

PetaBricks and Julia

Kathleen C. Alexander

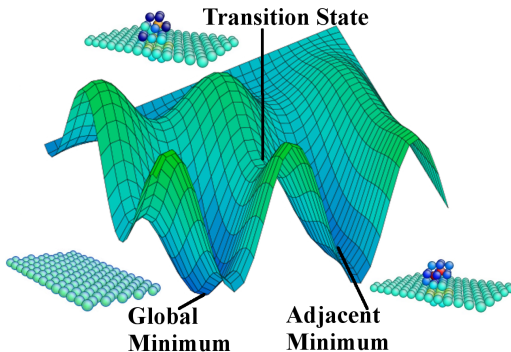
Massachusetts Institute of Technology

December 11th, 2013

Motivation

The Programmer's Dilemma

a personal example— energy landscapes



The Programmer's Dilemma

which algorithm is best?

```
1 #ifndef EIGENFINDHEADERDEF
2 #define EIGENFINDHEADERDEF
3
4 #include <iostream>
5 #include <cmath>
6 #include "Vector.hpp"
7 #include "Matrix.hpp"
8 #include "BLAS.hpp"
9
10 #include "Aton.hpp"
11 #include "RW.hpp"
12 #include "RunLammps.hpp"
13
14 using namespace std;
15
16 double LanczosNew(Vector &eigenvector, int lanczos_number,
17 vector<Atom> atoms_old, ofstream &outfile, char* filename,
18 char* newfile, int mynode, char* hess_lammps, char* hess_data,
19 int dim, int local_count);
20 double LanczosHess(Vector &eigenvector, Matrix &hess, int lanczos_number);
21 void getForces(vector<Atom> &atoms_old, Vector r, Vector &forces_new,
22 ofstream &outfile, char* filename, char* newfile, int mynode, char* hess_lammps,
23 char* hess_data, int dim);
24 void getForces2ndOrder(vector<Atom> &atoms_old, Vector r, Vector &forces_new,
25 ofstream &outfile, char* filename, char* newfile, int mynode, char* hess_lammps,
26 char* hess_data, int dim);
27 void getForces(vector<Atom> &atoms_new, Vector &forces_new,
28 ofstream &outfile, char* filename, char* newfile, int mynode, char* hess_lammps,
29 char* hess_data, int dim);
30 double TDISPowerMethod(int size, Vector &a, Vector &b, Vector &q, int freq, float shift);
31 double QRSolver(Vector &a, Vector&b, Vector &y, int size);
32 void QRdecomp(Matrix A, Matrix &Q, Matrix &R, int size);
33 double Householder(Vector &w, Matrix &A, int index, int size);
```

The Programmer's Dilemma

which algorithm is best?

```
1 #ifndef EIGENFINDHEADERDEF
2 #define EIGENFINDHEADERDEF
3
4 #include <iostream>
5 #include <cmath>
6 #include "Vector.hpp"
7 #include "Matrix.hpp"
8 #include "BLAS.hpp"
9
10 #include "Atom.hpp"
11 #include "RW.hpp"
12 #include "RunLammps.hpp"
13
14 using namespace std;
15
16 double LanczosNew(Vector &eigenvector, int lanczos_number,
17                 vector<Atom> atoms_old, ofstream &outfile, char* filename,
18                 char* newfile, int mynode, char* hess_lammps, char* hess_data,
19                 int dim, int local_count);
20 double LanczosHess(Vector &eigenvector, Matrix &hess, int lanczos_number);
21 void getForces(vector<Atom> &atoms_old, Vector r, Vector &forces_new,
22              ofstream &outfile, char* filename, char* newfile, int mynode, char* hess_lammps,
23              char* hess_data, int dim);
24 void getForces2ndOrder(vector<Atom> &atoms_old, Vector r, Vector &forces_new,
25                       ofstream &outfile, char* filename, char* newfile, int mynode, char* hess_lammps,
26                       char* hess_data, int dim);
27 void getForces(vector<Atom> &atoms_new, Vector &forces_new,
28               ofstream &outfile, char* filename, char* newfile, int mynode, char* hess_lammps,
29               char* hess_data, int dim);
30 double TDISPowerMethod(int size, Vector &a, Vector &b, Vector &q, int freq, float shift);
31 double QRSolver(Vector &a, Vector&b, Vector &y, int size);
32 void QRdecomp(Matrix A, Matrix &Q, Matrix &R, int size);
33 double Householder(Vector &w, Matrix &A, int index, int size);
```

Goal: determine the best algorithm for the application—
which may be machine dependent

