Asynchronous Transfer Mode (ATM)

- Background
- Physical and ATM Layers
- AALs
- Applications

ATM Quick Highlights

- The Telecom Industry's thrust into multi-media data networking
- Comm unit is small, fixed-sized "cell" (53 bytes)
- Built to provide Quality-of-service
- Connection-oriented
- Designed to run over SONET/SDH

ATM History

- 1980 — ISDN
- 1985 — B-ISDN Study Group; ITU chooses ATM approach B-ISDN
- 53-byte cell chosen
- 1990 — ATM Forum founded
- 1995 — Anchorage Accord Stabilizes Specifications
- 2000

The B-ISDN ATM Reference Model

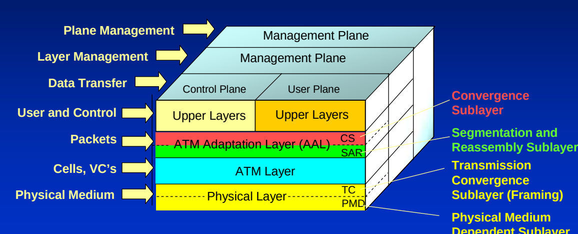

ATM
- Background

Why?

- Why a small cell instead of a large packet?
  - Queue delays tend to grow as packet size grows. A small cell helps maintain streamlined flows.
  - No/little performance loss due to padding large fields
  - Small cells better for voice
  - No need for in-route fragmentation

ATM
- Background

Why?

- Why a fixed cell size instead of variable-size packets?
  - Switch architecture can be optimized to the fixed size, so switching can be done in hardware
  - Scalable parallel switch designs

ATM
- Background

Why?

- Why 53 bytes?
  - US wanted 64 payload bytes, Europe wanted 32
  - Compromised on 48
  - +5 header = 53

ATM
- Background

Why?

- Why start out with 9% Overhead?

Overhead isn't everything…

Ethernet / SS10: **9 Mb/s BW,** **900 μsecs ovhd**
ATM Synoptics: **78 Mb/s BW,** **1,250 μsecs ovhd.**
NFS trace over 1 week: 95% msgs < 200 bytes

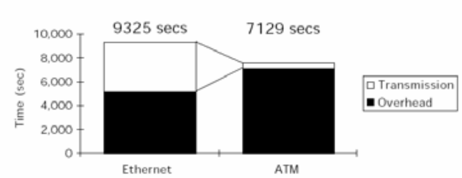

Figure from CMU Test Run 1997http://www-2.cs.cmu.edu/afs/cs.cmu.edu/academic/class/15347-s98/public/lectures/lect24.pdf

ATM + and –

+
- QoS
- Multimedia Support
- Hardware Switching -> High Speed
- Connection-Oriented (-?)

-
- IP Support
- LAN arena dominated by huge installed Ethernet base
- Ethernet growing toward MAN, WAN
- Connection-Oriented (+?)
- Living up to the hype of the early 90's

ATM Architecture

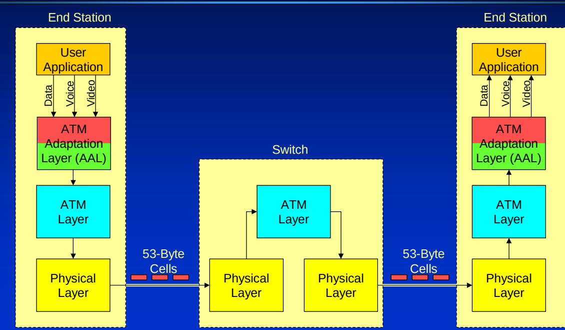

ATM

- Background
- Physical and ATM Layers
  - Cells, Formats, and Addressing
  - Virtual Circuits
  - Switches and Media
  - Interfaces
- AALs
- Applications

