 48-octet alignment of the CPCS_PDU trailer.

CPCS structure and coding. The CPCS functions require an 8-octet CPCS_PDU trailer. The CPCS_PDU trailer is always located in the last 8 octets of the last SAR_PDU of the CPCS_PDU. Therefore, a padding field provides for a 48-octet alignment of the CPCS_PDU. The

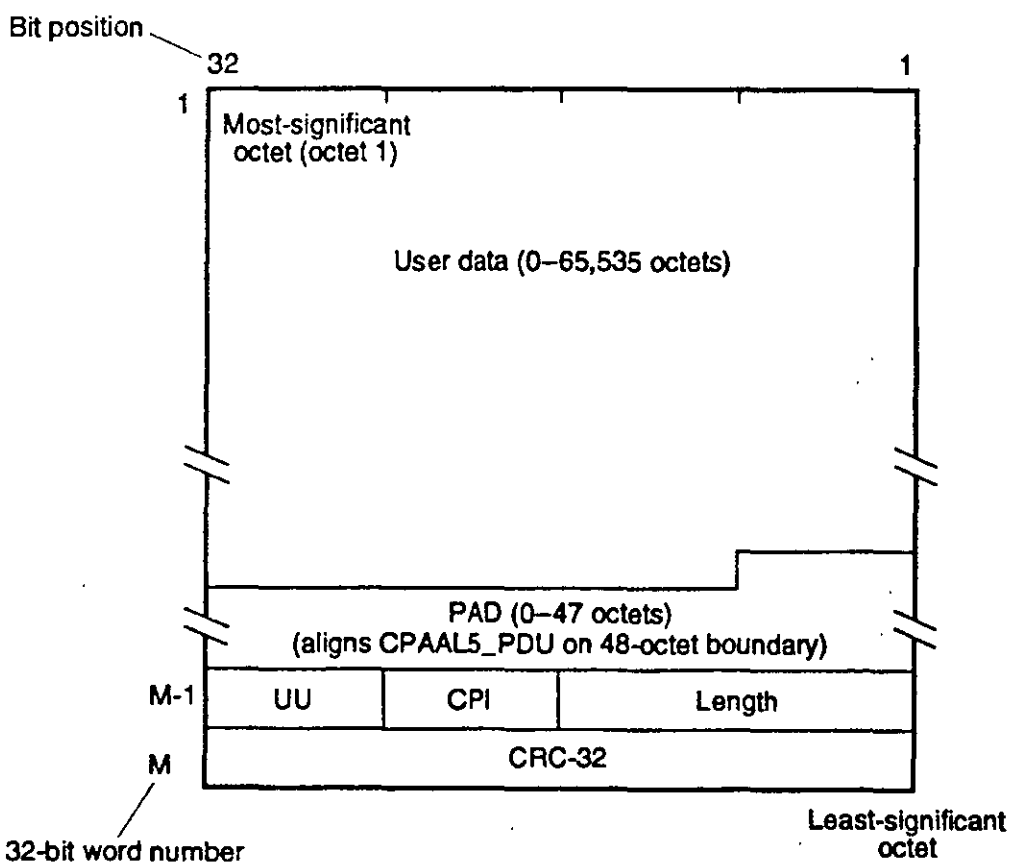

Figure 3.17 CPAAL5_PDU.

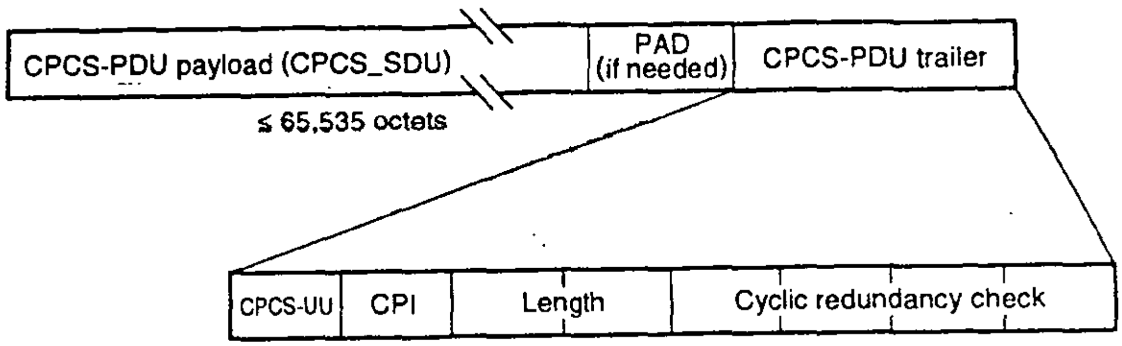

Figure 3.18 CPAAL5_PDU, another view.

