

A Computational Mathematician's Guide to High Performance Computing

Shane Sawyer

Research Associate, Joint Institute for Computational Sciences
University of Tennessee, Knoxville

shane-sawyer@tennessee.edu

February 18, 2015



Acknowledgements

Special thanks to colleagues who provided valuable feedback:
Dr. Vince Betro, Dr. Travis Thompson, Dr. Ryan Glasby,
and Dr. Taylor Erwin.



Overview of talk

- 1 Introduction** What is high performance computing
- 2 Programming Models** Threads, Processes, Distributed/Shared Memory
- 3 Hardware** Multicore processors, memory hierarchies, accelerators
- 4 Libraries** Existing libraries that simplify development and deliver performance
- 5 Optimization** When to optimize, available tools
- 6 Taking the Next Step** Where to find additional resources



Introduction

1 Introduction

- Defining High Performance Computing

2 Programming Models

- Nomenclature
- Models
- Examples

3 Hardware

- CPUs
- Memory

4 Libraries

- blaze-lib
- PETSc
- deal.II

5 Optimization

- Intel VTune
- ITAC
- TAU

6 Additional Resources



High performance computing is ubiquitous.



Titan supercomputer at ORNL.
Image courtesy of Oak Ridge
National Laboratory.

- Broadly defined: it is a collection of resources that offer more performance than desktops.
- Computational tasks that are too large (memory/operations) for a single resource.

Clusters vs. Supercomputers

■ Clusters

- Collection of resources (servers, desktops, ...)
- Interconnect (ethernet, InfiniBand, ...)
- Commodity Linux distributions

■ Supercomputer


- Collection of specialized resources. Typically higher density.
- High-speed interconnects. Higher performance networking topologies.
- Customized compilers, tools, and OS.



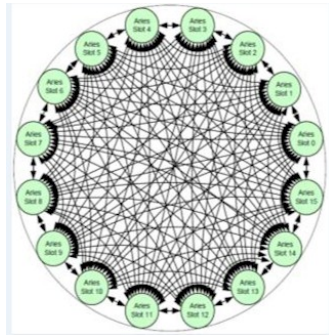
Clusters vs. Supercomputers

ACCRE Cluster Diagram

Blue = Gigabit Ethernet, Red = 10 Gigabit Ethernet



Network diagram from ACCRE,
Vanderbilt University.



Cray's Dragonfly topology.
Image credit: Timothy
Prickett Morgan, The
Register.



Programming Models

- 1 Introduction
 - Defining High Performance Computing
- 2 Programming Models
 - Nomenclature
 - Models
 - Examples
- 3 Hardware
 - CPUs
 - Memory
- 4 Libraries
 - blaze-lib
 - PETSc
 - deal.II
- 5 Optimization
 - Intel VTune
 - ITAC
 - TAU
- 6 Additional Resources



Nomenclature

- **Threads** A single stream of execution.
- **Processes** Complete program with address space, code, I/O handles, ...
- **Shared Memory** Single pool of memory shared by resources. Explicit protection.
- **Distributed Memory** Memory is spread across resources. Explicit exchange of information.



Models

- 1 OpenMP** An API, compiler directives, and runtime engine for shared memory, multithreading.
- 2 MPI** A library and runtime for distributed memory parallel programming. Explicit message exchange.
- 3 Hybrid** Use OpenMP on node and MPI between nodes for communication.
- 4 Heterogeneous** Conventional resources with accelerators (GPU, Xeon Phi, FPGA, ...)



OpenMP Example

Dot product of two vectors.

```
1 #include <omp.h>
2 int main(void)
3 {
4     double *a, *b, dotp;
5     int i;
6     // ... initialize and allocate a and b
7
8
9     #pragma omp parallel for shared(a,b,N) \
10     private(i) reduction(+ : dotp)
11     for (i=0; i < N; ++i)
12         dotp += a[i]*b[i];
13
14
15     return 0;
16 }
```