Parallel Programming

```

1 #ifndef EIGENFINDHEADERDEF
2 #define EIGENFINDHEADERDEF
3
4 #include <iostream>
5 #include <cmath>
6 #include "Vector.hpp"
7 #include "Matrix.hpp"
8 #include "BLAS.hpp"
9
10 #include "Atom.hpp"
11 #include "RW.hpp"
12 #include "RunLammps.hpp"
13
14 using namespace std;
15
16 double LanczosNew(Vector &eigenvector, int lanczos_number,
17                 vector<Atom> atoms_old, ofstream &outfile, char* filename,
18                 char* newfile, int nynode, char* hess_lammps, char* hess_data,
19                 int dim, int local_count);
20 double LanczosHess(Vector &eigenvector, Matrix &hess, int lanczos_number);
21 void getForces(vector<Atom> &atoms_old, Vector r, Vector &forces_new,
22              ofstream &outfile, char* filename, char* newfile, int nynode, char* hess_lammps,
23              char* hess_data, int dim);
24 void getForces2ndOrder(vector<Atom> &atoms_old, Vector r, Vector &forces_new,
25                       ofstream &outfile, char* filename, char* newfile, int nynode, char* hess_lammps,
26                       char* hess_data, int dim);
27 void getForces(vector<Atom> &atoms_new, Vector &forces_new,
28               ofstream &outfile, char* filename, char* newfile, int nynode, char* hess_lammps,
29               char* hess_data, int dim);
30 double TDISPowerMethod(int size, Vector &a, Vector &b, Vector &Q, int freq, float shift);
31 double QRSolver(vector &a, Vector &b, Vector &y, int size);
32 void QRdecomp(Matrix A, Matrix &Q, Matrix &R, int size);
33 double Householder(Vector &w, Matrix &A, int index, int size);

```

- many parts of these algorithms can be written in parallel
- often they can be parallelized in many different ways
- optimizing these options is a challenge

Determine the best way to parallelize the program—
which will be machine dependent

Parallel Programming

```
1 #ifndef EIGENFINDHEADERDEF
2 #define EIGENFINDHEADERDEF
3
4 #include <iostream>
5 #include <cmath>
6 #include "Vector.hpp"
7 #include "Matrix.hpp"
8 #include "BLAS.hpp"
9
10 #include "Atom.hpp"
11 #include "RW.hpp"
12 #include "RunLammps.hpp"
13
14 using namespace std;
15
16 double LanczosNew(Vector &eigenvector, int lanczos_number,
17                 vector<Atom> atoms_old, ofstream &outfile, char* filename,
18                 char* newfile, int nynode, char* hess_lammps, char* hess_data,
19                 int dim, int local_count);
20 double LanczosHess(Vector &eigenvector, Matrix &hess, int lanczos_number);
21 void getForces(vector<Atom> &atoms_old, Vector r, Vector &forces_new,
22              ofstream &outfile, char* filename, char* newfile, int nynode, char* hess_lammps,
23              char* hess_data, int dim);
24 void getForces2ndOrder(vector<Atom> &atoms_old, Vector r, Vector &forces_new,
25                       ofstream &outfile, char* filename, char* newfile, int nynode, char* hess_lammps,
26                       char* hess_data, int dim);
27 void getForces(vector<Atom> &atoms_new, Vector &forces_new,
28               ofstream &outfile, char* filename, char* newfile, int nynode, char* hess_lammps,
29               char* hess_data, int dim);
30 double TDSPowerMethod(int size, Vector &a, Vector &b, Vector &Q, int freq, float shift);
31 double QRSolver(Vector &a, Vector &b, Vector &y, int size);
32 void QRdecomp(Matrix A, Matrix &Q, Matrix &R, int size);
33 double Householder(Vector &w, Matrix &A, int index, int size);
```

- many parts of these algorithms can be written in parallel
- often they can be parallelized in many different ways
- optimizing these options is a challenge

Determine the best way to parallelize the program—
which will be machine dependent

