

# Petabricks: A language and compiler for algorithmic choice

J. Ansel et al. ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI, 2009.

# Motivation – Many algorithms with various trade offs

## Sorting

MergeSort	QuickSort	InsertionSort	RadixSort
Fast for medium size input	Fastest for medium size input	Fastest for small input	Fastest for largest input
Highly parallelizable	Exploit spatial locality	$O(1)$ Memory	$O(wn)$ runtime

- But gets more complicated!
- Mixing the algorithms yield better results
- e.g QuickSort then cut off to InsertionSort once list is small enough
  - Requires knowing an optimal cut-off point!
  - Differs from an architecture to an architecture
- Choices are left to the developer, complex, time-consuming, error prone

## Hence the need for Autotuning

# PetaBricks – A language and a compiler

- Self-tuning compiler for bespoke architecture
- A language that allow expressing choice in algorithms
- Implicitly parallelizable
- Auto-select the desired trade-off between accuracy and performance

# Sample Code - Sorting

```
1  #define SORT Sort
2
3  // #include "Bitonicsort.pbcc"
4  #include "Insertionsort.pbcc"
5  #include "Mergesort2.pbcc"
6  #include "Quicksort2.pbcc"
7  #include "Selectionsort.pbcc"
8
9  function Sort
10 from in[n]
11 to out[n]
12 {
13     Mergesort2(out, in);
14 } or {
15     Mergesort4(out, in);
16 } or {
17     Quicksort2(out, in);
18 } or {
19     Insertionsort(out, in);
20 } or {
21     Selectionsort(out, in);
22 }
23 // or {
24 //     Bitonic(out, in);
25 // }
26
```

- Define multiple functions
- Functions have tunable variables
- PetaBricks choose optimal combination

# PetaBricks effect

- Source code is compiled into a binary
- The binary is auto-tuned on the architecture of the system
- The compiler produces a final binary contains the optimal configuration


# Tuned variables - Sample

Branch: master ▾

[petabricks](#) / [examples](#) / [sort](#) / [Sort.cfg.X1](#)

