Ethernet

Círculo de tecnologia Eth

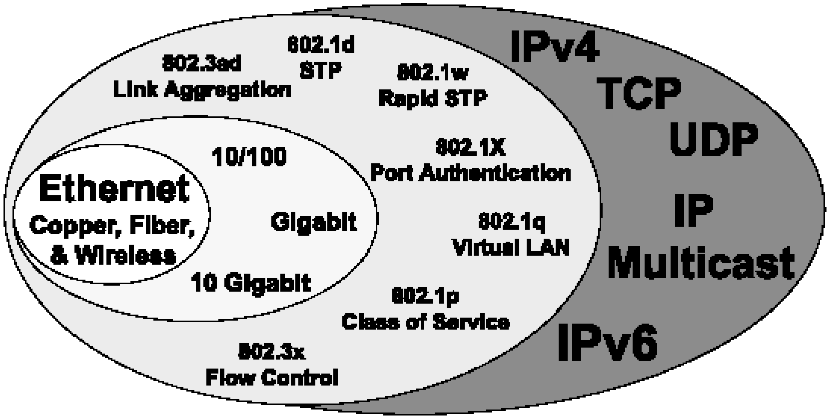

- Ethernet technology is onipresente.
- Mais de 85% de todas as conexões instaladas são Eth (IDC-1997)
- Isso significa mais de 118 milhões de PCs,workstations e servers.
- As conexões restantes são uma combinação de
  - Token Ring,
  - Fiber Distributed Data Interface (FDDI),
  - Asynchronous Transfer Mode (ATM)
  - Outras tecnologias

A tecnologia dominante

- Todos os sistemas populares e aplicações são Ethernet-compatible, como também são as camadas superiores dos protocolos:
  - Transmission Control Protocol/Internet Protocol (TCP/IP),
  - IPX,
  - NetBEUI
  - DECnet.
- O ano de 1998 foi marcante para o Eth no mercado de NICs
- 86% das vendas
- 48 milhões de unidades

- Em termos de portas foram excedidas 48 milhões em hubs/switches.
- Em contraste, ATM, FDDI/CDDI e Token Ring

Combinados ficaram próximos de 5 milhões em1998, 10 % do total.

- Em termos de portas de concentradores (hubs/switches) ATM, FDDI/CDDI e Token Ring chegaram a 4 milhões, 7 % do total.

A banda e o eth

figure 1. Ethernet bandwidth evolution since 1973

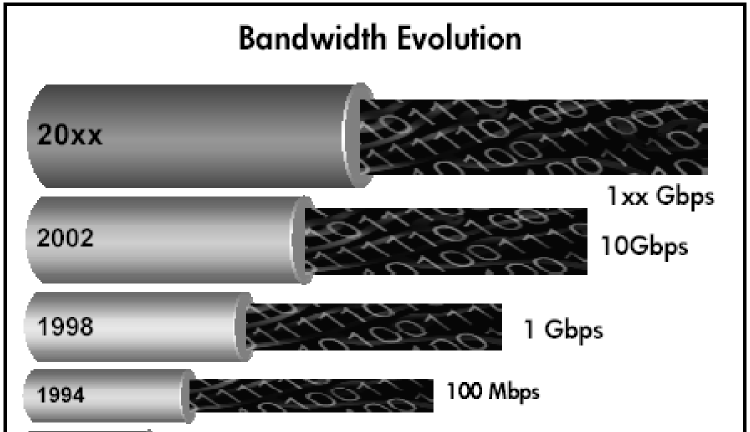

- Novos suplementos recebem uma designação de letra quando são criados.
- Quando o suplemento tiver completado o processo de padronização,
- ele se torna parte do padrão básico
- e não é mais publicado como um documento suplementar separado.
- Por outro lado, você às vezes encontrará alguma literatura referindo-se ao equipamento Ethernet
- com a letra do suplemento em que a variedade foi desenvolvida inicialmente

Suplementos IEEE

- (por exemplo, IEEE 802.3u pode ser usado como referência ao Fast Ethernet).
- A Tabela relaciona diversos suplementos e a que eles se referem.
- As datas indicam quando ocorreu a aceitaçao formal do suplemento ao padrão.

- Suplemento Descrição
- 802.3a- 1985 1OBASE2 thin Ethernet
- 802.3c- 1985 10 Mbps repeater specifications, clause 9
- 802.3d- 1987 FOIRL fiber link
- 802.3i- 1990 10BASE-T twisted-pair
- 802.3j- 1993 1OBASE-F fiber optic
- 802.3u- 1995 100BASE-T Fast Ethernet and Auto-Negotiation
- 802.3x- 1997 Full-Duplex standard
- 802.3z- 1998 1000BASE-X Gigabit Ethernet
- 802.3ab- 1999 1000BASE-T Gigabit Ethernet over twisted-pair
- 802.3ac- 1998 Frame size extension to 1522 bytes for VLAN tag
- 802.3ad- 2000 Link aggregation for parallel links