Background

Petabricks – Algorithmic Choice

PetaBricks was developed to alleviate some of the optimization responsibility from the programmer

the transform

```
1 #ifndef SORT_PBCC
2 #define SORT_PBCC
3
4 #define SORTSUBARRAY SortSubArray
5 #include "Quicksort.pbcc"
6 #include "Insertionsort.pbcc"
7 #include "Radixsort.pbcc"
8 #include "Parallel_Mergesort.pbcc"
9 #include "Mergesort.pbcc"
10 #include "Selectionsort.pbcc"
11
12 transform SortSubArray
13 from IN[n], Pos
14 to OUT[n], TEMP[n]
15 {
16 //to (OUT out, TEMP temp) from (IN in, Pos p)
17 //{
18 // Parallel_MergesortSubArray( out, temp, in, p);
19 //}
20
21 rule MergeSort2
22 to (OUT out, TEMP temp) from (IN in, Pos p)
23 {
24   MergesortSubArray<1>(out, temp, in, p);
25 }
26
27 rule MergeSort4
28 to (OUT out, TEMP temp) from (IN in, Pos p)
29 {
30   MergesortSubArray<2>(out, temp, in, p);
31 }
32
33 rule MergeSort8
34 to (OUT out, TEMP temp) from (IN in, Pos p)
35 {
```

Petabricks – Algorithmic Choice

PetaBricks was developed to alleviate some of the optimization responsibility from the programmer

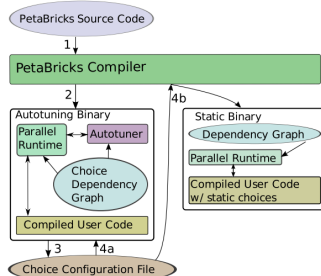
the transform

```

1 #ifndef SORT_PBCC
2 #define SORT_PBCC
3
4 #define SORTSUBARRAY SortSubArray
5 #include "Quicksort.pbcc"
6 #include "Insertionsort.pbcc"
7 #include "Radixsort.pbcc"
8 #include "Parallel_Mergesort.pbcc"
9 #include "Mergesort.pbcc"
10 #include "Selectionsort.pbcc"
11
12 transform SortSubArray
13 from IN[n], Pos
14 to OUT[n], TEMP[n]
15 {
16 //to (OUT out, TEMP temp) from (IN in, Pos p)
17 //{
18 // Parallel_MergesortSubArray( out, temp, in, p);
19 //}
20
21 rule MergeSort2
22 to (OUT out, TEMP temp) from (IN in, Pos p)
23 {
24   MergesortSubArray<1>(out, temp, in, p);
25 }
26
27 rule MergeSort4
28 to (OUT out, TEMP temp) from (IN in, Pos p)
29 {
30   MergesortSubArray<2>(out, temp, in, p);
31 }
32
33 rule MergeSort8
34 to (OUT out, TEMP temp) from (IN in, Pos p)
35 {

```

compiling framework

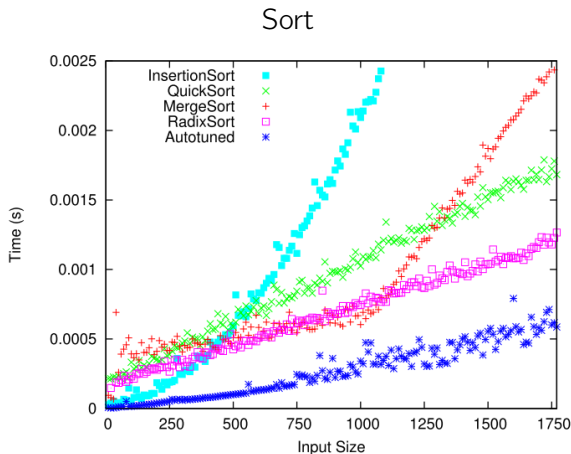


Ansel, et al. ACM SIGPLAN Conference (2009).

Petabricks – Autotuning

the autotuner determines the best configuration for the machine under the tuning constraints


Petabricks – Autotuning



Ansel, et al. ACM SIGPLAN Conference (2009).

Petabricks – Autotuning


Eigen Problem



Ansel, et al. ACM SIGPLAN Conference (2009).

Petabricks – Autotuning

Matrix Multiply

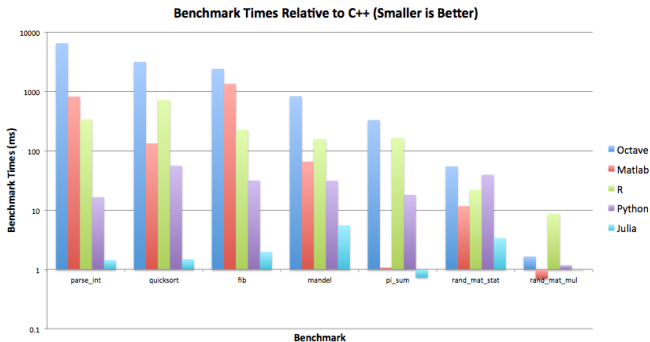


Ansel, et al. ACM SIGPLAN Conference (2009).

