OpenTuner: An Extensible Framework for Program Autotuning

Jason Ansel

Shoaib Kamil

Kalyan Veeramachaneni

Jonathan Ragan-Kelley

Jeffrey Bosboom

Una-May O'Reilly

Saman Amarasinghe

Massachusetts Institute of Technology
Cambridge, MA

{jansel, skamil, kalyan, jrk, jbosboom, unamay, saman}@csail.mit.edu

ABSTRACT

Program autotuning has been shown to achieve better or more portable performance in a number of domains. However, autotuners themselves are rarely portable between projects, for a number of reasons: using a domain-informed search space representation is critical to achieving good results; search spaces can be intractably large and require advanced machine learning techniques; and the landscape of search spaces can vary greatly between different problems, sometimes requiring domain specific search techniques to explore efficiently.

This paper introduces OpenTuner, a new open source framework for building domain-specific multi-objective program autotuners. OpenTuner supports fully-customizable configuration representations, an extensible technique representation to allow for domain-specific techniques, and an easy to use interface for communicating with the program to be autotuned. A key capability inside OpenTuner is the use of ensembles of disparate search techniques simultaneously; techniques that perform well will dynamically be allocated a larger proportion of tests. We demonstrate the efficacy and generality of OpenTuner by building autotuners for 7 distinct projects and 16 total benchmarks, showing speedups over prior techniques of these projects of up to 2.8× with little programmer effort.

1. INTRODUCTION

Program autotuning is increasingly being used in domains such as high performance computing and graphics to optimize programs. Program autotuning augments traditional human-guided optimization by offloading some or all of the search for an optimal program implementation to an automated search technique. Rather than optimizing a program directly, the programmer expresses a search space of possible implementations and optimizations. Autotuning

can often make the optimization process more efficient as autotuners are able to search larger spaces than is possible by hand. Autotuning also provides performance portability, as the autotuning process can easily be re-run on new machines which require different sets of optimizations. Finally, multi-objective autotuning can be used to trade off between performance and accuracy, or other criteria such as energy consumption and memory usage, and provide programs which meet given performance or quality of service targets.

While the practice of autotuning has increased in popularity, autotuners themselves often remain relatively simple and project specific. There are three main challenges which make the development of autotuning frameworks difficult.

The first challenge is using the right configuration representation for the problem. Configurations can contain parameters that vary from a single integer for a block size to a much more complex type such as an expression tree representing a set of instructions. The creator of the autotuner must find ways to represent their complex domain-specific data structures and constraints. When these data structures are naively mapped to simpler representations, such as a point in high dimensional space, locality information is lost which makes the search problem much more difficult. Picking the right representation for the search space is critical to having an effective autotuner. To date, all autotuners that have used a representation other than the simplest ones have had custom project-specific representations.

The second challenge is the size of the valid configuration space. While some prior autotuners have worked hard to prune the configuration space, we have found that for many problems excessive search space pruning will miss out on non-intuitive good configurations. We believe providing all the valid configurations of these search spaces is better than artificially constraining search spaces and possibly missing optimal solutions. Search spaces can be very large, up to 10^{3600} possible configurations for one of our benchmarks. Full exhaustive search of such a space will not complete in human lifetimes! Thus, intelligent machine learning techniques are required to seek out a good result with a small number of experiments.

The third challenge is the landscape of the configuration space. If the configuration space is a monotonic function, a search technique biased towards this type of search space

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

PACT'14, August 24–27, 2014, Edmonton, AB, Canada.

Copyright 2014 ACM 978-1-4503-2809-8/14/08 ...\$15.00.

<http://dx.doi.org/10.1145/2628071.2628092>.

(such as a hill climber) will be able to find the optimal configuration. If the search space is discontinuous and haphazard an evolution algorithm may perform better. However, in practice search spaces are much more complex, with discontinuities, high dimensionality, plateaus, hills with some of the configuration parameters strongly coupled and some others independent from each other. A search technique that is optimal in one type of configuration space may fail to locate an adequate configuration in another. It is difficult to provide a robust system that performs well in a variety of situations. Additionally, many application domains will have domain-specific search techniques (such as scheduling or blocking heuristics) which may be critical to finding an optimal solution efficiently. This has caused most prior autotuners to use customized search techniques tailored to their specific problem. This requires machine learning expertise in addition to the individual domain expertise to build an autotuner for a system. We believe that this is one of the main reasons that, while autotuners are recognized as critical for performance optimization, they have not seen commodity adoption.

In this paper we present OpenTuner, a new framework for building domain-specific program autotuners. OpenTuner features an extensible configuration and technique representation able to support complex and user-defined data types and custom search heuristics. It contains a library of predefined data types and search techniques to make it easy to setup a new project. Thus, OpenTuner solves the custom configuration problem by providing not only a library of data types that will be sufficient for most projects, but also extensible data types that can be used to support more complex domain specific representations when needed.

A core concept in OpenTuner is the use of *ensembles* of search techniques. Many search techniques (both built in and user-defined) are run at the same time, each testing candidate configurations. Techniques which perform well by finding better configurations are allocated larger budgets of tests to run, while techniques which perform poorly are allocated fewer tests or disabled entirely. Techniques are able to share results using a common results database to constructively help each other in finding an optimal solution. algorithms add results from other techniques as new members of their population. To allocate tests between techniques we use an optimal solution to the *multi-armed bandit* problem using area under the curve credit assignment. Ensembles of techniques solve the large and complex search space problem by providing both a robust solutions to many types of large search spaces and a way to seamlessly incorporate domain specific search techniques.