10BASE-2

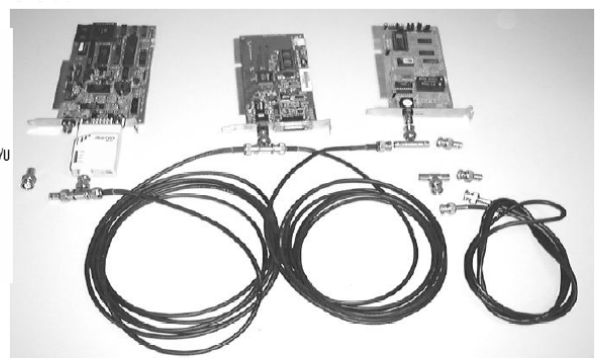

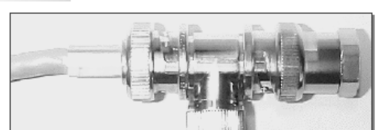

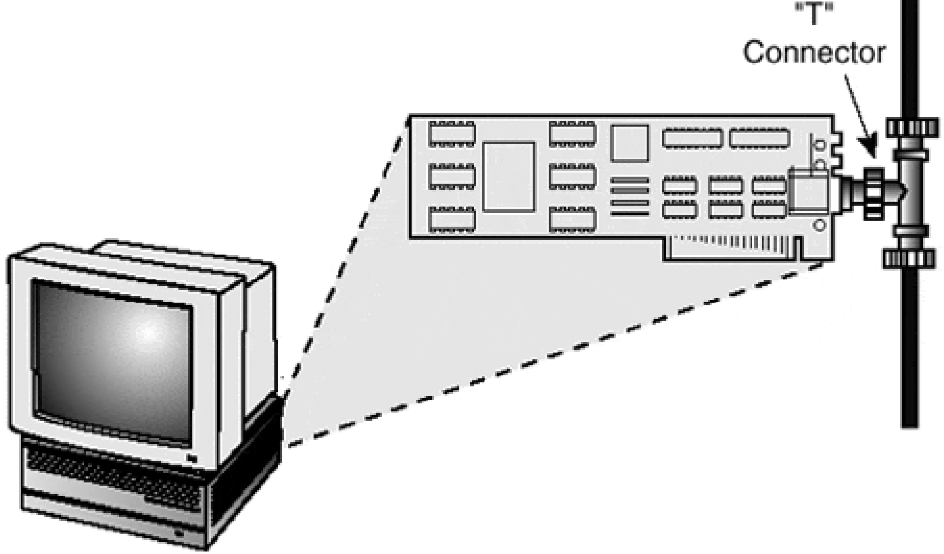

As subcamadas do Eth

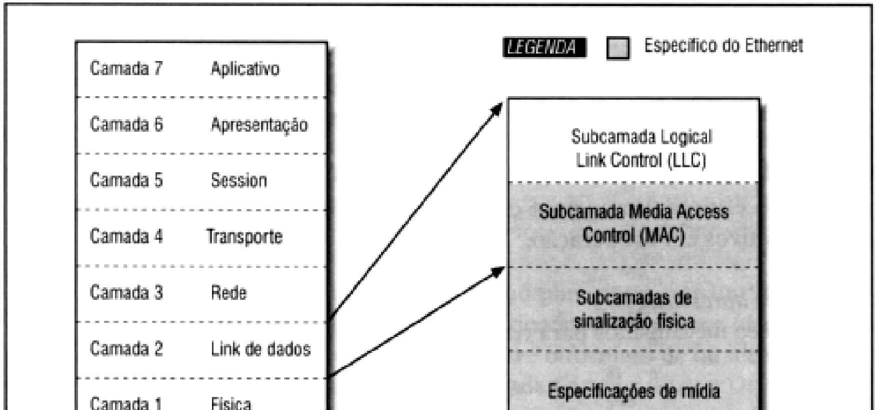

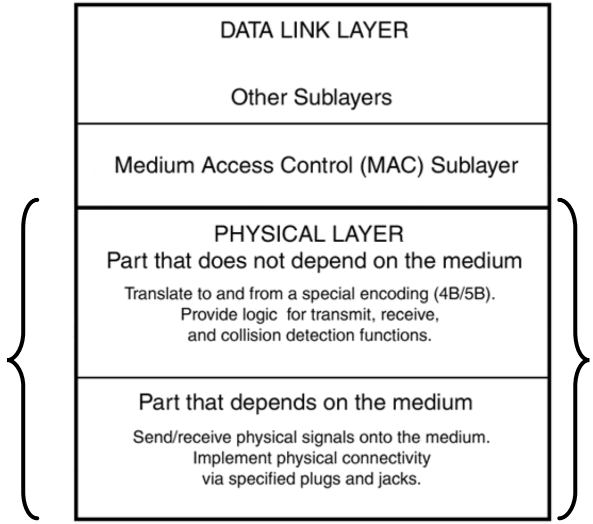

Nome das PDUs

Application Layer	
Presentation Layer	
Session Layer	"Messages"
Transport Layer	TCP Segment(s)
Network Layer	IP Datagram or Packet
Data Link Layer	Frame

All Ethernet systems use Manchester encoding due to its simplicity.

The high signal is + 0.85 volts and the low signal is - 0.85 volts, giving a DC value of 0 volts.

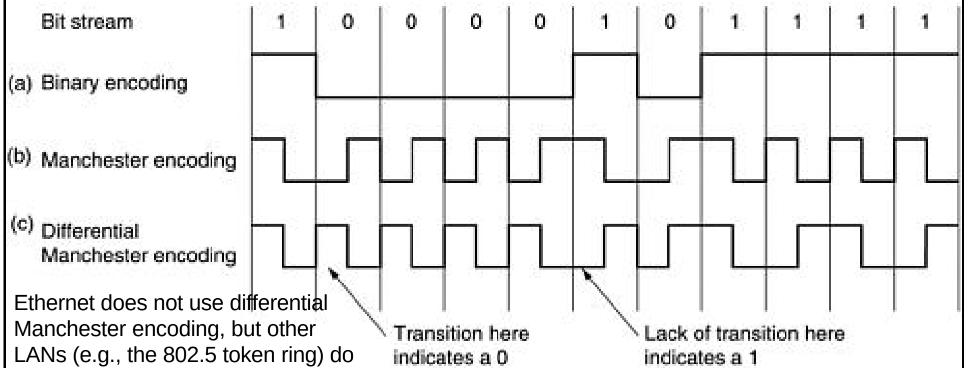

Ethernet does not use differential Manchester encoding, but other LANs (e.g., the 802.5 token ring) do use it.

O frame

7 bytes	1 byte	6 bytes	6 bytes	2 bytes	46-1500 bytes	4 bytes
Preamble	Start of Frame Delimiter (SFD)	Destination Mac Address	Source Mac Address	Length of Data Field	Data Field	Frame Check Sequence (FCS) (cyclical redundancy check)

minimum of 64 bytes, maximum of 1518 bytes

O tamanho do campo de dados é usado agora, ao contrário do eth original, para especificar quanto de dados existe no campo a seguir

Quando forem 1500 ou menos, o campo é usado para designar o

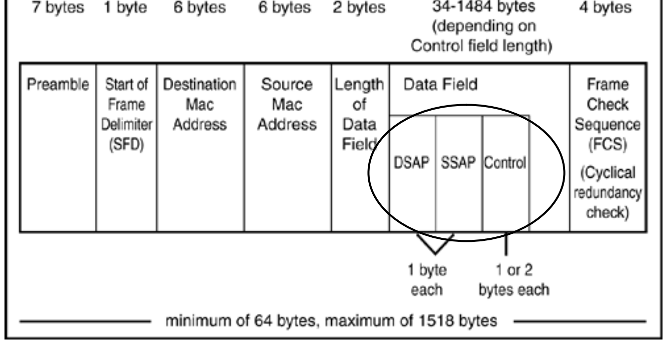

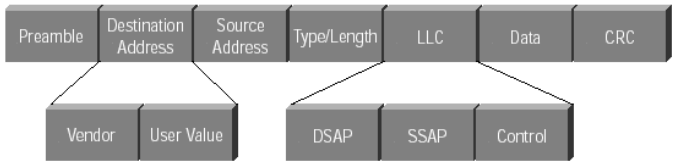

DSAP (1 byte)
SSAP (1 byte)
Control (1 or 2 bytes)

Campo de dados

- *(46 a 1500 bytes)*
- Esse campo carrega o datagrama
- A unidade máxima de transferência (MTU) da Ethernet é 1500 bytes.
- Isso significa que, se o datagrama IP exceder 1500 bytes,
- O host terá de fragmentar o datagrama,
- 0 tamanho mínimo do campo de dados é 46 bytes.
- Isso significa que, se um datagrama IP tiver menos do que 46 bytes,

- Quando se usa enchimento, os dados passados à camada de rede contêm o recheio, bem como um datagrama IP.
- A camada de rede usa o campo de comprimento do cabeçalho do datagrama IP para remover o enchimento.

Endereço de destino

- *(6 bytes).* Esse campo contém o endereço de LAN do adaptador de destino,
- ou seja, BB-BB-BB-BB-BB-BB.
- Quando o adaptador B recebe um quadro Ethernet com um endereço de destino *que não é* seu próprio endereço físico, BB-BB-BB-BB-BB-BB
- nem o endereço de broadcast da LAN (FFFFFFFFFFFF)
- ele descarta o quadro.
- Caso contrário, ele passa o conteúdo do

- *(6 bytes)*
- Esse campo contém o endereço de LAN do adaptador
- que transmite o quadro para a LAN,
- Ex: AA-AA-AA-AA-AA-AA.

Campo de tipo

- *(2 bytes). 0* campo de tipo permite que a Ethernet 'multiplexe' os protocolos de camada de rede.
- os hosts podem usar outros protocolos de camada de rede além do IP.
- De fato, um dado host pode suportar múltiplos protocolos de camada de rede e usar diferentes protocolos para diferentes aplicações.
- Por essa razão, quando o quadro Ethernet chega ao adaptador B, este precisa saber para qual protocolo de camada de rede ele deve passar (isto é, demultiplexar) o conteúdo do campo de dados.
- 0 IP e outros protocolos de camada de rede (por exemplo, Novell IPX ou AppleTalk) têm seu próprio número de tipo padronizado.
- Além disso, o protocolo ARP tem seu próprio número de tipo. Note que o campo de tipo é análogo ao campo de protocolo no datagrama de camada de rede e aos campos de número de

```
DLC:  Frame 1 arrived at  12:09:34.0000; frame size is 60
(003C hex) bytes.
      DLC:  Destination = BROADCAST FFFFFFFFFFFF, Broadcast
      DLC:  Source      = Station Exceln201982
      DLC:  Ethertype   = 0806 (ARP)
      DLC:
ARP: ----- ARP/RARP frame -----
      ARP:
          (28 byte ARP message)
      ARP:
      ARP: 18 bytes frame padding
      ARP:
HEX
ff ff ff ff ff ff 08 00 14 20 19 82 08 06 00 01
08 00 06 04 00 01 08 00 14 20 19 82 81 54 19 02
00 00 00 00 00 00 81 54 19 fe 01 01 00 00 26 3d
ea d9 00 00 00 00 6b 69 6c 6c 6a 6f
```


Ethernet Types

Hex Value	Decimal Value	Content of the Data Field
08-00	2048	IP datagram
08-06	2054	Address Resolution Protocol (ARP) message
0B-AD	2989	Banyan VINES
80-9B	32923	AppleTalk data units
80-D5	32981	IBM SNA services over Ethernet
81-37, 81-38	33079, 33080	NetWare data units
86-DD	34525	IP Version 6 datagram
60-03	24579	DECnet Phase IV

- *(4 bytes).* finalidade do campo de CRC é permitir que o adaptador receptor, o adaptador B,
- detecte se algum erro foi introduzido no quadro,
- isto é, <u>se os bits do quadro foram trocados.</u>
- Dentre as causas de erro nos bits estão
  - a atenuação da força do sinal
  - e a energia eletromagnética ambiente que perpassa os cabos e as placas de interface Ethernet.

