Network Analysis Course

G. A. Marin

Course Goal: This is a course in data communication networking intended for students in CS and Engineering whose long-term goals may include design and development of networking hardware or software products as well as design and implementation of enterprise or service-provider networks. It introduces the concept of good network design using protocol and function layering. It introduces the principal system considerations such as cabling, modulation, topology, hardware, and software. It interweaves the corresponding control issues throughout including throughput, delay, bandwidth management, congestion control, error control, sliding windows, retransmission strategies, contention resolution. In order to do the latter it introduces a number of methods from Probability Theory, Statistics, and Queuing Theory that can be used to tackle these problems.

History

- Spring 2000 and Spring 2001 offerred the “Network Analysis” course out of Tannenbaum and concentrated on working text problems.
- Findings:
  - Many students did not know how to use a simple random variable, X, or compute E(X).
  - Virtually no students understood a simple Binomial distribution (certainly had no instinct for when/how to apply).
  - Could not begin to explain any problem that might require a continuous probability distribution (such as talking about arrivals having an exponential distribution).
  - Students challenged by mathematical notation.

Major Course Constraint

- We (UCF) cannot assume students have had the Network Lab course (or any other networking course).
  - Network Lab course required for IT majors and most do not have calculus.
  - This course requires 2 semesters of calculus plus a course in statistics.
- Must repeat (filtered) some networking material.
  - Some students do take both courses.

Overview of Topics

Overview of Topics Continued

- Continuous Random Variables.....................Chapter 3 Trivedi (3.1-3.4,3.5)
- Expectation...Chapter 4 Trivedi (4.1-4.3)
- Stochastic Processes and Markov Chains....Chapter 6,7 Trivedi (excerpts)
- Queuing Analysis.......................................Chapter 8 Stallings
- TEST
- Congestion Control....................................Chapter 10 Stallings
- Link-level Performance Issues...................Chapter 11 Stallings
- TCP Traffic Control.....................................Chapter 12 Stallings
- Traffic & Congestion Control in ATM...........Chapter 13 Stallings

Introduction and Basics

Topics

- Important networking ideas: 1970’s through 1990’s.
- Networking Basics
- Connection-oriented vs connectionless
- Protocols and Building Blocks
- OSI and Course Reference Models
- Example Networks
- Overview of ATM

Frame Relay

- Also connection-oriented standard for moving bits generally across a public network.
- Takes advantage of fact that leased lines are now fast, digital, reliable --> simple protocols.
- Predominantly a virtual leased line (PVC) today.
- Frames can be up to 1600 bytes and carry a number (dlci) identifying the virtual circuit.
- Contract with carrier for an average service
- burst allowed as in SMDS
- Generally operates at T1 (1.5mbps) or above
- FR determines start and end of frame, detects some transmission errors, discards bad frames.
- No flow control yet defined only a CI bit.

