



The University of Sydney
AUSTRALIA

School of Electrical and Information Engineering

Advanced Communication Networks

Chapter 9

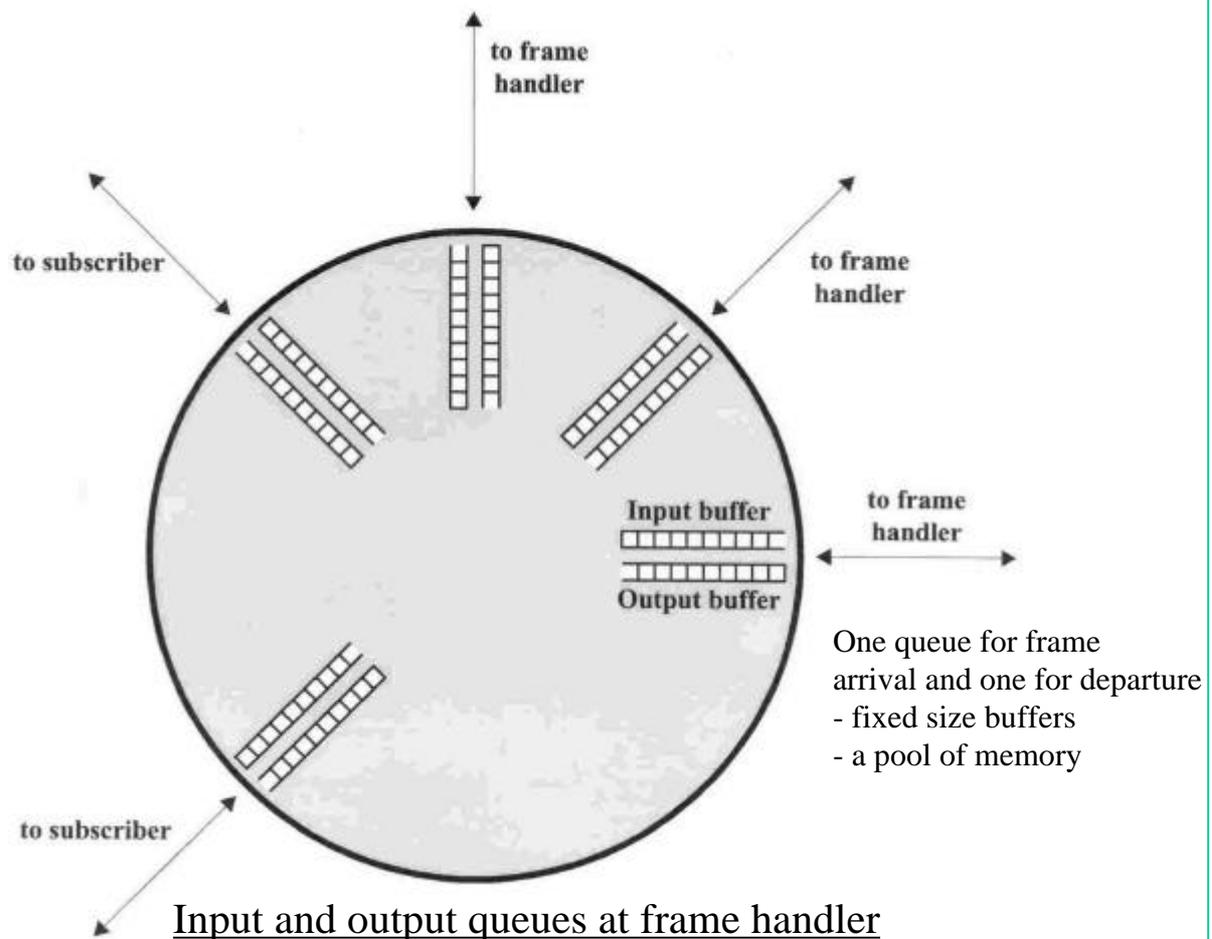
Frame Relay Congestion Control

Based on chapter 13 of Stallings ISDN-4e book

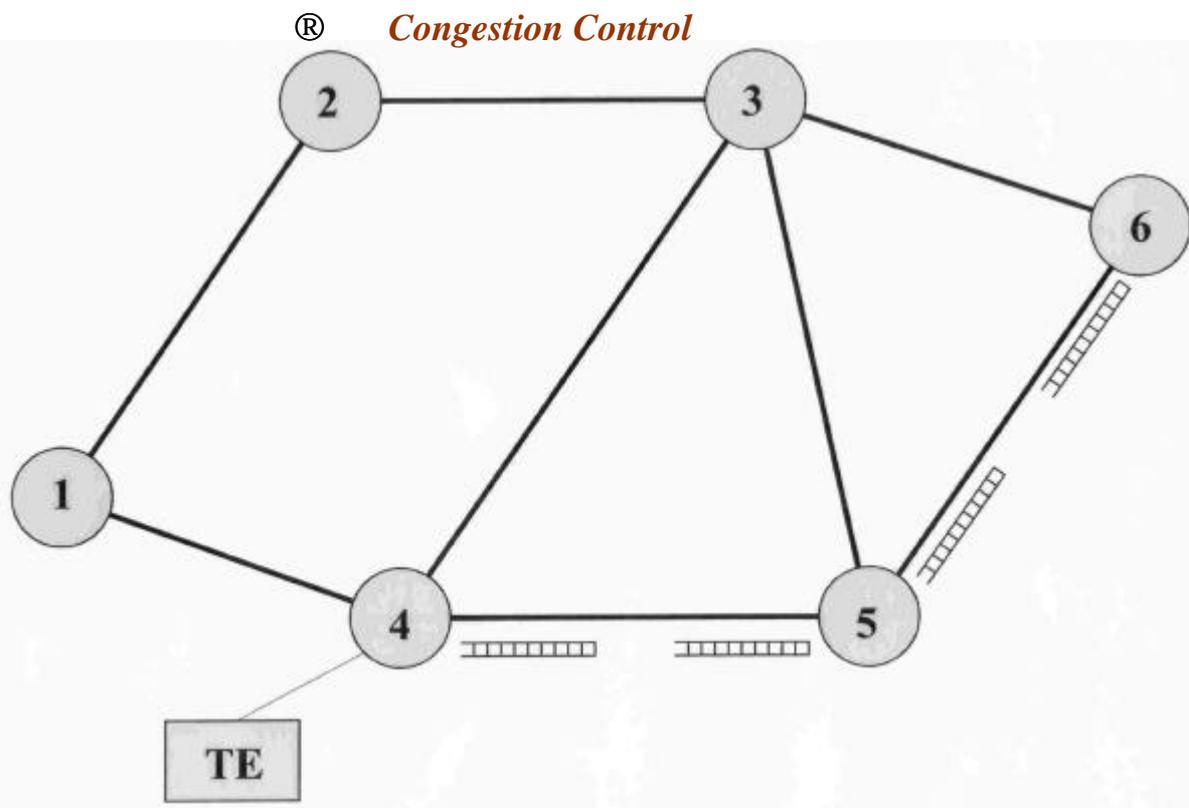
Abbas Jamalipour

9.1 Congestion in Frame Relay Networks

- With frame relay there is no flow control and error control mechanism between user and network.
 - LAPF frame has no control field and hence no sequence number.
 - It is an efficient data transfer but network can be congested.
 - ITU-T and ANSI have proposed a variety of congestion control techniques.