CPCS_PDU trailer, the padding field, and the CPCS_PDU payload make up the CPCS_PDU.

The coding of the CPCS_PDU conforms to the coding conventions specified in 2.1 of Recommendation I.361. See Figs. 3.17 and 3.18.

1. *CPCS_PDU payload.* The CPCS_PDU payload is the CPCS_SDU.
2. *Padding (Pad) field.* Between the end of the CPCS_PDU payload and the CPCS_PDU trailer, there will be from 0 to 47 unused octets. These unused octets are called the padding (Pad) field; they are strictly used as filler octets and do not convey any information. Any coding is acceptable. This padding field complements the CPCS_PDU (including CPCS_PDU payload, padding field, and CPCS_PDU trailer) to an integral multiple of 48 octets.
3. *CPCS User-to-User Indication (CPCS-UU) field.* The CPCS-UU field is used to transparently transfer CPCS user-to-user information.
4. *Common Part Indicator (CPI) field.* One of the functions of the CPI field is to align the CPCS_PDU trailer to 64 bits. Other functions are for further study. Possible additional functions may include identification of Layer Management messages. When only the 64-bit alignment function is used, this field is coded as zero.
5. *Length field.* The Length field is used to encode the length of the CPCS_PDU payload field. The Length field value is also used by the receiver to detect the loss or gain of information. The length is binary coded as number of octets. A Length field coded as zero is used for the abort function.
6. *CRC field.* The CRC-32 is used to detect bit errors in the CPCS_PDU. The CRC field is filled with the value of a CRC calculation which is performed over the entire contents of the CPCS_PDU, including the CPCS_PDU payload, the Pad field, and the first 4 octets of the CPCS_PDU trailer. The CRC field shall contain the 1s complement of the sum (modulo 2) of

a. The remainder of $xk^*(x^{31} + x^{30} + \ldots + x + 1)$ divided (modulo 2) by the generator polynomial, where k is the number of bits of the information over which the CRC is calculated.

b. The remainder of the division (modulo 2) by the generator polynomial of the product of x^{32} and the information over which the CRC is calculated.

The CRC-32 generator polynomial is:

$$G(x) = x^{32} + x^{26} + x^{23} + x^{22} + x^{16} + x^{12} + x^{11} + x^{10} + x^{8} + x^{7} + x^{5} + x^{4} + x^{2} + x + 1$$

The result of the CRC calculation is placed with the least significant bit right-justified in the CRC field.

As a typical implementation at the transmitter, the initial content of the register of the device computing the remainder of the division is preset to all 1s and is then modified by division by the generator polynomial (as described above) of the information over which the CRC is to be calculated; the 1s complement of the resulting remainder is put into the CRC field.

As a typical implementation at the receiver, the initial content of the register of the device computing the remainder of the division is preset to all 1s. The final remainder, after multiplication by x^{32} and then division (modulo 2) by the generator polynomial of the serial incoming CPCS_PDU, will be (in the absence of errors)

$$C(x) = x^{31} + x^{30} + x^{26} + x^{25} + x^{24} + x^{18} + x^{15} + x^{14} + x^{12} + x^{11} + x^{10} + x^{8} + x^{6} + x^{5} + x^{4} + x^{3} + x + 1$$

3.6 Frame Relay Service-Specific AAL

The Frame Relay Service-Specific ATM Adaptation Layer Convergence Sublayer (FR-SSCS) is positioned in the upper part of the ATM Adaptation Layer; it is located above the CPAAL5, as shown in Figs. 3.19 and 3.20. It is an example of an SSP. The purpose of the FR-SSCS protocol at an ATM CPE (that is, user's equipment) is to emulate the Frame Relaying Bearer Service (FRBS) in an ATM-based network (Fig. 3.19). On network nodes, the FR-SSCS is used for interworking between an ATM-based network and a Q.922-based Frame Relaying Network (Fig. 3.20).

