

Welcome to CS 756, Spring 2004

- ❑ Instructor: Yih (Ian) Huang
- ❑ Office: ST II, Rm 443
- ❑ E-mail: huangyih@cs.gmu.edu
- ❑ Office hours: Wednesday 1:00pm --- 3:00pm
- ❑ Course home page:
<http://www.cs.gmu.edu/~huangyih/756>
- ❑ We will use emails for communication; you must have a GMU account and check the account for course messages periodically, if not daily.

Internet Course Delivery

- ❑ Classes are delivered live over the Internet and also recorded for playback.
 - You can attend classes remotely before midterm.
 - [Classroom presence is mandatory after midterm.](#)
- ❑ Minimum requirements
 - Windows system 200MHz with a sound card.
 - Internet explorer
 - A clear network path (no firewall or congestion)
 - ☞ Direct dial-up to GMU recommended

Web Cites

- ❑ URL: netlab.gmu.edu/disted
 - Class delivery and playback are protected by passwords.
 - Following a link on the page to obtain your password.
 - Your osf1 account will be your login ID.
 - Down load client software from netlab.gmu.edu/NEW

CS 756

3

Course Outline

- ❑ Mathematic tools
 - Queueing theory
 - Random variable generations
- ❑ Network performance issues
 - TCP performance modeling
 - Realtime, multimedia rate control
 - Internet congestion control
 - Internet topology/traffic characterizations
 - Quality of Service (QoS) architectures
 - Discrete event simulation

CS 756

4

What You Should Know Already

- ❑ OSI 7-layer reference model
- ❑ Internet architecture
- ❑ Important Internet protocols, particularly, TCP and UDP
- ❑ Calculus-based probability theory
- ❑ Basic graph theory

**These are serious requirements of a very demanding class.
Do make sure you are ready.**

CS 756

5

Course Format

- ❑ Lectures
 - The instructor give lectures before the midterm.
 - These lectures build the math foundations and discuss interesting network performance issues for your to select a project topic.
 - Midterm (tentatively March 13th) covers all lectures.

CS 756

6

Projects

- ❑ Two persons team up to investigate a network performance issue.
- ❑ 2-page proposal due the day of midterm exam.
- ❑ Each team delivers an one-hour presentation after the midterm.
 - To encourage student interactions, classroom attendance is **mandatory** after midterm.
- ❑ Each team submits a project report in the final week.
- ❑ There will be no final exam.

CS 756

7

- ❑ You shape your project; possibilities include
 - Extensive survey of a subject
 - Performance study by simulation
 - Performance study by analysis/measures
- ❑ A reading list available on the course page.
 - Use it to find the topics that interest you.
 - You are by no means restricted to the list.
 - You are encouraged to find combinations and/or come up with novel ideas.
- ❑ Keys to success: **start early, have fun**

CS 756

8

Project Proposal

- ❑ A proposal contains three parts.
 1. Title Page: project title and team members.
 2. Proposal body: 1 to 2 pages, 12 pt, 1.5 spacing.
 3. Copies of 2 to 5 related papers
- ❑ Include in the proposal body:
 - Subject of interests
 - Nature of the study (performance evaluation, survey, theory study, etc.)
 - Performance metrics
 - Evaluation methods and tools
 - Expected results

CS 756

9

Presentation Format

- ❑ 30 to 40 minutes talk, 10 to 15 minutes discussion
- ❑ Must haves
 - Present your subject(s) clearly to class
 - Define your goals
- ❑ Optional, depending on the project.
 - Performance metrics of interests
 - Performance evaluation method
 - Results or a study plan
 - Math models, proofs, etc.
 - Comments to your proposal

CS 756

10

Project Report

- ❑ Required: title page and report body
- ❑ Optional: more reading materials (only those not included in the proposal)
- ❑ Report body
 - Times-Roman, 12 points, 1.5 spacing
 - 8 to 15 text pages, excluding plots and figures
 - All plots and figures placed after texts

