CS 322: Introduction to Scientific Computing
Spring 2003
Problem Set 4

Handed out: Wed., Mar. 26.

Due: Wed., Apr. 2 in lecture.

The policies for this (and other problem sets) are as follows:

- You should hand in your on-time problem set in lecture in the box at the front of the room on the day it is due. Problem sets handed in elsewhere (TA's office, Upson 303, etc.) will be considered late. See the next bullet.
- Late papers may be handed in up to 24 hours late. For instance, this problem set may be handed in up to 11:00 am on Apr. 3. You can hand in a late paper in Upson 303. Late papers get an automatic deduction of 10%. The full late penalty is applied even if you turn in part of the solution on time.
- Problem sets may be done individually or in teams of two. Put your name or names on the front page. Re-read the academic integrity statement on the web for the policy concerning working in larger groups.
- Problem sets count for 20% of the final course grade. The lowest scoring problem set will be dropped.
- You need Matlab for some of the questions. Matlab is available in the following CIT labs: Upson and Carpenter.
- If you need clarification for a homework question, please either ask your question in section, lecture, or office hours or else post it to the newsgroup cornell.class.cs322. The professor reads this newsgroup and will post an answer.
- Write your section number (like this: "Section 2") at the top of the front page of your paper, and circle it. This is the section where your graded paper will be returned. As a reminder, Section 1 is Th 12:20, Section 2 is Th 3:35, Section 3 is F 2:30, Section 4 is F 3:35.

1. Carry out on paper Gaussian elimination with and without partial pivoting on the following linear system:

$$\begin{pmatrix} 10^{-5} & 1 \\ 1 & 1 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \begin{pmatrix} 1 \\ 3 \end{pmatrix}.$$

   [**Note correction, posted 3/27/03, to (1,1) entry of the matrix.**] Carry out this solution process using four-digit decimal arithmetic, i.e., after every arithmetic operation, round the answer so as to carry only four significant decimal digits. The answer without pivoting should be much less accurate than the answer with pivoting. (Use Matlab to find the exact answer to this system of equations. In matlab, `format long e` shows all significant digits in printouts.)

   In the case of no pivoting, identify a particular operation in which catastrophic cancellation error occurs that is responsible for the inaccurate answer.

2. Cubic spline interpolation is generally a well-conditioned problem, except in the case that two of the knots are very close. In this case, it is ill-conditioned. Carry out two or three experiments using the cubic spline functions from PS2 to demonstrate these claims. To talk about conditioning, you must have a well-defined notion of what is the "input" and "output" are. The "input" for the cubic spline problem is the list of datapoints and the values for the end-conditions. The "output" for the cubic spline function is the interpolant. For measuring sensitivity of the output, you can treat the output as a long vector by evaluating the interpolant at a grid of closely space points and storing the interpolant values at those points in a vector. To answer this question, you must compare the solution for a given input to the solution for a slightly perturbed version of that same input. Hand in a description of your experiments, listings of relevant m-files and Matlab printouts, and what you discovered about sensitivity of the output with respect to the input.

3. Given three $n \times n$ matrices B, C, D, suppose we suspect that D is actually a linear combination of B and C. To verify whether this is true, we could solve the problem of minimizing $\|D - sB - tC\|$, where the minimum is taken over all scalars s, t. Suppose the norm used is the Frobenius norm. In this case, show that the problem is really linear least-squares by explaining how to rewrite it in the standard form for linear least-squares, that is, in the form of minimizing $\|A\mathbf{x} - \mathbf{b}\|_2$ where $A, \mathbf{x}, \mathbf{b}$ involve B, C, D, s, t. Furthermore, explain what the assumption that A has full-rank means in terms of the original problem data B, C, D.