- A frame relay network is a form of packet-switching network with the *packets* in layer 2 frames.
 - Thus, a key network design issue is the *congestion control*.
 - Thus, a frame relay network is a network of queues.
 - Queues are in frame handlers for each outgoing link.
 - When the arrival rate of frames $>$ transmission rate of frames \rightarrow infinite delay
 - Even if the arrival rate of frames $<$ transmission rate of frames \rightarrow growing queue length when the arrival rate approaches the transmission rate.

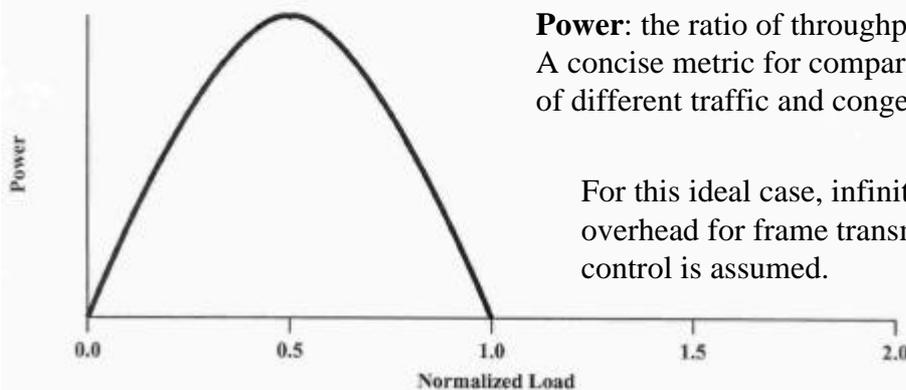
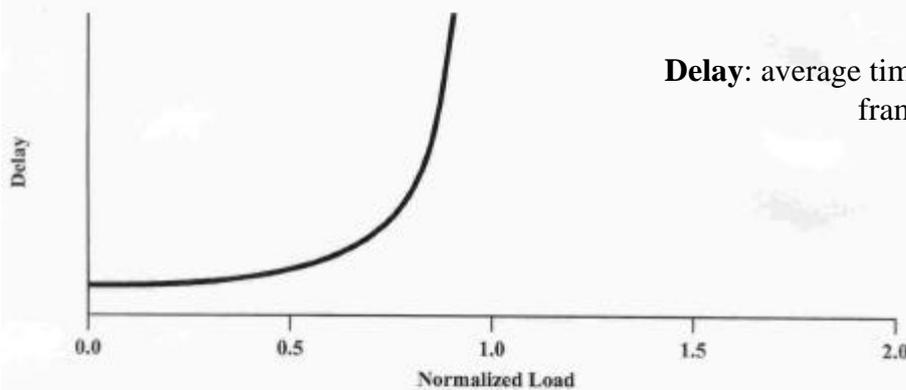
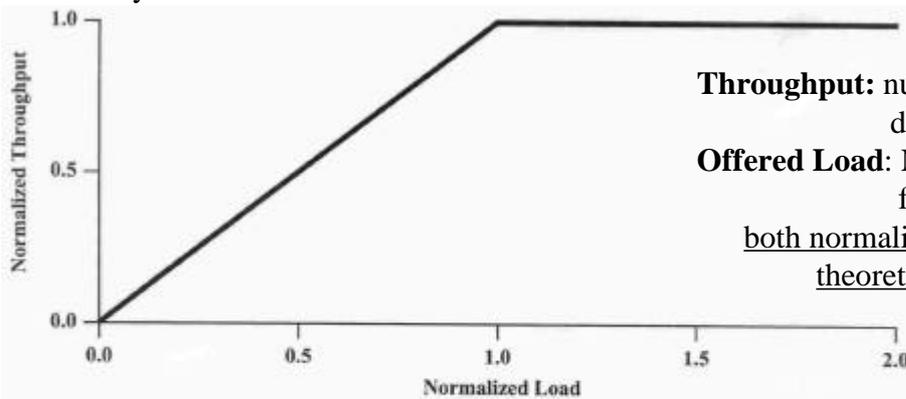