Introduction to ATM

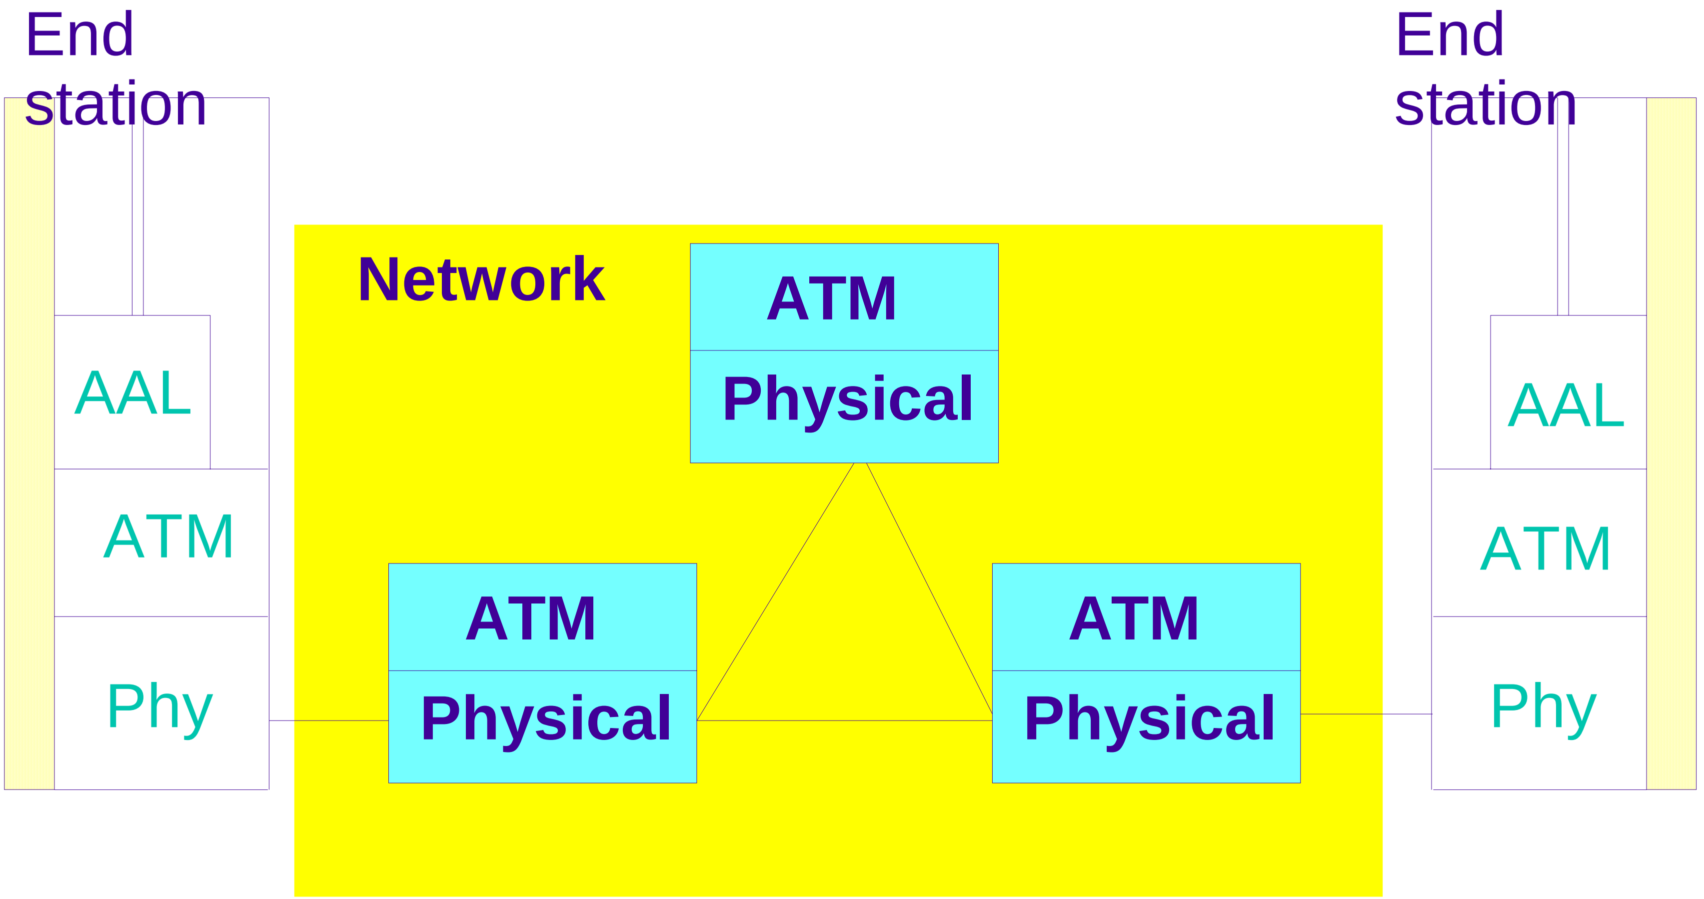

Chapters 2 and 3 - Stallings

Protocols and the TCP/IP Suite

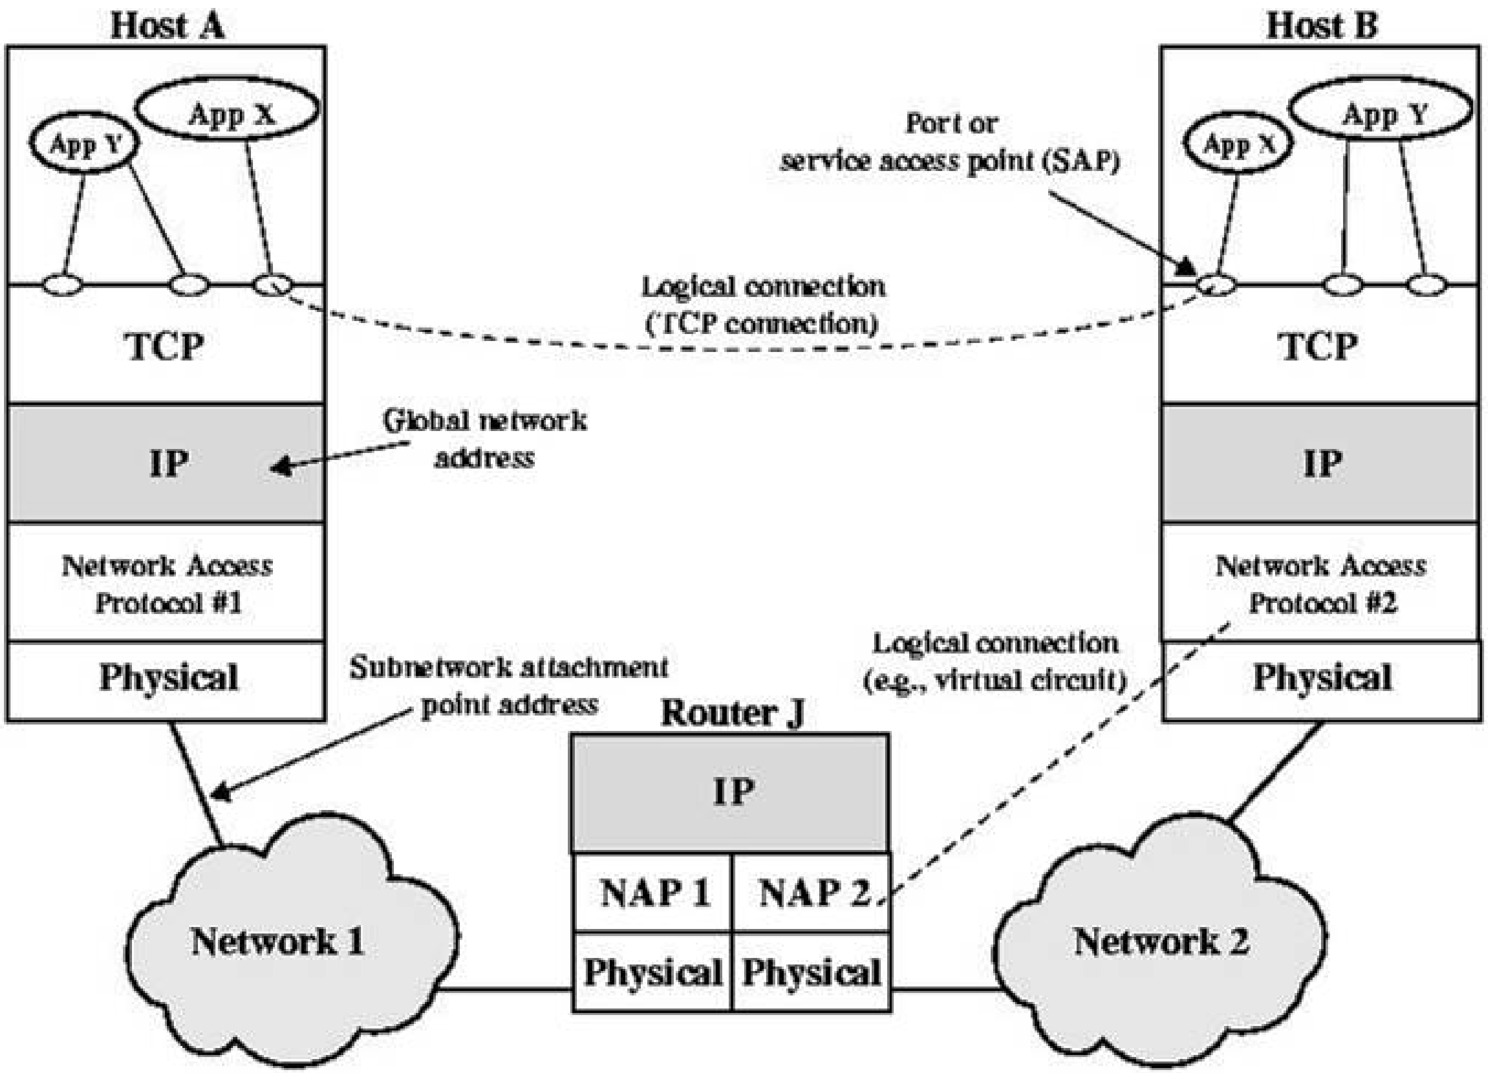

Figure 2.3 TCP/IP Concepts

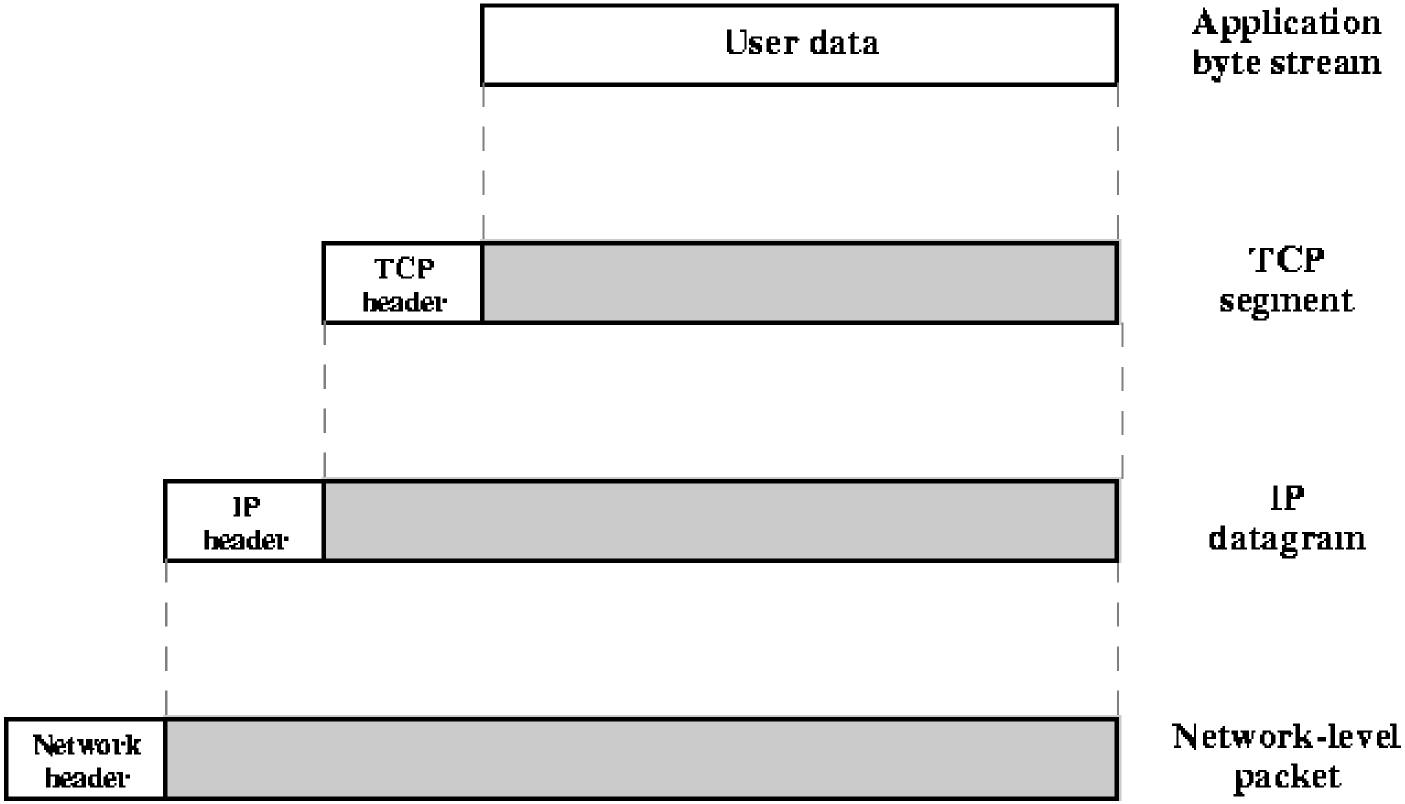

Figure 2.4 Protocol Data Units (PDUs) in the TCP/IP Architecture

TCP and UDP

- TCP:
  - connection-oriented
  - Reliable packet delivery in sequence
- UDP:
  - connectionless (datagram)
  - Unreliable packet delivery
  - Packets may arrive out of sequence or duplicated

AND SO ON...

Routers

- Provide link between networks
- Accommodate network differences:
  - Addressing schemes
  - Maximum packet sizes
  - Hardware and software interfaces
  - Network reliability