1.1 Contributions

This paper makes the following contributions:

- To the best of our knowledge, OpenTuner is the first to introduce a general framework to describe complex search spaces for program autotuning.
- OpenTuner introduces the concept of ensembles of search techniques to program autotuning, which allow many search techniques to work together to find an optimal solution.
- OpenTuner provides more sophisticated search techniques than typical program autotuners. This enables

expanded uses of program autotuning to solve more complex search problems and pushes the state of the art forward in program autotuning in a way that can easily be adopted by other projects.

- We demonstrate the versatility of our framework by building autotuners for 7 distinct projects and demonstrate the effectiveness of the system with 16 total benchmarks, showing speedups over existing techniques of up to $2.8\times$.
- We show that OpenTuner is able to succeed both in massively large search spaces, exceeding 10^{3600} possible configurations in size, and in smaller search spaces using less than 2% of the tests required for exhaustive search.

2. RELATED WORK

Package	Domain	Search Method
Active Harmony [31]	Runtime System	Nelder-Mead
ATLAS [34]	Dense Linear Algebra	Exhaustive
FFTW [14]	Fast Fourier Transform	Exhaustive/Dynamic Prog.
Insieme [19]	Compiler	Differential Evolution
OSKI [33]	Sparse Linear Algebra	Exhaustive+Heuristic
PATUS [9]	Stencil Computations	Nelder-Mead or Evolutionary
PetaBricks [4]	Programming Language	Bottom-up Evolutionary
Sepya [21]	Stencil Computations	Random-Restart Gradient Ascent
SPIRAL [28]	DSP Algorithms	Pareto Active Learning

Figure 1: Summary of selected related projects using autotuning

A number of offline empirical autotuning frameworks have been developed for building efficient, portable libraries in specific domains; selected projects and techniques used are summarized in Figure 1. ATLAS [34] utilizes empirical autotuning to produce an optimized matrix multiply routine. FFTW [14] uses empirical autotuning to combine solvers for FFTs. Other autotuning systems include SPIRAL [28] for digital signal processing PATUS [9] and Sepya [21] for stencil computations, and OSKI [33] for sparse matrix kernels.

The area of iterative compilation contains many projects that use different machine learning techniques to optimize lower level compiler optimizations [2, 15, 26, 1]. These projects change both the order that compiler passes are applied and the types of passes that are applied.

In the dynamic autotuning space, there have been a number of systems developed [18, 17, 22, 6, 8, 5] that focus on creating applications that can monitor and automatically tune themselves to optimize a particular objective. Many of these systems employ a control systems based autotuner that operates on a linear model of the application being tuned. For example, PowerDial [18] converts static configuration parameters that already exist in a program into dynamic knobs that can be tuned at runtime, with the goal of trading QoS guarantees for meeting performance and power usage goals. The system uses an offline learning stage to construct a linear model of the choice configuration space which can be subsequently tuned using a linear control system. The system employs the heartbeat framework [16] to provide feedback to the control system. A similar technique is employed in [17], where a simpler heuristic-based controller dynamically adjusts the degree of loop perforation performed on a target application to trade QoS for performance.

Figure 2: Overview of the major components in the OpenTuner framework.

Related to our Mario benchmark, Murphy [23] presents algorithms for playing NES games based on example human inputs. In contrast, this paper’s Mario example has no knowledge of the game except for pixels moved to the right.

3. THE OPENTUNER FRAMEWORK

Our terminology reflects that the autotuning problem is cast as a search problem. The search space is made up of *configurations*, which are concrete assignments of a set of *parameters*. Parameters can be *primitive* such as an integer or *complex* such as a permutation of a list. When the performance, output accuracy, or other metrics of a configuration are measured (typically by running it in a domain-specific way), we call this measurement a *result*. *Search techniques* are methods for exploring the search space and make requests for measurement called *desired results*. Search techniques can change configurations using a user-defined *configuration manipulator*, which also includes *parameters* corresponding directly the parameters in the configuration. Some parameters include *manipulators*, which are opaque functions that make stochastic changes to a specific parameter in a configuration.

Figure 2 provides an overview of the major components in OpenTuner. The *search* process includes techniques, which use the user defined configuration manipulator in order to read and write configurations. The *measurement* processes evaluate candidate configurations using a user defined measurement function. These two components communicate exclusively through a *results database* used to record all results collected during the tuning process, as well as providing ability to perform multiple measurements in parallel.

3.1 OpenTuner Usage

To implement an autotuner with OpenTuner, first, the user must define the search space by creating a *configuration manipulator*. This configuration manipulator includes a set of parameter objects which OpenTuner will search over. Second, the user must define a *run* function which evaluates the fitness of a given configuration in the search space to produce a result. These must be implemented in a small Python program in order to interface with the OpenTuner API.