- If frames arrive too fast for the node to process them (e.g. routing) or faster than frames can be cleared from the outgoing buffers, then eventually frames will arrive for which no memory is available.
- At such a saturation point, there are two general strategies:
 - discarding any incoming frame when there is no buffer space available
 - a self-defeating approach
 - retransmission of the discarded frames add to the network congestion
 - some mechanism to restrict the rate of inserting new frames into the network
 - e.g., enabling each node experiencing queue saturation to do some sort of flow control over its neighboring for managing the traffic
 - In the figure below, if node 6 restrain the flow of frames from node 5, this causes the output buffer in node 5 for the link to node 6 to fill up.
 - Thus, congestion quickly propagates throughout a region or all of the network.
 - Therefore, need for a more sophisticated flow control



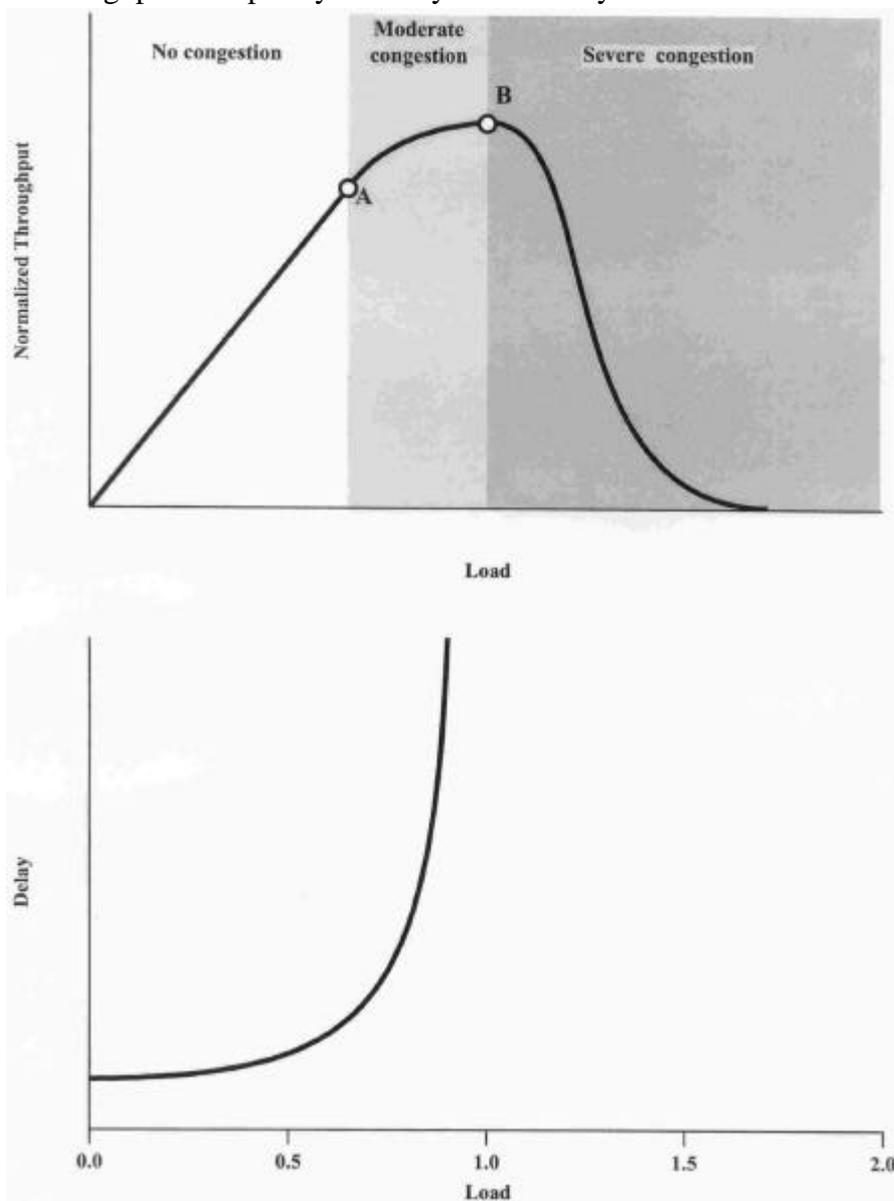
Ideal Network Performance

- At low load, delay is only propagation delay + processing delay of nodes
- At high load, delay includes also the queuing delays
- Even in an ideal network, if the input load exceeds network capacity, the queue sizes grow toward infinity and hence infinite queuing delay.
- The ideal performance in figure below are the goal of all traffic and congestion control techniques. They are upper bound limits of network performance.
- Typically, a congestion control that results in higher throughput also results in higher delay.



Practical Network Performance

- In a real, buffers are finite (buffer overflow is possible) and congestion control requires usage of network capacity for exchange of control signals.
- Without any congestion control, as the traffic load increases moderate congestion happens to some nodes and severe congestion to some others. For balancing the traffic network reroute the frames: more signaling and overhead and less data frame capacity
- With finite buffers, when buffers become full, the new frames will be discarded and more retransmission happens: capacity tends to become zero!
- Congestion control techniques are required: to limit queue lengths at the nodes so as to avoid throughput collapse by so many unnecessary retransmissions.



