

CS4/591:Special Topics: High Performance Computing

Last changed on May 3, 2024

Lectures will be held in the CARC workshop room (1601 Central Ave).

Contents

1	Description	2
2	Contacts and Office Hours	3
3	Format	3
3.1	Classroom and Workstations	3
3.2	Curriculum Format	3
3.3	Individual Homework	3
3.4	Team Projects	4
3.5	Presentations	4
3.6	Exams	4
4	Resources	4
4.1	Slack Workspace	4
4.2	Center for Advanced Research Computing (CARC)	4
4.3	Class Cluster	4
5	Schedule	5
5.1	Week 1 (Jan 15 - 19): Introduction	5
5.2	Week 2 (Jan 22 - 26)	5
5.3	Week 3 (Jan 28 - Feb 2)	6
5.4	Week 4 (Feb 5 - 9)	6
5.5	Week 5 (Feb 12 - 16)	6
5.6	Week 6 (Feb 19 - 23)	7
5.7	Week 7 (Feb 26 - March 1)	7
5.8	Week 8 (March 4 - 8)	7
5.9	Week 9 (March 18 - 22)	8
5.10	Week 10 (March 25 - 29)	8
5.11	Week 11 (April 1 - 5)	8
5.12	Week 12 (April 8 - 12)	9
5.13	Week 13 (April 15 - 19)	9
5.14	Week 14 (April 22 - 26)	9
5.15	Week 15 (April 29 - May 3)	9
5.16	Final Exam Week (May 4 - 11)	10
6	Assignments and Grading	10
6.1	Late Projects	10

7 Attendance	11
8 Cheating	11
8.1 UNM Policy	11
8.2 Why it's a bad idea	11
8.3 What I do allow	11
8.4 What I don't allow	12
9 Innovative Teaching Study	12
10 UNM Resources and Student Rights	12
10.1 Student Health and Counselling (SHAC)	12
10.2 Center for Academic Program Support	12
10.3 Accessibility Resource Center	12
10.4 Equal Educational Opportunities	13

1 Description

This course is a senior/grad level introduction to high performance computing (HPC). The goal of the course is to familiarise students with building a cluster and how modern hardware architectures impact software performance. The skills taught in this course will prepare students to enter the HPC workforce as systems engineers and to better understand how to optimise code for use on HPC systems as end users. Topics include basics of Linux and its administration, building a high-performance computing cluster, optimisation of compiled code for HPC architectures (caching and accelerators), deployment and use of industry standard HPC tools such as spack, Slurm, wawulf, and hardware aware deployment. We will use benchmarks such as stream and HPL to understand the hardware. This course will not go into detail on MPI since that is covered in the Parallel Processing Course. Most of the work for this course will be team project based since that reflects the reality of devops for most professional systems analysts and programmers.

This syllabus is intended to make course expectations and policies clear so that you, the student, have the best chance possible of meeting those expectations. If anything is unclear please let me know so I can update the syllabus accordingly.

This syllabus is likely to be revised throughout the semester.

2 Contacts and Office Hours

Instructor: **Prof. Matthew Fricke**
Email: mfricke@unm.edu
Website: <http://www.cs.unm.edu/~mfricke>
Office Hour: Friday 2-3pm. Farris 3330.

Teaching Assistant: **Nicholas Bacon**
Email: nbacon@unm.edu
Office Hour: Noon-1pm, Wednesdays. Farris 3360

Faculty Assistant: **Ryan Scherbarth**
Email: rscherbarth@unm.edu
Office Hour: Mondays, 2-3pm. Farris 3360.

3 Format

This is a three credit hour course consisting of three 50 minute lectures per week to be held on Monday, Wednesday, and Friday at 11:00pm.

3.1 Classroom and Workstations

Classes will be held in the CARC workshop room. This workshop room is equipped with Linux workstations that you will be expected to use for in class assignments. You are welcome to bring a laptop to class, but due to the variety of possible OS configurations I will not be able to help you carry out the in class assignments on your laptop.

Adjacent to the workshop room is the Roadrunner room. The Roadrunner room houses the class clusters. The Roadrunner room also houses sensitive equipment for CARC. You will not be permitted access to the Roadrunner room unless I am present. After the first 2 weeks physical access should become unnecessary.

