CSci 8314 Fall 2026
Sparse Matrix Computations

General Information This course is an introduction to sparse matrix techniques and their applications. It will cover sparse matrices in the context of linear systems of equations and eigenvalue problems, as well as in sparse graphs, networks, embeddings, etc. The course starts with a general discussion of sparse matrices, their origins and how they are stored and exploited. Then it will briefly cover direct solution methods and iterative methods for solving sparse linear systems of equations and sparse eigenvalue problems. It will end with a coverage of various other topics related to sparsity, e.g., graph-based algorithms in machine learning, and basic optimization and techniques for nonlinear equations. This course combines theory (some complexity, convergence theory) and practice (implementing sparse matrix techniques, efficiency, applications). One of the requirements of the course is the completion of a term project, possibly using available software packages.
The course will consist of lectures given by the instructor and in class discussions. It qualifies as a 50% project toward a computer science plan C (coursework only) Masters.

- **Class Schedule:** 2:30pm - 3:45 pm M-W, Keller Hall 3-125
- **Instructor:** Yousef Saad `http://www.cs.umn.edu/~saad`
  [e-mail: saad@umn.edu]
  Office: 5 -225B Keller Hall – Office Phone: (612) 624 – 7804.
- **Prerequisite:** csci 5304 (or equivalent)
- **Office hours:** See class web-page

Textbook Part of the lectures will be based on my book: *"Iterative methods for sparse linear systems (2nd edition)"* which is posted (free download) – see

`http://www.cs.umn.edu/~saad/books`.

This will be supplemented by articles and lecture notes.
Supplemental reading: *Direct methods for sparse linear systems*, by T. A. Davis, SIAM publishing, 2006. (see also accompanying software CSparse noted in the resources link in the class web-site (see below). *You need not purchase the books*

listed above. The first book is posted for free and will be only used for part of the material.

Lecture Notes Lecture notes on some of the chapters will be posted regularly, usually in advance of the lecture. These will be posted on github - see canvas for link. The lecture notes will be posted by topic rather than lecture by lecture.

Project To pass this course, each student will need to complete a class project. This will generally be on the use of sparse matrix techniques in specific application areas. You can for example write a comprehensive survey on a selected topic related to the course. But you can also write a paper on how to exploit a specific class of sparse matrix methods in your own research. In the past the projects ranged from exhaustive survey articles on a specific topic (e.g., 'sparse matrix techniques in text mining') to implementations /comparisons of (a) specific method (s) say for solving sparse linear systems; or a specific theme in your own research that illustrates the use of sparse matrix methods. Minimal requirements, as well as the details on grading criteria, will be specified later in the semester.

Student Presentations: The current plan is to have each student who takes the class for credit give a 25mn presentation on their project. (Alternatively a few students will provide longer reports but will not present their projects). You can work alone or in a team of 2 [more if justified]. Since these presentations (and reports) constitute supplemental course material for the class, a good part of the grade will be based on their clarity.

Evaluation The evaluation of your performance in this class will be based on the following:

- 60 % : *Term Project* - (based on report and presentation components)
- 24 % : *Three quizzes* at 8 % each [actually three best out of 4 quizzes given.] Details below
- 16 % : *Class participation and attendance*

Quizzes will be at most 15mn long each and will be related to the course material as well as the in-class exercises. There will be no make-up quizzes. The final letter grades will be assigned based on the following scale, where T is the total score (out of 100) you achieved in the class.

A : $100 \geq T \geq 92$	**A-** : $92 > T \geq 87$	**B+** : $87 > T \geq 83$
B : $82 > T \geq 77$	**B-** : $77 > T \geq 72$	**C+** : $72 > T \geq 65$
C : $65 > T \geq 60$	**C-** : $60 > T \geq 55$	**D+** : $55 > T \geq 50$
D : $50 > T \geq 40$	**F** : $40 > T$	

Topics to be covered. The topics to be covered will depend on participants and their interests and will be discussed and finalized during the first two lectures. However the following is a list of *core topics* in sparse matrix computations which will certainly be taught:

1. Sparse matrices and their origin. Graph representation of sparse matrices, sparse graphs, Discretization of Partial Differential Equations. Electrical networks, Information retrieval, ...
2. Storage schemes for sparse matrices. Regular and irregular structures.
3. Direct solution methods; Variants of Gaussian Elimination; Permutations and orderings; Band and envelope methods; Cuthill-Mc Kee and reverse Cuthill-Mc Kee orderings; Graph representation; Elimination tree; The frontal and multifrontal approaches; Minimal degree and nested dissection orderings.
4. More on Graphs: Graph Laplaceans, graph partitioning, clustering; Application: image segmentation; Networks and graphs; measures of node centrality.
5. Iterative methods; Projection methods; One-dimensional case: steepest descent, minimal residual methods; Krylov subspace methods; Conjugate gradient (CG) method; basic convergence theory; Connection to Lanczos tridiagonalization and to orthogonal polynomials; In brief: The idea of preconditioning.
6. Eigenvalue problems. Types of problems; Subspace iteration; Krylov methods; Arnoldi's method; The Lanczos algorithm; Nonsymmetric Lanczos.

A request: If you attend class you agree not to use cell-phones during lectures. Laptops are fine as long as you access class-related material while the lecture is ongoing.