9.2 Approaches to Congestion Control

- I.370: Objectives for frame relay congestion control
 - Minimize frame discard
 - Maintain, with high probability and minimum variance, an agreed QoS
 - Minimize the possibility that one end user can monopolize network resources at the expense of other end users
 - Be simple to implement, and place little overhead on either end user or network
 - Create minimal additional network traffic
 - Distribute network resources fairly among end users
 - Limit spread of congestion to other networks and elements within the network
 - Operate effectively regardless of the traffic flow in either direction between end users
 - have minimum interaction or impact on other systems in the frame-relaying network
 - Minimize the variance in QoS delivered to individual frame relay connections during congestion
- Congestion control in a frame relay network is difficult because the frame relay protocol streamlined to maximize throughput and efficiency.
 - Frame handler cannot control frames flow using typical sliding-window flow control protocol, such as found in LAPD.
- Therefore, congestion control becomes responsibility of both the network and the end users:
 - network monitors the degree of congestion
 - end users control the congestion by limiting the flow of traffic
- Congestion control techniques:

Discard Strategy, Congestion Avoidance, Congestion Recovery

Technique	Type	Function	Key Elements
Discard control	Discard strategy	Provides guidance to network concerning which frames to discard	DE bit
Backward explicit Congestion Notification	Congestion avoidance	Provides guidance to end systems about congestion in network	BECN bit or CLLM message
Forward explicit Congestion Notification	Congestion avoidance	Provides guidance to end systems about congestion in network	FECN bit
Implicit congestion notification	Congestion recovery	End system infers congestion from frame loss	Sequence numbers in higher-layer PDU

9.3 Traffic Rate Management

- The last option in any congestion control scheme is to discard frames.
 - Because the frame handlers' buffers are finite.
 - The simplest way is to discard frames arbitrarily w/o regard to their source
- To have a fairer resource management, frame relay use CIR
 - CIR is the “committed information rate”: the rate in bps at which the network agrees to support for a particular frame-mode connection.
 - Data transmission at a rate $>$ CIR should be discarded in the case of congestion
 - In some situation, network may provide less data rate than CIR!
 - The frame relay node should manage so that the aggregate of CIRs of all connections of all attached end systems to the node does not exceed the capacity of the node.
 - Moreover, the aggregate of the CIRs should not exceed the physical data rate across the user-network interface—**The Access Rate**

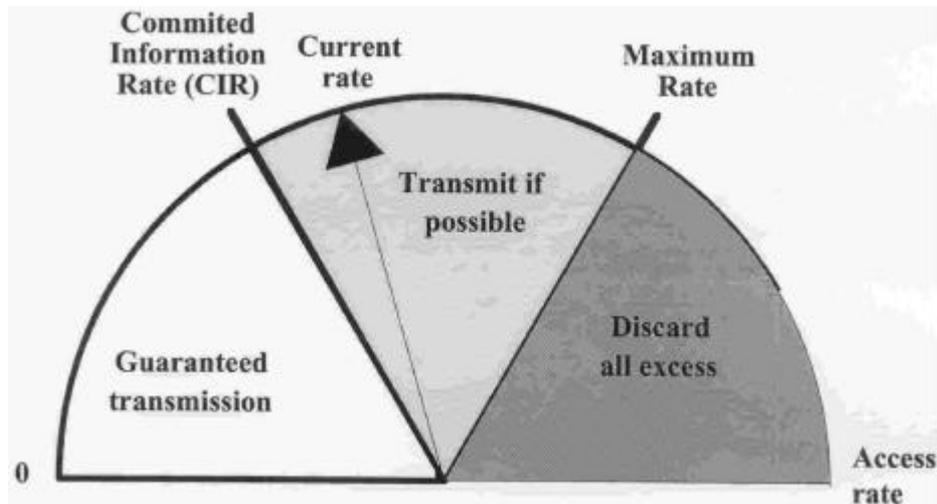
$$\sum \text{CIR}_{i,j} \leq \text{AccessRate}_j$$

$\text{CIR}_{i,j}$: Committed information rate for connection i on channel j

AccessRate_j : Data rate of user access channel j (D, B, or H)

- For Permanent frame relay connections, CIR is agreed at the time of establish
- For switched connections, the CIR parameter is negotiated.
 - Link layer Core Parameters Information Element includes throughput and minimum acceptable throughput parameters.
 - In SETUP message, throughput parameter represents the requested TH by calling user.
 - Called user or network can reduce this and return the CIR in the CONNECT message.
 - CIR may not be less than the minimum acceptable TH declared by calling user.
 - If cannot support this level of TH, the connection request is rejected.
- CIR is a way of discriminating among frame in determining which frame to discard in the case of congestion.
- CIR does not provide much flexibility in dealing with traffic rates.
- A frame handler measures traffic over each logical connection for a time interval specific to the connection and then makes a decision based on the amount of data received during that interval.
- Two additional parameters, assigned on permanent connections and negotiated on switched connections, are required:
 - **Committed Burst Size (B_c)**
 - **Excess Burst Size (B_e)**

- The frame handler to which the user's station attaches performs a metering function.
 - If the user is sending data at less than the CIR, the incoming frame handler does not alter the DE bit.
 - If the rate exceeds the CIR, the incoming frame handler will set the DE bit on the excess frames and then forward them.
 - A maximum rate is also defined, such that any frames above the maximum rate are discarded at the entry frame handler.