3.2 Curriculum Format

Classes will be taught using a mixture of lectures and hands on activities. The material you will learn is divided into two areas: 1) Linux and HPC Cluster Management and 2) High performance application management and benchmarking. Material on each topic will be presented each week. At first you will conduct your benchmarking on the UNM CARC clusters, but once your own clusters are deployed you will include them in your application deployment and benchmarking tasks.

3.3 Individual Homework

During the course you will be assigned homework to complete individually. See the "What I allow" section for guidelines, but you must write up the assignment in your own words, plots and/or code (i.e don't copy).

3.4 Team Projects

You will be assigned to a team for each project. Team membership will change for each project. Random assignment is essential for determining final grades. Requests to work with particular students will not be entertained. Projects will be assigned on Fridays. Your team will be given two weeks to complete your project with a progress update due on the first Friday.

3.5 Presentations

During the last week of the semester and the final exam week teams present the report from one of their projects. The audience will be HPC experts from industry, academia, and the national labs.

3.6 Exams

From time to time I will give in class exams. Exams will be announced at least 2 weeks ahead.

4 Resources

4.1 Slack Workspace

The slack workspace serves several purposes. 1) I will make announcements about the class on slack. There is no guarantee that I will also make those announcements in class. I may also make announcements in class that are not repeated on slack. You are responsible for both. 2) Slack allows class-wide discussion. Feel free to post questions about the material presented in class, the code you are writing for your assignments, and logistic questions related to the class.

4.2 Center for Advanced Research Computing (CARC)

CARC is the UNM high performance computing center. All students will be provided with a CARC account so you can compare the performance of the code you build on different hardware.

4.3 Class Cluster

Students will be assigned to teams. Each team will be responsible for assembling and configuring their own small cluster using parts provided by CARC. The cluster will be housed in the Roadrunner room at CARC but teams will be able to access their assigned cluster remotely. These clusters will be used over the course of the semester. Team composition will change each project.

5 Schedule

This is the planned schedule, any changes will be recorded here.

The schedule will follow two parallel tracks. 1) Building and Administering an HPC cluster, and 2) HPC Applications. Each week we will cover topics in each. This is so you don't have to wait until you have deployed your clusters to start learning about how they are used. The goal is to have these two track merge at the end of the semester.

For UNM semester dates and deadlines are here:

<https://registrar.unm.edu/semester-deadline-dates/spring-2024.html>

5.1 Week 1 (Jan 15 - 19): Introduction

General

- Presurvey: [Presurvey](#)
- Course Introduction
- Syllabus Review
- Running on CARC Systems
- A Brief History of HPC

Lecture Slides

- [Lecture 1: A Brief History](#)

Assignments

- Homework 1 Assigned. Complete the [Presurvey](#). Due by midnight on Friday the 26th of January.

5.2 Week 2 (Jan 22 - 26)

General

- git
- Jupyter Notebooks, Overleaf, and Matplotlib

System Administration

- IPMI, iDRAC, and the BIOS
- MBR and UEFI

Applications

- CPU Architectures (Vector Operations)

Lecture Slides

- [Lecture 2: iDRAC](#)
- [Lecture 3: Configure Your Linux Machine](#)

Assignments

- Homework 1 [Presurvey](#). Due by Midnight.

5.3 Week 3 (Jan 28 - Feb 2)

Jan 26th - Last Day to Add Classes

System Administration

- Linux: Devices, Disks

Applications

- Instruction Execution and Pipelining
- Memory and Cache

Lecture Slides

- [Lecture 4: Vector Operations and User Accounts](#)
- [Lecture 5: Devices](#)

Assignments

- Homework 2 on Vector Instructions Assigned. Due Midnight on Monday, February 5th.

Feb 2nd - Last day to DROP without "W" grade and receive a 100% tuition refund

5.4 Week 4 (Feb 5 - 9)

System Administration

- Linux: Filesystems
- Linux: Networks and Network Configuration
- Linux: Infiniband

Applications

- Compiler Optimisations

Lecture Slides

- [Lecture 6: Networking](#)
- [Lecture 7: Filesystems](#)

Assignments

5.5 Week 5 (Feb 12 - 16)

System Administration

- Linux: The Kernel

Applications

- Timing and Profiling (gcov and gprof)
- Top Down Optimisation

Lecture Slides

- [Lecture 8: Anatomy of a Machine Room](#)
- [Lecture 9: Booting and the Kernel](#)