Cell Format

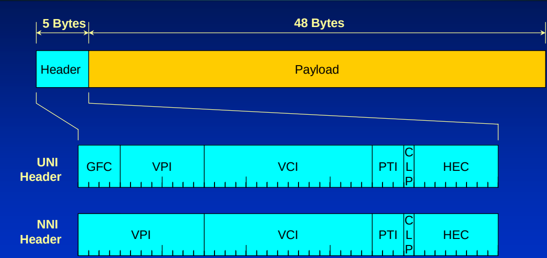

ATM
- Physical and ATM Layers

Header Error Control (HEC)

HEC covers the header only, not the payload -- the goal is to ensure correct delivery

First Four Cell Header Bytes

bbbbbbbb	bbbbbbbb	bbbbbbbb	bbbbbbbb

/ $x^8 + x^2 + x + 1$ ➡ **Remainder**
+ <u>01010101</u>
HEC

If P(bit error) = 10^{-8}, then P(Undetected header error) = $\sim 10^{-20}$:
at OC-3, about 1 per 90,000 yrs

HEC also assists in synchronizing:
- **Look at 53-byte sequences until you find one where the HEC field works correctly**
- **If this holds up for D sequences in a row, assume you've synched**

- **p(bad synch) = $\sim 2^{-8D}$**

ATM
- Physical and ATM Layers

UNI Header Fields

GFC	VPI	VCI	PTI	CLP	HEC

GFC - General Flow Control
Only used between host and network. Overwritten by first switch.

VPI - Virtual Path ID

VCI - Virtual Circuit ID

PTI - Payload type ID

CLP - Cell Priority (ID's cells for deletion when congestion experienced)

HEC - Header Checksum (all 1-bit errors corrected, 90% of multi-bit errors detected)

ATM
- Physical and ATM Layers

NNI Header Fields

VPI	VCI	PTI	CLP	HEC

VPI - Virtual Path ID

VCI - Virtual Circuit ID

PTI - Payload type ID

CLP - Cell Priority (used to ID cells for deletion when congestion experienced)

HEC - Header Checksum (all 1-bit errors corrected, 90% of multi-bit errors detected)

ATM
- Physical and ATM Layers

PTI Field Codes

PTI	Meaning	
000	User Data Cell Type 0	No congestion experienced
001	User Data Cell Type 1	No congestion experienced
010	User Data Cell Type 0	Congestion experienced
011	User Data Cell Type 1	Congestion experienced
100	Maintenance info between adjacent switches	
101	Maintenance info between source and destination switches	
110	Resource management cell (for ABR congestion control)	
111	Reserved	

Used by AAL5 to denote end of message

0	b	b

Explicit Forward Congestion Indicator (EFCI) Set by Congested switch

ATM Addressing

E.164: Telephone Numbers - up to 15 digits

ATM End-System Addresses (AESA) -- 20-byte Addresses

Address Type / US Code / Format indicator / End-System Sub-address

Example Format

39	840F	80	Org. ID	Specified by Organization	End-System Address	

Address Aggregation (Hierarchical Routing)

39.840F.80.809134.0001 *BigCorp, Inc.*

Plant 1 39.840F.80.809134.0001.01

39.840F.80.809134.0001.0F *Plant 15*

39.840F.80.809134.0001.0F.01 *Bldg 1*

39.840F.80.809134.0001.0F.02 *Bldg 2*

ATM

- Background
- Physical and ATM Layers
  - Cells, Formats, and Addressing
  - Virtual Circuits
  - Switches and Media
  - Interfaces
- AALs
- Applications

Virtual Paths, Virtual Circuits

A Link can Carry Multiple Virtual Paths, which can Carry Multiple Virtual Circuits

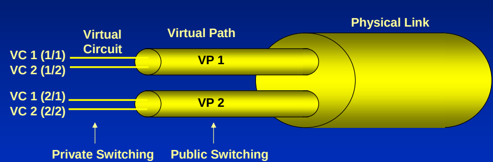

Advantage of Two-Level Approach:
A Bundle of VC's can be switched by switching a single VP

Virtual Circuits

- Normally Unicast, but one-way Multicasting Supported
- Unidirectional, but a pair can be created with same ID -- effectively full-duplex
- Customers can lease a VP, then allocate VC's within it ("Permanent VP")
- Types of VC's:
  - Standard VC ("PVC") – Static route
  - Soft VC – Route can be changed in event of failure
  - Signalled VC ("SVC") – Demand connection initiated by used

