

Language and Compiler Support for Auto-Tuning Variable-Accuracy Algorithms

Jason Ansel, Yee Lok Wong, Cy Chan, Marek Olszewski, Alan Edelman, Saman Amarasinghe
Computer Science and Artificial Intelligence Laboratory
Massachusetts Institute of Technology
Cambridge, MA
{jansel, ylwong, cychan, mareko, edelman, saman}@csail.mit.edu

Abstract—Approximating ideal program outputs is a common technique for solving computationally difficult problems, for adhering to processing or timing constraints, and for performance optimization in situations where perfect precision is not necessary. To this end, programmers often use approximation algorithms, iterative methods, data resampling, and other heuristics. However, programming such *variable accuracy* algorithms presents difficult challenges since the optimal algorithms and parameters may change with different accuracy requirements and usage environments. This problem is further compounded when multiple variable accuracy algorithms are nested together due to the complex way that accuracy requirements can propagate across algorithms and because of the size of the set of allowable compositions. As a result, programmers often deal with this issue in an ad-hoc manner that can sometimes violate sound programming practices such as maintaining library abstractions.

In this paper, we propose language extensions that expose trade-offs between time and accuracy to the compiler. The compiler performs fully automatic compile-time and install-time autotuning and analyses in order to construct optimized algorithms to achieve any given target accuracy. We present novel compiler techniques and a structured genetic tuning algorithm to search the space of candidate algorithms and accuracies in the presence of recursion and sub-calls to other variable accuracy code. These techniques benefit both the library writer, by providing an easy way to describe and search the parameter and algorithmic choice space, and the library user, by allowing high level specification of accuracy requirements which are then met automatically without the need for the user to understand any algorithm-specific parameters. Additionally, we present a new suite of benchmarks, written in our language, to examine the efficacy of our techniques. Our experimental results show that by relaxing accuracy requirements, we can easily obtain performance improvements ranging from 1.1x to orders of magnitude of speedup.

I. INTRODUCTION

Traditionally, language designers and compiler writers have operated under the assumption that programs require a fixed, strongly defined behavior; however, this is not always the case in practice. For certain classes of applications, such as NP-hard problems or problems with tight computation or timing constraints, programmers are often willing to sacrifice some level of accuracy for faster performance. In this paper, we broadly define these types of programs as *variable accuracy* algorithms. There are many different classes of variable accuracy algorithms spanning many different fields.

One class of variable accuracy algorithms are *approximation algorithms* in the area of soft computing [1]. Approximation

algorithms are used to find approximate solutions to computationally difficult tasks with results that have provable quality. For many computationally hard problems, it is possible to find such approximate solutions asymptotically faster than it is to find an optimal solution. A good example of this is BINPACKING. Solving the BINPACKING problem is NP-hard, yet arbitrarily accurate solutions may be found in polynomial time [2]. Like many soft computing problems, BINPACKING has many different approximation algorithms, and the best choice often depends on the level of accuracy desired.

Another class of variable accuracy algorithms are *iterative algorithms* used extensively in the field of applied mathematics. These algorithms iteratively compute approximate values that converge toward an optimal solution. Often, the rate of convergence slows dramatically as one approaches the solution, and in some cases a perfect solution cannot be obtained without an infinite number of iterations [3]. Because of this, many programmers create *convergence criteria* to decide when to stop iterating. However, deciding on a convergence criteria can be difficult when writing modular software because the appropriate criteria may not be known to the programmer ahead of time.

A third class of variable accuracy algorithms are algorithms in the signal processing domain. In this domain, the accuracy of an algorithm can be directly determined from the problem specification. For example, when designing digital signal processing (DSP) filters, the type and order of the filter can be determined directly from the desired sizes of the stop, transition and pass-bands as well as the required filtering tolerance bounds in the stop and pass-bands. When these specifications change, the optimal filter type may also change. Since many options exist, determining the best approach is often difficult, especially if the exact requirements of the system are not known ahead of time.

A key challenge when writing and using variable accuracy code arises from maintaining the abstraction boundary between the *library writer* and the *library user*¹. The library writer understands the algorithm and the various choices and parameters affecting accuracy, but does *not* know the accuracy requirements for each use case. Because special-casing an algorithm for each foreseeable accuracy requirement can be

¹We use the term library loosely to mean any abstraction

extremely tedious and error-prone, library writers often follow the practice of exposing many internal algorithmic parameters to the interface of their library, potentially breaking any useful abstractions in the variable accuracy library. Unfortunately, while the library user knows the application’s requirements, he/she may not know how the exposed implementation-specific parameters and algorithmic choices of the library impact these accuracy requirements.

This practice is exemplified by the `fmincon()` function in Matlab, which attempts to find the minimum of a user-specified nonlinear multivariate function subject to a set of specified constraints. `fmincon()` takes accuracy and optimization options specified by an `options` structure. This structure contains 42 fields that the user can set to specify various options such as which of three algorithms to use, how many iterations to run, and what tolerances to use. Additionally, there are a number of options specific to each of the three algorithms, some of which further affect additional algorithmic choices. For example, the value specified in the `PrecondBandWidth` option used by the `trust-region-reflective` algorithm will indirectly affect both the number of preconditioning iterations performed, as well as the type of factorization algorithm used during the preconditioning phase. This option can have a dramatic effect on the performance and accuracy of the solver.

Yet another challenge arises when a program is constructed by composing multiple variable accuracy modules, or by recursively calling variable accuracy functions. In this type of program, the best top-level performance and accuracy may be obtained by using higher or lower accuracies for intermediate components of the algorithm. For example, at each recursive level of a multigrid solver, it may be possible to solve to a lower accuracy while still meeting the accuracy requirement of the final solution. Manually determining what accuracy level to use at each recursive step can be extremely onerous because of the inter-dependencies between accuracy choices.

In this paper we propose a novel set of language extensions and an accuracy-aware compiler to address the challenges in writing variable accuracy code. With our extensions, accuracy time trade-offs are made visible to the compiler, enabling it to perform empirical autotuning to build optimized algorithms for each accuracy level required by a user. As we will show, these extensions simplify writing variable accuracy code both for the library writer and for the library user.

For the library writer, our compiler automates the otherwise tedious search over both the algorithmic search space and the parameter space to find algorithms with the best performance that meet each required level of accuracy. This is done without forcing the library writer to sacrifice control over how the algorithm operates or how accuracy is achieved. For the library user, our extensions allow the specification of top-level accuracy requirements without the library user needing to understand any parameters and choices that are specific to the implementation details of the algorithm. This helps create a better abstraction barrier that separates the details of the algorithms in a library from the requirements of the library

user.

We have implemented our language extensions in the context of the PetaBricks programming language and compiler [4], [5], [6]. PetaBricks is a programming language that allows algorithmic choice to be expressed at the language level. PetaBricks automatically selects and constructs algorithms optimized for each target architecture using empirical autotuning methods. More background on the PetaBricks language and compiler is provided in Section II.

A. Contributions

We make the following contributions:

- We have introduced a novel programming language that incorporates algorithmic accuracy choices. This includes support for multiple accuracy metrics, which provide a general-purpose way for users to define arbitrary accuracy requirements in any domain.
- We have developed a technique for mapping variable accuracy code so that it can be efficiently autotuned without the search space growing prohibitively large.
- We have implemented a language, compiler, and runtime system to demonstrate the effectiveness of our techniques.
- We have implemented a suite of six benchmarks, representative of commonly used algorithms with variable accuracy requirements.
- We demonstrate the importance of algorithmic choice by showing that for different accuracy levels our autotuner chooses different algorithmic techniques.
- We show that by using variable accuracy one can get a speedup up to four orders of magnitude over using the highest accuracy code.