Assignments

5.6 Week 6 (Feb 19 - 23)

System Administration

- Linux: Systemd and System V

Applications

- Compiler Optimisation and Profiling

Lecture Slides

- [Lecture 10: Systemd](#)
- [Lecture 11: Compiler Optimisation](#)

Assignments

5.7 Week 7 (Feb 26 - March 1)

System Administration

- Linux: Monitoring Processes and Resources

Applications

- High Performance Languages: C, FORTRAN, and Julia
- Open Multi-Processing (Shared MP)

Lecture Slides

- [Lecture 12: Shared Memory Parallelism](#)

Assignments

5.8 Week 8 (March 4 - 8)

System Administration

- Slurm (end user)

Lecture Slides

- [Lecture 13: Process and Resource Monitoring](#)
- [Lecture 14: Slurm](#)
- [Sample Midterm](#)

Spring Break (March 10 - 17)

5.9 Week 9 (March 18 - 22)

System Administration

- Warewulf

Lecture Slides

- [Lecture 15: Warewulf](#)

Assignments

5.10 Week 10 (March 25 - 29)

System Administration

- Slurm Installation and Administration
- Linux: Shell Scripting, AWK, and Sed

Applications

- Spack and LMOD

Lecture Slides

- [Lecture 16: Slurm Installation and Configuration](#)
- [Lecture 17: Spack and LMod](#)

Assignments

5.11 Week 11 (April 1 - 5)

System Administration

- Distributed Computing (MPI and Job Arrays)

Lecture Slides

- [Lecture 18: Distributed Computing](#)

Assignments

Project 1 Assigned
Due by 11:59pm on April 22nd.

- [Project 1: High Performance Linpack](#)

5.12 Week 12 (April 8 - 12)

Applications

- Scaleable Parallel Processing (Amdahl and Gustafson)
- Strategies for Parallelism and Halo Exchange

Lecture Slides

- [Lecture 19: Scaling](#)

Assignments:

- Project 1 due by 11:59pm on April 22nd.

April 12th - Last Day to Drop without Dean's Permission

5.13 Week 13 (April 15 - 19)

Applications

- CUDA Programming and GPUs

Lecture Slides

- [Lecture 20: GPU Computation](#)

Assignments:

- Project 1 due by 11:59pm on April 22nd.

5.14 Week 14 (April 22 - 26)

System Administration

- Containers

Applications

- Benchmarks: HPCG

Lecture Slides

- [Lecture 21: HPCG and Singularity](#)

5.15 Week 15 (April 29 - May 3)

System Administration

- RAID

Lecture Slides

- [Lecture 22: Network Throughput and RAID](#)
- [Sample Final](#)

Assignments

Project 2 due by 11:59pm on May 4th. Homework 6 due by 11:59pm on May 4th.

- [Project 2: High Performance Conjugate Gradient](#)

May 8 - Last day to DROP with Dean's Permission

5.16 Final Exam Week (May 4 - 11)

Final Exam. Wednesday May 8th. CARC Workshop Room

- Group A: 10am-11am.
- Group B: 11am-noon.

Sunday, May 17th, 2:30pm - Upload of Final Grades Expected by UNM.

6 Assignments and Grading

I use a 10 point fractional grading scale:

A+	A	A-	B+	B	B-	C+	I
98-100	93-97	90-92	87-89	83-86	80-82	77-79	Incomplete
C	C-	D+	D	D-	F		W
73-76	70-72	67-69	63-66	60-62	< 60		Withdrew

A+, A, and A- indicate excellent performance.

B+, B, and B- indicate good performance.

C+, C, and C- indicate satisfactory performance.

D+, D, and D- indicate less than satisfactory performance.

F indicates unsatisfactory performance.

W indicates withdrawal from the course.

The course grade is comprised of the following:

Homework: 30% (6 homeworks 5% each)

Projects: 40% (2 team projects 20% each)

Exams: 30% (2 exams 15% each)

6.1 Late Projects

Late homework and projects will not be accepted unless accompanied by an absence note provided by the Dean of Students Office for the period beginning on or before the *day the assignment was due* until the *day before* the assignment

is actually submitted. I do this because allowing late submissions without good cause results in more stress for students than having strict deadlines.

7 Attendance

- Please see the UNM policy on class attendance: [Class Attendance Policy](#)
- Class attendance will not be recorded.