ATM
- Physical and ATM Layers

VC Connection Messages

Message	Meaning (if from host)	Meaning (if from Network)
SETUP	Call request	Incoming call
CALL PROCEEDING	ACK Incoming call	ACK Call request
CONNECT	Incoming call accepted	Call request accepted
CONNECT ACK	ACK Call request accepted	ACK Incoming call accepted
RELEASE	Terminate request	Terminate req from remote
RELEASE COMPLETE	ACK Terminate from remote	ACK terminate request

ATM
- Physical and ATM Layers

Virtual Circuit Setup Process

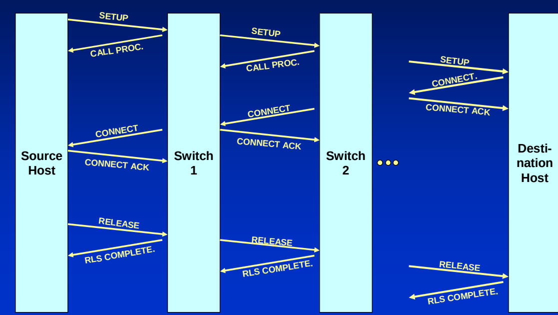

ATM
- Physical and ATM Layers

ATM Broadcasting /Multicasting

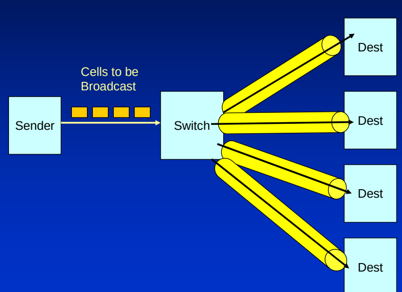

ATM
- Physical and ATM Layers

Operation Administration and Maintenance (OA&M)

Supervison, Maintainance, Testing, Performance Measurement, Loopbacks

Organized into levels according to Network Segment Type

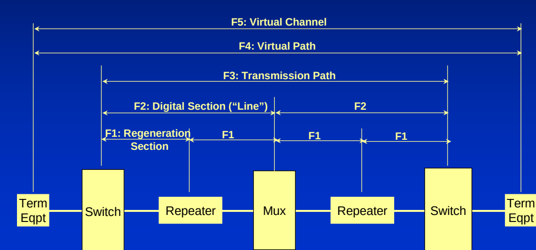

Preassigned VPI and VCI Numbers

VPI	VCI	Meaning
0	0	Unassigned
0	1	Metasignaling
0	3	F4 Flow (segment)
0	4	F4 Flow (end-to-end)
0	5	Signaling
0	15	SMDS
0	16	Intermediate Layer Management Interface (ILMI)

ATM

- Background
- Physical and ATM Layers
  - Cells, Formats, and Addressing
  - Virtual Circuits
  - Switches and Media
  - Interfaces
- AALs
- Applications

Switch Operation

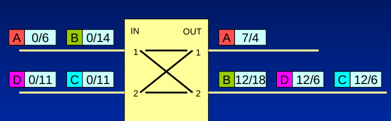

Switching Table

IN Port	IN ID	OUT Port	OUT ID
1	0/6	1	7/4
1	0/14	2	12/18
2	0/11	2	12/6

Label Switching

Transmission Media

- Baseline is SONET (B-ISDN)
- Short runs, Cat 5 TP OK
- All runs point-to-point

SONET / SDH

Bell: Synchronous Optical Network (SONET)
CCIT: Synchronous Digital Hierarchy (SDH)
Only minor differences

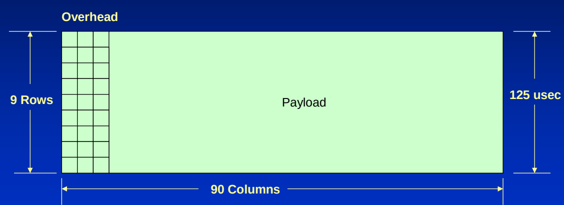

- Synchronous, frame-oriented, TDM
- For basic SONET:
  - 6480 bits/125usec = 51.84 Mbps total (“STS-1”)
  - User data rate =~ 50Mbps