Course Mission Statements

- ❑ Awareness of basic math tools in studying network performance
- ❑ Understanding fundamental factors of Internet performance
- ❑ Independent survey of literatures and standards
- ❑ Applying above knowledge/skills to a topic of your choice
- ❑ Development of verbal and written communication skills

Grading

- Homework 10%
- Midterm 40%
- Paper presentation 25%
 - Grading is based on how much the class learns from you.
- Project report 25%
 - Grading based on the quality of your work

Homework #1

- Download the reading list from the course web page.
- Read [SRC184] (SRC-ell-84) and [BIC101] (B-ell-C-ell-zero-one).
- Write a 1-to-2 page summary (Time New Roman, 12 points, 1.5 spacing) of the two papers.
 - Feel free to disagree with the authors.
 - *Your thinking* is worth more points than merely repeating the thoughts of the authors.
- Due 7:20pm Feb. 4th.

On-line Literatures

- ❑ You can download some (but not all) listed papers from www.cs.gmu.edu/~huangyih/docs/.
 - For example, to download [BICI01], enter the URL to Netscape or Explorer
www.cs.gmu.edu/~huangyih/docs/BICI01.pdf
 - Most of the files are pdf, but you should try ps too if a pdf version is not found.
- ❑ IEEE literature database: iel.ihs.com

A Review of Probability

These slides are created by Dr. Yih Huang of George Mason University. Students registered in Dr. Huang's courses at GMU can make a single machine-readable copy and print a single copy of each slide for their own reference, so long as each slide contains the copyright statement, and GMU facilities are not used to produce paper copies. Permission for any other use, either in machine-readable or printed form, must be obtained from the author in writing.

Disclaimer

This review is concise and informal.
Please refer to your probability
textbooks for rigid discussions.

Random Variables

- ❑ Random variables contain the outcomes of indeterministic events.
- ❑ Example:
 - X : the number A s of CS756, Spring 2003
 - Y : the time Dr. Huang walks into the classroom (using class start time as 0).
- ❑ Random variable X is said to be **discrete**.
- ❑ Random variable Y is said to be **continuous**.

Probability Functions

Purpose: Describing the chances, from 0 to 1, of potential outcomes.

- ❑ A **probability mass function** gives the probability of a specific outcome for a discrete random variable.

$$\begin{aligned}h(n) &= P\{X = n\} \\ &= P\{\text{we will have exactly } n \text{ As in CS656}\}\end{aligned}$$

- ❑ A **probability density function** (pdf) gives the probability of a range of outcome for a continuous random variable.

- Let $f(t)$ be the pdf of Y .
- Then the probability for Dr. Huang to be late is given by

$$\int_0^{\infty} f(t) dt$$

CS 756

19

Distributions

- ❑ Random variables can be categorized into several important families, each of which possesses a distinguished pattern about how chances are “distributed” among potential outcomes.
- ❑ **Uniform distribution:** every outcome has the same chance. $p(x) = 1/N$
- ❑ Consider a uniform random variable X .
 - If X is discrete and has N possible outcomes, then its pmf $f(x) = 1/(b-a)$
 - If X is continuous and has the outcome range of a to b , then its pdf


CS 756

20

Poisson Distribution

A random variable X , taking on one of the values $0, 1, 2, \dots$, is Poisson distributed with mean $\lambda > 0$, if

$$p(i) = P\{X = i\} = e^{-\lambda} \frac{\lambda^i}{i!}$$




CS 756

21

Normal Distribution

$$f(x) = \frac{1}{\sqrt{2\pi}\sigma} e^{-(x-\mu)^2/2\sigma^2}, \quad -\infty < x < \infty,$$

where μ is the average and σ^2 the variance.



