

# Some notes on efficient computing and high performance computing environments

---

Abhi Datta<sup>1</sup>, Sudipto Banerjee<sup>2</sup> and Andrew O. Finley<sup>3</sup>

July 31, 2017

<sup>1</sup>Department of Biostatistics, Bloomberg School of Public Health, Johns Hopkins University, Baltimore, Maryland.

<sup>2</sup>Department of Biostatistics, Fielding School of Public Health, University of California, Los Angeles.

<sup>3</sup>Departments of Forestry and Geography, Michigan State University, East Lansing, Michigan.

## **Code implementation**

---

Very useful libraries for efficient matrix computation:

1. Fortran BLAS (Basic Linear Algebra Subprograms, see Blackford et al. 2001). Started in late 70s at NASA JPL by Charles L. Lawson.
2. Fortran LAPACK (Linear Algebra Package, see Anderson et al. 1999). Started in mid 80s at Argonne and Oak Ridge National Laboratories.

Modern math software has a heavy reliance on these libraries, e.g., Matlab and *R*. Routines are also accessible via C, C++, Python, etc.

Many improvements on the standard BLAS and LAPACK functions, see, e.g.,

- Intel Math Kernel Library (MKL)
- AMD Core Math Library (ACML)
- Automatically Tuned Linear Algebra Software (ATLAS)
- Matrix Algebra on GPU and Multicore Architecture (MAGMA)
- OpenBLAS <http://www.openblas.net>
- vecLib (for Mac users only)

## Key BLAS and LAPACK functions used in our setting.

Function	Description
dpotrf	LAPACK routine to compute the Cholesky factorization of a real symmetric positive definite matrix.
dtrsv	Level 2 BLAS routine to solve the systems of equations $\mathbf{Ax} = \mathbf{b}$ , where $\mathbf{x}$ and $\mathbf{b}$ are vectors and $\mathbf{A}$ is a triangular matrix.
dtrsm	Level 3 BLAS routine to solve the matrix equations $\mathbf{AX} = \mathbf{B}$ , where $\mathbf{X}$ and $\mathbf{B}$ are matrices and $\mathbf{A}$ is a triangular matrix.
dgemv	Level 2 BLAS matrix-vector multiplication.
dgemm	Level 3 BLAS matrix-matrix multiplication.

# Computing environments

---

Consider different environments:

1. A **distributed system** consists of multiple autonomous computers (nodes) that communicate through a network. A computer program that runs in a distributed system is called a distributed program. Message Passing Interface (MPI) is a specification for an Application Programming Interface (API) that allows many computers to communicate.

Consider different environments:

1. A **distributed system** consists of multiple autonomous computers (nodes) that communicate through a network. A computer program that runs in a distributed system is called a distributed program. Message Passing Interface (MPI) is a specification for an Application Programming Interface (API) that allows many computers to communicate.
2. A **shared memory multiprocessing system** consists of a single computer with memory that may be simultaneously accessed by one or more programs running on multiple Central Processing Units (CPUs). OpenMP (Open Multi-Processing) is an API that supports shared memory multiprocessing programming.
3. A **heterogeneous system** uses more than one kind of processor, e.g., CPU & (Graphics Processing Unit) GPU or CPU & Intel's Xeon Phi Many Integrated Core (MIC).





Which environments are right for large  $n$  settings?

- MCMC necessitates iterative evaluation of the likelihood which requires operations on large matrices.
- A specific hurdle is **factorization** to computing determinant and inverse of large dense covariance matrices.
- We try to model our way out and use computing tools to overcome the complexity (e.g., covariance tapering, Kaufman et al. 2008; low-rank methods, Cressie and Johannesson 2008; Banerjee et al. 2008, etc.).
- Due to **slow network communication** and transport of submatrices among nodes distributed systems are not ideal for these types of iterative large matrix operations.

- My lab currently favors **shared memory multiprocessing** and **heterogeneous** systems.
- Newest unit is a Dell Poweredge with 384 GB of RAM, 2 threaded 10-core Xeon CPUs, and 2 Intel Xeon Phi Coprocessor with 61-cores (244 threads) running a Linux operating systems.
- Software includes OpenMP coupled with Intel MKL. MKL is a library of highly optimized, extensively threaded math routines designed for Xeon CPUs and Phi coprocessors (e.g., BLAS, LAPACK, ScaLAPACK, Sparse Solvers, Fast Fourier Transforms, and vector RNGs).



So what kind of speed up to expect from threaded BLAS and LAPACK libraries.



## R and threaded BLAS and LAPACK

- Many core and contributed packages (including *spBayes*) call BLAS)and LAPACK Fortran libraries.
- Compile *R* against threaded BLAS and LAPACK provides substantial computing gains:
  - processor specific threaded BLAS/LAPACK implementation (e.g., MKL or ACML)
  - processor specific compilers (e.g., Intel's *icc/ifort*)

For Linux/Unix: compiling *R* to call MKL's BLAS and LAPACK libraries (rather than stock serial versions). Change your `config.site` file and `configure` call in the R source code directory.

### My `config.site` file

```
CC=icc
CFLAGS="-g -O3 -wd188 -ip -mp"
F77=ifort
FLAGS="-g -O3 -mp -openmp"
CXX=icpc
CXXFLAGS="-g -O3 -mp -openmp"
FC=ifort
FCFLAGS="-g -O3 -mp -openmp"
ICC_LIBS=/opt/intel/composerxe/lib/intel64
IFC_LIBS=/opt/intel/composerxe/lib/intel64
LDFLAGS="-L$ICC_LIBS -L$IFC_LIBS -L/usr/local/lib"
SHLIB_CXXLD=icpc
SHLIB_CXXLDLDFLAGS=-shared
```

Compiling *R* to call MKL's BLAS and LAPACK libraries (rather than stock serial versions). Change your `config.site` file and `configure` call in the R source code directory.