SONET Multiplexing

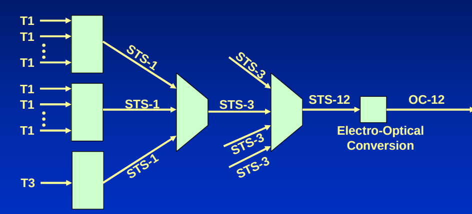

SONET/SDH Data Rates

SONET		SDH	Data Rate	
Electrical	Optical	Optical	Gross	User
STS-1	OC-1		51.84	49.5
STS-3	OC-3	STM-1	155.52	148.6
STS-9	OC-9	STM-3	466.56	445.8
STS-12	OC-12	STM-4	622.08	594.4
STS-18	OC-18	STM-6	933.12	891.6
STS-24	OC-24	STM-8	1244.16	1188.9
STS-36	OC-36	STM-12	1866.24	1783.3
STS-48	OC-48	STM-16	2488.32	2377.7

STS-1 Mux’ed

“OC-x” means multiple users, muxed
“OC-xc” means one user – slightly higher User B/W

ATM over SONET

- ATM designed to run over SONET OC-3c
- Basic: 155.52 Mbps gross rate
  - Usually quoted as 155 Mbps
- New generation runs at OC-12 (622 Mbps), OC-48 (2.4 Gbps)

ATM

- Background
- Physical and ATM Layers
  - Cells, Formats, and Addressing
  - Virtual Circuits
  - Switches and Media
  - Interfaces
- AALs
- Applications

ATM Interfaces

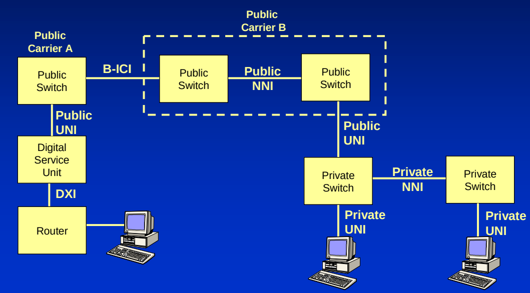

Broadband Inter-Carrier Interface (B-ICI)

- Public Network-to-Network Interface
- Based on Broadband ISDN User-Part (B-ISUP) messages

NNI

- Switch-to-switch interface protocol
- Two versions: Public and private (similar, more flexibility in private version)
- NNI Includes:
  - Routing protocol (Link-sate/OSPF)
  - Signaling protocol for link setup/teardown

ATM
- Physical and ATM Layers

UNI

- Protocol for interfacing with user equipment
- Follows ITU-T Q.2931 message format

ATM
- Physical and ATM Layers

DXI

"Frames In, Cells Out"

Provides Frame-Based Access to an ATM Network

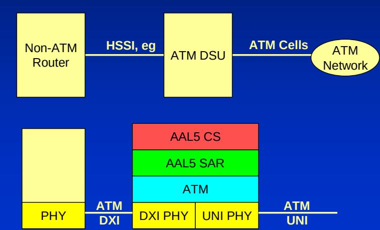

ATM
- Physical and ATM Layers

DXI Variants

Mode	AALs Supported	VCs Supported
1a	AAL5	1023
1b	AAL5	1023
	AAL3/4	1
2	AAL3/4,5	16.6M

1	2	0-9232	2	1
Flag	Header	Body	FCS	Flag

1	2	0-9224	2	1
Flag	Header	Body	FCS	Flag

1	2	0-65535	4	1
Flag	Header	Body	FCS	Flag

ATM
- AAL Layer

ATM

- Background
- Physical and ATM Layers
- AALs
  - AAL Overview
  - AAL 1 and 2
  - AAL 3/4
  - AAL 5
- Applications

ATM
 - AAL Layer

Why AAL?