Processo de detecção

- Quando o hospedeiro A monta o quadro Ethernet, ele calcula um campo de CRC,
- que é obtido com base em um mapeamento dos outros bits presentes no quadro (exceto os bits do preâmbulo).
- Quando o host B recebe o quadro, ele aplica o mesmo mapeamento ao quadro
- e verifica se o resultado do mapeamento é igual ao que está no campo de CRC.
- Essa operação é chamada, no hospedeiro receptor, de **verificação de CRC.**
- Se a verificação de CRC falhar (isto é, se o resultado do mapeamento não for igual ao conteúdo do campo de CRC),
- o hospedeiro B saberá que há um erro no quadro.
- O quadro é dropado
- O emissor não é avisado
  - Sem conexão,

Tipos de erros

Bad FCS e Misaligned Frames

- Frame Check Sequence (FCS) error.
- A camada MAC computa um valor de cyclic redundancy check (CRC),
- Baseado no conteúdo do frame
- E toca o valor no campo FCS
- É possivel que o valor tenha saído errado da origem, devido a problemas de HW
- Também é possível que o NIC que manda o frame não esteja transmitindo corretamente os bits.
- Mas a maioria dos erros é devida a ruídos , interferencias e outros
- Quando voce estiver monitorando o nível de BAD FCS e essa coisa passar de 2 a 3% da banda utilizada, você tem que encontrar a fonte dos erros e tomar

- Um *runt* é um frame eth menor que o tamanho mínimo de 64 bytes
- Lembre-se que uma NIC deve transmitir um pacote por um período de tempo igual ao RTT antes de parar a transmissão
- Caso contrário, a NIC não pode detectar a colisão
- O tempo máximo de propagação para o eth é de 51.2 microssegundos
- Que é o total de tempo necessário para transmitir 64 bytes
- O preambulo nao esta computado nesse calculo

Razoes para o *Runt*

- Colisoes
- Placas histericas
- Portas detonadas

- Erros de topologia

- Algumas vezes a NIC produz frames maiores que o m’aximo permitido
- O *giant frame error* eh o oposto ao runt error

Razoes para o gigante

- A defective NIC that is transmitting continuously.
- Os bits que indicam o tamanho do frame podem estar errados
- Ruido no cabo pode ser interpretado como parte do frame, mas isso normalmente eh detectado pelo FCS

- O termo *jabber* é muitas vezes susado par referenciar frames muito grandes
- Mas na real, ‘e um termo gen’erico para indicar um dispositivo mal comportado (not following the rules)
- Uma placa pode mandar os frames muito grandes
- Ou mandar ininterruptamente,
- Esse tipo de erro pode derrubar qualquer rede, porque as placas boas ficam esperando indefinidamente pelo acesso

Preâmbulo

- *(8 bytes). 0* quadro Ethernet começa com um campo de preâmbulo de 8 bytes.
- Cada um dos primeiros 7 bytes do preâmbulo tem um valor de 10101010;
- o último byte é 10101011.
- Os primeiros 7 bytes do preâmbulo servem para `despertar' os adaptadores receptores e sincronizar o relógio deles com o relógio do remetente.
- Por que os relógios poderiam estar fora de sincronização?
- Tenha em mente que o adaptador A visa a transmitir o quadro a 10 Mbps, 100 Mbps ou 1 Gbps, dependendo do tipo de

- Dependendo da ferramenta usada para monitorar a rede
- O <u>numero de erros tipos diferentes</u> pode variar
- Por exemplo, erros de frame desalinhado usualmente possuem tambem FCS errado.
- Alguns analizadores podem computar 2 erros para o mesmo evento
- Enquanto outros podem computar somente um ou somente outro
- Leia o manual da ferramenta, se voce for capaz!!!

Broadcast Storms

- *Broadcast storms* ocorrem quando algum dispositivo da rede gera trafego
- Que induz os demais a gerar mais trafego ainda
- Embora esse trafego adicional possa ser devido a falhas fisicas,
- Nos dispositivos ou no cabeamento,
- Normalmente eles sao causados pelos protocolos de nivel mais alto

- Uma broadcast storm é coisa feia…
- O maior problema para detectar é ter acesso a rede
- Broadcast storms deixam a rede lenta, e podem derrubar a miserável
- <u>Um tráfego normal de broadcast é de 100 frames por segundo ou menos</u>
- Quando esse valor aumenta para mais de 100 frames/segundo, pode haver um problema
- Uma placa precisa ser derretida
- Ou o dominio de colisao precisa ser motoserrado

Monitorando os Erros

- Existem várias ferramentas para monitorar erros
- Um *network analyzer*, por exemplo, Network Sniffer from Network General,
- Mostra informações dos conteúdos
  - dos frames que contem erros,
  - runts,
  - jabbers
  - CRC,
  - E alignment errors.
- Para gerência centralizada em uma rede séria, nada melhor que o SNMP e o RMON

Preço dos eth

Ratio	2.0	1.5

Switch Per-Port Average Price	1996	1998
Ethernet	$440	$215
Fast Ethernet	$716	$432
Ratio	1.6	2.0

NIC Per-Port Average Price	1996	1998
Ethernet	$ 68	$ 35
Fast Ethernet	$122	$ 74
Ratio	1.8	2.2

Source: IDC, 1996

Algoritmo do CSMA/CD

Carrier Sense Multiple Access / Collision Detection

- É um método de acesso ao meio com contenção (quando detecta colisão, contém a transmissão)
- Perde até 40% da eficiência em grandes tráfegos
- Não utiliza níveis de prioridade

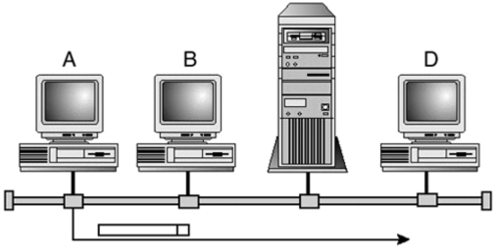

Um sistema que queira transmitir precisa esperar até que o barramento esteja em silêncio

Ou seja, ninguém mais está transmitindo

Somente é permitido a um frame atravessar o barramentto em um dado instante.

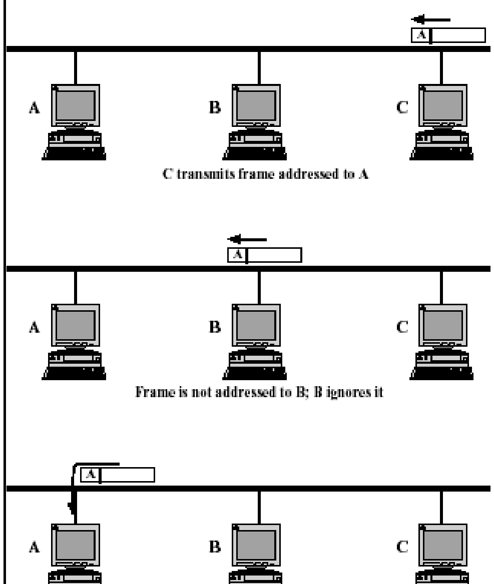

Suas transmissões irão colidir

E cada uma delas deverá esperar um tempo aleatório

Antes de tentar novamente

Esse conjunto de regras é o *Carrier Sense Multiple Access with Collision Detection* (CSMA/CD).

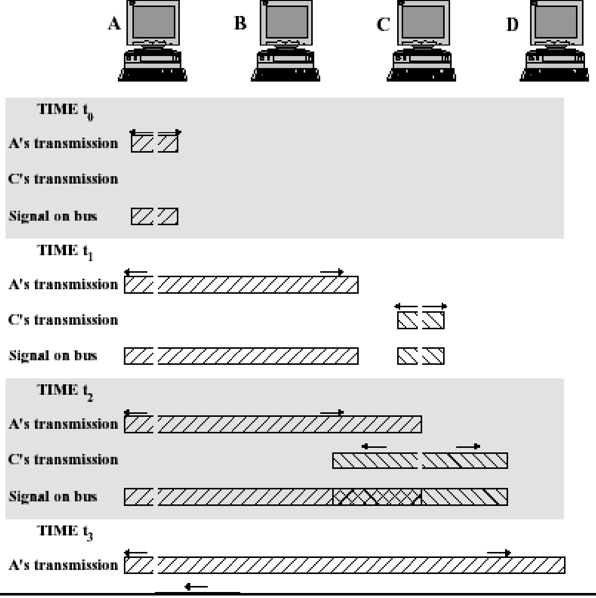

Todos os sistemas ouvem o meio e percebem cada frame que está sendo transmitido.