II. PETABRICKS LANGUAGE BACKGROUND

The PetaBricks language provides a framework for the programmer to describe multiple ways of solving a problem while allowing the autotuner to determine which of those ways is best for the user’s situation [4]. It provides both algorithmic flexibility (multiple algorithmic choices) as well as coarse-grained code generation flexibility (synthesized outer control flow).

At the highest level, the programmer can specify a *transform*, which takes some number of inputs and produces some number of outputs. In this respect, the PetaBricks transform is like a function call in a procedural language. The major difference is that we allow the programmer to specify multiple pathways to convert the inputs to the outputs for each transform. Pathways are specified in a dataflow manner using a number of smaller building blocks called *rules*, which encode both the data dependencies of the rule and C++-like code that converts the rule’s inputs to outputs.

Figure 1 presents an example PetaBricks program, `kmeans`. This `kmeans` program groups the input `Points` into a number of clusters and writes each points cluster to the output `Assignments`. Internally the program uses the intermediate data `Centroids` to keep track of the current center of each cluster. The transform header declares each of these data

```

1 transform kmeans
2 from Points[n,2] // Array of points (each column
3 // stores x and y coordinates)
4 through Centroids[sqrt(n),2]
5 to Assignments[n]
6 {
7 // Rule 1:
8 // One possible initial condition: Random
9 // set of points
10 to(Centroids.column(i) c) from(Points p) {
11 c=p.column(rand(0,n))
12 }
13
14 // Rule 2:
15 // Another initial condition: Centerplus initial
16 // centers (kmeans++)
17 to(Centroids c) from(Points p) {
18 CenterPlus(c, p);
19 }
20
21 // Rule 3:
22 // The kmeans iterative algorithm
23 to(Assignments a) from(Points p, Centroids c) {
24 while (true) {
25 int change;
26 AssignClusters(a, change, p, c, a);
27 if (change==0) return; // Reached fixed point
28 NewClusterLocations(c, p, a);
29 }
30 }
31 }

```

Fig. 1. Pseudocode for kmeans

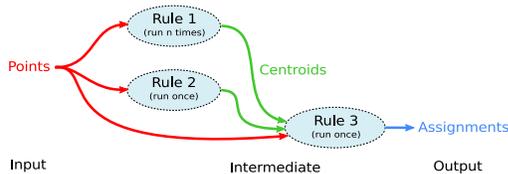


Fig. 2. Dependency graph for kmeans example. The rules are the vertices while each edge represents the dependencies of each rule. Each edge color corresponds to each named data dependence in the pseudocode.

structures as its inputs (Points), outputs (Assignments), and intermediate or “through” data structures (Centroids). The rules contained in the body of the transform define the various pathways to construct the Assignments data from the initial Points data. The transform can be depicted using the dependence graph shown in Figure 2, which indicates the dependencies of each of the three rules.

The first two rules specify two different ways to initialize the Centroids data needed by the iterative kmeans solver in the third rule. Both of these rules require the Points input data. The third rule specifies how to produce the output Assignments using both the input Points and intermediate Centroids. Note that since the third rule depends on the output of either the first or second rule, the third rule will not be executed until the intermediate data structure Centroids has been computed by one of the first two rules. Additionally, the first rule provides an example of how the autotuner can synthesize outer control flow. Instead of explicitly looping over every column of Centroids 2D array, the programmer can specify a computation

```

1 transform kmeans
2 accuracy_metric kmeansaccuracy
3 accuracy_variable k
4 from Points[n,2] // Array of points (each column
5 // stores x and y coordinates)
6 through Centroids[k,2]
7 to Assignments[n]
8 {
9 // (Rules 1 and 2 same as in Figure 1) ...
10
11 // Rule 3:
12 // The kmeans iterative algorithm
13 to(Assignments a) from(Points p, Centroids c) {
14 for_enough {
15 int change;
16 AssignClusters(a, change, p, c, a);
17 if (change==0) return; // Reached fixed point
18 NewClusterLocations(c, p, a);
19 }
20 }
21
22 transform kmeansaccuracy
23 from Assignments[n], Points[n,2]
24 to Accuracy
25 {
26 Accuracy from(Assignments a, Points p){
27 return sqrt(2*n/SumClusterDistanceSquared(a,p));
28 }
29 }
30 }

```

Fig. 3. Pseudocode for variable accuracy kmeans. The new variable accuracy code is highlighted in light blue.

that is done for each column of the output (using the “column” keyword). The order over which these columns are iterated, and the amount of parallelism to use, is then synthesized and tuned by the compiler and autotuner.

To summarize, when our transform is executed, the cluster centroids are initialized either by the first rule, which performs random initialization on a per-column basis with synthesized outer control flow, or the second rule, which calls the CenterPlus algorithm. Once Centroids is generated, the iterative step in the third rule is called.

III. PETABRICKS LANGUAGE EXTENSIONS FOR VARIABLE ACCURACY

At a high level, our language extensions extend the idea of algorithmic choice to include choices between different accuracies. The extensions also allow the user to specify how accuracy should be measured. Our new accuracy-aware autotuner then searches to optimize for both time and accuracy. The result is code that probabilistically meets users’ accuracy needs. Optionally, users can request hard guarantees that utilize runtime checking of accuracy.

Figure 3 presents our kmeans example with our new variable accuracy extensions. The updates to the code are highlighted in light blue. The example uses three of our new variable accuracy features.

First the `accuracy_metric`, on line 2, defines an additional transform, `kmeansaccuracy`, which computes the accuracy of a given input/output pair to kmeans. PetaBricks uses this transform during autotuning and sometimes at

runtime to test the accuracy of a given configuration of the kmeans transform. The accuracy metric transform computes the $\sqrt{\sum \frac{2n}{D_i^2}}$, where D_i is the Euclidean distance between the i -th data point and its cluster center. This metric penalizes clusters that are sparse and is therefore useful for determining the quality of the computed clusters. Accuracy metric transforms such as this one might typically be written anyway for correctness or quality testing, even when programming without variable accuracy in mind.

The `accuracy_variable` k , on line 3 controls the number of clusters the algorithm generates by changing the size of the array `Centroids`. The variable k can take different values for different input sizes and different accuracy levels. The compiler will automatically find an assignment of this variable during training that meets each required accuracy level.

The `for_enough` loop on line 26 is a loop where the compiler can pick the number of iterations needed for each accuracy level and input size. During training the compiler will explore different assignments of k , algorithmic choices of how to initialize the `Centroids`, and iteration counts for the `for_enough` loop to try to find optimal algorithms for each required accuracy.

The next section goes on to explain each new variable accuracy feature in more detail.

A. Variable Accuracy Extensions

In order to support variable accuracy we made the following extensions to PetaBricks:

The `accuracy_metric` keyword in the `transform` header allows the programmer to specify the name of another user-defined transform to compute accuracy from an input/output pair. This allows the compiler to test the accuracy of different candidate algorithms during training. It also allows the user to specify a domain specific accuracy metric of interest to them.

The `accuracy_variable` keyword in the `transform` header allows the user to define one or more algorithm-specific parameters that influence the accuracy of the program. These variables are set automatically during training and are assigned different values for different input sizes. The compiler explores different values of these variables to create candidate algorithms that meet accuracy requirements while minimizing execution time.

The `accuracy_bins` keyword in the `transform` header allows the user to define the range of accuracies that should be trained for and special accuracy values of interest that should receive additional training. This field is optional and the compiler can add such values of interest automatically based on how a transform is used. If not specified, the default range of accuracies is 0 to 1.0.

The `for_enough` statement defines a loop with a compiler-set number of iterations. This is useful for defining iterative algorithms. This is syntactic sugar for adding an `accuracy_variable` to specify the number of iterations of a traditional loop.