Find file

Copy path

 **Marek Olszewski** New cross tuning configs

6eafee1 on Mar 30, 2009

0 contributors

52 lines (51 sloc) | 3.31 KB

Raw

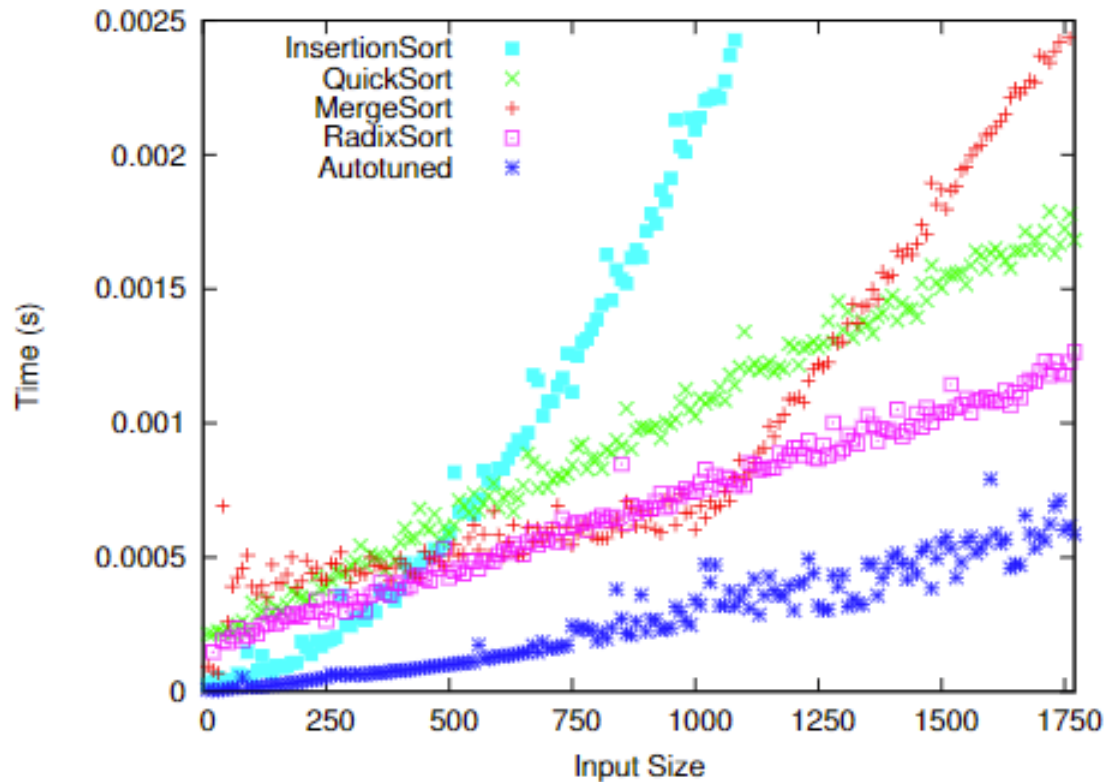
Blame

History



```
1 Copy1D_sequential_cutoff = 1455 # valid range: 0 to 2147483647
2 Copy1D_split_size = 1714 # valid range: 1 to 2147483647
3 Copy2D_sequential_cutoff = 568 # valid range: 0 to 2147483647
4 Copy2D_split_size = 910 # valid range: 1 to 2147483647
5 InsertionsortSubArray_sequential_cutoff = 0 # valid range: 0 to 2147483647
6 Insertionsort_sequential_cutoff = 0 # valid range: 0 to 2147483647
7 Merge16_sequential_cutoff = 1455 # valid range: 0 to 2147483647
8 Merge2_Parallel_Cutoff = 949 # valid range: 100 to 100000
9 Merge2_Parallel_sequential_cutoff = 1373 # valid range: 0 to 2147483647
10 Merge4_sequential_cutoff = 1334 # valid range: 0 to 2147483647
11 Merge8_sequential_cutoff = 569 # valid range: 0 to 2147483647
12 Mergesort16_sequential_cutoff = 1497 # valid range: 0 to 2147483647
13 Mergesort4_sequential_cutoff = 1497 # valid range: 0 to 2147483647
14 Mergesort8_sequential_cutoff = 1458 # valid range: 0 to 2147483647
15 MergesortSubArray16_sequential_cutoff = 1455 # valid range: 0 to 2147483647
16 MergesortSubArray4_sequential_cutoff = 1497 # valid range: 0 to 2147483647
17 MergesortSubArray8_sequential_cutoff = 1497 # valid range: 0 to 2147483647
18 Parallel_MergesortSubArray_sequential_cutoff = 945 # valid range: 0 to 2147483647
19 Parallel_Mergesort_sequential_cutoff = 1497 # valid range: 0 to 2147483647
20 QuickSort_sequential_cutoff = 1497 # valid range: 0 to 2147483647
21 QuicksortSubArray_sequential_cutoff = 949 # valid range: 0 to 2147483647
22 RadixsortSubArray_sequential_cutoff = 1714 # valid range: 0 to 2147483647
```