The FR-SSCS protocol provides for the transport of variable-length frames with error detection.* The FR-SSCS provides its service over

*This discussion is based on Ref. 4.

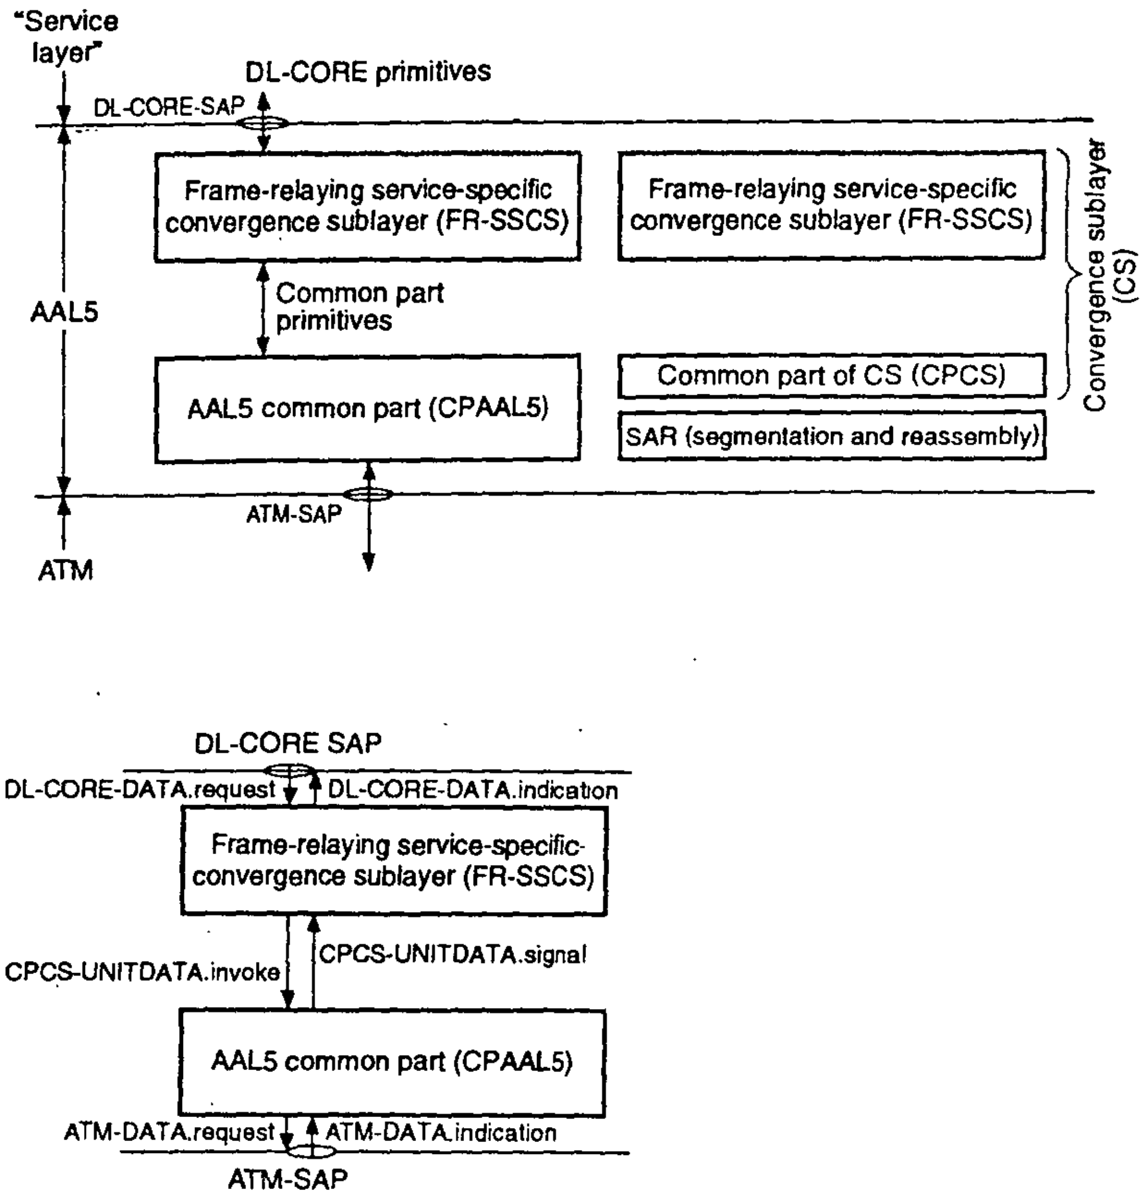

Figure 3.19 AAL5 for interworking of frame relay and ATM (in CPE).

preestablished connections with negotiated traffic parameters. An FR-SSCS connection represents the segment of an end-to-end frame relay (FR) connection over B-ISDN. At an ATM-based B-TE, the FR-SSCS connection is terminated at the point of termination of the FR-SSCS service and represents one end of the FR connection. Optionally, multiplexing may be performed at the FR-SSCS, allowing various FR-SSCS connections to be associated with a single CPAAL5 connection (and with the corresponding ATM connection). FR-SSCS connections within a CPAAL5 connection are uniquely identified by data link connection identifiers (DLCIs). The establishment (or provisioning) and initialization of an FR-SSCS connection is performed by interaction with FR-SSCS Layer Management (MFR-SSCS) entities. The traffic parameters of each FR-SSCS connection are determined at the time of its estab-

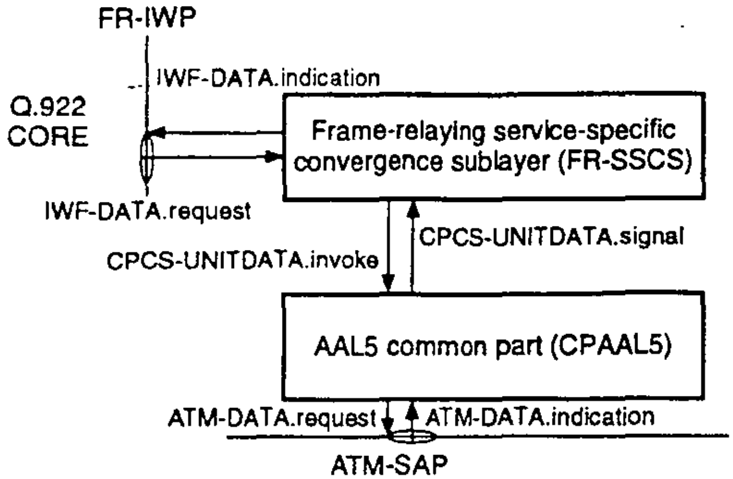

Figure 3.20 AAL5 for interworking of frame relay and ATM (in a network element supporting an interworking function). FR-IWP = frame relay interworking point.

lishment. The negotiated traffic parameters are bounded by the ATM Layer/CPAAL5 connection characteristics.

The FR-SSCS can indicate to its user that the receiver FR_SSCS_PDU has experienced congestion (forward congestion) or that an FR_SSCS_PDU traveling in the opposite (sending) direction has experienced congestion (backward congestion). The FR-SSCS allows for two discard eligibility priorities. The FR-SSCS user can request the discard eligibility (loss priority) associated with each FR_SSCS_SDU. The FR-SSCS uses the CPAAL5 message-mode service without the corrupted data delivery option and preserves the FR_SSCS_SDU sequence integrity.

The MFR-SSCS is responsible for the following actions: assignment of the FR_SSCS association necessary for the establishment or provisioning of FR-SSCS connections between peer FR-SSCS entities, resetting the parameters and state variables associated with a FR-SSCS connection when required, releasing the association created for a FR-SSCS connection between peer FR-SSCS entities, and performance monitoring of the quality of the FR-SSCS connection service provided through notification of errors (i.e., FR_SSCS_PDU discards resulting from errors in the FR_SSCS_PDU).