The semantics for calling variable accuracy transforms is also extended. When a variable accuracy transform calls another variable accuracy transform (including recursively), the required sub-accuracy level is determined automatically by the compiler. This is handled by expanding each sub-call into an `either ... or` statement which allows the compiler to call the variable accuracy transform with different sub-accuracies.

The keyword `verify_accuracy` in the rule body directs the compiler to insert a runtime check for the level of accuracy attained. If this check fails the algorithm can be retried with the next higher level of accuracy or the user can provide custom code to handle this case. This keyword can be used when strict accuracy guarantees, rather than probabilistic guarantees, are desired for all program inputs.

B. Accuracy Guarantees

PetaBricks supports the following three types of accuracy guarantees:

Statistical guarantees are the most common technique used, and the default behavior of our system. They work by performing off-line testing of accuracy using a set of program inputs to determine statistical bounds on an accuracy metric to within a desired level of confidence.

Runtime checking can provide a hard guarantee of accuracy by testing accuracy at runtime and performing additional work if accuracy requirements are not met. Runtime checking can be inserted using the `verify_accuracy` keyword. This technique is most useful when the accuracy of an algorithm can be tested with low cost and may be more desirable in case where statistical guarantees are not sufficient.

Domain specific guarantees are available for many types of algorithms. In these cases, a programmer may have additional knowledge, such as a lower bound accuracy proof or a proof that the accuracy of an algorithm is independent of data, that can reduce or eliminate the cost of runtime checking without sacrificing strong guarantees on accuracy.

As with variable accuracy code written without language support, deciding which of these techniques to use with what accuracy metrics is a decision left to the programmer.

IV. PETABRICKS COMPILER INFRASTRUCTURE BACKGROUND

Figure 4 displays the general flow for the compilation of a PetaBricks transform. Compilation is split into two representations. The first representation operates at the rule level, and is similar to a traditional high level sequential intermediate representation. The second representation operates at the transform level, and is responsible for managing choices and for code synthesis.

The main transform level representation is the *choice dependency graph*, which is the primary way that choices are represented in PetaBricks. At a high level, the information contained in the choice dependency graph is similar to the dependency graph shown for our example program in Figure 2, however, the data is represented as an “inverse”

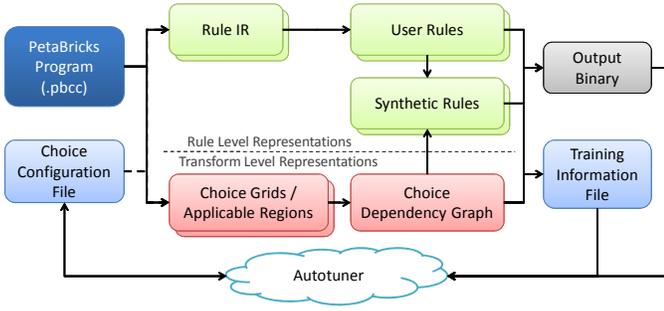


Fig. 4. Flow for the compilation of a PetaBricks program with a *single* transform. (Additional transforms would cause the center part of the diagram to be duplicated.)

```

1 population = [ ... ]
2 mutators = [ ... ]
3 for inputsize in [1, 2, 4, 8, 16, ..., N]:
4   testPopulation(population, inputsize)
5   for round in [1, 2, 3, ..., R]:
6     randomMutation(population, mutators, inputsize)
7     if accuracyTargetsNotReached(population):
8       guidedMutation(population, mutators, inputsize)
9     prune(population)

```

Fig. 5. Pseudocode showing the high level flow of our autotuning algorithm. The individual phases are described in Section V-C.

of that graph: data dependencies (previously represented by edges) are represented by vertices, while rules (previously represented by vertices) are represented by graph hyperedges. Additionally, data may be split into multiple vertices in the choice dependency graph if the transform contains rules that operate on just subregions of that data. The PetaBricks compiler uses this graph to manage code choices and to synthesize the outer control flow of the rules.

The final phase of compilation generates an *output binary* and a *training information file* containing static analysis information. These two outputs are used by the autotuner to search the space of possible algorithmic choices. The autotuner creates a *choice configuration file*, which can either be directly used by the output binary to run with dynamic switching between choices or fed back into the compiler to allow additional optimizations.

V. VARIABLE ACCURACY PETABRICKS AUTOTUNER

Figure 5 presents high level pseudocode for our autotuning algorithm. The autotuner follows a genetic algorithm approach to search through the available choice and accuracy space. It maintains a population of candidate algorithms which it continually expands using a set of mutators to allow the population to evolve more optimal algorithms. The input sizes used for testing during this process grow exponentially, which naturally exploits any optimal substructure inherent to most programs.

We will now discuss two important building blocks in our autuner, *choice configuration files* and *mutator functions*, before discussing the individual phases of the autotuner.

A. Choice Configuration Files

The PetaBricks compiler and autotuner represents different possible candidate algorithms through configuration files representing an assignment of decisions to all available choices. Broadly, one can divide the choices contained in the configuration file into the following categories.

- **Decision trees** to decide which algorithm to use for each choice site, accuracy, and input size.
- **Cutoffs values.** For example, switching points from a parallel work stealing scheduler to sequential code or the blocking sizes for data parallel operations.
- **Switches.** For example, the type of storage for intermediate data.
- **Synthesized functions.** A function from a dynamic input size to a value. For example, how many iterations in a `for_enough` loop, which may be different depending on the size of data the loop is operating on. The autotuner sets a fixed number of values of this function and other values are filled in through interpolation.
- **User defined parameters.**

B. Mutator Functions

Abstractly, a mutator function creates a new algorithm configuration from an existing configuration and a current training input size. The set of mutator functions is different for each program, and is generated fully automatically with information from static analysis. The created mutators can be divided into three categories:

- **Decision tree manipulation mutators** either add, remove, or change levels to a specific decision tree represented in the configuration file. When adding new levels, the cutoff point is initially set to $\frac{3N}{4}$. This leaves the behavior for smaller inputs the same, while changing the behavior for the current set of inputs being tested. This cutoff point can later be changed, as a *log-normal random scaling mutator* is introduced for each active cutoff value in the decision tree. Initially decision trees are very simple, set to use just a single algorithm. At this point, the only mutators that can be applied simply change this single algorithm or add a level to the decision tree. As the decision tree becomes more complex, more and more mutators can legally be applied and the search space grows.
- **Log-normal random scaling mutators** scale a configuration value by a random number taken from a log-normal distribution with scale of 1. This type of mutator is used to change cutoff values that are compared to data sizes. For example, blocking sizes, cutoffs in decision trees, and cutoffs to switch between sequential and parallel code. The intuition for why a log-normal distribution is used comes from the observation that small changes have larger effects on small values than large values. For example, most readers would expect changing a blocking size from 5 to 6 to have a much larger impact that changing a blocking size from 105 to

106. This observation is a general trend that applies to most configuration values dealing with sizes. We have confirmed this intuition experimentally by observing much faster convergence times with this type of scaling.

- **Uniform random mutators** replace an existing configuration value with a new value taken from a discrete uniform random distribution containing all legal values for the configuration item. This type of mutator is used for choices where there are a relatively small number of possibilities. An example of this is deciding the scheduling strategy for a specific block of code or algorithmic choices.

C. Autotuning Phases

1) *Population Testing*: We represent both time and accuracy by using least squares to fit a normal distribution to the observed data of empirical test runs. This fitting allows us to give statistical bounds (for example with a 95% confidence) for accuracy and time. When the programmer elects to use statistical accuracy guarantees, this alone is sufficient to guarantee accuracy. When the programmer elects to have runtime verification of accuracy, we use the same approach for autotuning to find a tuned algorithm that is verified using a runtime-only check (this check must be disabled during autotuning to allow exploration of the choice space).