# Sorting - Results



**Figure 14.** Performance for sort on 8 cores.

Figure taken from paper in review

# PetaBricks Components Graph

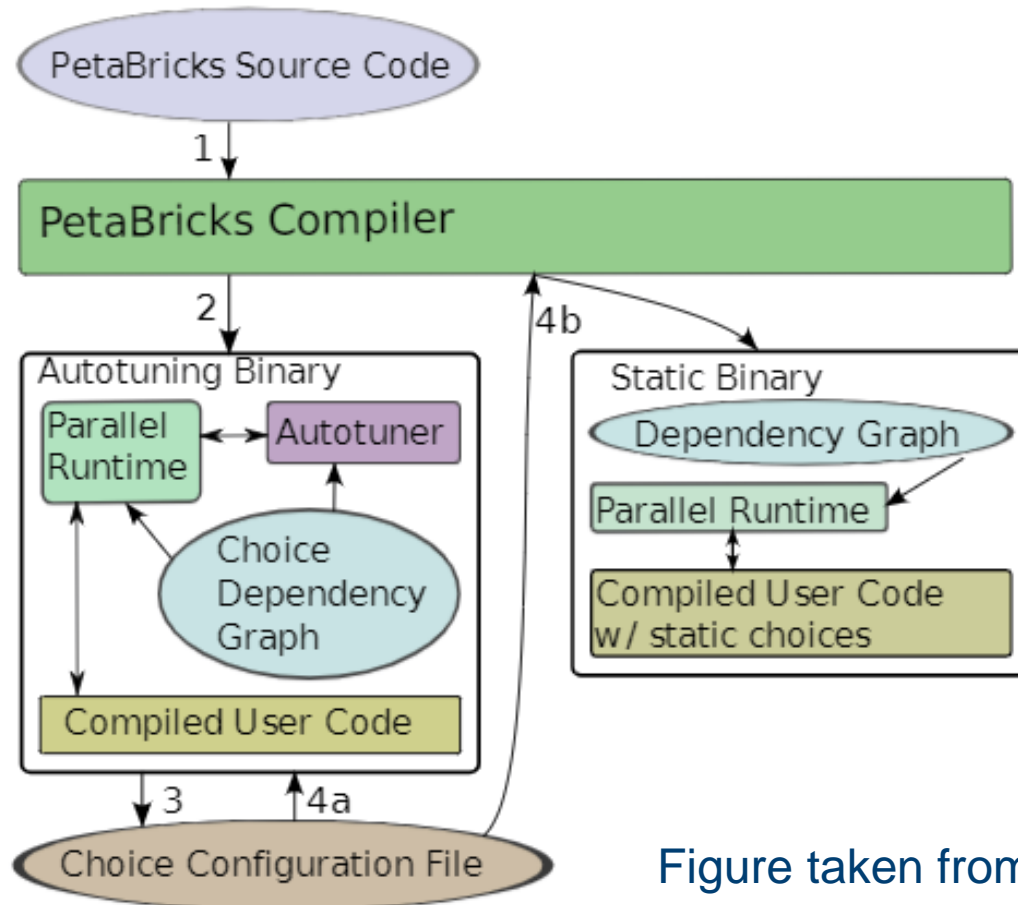


Figure taken from paper in review



# PetaBricks Internals – Source2Source Compiler

- Compiles from PetaBricks to C++
- Input parsed into syntax tree
- Construct a choice grid for matrix type
- Build a choice dependency graph

# Compilation – Rolling Sum example

```
transform RollingSum
from A[n]
to B[n]
{
  //rule0: sum all elements to the left
  to(B.cell(i) b) from(A.region(0, i) in) {
    b=sum(in);
  }

  //rule1: use the previously computed value
  to(B.cell(i) b) from(A.cell(i) a,
                        B.cell(i-1) leftSum) {
    b=a+leftSum;
  }
}
```

Figures taken from paper in review

# Compilation – Applicable Region

```
transform RollingSum
from A[n]
to B[n]
{
  //rule0: sum all elements to the left
  to(B.cell(i) b) from(A.region(0, i) in) {
    b=sum(in);
  }

  //rule1: use the previously computed value
  to(B.cell(i) b) from(A.cell(i) a,
    B.cell(i-1) leftSum) {
    b=a+leftSum;
  }
}
```

Figures taken from paper in review

# Compilation – Applicable Region

```
transform RollingSum
from A[n]
to B[n]
{
  //rule0: sum all elements to the left
  to(B.cell(i) b) from(A.region(0, i) in) {
    b=sum(in);
  }

  //rule1: use the previously computed value
  to(B.cell(i) b) from(A.cell(i) a,
                    B.cell(i-1) leftSum) {
    b=a+leftSum;
  }
}
```