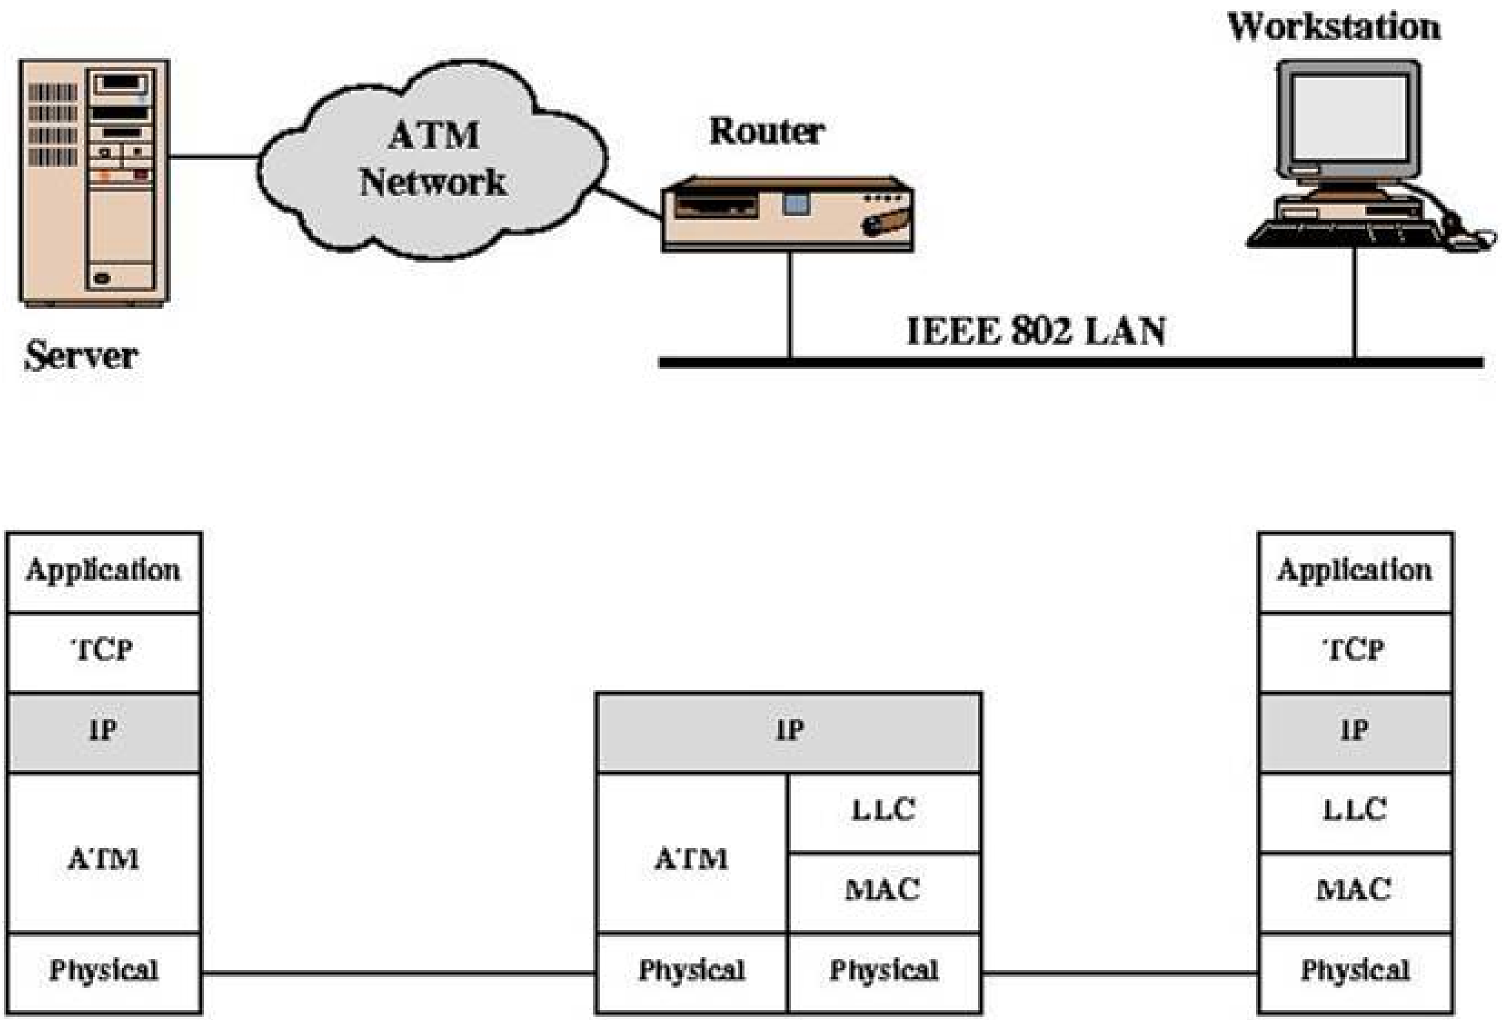

Figure 2.7 Configuration for TCP/IP Example

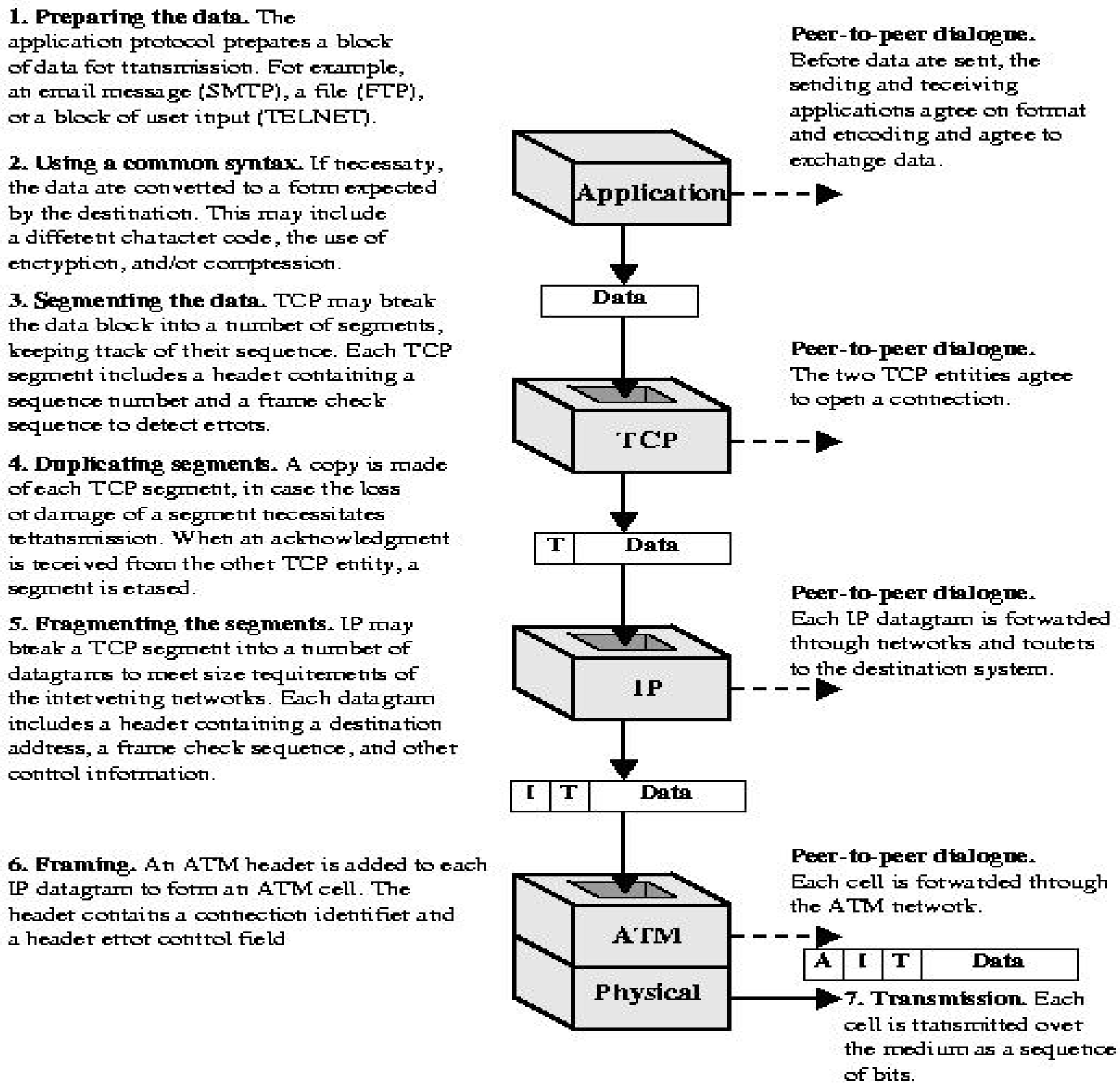

Figure 2.8 Operation of TCP/IP: Action at Sender

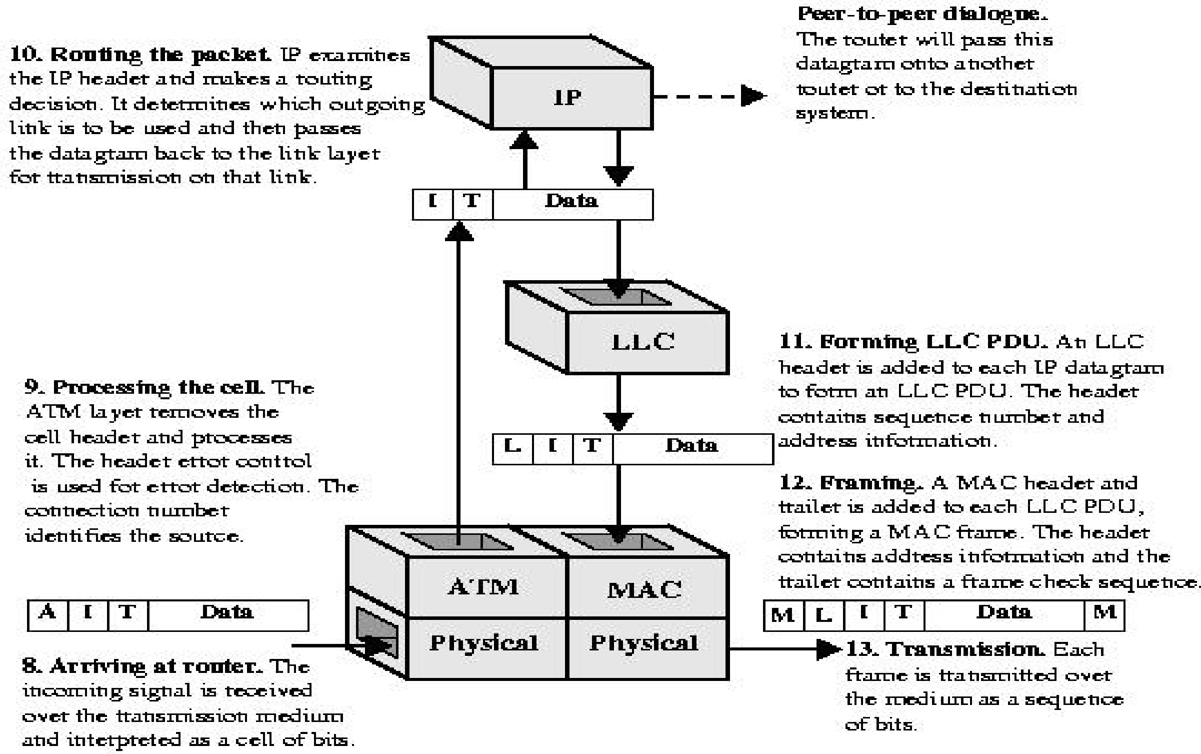

Figure 2.9 Operation of TCP/IP: Action at Router

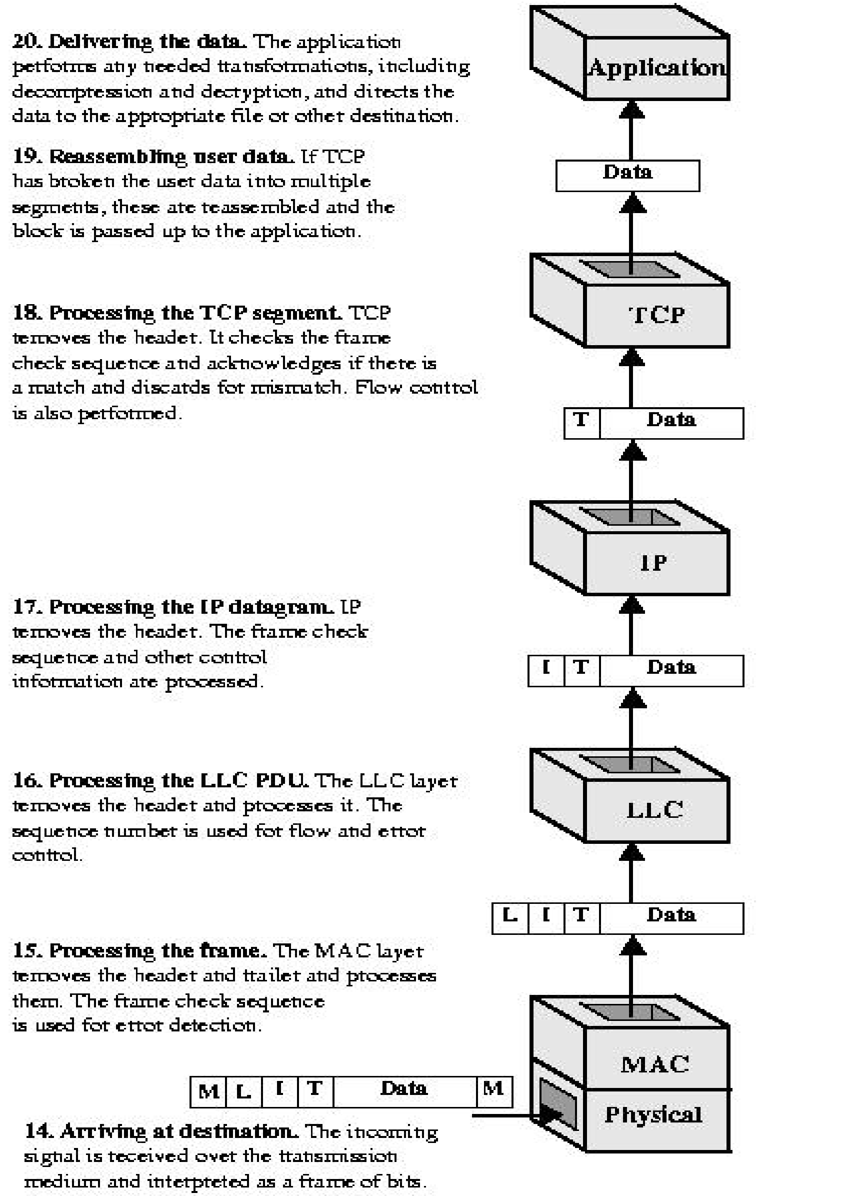

Figure 2.10 Operation of TCP/IP: Action at Receiver

Chapter 4

Frame Relay

Introduction

- Circuit Switching Networks
- Packet-Switching Networks
  - Switching Technique
  - Routing
    - Datagram
    - Virtual Circuit
- Frame Relay Networks
  - Architecture
  - User Data Transfer
  - Call Control

Frame Switching: Making a network look like a link.