An important decision that must be made is how many times to test each candidate algorithm. With too few tests, random deviations may cause non-optimal decisions to be made, while with too many tests, autotuning will take an unacceptably long time. Our autotuner dynamically decides how many tests to run between a configurable upper and lower limit. We first run the minimum number of tests, then, as candidates are compared we use the following heuristic to decide when to dynamically run more tests. When comparing two candidate algorithms, C_1 and C_2 , we perform the following steps:

- 1) Use statistical hypothesis testing (a t-test) to estimate the probability $P(\text{observed results} \mid C_1 = C_2)$. If this results in a p -value less than 0.05, we consider C_1 and C_2 different and stop.²
- 2) Use least squares to fit a normal distribution to the percentage difference in the mean performance or accuracy of the two algorithms. If this distribution estimates there is a 95% probability of less than a 1% difference, consider the two algorithms the same and stop.²
- 3) If both candidate algorithms have reached the maximum number of tests, consider the two algorithms the same and stop.
- 4) Run one additional test on either C_1 or C_2 . Decide which candidate to test based on the highest expected reduction in standard error and availability of tests without exceeding the maximum.
- 5) Go to step 1.

² The constants shown here are configurable, and can be changed based on the needs of the user. We present a typical value for clarity.

This heuristic results in a good balance, where additional tests are run only where they are needed. It also gives our autotuner the ability to adapt to variance in the environment. As an interesting anecdotal experiment, if one begins to rapidly move the mouse cursor during autotuning, the increased variance from this I/O activity causes the average number of trials to approximately triple.

2) *Random Mutation*: The random mutation phase of autotuning (line 6 in Figure 5), attempts to add new members to the population using the mutators for the current algorithm. Repeatedly (for a configurable number of *attempts*), it picks a random candidate algorithm (the *parent*) from the population and a random mutator from the mutator pool and uses that mutator to create a new candidate algorithm (the *child*). Parent algorithms remain in the population, and are only removed by the pruning phase. Upon creation, a child is tested against its parent. If the child is better than the parent either in time or in accuracy it is added to the population. Child algorithms may provide different accuracies than the parent when the mutator applied makes changes that affect accuracy. This difference in accuracy is accounted for during pruning as the autotuner attempts to maintain a variety of algorithms that are optimal for each level of accuracy.

3) *Guided Mutation*: Infrequently, the random mutation process may not produce any candidate algorithms that meet the accuracy requirements given by the user. This most commonly happens with the initial population, after only 1 round of random mutation. In this case we use a guided mutation process to attempt to create a candidate algorithm that meets the user's requirements. The guided mutation performs hill climbing on the accuracy variables. If the required accuracy still cannot be attained after guided mutation, an error is reported to the user.

4) *Population Pruning*: The pruning phase of autotuning (line 9 in Figure 5), removes unfit candidates from the population. For each accuracy bin required by the user, the pruning keeps the fastest K algorithms that meet the accuracy requirement, where K is a configurable parameter of the autotuner. Any algorithms not included in these sets are removed from the population.

To understand the pruning process, one could imagine an optimal frontier of algorithms. This optimal frontier is a curve on a 2D plane where one axis is time and the other is accuracy. For each possible accuracy, there exists a single algorithm that provides at least that accuracy in the lowest possible time. The multiple accuracy bins in our autotuner store a discretized version of this optimal frontier of algorithms.

Since comparisons between candidates can result in additional trials, determining the fastest K algorithms in each accuracy bin can be expensive. To reduce the number of comparisons we use a technique similar to partition from quick sort, where we guess the pivot item based on an rough ordering using tests that have already been run.

VI. BENCHMARKS

In this section, we describe a number of benchmarks that we implemented to test the efficacy of our system.

A. Bin Packing

Bin packing is a classic NP-hard problem where the goal of the algorithm is to find an assignment of items to unit sized bins such that the number of bins used is minimized, no bin is above capacity, and all items are assigned to a bin. It is an interesting problem because, while finding the optimal assignment is NP-hard, there are a number of polynomial time approximation algorithms that each provides different levels of approximation and performance.

The bin packing benchmark demonstrates the ability of our system to handle a large number of algorithmic choices. Variable accuracy is attained primarily through using different algorithms. We implemented 13 well known bin packing algorithms: AlmostWorstFit, AlmostWorstFitDecreasing, BestFit, BestFitDecreasing, FirstFit, FirstFitDecreasing, LastFit, LastFitDecreasing, ModifiedFirstFitDecreasing, NextFit, NextFitDecreasing, WorstFit, and WorstFitDecreasing.

To train this benchmark, we generate training data by dividing up full bins into a number of items such that the resulting distribution of item sizes matches that of a distribution of interest to us. Using this method, we can construct an accuracy metric that measures the relative performance of an algorithm to the optimal packing at training time, without the need for an exponential search. In this way, we are able to efficiently autotune the benchmark for a particular distribution of item sizes with an effective accuracy metric.

B. Clustering

Clustering divides a set of data into clusters based on similarity. The problem of k -clustering is NP-hard when k is fixed. Clustering is a common technique for statistical data analysis in areas including machine learning, pattern recognition, image segmentation and computational biology. We implemented a variant of Lloyd’s algorithm for k -means clustering. Our implementation gives the autotuner multiple choices for creating initial assignments of clusters, including random and k -means++ [7].

In our PetaBricks transform, the number of clusters, k , is the accuracy variable to be determined on training. Several algorithmic choices are implemented in our version of k -means clustering: The initial set of k cluster centers are either chosen randomly among the n data points, or according to the k -means++ algorithm [7], which chooses subsequent centers from the remaining data points with probability proportional to the distance squared to the closest center. Once the initial cluster centers are computed, the final cluster assignments and center positions are determined by iterating, either until a fixed point is reached or in some cases when the compiler decides to stop early.

The training data is a randomly generated clustered set of n points in two dimensions. First, \sqrt{n} “center” points

are uniformly generated from the region $[-250, 250] \times [-250, 250]$. The remaining $n - \sqrt{n}$ data points are distributed evenly to each of the \sqrt{n} centers by adding a random number generated from a standard normal distribution to the corresponding center point. The optimal value of $k = \sqrt{n}$ is not known to the autotuner.

Rather than assigning a fixed k through a heuristic (such as the commonly used $k = \sqrt{n}$), we define k as an accuracy variable and allow the autotuner to set it. This allows the number of clusters to change based on how compact clusters the user of the algorithm requests through the accuracy requirement. The accuracy metric used is $\sqrt{\frac{2n}{\sum D_i^2}}$, where D_i is the Euclidean distance between the i -th data point and its cluster center. The reciprocal is chosen such that a smaller sum of distance squared will give a higher accuracy.

C. 3D Variable-Coefficient Helmholtz Equation

The variable coefficient 3D Helmholtz equation is a partial differential equation that describes physical systems that vary through time and space. Examples of its use are in the modeling of vibration, combustion, wave propagation, and climate simulation. It can be expressed as:

$$\alpha(a\phi) - \beta \nabla \cdot (b \nabla \phi) = f,$$

where α and β are constants, a , b , and f are scalar valued functions over the area of interest, and ϕ is the unknown quantity we are solving for.

As an accuracy metric, we used the ratio between the RMS error of the initial guess fed into the algorithm and the RMS error of the guess afterwards. During training, the values of a and b were taken from the uniform distribution between 0.5 and 1 to ensure the system is positive-definite.

This benchmark utilizes multiple resolution levels, where each recursive call works on a problem with half as many points in each dimension. Since this is a three dimensional problem, every time a recursive call is made, the amount of data decreases by a factor of eight, possibly changing key performance parameters such as iteration strategy and number.