MPI Example

Dot product of two vectors.

```
1 #include <mpi.h>
2 int main(int argc, char* argv[])
3 {
4     double *a, *b, dotp, temp;
5     int i;
6     // initialize MPI
7     MPI_Init(&argc,&argv);
8     MPI_Comm_rank(MPI_COMM_WORLD,&my_rank);
9     MPI_Comm_size(MPI_COMM_WORLD,&np);
10
11     // ... initialize and allocate a and b, determine start and end
12     dotp = temp = 0.;
13     for ( i=START; i < END; ++i )
14         temp += a[i] * b[i];
15
16     MPI_Reduce(&temp,&dotp,1,MPI_DOUBLE,MPI_SUM,root,MPI_COMM_WORLD);
17
18     MPI_Finalize();
19     return 0;
20 }
```



Hardware

- 1** Introduction
 - Defining High Performance Computing
- 2** Programming Models
 - Nomenclature
 - Models
 - Examples
- 3** Hardware
 - CPUs
 - Memory
- 4** Libraries
 - blaze-lib
 - PETSc
 - deal.II
- 5** Optimization
 - Intel VTune
 - ITAC
 - TAU
- 6** Additional Resources



Multicore Processors are everywhere.



Quad-core AMD Opteron processor.

Image credit: American Micro
Devices, Inc.

- "The Free Lunch is Over" – Herb Sutter.
- Proliferation of multicore processors.
- Algorithms pushed towards parallelization.



Granularity

- **Socket** The physical packaging of cores with cache and interconnect.
- **Core** A complete processing element.
- **SIMD** Registers and execution units that allow one operation performed on multiple data in one tick.



NUMA and the Latency Hierarchy

NUMA – Non-Uniform Memory Access




Image credit: Jon Stokes, ArsTechnica.



NUMA and the Latency Hierarchy

Memory Hierarchy

- 1 Registers – 1 cycle
- 2 Cache – L1 (4 cycles) → L2 (10 cycles) → L3 (40-75 cycles)
- 3 RAM – 100ns
- 4 Disk – 2ms



Data Locality

- Goal: Maintain high FP intensity
- Algorithms: Reuse data in cache
- Reflected in many modern linear algebra packages
- Example: create a tiling of matrices for multiplication.



Matrix Tiling




Image credit: Nvidia.

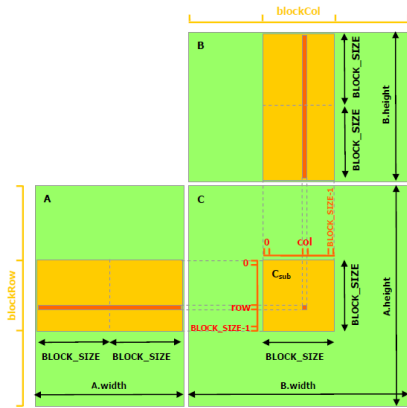


Image credit: Nvidia.



Libraries

1 Introduction

- Defining High Performance Computing

2 Programming Models

- Nomenclature
- Models
- Examples

3 Hardware

- CPUs
- Memory

4 Libraries

- blaze-lib
- PETSc
- deal.II

5 Optimization

- Intel VTune
- ITAC
- TAU

6 Additional Resources



Libraries for Computational Mathematics

- Parallel programming is challenging – changing technology, effort spent on low-level details
- Using libraries allows quicker development time, less debugging, abstracts communication details, lets application writers focus on their problem
- Plethora of excellent libraries for computational mathematics: linear algebra, nonlinear solvers, graph partitioning, PDEs, ...



Linear Algebra

- Modern libraries take advantage of hardware advances: ATLAS, Intel Math Kernel Library (MKL), PLASMA (ICL), blaze-lib
- Libraries for Heterogenous systems:
 - Intel Xeon Phi: MKL (with automatic offload support), MAGMA-MIC (ICL)
 - Nvidia GPU: CUBLAS, MAGMA (ICL)
- Parallel libraries: PETSc, Trilinos
- PDEs: deal.II, FEniCS, libMesh



blaze-lib: CG Example

```

100  const size_t NN( N*N );
101
102  blaze::CompressedMatrix<double,rowMajor> A( NN, NN );
103  blaze::DynamicVector<double,columnVector> x( NN, 1.0 ), b( NN, 0.0 ),
104                                             r( NN ), p( NN ), Ap( NN );
105  double alpha, beta, delta;
106
107  // ... Initializing the sparse matrix A
108
109  // Performing the CG algorithm
110  r = b - A * x;
111  p = r;
112  delta = (r,r);
113
114  for( size_t iteration=0UL; iteration<iterations; ++iteration )
115  {
116      Ap = A * p;
117      alpha = delta / (p,Ap);
118      x += alpha * p;
119      r -= alpha * Ap;
120      beta = (r,r);
121      if( std::sqrt( beta ) < 1E-8 ) break;
122      p = r + ( beta / delta ) * p;
123      delta = beta;
124  }