Figure 3 shows an example of using OpenTuner to search over the space of GCC compiler flags in order to minimize

```

import opentuner
from opentuner import ConfigurationManipulator
from opentuner import EnumParameter
from opentuner import IntegerParameter
from opentuner import MeasurementInterface
from opentuner import Result

GCC_FLAGS = [
    'align-functions', 'align-jumps', 'align-labels',
    'branch-count-reg', 'branch-probabilities',
    # ... (176 total)
]

# (name, min, max)
GCC_PARAMS = [
    ('early-inlining-insns', 0, 1000),
    ('gcse-cost-distance-ratio', 0, 100),
    # ... (145 total)
]

class GccFlagsTuner(MeasurementInterface):

    def manipulator(self):
        """
        Define the search space by creating a
        ConfigurationManipulator
        """
        manipulator = ConfigurationManipulator()
        manipulator.add_parameter(
            IntegerParameter('opt_level', 0, 3))
        for flag in GCC_FLAGS:
            manipulator.add_parameter(
                EnumParameter(flag,
                    ['on', 'off', 'default']))
        for param, min, max in GCC_PARAMS:
            manipulator.add_parameter(
                IntegerParameter(param, min, max))
        return manipulator

    def run(self, desired_result, input, limit):
        """
        Compile and run a given configuration then
        return performance
        """
        cfg = desired_result.configuration.data
        gcc_cmd = 'g++ raytracer.cpp -o ./tmp.bin'
        gcc_cmd += ' -O{0}'.format(cfg['opt_level'])
        for flag in GCC_FLAGS:
            if cfg[flag] == 'on':
                gcc_cmd += ' -f{0}'.format(flag)
            elif cfg[flag] == 'off':
                gcc_cmd += ' -fno-{0}'.format(flag)
        for param, min, max in GCC_PARAMS:
            gcc_cmd += ' --param {0}={1}'.format(
                param, cfg[param])

        compile_result = self.call_program(gcc_cmd)
        assert compile_result['returncode'] == 0
        run_result = self.call_program('./tmp.bin')
        assert run_result['returncode'] == 0
        return Result(time=run_result['time'])

if __name__ == '__main__':
    argparser = opentuner.default_argparser()
    GccFlagsTuner.main(argparser.parse_args())

```

Figure 3: GCC/G++ flags autotuner using OpenTuner.

execution time of the resulting program. In Section 4, we present results on an expanded version of this example which obtains up to 2.8x speedup over -O3.

This example tunes three types of flags to GCC. First it chooses between the four optimization levels -O0, -O1, -O2, -O3. Second, for 176 flags listed on line 8 it decides between turning the flag on (with -fFLAG), off (with -fno-FLAG), or omitting the flag to allow the default value to take precedence. Including the default value as a choice is not necessary for completeness, but speeds up convergence and results in shorter command lines. Finally, it assigns a bounded integer value to the 145 parameters on line 15 with the --param NAME=VALUE command line option.

The method `manipulator` (line 23), is called once at startup and creates a `ConfigurationManipulator` object which defines the search space of GCC flags. All accesses to configurations by search techniques are done through the configuration manipulator. For optimization level, an `IntegerParameter` between 0 and 3 is created. For each flag, a `EnumParameter` is created which can take the values `on`, `off`, and `default`. Finally, for the remaining bounded GCC parameters, an `IntegerParameter` is created with the appropriate range.

The method `run` (line 40) implements the measurement function for configurations. First, the configuration is realized as a specific command line to `g++`. Next, this `g++` command line is run to produce an executable, `tmp.bin`, which is then run using `call_program`. `call_program` is a convenience function which runs and measures the execution time of the given program. Finally, a `Result` is constructed and returned, which is a database record type containing many other optional fields such as `time`, `accuracy`, and `energy`. By default OpenTuner minimizes the `time` field, but this can be customized.

3.2 Search Techniques

To provide robust search, OpenTuner includes techniques that can handle many types of search spaces and runs a collection of search techniques at the same time. Techniques which perform well are allocated more tests, while techniques which perform poorly are allocated fewer tests. Techniques share results through the results database, so that improvements made by one technique can benefit other techniques. This sharing occurs in technique-specific ways; for example, evolutionary techniques add results found by other techniques as members of their population. OpenTuner techniques are meant to be extended. Users can define custom techniques which implement domain-specific heuristics and add them to *ensembles* of pre-defined techniques.

Ensembles of techniques are created by instantiating a *meta technique*, which is a technique made up of a collection of other techniques. The OpenTuner search driver interacts with a single *root* technique, which is typically a meta technique. When the meta technique gets allocated tests, it incrementally decides how to divide these tests among its sub-techniques. OpenTuner contains an extensible class hierarchy of techniques and meta techniques, which can be combined together and used in autotuners.

3.2.1 AUC Bandit Meta Technique

In addition to a number of simple meta techniques, such as round robin, OpenTuner’s core meta technique used

in results is the *multi-armed bandit with sliding window, area under the curve credit assignment* (AUC Bandit) meta technique. A similar technique was used in [25] in the different context of online operator selection. It is based on an optimal solution to the multi-armed bandit problem [12]. The multi-armed bandit problem is the problem of picking levers to pull on a slot machine with many arms each with an unknown payout probability. The sliding window makes the technique only consider a subset of history (the length of the window) when making decisions. This meta-technique encapsulates a fundamental trade-off between *exploitation* (using the best known technique) and *exploration* (estimating the performance of all techniques).