Additionally, there is a lot of state data that is transformed (either averaged down or interpolated up) between levels of recursion due to the presence of the variable coefficient arrays a and b . The overhead of making recursive calls influences the decision of when it is optimal to transition from recursing further to smaller problem sizes or to stop and solve the problem as best we can on the current problem size using a direct or iterative method.

D. Image Compression

Our image compression benchmark performs Singular Value Decomposition (SVD) on an $m \times n$ matrix. SVD is major component in some image compression algorithms [8]. For any $m \times n$ real matrix A with $m \geq n$, the SVD of A is $A = U\Sigma V^T$. The columns u_i of the matrix U , the columns v_i of V , and the diagonal values σ_i of Σ (singular values) form the best rank- k approximation of A , given by $A_k = \sum_{i=1}^k \sigma_i u_i v_i^T$. Only the first k columns of U and V and the

first k singular values σ_i need to be stored to reconstruct the image approximately.

The SVD of a square matrix A can be computed using the eigenvalues and eigenvectors of the matrix $H = [0 \ A^T; A \ 0]$. The number of singular values, k , to be used in the approximation is the accuracy variable to be determined by the PetaBricks autotuner. The transform for matrix approximation consists of a hybrid algorithm for finding all eigenvalues and eigenvectors, which combines Divide and Conquer, QR Iteration, and Bisection method. Another algorithmic choice exposed is using the Bisection method for only k eigenvalues and eigenvectors. The accuracy metric used is the ratio between the RMS error of the initial guess (the zero matrix) to the RMS error of the output compared with the input matrix A , converted to log-scale.

E. 2D Poisson's Equation

The 2D Poisson's equation is an elliptic partial differential equation that describes heat transfer, electrostatics, fluid dynamics, and various other engineering disciplines. The continuous and discrete versions are:

$$\nabla^2 \phi = f \quad \text{and} \quad Tx = b,$$

where T , x , and b are the finite difference discretization of the Laplace operator, ϕ , and f , respectively.

To build an autotuned multigrid solver for Poisson's equation, we consider the use of three basic algorithmic building blocks: one direct (band Cholesky factorization through LAPACK's DPBSV routine), one iterative (Red-Black Successive Over Relaxation), and one recursive (multigrid V-cycles).

As an accuracy metric, we used the ratio between the RMS error of the initial guess fed into the algorithm and the RMS error of the guess afterwards. The metric was defined this way to yield a higher accuracy value whenever the error is more greatly reduced. During training, the right hand side vector b was taken to be uniform over the interval $[-2^{31}, 2^{31}]$.

Of these three algorithmic choices, only two of them are variable accuracy, while the third (direct) solves the problem to machine-precision (assuming reasonably behaved inputs). This variable accuracy algorithm is one of the more complex of our benchmarks due to its ability to tune the number of iterations at each level of recursion thus making varying accuracy guarantees during each stage of execution. For example, we may want to iterate many times at a lower recursion level to obtain a high-fidelity estimate before interpolating up to the initial problem size to save on expensive operations at the highest grid resolution. On the other hand, if our final accuracy requirements are high enough, it may not pay to do many iterations at a lower recursion level if we are going to have to do many expensive full resolution iterations either way. It is this kind of trade-off that our variable accuracy auto-tuner excels at exploring.

F. Preconditioned Iterative Solvers

Solving a linear system of equations $Ax = b$ is a common problem in both scientific research and real-world applications

such as cost optimization and asset pricing. Iterative methods are often used to provide approximate solutions as direct solvers are usually too slow to produce exact solutions. Preconditioning is a technique that speeds up the convergence of an iterative solver.

The convergence of a matrix iteration depends on the properties of the matrix A , one of which is called the condition number. A preconditioner P of a matrix A is a matrix that if well chosen, the condition number of $P^{-1}A$ is smaller than that of A . Solving the preconditioned system $P^{-1}Ax = P^{-1}b$ gives the same solution, but the rate of convergence improves. Achieving a faster convergence rate while keeping the operation of P^{-1} simple to compute is the key to finding a good preconditioner.

Our preconditioner PetaBricks transform implements three choices of preconditioners and solves the system. The first choice is the Jacobi preconditioner coupled with Preconditioned Conjugate Gradient (PCG). The preconditioner is chosen to be the diagonal of the matrix $P = \text{diag}(A)$. Another choice is to apply the polynomial preconditioner $P^{-1} = p(A)$, where $p(A)$ is an approximation of the inverse of A by using a few terms of the series expansion of A^{-1} , and solve the preconditioned system with PCG. We also implemented the Conjugate Gradient method (CG) which solves the system without any preconditioning. The accuracy metric is the ratio between the RMS error of the initial guess Ax_{in} to the RMS error of the output Ax_{out} compared to the right hand side vector b , converted to log-scale.

VII. EXPERIMENTAL RESULTS

In this section, we present experimental results for a suite of new variable accuracy PetaBricks benchmarks. We performed all tests on a 3.16 GHz 8-core (dual-Xeon X5460) system. All benchmarks are automatically parallelized by the PetaBricks compiler and were run and trained using 8 threads.

Figures 6(a)-6(f) show the speedups that are attainable when a user is in a position to use an accuracy lower than the maximum accuracies of our benchmarks. On the largest tested input size, for benchmarks such as Clustering and Preconditioner speedups range from 1.1 to 9.6x; for benchmarks such as Helmholtz, Image Compression, and Poisson speedups range from 1.3 to 34.6x; and for the Bin Packing benchmark speedups ranged from 1832 to 13789x. Such dramatic speedups are a result of algorithmic changes made by our autotuner that can change the asymptotic performance of the algorithm ($O(n)$ vs $O(n^2)$) when allowed by a change in desired accuracy level. Because of this, speedup can become a function of input size and will grow arbitrarily high for larger and larger inputs. These speedups demonstrate some of the performance improvement potentials available to programmers using our system.

A. Analysis

This section provides more detailed analysis of the impact of accuracy on algorithmic choice and of programmability. Due to space constraints we provide analysis for only a subset of

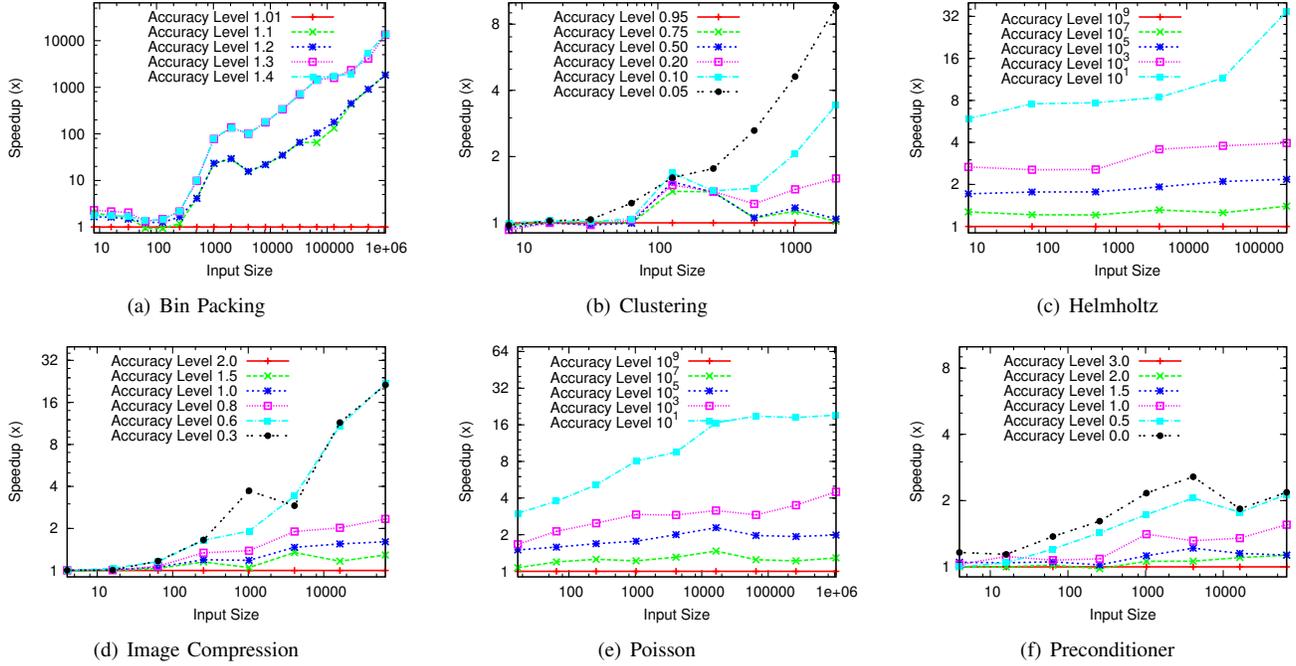


Fig. 6. Speedups for each accuracy level and input size, compared to the highest accuracy level for each benchmark. Run on an 8-way (2×4 -core Xeon X5460) system.