Julia


- Julia was developed to bridge the gap between interpreted and compiled scientific computing
- streamlining parallelization techniques has been a priority

Julia



<http://forio.com/julia/julia>

Julia



<http://forio.com/julia/julia>

Question: is there room for overlap between the PetaBricks and Julia approaches?

Approach

Options for Implementation

Julia in PetaBricks

- can utilize PetaBricks autotuner and compiler
- PetaBricks compiler needs to interpret Julia

Options for Implementation

Julia in PetaBricks

- can utilize PetaBricks autotuner and compiler
- PetaBricks compiler needs to interpret Julia

PetaBricks in Julia

- can run PetaBricks binaries inside Julia
- no PetaBricks shared object files, functions require disk i/o
- doesn't take advantage of JuliaLang

Options for Implementation

Julia in PetaBricks

- can utilize PetaBricks autotuner and compiler
- PetaBricks compiler needs to interpret Julia

PetaBricks in Julia

- can run PetaBricks binaries inside Julia
- no PetaBricks shared object files, functions require disk i/o
- doesn't take advantage of JuliaLang

Julia + OpenTuner

- apply PetaBricks framework to Julia
- utilize OpenTuner to optimize Julia

Approach Used Here

PetaBricks in Julia

- can run PetaBricks binaries inside Julia
- no PetaBricks shared object files, functions require disk i/o
- doesn't take advantage of JuliaLang

Approach Used Here

PetaBricks in Julia

- can run PetaBricks binaries inside Julia
- no PetaBricks shared object files, functions require disk i/o
- doesn't take advantage of JuliaLang

⇒ most naive approach possible:

→ compile PetaBricks executable, exe

→ `julia > run('$exe $in $out')`

Approach Used Here

PetaBricks in Julia

- can run PetaBricks binaries inside Julia
- no PetaBricks shared object files, functions require disk i/o
- doesn't take advantage of JuliaLang

⇒ most naive approach possible:

→ compile PetaBricks executable, exe

→ `julia > run('$exe $in $out')`

⇒ compare with PetaBricks and Julia alone

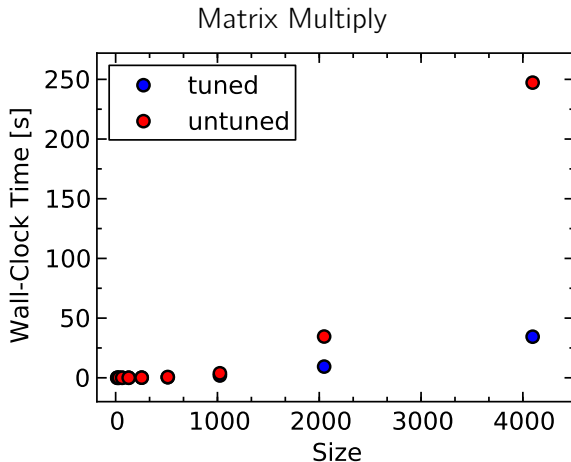
→ lower bound of performance improvement

→ is there proof of benefit?

Results

PetaBricks- Tuning Improvements

performance improvement— tuned and untuned PetaBricks



Comparing PetaBricks with Julia - Apples to Apples

PetaBricks

- functions read in ASCII files and output same
- determines parallelization during autotuning
- autotuning can take days

Julia

- JIT for each independent execution
- can addprocs(n), but may not parallelize
- can be used interactively

Comparing PetaBricks with Julia - Apples to Apples

PetaBricks

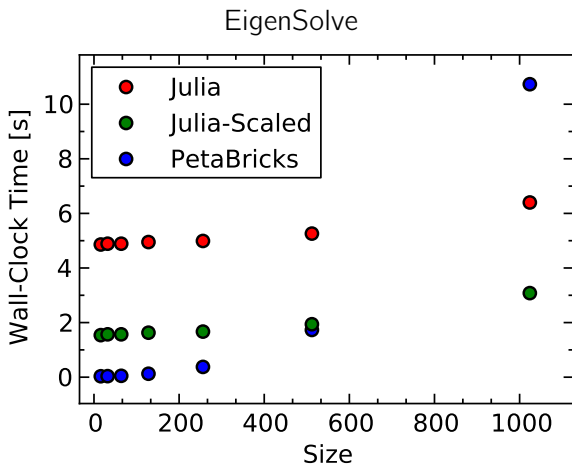
- functions read in ASCII files and output same
- determines parallelization during autotuning
- autotuning can take days

PetaBricks

- JIT for each independent executable
- can addprocs(n), but may not parallelize
- can be used interactively