The AUC Bandit meta technique assigns each test to the technique, t , defined by the formula $\arg \max_t (AUC_t + C \sqrt{\frac{2 \lg |H|}{H_t}})$ where $|H|$ is the length of the sliding history window, H_t is the number of times the technique has been used in that history window, C is a constant controlling the exploration/exploitation trade-off, and AUC_t is the credit assignment term quantifying the performance of the technique in the sliding window. The second term in the equation is the exploration term which gets smaller the more often a technique is used.

The area under the curve credit assignment mechanism, based on [13], draws a curve by looking at the history for a specific technique and looking only at if a technique yielded a new global best or not. If the technique yielded a new global best, a upward line is drawn, otherwise a flat line is drawn. The area under this curve (scaled to a maximum value of 1) is the total credit attributed to the technique. This credit assignment mechanism can be described more precisely by the formula: $AUC_t = \frac{2}{|V_t|(|V_t|+1)} \sum_{i=1}^{|V_t|} i V_{t,i}$ where V_t is the list of uses of t in the sliding window history. $V_{t,i}$ is 1 if using technique t the i th time in the history resulted in a speedup, otherwise 0.

3.2.2 Other Techniques

OpenTuner includes implementations of the techniques: differential evolution; many variants of Nelder-Mead search and Torczon hillclimbers; a number of evolutionary mutation techniques; pattern search; particle swarm optimization; and random search. OpenTuner also includes a bandit mutation technique which uses the same AUC Bandit method to decide which manipulator function across all parameters to call on the best known configuration. These techniques span a range of strategies and are each biased to perform best in different types of search spaces. They also each contain many settings which can be configured to change their behavior. Each technique has been modified so that with some probability it will use information found by other techniques if other techniques have discovered a better configuration.

The default meta technique, used in results in this paper and meant to be robust, uses an AUC Bandit meta technique to combine greedy mutation, differential evolution, and two hill climber instances.

3.3 Configuration Manipulator

The configuration manipulator provides a layer of abstraction between the search techniques and the raw configuration structure. It is primarily responsible for managing a list of parameter objects, each of which can be

Figure 4: Hierarchy of built in parameter types. User defined types can be added at any point below Primitive or Complex in the tree.

used by search techniques to read and write parts of the underlying configuration.

The default implementation of the configuration manipulator uses a fixed list of parameters and stores the configuration as a dictionary from parameter name to parameter-dependent data type. The configuration manipulator can be extended by the user either to change the underlying data structure used for configurations or to support a dynamic list of parameters that is dependent on the configuration instance.

3.3.1 Parameter Types

OpenTuner has a hierarchy of built-in parameter types. Each parameter type is responsible for interfacing between the raw representation of a parameter in the configuration and the standardized view of that parameter presented to the search techniques. Parameter types can be extended both to change the underlying representation, and to change the abstraction provided to search techniques to cause a parameter to be search in different ways.

From the viewpoint of search techniques there are two main types of parameters, each of which provides a different abstraction to the search techniques:

Primitive parameters present a view to search techniques of a numeric value with an upper and lower bound. These upper and lower bounds can be dependent on the configuration instance.

The built in parameter types `Float` and `LogFloat` (and similarly `Integer` and `LogInteger`) both have identical representations in the configuration, but present a different view of the underlying value to the search techniques. `Float` is presented directly to search techniques, while `LogFloat` presents a log scaled view of the underlying parameter to search techniques. To a search technique, halving and doubling a log scaled parameter are changes of equal magnitude. Log scaled variants of parameters are often better for parameters such as block sizes where fixed changes in values have diminishing effects the larger the parameter becomes. `PowerOfTwo` is a commonly used special case, similar to `LogInteger`, where the legal values of the parameter are restricted to powers of two.

Complex parameters present a more opaque view to search techniques. Complex parameters have a variable set of manipulation operators (manipulators) which make stochastic changes to the underlying parameter. These manipulators are arbitrary functions defined on the parameter which can make high level type dependent changes. Complex parameters are meant to be easily extended to add domain specific structures to the search space. New complex parameters define custom manipulation operators for use by the configuration manipulators.

The built in parameter types `Boolean`, `Switch`, and `Enum` could theoretically also be represented as primitive parameters, since they each can be translated directly to a small integer representation. However, in the context of search techniques they make more sense as complex parameters. This is because, for primitive parameters, search techniques will attempt to follow gradients. These parameter types are unordered collections of values for which no gradients exist. Thus, the complex parameter abstraction is a more efficient representation to search over.

The `Permutation` parameter type assigns an order to a given list of values and has manipulators which make various types of random changes to the permutation. A `Schedule` parameter is a `Permutation` with a set of dependencies that limit the legal order. Schedules are implemented as a permutation that gets topologically sorted after each change. Finally, a `Selector` parameter is a special type of tree which is used to define a mapping from an integer input to an enumerated value type.

In addition to these primary primitive and complex abstractions for parameter types, there are a number of derived ways that search techniques will interact with parameters in order to more precisely convey intent. These are additional methods on parameter which contain default implementations for both primitive and complex parameter types. These methods can optionally be overridden for specific parameters types to improve search techniques. Parameter types will work without these methods being overridden, but implementing them can improve results.