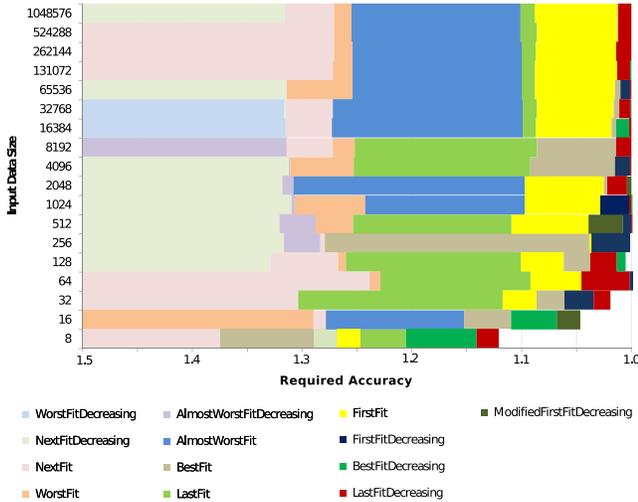


Fig. 7. Best algorithm for each accuracy and input size in the Bin Packing benchmark. By best we mean on the optimal frontier (there exists no algorithm with better performance and accuracy for a given input size on average). Accuracy is defined as the number of bins over the optimal number of bins achievable. Lower numbers represents a higher accuracy.

our benchmarks. We observed similar behaviours for the other benchmarks.

1) *Bin Packing*: Figure 7 depicts the results of autotuning the Bin Packing benchmark for various desired accuracy levels (average number of bins used over the optimal). For any desired accuracy between 1 and 1.5, the figure indicates the approximation algorithm that performs fastest on average, for input data sizes between 8 and 2^{20} generated

by our training data generator. The results show that each of the 13 approximation algorithms used by the benchmark perform fastest for some areas of the accuracy/data size space. This presents a major challenge to developers seeking high performance when using today’s programming languages since there exists no clear winner among the algorithms. Instead, the best choice will depend on the desired accuracy and input size. Thus, when writing a Bin Packing library, today’s high performance programmers have the option of either producing a brittle special-casing of the algorithmic choices manually (which would be very tedious given the number of well performing choices), or break the algorithm’s abstraction to let the user specify which choice to go with. Either of the two options are undesirable.

It is also interesting to note the relatively poor performance of *ModifiedFirstFitDecreasing*, despite the fact that it has the best provable accuracy bounds out of the set algorithms. It is best in only three small areas in the accuracy/data size space. Additionally, despite the fact that it is provably guaranteed to be within $71/60$ ($1.18\times$) of optimal, it is never the best performing algorithm when a probabilistic bound of worse than $1.07\times$ accuracy is desired. This result highlights the advantages of using an empirical approach to determining optimal algorithms when probabilistic guarantees on accuracy are permissible.

2) *Clustering*: Figure 8 illustrates the results of autotuning our k-means benchmark on our sample input of size $n = 2048$. The results show interesting algorithmic choices and number of clusters k chosen by the autotuner. For example, at accuracies greater than 0.2, the autotuned algorithm correctly

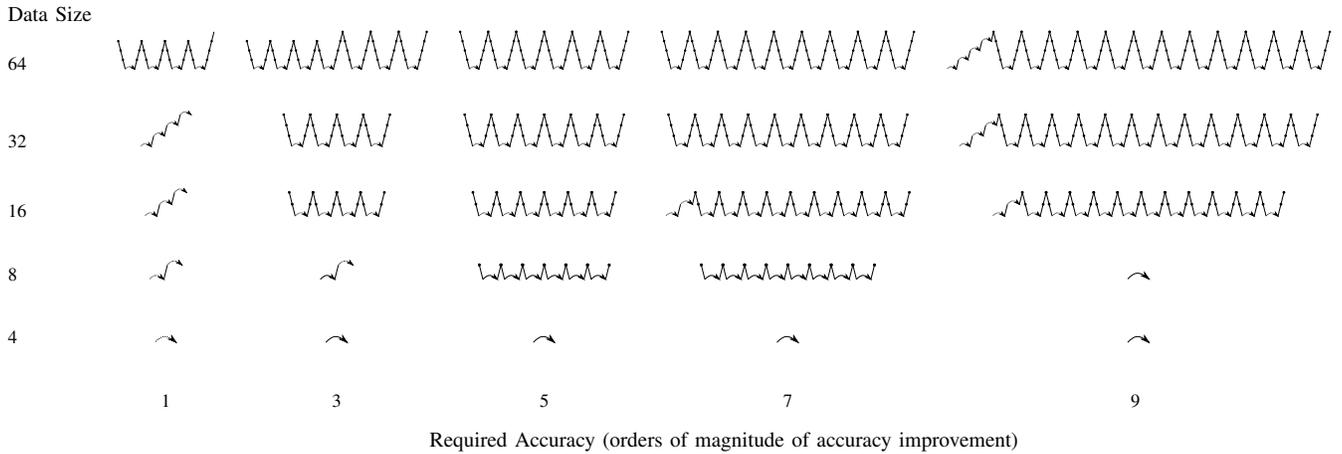


Fig. 9. Resulting cycle shapes for Helmholtz after tuning for different input data sizes and required accuracies. The solid arrows at the bottom of the cycles represent shortcut calls to the direct solver, while the dashed arrows represents an iterative solve using SOR. The dots present in the cycle represent single relaxations. Note that some paths in the full multigrid cycles skip relaxations while moving to a higher grid resolution. The notation of these diagrams is similar to that of [9].

Accuracy	k	Initial Center	Iteration Algorithm
0.10	4	random	once
0.20	38	k-means++	25% stabilize
0.50	43	k-means++	once
0.75	45	k-means++	once
0.95	46	k-means++	100% stabilize

Fig. 8. Algorithm selection and initial k value results for autotuned k-means benchmark for various accuracy levels with $n=2048$ and $k_source=45$

uses the accuracy metric (based on Euclidean distances between data points and cluster centers) to construct an algorithm that picks a k value that is close to 45, which is the number of clusters generated by our training data (which is not known to the autotuner).