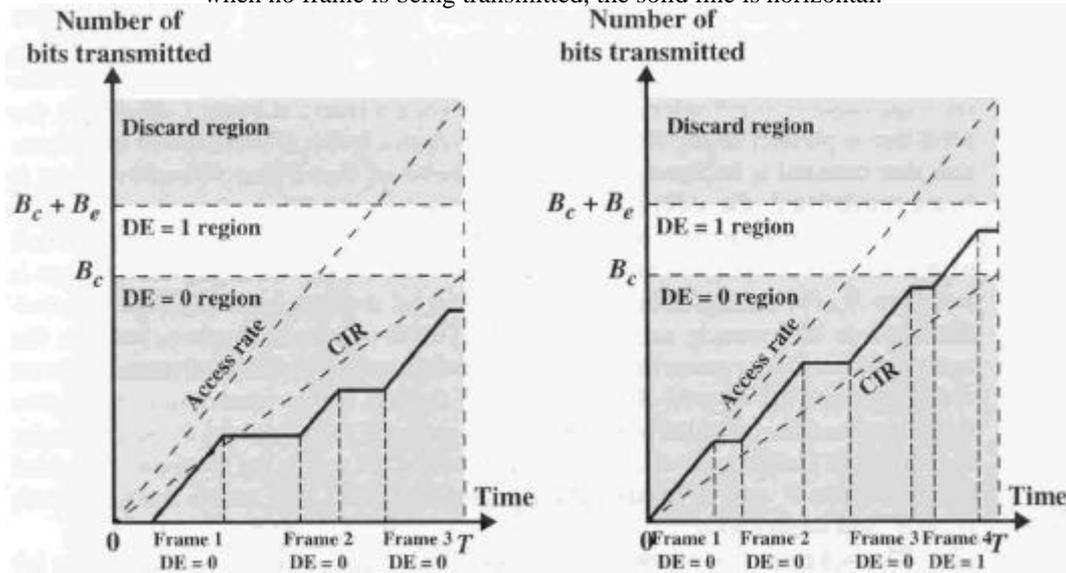


- **Committed Burst Size (B_c)**
 - is the maximum amount data that the network agrees to transfer, under normal conditions, over a measurement interval T . These data may or may not be contiguous (appear in one or several frames)
- **Excess Burst Size (B_e)**
 - is the maximum amount of data in excess of B_c that the network will attempt to transfer, under normal conditions, over a measurement interval T . These data are uncommitted. The data that represent B_e are delivered with lower probability than the data within B_c .
- The quantities B_c and CIR are related.
 - B_c is the amount of committed data for transmission over T
 - CIR is the rate at which committed data may be transmitted.
 - hence

$$T = B_c / CIR$$

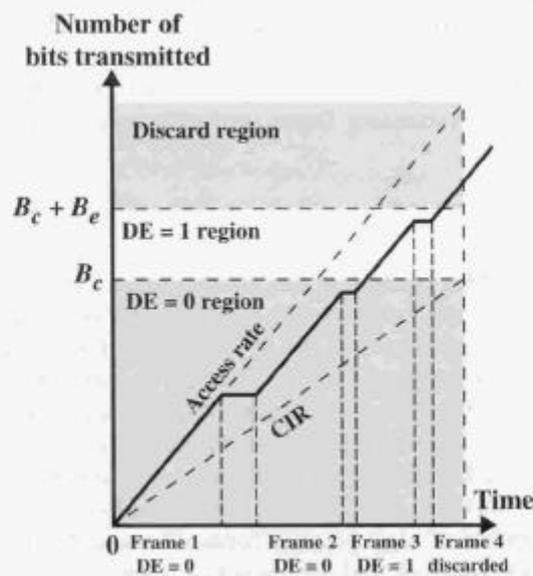
Relationships Among Congestion Parameters

- *solid lines*: cumulative number of information bits transferred over a given connection since time $T = 0$
- *dashed lines* (access rate): data rate over the channel containing the connection
- *dashed lines* (CIR): combined information rate over the measurement interval T
 - when a frame being transmitted, solid line is parallel to Access Rate line. When a frame is being transmitted, that channel is dedicated to the transmission of of that frame.
 - when no frame is being transmitted, the solid line is horizontal.



(a) All frames within CIR

(b) One frame marked DE



(c) One frame marked DE; one frame discarded

(a) All frames within CIR

- Three frames are transmitted within T , the total no of bits $< B_c$
- If the actual transmission rate exceeds CIR, there is no problem since the frame handler is only concerned with the cumulative no. of bits over T .