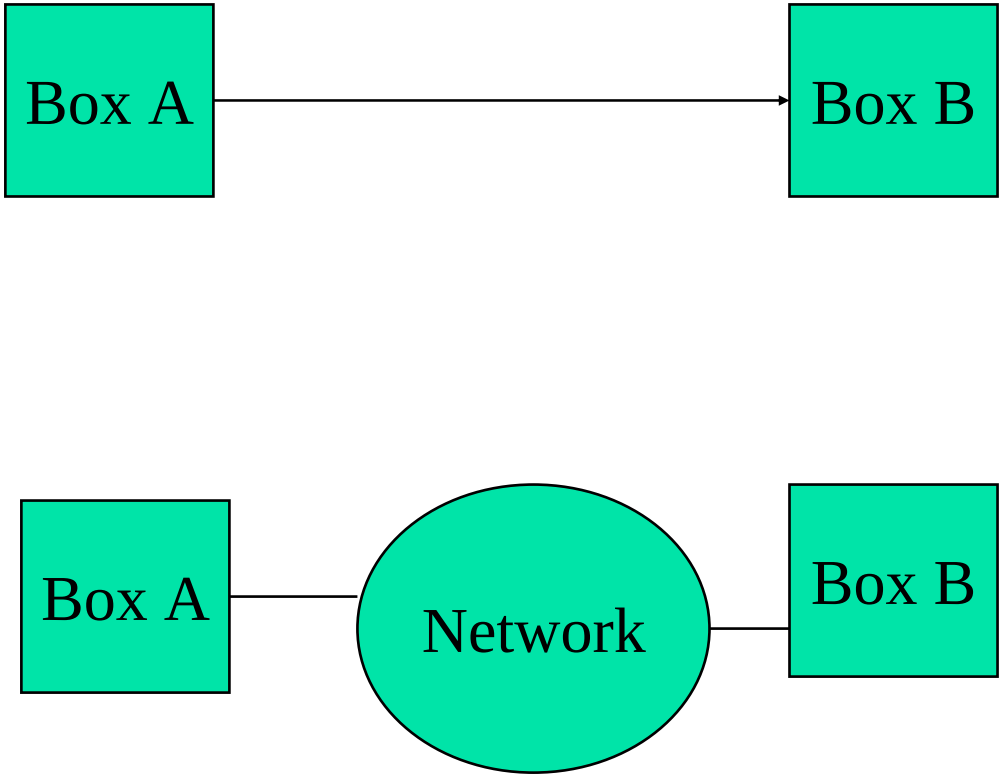

Two Possible Approaches

- Terminate the link control protocol (much like router) and use network headers to route across the network.
  - Problem: Must understand various layer 3's
- Put our own network headers onto whatever is sent by Box A (another envelope), route through the network, then strip off headers and send to Box B.
  - Frame Switching does this.

Frame Switching Problems

- Link control protocols have very short time-outs and network delay is irregular.
- Many link control protocols use polling; sending polls across a network increases network overhead.
- The network's internal routing protocol must use a known address structure and these internal addresses must somehow be associated with external box addresses.
- Network will not be able to recognize priorities or TOS bits in network layer header.

Frame Relay Concept

- The concept of frame relay is simple. A network is interposed between devices communicating on a link, but the devices are not aware that this has happened. Because of the problems already discussed, this is not simple in implementation. However, it is a simple concept.

Introduction to Probability

Includes:

- Sample Space
- Events
- Mutually Exclusive and Collectively Exhaustive
- Probability and Probability Space
- Permutations and Combinations
- Conditional Probability
- Independent Events
- Bayes' Rule
- Bernoulli Trials

Discrete Random Variables

Includes

- Random Variable
- Probability Mass Function
- Probability of Sets and Intervals
- Distribution Function Properties
- Bernoulli Distribution
- Binomial, Geometric, Poisson, Uniform, and Constant Discrete Distributions
- Joint PMF and Independence
- Sum of Independent Random Variables
- Probability Generating Function

Continuous Random Variables

Includes

- Cumulative Distribution Function
- Probability Density Function
- Exponential Distribution
- Hypoexponential, Erlang, Normal, Gamma, Uniform
- Joint Distributions
- Independent Random Variables...

Expectation

Includes

- Expected Value and High Moments
- Exponential, Poisson, Binomial, Geometric, Continuous Uniform
- Moments of Jointly Distributed (simple) Random Variables
- Marginal Distributions
- Linearity of Expectation
- Independence and Expectation

Example

	Y=1	Y=2	Y=3
X=1	1/12	1/6	1/12
X=2	1/6	1/4	1/12
X=3	1/12	1/12	0

What is $E(XY^2)$? What are $E(X)$ and $E(Y)$?

(Also consider Example 4.7 on page 201.)

Stochastic Processes

Includes

- Definition and Classification
- Cumulative Distribution
- Nth-order Joint Distribution Function
- Stationary
- Independent Process
- Markov Chain

Discrete Time Markov Chain

These are particularly important application processes defined by:

(a) discrete index set $N = 0,1,\ldots$

(b) discrete state space $E_1, E_2, \ldots$

(c) the one-step dependence defined by the transition probabilities:

$p_{i,j} = P(X_{n+1} = j \mid X_n = i)$ for n $\in$ N. We assume no dependence on *n*.

Example: Freddie the frog spends all day jumping back and forth among 4 lilly-pads that are numbered 1,2,3,4. When he is on pad 1, he always jumps to pad 2. When he is on pad 2 or 3, the probability is $\frac{1}{2}$ that he will jump one pad in either direction. When he is on pad 4, he always jumps to pad 3. Suppose Freddie starts on pad 2 one morning. What are the probabilities that he will be on Pads 1,2,3,4 after 3 jumps?