Cada frame possui em seu cabeçalho os endereços de origem e destino

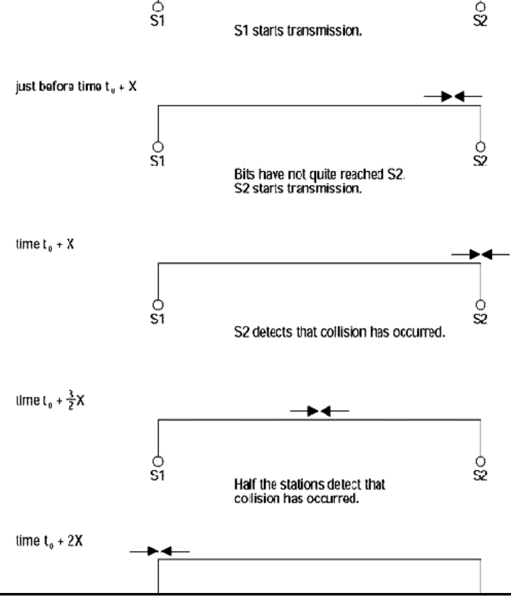

Algoritmo do CSMA/CD

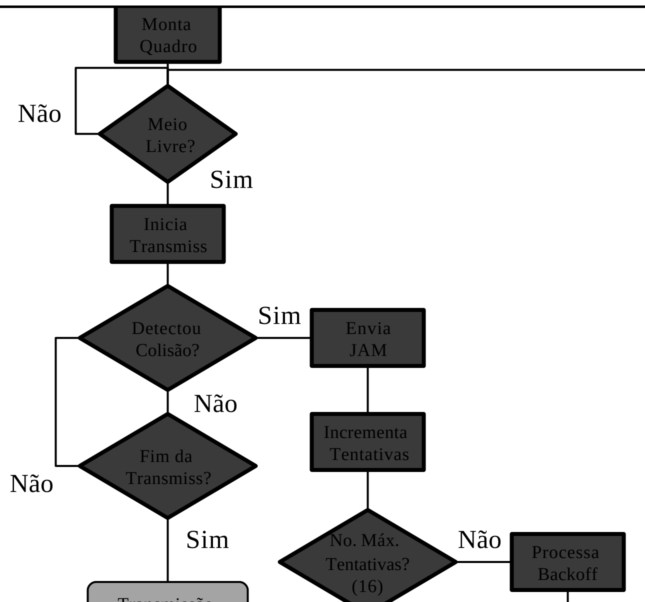

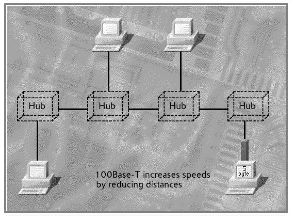

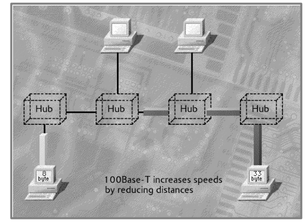

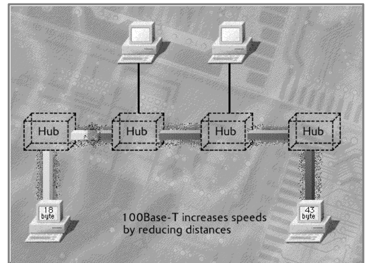

10BASE2 Collision Domain Parameters	
Parameter or Characteristic	**Value**
Topology	Branching bus
Type of segment cable	Thin 50-ohm coax
Connector	BNC
Maximum length of a segment	185 meters
Maximum propagation delay for a segment	950 nanoseconds
Maximum number of nodes (stations and repeaters) that can be attached to a segment	30
Minimum distance between two nodes	.5 meter
Maximum collision domain diameter	925 meters
Maximum number of cable segments traversed between a source and destination	5
Maximum number of repeaters between a source and	4

- O padrão eth original, com cabo coaxial, só conseguia transmissão half duplex.
- Isso pq o sinal em banda básica ocupava toda a banda passante.

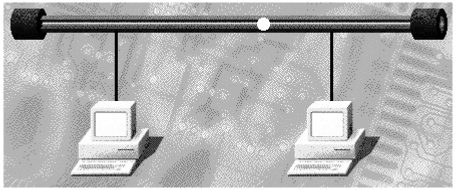

10Mbps Twisted-Pair Transmission

- 10Base-T carries data across two twisted pairs of wire. Data is transmitted from a system on one pair and is received on the other pair. The term *twisted-pair* refers to two wires wrapped around each other. For 10BASE-T, the same signal is sent across both of the wires in a pair but with reverse polarity. This means that when a positive voltage is placed on one wire, an equal negative voltage is placed on the other.
- There is a good reason for this. A current sent down a copper wire creates an electromagnetic field that induces currents in nearby wires. Some people experience this effect (called *crosstalk*) during a telephone call when they suddenly begin to hear a conversation taking place on another set of wires.
- When two wires carrying these reversed signals are twisted around one another, the electromagnetic fields around each wire come very close to canceling each other out. The arrangement is called a *balanced cable*. The more twists per foot, the better this works. The number of twists ranges from 2 to 12 per foot. Figure 6.14 illustrates the relationship between the pairs. The wire labeled Tx+ carries the transmitted signal, and Tx- carries the inverted transmit signal. The wire labeled Rx+ carries the signal that was transmitted from the remote end, and Rx- carries the inverted signal.

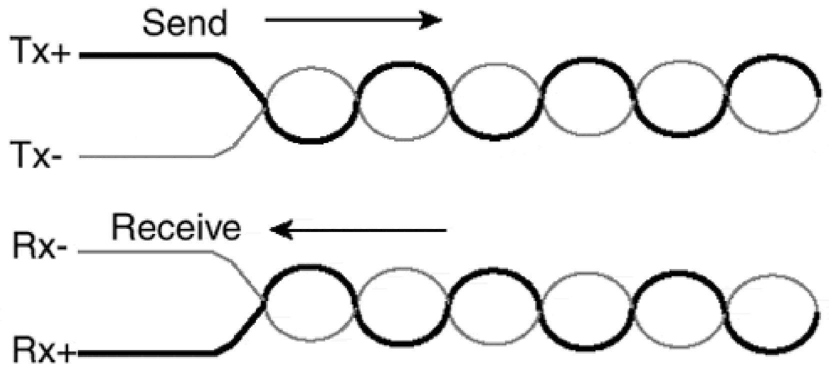

- There is another good reason for this paired transmission. If a noise spike occurs along the way, both signals in a pair will be subjected to the same spike. On reception, the inverted signal on the incoming Rx- cable is inverted again. This causes the noise spike to be inverted. The inverted Rx- signal then is added to the Rx+ signal, and the result is that the noise

The two twisted pairs are terminated by a RJ-45 plug at each end.
This plug has eight contacts, numbered 1 to 8.
Only four are needed, and the ones that are used are:

1	Tx+
2	Tx-
3	Rx+
6	Rx-

Pin, wire, and jack connections.

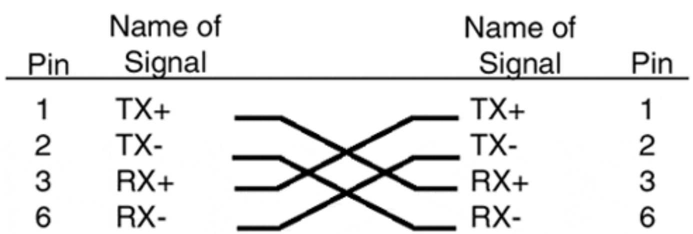

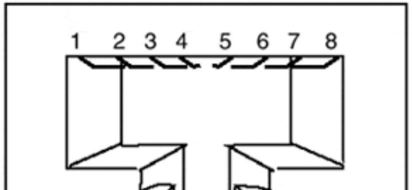

CABLE TYPE	SPECIFICATION (MHz)	CONNECTOR TYPE	STANDARD STATUS	TESTING	COMMENT
Category 5	100	RJ-45	Installation defined in ANSI/TIA/EIA568-A: 1995. Testing defined by TSB67 and TSB95 (in final review)	TSB67 TSB95	First defined in 1991, redefined in 1995. Cable systems installed before 1995 contain non-standard hardware. Because ANSI/TIA/EIA568-A: 1995 does not specify return loss and ELFEXT performance, Category 5 installations require additional testing per TSB95 or IEEE802.3ab.
Category 5-e (enhanced)	100	RJ-45	In development by TIA TR-42 as Standards Proposal 4195-B. Will be Addendum 5 to TIA/EIA568-A.	Will be Addendum 5 to TIA/EIA 568-A.	Recommended for new installations to insure that installers incorporate Return Loss and ELFEXT requirements in their certification.
Category 6	250	RJ-45	In development by TIA TR-42.7.1 and ISO/IEC/SC25/WG3		Products labeled Category 6 today may not be in compliance with the final specification.
Category 7	600	RJ-45 (2 pr) 4 pr TBD	In development by ISO/IEC/SC25/WG3		Expected in ISO/IEC 11801-2000

Crossover cable wires.