(b) One frame marked DE

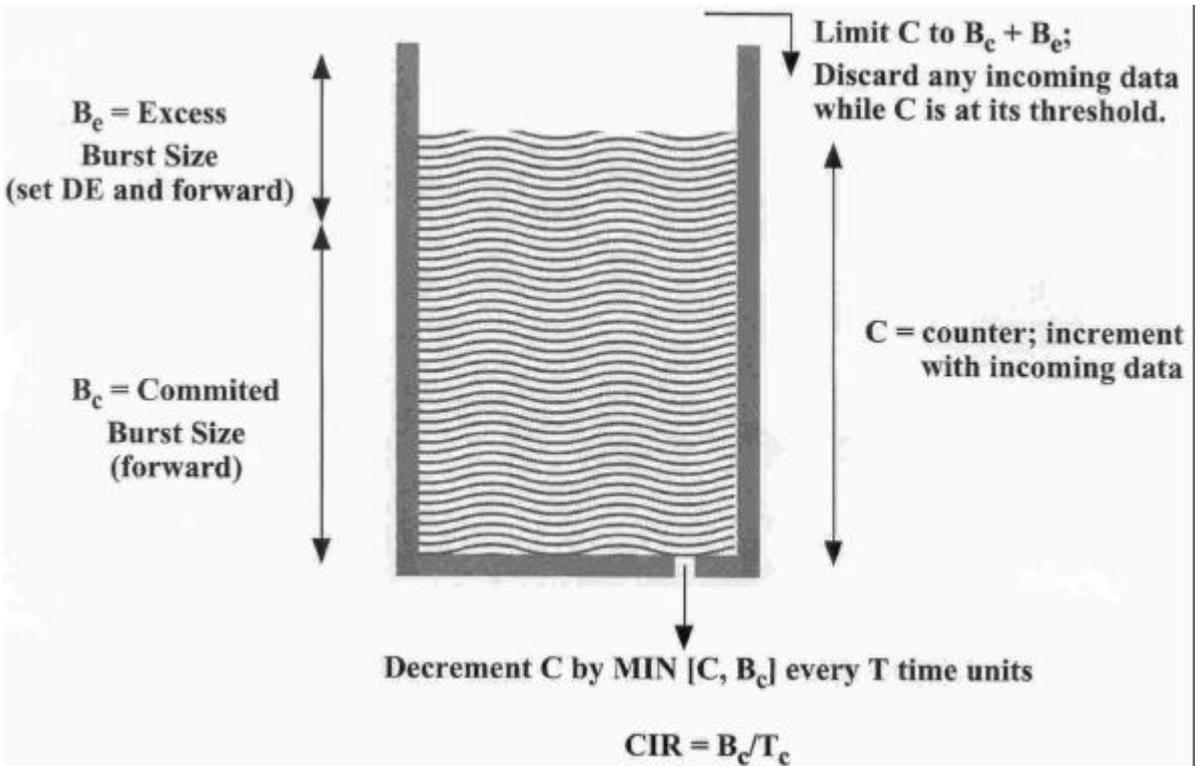
- The last frame transmitted during the interval causes the cumulative no of bits transmitted to exceed B_c
- Thus, the DE bit of that frame is set by the frame handler.

(c) One frame marked DE, one frame discarded

- The third frame exceeds B_c and so is labeled for potential discard.
- The fourth frame exceeds $B_c + B_e$ and is discarded.

• This is an example of a **Leaky Bucket** algorithm (figure below).

- Frame handler records the cumulative amount of data sent over a connection in counter C . The counter is decremented at a rate of B_c every T time units. Whenever the counter exceeds B_c but is less than $B_c + B_e$, incoming data are in excess of the committed burst size and are forwarded with the DE bit set. If the counter reaches $B_c + B_e$ all incoming frames are discarded until the counter has been decremented.



CIR Levels

- Priority allocation to different traffic streams multiplexed over the same access channel is possible by the use of CIR values.
 - An application with higher throughput requirements can be assigned a larger CIR than an application with a lower throughput requirement.
 - Total data rate allocated must be less than access rate of the channel.
 - We can establish a frame relay connection with CIR=0
 - no commitment of network on delivery of frames on that connection
 - a typical scenario is to have a number of high-priority connections with their required CIRs multiplexed with a several lower-priority with CIR=0
 - In the case of CIR=0, the measurement intervals is defined by

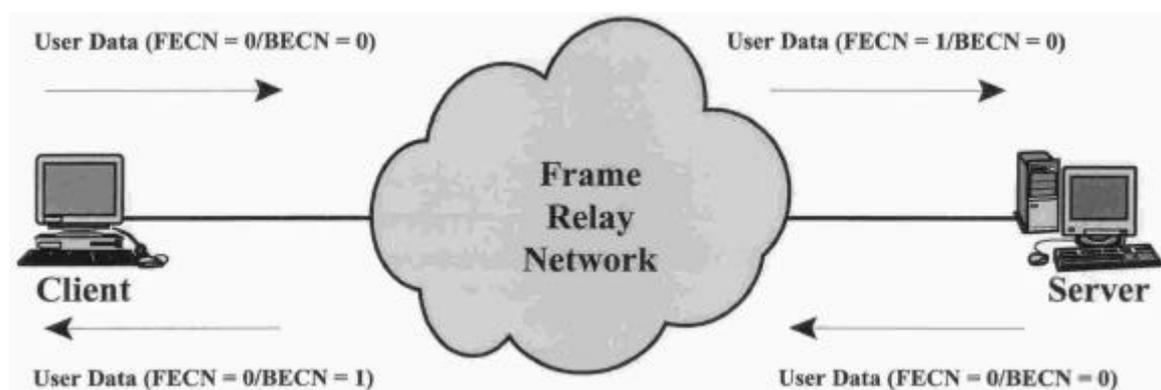
$$T = B_e / \text{Access Rate}$$

End-User Control of DE Bit

- DE bit may also be set by the end system that originated the frame transmission.
- Allows the user to decide which frames are less important and should be more vulnerable to discard.
- When a user marks a frame with the DE bit, the network considers the frame to be within B_e but not part of the CIR.