8 Cheating

8.1 UNM Policy

- See the UNM policy on academic dishonesty: [Academic Dishonesty](#)

8.2 Why it's a bad idea

Cheating is harmful to you, the University, and your fellow students:

- The University warrants that receipt of grades and a degree accurately reflects the knowledge and skills of the recipient. Cheating undermines the value of the degrees and grades awarded by the University.
- Cheating makes assignments and exams look easier than they really are. This encourages instructors to make assignments and exams harder than they would otherwise.
- The desire to cheat indicates that there is an issue with developing the desired skills that needs to be resolved. Please approach lab instructors or the professor to discuss problems with the material and how to honestly do better in the class.
- Cheating indicates that the material is not being sufficiently mastered, this will likely result in difficulties when you take classes in the future that require this course.
- Trading your integrity for a better grade is something that will stay with you for the rest of your life, whether or not you are discovered.
- If you are caught cheating you will receive a failing grade for the course. However, the University may take further action including dismissal from the University.

8.3 What I do allow

You may use any resources to complete the projects including learning from other students and taking information from the internet and getting help from artificial intelligences, such as large language models. I encourage learning from your classmates.

8.4 What I don't allow

Quizzes will be used to test whether you have developed an individual understanding HPC fundamentals covered in the course. No collaboration or exam aids such as notes or devices may be used during the exam. Failure to follow this policy will result in an F on the exam or quiz, possibly the course, and may result in disciplinary action by the University. You may not copy quiz answers from other students.

9 Innovative Teaching Study

This class is paid for by a Department of Education Grant and funds from the School of Engineering Innovative Teaching Program. As part of those programs we will be collecting information about the class and your interactions. This requires IRB review and your consent. I will be distributing consent forms at the start of the semester. There is no obligation to participate and the data gathered will have no bearing on your grade in the class.

10 UNM Resources and Student Rights

10.1 Student Health and Counselling (SHAC)

“SHAC provides quality health and counseling services to all UNM students to foster student success.”

<https://shac.unm.edu/>

505.277.3136 (24-hr number)

10.2 Center for Academic Program Support

<https://caps.unm.edu/index.php>

“CAPS is the University of New Mexico’s learning center. We are a free service that provides academic support for undergraduate students at UNM through peer-tutoring. Our tutors are UNM students who are trained to create comfortable and welcoming spaces for students to learn and thrive in. We are passionate about helping students achieve individual academic success at UNM through peer tutoring and interactions.”

10.3 Accessibility Resource Center

<https://arc.unm.edu/>

In accordance with University Policy 2310 and the Americans with Disabilities Act (ADA), academic accommodations may be made for any student who notifies the instructor of the need for an accommodation. It is imperative that

you take the initiative to bring such needs to my attention, as I am not legally permitted to inquire. Students who may require assistance in emergency evacuations should contact the instructor as to the most appropriate procedures to follow. Contact Accessibility Resource Center at 277-3506 for additional information.

If you need an accommodation based on how course requirement interact with the impact of a disability, you should contact me to arrange an appointment as soon as possible. At the appointment we can discuss the course format and requirements, anticipate the need for adjustments and explore potential accommodations. I rely on the Disability Services Office for assistance in developing strategies and verifying accommodation needs. If you have not previously contacted them I encourage you to do so.

10.4 Equal Educational Opportunities

Our classroom and our university should always be spaces of mutual respect, kindness, and support, without fear of discrimination, harassment, or violence. Should you ever need assistance or have concerns about incidents that violate this principle, please let me know and/or access the resources available to you on campus, especially the LoboRESPECT Advocacy Center and the support services listed on its website (<http://loborespect.unm.edu/>).

Several Federal civil rights laws prohibit discrimination in programs or activities that receive Federal funds from the Department of Education. These laws prohibit discrimination on the basis of race, colour, and national origin (Title VI of the Civil Rights Act of 1964), sex (Title IX of the Education Amendments of 1972), age (Age Discrimination Act of 1975) and disability (The Americans with Disabilities Act of 2008). The Office for Civil Rights enforces these rights: <https://www2.ed.gov/about/offices/list/ocr/know.html>.

Please note that because the teaching assistants and I are considered “responsible employees” by the Department of Education, we are required to report any disclosure of gender discrimination (including sexual harassment, sexual misconduct, and sexual violence) to the university’s Title IX coordinator. For more information on the campus policy regarding sexual misconduct, please see: [UNM Policy](#).

UNM’s Title IX Coordinator, Angela Catena, can be reached at acatena@unm.edu.