```



PETSc

- Library for large-scale scientific computation
- Large collection of parallel functions for linear solvers, nonlinear solvers, ODE integrators
- Abstracts communication details from user; focus on solving the problem
- Well documented; large collection of online examples/tutorials



PETSc: Parallel Linear Solve Example

```

1  #include <petscksp.h>
2  PETSC_EXTERN PetscErrorCode PCCreate_Jacobi(PC);
3
4  int main(int argc, char **args)
5  {
6      Vec          x,b,u; /* approx solution, RHS, exact solution */
7      Mat          A;     /* linear system matrix */
8      KSP          ksp;   /* linear solver context */
9      PetscReal    norm;  /* norm of solution error */
10     PetscInt     i,j,Ii,J,Istart,Iend,m = 8,n = 7,its;
11     PetscScalar  v,one = 1.0,neg_one = -1.0;
12     PC           pc;     /* preconditioner context */
13
14     PetscInitialize(&argc,&args,(char*)0,help);
15     PetscOptionsGetInt(NULL,"-m",&m,NULL);
16     PetscOptionsGetInt(NULL,"-n",&n,NULL);
17
18     MatCreate(PETSC_COMM_WORLD,&A);
19     MatSetSizes(A,PETSC_DECIDE,PETSC_DECIDE,m*n,m*n);
20     MatSetFromOptions(A);
21     MatSetUp(A);

```



PETSc: Parallel Linear Solve Example

```

22     MatGetOwnershipRange(A,&Istart,&Iend);
23
24     for (Ii=Istart; Ii<Iend; Ii++) {
25         v = -1.0; i = Ii/n; j = Ii - i*n;
26         if (i>0)    {J = Ii - n; MatSetValues(A,1,&Ii,1,&J,&v,INSERT_VALUES)}
27         if (i<m-1) {J = Ii + n; MatSetValues(A,1,&Ii,1,&J,&v,INSERT_VALUES)}
28         if (j>0)    {J = Ii - 1; MatSetValues(A,1,&Ii,1,&J,&v,INSERT_VALUES)}
29         if (j<n-1) {J = Ii + 1; MatSetValues(A,1,&Ii,1,&J,&v,INSERT_VALUES)}
30         v = 4.0; MatSetValues(A,1,&Ii,1,&Ii,&v,INSERT_VALUES);
31     }
32
33     MatAssemblyBegin(A,MAT_FINAL_ASSEMBLY);
34     MatAssemblyEnd(A,MAT_FINAL_ASSEMBLY);
35
36     VecCreate(PETSC_COMM_WORLD,&u);
37     VecSetSizes(u,PETSC_DECIDE,m*n);
38     VecSetFromOptions(u);
39     VecDuplicate(u,&b);
40     VecDuplicate(b,&x);
41     VecSet(u,one);
42     MatMult(A,u,b);

```



PETSc: Parallel Linear Solve Example

```

43 KSPCreate(PETSC_COMM_WORLD,&ksp);
44 KSPSetOperators(ksp,A,A);
45 PCRegister("ourjacobi",PCCreate_Jacobi);
46 KSPGetPC(ksp,&pc);
47 PCSetType(pc,"ourjacobi");
48 KSPSetFromOptions(ksp);
49
50 KSPSolve(ksp,b,x);
51
52 VecAXPY(x,neg_one,u);
53 VecNorm(x,NORM_2,&norm);
54 KSPGetIterationNumber(ksp,&its);
55 PetscPrintf(PETSC_COMM_WORLD,"Norm of error %g iterations %D\n",
56             (double)norm,its);
57
58 KSPDestroy(&ksp);
59 VecDestroy(&u); VecDestroy(&x);
60 VecDestroy(&b); MatDestroy(&A);
61 PetscFinalize();
62 return 0;
63 }

```



deal.II




Image Credit: Dr. Wolfgang
Bangerth.