9.4 Explicit Congestion Avoidance

- Purpose
 - To use as much of the available capacity in a frame relay network as possible having congestion controlled in a fair manner.
- How
 - Network alerts end systems to growing congestion in network and the end systems take steps to reduce the offered load to the network.
- When network detects increasing congestion, signals to those end users with frame relay connection that are affected by congestion.
 - Using one of two bits in LAPF address field of each frame,
 - or, by using a special message in an XID LAPF frame.
 - The bits constitute signals from network to end user.
 - **Backward Explicit Congestion Notification (BECN) bit**
 - notifies the user that congestion avoidance procedures should be initiated where applicable for traffic in the *opposite direction* of the received frame. Indicates that the frames that the user transmits on this logical connection may encounter congested resources.
 - **Forward Explicit Congestion Notification (FECN) bit**
 - notifies the user that congestion avoidance procedures should be initiated where applicable for traffic in the *same direction* as the received frame. Indicates that this frame on this logical connection has encounter congested resources.
 - **Consolidated Link Layer Management (CLLM) message**
 - notifies the user that congestion avoidance procedures should be initiated where applicable for traffic in the *opposite direction* of the received frame. Indicates that the frames that the user transmits on this logical connection may encounter congested resources.



Network Notification of Congestion

- Network detects and signals congestion using monitoring of queue size for each frame handler.
- Frame relay monitors size of each of its queues.
 - During a cycle begins when outgoing circuit goes from idle (empty queue) to busy.
 - Averaging queue size during the previous cycle and current one
 - If the average size > a threshold value, the congestion avoidance bits should be set on some or all logical connections that use that circuit.
 - Having the average on two cycles instead of one, the system avoids reacting to temporary surges that would not necessarily produce congestion.

The algorithm makes use of the following variables:

t = current time
 t_i = time of i^{th} arrival or departure event
 q_i = number of frames in the system after the event
 T_0 = time at the beginning of the previous cycle
 T_1 = time at the beginning of the current cycle

The algorithm consists of three components:

1. Queue Length Update: Beginning with $q_0 := 0$
If the i^{th} event is an arrival event, $q_i := q_{i-1} + 1$
If the i^{th} event is a departure event, $q_i := q_{i-1} - 1$

2. Queue Area (integral) update:

$$\text{Area of the previous cycle} = \sum_{t_i \in [T_0, T_1]} q_{i-1}(t_i - t_{i-1})$$

$$\text{Area of the current cycle} = \sum_{t_i \in [T_1, t]} q_{i-1}(t_i - t_{i-1})$$

3. Average Queue Length Update

Average queue length over the two cycles

$$= \frac{\text{Area of the two cycles}}{\text{Time of the two cycles}} = \frac{\text{Area of the two cycles}}{t - T_0}$$

Forward Explicit Congestion Notification

- After receiving the FECN bit set by the end user, the receiving system should try to reduce the flow of data from the sending system.
- The mechanism for doing so must be above the level of frame relay, which provides no direct flow control facilities.
- Generally, the receiving end system should do as following
 - Compute the fraction of frames for which the FECN bit is set over the interval.
 - If more frames have the FECN bit set than have a FECN bit of zero, then reduce the flow of frames from the sending system.
 - If the congestion condition persists, institute additional reductions.
 - When the congestion condition ends, gradually increase the flow of frames.
- This strategy reacts slowly to congestion notifications because:
 - the end system doesn't react immediately to a particular FECN bit but waits until the average behavior of the system over an interval indicates congestion.
 - the end system doesn't immediately reduce its outgoing flow but rather signals its peers to reduce the incoming flow.
- Two methods of actual control of the information rate
 - **rate-based control** can provide a more precise control of information flow because it is based on the actual information rate in bits per second.
 - **window-based control** can provide only an approximate control over information rate, because frame relay does not require fixed-size frames.

See Rate-Based and Window-Based Control algorithms in pp. 360-361 of the text

Backward Explicit Congestion Notification

- This can be achieved with either the BECN bit in the LAPF address field or a consolidated link layer management message.
- After receiving the BECN bit set by the end system, the receiving system should try to reduce the flow of data transmitted on that connection.
- Generally, the receiving end system should do as following
 - When first frame with BECN bit set is received, reduce information rate to CIR
 - If additional consecutive frames with the BECN bit set are received, then institute additional reductions.
 - If a consecutive sequence of frames with the BECN bit set to zero are received, then gradually increase the flow of frames.

- This strategy reacts rapidly to congestion notifications because:
 - the end system react immediately to a single BECN bit.
 - the end system immediately reduce its outgoing flow rather than signaling its peers to reduce the incoming flow.
- Two methods of actual control of the information rate
 - **Rate-Based Control**
 - **Window-Based Control**

