

Optimization Lab

Goals:

- See how compiler options help you optimize the performance of hand-coded routines.
- See how performance can be improved by calling numerical libraries.

You will be working with three different versions of a code to solve a system of linear equations via LU factorization. The tarball contains all three codes, together with a makefile to compile them and a script to submit them to the scheduler. Embedded in the script are instructions that time each code and put the timing data into an output file.

Unpack the source code with:

```
$ cd ~  
$ tar xvfz ~/tg459572/LABS/ludecomp.tgz
```

You will find a directory called ludecomp. In it are these three code versions:

- nr.c – uses code copied from a book called Numerical Recipes. It does not require any external libraries.
- gsl.c – calls the GNU Scientific Library. You need to use the module command to link and load this library in the TACC environment.
- lapack.c – calls the standard interface to LAPACK. This call may be linked against any compatible, optimized library that performs linear algebra: either the Math Kernel Library if you are using the Intel compiler, or one of the two PGI LAPACK libraries if you are using the PGI compiler on Ranger. Load the appropriate module for whichever compiler you choose.

The Makefile has many targets. Here are the two you might want to use the most:

- make - This makes all three versions of the program: nr, gsl, and lapack. If make fails, it is likely because it wants libraries that are not currently loaded.
- make clean - This deletes all binaries you compiled for a clean start if you are switching from PGI to Intel, or vice versa.

To get started on Ranger, here are the steps to follow:

1. Add GSL to the currently-loaded modules: `module load gsl`
2. In the directory `~/ludecomp`, type: `make`
3. Submit the `job.sge` script with your account designated. A way to do this is `qsub -A your_account job.sge`
4. If you like, you can view the results of your runs: `less results.txt`

Everybody's results will be recorded on The Eric Chen Scoreboard, http://consultrh5.cac.cornell.edu/emc256/intro_to_ranger/.

Our results so far are based on the PGI compiler and a “vanilla” LAPACK. Let’s compare them to equivalents based on the Intel compiler and MKL.

1. Switch to the Intel compiler suite: `module swap pgi intel`
2. Add MKL to the currently-loaded modules: `module load mkl`
3. Edit the first few lines of the Makefile so that the COMPILER is `icc`, LAPACKLIB is `mkl`, and FFLAGS is “`-g -xW -ipo`”
4. Again, submit the `job.sge` script with your account designated.
5. Examine `results.txt`, where the new results will be appended after the job runs; compare the times the codes took to run.

Now take a minute to look at the codes and evaluate each based on these criteria:

1. How many lines of code did it take to implement the solution?
2. How fast does it run? Look at `results.txt` when it finishes.
3. For each version, how hard would it be to swap in a different algorithm by, for instance, substituting an iterative solver, or using a sparse-matrix solver?
4. Can any of these codes run multithreaded? Can they run distributed, using MPI? You may need to Google the libraries to figure this out.

We’ve already seen what the two compilers can do with `-g`, which is the debug suite of options. Next let’s experiment with the compilers’ optimization options, to assess how they affect running times.

1. Use the `module` command to select a compiler (PGI or Intel), then to load a matching math library that supports LAPACK (ACML for PGI; MKL for Intel).
2. Edit the first few lines of the Makefile so that the COMPILER matches your module choice: `pgcc` for PGI, `icc` for Intel. Choose LAPACKLIB appropriately as well. Some possible FFLAGS are listed below.
3. Compile the three codes: `make`
4. Submit the codes to the scheduler: `qsub -A your_account job.sge`
5. Again examine `results.txt`.
6. Try some other choices of compiler and optimizations and see what is fastest. For the codes that call libraries, how does your choice of options affect the performance? (Remember, you’re not compiling the libraries!)

Here are the recommended choices.

PGI: `-O3 -fast -tp barcelona-64`

Intel: `-O3 -ipo -xW`

Don’t exceed `-O2` without checking whether your output is correct.

For **PGI**, here are some other common compiler options to try:

- O3 - Performs some compile time and memory intensive optimizations in addition to those executed with `-O2`, but may not improve performance for all programs.
- Mipa=fast,inline - Creates inter-procedural optimizations. There is a loader problem with this option.

- tp barcelona-64 - Includes specialized code for the AMD Barcelona chip.
- fast - A catchall for: -O2 -Munroll=c:1 -Mnoframe -Mlre -Mautoinline
-Mvect=sse -Mscalarsse -Mcache_align -Mflushz
- g, -gopt - Produces better debugging information (without disabling optimization).
- mp - Enables the parallelizer to generate multi-threaded code based on the OpenMP directives.
- Minfo=mpi,ipa - Provides information about OpenMP, and inter-procedural optimization.

For **Intel**, here are some other common compiler options to try:

- O3 - More than O2, but maybe not faster.
- ipo - Creates inter-procedural optimizations.
- vec_report[0|..|5] - Controls the amount of vectorizer diagnostic information.
- xW - Includes specialized SSE and SSE2 instructions (recommended for Ranger).
- xO - Includes specialized SSE3 instructions.
- xSSE4.2 - Recommended for Lonestar so that it uses SSE4.2.
- fast - Includes: -ipo, -O2, -static [DO NOT USE] static load not allowed because only dynamic loading is allowed.
- g - Produces better debugging information.
- fp-stack-check - catches errors in the floating point stack, at the cost of speed
- openmp - Enable OpenMP directives
- openmp_report[0|1|2] - OpenMP parallelizer diagnostic level.

Extra credit: The ACML_MP and MKL libraries have built-in support for OpenMP multithreading. To enable it, you simply set the following environment variable:
export OMP_NUM_THREADS=4

Try this! Specify either acml_mp or mkl as the LAPACKLIB in the Makefile, run make, uncomment the above line in job.sge, and submit the job. Does the time improve or not, when compared to leaving this variable unset? What if you try a different number of threads? (Remember, Ranger nodes have 4 processors and a total of 16 cores. Note, the default value of OMP_NUM_THREADS is 1, so that when 16 MPI processes share a node they won't create chaos by forking 16xN threads.)