An “impedance matcher”
between ATM and
higher-level protocols
with variable-length cells

Chart © Glenn W Cox, 2001-2004

ATM
 - AAL Layer

The Evolution Of the AAL’s

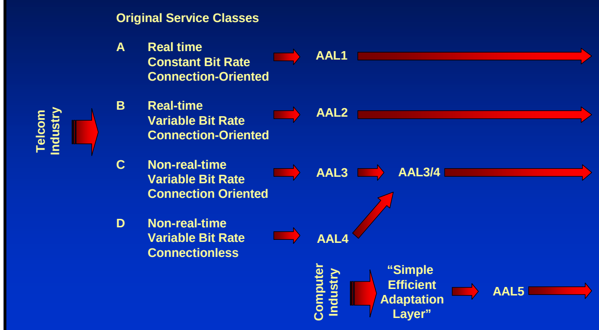

Chart © Glenn W Cox, 2001-2004

ATM
 - AAL Layer

Cell Formation

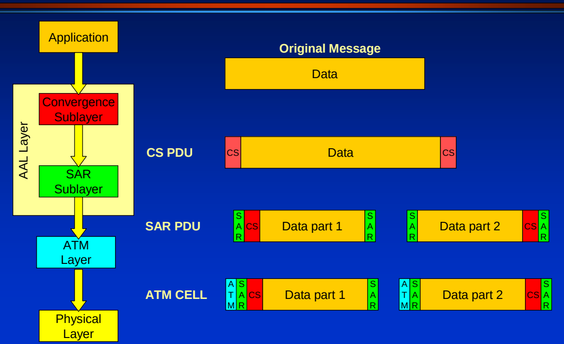

Chart © Glenn W Cox, 2001-2004

ATM
 - AAL Layer

ATM

- Background
- Physical and ATM Layers
- AALs
  - AAL Overview
  - AAL 1 and 2
  - AAL 3/4
  - AAL 5
- Applications

Chart © Glenn W Cox, 2001-2004

AAL 1

- Designed to support Class A traffic (voice)
- Voice has good error tolerance -> No bit error control (CRC) needed
- Sequence numbers needed to ID missing cells

AAL 1 Convergence Sublayer

- Detects lost cells
- Detects mis-delivered (“misinserted”) cells
- Smooths incoming traffic to minimize jitter
- Breaks bit stream into 47/46-byte segments for SAR sublayer
- Does not add headers or trailers

AAL 1 SAR PDU (non-pointer type)

- Adds sequence # with protection (checksum)
- Adds parity bit (even) over header

AAL1 SAR PDU (Pointer Type)

- Pointer field gives offset to start of next message (0-92 bytes)

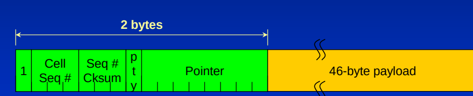

AAL 2

- Designed to support Variable Bit Rate ("Bandwidth on Demand")
- Provides for partial payloads to support low-rate data with low latency
- Error protection over full PDU
- Simple flag to indicate position in message

AAL2 SAR PDU

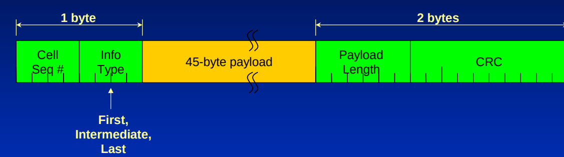

ATM

- Background
- Physical and ATM Layers
- AALs
    - AAL Overview
    - AAL 1 and 2
    - AAL 3/4
    - AAL 5
- Applications

AAL 3/4

- Originally 2 separate AALs:
    - AAL3: Connection-oriented packet svcs (X.25)
    - AAL4: Connectionless svcs (IP)
- Eventually combined into a single type for all data service
- Data support overtaken by AAL5