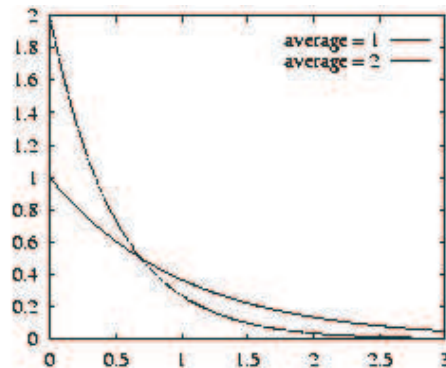
CS 756

22

Exponential Distribution

$$f(x) = \begin{cases} \lambda e^{-\lambda x}, & \text{if } x \geq 0 \\ 0, & \text{if } x < 0 \end{cases}$$

where λ is the average.



CS 756

23

The Memoryless (Markov) Property of Exponential Distribution

$$P\{X > r+t \mid X > t\} = P\{X > r\}, \quad r, t > 0$$

Example:

- Assume that phone call times are exponentially distributed.
- I arrive at two phone booths; both are busy.
- A person nearby tells me that one person just started her call, while the other has been on the phone for 30 minutes.
- Which booth should I wait for ?

CS 756

24

Poisson and Exponential Distributions

- ❑ When the interarrival times of events are exponentially distributed, the number of events in a fixed time interval is Poisson distributed, with the same average rate.
- ❑ The reverse is also true.

CS 756

25

Little's Theorem

- ❑ Consider a system where customers come and leave. We have
$$N = \lambda T$$
 - N is the average number of customers in the system
 - T is the average time per customer in system
 - λ is the customer arrival rate
- ❑ Note that the theorem does not assume particular distributions about
 - Customer interarrival times
 - Customer times in system

CS 756

26

An Intuitive “Proof”

- ❑ A customer arrives at the system.
- ❑ After a period of time T , he/she leaves the system.
- ❑ How many new customers have arrived during the period ?

CS 756

27

Example

- ❑ Consider a 24-hour barbershop where there is only one barber.
- ❑ At 5:30pm
 - $N = 6$
 - $\lambda = 0.33$ customers per second
 - $T = ?$
- ❑ At 1:00am
 - $N = 2$ customers
 - $T = 20$ minutes
 - $\lambda = ?$

CS 756

28

The Power of Analysis: A Classic Example

- ❑ ALOHA: a MAC protocol developed at University of Hawaii for ground-based radio (1971).
 - seven campuses on four of the islands
 - main computer center on Oahu
 - other campuses need access to the main computer
 - telephone connections expensive and unreliable
- ❑ Each station equipped with FM radio transceiver.
- ❑ Transmitters send packets to a hub (located on Oahu) on a shared frequency and the hub retransmits received packets on a second frequency.

CS 756

29

- ❑ The retransmissions are for collision detection:
 - A transmitter listens to the retransmission of its packet from the hub.
 - If two or more transmissions to the hub collide, then the hub discards the corrupted transmission and none of the colliding transmitters would receive the retransmission of their respective packets from the hub.
 - When this happens, each colliding transmitter waits a random time and then tries to send its packet again.
- ❑ Collisions can occur whenever any part of a transmitted packet overlaps.

CS 756

30

Performance Analysis

- ❑ P_0 : probability that a frame does not suffer a collision
- ❑ S : number of new frames generated per frame time
- ❑ G : number of new and retransmitted frames generated per frame time
- ❑ $S = GP_0$
- ❑ Assume both S and G are Poisson distributed
- ❑ Probability that k frames are generated within a given frame time is