As an example, a common operation in many search techniques is to add the difference between configuration *A* and *B* to configuration *C*. This is used both in differential evolution and many hill climbers. Complex parameters have a default implementation of this indent which compares the value of the parameter in the 3 configurations: if $A = B$, then there is no difference and the result is *C*; similarly, if $B = C$, then *A* is returned; otherwise a change should be made so random manipulators are called. This works in general, but for individual parameter types there are often better interpretations. For example, with permutations one could calculate the positional movement of each item in the list and calculate a new permutation by applying these movements again.

3.4 Objectives

OpenTuner supports multiple user-defined objectives. Result records have fields for `time`, `accuracy`, `energy`, `size`, `confidence`, and user defined data. The default objective is to minimize time. Many other objectives are supported, such as: maximize accuracy; threshold accuracy while minimizing time; and maximize accuracy then minimize size. The user can easily define their own objective by defining comparison operators and display methods on a subclass of `Objective`.

3.5 Search Driver and Measurement

OpenTuner is divided into two submodules, search and measurement. The search driver and measurement driver in each of these modules orchestrate most of the framework of the search process. These two modules communicate only through the results database. The measurement module is minimal by design and is primarily a wrapper around the user defined measurement function which creates results from configurations.

This division between search and measurement is motivated by a number of different factors:

- To allow parallelism between multiple search measurement processes, possibly across different machines. Parallelism is most important in the measurement processes since in most autotuning applications measurement costs dominate. To allow for parallelism the search driver will make multiple requests for desired results without waiting for each request to be fulfilled. If a specific technique is blocked waiting for results, other techniques in the ensemble will be used to fill out requests to prevent idle time.
- The separation of the measurement modules is desirable to support online learning and sideline learning. In these setups, autotuning is not done before deployment of an application, but is done online as an application is running or during idle time. Since the measurement module is minimal by design, it can be replaced by a domain specific online learning module which periodically examines the database to decide which configuration to use and records performance back to the database.
- Finally, in many embedded or mobile settings which require constrained measurement environments it is desirable to have a minimal measurement module which can easily be re-implemented in other languages without needing to modify the majority of the OpenTuner framework.

3.6 Results Database

The results database is a fully featured SQL database. All major database types are supported, and SQLite is used if the user has not configured a database type so that no setup is required. It allows different techniques to query and share results in a variety of ways and is useful for introspection about the performance of techniques across large numbers of runs of the autotuner.

4. EXPERIMENTAL RESULTS

Project	Benchmark	Possible Configurations
GCC/G++ Flags	<i>all</i>	10^{806}
Halide	Blur	10^{25}
Halide	Wavelet	10^{32}
Halide	Bilateral	10^{176}
HPL	<i>n/a</i>	$10^{9.9}$
PetaBricks	Poisson	10^{3657}
PetaBricks	Sort	10^{90}
PetaBricks	Strassen	10^{188}
PetaBricks	TriSolve	10^{1559}
Stencil	<i>all</i>	$10^{6.5}$
Unitary	<i>n/a</i>	10^{21}
Mario	<i>n/a</i>	10^{6328}

Figure 5: Search space sizes in number of possible configurations, as represented in OpenTuner.

We validated OpenTuner by using it to implement autotuners for seven distinct projects. This section describes these seven projects, the autotuners we implemented, and presents results comparing to prior practices in each project.

Figure 5 lists, for each benchmark, the number of distinct configurations that can be generated by OpenTuner. This

measure is not perfect because some configurations may be semantically equivalent and the search space depends on the representation chosen in OpenTuner. It does, however, provide a sense of the relative size of each search space, which is useful as a first approximation of tuning difficulty. For many of these benchmarks, the size of the search space makes exhaustive search intractable. Thus, we compare to the best performance obtained by the benchmark authors when optimality is impossible to quantify.

For each benchmark, we keep the environment as constant as possible, but do not use the same environment for all benchmarks, due to particular requirements for each benchmark, and because one (the stencil benchmark) is based on data from an environment outside our control.

4.1 GCC/G++ Flags

The GCC/G++ flags autotuner is described in detail in Section 3.1. There are a number of features that were omitted from the earlier example code for simplicity, which are included in the full version of the autotuner.

First, we added error checking to gracefully handle the compiler or the output program hanging, crashing, running out of memory, or otherwise going wrong. Our tests uncovered a number of bugs in GCC which triggered internal compiler errors and we implemented code to detect, diagnose, and avoid error-causing sets of flags. We are submitting bug reports for these crashes to the GCC developers.

Second, instead of using a fixed list of flags and parameters (which the example does for simplicity), our full autotuner automatically extracts the supported flags from `g++ --help=optimizers`. Parameters and legal ranges are extracted automatically from `params.def` in the GCC source code.

Additionally, there were a number of smaller features such as: time limits to abort slow tests which will not be optimal; use of `LogInteger` parameter types for some values; a `save_final_config` method to output the final flags; and many command line options to control autotuner behavior.

We ran experiments using `gcc 4.7.3-1ubuntu1`, on an 8 total core, 2-socket Xeon E5320. We allowed flags such as `-ffast-math` which can change rounding or NaN behavior of floating-point numbers and have small impacts on program results. We still observe speedups with these flags removed.