At accuracy 0.1, the autotuner determines 4 to be the best choice of k and chooses to start with a random cluster assignment with only one level of iteration. While this is a very rough estimate of k and a very rough cluster assignment policy, it is sufficient to achieve the desired low level of accuracy. To achieve accuracy 0.2, the autotuner uses 38 clusters, which is slightly less than the predetermined value. Our autotuned algorithm determines the initial cluster centers by k-means++, and iterates until no more than 25% of the cluster assignments change. For accuracy 0.5 and 0.75, the k s picked by the autotuner algorithm are 43 and 45 respectively, which are only slightly smaller or equal to the predetermined k . The initial centers are decided by k-means++ and only one iteration is used. By successfully finding a number of clusters that is close to the predetermined k and picking good initial centers, only one iteration is needed on average during training to achieve a high level of accuracy. Finally, to achieve the highest accuracy of 0.95, the algorithm uses k value of 46. Initial centers are determined by k-means++ and iterations are performed until a fixed point is reached. It is interesting to note that on average, the autotuner finds that a value of k that is one higher than

the k used to generate the data, is best to minimize the user specified accuracy metric,

3) *3D Variable-Coefficient Helmholtz Equation*: Figure 9 presents the multigrid cycle shapes chosen by our autotuner for the Helmholtz algorithm. The shapes depict execution traces of the algorithm for varying accuracy levels and input sizes. The results show that compiler is able to do a good job searching the space of possible cycle shapes despite having to make difficult time-accuracy trade-offs at every stage of the recursion.

The asymmetry in some of the figures are due to a property of the algorithm that allows for an *estimation phase*, during which work is done to converge towards the solution at smaller problem sizes before work is expended at the largest problem size. This is done to provide the Helmholtz solver with an initial guess closer to the optimal solution, which in some cases can pay for itself in saved iterations at the most expensive level.

We found that in most cases with a large enough input, the accuracy levels used for recursive calls could be well below the desired accuracy of the final output given enough repetitions of the recursive call. Further, the depth of the V-cycles can be truncated by substituting the traditional deep V-cycle shape with a shallow V-cycle with an iterative bottom solve. This performance enhancing substitution is made possible by the autotuner's awareness of the different algorithmic choices available to achieve each desired accuracy level.

Additionally, it is interesting to notice the effect of the desired accuracy on the shape of the cycles. For the lowest accuracy requirement, we find that the result of the estimate phase (the non-symmetric, leftmost part of the cycle shape) is sufficiently accurate to meet the accuracy requirement. When solving for 3 orders of magnitude of accuracy improvement, the algorithm begins to rely more on the solve phase. At 5 and 7 orders of magnitude of accuracy improvement, the algorithm decides not to perform any estimation at all the majority of the

time. This result stands in contrast to the results at 9 orders of magnitude of accuracy improvement, where for data sizes equal to and greater than 16, the algorithm performs extensive estimation through multiple SOR relaxations at different levels of recursion. Additionally, for this accuracy at input size 8, it is also interesting to see that the algorithm abandons the use of recursion completely, opting instead to solve the problem with the ideal direct solver.

4) *Programmability*: While determining the programmer productivity of a new language can be quite challenging, our anecdotal experience has shown that our extensions greatly simplify the task of programming variable accuracy code. We have written a variable accuracy version of the 2D Poisson’s equation solver benchmark in the PetaBricks language [5] both before and after we added our new variable accuracy language constructs. We found that our new language features greatly simplified the benchmark, resulting in a 15.6x reduction in code size.

In the original PetaBricks language, we were able to leverage the language’s autotuner to perform the search through the accuracy performance space. However, unlike in the new code, much of the heavy lifting during the training stage had to be performed by code written by the programmer. For example, the original code contained specialized transforms used only during training that predetermined the levels of accuracy required at each recursive step in the multigrid algorithm. These transforms stored this information in a file which was used during subsequent non-training runs. Additionally, we were able to eliminate a substantial amount of code duplication because we were able to represent variable accuracy directly instead of being forced to represent it as algorithmic choices. Finally, we should note that had the original code been written in a language without autotuning support, the code would have no doubt been even more complex if it were to not expose the numerous choices in the multigrid solver to the user.

The amount of time spent training is heavily dependent on both the configuration used for the autotuner and the time complexity of the application being tuned. The configuration used for the autotuner determines the number of times the program will be run to measure performance and accuracy metrics for different parameter values. A larger number of runs will be chosen automatically in cases where the timing or accuracy of the result has larger variance – this is highly benchmark dependent. Total training times for the various benchmark configurations shown in Figure 5 ranged from fifteen minutes to two hours per benchmark, where the longest training times resulted from running our most complex benchmarks on very large inputs.

VIII. RELATED WORK

Techniques such as Loop Perforation [10], Code Perforation [11], and Task Skipping [12], [13] automatically transform existing computations and/or programs to achieve higher performance. The resulting new computations may skip subcomputations (for example loop iterations or tasks) that may not be needed to achieve a certain level of accuracy.

The computation may perform less computational work and therefore execute more quickly and/or consume less energy. While this approach can be performed robustly in many cases, it is not sound and therefore may require additional programmer time to verify the correctness of the perforated code (should such verification be needed or desired). In contrast, our system provides a new language and compiler that enables programmers to safely write programs with variable accuracy in mind right from the start. In addition to altering loop iterations, our language allows programmers to specify entirely different algorithmic choices and data representations that may be optimal for different accuracies.

PowerDail [14] is a system that converts static configuration parameters that already exist in a program into dynamic knobs that can be tuned at runtime. Their system can then change these knobs at runtime to make the program meet performance and power usage goals. They use an application wide quality of service metric to measure the loss or gain in accuracy.

Our work also bears similarities to the Green system [15], whose primary goal is to lower the power requirements of programs. Green uses pragma-like annotations to allow multiple versions of a function that have different power requirements and resulting accuracies. Green uses a global quality of service metric to monitor the impact of running the various approximate versions of the code. PetaBricks differs from Green in that it supports multiple accuracy metrics per program, allows the definition of a much larger class of algorithmic choices, has parallelism integrated with its choice model, and contains a robust genetic autotuner.

Finally, there exists a large variety of work related to PetaBrick’s autotuning approach of optimizing programs. For example, a number of empirical autotuning frameworks have been developed for building efficient, portable libraries in specific domains. PHiPAC [16] is an autotuning system for dense matrix multiply. ATLAS [17] utilizes empirical autotuning to produce a cache-contained matrix multiply. FFTW [18] uses empirical autotuning to combine solvers for FFTs. A system by Kessler *et al.* [19], [20] automatically composes algorithms using empirical techniques. Other autotuning systems include SPARSITY [21] for sparse matrix computations, SPIRAL [22], [23], [24] for digital signal processing, UHFFT [25] for FFT on multicore systems, and OSKI [26] for sparse matrix kernels. ActiveHarmony [27], [28] provides a general framework for tuning configurable libraries and exploring different compiler optimizations. In addition to these systems, various performance models and tuning techniques [29], [30], [31], [32] have been proposed to evaluate and guide automatic performance tuning.

IX. CONCLUSIONS

We have presented a new programming model where trade-offs between time and accuracy are exposed at the language level to the compiler. To the best of our knowledge, this is the first programming language that incorporates a comprehensive solution for choices relating to algorithmic accuracy. We have developed novel techniques to automatically search

the space of algorithms and parameters to construct an optimized algorithm for each accuracy level required. We have implemented a suite of 6 benchmarks that are representative of commonly used algorithms that leverage variable accuracy. Using these benchmarks, we have provided evidence of the importance of exposing accuracy and algorithmic choices to the compiler when autotuning variable accuracy programs.

Using our new programming model, library writers are able to hide the implementation details of their variable accuracy algorithms without limiting the user's choice of desired accuracy and the resulting performance. Moreover, library users do not need to burden themselves with learning the implementation specific parameters that would otherwise have been exposed. Additionally, our extensions allow variable accuracy algorithms to adapt to new environments not known to the original programmer. These new environments include new architectures, where different relative costs of operations may change the accuracy/performance trade-offs of the underlying choices. In addition, our system allows variable accuracy algorithms to adapt to changes in required accuracy and accuracy metrics. This ability to adapt extends the lifetime of programs, since the program can automatically change to meet the needs of future generations of users.