- Modern C++ based library for building applications to solve PDEs with finite elements
- Support for arbitrary degree, adaptive refinement, 1/2/3 spatial dimensions
- Interfaces to a variety of libraries: ARPACK, PETSc, Trilinos, SLEPc, MPI, p4est, METIS, ...
- Well documented code; greater than 50 tutorial programs; online collection of video lectures



Optimization

- 1** Introduction
 - Defining High Performance Computing
- 2** Programming Models
 - Nomenclature
 - Models
 - Examples
- 3** Hardware
 - CPUs
 - Memory
- 4** Libraries
 - blaze-lib
 - PETSc
 - deal.II
- 5** Optimization
 - Intel VTune
 - ITAC
 - TAU
- 6** Additional Resources



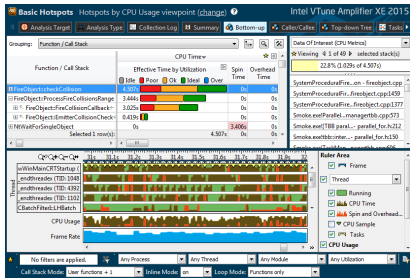
Optimization

- Getting the most performance out of available hardware; being able to scale efficiently to larger resources
- Libraries are typically optimized; application code can be the bottleneck
- “Premature optimization is the root of all evil” – Donald Knuth; measure code performance and look for critical sections
- Several available tools to provide metrics on performance including CPU, cache utilization, memory bandwidth, communication, ...



Intel VTune

Intel VTune



- Provides CPU metrics, cache misses, thread synchronization information, ...
- Identifies which functions use the most CPU time

Image Credit: Intel.



Intel Trace Analyzer and Collector




Image Credit: Intel.

- Collects and reports on MPI communication patterns
- Aids in finding bottlenecks and load balancing issues

Tuning and Analysis Utilities

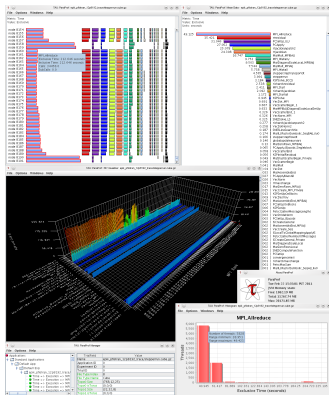


Image Credit: PRACE.

- Open source resource providing similar information as VTune and ITAC, but with a steeper learning curve
- Accesses hardware counters to provide hardware metrics; can instrument MPI calls to trace communication patterns



Additional Resources

- 1** Introduction
 - Defining High Performance Computing
- 2** Programming Models
 - Nomenclature
 - Models
 - Examples
- 3** Hardware
 - CPUs
 - Memory
- 4** Libraries
 - blaze-lib
 - PETSc
 - deal.II
- 5** Optimization
 - Intel VTune
 - ITAC
 - TAU
- 6** Additional Resources



Additional Resources

- **Software Carpentry** Lessons on shell, source control, Python, R, SQL
<http://software-carpentry.org/>
- **HPC Beginner's Guide** More in-depth introduction
<http://tinyurl.com/korh48z>
- **LLNL Training** Great collection of tutorials and presentations including: MPI, OpenMP, TAU, Python
<http://tinyurl.com/3zxaw6>
- **deal.II** Video lectures introducing deal.II usage
<http://www.math.tamu.edu/~bangerth/videos.html>



Additional Resources

- **NICS HPC Seminar** videos, slides, and on campus at Claxton; introduces HPC basics.
<https://www.nics.tennessee.edu/hpc-seminar-series>
- **MOOCs**
 - 1 High Performance Scientific Computing – Dr. Randall LeVeque; covers OpenMP, MPI, Python, Fortran. Starts this Friday!
<https://www.coursera.org/course/scicomp>
 - 2 Heterogeneous Parallel Computing – Dr. Wen-mei Hwu; covers common parallel algorithm patterns with CUDA
<https://www.coursera.org/course/hetero>



Contacts at JICS

- Dr. Glenn Brook – Director of AACE and PI of Beacon Project
glenn-brook@tennessee.edu
- Dr. Lonnie Crosby – Scientific Computing Group Lead and XSEDE ERST
lcrosby1@utk.edu
- Dr. Ryan Glasby – Computational Engineer and CFD Group Lead
ryan-glasby@tennessee.edu



Thank You

Shane Sawyer

shane-sawyer@tennessee.edu