### My configure script

```
MKL_LIB_PATH="/opt/intel/composerxe/mkl/lib/intel64"
```

```
export LD_LIBRARY_PATH=$MKL_LIB_PATH
```

```
MKL="-L${MKL_LIB_PATH} -lmkl_intel_lp64  
      -lmkl_intel_thread  
      -lmkl_core -liomp5  
      -lpthread -lm"
```

```
./configure --with-blas="$MKL" --with-lapack  
            --prefix=/usr/local/R-mkl
```

For many BLAS and LAPACK functions calls from *R*, expect near linear speed up ...

```
Terminal - andy@quercus:~/mic/samples/mkl

File Edit View Terminal Tabs Help

1 [|||||] 36.5% 11 [|||||] 100.0% 21 [|||||] 2.0% 31 [|||||]
2 [|||||] 3.9% 12 [|||||] 100.0% 22 [|||||] 3.2% 32 [|||||]
3 [|||||] 0.0% 13 [|||||] 0.0% 23 [|||||] 100.0% 33 [|||||]
4 [|||||] 0.0% 14 [|||||] 100.0% 24 [|||||] 100.0% 34 [|||||]
5 [|||||] 100.0% 15 [|||||] 100.0% 25 [|||||] 0.0% 35 [|||||]
6 [|||||] 0.0% 16 [|||||] 0.0% 26 [|||||] 100.0% 36 [|||||]
7 [|||||] 100.0% 17 [|||||] 100.0% 27 [|||||] 0.0% 37 [|||||]
8 [|||||] 100.0% 18 [|||||] 0.0% 28 [|||||] 0.0% 38 [|||||]
9 [|||||] 71.4% 19 [|||||] 100.0% 29 [|||||] 2.0% 39 [|||||]
10 [|||||] 100.0% 20 [|||||] 100.0% 30 [|||||] 0.0% 40 [|||||]

Mem[|||||] 34231/3075596 Tasks: 52, 36 thr; 21 running
Swap[|||||] 0/409596 Load average: 0.37 7.16 4.04
UpTime: 8 days, 23:13:24

PID USER PRI NI VIRT RES SHR S CPU% MEM% TIME+ Command
22991 andy 20 0 36.9G 25.4G 6124 S 1802 6.7 1:43.90 ./a.out 40000 20
23007 andy 20 0 36.9G 25.4G 6124 R 100 6.7 0:04.92 ./a.out 40000 20
23009 andy 20 0 36.9G 25.4G 6124 R 100 6.7 0:04.92 ./a.out 40000 20
23011 andy 20 0 36.9G 25.4G 6124 R 100 6.7 0:04.92 ./a.out 40000 20
23015 andy 20 0 36.9G 25.4G 6124 R 100 6.7 0:04.92 ./a.out 40000 20
23016 andy 20 0 36.9G 25.4G 6124 R 100 6.7 0:04.92 ./a.out 40000 20
23020 andy 20 0 36.9G 25.4G 6124 R 100 6.7 0:04.92 ./a.out 40000 20
23018 andy 20 0 36.9G 25.4G 6124 R 100 6.7 0:04.92 ./a.out 40000 20
23019 andy 20 0 36.9G 25.4G 6124 R 100 6.7 0:04.92 ./a.out 40000 20
23008 andy 20 0 36.9G 25.4G 6124 R 100 6.7 0:04.92 ./a.out 40000 20
23013 andy 20 0 36.9G 25.4G 6124 R 100 6.7 0:04.92 ./a.out 40000 20
23018 andy 20 0 36.9G 25.4G 6124 R 100 6.7 0:04.92 ./a.out 40000 20
23023 andy 20 0 36.9G 25.4G 6124 R 100 6.7 0:04.92 ./a.out 40000 20
23012 andy 20 0 36.9G 25.4G 6124 R 100 6.7 0:04.92 ./a.out 40000 20
23014 andy 20 0 36.9G 25.4G 6124 R 100 6.7 0:04.92 ./a.out 40000 20
23017 andy 20 0 36.9G 25.4G 6124 R 100 6.7 0:04.92 ./a.out 40000 20
23021 andy 20 0 36.9G 25.4G 6124 R 100 6.7 0:04.92 ./a.out 40000 20
23006 andy 20 0 36.9G 25.4G 6124 R 100 6.7 0:04.92 ./a.out 40000 20
23022 andy 20 0 36.9G 25.4G 6124 R 99.0 6.7 0:04.92 ./a.out 40000 20
22125 andy 20 0 300M 20628 10552 S 3 0 0.0 1:12.73 /opt/intel/mic/bin/micsnc-gui
12347 andy 20 0 110M 2288 1272 R 3 0 0.0 0h44:24 http
```

Mac users, see some vecLib linking hints at  
<http://blue.for.msu.edu/GWEDA17/BLASHints.txt>

Linux/Unix users, compile R with MKL as described above, or compile OpenBLAS and just redirect R's default libRblas.so to libopenblas.so, e.g.,

```
/usr/local/lib/R/lib/libRblas.so ->  
    /opt/OpenBLAS/lib/libopenblas.so
```

See <http://blue.for.msu.edu/comp-notes> for some simple examples of C++ with MKL and Rmath libraries along with associated Makefile files.