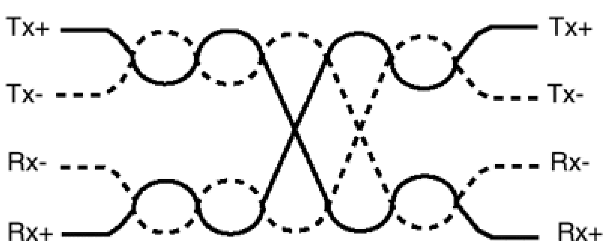

- An ordinary garden-variety port that transmits on pins 1 and 2 and receives on pins 3 and 6 is called a *medium dependent interface* (MDI) port.
- Stations have MDI ports.

- You might expect that a crossover cable would be needed between a station and a hub, but hub ports are structured so that the crossover is performed inside the port. These ports are said to be of type *medium dependent interface crossover* (MDI-X). A straight-through cable is used between an MDI port in an end system and an MDI-X port in a hub.
- What do you do when a hub has to be connected to another hub? A straight-through cable won't work between two MDI-X ports. Transmit would be connected to transmit and receive to receive. Fortunately, most hubs have a special MDI port that is provided for this purpose. A straight-through cable is used to connect an MDI port on one hub to an MDI-X port on the other hub.
- Users prefer not to waste a port on a hub that does not need to be connected to another hub. Many hubs have a special port that is marked MDI/MDI-X. It can be set to MDI or MDI-X by pushing or releasing a button or by moving a slider. Choosing MDI-X lets you connect an end-user station to the port via a straight-through cable. Choosing MDI enables you to connect to an MDI-X port on another

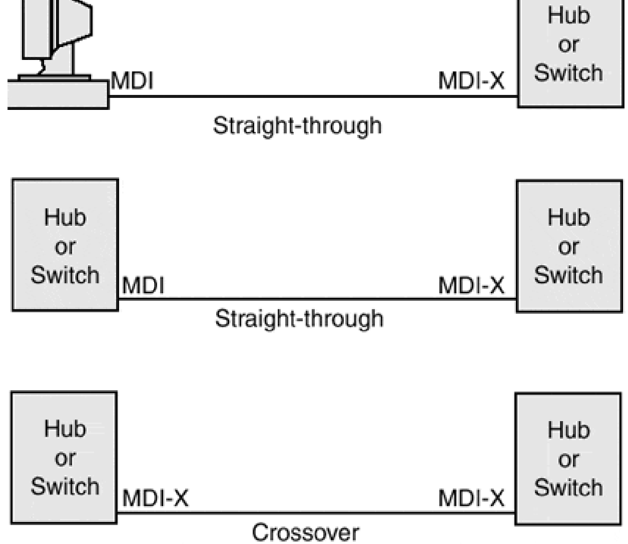

Splines

- Alguns fabricantes de cabos de high-performance Category 6 isolam os 4 pares de cobre
- Usando separadores chamados *splines.*
- Os *splines* mantêm o espaçamento entre os pares
- O uso de splines reduz a atenuação e os crosstalk:
  - near end crosstalk,
  - far end crosstalk.

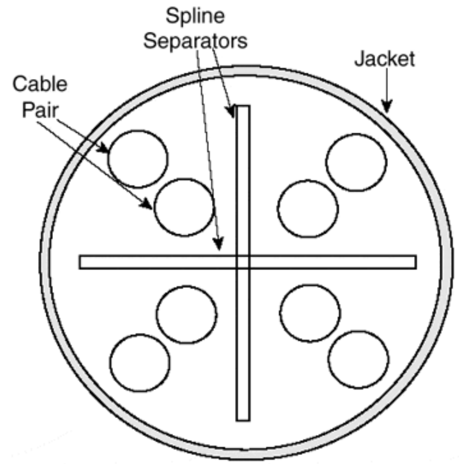

Near end crosstalk (NEXT)

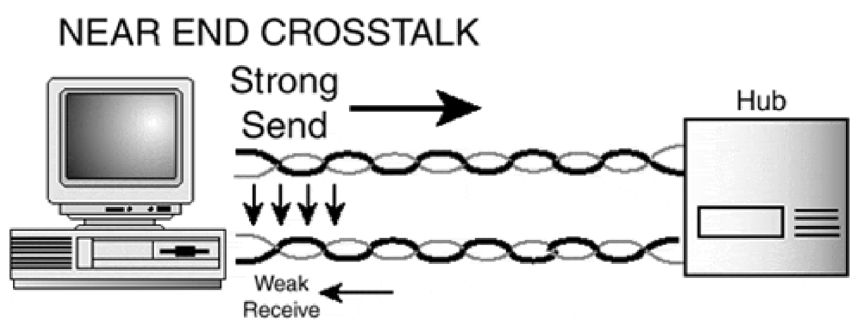

Um sinal forte, na saída da interface

Um sinal fraco chegando, mesmo assim introduz um FEXT no par de cobre vizinho

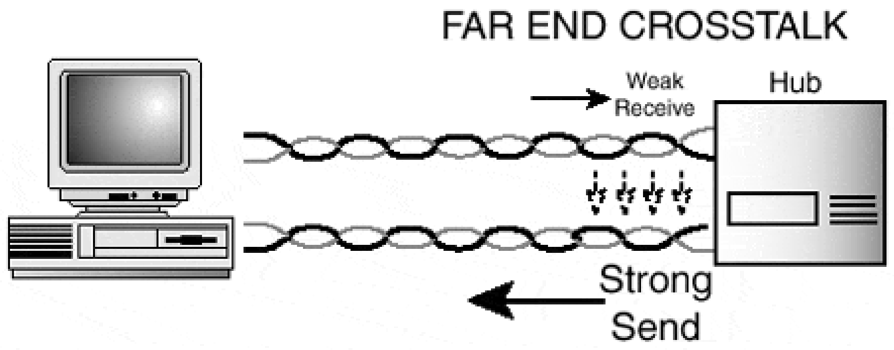

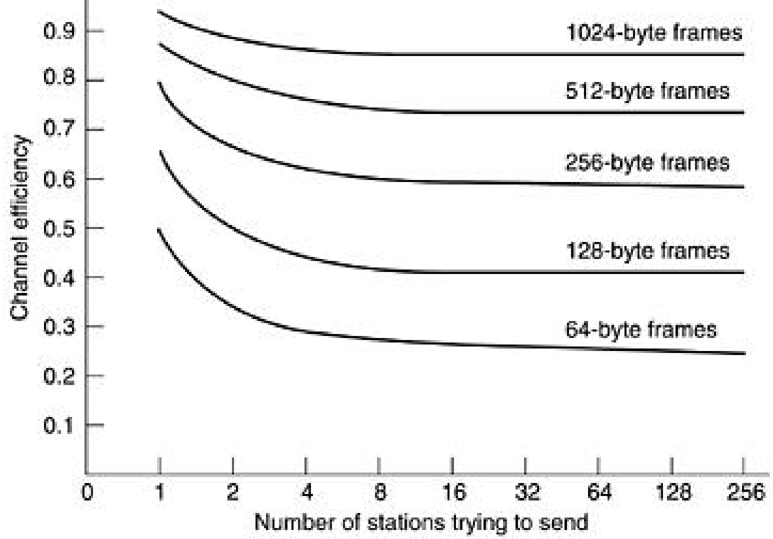

Sinalização em UTP cat 3

- Com o advento do par trançado, foi possível usar um par para transmitir e outro para receber.
- Ainda não foi possível transmitir em full-duplex,
- pois o padrão UTP cat 3 emulava a transmissão em banda básica (coax)

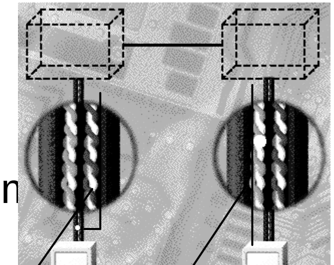

Repeaters

Um segmento de LAN deve ser limitado para que o sinal seja sempre legível
O tamanho da LAN pode ser aumentado com o uso de repeaters

Um repeater recebe os zeros e ums de um segmento e os propaga para dentro de outros segmentos de cabo com força renovada
Um repetidor é um dispositivo de camada física
Nota:

Hub ou concentradores são outros nomes dos repeaters
Esses nomes são usados para repeaters que conectam mais de dois segmentos

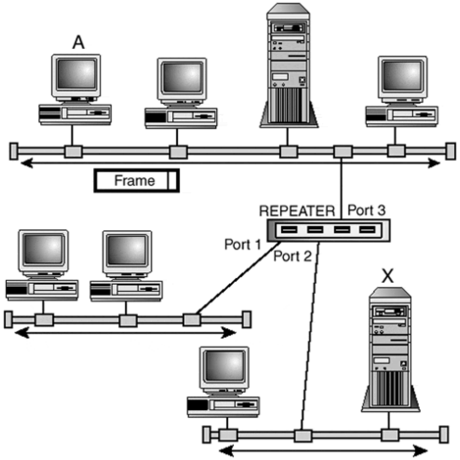

UTP e o hub

- The change to twisted-pair cabling and a star topology gave hubs a central role. The hub in Figure 3.3 repeats the bits in a frame transmitted by station A to all the other stations connected to the hub.
- ***Figure 3.3. Propagating a frame to all stations.***
- An important feature of this topology is that there are only two nodes on each cable segment: a computer and the repeater. There are four segments in Figure 3.3, and each connects a computer to the hub.
- **Note**
- All that the twisted-pair hub in Figure 3.3does is to repeat bits from one segment onto another. So, why isn't the device called a twisted-pair repeater?
- The answer is "marketing." Early repeaters connected two or three coax segments. Vendors wanted to describe their products using a word that indicated that their products were new. The "hub" products looked different and connected many twisted-pair segments.
- Today, the term "hub" sometimes is used for a chassis that contains repeaters along with other network devices. This is unfortunate because it causes confusion and damages the meaning of a convenient piece of networking terminology.
- In this book, "hub" will be used to denote a repeater device.

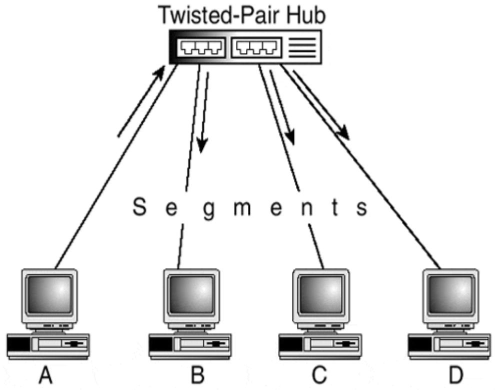

- **Improving Performance with Bridges**
- The Ethernet specifications allow 1,024 stations to be attached to a LAN. But when Ethernet bandwidth is a shared commodity, poor performance will make users groan and gnash their teeth long before anywhere near that number of systems has been attached to their LAN.
- A *bridge* is a Layer 2 device that takes a lot of the pain out of LAN growth.The three LAN segments in Figure 3.4 are connected by a bridge. Whenever possible, a bridge blocks frames from reaching segments that have no realneed to carry them.
- ***Figure 3.4. Three segments connected by a bridge.***

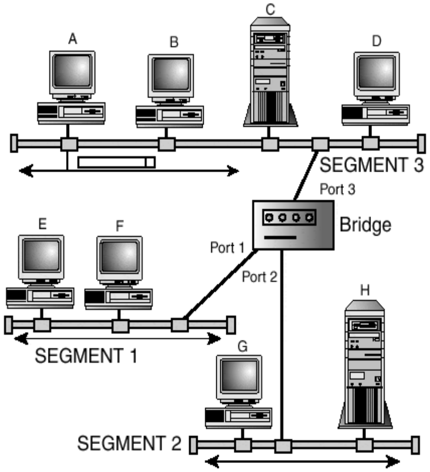

Collision Domains

- The LAN segments in Figure 3.2 and Figure 3.3 are connected to one another by repeaters. For either of these LANs, a frame transmitted by any station on the LAN will be seen by all stations on the LAN. If two stations send at the same time, their frames will collide. For this reason, a set of segments connected by repeaters is called a *collision domain*.
- **Note**
- At any given time, one station (at most) in a collision domain can transmit a frame successfully. If one station talks, the others should be listening. For this reason, CSMA/CD also is called half-duplex Ethernet transmission.
- The bridge in Figure 3.4 splits its LAN into three separate collision domains. Frames with local destinations do not need to be forwarded onto other segments. When a frame must be forwarded, the bridge behaves like a good citizen of the destination collision domain: It listens to check whether the destination medium is available before transmitting the frame. If the medium is busy, the bridge can hold the frame in buffer memory until the medium becomes quiet again.
- Administrators appreciated the fact that a bridge could be installed by hooking up the cables and plugging in the power. A bridge eavesdrops on traffic originating on each segment to discover the MAC addresses of the stations on the segment. The bridge creates a table that maps each MAC address to the port through which it is reached. When a MAC address is listed in the bridge table, the bridge will be capable of forwarding frames addressed to that MAC address onto the correct segment. Chapters 12-17 contain a full description of what bridges do and how they do it.

- **Leaping to Higher Performance with Switches**
- In 1993, a company named Kalpana introduced LAN *switches*. This was an astonishing event, startling many LAN experts who wondered why they had not thought of it first. A switch is a multiport bridge that can forward several frames at the same time.
- Eventually, LAN switches evolved into the popular *Layer 2 switches* in use today. Figure 3.5 shows a set of stations connected to a Layer 2 switch. Twisted-pair or fiber optic cable are the media normally used with a Layer 2 switch.

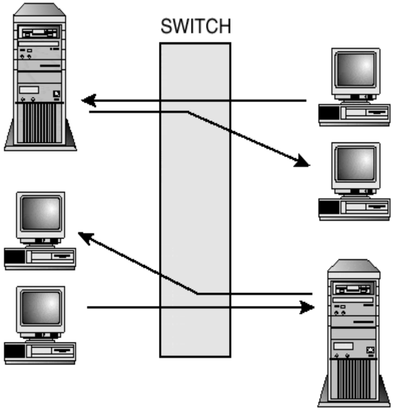

- In Figure 3.5, each link connecting a station to a switch is a separate segment that is bridged to all other segments. The internal architecture of the switch allows many frames to be in transit at the same time, which greatly increases the LAN bandwidth.
- **Note**
- Just as "twisted-pair hub" is an up-market name for a modern repeater, "Layer 2 switch" is an up-market name for a modern bridge. A Layer 2 switch can have numerous ports and uses better hardware technology than the bridges of long ago, but functionally, it is a bridge.

10BASE-T Collision Domain Parameters

Parameter or Characteristic	**Value**
Topology	A tree of stars
Patch cords	Up to 10 meters total
Type of segment cable	2 twisted pairs, Category 3 or better
Connector	RJ-45 jack and plug
Maximum length of a segment	100 meters
Maximum number of nodes that can be attached to a segment	2
Maximum collision domain diameter	500 meters
Maximum number of cable segments traversed between a source and destination	5
Maximum number of repeaters (hubs) between	4

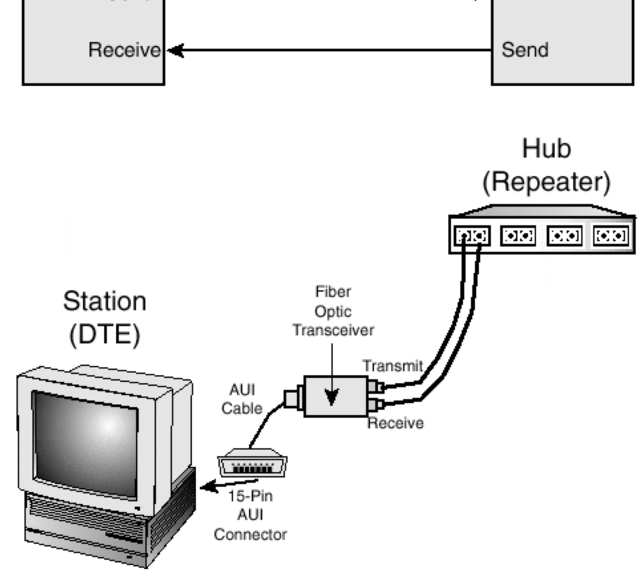

Switch

- Usando-se switch, conseguimos transmitir em diferentes domínios de colisão.
- Isso porque cada quadro de entrada é armazenado em um buffer em uma placa de linha no switch.