Finding Freddie

Let $X_n(i)$ be a random variable giving Freddie's location on pad $i, (i = 1, 2, 3, 4)$ after n jumps. Thus, $\{X_n\}$ is a discrete-time Markov Chain. We can find Freddie's probable location by noting that

$$P(X_n = j) = \sum_{i=1}^{4} P(X_n = j \mid X_{n-1} = i) P(X_{n-1} = i) \text{ for integers } n \text{ and } j \geq 1$$

$$= \sum_{i=1}^{4} p_{i,j} P(X_{n-1} = i), \text{ because the transition probabilities do not change}$$

depending on the hop count, $= \sum_{i=1}^{4} p_{i,j} p_i^{(n-1)}$ if we let $p_i^{(n)}$ be the probability that Freddie is on pad i after n hops.

Finding Freddie (continued)

If we know Freddie's location after $n-1$ hops, we can write Freddie's location probabilities after n hops using matrix multiplication:

$$[p_1^{(n)}\ p_2^{(n)}\ p_3^{(n)}\ p_4^{(n)}] = [p_1^{(n-1)}\ p_2^{(n-1)}\ p_3^{(n-1)}\ p_4^{(n-1)}] \begin{bmatrix} p_{1,1} & p_{1,2} & p_{1,3} & p_{1,4} \\ p_{2,1} & p_{2,2} & p_{2,3} & p_{2,4} \\ p_{3,1} & p_{3,2} & p_{3,3} & p_{3,4} \\ p_{4,1} & p_{4,2} & p_{4,3} & p_{4,4} \end{bmatrix}.$$

The matrix $P = \left[\, p_{i,j} \,\right]$ for i and $j = 1, 2, 3, 4$ is called the transition probability matrix. It has the wonderfully useful property that if $P^{(n)}$ represents the matrix of probabilities that Freddie moves from pad i to pad j in n steps, then $P^n = P^n$. For example, the probabilities that Freddie jumps from pad i to pad j in 2 hops are given by:

Two-hop transition probabilities

$$P^2 = \begin{bmatrix} p_{1,1} & p_{1,2} & p_{1,3} & p_{1,4} \\ p_{2,1} & p_{2,2} & p_{2,3} & p_{2,4} \\ p_{3,1} & p_{3,2} & p_{3,3} & p_{3,4} \\ p_{4,1} & p_{4,2} & p_{4,3} & p_{4,4} \end{bmatrix} \square \begin{bmatrix} p_{1,1} & p_{1,2} & p_{1,3} & p_{1,4} \\ p_{2,1} & p_{2,2} & p_{2,3} & p_{2,4} \\ p_{3,1} & p_{3,2} & p_{3,3} & p_{3,4} \\ p_{4,1} & p_{4,2} & p_{4,3} & p_{4,4} \end{bmatrix}.$$

Back to Freddie

In his case the transition matrix is

$$P = \begin{bmatrix} 0 & 1 & 0 & 0 \\ \frac{1}{2} & 0 & \frac{1}{2} & 0 \\ 0 & \frac{1}{2} & 0 & \frac{1}{2} \\ 0 & 0 & 1 & 0 \end{bmatrix} \Rightarrow P^2 = \begin{bmatrix} \frac{1}{2} & 0 & \frac{1}{2} & 0 \\ 0 & \frac{3}{4} & 0 & \frac{1}{4} \\ \frac{1}{4} & 0 & \frac{3}{4} & 0 \\ 0 & \frac{1}{2} & 0 & \frac{1}{2} \end{bmatrix}$$

After 3 hops

$$P^3 = P \cdot P^2 = \begin{bmatrix} 0 & \frac{3}{4} & 0 & \frac{1}{4} \\ \frac{3}{8} & 0 & \frac{5}{8} & 0 \\ 0 & \frac{5}{8} & 0 & \frac{3}{8} \\ \frac{1}{4} & 0 & \frac{3}{4} & 0 \end{bmatrix}$$

If Freddie's starting probabilities are $[p_1^{(0)} \; p_2^{(0)} \; p_3^{(0)} \; p_4^{(0)}]$, then we can find his "lilypad probabilities" after n steps as:

$$[p_1^{(n)} \; p_2^{(n)} \; p_3^{(n)} \; p_4^{(n)}] = [p_1^{(0)} \; p_2^{(0)} \; p_3^{(0)} \; p_4^{(0)}] P^n.$$

Where is he after 3 steps in this case?

Answer: $$[p_1^{(3)} \; p_2^{(3)} \; p_3^{(3)} \; p_4^{(3)}] = [0 \; 1 \; 0 \; 0] \begin{bmatrix} 0 & \frac{3}{4} & 0 & \frac{1}{4} \\ \frac{3}{8} & 0 & \frac{5}{8} & 0 \\ 0 & \frac{5}{8} & 0 & \frac{3}{8} \\ \frac{1}{4} & 0 & \frac{3}{4} & 0 \end{bmatrix} = \begin{bmatrix} \frac{3}{8} & 0 & \frac{5}{8} & 0 \end{bmatrix}.$$

Time Homogeneous

Given any two points in time $0 < t_1 < t_2$, if the Markov Process has the property that $P\left[X_{t_2} = x_2 \mid X_{t_1} = x_1\right] = P\left[X_{t_2 - t_1} = x_2 \mid X_0 = x_1\right]$, then the process is said to be time homogeneous.

Example. Repeated failures of lightbulbs when time to replace is negligible.

Written more simply this means that $P\left[Y \le r + t \mid Y \ge t\right] = P\left[Y \le r\right]$ if Y is a time homogeneous Markov Process. This is what we have called the memoryless property. Using it one can show that that waiting time in any state of a time-homogeneous Markov Process has an exponential distribution.

Where is Freddie Long-term?

In general, we are more interested in where Freddie will be "in the long run" than after 2, 4, 8, ... hops in particular. Thus, we recall the set of probabilities

$p_j^{(n)} = P\left[X_n = j\right] =$ the probability that Freddie is on pad j after n hops.

We then denote the "limiting state probabilities" by

$v_j = \lim_{n\to\infty} p_j^{(n)}, \; j = 0,1,\ldots.$ Many Markov Chains (not all) have these limiting probabilities, which are important from a system analysis standpoint.

When these probabilities exist, they are said to form the "steady-state probability vector $\underline{v} = \left[v_1 \; v_2 \; \ldots\right]$.

To predict when they exist and to find them, we need definitions that are used to classify Markov Chains.

Transient vs Recurrent

- A state i is transient (or non-recurrent) iff there is a positive probability that the process will NOT return to this state.
- A state i is recurrent iff, starting from state i, the probability that the process eventually returns to state i is one.

If we let f_{ij} be the probability that we will "ever" visit state j starting from state i, and $f_{ij}^{(n)}$ be the probability that the first visit to state j, starting from i, occurs in exactly n steps, then $f_{ij} = \sum_{n=1}^{\infty} f_{ij}^{(n)}$.

