The Importance of Highly Efficient Computational Kernels for All Block Sizes

Barry Smith

http://www.mcs.anl.gov/petsc/petsc.html

ARGONNE NATIONAL LABORATORY • UNIVERSITY OF CHICAGO

Clique (supernode) example: ILU preconditioner

Identical nodes: Matrix-vector product example

I-node Example: Matrix-vector product

Demonstrates that even small blocks can drastically improve performance, even without assembler code.

Example with I-nodes of only size 3.

FLOPS

	IBM rs6000	DEC alpha
Basic code	14	15
I-node code	21	30

Small Block Sizes

Small blocks naturally occur in sparse matrix computation.

If they are not taken advantage of effectively, floating point performance will be far below machine peak.

Since the blocks are very small, the computational kernels must be **very** light-weight and have little overhead.

Basic Requirements

- **No** function overloading.
- **No** error checking.
- **No** character string arguments.
- Should be inlineable.
- Should allow caching of reusable structure information.
- Separate operations for stride 1.
- Many of the "BLAS" operations.