ATM
- AAL Layer

AAL3/4 CS PDU

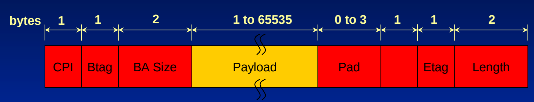

CPI - Common Part Indicator	**Message type, Units for BA Size and Length**
Btag, Etag	**Identical sentinel bytes, Incremented for each new message**
BA Size	**Estimated payload size (for buffer allocation)**
Pad	**Bytes added to to make Payload a multiple of 4 bytes**
Length	**True payload size**

ATM
- AAL Layer

AAL3/4 SAR PDU

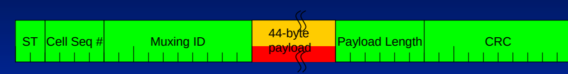

ST - Segment Type	**Indicates if this cell is from the middle (00), end (01) or beginning (10) of this message or if this is a single-cell message (11)**
Muxing ID	**ID of the session that this cell belongs to (the CS may be handling multiple sessions simultaneously**

ATM
- AAL Layer

ATM

- Background
- Physical and ATM Layers
- AALs
  - AAL Overview
  - AAL 1 and 2
  - AAL 3/4
  - AAL 5
- Applications

ATM
- AAL Layer

AAL 5

- Pushed by computer industry as a lower-overhead data format
- The idea: Instead of using some of the 48-byte cell payload for SAR info, steal a bit from the cell to denote end of message

Efficiency:

AAL3/4: 4 bytes per message + 4 bytes per cell => 44 User Data Bytes / Cell

AAL5: 8 bytes per message => 48 User Data Bytes / Cell, 8% improvement

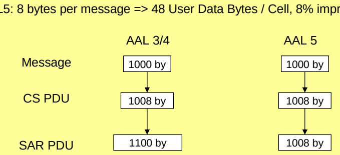

ATM
- AAL Layer

AAL5 CS PDU

1 to 65535 bytes	0 to 47	1	1	2	4
Payload	Pad	UU		Length	CRC

Pad	**Inserted to make entire CS PDU a multiple of 48 bytes**
UU	**"User - to - User" field. Available for use by higher levels.**
Length	**Payload length, not counting padding**

ATM
- AAL Layer

AAL5 SAR

- Simply breaks CS PDU into 48-byte chunks and passes them to ATM Layer.
- No overhead bytes added.

ATM
- Applications

ATM

- Background
- Physical and ATM Layers
- AALs
- Applications
  - QoS
  - LANE
  - IP over ATM, MPOA

ATM
- Applications

ATM QoS

- A (The?) Major ATM selling point vs. Best-Effort
- The idea:
  - At VC setup, sender specifies level and quality of service required, also planned traffic profile.
  - While establishing VC, network attempts to allocate resources to meet requirements.
  - Requirements are agreed to or available capability is passed back to sender.
  - During transmission, network enforces traffic profile

ATM
- Applications

ATM Classes of Service

Class	Description	CLR- Cell Loss Ratio	CTD - Cell Transfer delay	CDV - Cell Delay Variation	PCR - Peak Cell Rate	SCR - Sustained Cell Rate	BT - Burst Tolerance	Flow Control
CBR	Constant bit rate (e.g., Phone traffic)	Y	Y	Y	Y			
VBR-RT	Variable bit rate, Real-Time (e.g., Interactive Compressed Video)	Y	Y	Y	Y	Y	Y	
VBR-NRT	Variable bit rate, Non-Real-Time (e.g, Multimedia email)	Y		Y	Y	Y	Y	
ABR	Available bit rate (e.g., File Xfer, email)	Y						Y
UBR	Unspecified bit rate (e.g., TCP/IP)				Y			

ATM
- Applications

QoS Parameters

Parameter	*Meaning*
PCR - Peak Cell Rate	Max rate req'd
SCR - Sustained Cell Rate	Avg rate req'd
MCR - Minimum Cell Rate	Min acceptable rate (Used in ABR Service)
CDVT - Cell delay variation tolerance	Max acceptable jitter
CLR - Cell Loss Ratio	Fraction of cells lost or late
CTD - Cell Transfer Delay	Delivery time (mean and max)
CDV - Cell Delay Variation	Measured jitter
CER - Cell Error Rate	Fraction with one or more errors
SECBR - Severely-Errored Cell Block Ratio	Fraction of "M"-cell blocks with "N" or more errored cells
CMR - Cell Misinsertion Rate	Fraction delivered to wrong destination
BT - Burst Tolerance	Max Burst That can be sent at Peak Rate