Service provided by the FR-SSCS. The FR-SSCS provides services to (1) the core service user (upper layer) at ATM-based B-TEs or (2) the Q.922-CORE Data Link Layer (Q.922-DLL) on network nodes at interworking functions (IWFs) points. Only item (1) is covered here.

The FR-SSCS provides the capability to transfer variable-length octet-aligned FR_SSCS_SDUs from one or more FR_SSCS users. The FR-SSCS Sublayer preserves the FR_SSCS_SDU sequence integrity within an FR-SSCS connection. During this process, FR_SSCS_SDUs may be lost or corrupted. Lost or corrupted FR_SSCS_SDUs are not recovered by the FR-SSCS. The FR-SSCS uses CPAAL5 message-mode service without the corrupted data delivery option.

FR-SSCS functions. The functions provided by the FR_SSCS include

Multiplexing/demultiplexing. This function provides for the optional multiplexing and demultiplexing of FR-SSCS connections into a single CPAAL5 connection. The number of FR-SSCS connections supported over a CPAAL5 connection is defined at connection establishment or provisioning. The default number of FR-SSCS connections when multiplexing is not supported is 1. Within a given FR-SSCS connection, sequence integrity is preserved.

Inspection of the FR_SSCS_PDU length. This function inspects the FR_SSCS_PDU to ensure that it consists of an integral number of octets and to ensure that it is neither too long nor too short.

Congestion control. These functions provide the means to notify the end user that congestion avoidance procedures should be initiated,

TABLE 3.6 DL-CORE Primitives

DL-CORE-DATA.request (DL_CORE_User_Data, Discard_Eligibility, DL_CORE_Service_User_Protocol_Control_Information)	This primitive is received from the FR-SSCS user to request the transfer of an FR_SSCS_SDU over the associated FR-SSCS connection.
DL-CORE-DATA.indication (DL_CORE_User_Data, Congestion_Encountered_Backward, Congestion_Encountered_Forward, DL_CORE_Service_User_Protocol_Information)	This primitive is used to the FR-SSCS user to indicate the arrival of an FR_SSCS_SDU from the associated connection.

Description of parameters:

DL_CORE_User_Data: This parameter specifies the FR_SSCS_SDU transported between the FR-SSCS user and the FR-SSCS. This parameter is octet-aligned and can range from 1 to a maximum of at least 4096 octets in length.

Discard_Eligibility: This parameter indicates the loss priority assigned to the FR_SSCS_SDU. Two levels of priority are identified: High and Low. A value of High indicates that the FR_SSCS_SDU may experience a better quality of service with respect to loss (i.e., minimal loss) than if the Discard_Eligibility parameter were set to Low.

DL_CORE_Service_Protocol_Information: This parameter specifies a 1-bit FR-SSCS/Q.922-DLL user control information to be transparently transferred between FR-SSCS/Q.922-DLL users.

Congestion_Encountered_Backward: This parameter indicates that an FR_SSCS_SDU has experienced congestion in the opposite (sending) direction, and therefore that an FR_SSCS_SDU sent on the corresponding connection may encounter congested resources. This parameter may take on two values: True or False. A value of True indicates that an FR_SSCS_SDU has experienced congestion in the opposite (sending) direction of the connection.

Congestion_Encountered_Forward: This parameter indicates that the received FR_SSCS_SDU has experienced congestion. This parameter may take two values: True or False. A value of True indicates that the FR_SSCS_SDU has experienced congestion.

where applicable (congestion control forward and congestion control backward). In addition, the functions provide the means for the end user and/or the network to indicate what frames should be discarded in a congestion situation.

Primitives. The information exchanged between the FR-SSCS and its user (for ATM-based B-TEs) is modeled by the primitives of Table 3.6 (which are the same DL-CORE primitives in Annex C of ITU-T Recommendation I.233.1).

Services expected from the CPAAL5. The FR-SSCS expects the CPAAL5 to provide the capability to transfer variable-length (from 3 to a maximum of at least 4100 octets) octet-aligned FR_SSCS_SDUs, with error detection and in sequence, between communicating FR-SSCS entities. Lost or corrupted FR_SSCS_PDUs are not expected to be recovered by the CPAAL5. Multicast services, derived from the ATM Layer, are expected.