For target programs to optimize we used: a fast Fourier transform in C, `fft.c`, taken from the SPLASH2 [35] benchmark suite; a C++ template matrix multiply, `matrix-multiply.cpp`, written by Xiang Fan [11] (version 3); a C++ ray tracer, `raytracer.cpp`, taken from the scratchpixel website [27]; and a genetic algorithm to solve the traveling salesman program in C++, `tsp_ga.cpp`, by Kristoffer Nordkvist [24], which we modified to run deterministically. These programs were chosen to span a range from highly optimized codes, like `fft.c` which contains cache aware tiling and threading, to less optimized codes, like `matrixmultiply.cpp` which contains only a transpose of one of the inputs.

Figure 6 shows the performance for autotuning GCC flags on these four different sample programs. Final speedups ranged from $1.15\times$ for FFT to $2.82\times$ for matrix multiply.

Figure 7 shows a separate experiment to try to understand which flags contributed the most to the speedups obtained. Full GCC command lines found contained over 250 options

Figure 6: **GCC/G++ Flags:** Execution time (lower is better) as a function of autotuning time. Aggregated performance of 30 runs of OpenTuner, error bars indicate median, 20th, and 80th percentiles. Note that in (b) the O1/O2/O3 and in (c) the O2/O3 lines are on top of each other and may be difficult to see.

and are difficult understand by hand. To approximate the importance of each flag, we measured the slowdown resulting from removing it from the autotuned GCC command. We normalize all slowdowns to sum to 100%. This experiment was conducted on a single run of the autotuner. This measurement of importance is imperfect because flags can interact in complex ways and are not independent. We note that the total slowdown from removing each flag independently is larger than the slowdown of removing all flags.

The benchmarks show interesting behavior and varying sets of optimal flags. The matrix multiply benchmark reaches optimal performance through a few large steps, and its speedup is dominated by 3 flags: `-fno-exceptions`, `-fwrapv`, and `-funsafe-math-optimizations`. On the other hand, TSP takes many small, incremental steps and its speedup is spread over a large number of flags with smaller effects. The FFT benchmark is heavily hand optimized and performs more poorly at `-O3` than at `-O2`; the flags found for it are primarily disabling, rather than enabling, various compiler optimizations which interact poorly with its hand-written optimizations. While there are some patterns, each benchmark requires a different set of flags to get the best performance.

4.2 Halide

Halide [29, 30] is a domain-specific language and compiler for image processing and computational photography, specifically targeted towards image processing *pipelines* (graphs) that contain many stages. Halide separates the expression of the kernels and pipeline itself from the pipeline’s *schedule*, which defines the order of execution and placement of data by which it is mapped to machine execution. The schedule dictates how the Halide compiler synthesizes code for a given pipeline. This allows expert programmers to easily define and explore the space of complex schedules which result in high performance.

Until today, production uses of Halide have relied on hand-tuned schedules written by experts. The autotuning problem in Halide is to *automatically* synthesize schedules and search for high-performance configurations of a given pipeline on a given architecture. The Halide project previously integrated a custom autotuner to automatically search over the space of schedules, but it was removed because it became too complex to maintain and was difficult to use in practice.¹

¹Unfortunately, the original Halide autotuner cannot be used as a baseline to compare against, because the Halide language and compiler have changed too much since its removal to make direct comparison meaningful.

Figure 7: Top 10 most important GCC flags for each benchmark. Importance is defined as the slowdown when the flag is removed from the final configuration. Importance is normalized to sum to 100%.

Figure 8: **Halide**: Execution time (lower is better) as a function of autotuning time. Aggregated performance of 30 runs of OpenTuner, error bars indicate median, 20th, and 80th percentiles.

Halide’s model of schedules is based on defining three things for each function in a pipeline:

1. In what order pixels within the stage should be evaluated.
2. At what granularity pixels should be interleaved between producer and consumer stages.
3. At what granularity pixels should be stored for reuse.

In the Halide language, these choices are specified by applying schedule operators to each function. The composition of potentially many schedule operators for each function together define the organization of computations and data for a whole pipeline.

For example, the best hand-tuned schedule (against which we compare our autotuner) for a two-stage separable blur pipeline is written as follows:

```
blur_y.split(y, y, yi, 8)
  .parallel(y)
  .vectorize(x, 8);
blur_x.store_at(blur_y, y)
  .compute_at(blur_y, yi)
  .vectorize(x, 8);
```

`blur_y(x, y)` and `blur_x(x, y)` are the Halide functions which make up the pipeline. The scheduling operators available to the autotuner are the following:

- **split** introduces a new variable and loop nest by adding a layer of blocking. Recursive splits make the search space theoretically infinite; we limit the number of splits to at most 4 per dimension per function, which is sufficient in practice. We represent each of these splits as a **PowerOfTwoParameter**, where setting the size of the split to 1 corresponds to not using the split operator.

- **parallel**, **vectorize**, and **unroll** cause the loop nest associated with a given variable in the function to be executed in parallel (over multiple threads), vectorized (e.g., with SSE), or unrolled. OpenTuner represents each of these operators as a parameter per variable per function (including variables introduced by splits). The **parallel** operator is a **BooleanParameter**, while **vectorize** and **unroll** are **PowerOfTwoParameters** like splits.