See Rate-Based and Window-Based Control algorithms in pp. 362-364 of the text

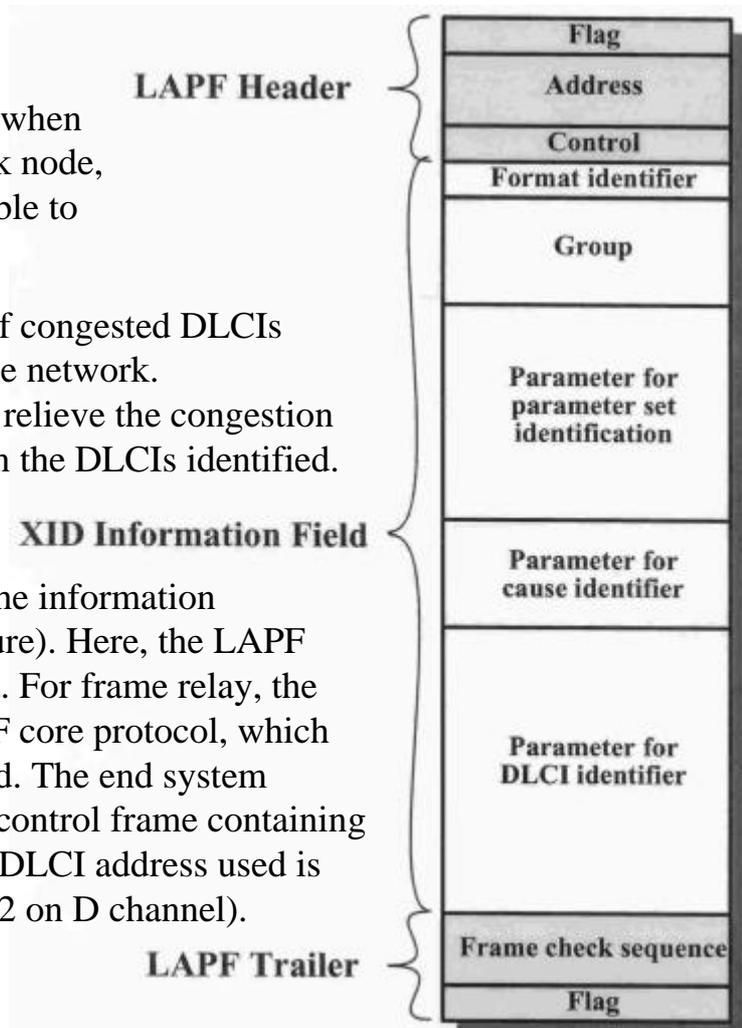
Consolidated Link Layer Management

- CLLM is a variation of backward explicit congestion notification that uses a message rather than the BECN bit to signal congestion.

CLLM technique can be used when congestion occurs at a network node, but no reverse traffic is available to carry BECN indication.

CLLM messages carry a list of congested DLCIs to reduce the traffic load on the network. The end station is expected to relieve the congestion by limiting the data transfer on the DLCIs identified.

CLLM message is carried in the information field of an XID LAPF frame (see the figure). Here, the LAPF header contains a control field. For frame relay, the network implements the LAPF core protocol, which does not include a control field. The end system recognize it as an XID LAPF control frame containing a CLLM by the address used. DLCI address used is decimal 1007 (eq. To SAPI=62 on D channel).



- The first field, format identifier marks XID as a CLLM message, followed by:
 - **Group:** indicates that this information field contains parameters beyond the scope of the HDLC-specific parameters; length: the length of following param.
 - **Parameter set identification:** indicates that this message contains parameters for I.122, which is the frame-mode bearer service.
 - **Cause identifier:** identifies the cause of this message as determined by the congested network node that originated the message.
 - **DLCI identifier:** a list of the DLCIs of the frame relay bearer connections that are congested.

Octet	8	7	6	5	4	3	2	1	Field name
1	0	1	1	1	1	1	1	0	Flag
2	1	1	1	1	1	0	R	0	Address octet 1 (R indicates response)
3	1	1	1	1	0	0	0	1	Address octet 2
4	1	0	1	0	1	1	1	1	XID control field
5	1	0	0	0	0	0	1	0	Format identifier (130)
6	0	0	0	0	1	1	1	1	Group identifier = 15 (private parameter negotiation)
7									Group length octet 1
8									Group length octet 2
9	0	0	0	0	0	0	0	0	Parameter identifier = 0 (parameter set identification)
10	0	0	0	0	0	1	0	0	Parameter length = 4
11	0	1	1	0	1	0	0	1	Parameter value = 105 (IA5 character "I")
12	0	0	1	1	0	0	0	1	Parameter value = 49 (IA5 character "1")
13	0	0	1	1	0	0	1	0	Parameter value = 50 (IA5 character "2")
14	0	0	1	1	0	0	1	0	Parameter value = 50 (IA5 character "2")
15	0	0	0	0	0	0	1	0	Parameter identifier = 2 (cause identifier)
16	0	0	0	0	0	0	0	1	Parameter length = 1
17									Cause value
18	0	0	0	0	0	0	1	1	Parameter identifier = 3 (DLCI identifiers)
19									Parameter length = 2n
20									DLCI value octet 1 (1st DLCI)
21									DLCI value octet 2 (1st DLCI)
.									.
.									.
.									.
2n + 17									DLCI value octet 1 (nth DLCI)
2n + 18									DLCI value octet 2 (nth DLCI)
2n + 19									FCS octet 1
2n + 20									FCS octet 2
2n + 21	0	1	1	1	1	1	1	0	Flag

-short term network congestion due to excessive traffic

-long term network congestion due to excessive traffic

-short term facility failure

-long term facility or failure

-short term maintenance action

-long term maintenance action

-unknown-short term

-unknown-long term

CLLM message for B or H channel

9.5 Implicit Congestion Control

- Implicit signaling occurs when the network discards a frame, and this fact is detected by the end user at a higher end-to-end layer such as Q.922.
- When this occurs, end user software may deduce that congestion exists.
- In a data link control protocol (e.g. Q.922) which uses a sliding window flow- and error-control, the protocol detects the loss of an I-frame by:
 - when a frame is dropped by the network, the following frame will generate a REJ frame from the receiving endpoint.
 - When a frame is dropped by the network, no acknowledgement is returned from the other end system. Eventually, the source end system will time out and transmit a command with P bit set to 1. The subsequent response with the F bit set to 1 should indicate that the receive sequence number N(R) from the other side is less than the current send sequence number.
- Once congestion is detected, the protocol uses flow control to recover from the congestion.
- Q.922 suggests that a user that is capable of varying the flow-control window size use this mechanism in response to implicit signaling.