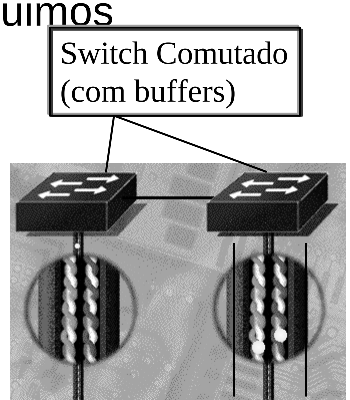

- The advantages of using switches include
- Full-duplex operation for each port
- Bandwidth proportional to the number of ports
- No collisions
- No 5-4-3 rule

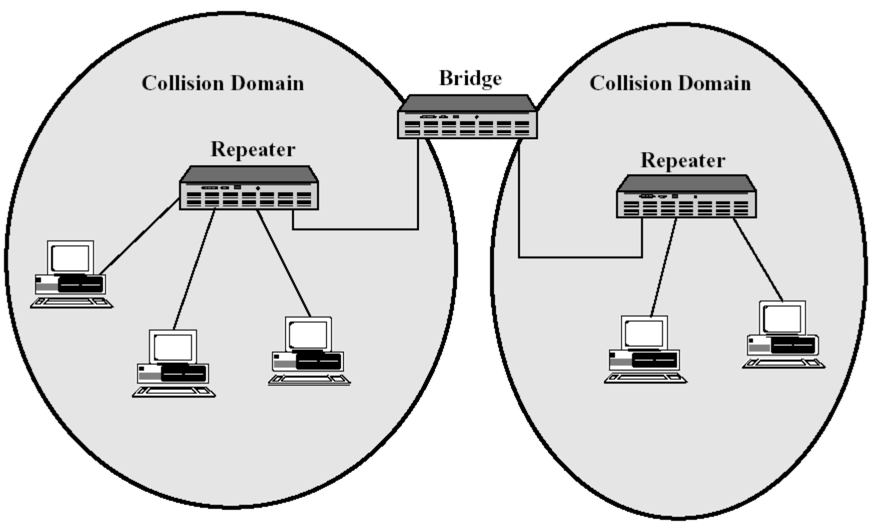

- Unindo Switch com UTP cat 5
- Os “Hubs” e placas ficam mais caros, mas todas as estações podem transmitir e receber ao mesmo tempo
- Além disso, a ausência de colisão permite uma distãncia maior (até 2 Km com fibra).

- **Switches and Full-Duplex Operation**
- The appeal of multiport switches was enhanced by an addition to the 802.3 standard that gave performance a big boost. As shown in Figure 3.6, when a single station is connected to a switch port, the link between the station and the switch can be used for full-duplex communication. Both systems can transmit and receive at the same time. For example, across a 100Mbps link, the station can transmit frames to the switch at 100Mbps and receive frames from the switch at 100Mbps.

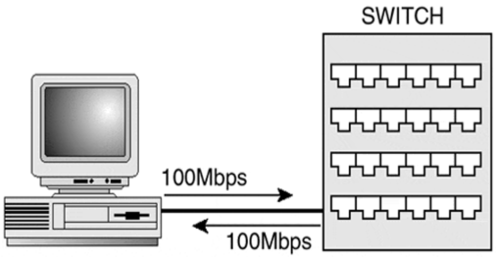

- The full-duplex Ethernet MAC protocol is very simple: Either party can send a frame whenever it pleases. CSMA/CD is not needed when full-duplex communication is used.
- **Note**
- Full-duplex communication can be used between any systems that are not repeaters. For example, a full-duplex link can be set up between two hosts, two switches (bridges), a switch and a router, or a pair of routers.

- Serve para backbones
- Ligações entre servidores de missão crítica

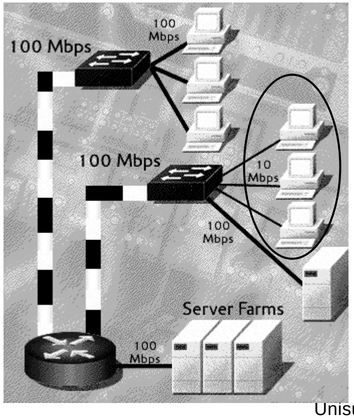

Requer upgrade de Hubs e Placas. Upgrade também de cabos para categoria 5

Sharing the LAN

- LANs are data communications workhorses. A LAN frame can carry traffic that belongs to any type of higher-layer protocol. It is not at all unusual to see a mixture of protocols such as TCP/IP, NetWare IPX/SPX, DECnet, and AppleTalk happily sharing a LAN. Figure 3.7shows how higher-layer protocols ride on top of Layer 1 and Layer 2 LAN protocols. Many hosts send and receive traffic for several different protocols through a single LAN adapter

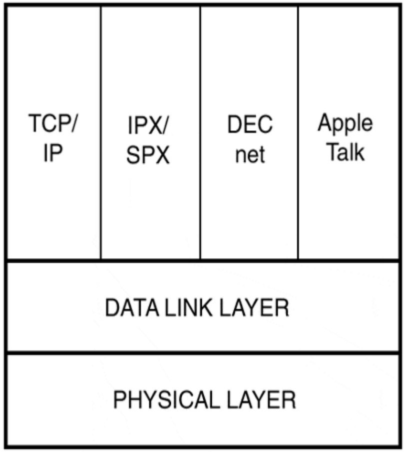

The Role of Routers

- It did not take long before users wanted to exchange information with servers located on external LANs dotted across an enterprise—and, with the advent of the Internet, with servers located around the world. Routers make this communication possible.
- Figure 3.8 shows two LANs connected by a router. The router also connects these LANs to a long-distance line. This could lead to another site within a company or to the Internet.

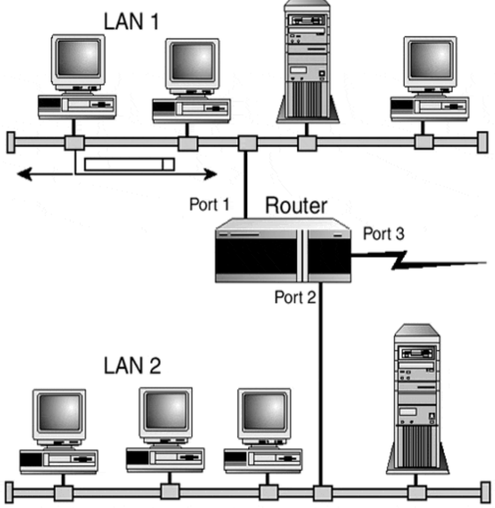

router is a Layer 3 device.

- A router is a Layer 3 device. Routers have lots of good features. To mention just a few:
- They can connect different types of LANs gracefully.
- Unlike a bridge, a router does not forward LAN broadcast traffic or local LAN multicasts. A sizeable amount of bandwidth can be saved. This is particularly important when the peak load on one of the LANs is close to the LANs capacity.
- They can connect a set of LANs to a set of WAN links of differenttypes—for example, dial-up, leased line, frame relay, or Asynchronous Transfer Mode (ATM).
- They can perform security screening and keep risky traffic off a LAN.
- An important thing to keep in mind is that when data from a LAN reaches a router, it has passed through a doorway and left the LAN. The router strips off the MAC frame header and trailer. The protocol data will be encapsulated in a new header and trailer before it is forwarded.