The FR-SSCS entity expects the CPAAL5 to provide each FR_SSCS_PDU (CPAAL5_SDU) with the CP_Congestion_Indication (True or False) set to the value of the Congestion_Indication received by the ATM Layer with the last ATM_SDU conforming to the CPAAL5_SDU; and with the CP_Loss_Priority set to either Low, if any of the ATM_SDUs conforming to the CPAAL5_SDU was received with the Received_Loss_Priority parameter set to Low, or High otherwise.

The FR-SSCS entity passes each FR_SSCS_PDU (CPAAL5_SDU) with the CP_Loss_Priority set to the value of the Discard_Eligibility parameter received from the upper layer or the Q.922-DLL (High or Low), the CP_Congestion_Indication (True or False) always set to False, and the User_User_Indication parameter always set to zero.

3.7 Signaling ATM Adaptation Layer (SAAL)

This section describes the Signaling ATM Adaptation Layer (SAAL) for use at the UNI. SAAL is used in the Control Plane. (This topic could also have been treated in the next chapter, but it was decided to include it here with other AALs.)

The SAAL resides between the ATM Layer and Q.2931 in the user's equipment, specifically in the software implementing the Control Plane (i.e., the signaling capability). The purpose of the SAAL is to provide reliable transport of Q.2931 messages between peer Q.2931 entities (e.g., ATM switch and host) over the ATM Layer. The SAAL is composed of two sublayers, a Common Part and a Service-Specific