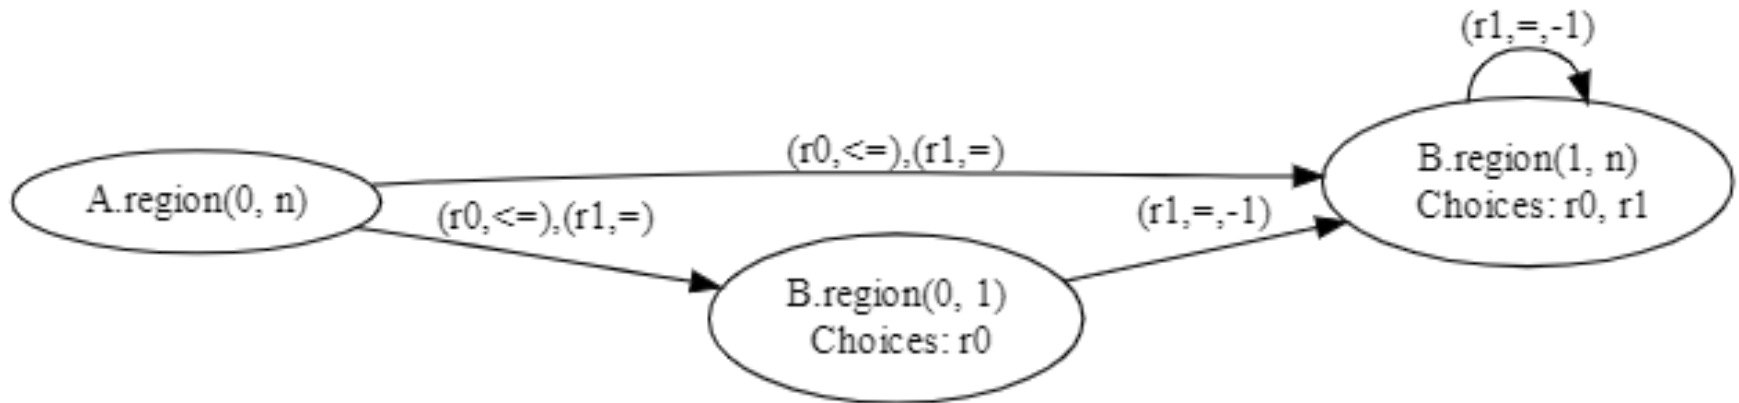
Figures taken from paper in review

# Compilation - Choice Grid Analysis

- Split the data into matrix (grids)
- Map data to rules
- E.g.  $[0, 1) = \{\text{rule } 0\}$  only  
while  $[1, n) = \{\text{rule } 0 \text{ or } 1\}$
- Rule priority is applied here as well

Figures taken from paper in review

# Compilation - Choice Dependency Graph



Figures taken from paper in review

# Code Generation

- Two modes
  - Default – choices and autotuner information are embedded in the output code
  - Second mode – code generation with all choices eliminated based on autotuner results
  - Second mode is useful to produce an intermediate code for C++ to compile – it is more efficient when choices are eliminated

# Auto Tuning System

- Tuning is done by running search on the available configurations
- The available configuration is described using the choice dependency graph
- Using bottom-up approach, works on smaller input and works it way up to large input



# Runtime library

- Dynamically schedule tasks to distribute workload
- When task reach tunable cut-off point they stop calling the scheduler and execute sequentially
- Maximize locality using Cilk, task stealing protocol;
  - Thread operates on top of its dequeue
  - When it runs out of tasks
  - Select a random victim to steal work from bottom of their dequeue

# PetaBricks – Other features

- Calling external libraries and other languages
- Template Transformation – Similar to C++
- Rule priorities and where clauses to manual tune edge cases
- Deadlock and race conditions prevention using the dependency graph
- Automated Consistency Checking - advantage of choices, you can run multiple versions and check their results for consistency

# Evaluation - Performance






Figure 14. Performance for sort on 8 cores.

# Evaluation – Performance other algorithms




**Figure 11.** Performance for algorithms to solve Poisson's equation up to an accuracy of  $10^9$  using 8 cores. The iterated SOR algorithm uses the corresponding optimal weight  $\omega_{opt}$  for each of the different input sizes



**Figure 12.** Performance for Eigenproblem on 8 cores. "Cutoff 25" corresponds to the hard-coded hybrid algorithm found in LAPACK.

# Evaluation – Parallelism



**Figure 16. Parallel scalability.** Speedup as more worker threads are added. Run on an 8-way (2 processor  $\times$  4 core) x86\_64 Intel Xeon System.


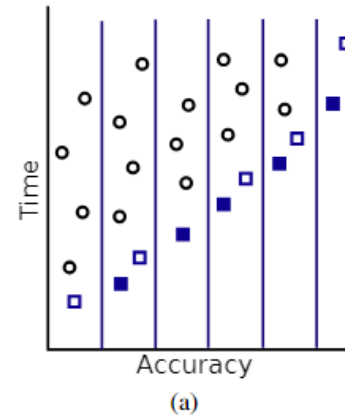
# Evaluation – Accuracy for performance

POISSON<sub>*i*</sub>(*x*, *b*)

- 1: **either**
- 2:   Solve directly
- 3:   Iterate using SOR <sub>$\omega_{opt}$</sub>  until accuracy  $p_i$  is achieved
- 4:   For some  $j$ , iterate with MULTIGRID <sub>$j$</sub>  until accuracy  $p_i$  is achieved
- 5: **end either**

MULTIGRID<sub>*i*</sub>(*x*, *b*)

- 1: **if**  $N = 3$  **then**
- 2:   Solve directly
- 3: **else**
- 4:   Compute one iteration of SOR<sub>1.15</sub>
- 5:   Compute the residual and restrict to half resolution
- 6:   On the coarser grid, call POISSON<sub>*i*</sub>
- 7:   Interpolate result and add correction term to current solution
- 8:   Compute one iteration of SOR<sub>1.15</sub>
- 9: **end if**



**Figure 10.** Pseudo code for family of functions POISSON<sub>*i*</sub> and MULTIGRID<sub>*i*</sub> where *i* is the required accuracy, as used in the benchmark.