- **reorder** / **reorder_storage** take a list of variables and reorganizes the loop nest order or storage order for those variables. We encode a single **reorder** and a single **reorder_storage** for each function. We represent **reorder_storage** as a **PermutationParameter** over the original (pre-split) variables in the function. The **reorder** operator for each function is a **ScheduleParameter** over all the post-split variables, with the permutation constraint that the split children of a single variable are ordered in a fixed order relative to each other, to remove redundancy in the choice space.

Together, these operators define the order of execution (as a perfectly nested loop) and storage layout *within* the domain of each function, independently. Finally, the granularity of interleaving and storage *between* stages are specified by two more operators:

- **compute_at** specifies the granularity at which a given function should be computed within its consumers.
- **store_at** specifies the granularity at which a given function should be stored for reuse by its consumers.

Figure 9: **High Performance Linpack:** Execution time (lower is better) as a function of autotuning time. Aggregated performance of 30 runs of OpenTuner, error bars indicate median, 20th, and 80th percentiles.

Both granularity choices are specified as a level in the loop nest of a downstream function whose scope encloses all uses of the producer. These producer-consumer granularity choices are the most difficult parameters to encode and search, because the valid choices for these parameters are inter-dependent with both granularity and order choices for all other functions in the pipeline. Naïvely encoding these parameters without respecting this interdependence creates numerous invalid schedules—which are meaningless in Halide’s schedule language and rejected by the compiler—for each valid schedule.

The result of all these operators is a schedule space that is an exceedingly complex tree; transformations can add or delete subtrees, for example. Searching over such a space is far more difficult when it must be encoded in simple integer parameters, so we take advantage of OpenTuner’s ability to encode more complex parameter types. We represented the set of all compute and storage interleaving choices for a whole pipeline as a single instance of a custom, domain-specific parameter type. The parameter extends a `ScheduleParameter`, encoding open and close pairs for each scope in the scheduled pipeline’s loop nest, with a correction step which projects any invalid permutations back into the space of meaningful schedules.

Figure 8 presents results for three benchmarks drawn from Halide’s standard demo applications: a two-stage blur, an inverse Daubechies wavelet transform, and a fast bilateral filter using the bilateral grid. For all three of these examples OpenTuner is able to create schedules that match or beat the best hand optimized schedules shipping with the Halide source code. Results were collected on three IvyBridge-generation Intel processors² using the latest release of Halide.

4.3 High Performance Linpack

The High Performance Linpack benchmark [10] is used to evaluate floating point performance of machines ranging from small multiprocessors to large-scale clusters, and is the evaluation criterion for the Top 500 [32] supercomputer benchmark. The benchmark measures the speed of

solving a large random dense linear system of equations using distributed memory. Achieving optimal performance requires tuning about fifteen parameters, including matrix block sizes and algorithmic parameters. To assist in tuning, HPL includes a built in autotuner that uses exhaustive search over user-provided discrete values of the parameters.

We run HPL on a 2.93 GHz Intel Sandy Bridge quad-core machine running Linux kernel 3.2.0, compiled with GCC 4.5 and using the Intel Math Kernel Library (MKL) 11.0 for optimized math operations. For comparison purposes, we evaluate performance relative to Intel’s optimized HPL implementation³. We encode the input tuning parameters for HPL as naïvely as possible, without using any machine-specific knowledge. For most parameters, we utilize `EnumParameter` or `SwitchParameter`, as they generally represent discrete choices in the algorithm used. The major parameter that controls performance is the blocksize of the matrix; this we represent as an `IntegerParameter` to give as much freedom as possible for the autotuner for searching. Another major parameter controls the distribution of the matrix onto the processors; we represent this by enumerating all 2D decompositions possible for the number of processors on the machine.

Figure 9 shows the results of 30 tuning runs using OpenTuner, compared with the vendor-provided performance. The median performance across runs, after 1200 seconds of autotuning, exceeds the performance of Intel’s optimized parameters. Overall, OpenTuner obtains a best performance of 86.5% of theoretical peak performance on this machine, while exploring a miniscule amount of the overall search space. Furthermore, the blocksize chosen is not a power of two, and is generally a value unlikely to be guessed for use in hand-tuning.

4.4 PetaBricks

PetaBricks [3] is an implicitly parallel language and compiler which incorporates the concept of algorithmic choice into the language. The PetaBricks language provides a framework for the programmer to describe multiple ways of solving a problem while allowing the autotuner to determine which combination of ways is best for the user’s situation. The search space of PetaBricks programs is both over low level optimizations and over different algorithms. The autotuner is used to synthesize *poly-algorithms* which weave many individual algorithms together by switching dynamically between them at recursive call sites.

The primary components in the search space for PetaBricks programs are *algorithmic selectors* which are used to synthesize instances of poly-algorithms from algorithmic choices in the PetaBricks language. A selector is used to map input sizes to algorithmic choices, and is represented by a list of cutoffs and choices. As an example, the selector [InsertionSort, 500, QuickSort, 1000, MergeSort] would correspond to synthesizing the function:

```
void Sort(List& list) {
    if(list.length < 500)
        InsertionSort(list);
    else if( list.length < 1000)
        QuickSort(list);
    else
        MergeSort(list);
}
```

²Blur was tuned on a Core i5-3550, wavelet on a Core i7-3770, and bilateral grid on a 12 cores of a dual-socket Xeon E5-2695 v2 server.