Thus, a state is recurrent if $f_{ii} = 1$ and transient if $f_{ii} < 1$.

NOTE: A state is "absorbing" if the transition probability $p_{ii} = 1$.

Recurrent: null vs non-null

The mean recurrence time of state i is $\mu_i = \sum_{n=1}^{\infty} n f_{ii}^{(n)}$. A recurrent state i is said to be recurrent non-null (or positive recurrent) if $\mu_i < \infty$ and is said to be recurrent null if its mean recurrence time is infinite.

Periodic vs Aperiodic

For any recurrent state i the probability $p_{ii}^{(n)}$ must be positive for at least one value of $n \geq 1$. The period of state i is defined as

$d_i =$ the greatest common divisior of the set of all positive integers n such that $p_{ii}^{(n)} > 0$.

A recurrent state i is periodic if $d_i < 1$ and aperiodic if $d_i = 1$.

Irreducible Markov Chain

A Markov Chain is irreducible if every state can be reached from every other state in a finite number of steps. Mathematically, this is written as follows:

For every pair of states i and j, there is an integer $n \geq 1$, such that $p_{ij}^{(n)} > 0$.

Key Theorems (p351, Trivedi)

Theorem 7.2. For an aperiodic Markov chain, the lmits $v_j = \lim_{n\to\infty} p_j^{(n)}$ exist.

Theorem 7.3. For any irreducible, aperiodic Markov chain, the limiting state probabilities $v_j = \lim_{n\to\infty} p_j^{(n)} = \lim_{n\to\infty} p_{ij}^{(n)}$ exist and are independent of the initial probability vector.

Theorem 7.4. For an irreducible, aperiodic Markov chain, with all states recurrent, non-null, the limiting probability vector $\underline{v}$ is the unique stationary probability vector that satisfies the matrix equation $\underline{v} = \underline{v}P$, where P is the original transition probability matrix.

So Where in the #@%?$$& is Freddie?

Think: "Does Freddie love RAIN?" which stands for (Recurrent, Aperiodic, Irreducible, Non-null)?

Clearly, Freddie's Markov Chain is irreducible because every lilly pad can be reached from every other lilly-pad. Furthermore, it is recurrent because there is no set of states where Freddie can get "stuck" so that he could not return to his starting place.

BUT this chain is periodic with period 2 (which can be seen be looking at P^2 and P^3). So...Freddie no longer serves us as a good example.

Freddie Cooperates

Sensing our difficulty Freddie modifies his behavior as follows:

When on pad 1 or pad 4, the probability is $\frac{1}{2}$ that he remains there for the next transition by "hopping in place." The probability he moves from pad 1 to pad 2 is $\frac{1}{2}$ as is the probability that he will hop from pad 4 to pad 3. Otherwise Freddie's behavior does not change.

Homework: Convince me that "Freddie now likes RAIN" and find his steady-state probability vector. Also, work the odd numbered problems in the Stallings text pages 179-181. This homework due Tuesday, March 26.

Section 5: Queuing Analysis

Problems...

1. A collection of five routers is to be connected in a point-to-point subnet. Between each pair of routers, the designers may put a high-speed line, a medium-speed line, a low-speed line, or no line. If it takes 100ms of computer time to generate and inspect each topology, how long will it take to inspect all of them to find the one that best matches the expected load?

-

2. A group of routers are interconnected in a centralized binary tree, with a router at each tree node. Router i communicates with router j by sending a message to the root of the tree. The root then sends the message back down to j. Derive an approximate expression for the mean number of hops per message for large n, assuming that all router pairs are equally likely.

Problems...

1. In most networks, the data link layer handles transmission errors by requesting damaged frames to be retransmitted. If the probability of a frame's being damaged is p, what is the mean number of transmissions required to send a frame if acknowledgements are never lost? What is the variance of this number?

-

2. Imagine that you have trained your St. Bernard, Bernie, to carry a box of three 8mm tapes instead of a flask of brandy. (When your disk fills up, you consider it an emergency.) These tapes each contain 7 gigabytes. The dog can travel to your side, wherever you may be, at 18 km/hr. For what range of distances does Bernie have a higher data rate than a 155-Mbps ATM line?

Problems...

1. An upper layer message is split into 10 frames, each of which has an 80 percent chance of arriving undamaged. If no error control is done by the data link protocol, how many times must the message be sent on the average to get the entire thing through? What is the probability that the message must be sent exactly 4 times?

2. If the bit string 0111101111101111110 is bit stuffed (by the sender) what is the result?

3. A block of bits with n rows and k columns uses horizontal and vertical parity bits for error detection. Suppose that exactly 4 bits are inverted due to transmission errors. Derive an expression for the probability that the error will be undetected.

Problems...

Suppose you are writing a simulation program and you wish to generate a sequence of packet arrivals according to a Poisson distribution. Your compiler has available a random number generator that can generate numbers uniformly distributed between 0 and 1. How can you generate the simulated arrivals times of the packets?

-

2. Consider the following design problem concerning implementation of virtual circuit service. If virtual circuits are used internal to the subnet, each data packet must have a 3-byte header, and each router must tie up 8 bytes of storage for circuit identification. If datagrams are used internally, 15-byte headers are needed, but no router table space is required. Transmission capacity costs 1 cent per bytes, per hop. Router memory can be purchased for 1 cent per byte and is depreciated over two years (business hours only). The statistically average session runs for 1000 sec, in which time 200 packets are transmitted. The mean packet requires four hops. Which implementation is cheaper, and by how much?

Problems...

1. Consider the effect of using slow-start on a line with a 10 msec round-trip time and no congestion. The receive window is 24 kBytes and the maximum segment size is 2 kBytes. How long does it take before the first full window can be sent?

-

2. A TCP machine is sending windows of 65,535 bytes over a 1-Gbps channel that has a 10-msec one-way delay. What is the maximum throughput achievable? What is the max line efficiency?

-

3. For a 1-Gbps network operating over 4000 km, the delay is the limiting factor, not the bandwidth. Consider a MAN with the average source and destination 20 km apart. At what data rate does the round-trip propagation delay equal the transmission delay for a 1-kByte packet?

Student Experience (Positive)

- 3 or 4 of 30 students were excited to learn how these problems could be treated.
- Approximately 10 students really seemed to have “mastered” most of the material.
- A number of students said it was there most difficult class at the university but they learned a great deal.

Student Experience (Negative)

- Many students felt the math was too demanding.
- Students with no previous network experience were overwhelmed.
- Students especially did not like the Stallings book because it did not show them how to work key problems.