ATM
- Applications

Traffic Shaping

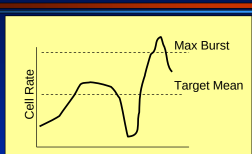

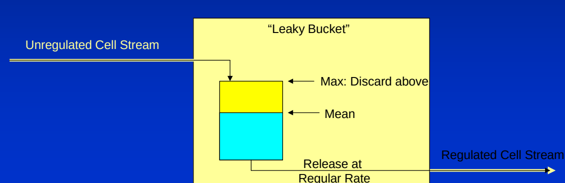

ATM
- Applications

Agenda

- Background
- Physical and ATM Layers
- AALs
- Applications
  - QoS
  - LANE
  - IP over ATM, MPOA

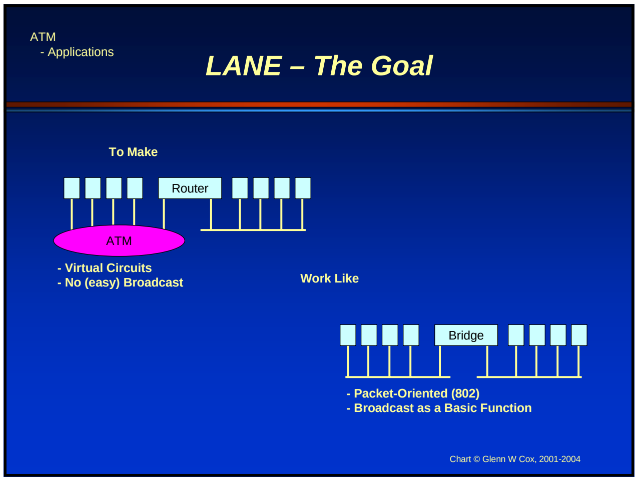
ATM
- Applications
LANE – The Goal
To Make
Router
ATM
- Virtual Circuits
- No (easy) Broadcast
Work Like
Bridge
- Packet-Oriented (802)
- Broadcast as a Basic Function
Chart © Glenn W Cox, 2001-2004

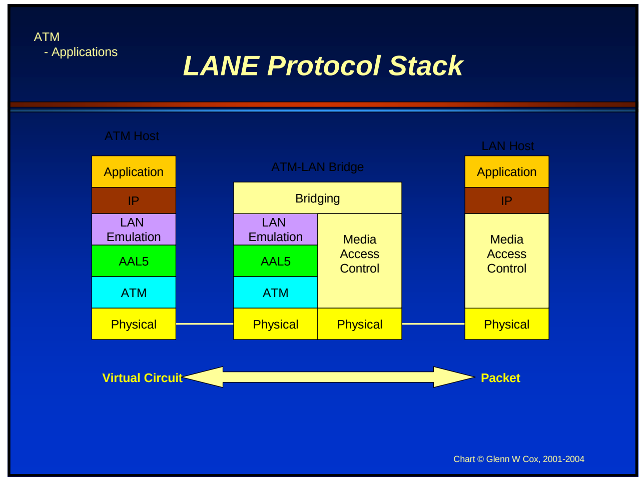
ATM
- Applications
LANE Protocol Stack
ATM Host
ATM-LAN Bridge
LAN Host
Application
IP
LAN Emulation
AAL5
ATM
Physical
Bridging
Media Access Control
Virtual Circuit
Packet
Chart © Glenn W Cox, 2001-2004