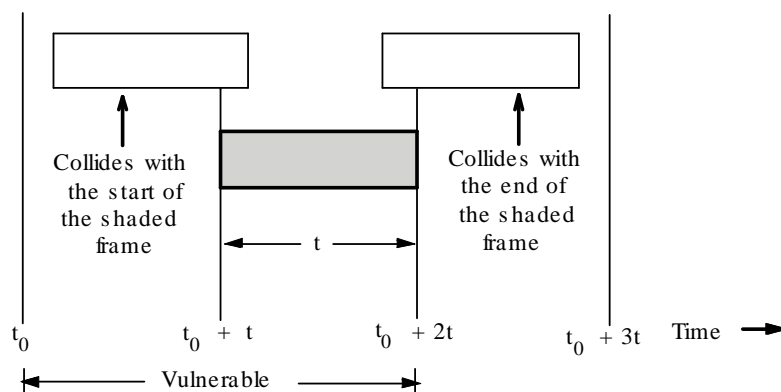
$$P[k] = \frac{G^k e^{-G}}{k!}$$

CS 756

31

Vulnerable Period

- ❑ For a given frame, the time window when no other frame may be transmitted to avoid collisions.
- ❑ Length ?



CS 756

32

❑ Let us compute P_0 :

$$\begin{aligned} P_0 &= P\{\text{no other frames in a 2 - frame interval}\} \\ &= \frac{(\lambda t)^k e^{-\lambda t}}{k!}, \text{ where } t=2, \lambda=G, k=0 \\ &= e^{-2G} \end{aligned}$$

❑ Thus, $S = Ge^{-2G}$

❑ Maximum throughput can be obtained by

$$S' = e^{-2G} - 2Ge^{-2G} = e^{-2G}(1 - 2G) = 0.$$

We have $G = 0.5$.

❑ That is, $S = GP_0 = 0.5e^{-2G} = 0.5e^{-1} \approx 0.184$

CS 756

33

Slotted ALOHA

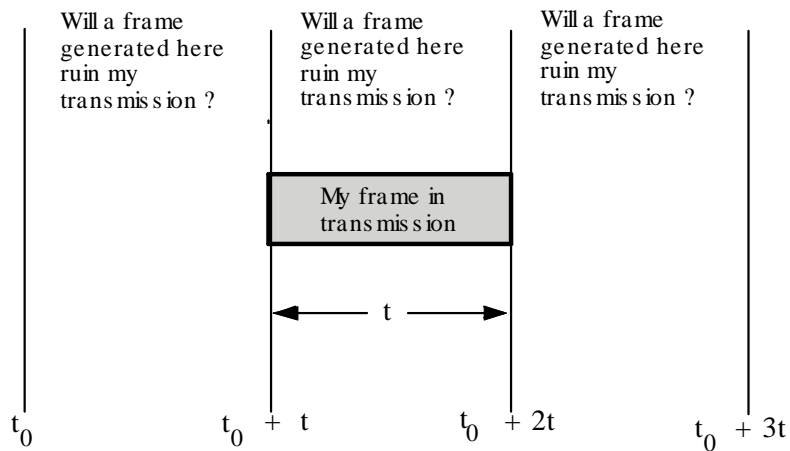
- ❑ Divide time into discrete intervals, called slots; each slot corresponding to one frame time
- ❑ Stations may send only at the beginning of a slot.
- ❑ if a frame is ready for transmission in the middle of a slot, it must be positioned to the beginning of the next slot.

CS 756

34

Vulnerable Period

□ Length ?



CS 756

35

- Now, the length of my vulnerable period is only:
- As such, $S = Ge^{-G}$
- That is to say, $S' = e^{-G} - Ge^{-G} = e^{-G}(1 - G) = 0$.
- We have $G=1$ and hence maximum $S = e^{-1} \approx 0.368$
- With appropriate discipline, the maximum throughput is doubled !

CS 756

36

A Second Example

- Consider the test for a certain disease with 1% error rate.
- Only 1 out of 10,000 has the disease.
- I'm tested "positive."
- What is the probability that I am ill ?

CS 756

37

Consider Intrusion Detection

- A test for IP packets to be "malicious" (used by hackers for whatever evil purposes) has the error rate of $e=1\%$.
- Only 1 out of $N=10,000$ packets is malicious.
- What is the false-alarm to real-alarm ratio ?

- Is this a fundamental limitation to intrusion detection systems ?

CS 756

38