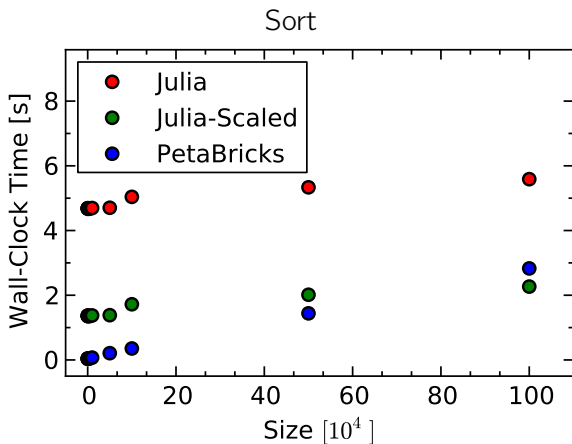
- make both programs do i/o
- run both programs from shell
- try addprocs(n) in Julia, with no other instructions
- subtract 'hello world' start-up time from Julia wall-clock

Comparing PetaBricks to Julia - EigenSolve



- Julia seems to do the best for large matrices
- however, the results were not comparable
- this test was not a good apples-to-apples performance test


Comparing PetaBricks with Julia - Sort




- Julia and PetaBricks converge for large vectors
- PetaBricks is better with shorter vectors
- effect of i/o not considered wrt performance

Comparing PetaBricks with Julia Matrix Multiply

i5-3339 (4 CPU)

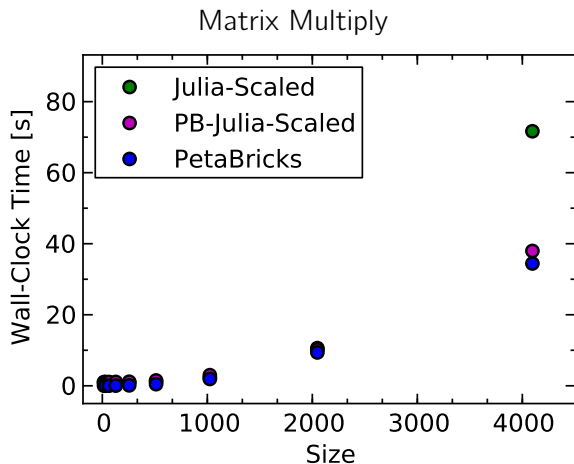


i7-3770 (8 CPU)



- Julia and PetBricks converge moderate matrix sizes on fewer cores
- PetaBricks is better with smaller lists and larger matrices
- using `addprocs(n)` with no other instruction does not utilize parallel functionality in Julia

Running PetaBricks from Julia

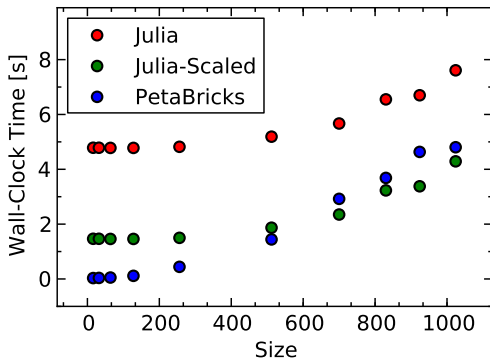


- Can get PetaBricks improvement by incorporating PetaBricks executable in Julia
- effect of i/o not considered wrt performance

Recommendations

Recommendations


Matrix Multiply



→ under many circumstances, Julia performs as well as PetaBricks without days of compilation

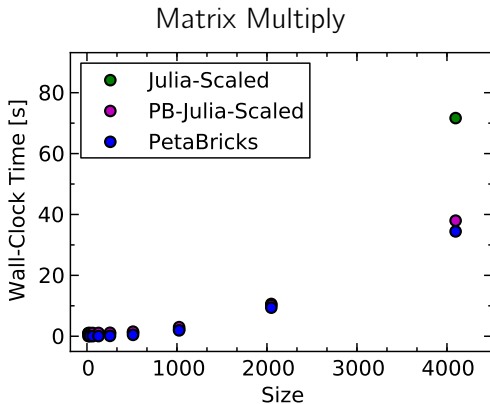
Recommendations

Matrix Multiply




→ there is room for improvement on the start-up time for Julia

Recommendations



→ PetaBricks performance
can be achieved by using
a shell command in Julia

Recommendations



Ansel, et. al. MIT CSAIL Technical Report
MIT-CSAIL-TR-2013-026 (2013).

→ implementing OpenTuner (when better documentation is available) with Julia may be a reasonable long term goal for performance gains of this kind