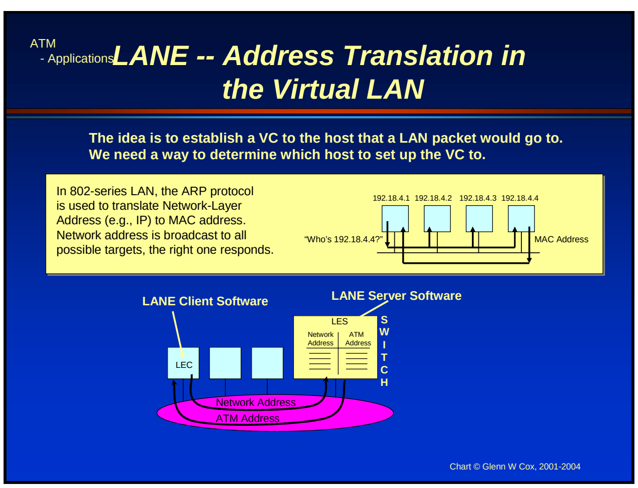
ATM
- Applications
LANE -- Address Translation in the Virtual LAN
The idea is to establish a VC to the host that a LAN packet would go to.
We need a way to determine which host to set up the VC to.
In 802-series LAN, the ARP protocol is used to translate Network-Layer Address (e.g., IP) to MAC address. Network address is broadcast to all possible targets, the right one responds.
192.18.4.1 192.18.4.2 192.18.4.3 192.18.4.4
"Who's 192.18.4.4?"
MAC Address
LANE Client Software
LANE Server Software
LES
Network Address
ATM Address
SWITCH
LEC
Network Address
ATM Address
Chart © Glenn W Cox, 2001-2004

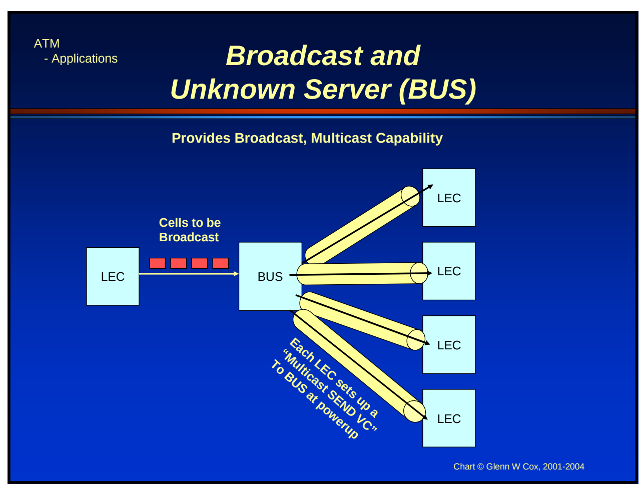
ATM
- Applications
Broadcast and Unknown Server (BUS)
Provides Broadcast, Multicast Capability
Cells to be Broadcast
LEC
BUS
LEC
LEC
LEC
LEC
Each LEC sets up a "Multicast SEND VC" To BUS at powerup
Chart © Glenn W Cox, 2001-2004

ATM
- Applications

ATM

- Background
- Physical and ATM Layers
- AALs
- Applications
  - QoS
  - LANE
  - IP over ATM, MPOA

ATM
- Applications

Classical IP over ATM

- Could use LANE to implement ATM-IP interoperability, but may be too much overhead for large networks
- Defined by IETF (RFC 1577)
- The key is IP <-> ATM address resolution
  - ATM Address Resolution Protocol (ATMARP)
  - Inverse ATMARP (InATMARP)

ATM
- Applications

IP over ATM Connection Setup Process

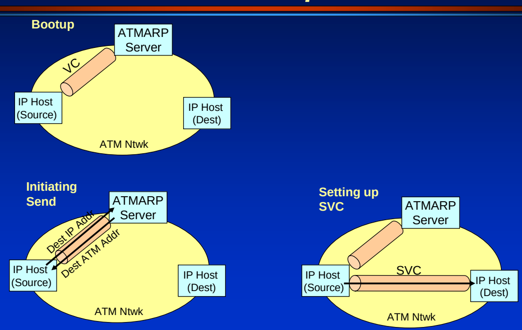

ATM
- Applications

MultiProtocol Over ATM (MPOA)

LANE supports a single legacy LAN format over ATM. MPOA extends this to multiple LAN types