# PetaBricks Today

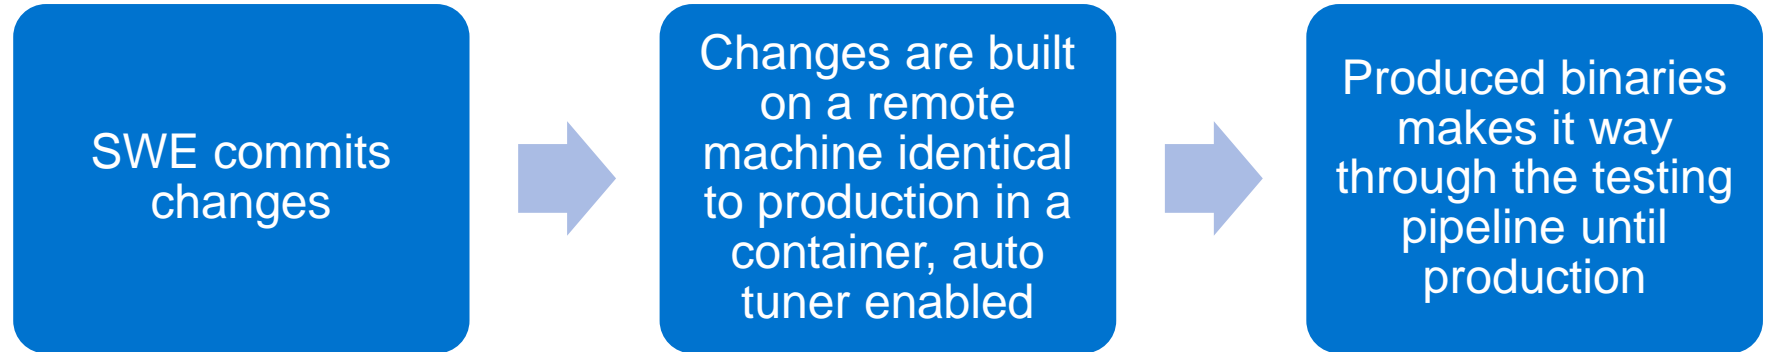
- 363 citations according to [Google Scholar](#)
- Experiments and benchmarked continued for 4 years looking into tuning variables in algorithms, portability, and study the trade-offs between accuracy and performance.
- [Main author](#)(during his PhD), Jason Ansel, Director of engineering at GoDaddy since 2013
- [GitHub repo](#), abandoned 6 years ago
- Ideas we can still use, with help of ML
  - Auto tuning between accuracy and performance
  - Auto tuning variables based on architecture

# Critique

<b>The Good</b>	<b>The Bad</b>
<b>First language that explores algorithmic choice with impressive results</b>	<b>Slow compiling time – impractical for software engineering tasks (target success metric)</b>
<b>Ease of selecting the trade-off between accuracy and performance</b>	<b>Complex code structure – harder to adhoc debug</b>
<b>Portability across architectures and future proof</b>	<b>Principles of auto tuning using simple parameter search might be too slow</b>



# Comment – Working in a heterogenous environment



- Optimisation could be running a different algorithm choice on a different machine, storing the metadata of binaries in a key-value store, and binaries in a replicated store
- Not every line of code has to be written in PetaBrick, just the lines that require high performance – e.g. a simple microservice that has no complex logic wouldn't benefit from PetaBrick optimisation

# Ideas for Future Work – GPUs?

- Extract the idea of choice and analysis per hardware architecture.
- Analyse the benefit of running algorithms on GPUs
- GPUs have different memory constraint than what traditional algorithms were designed for, auto tuning helps!
- Other work in this area already exist using ML for choosing GPU/CPU, but no (afaik) work exist to choose algorithm and tune it

# Questions?

**I HAVE NO IDEA  
WHAT I'M DOING**

A photograph of a dog, likely a Weimaraner, sitting in the driver's seat of a car. The dog is looking out the window with a somewhat confused or uncertain expression. The text "I HAVE NO IDEA WHAT I'M DOING" is overlaid in large, white, bold, sans-serif font across the top half of the image.

**MEMEBASE.COM**