

STOCHASTIC PERFORMANCE MODELING AND SCHEDULING IN COMPUTER SYSTEMS

Spring 2026

Instructor:	Ben Berg	Time:	Monday, Wednesday; 3:35 - 4:50
Email:	ben@cs.unc.edu	Room:	SN 115

Course Page: https://bsb20.github.io/690_s26/

Office Hours: TBA, or By Appointment in FB 336

Description This course will teach the fundamentals of stochastic performance modeling, and then apply those fundamentals to better understand the state-of-the-art in systems design and scheduling theory.

We will begin by developing an understanding of the stochastic processes used to model modern computer systems, and some methods for analyzing these models. We will see how these tools can be used to answer a variety of practical questions such as: How can we minimize latency or maximize throughput in a system? How many servers do I need to buy to meet my latency/throughput goals? Should I buy one faster server, or two slower servers? How should I balance load across the many servers in my system?

After looking at the basic tools required to answer some of these questions, students will examine some recent work in the area of performance modeling and system design, applying the fundamental concepts learned in lecture to understand current research.

Student Learning Outcomes: The goals of this class are to:

1. Develop an understanding of the basic tools of stochastic performance modeling
2. Be familiar with how problems in modern computer systems can be modeled and analyzed via stochastic processes
3. Contrast the style of analysis used by the performance modeling community with that of computer science research communities (e.g. the systems community, worst-case theory community)

Main References: The main text for the course is *Performance Modeling and Design of Computer Systems* (Harchol-Balter, Cambridge University Press). It is available for free online through the UNC Library.

Target Audience: This class is intended for graduate students and advanced undergraduates with an interest in a quantitative approach to building and optimizing computer systems. This might be a student with a strong theory background who is interested in how theoretical ideas can be deployed in modern computer systems. This might also be a systems student who wants to learn about a more formal approach to optimizing system performance. Students should have familiarity with applied probability at an undergraduate level, but a deeper background in stochastic processes is not required.

Prerequisites: STOR 435/535 or equivalent. Familiarity with the material in Chapter 3 of the course text book.

Lecture	Date	DoW	Topic	Work Due
1	1/7	W	Chpt 1: Motivating Examples on Queueing Theory.	
2	1/12	M	Chpt 2: Queueing Theory Notation/Vocabulary.	
3	1/14	W	Chpt 4: Simulating Random Variables	HW 1 Due 1/16
	1/19	M	MLK Day, NO CLASS	
4	1/21	W	Chpt 5: Convergence of Random Variables and Time Average versus Ensemble Average.	
5	1/26	M	Chpt 6: Operational Laws (Little's Law)	
6	1/28	W	Chpt 7: Modification Analysis	
7	2/2	M	Chpt 7: Modification Analysis	
8	2/4	W	Chpt 8: Discrete-time Markov Chains	HW 2 Due 2/6
	2/9	M	Well-being Day, NO CLASS	
9	2/11	W	Chpt 8: Discrete-time Markov Chains	
10	2/16	M	Chpt 9: Ergodicity - Finite-state DTMCs	
11	2/18	W	Chpt 10: Introduction to Transform Analysis	
12	2/23	M	Chpt 11: Exponential Distribution	
13	2/25	W	Chpt 11: Poisson Process	HW 3 Due 2/27
14	3/2	M	Chpt 12,13: M/M/1	
15	3/4	W	Chpt 14: M/M/k	
16	3/9	M	Transform Analysis II	
17	3/11	W	Chpt 15: Capacity Provisioning	HW 4 Due 3/13
	3/16	M	Spring Break, NO CLASS	
	3/18	W	Spring Break, NO CLASS	

18	3/23	M	Chpt 20: Pareto Distribution	
19	3/25	W	Chpt 21: Phase-type distributions + start Chpt23	
20	3/30	M	Chpt 23: M/G/1	
21	4/1	W	Chpt. 26 M/G/1 Busy Period and Transform	
22	4/6	M	Scheduling I	
23	4/8	W	Scheduling II	HW 5 Due 4/10
24	4/13	M	Scheduling III	
25	4/15	W	Scheduling IV	
26	4/20	M	Paper Presentations	
27	4/22	W	Paper Presentations	HW 6 Due 4/24
28	4/27	M	Paper Presentations	

Important Dates:

- HW1 Due - 1/16
- HW2 Due - 2/6
- HW3 Due - 2/27
- HW4 Due - 3/13
- HW5 Due - 4/10
- HW6 Due - 4/24
- Paper Presentations - 4/22 to 4/27
- Final Exam - TBD

Course Requirements: This course will consist of lectures, homework problem sets, paper discussions, and a final exam. The plan is to spend roughly the three quarters of the course learning a solid foundation in performance modeling, and then finish up with some paper reading.

