INDIANA UNIVERSITY

RESEARCH REPORTS AND INFORMATION ON HIGH PERFORMANCE COM-

The emily Sparse Matrix Visualization System

Solving linear systems of equations may be the most ubiquitous problem in scientific, engineering, and statistical computing. In spite of the tremendous growth in the size of computer memory and computational rates, many problems remain intractable because the linear subproblems are too large or have other unfavorable properties. Linear systems can be written as $Ax = b$, where A is an $n \times n$ matrix and b a given n-vector. Current research focuses on systems that are large, sparse, and unstructured. Large is a relative term. While a system of 30 unknowns was considered large in the 1940s, modern codes for three-dimensional fluid flows routinely have hundreds of thousands of unknowns.

The unstructured property of modern linear systems is usually linked to complicated physical domains, such as the interior of automobile engines or nuclear reactors. Simplex methods for linear programming also require solving highly unstructured linear systems. The lack of structure implies, among other things, the need for a common data structure capable of holding such systems and which can support the operations typically performed on the matrices.

The size of the problems also mandates using parallel machines. Because shared memory machine architectures usually are not scaleable (that is, able to grow with the size of the problem and still maintain high performance) another consideration is the distribution of the matrix data across distributed memory architectures. This requires partitioning the data, preferably with processors that receive equal amounts of data, and keeping to a minimum the amount of run-time communication between them.

To help understand the structure of large matrices we have developed a Motif-based program, emily, which can display extremely large sparse linear systems. Using a mouse, users can zoom in on selected regions of the matrix and expand them to see more detail. Up to 10 zoom windows can be active at one time. The windows that open display as row and column numbers the matrix coordinates of the region being viewed, allowing users to identify individual entries. emily then reads matrices stored in standard Harwell/Boeing format files. Users can also call emily from within Matlab, allowing them to perform manipulations on the matrices at a high level and quickly view the results.

Currently lacking is a universally good color map for the magnitude of the matrix entries. It helps to be able to change the background color and the scheme that maps magnitudes to colors, switching among various colors to bring out different features. The program uses tis own color map, so emily can allocate most of the color map for displaying matrix values without interference. Status areas show the current position in the matrix (as row and column matrix indices), a color bar that shows the range of colors and the numerical maximum and minimum values of the matrix, and the coordinates of the sub-matrix of interest. Researchers can use emily to view extremely large matrices; it employs a unique compression algorithm to display matrices of arbitrary size on a 1024x1024 display.

Two versions of emily exist: a Matlab version implemented using a cmex interface and a traditional command-line interface version. Both can accept filenames as input; the file is assumed to be in Harwell-Boeing format. The Matlab version can also handle either sparse or dense matrix Matlab input matrices.

For more information:

Randall Bramley

beyond boundaries