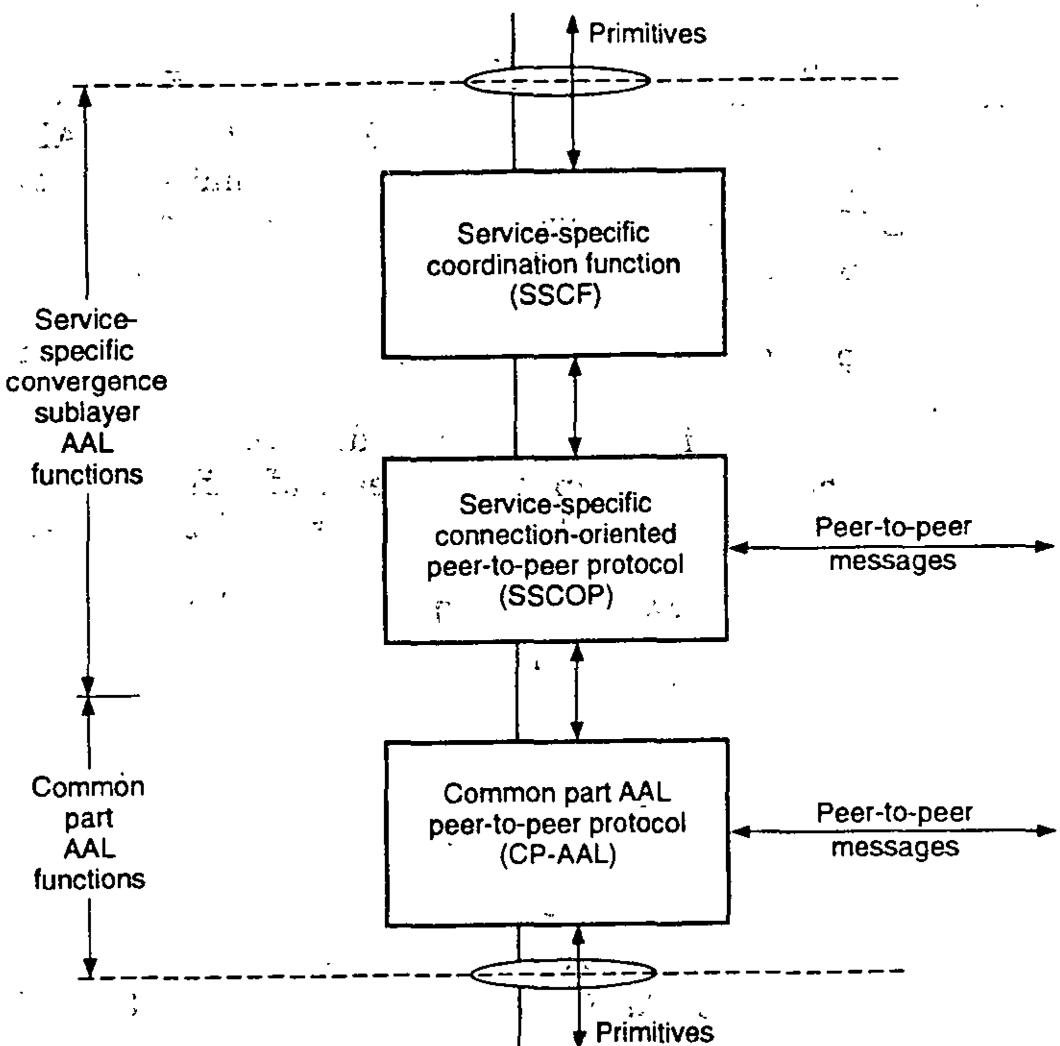

Figure 3.21 SAAL structure. (*Note:* This figure represents the allocation of functions and is not intended to illustrate sublayers as defined by OSI modeling principles.)

Part. The Service-Specific Part is further subdivided into a Service-Specific Coordination Function (SSCF) and a Service-Specific Connection-Oriented Protocol (SSCOP). Figure 3.21 illustrates the structure of the SAAL.[5]

The SAAL for supporting signaling uses the protocol structure illustrated in Fig. 3.21. The Common Part AAL protocol provides unassured information transfer and a mechanism for detecting corruption of SDUs. The AAL Type 5 Common Part protocol is used to support signaling. The AAL Type 5 Common Part protocol is specified in Draft Recommendation I.363.

The SAAL for supporting signaling at the UNI uses the AAL Type 5 Common Part protocol, discussed above, as specified in Ref. 7 with minor amendments.[8]

The Service-Specific Connection-Oriented Protocol (SSCOP) resides in the Service-Specific Convergence Sublayer (SSCS) of the SAAL. SSCOP is used to transfer variable-length service data units (SDUs) between users of SSCOP. SSCOP provides for the recovery of lost or corrupted SDUs. SSCOP is specified in ITU-T Recommendation Q.2110.[9]

The SAAL for supporting signaling utilizes SSCOP as specified in Q.2110.[9]

An SSCF maps the service of SSCOP to the needs of the SSCF user. Different SSCFs may be defined to support the needs of different AAL users. The SSCF used to support Q.93B at the UNI is specified in ITU-T Recommendation Q.2130.[10]

The external behavior of the SAAL at the UNI appears as if the UNI SSCF specified in Q.2130[10] were implemented.

References

1. CCITT I.362, *B-ISDN AAL Functional Description*, Geneva, Switzerland, 1992.
2. CCITT I.363, *B-ISDN AAL Specification*, Geneva, Switzerland, 1992.
3. ANSI T1.BCR-199x, *Broadband ISDN—ATM Adaptation Layer for Constant Bit Rate Services Functionality Specification*, New York, Nov. 13, 1992.
4. Bellcore, *Asynchronous Transfer Mode (ATM) and ATM Adaptation Layer (AAL) Protocols Generic Requirements*, TA-NWT-001113, Issue 2, July 1993.
5. B. Kittams, Bellcore, personal communication, May 1993.
6. T1S1.5/93-52, *Broadband Aspects of ISDN Baseline Document*, T1S1 Technical Subcommittee, August 1990, Chief Editor: Rajeev Sinha. Reissued February 1993, Chief Editor: Erwin Fandrich.
7. CCITT Document TD-XVIII/10 (AAL5), "AAL Type 5, Draft Recommendation Text for Section 6 of I.363," Geneva, Switzerland, Jan. 29, 1993.
8. ATM Forum, *ATM User-Network Interface Specification*, Version 3.0, August 1993.
9. ITU Document DT/11/3-28 [Q.SAAL1 (now Q.2110)], *Service-Specific Connection-Oriented Protocol (SSCOP) Specification*, Geneva, Switzerland, May 17, 1993.
10. ITU Document DT/11/3-XX [Q.SAAL2 (now Q.2130)], *Service-Specific Connection-Oriented Protocol (SSCOP) Specification*, Geneva, Switzerland, May 17, 1993.