Most of the course will therefore consist of lectures which will be accompanied by problem sets. Problem sets will account for 60% of your course grade.

Towards the end of the course, each student will lead a paper presentation. Your job is to pick between 1 and 3 papers on a single topic that is related to the course material. You will read the paper(s) and give an in-class presentation on the topic you chose. Your goal is to explain the contents of what you read in the context of the material from lecture. For example, if you choose a more systems-oriented paper, explain how you might formally model the problem addressed in the paper and explain why this is a hard theory problem. Alternately, if you choose a scheduling paper from the worst-case or real-time scheduling communities, compare the model being used to an analogous stochastic model, and explain how the differences between the models might change the results of the paper. Your paper presentation will count for 15% of your course grade.

There will be a final exam worth 15% of your course grade. Also, throughout the course, but particularly during the paper discussions, participation is critical and therefore will count for 10% of your course grade.

Grading: To Recap:

- Homework: 60%
- Paper Presentation: 15%
- Final Exam: 15%
- Participation: 10%

Collaboration Policy and Other Rules: Collaboration on homework with current classmates is allowed. This includes any discussion that occurs in office hours. However, everyone must turn in their own version of each assignment. Please document your collaborators at the top of each assignment.

In addition to all UNC policies (below), everyone must adhere to the [reasonable person principle](#).

University of North Carolina at Chapel Hill
Statements for Undergraduate Classes
Fall 2024

Attendance Policy

University Policy: As stated in the University's [Class Attendance Policy](#), no right or privilege exists that permits a student to be absent from any class meetings, except for these University Approved Absences:

1. Authorized University activities: [University Approved Absence Office \(UAAO\)](#) website provides information and [FAQs for students](#) and [FAQs for faculty](#) related to University Approved Absences
2. Disability/religious observance/pregnancy, as required by law and approved by the [Equal Opportunity and Compliance Office](#) (EOC)
3. Significant health condition and/or personal/family emergency as approved by the [Office of the Dean of Students, Gender Violence Service Coordinators](#), and/or the [Equal Opportunity and Compliance Office](#) (EOC).

Code of Conduct

All students are expected to adhere to University policy and follow the guidelines of the UNC Code of Conduct. Additional information can be found at [studentconduct.unc.edu](#).

Artificial Intelligence (AI) Use Policy – CAS units only

Instructors should specify the details of AI Use Policies for the particular course, either by indicating that:

Use of generative AI tools of any kind is not permitted in this course. Any use of these tools will be considered an instance of academic dishonesty and will be referred to the Honor System.

- or -

The following uses of generative AI tools are permitted in this course: *Categories of possible permitted use include, but are not limited to: topic selection, brainstorming and idea generation, research, source validation, outlining and planning, drafting, media creation, peer review, revising, and polishing.*

Syllabus Changes

Information for Students (to be included on the syllabus):

The instructor reserves the right to make changes to the syllabus including project due dates and test dates. These changes will be announced as early as possible.

Equal Opportunity and Compliance - Accommodations

Equal Opportunity and Compliance Accommodations Team ([Accommodations - UNC Equal Opportunity and Compliance](#)) receives requests for accommodations for disability, pregnancy and related conditions, and sincerely held religious beliefs and practices through the University's Policy on Accommodations. EOC Accommodations team determines eligibility and reasonable accommodations consistent with state and federal laws.

Counseling and Psychological Services (CAPS)

UNC-Chapel Hill is strongly committed to addressing the mental health needs of a diverse student body. The [Heels Care Network](#) website is a place to access the many mental health resources at Carolina. CAPS is the primary mental health provider for students, offering timely access to consultation and connection to clinically appropriate services. Go to the

[CAPS website](#) or visit their facilities on the third floor of the Campus Health building for an initial evaluation to learn more. Students can also call CAPS 24/7 at 919-966-3658 for immediate assistance.

Title IX and Related Resources

Any student who is impacted by discrimination, harassment, interpersonal (relationship) violence, sexual violence, sexual exploitation, or stalking is encouraged to seek resources on campus or in the community. Reports can be made [online to the EOC](#) or by contacting the [University's Title IX Coordinator](#), Elizabeth Hall, or the [Report and Response Coordinators](#) in the Equal Opportunity and Compliance Office. Please note that I am designated as a Responsible Employee, which means I must report to the EOC any information I receive about the forms of misconduct listed in this paragraph. If you'd like to speak with a confidential resource, those include Counseling and Psychological Services and the [Gender Violence Services Coordinators](#). Additional resources are available at [safe.unc.edu](#).