REFERENCES

- [1] L. A. Zadeh, "Fuzzy logic, neural networks, and soft computing," *Commun. ACM*, vol. 37, no. 3, pp. 77–84, 1994.
- [2] W. F. de la Vega and G. S. Lueker, "Bin packing can be solved within $1+\epsilon$ in linear time," *Combinatorica*, vol. 1, no. 4, pp. 349–355, 1981.
- [3] D. Young, *Iterative solution of large linear systems*. Dover Publications, 2003.
- [4] J. Ansel, C. Chan, Y. L. Wong, M. Olszewski, Q. Zhao, A. Edelman, and S. Amarasinghe, "Petabricks: A language and compiler for algorithmic choice," in *ACM SIGPLAN Conference on Programming Language Design and Implementation*, Dublin, Ireland, Jun 2009.
- [5] C. Chan, J. Ansel, Y. L. Wong, S. Amarasinghe, and A. Edelman, "Autotuning multigrid with petabricks," in *ACM/IEEE Conference on Supercomputing*, Portland, OR, Nov 2009.
- [6] J. Ansel and C. Chan, "Petabricks: Building adaptable and more efficient programs for the multi-core era," *Crossroads, The ACM Magazine for Students (XRDS)*, vol. 17, no. 1, pp. 32–37, Sep 2010.
- [7] D. Arthur and S. Vassilvitskii, "k-means++: the advantages of careful seeding," in *ACM-SIAM Symposium on Discrete Algorithms*, January 2007.
- [8] P. Waldemar and T. Ramstad, "Hybrid KLT-SVD image compression," in *IEEE International Conference on Acoustics, Speech, and Signal Processing*, Washington, DC, USA, 1997, p. 2713.
- [9] J. W. Demmel, *Applied Numerical Linear Algebra*. SIAM, August 1997.
- [10] S. Misailovic, S. Sidiroglou, H. Hoffman, and M. Rinard, "Quality of service profiling," in *International Conference on Software Engineering*, Cape Town, South Africa, May 2010.
- [11] H. Hoffmann, S. Misailovic, S. Sidiroglou, A. Agarwal, and M. Rinard, "Using code perforation to improve performance, reduce energy consumption, and respond to failures," Massachusetts Institute of Technology, Tech. Rep. MIT-CSAIL-TR-2209-042, Sep 2009.
- [12] M. Rinard, "Probabilistic accuracy bounds for fault-tolerant computations that discard tasks," in *ICS '06: Proceedings of the 20th annual international conference on Supercomputing*, 2006, pp. 324–334.
- [13] M. Rinard, "Using early phase termination to eliminate load imbalances at barrier synchronization points," in *ACM SIGPLAN Conference on Object-oriented Programming Systems and Applications*, 2007, pp. 369–386.
- [14] H. Hoffmann, S. Sidiroglou, M. Carbin, S. Misailovic, A. Agarwal, and M. Rinard, "Dynamic knobs for power-aware computing," in *Proceedings of the 15th International Conference on Architectural Support for Programming Languages and Operating Systems*, Newport Beach, CA, USA, Mar 2011.
- [15] W. Baek and T. Chilimbi, "Green: A framework for supporting energy-conscious programming using controlled approximation," in *ACM SIGPLAN Conference on Programming Language Design and Implementation*, June 2010.
- [16] J. Bilmes, K. Asanovic, C.-W. Chin, and J. Demmel, "Optimizing matrix multiply using PHI-PAC: a portable, high-performance, ANSI C coding methodology," in *Proceedings of the ACM/IEEE Conference on Supercomputing*. New York, NY, USA: ACM, 1997, pp. 340–347.
- [17] R. C. Whaley and J. J. Dongarra, "Automatically tuned linear algebra software," in *ACM/IEEE Conference on Supercomputing*. Washington, DC, USA: IEEE Computer Society, 1998, pp. 1–27.
- [18] M. Frigo and S. G. Johnson, "The design and implementation of FFTW3," *Proceedings of the IEEE*, vol. 93, no. 2, pp. 216–231, February 2005, invited paper, special issue on "Program Generation, Optimization, and Platform Adaptation".
- [19] C. W. Kessler and W. Lowe, "A framework for performance-aware composition of explicitly parallel components," in *PARCO*, ser. *Advances in Parallel Computing*, vol. 15. IOS Press, 2008, pp. 226–234.
- [20] J. Andersson, M. Ericsson, C. Kessler, and W. Lowe, "Profile-guided composition," *Lecture Notes in Computer Science*, vol. 4954, pp. 157–164, March 2008.
- [21] E. Im and K. Yelick, "Optimizing sparse matrix computations for register reuse in SPARSITY," in *Proceedings of the International Conference on Computational Science*. Springer, 2001, pp. 127–136.
- [22] M. Püschel, J. M. F. Moura, B. Singer, J. Xiong, J. R. Johnson, D. A. Padua, M. M. Veloso, and R. W. Johnson, "Spiral: A generator for platform-adapted libraries of signal processing algorithms," *IJHPCA*, vol. 18, no. 1, pp. 21–45, 2004.
- [23] Y. Voronenko, F. de Mesmay, and M. Püschel, "Computer generation of general size linear transform libraries," in *International Symposium on Code Generation and Optimization (CGO)*, 2009, pp. 102–113.
- [24] F. Franchetti, F. de Mesmay, D. McFarlin, and M. Püschel, "Operator language: A program generation framework for fast kernels," in *IFIP Working Conference on Domain Specific Languages (DSL WC)*, ser. *Lecture Notes in Computer Science*, vol. 5658. Springer, 2009, pp. 385–410.
- [25] A. Ali, L. Johnsson, and J. Subhlok, "Scheduling FFT computation on SMP and multicore systems," in *Proceedings of the ACM/IEEE Conference on Supercomputing*. New York, NY, USA: ACM, 2007, pp. 293–301.
- [26] R. Vuduc, J. W. Demmel, and K. A. Yelick, "OSKI: A library of automatically tuned sparse matrix kernels," in *Proceedings of the Scientific Discovery through Advanced Computing Conference*, ser. *Journal of Physics: Conference Series*. San Francisco, CA, USA: Institute of Physics Publishing, June 2005.
- [27] C. Țăpuș, I.-H. Chung, and J. K. Hollingsworth, "Active harmony: towards automated performance tuning," in *Proceedings of the 2002 ACM/IEEE conference on Supercomputing*, ser. *Supercomputing '02*. Los Alamitos, CA, USA: IEEE Computer Society Press, 2002, pp. 1–11.
- [28] A. Tiwari, C. Chen, C. Jacqueline, M. Hall, and J. K. Hollingsworth, "A scalable auto-tuning framework for compiler optimization," in *Proceedings of the 2009 IEEE International Symposium on Parallel & Distributed Processing*. Washington, DC, USA: IEEE Computer Society, 2009, pp. 1–12.
- [29] S. W. Williams, A. Waterman, and D. A. Patterson, "Roofline: An insightful visual performance model for floating-point programs and multicore architectures," EECS Department, University of California, Berkeley, Tech. Rep. UCB/EECS-2008-134, Oct 2008.
- [30] E. A. Brewer, "High-level optimization via automated statistical modeling," in *Proceedings of the ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*. New York, NY, USA: ACM, 1995, pp. 80–91.
- [31] M. G. Lagoudakis and M. L. Littman, "Algorithm selection using reinforcement learning," in *Proceedings of the International Conference on Machine Learning*. Morgan Kaufmann, 2000, pp. 511–518.
- [32] H. Yu, D. Zhang, and L. Rauchwerger, "An adaptive algorithm selection framework," in *Proceedings of the International Conference on Parallel Architectures and Compilation Techniques*. Washington, DC, USA: IEEE Computer Society, 2004, pp. 278–289.