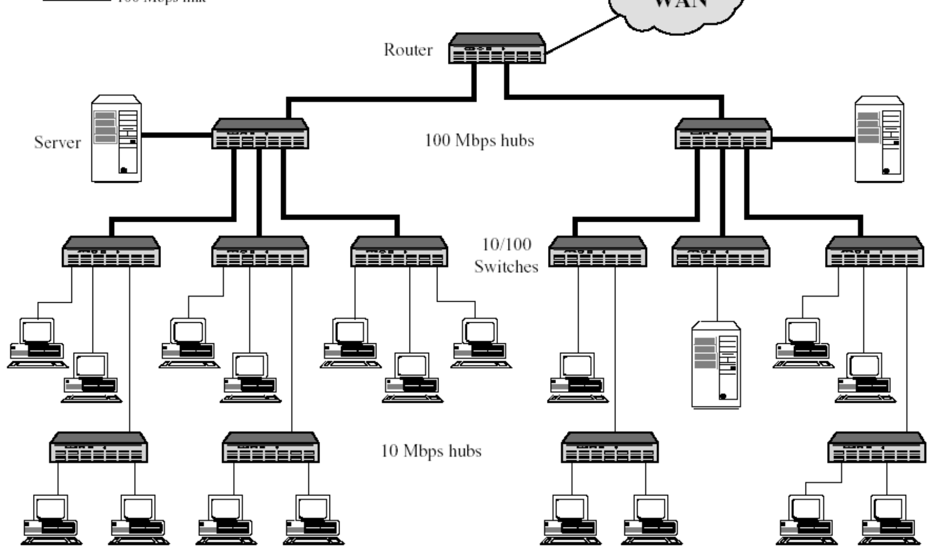

Backbone Gbps

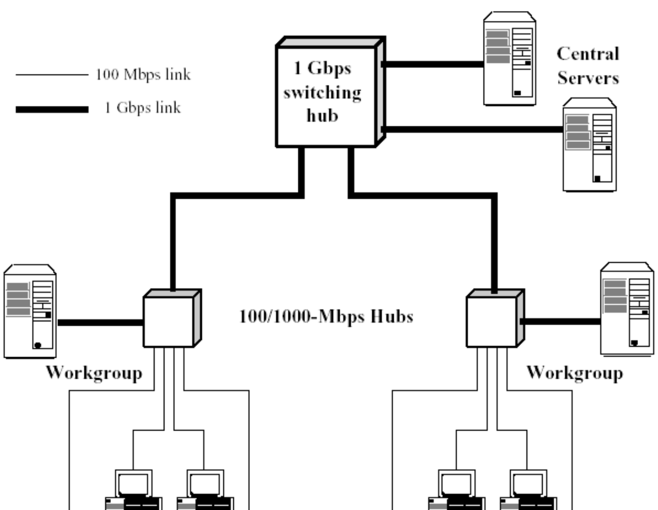

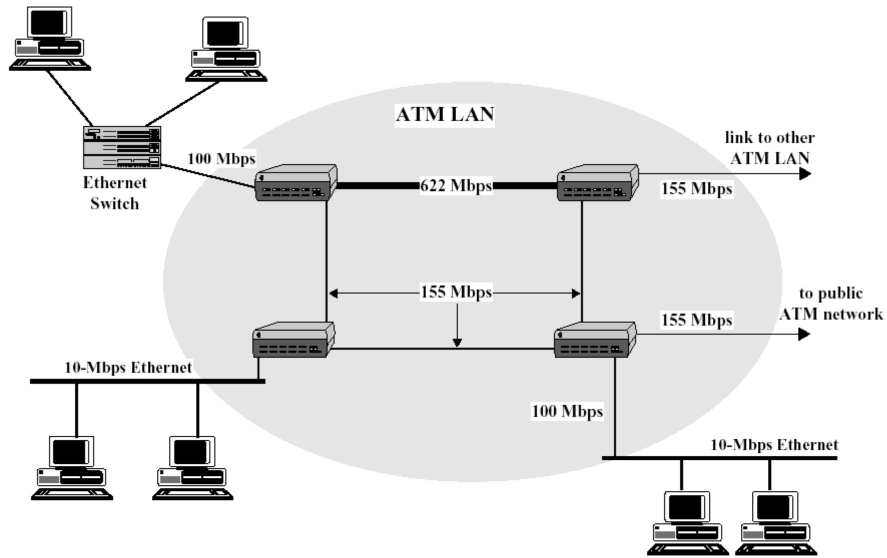

Mix ATM e outros

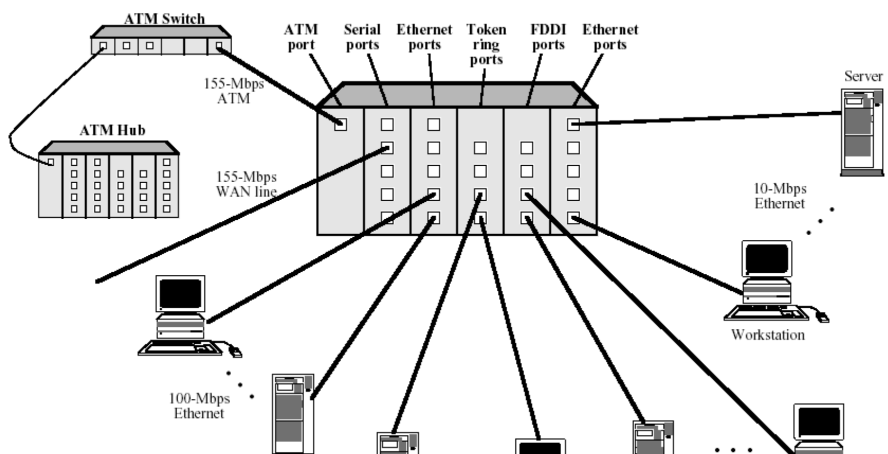

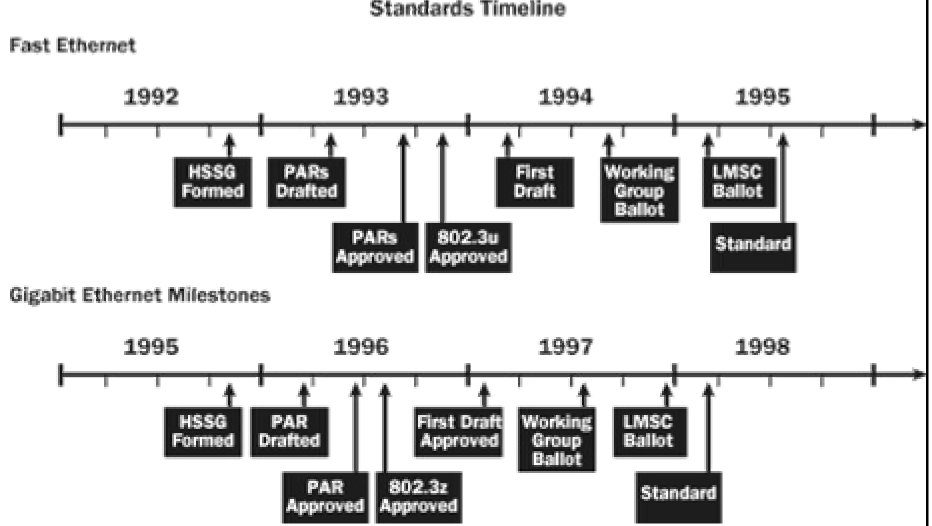

Classes de serviços com ETH

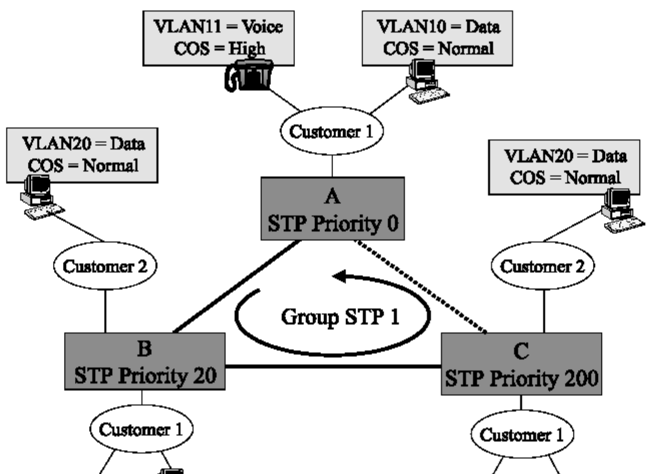

Quadro IEEE 802.3 sem cabeçalho de tag VLAN

Preâmbulo/SFD	Endereço de destino	Endereço de origem	Tamanho/ tipo	Dados	FCS

Preâmbulo/SFD	Endereço de destino	Endereço de origem	TPID	TCI	Tamanho/ tipo	Dados	FCS

Quadro IEEE 802.3 com cabeçalho de tag VLAN com 4 bytes

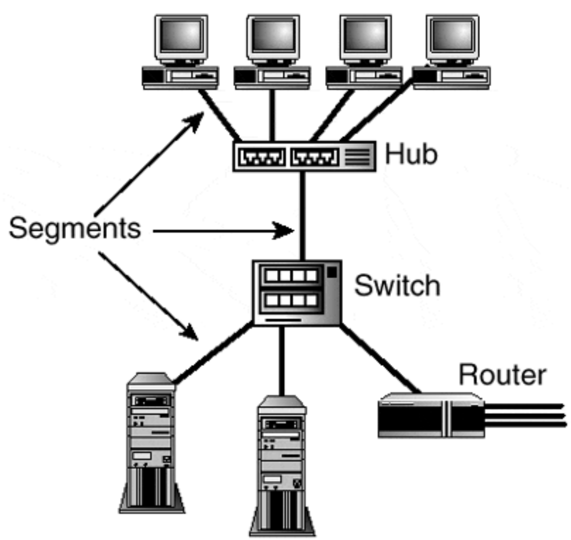

10baseT e os segmentos