³Available at <http://software.intel.com/en-us/articles/intel-math-kernel-library-linpack-download>.

Figure 10: **PetaBricks**: Execution time (lower is better) as a function of autotuning time. Aggregated performance of 30 runs of OpenTuner, error bars indicate median, 20th, and 80th percentiles.

where **QuickSort** and **MergeSort** recursively call **Sort** so the program dynamically switches between sorting methods as recursive calls are made on smaller and smaller lists. We used the general **SelectorParameter** to represent this choice type, which internally keeps track of the order of the algorithmic choices and the cutoffs. PetaBricks programs contain many algorithmic selectors and a large number of other parameter types, such as block sizes, thread counts, iteration counts, and program specific parameters.

Results using OpenTuner compared to the built-in PetaBricks autotuner are shown in Figure 10. The PetaBricks autotuner uses a different strategy that starts with tests on very small problem inputs and incrementally works up to full sized inputs [4]. In all cases, the autotuners arrive at similar solutions, and for Strassen, the exact same solution. For Sort and Tridiagonal Solver, OpenTuner beats the native PetaBricks autotuner, while for Poisson the PetaBricks autotuner arrives at a better solution, but has much higher variance.

The Poisson equation solver (Figure 10a) presents the most difficult search space. The search space for Poisson in PetaBricks is described in detail in [7]. It is a variable accuracy benchmark where the goal of the autotuner is to find a solution that provides 8-digits of accuracy while minimizing time. All points in Figure 10a satisfy the accuracy target, so we do not display accuracy. OpenTuner uses the *ThresholdAccuracyMinimizeTime* objective

described in Section 3.4. The Poisson search space selects between direct solvers, iterative solvers, and multigrid solvers where the shape of the multigrid V-cycle/W-cycle is defined by the autotuner. The optimal solution is a poly-algorithm composed of multigrid W-cycles. However, it is difficult to obtain 8 digits of accuracy with randomly generated multigrid cycle shapes, though it is easy with a direct solver (which solves the problem exactly). This creates a large “plateau” which is difficult for the autotuners to improve upon, and is shown near 0.16. The native PetaBricks autotuner is less affected by this plateau because it constructs algorithms incrementally bottom up; however the use of these smaller input sizes causes larger variance as mistakes early on get amplified.

4.5 Stencil

In [20], the authors describe a generalized system for autotuning memory-bound stencil computations on modern multicore machines and GPUs. By composing domain-specific transformations, the authors explore a large space of implementations for each kernel; the original autotuning methodology involves exhaustive search over thousands of implementations for each kernel.

We obtained the raw execution data, courtesy of the authors, and use OpenTuner instead of exhaustive search on the data from a Nehalem-class 2.66 GHz Intel Xeon machine, running Linux 2.6. We compare against the optimal performance obtained by the original autotuning

(a) Laplacian

(b) Divergence

Figure 11: **Stencil**: Execution time (lower is better) as a function of tests. Aggregated performance of 30 runs of OpenTuner, error bars indicate median, 20th, and 80th percentiles.

system through exhaustive search. The search space for this problem involves searching for parameters for the parallel decomposition, cache and thread blocking, and loop unrolling for each kernel; to limit the impact of control flow and cache misalignment, these parameters depend on one another (for example, the loop unrolling will be a small integer divisor of the thread blocking). We encode these parameters as `PowerOfTwoParameters` but ensure that invalid combinations are discarded.

Figure 11 shows the results of using OpenTuner for the Laplacian and divergence kernel benchmarks, showing the median performance obtained over 30 trials as a function of the number of tests (results for the gradient kernel are similar, so we omit them for brevity). OpenTuner is able to obtain peak performance on Laplacian after less than 35 tests of candidate implementations and 25 implementations for divergence; thus, using OpenTuner, less than 2% of the search space needs to be explored to reach optimal performance. These results show that even for problems where exhaustive search is tractable (though it may take days), OpenTuner can drastically improve convergence to the optimal performance with little programmer effort.

4.6 Unitary Matrices

As a somewhat different example, we use OpenTuner in an example from physics, namely the quantum control problem of synthesizing unitary matrices in $SU(2)$ in optimal time,

Figure 12: **Unitary**: Accuracy (higher is better) as a function of autotuning time. Aggregated performance of 30 runs of OpenTuner, error bars indicate median, 20th, and 80th percentiles.

using a finite control set composed of rotations around two non-parallel axes. (Such rotations generate the complete space $SU(2)$.)

Unlike other examples, which use OpenTuner as a traditional autotuner to optimize a program, the Unitary example uses OpenTuner to perform a search over the problem space as a subroutine at runtime in a program. The problem has a fixed set of operators (or controls), represented as matrices, and the goal is to find a sequence of operators that, when multiplied together, produce a given target matrix. The objective function is an accuracy value defined as a function of the distance of the product of the sequence to the goal (also called the *trace fidelity*).

Figure 12 shows the performance of the Unitary example on both easy and hard instances of the problem. For both types of problem instance OpenTuner is able to meet the accuracy target within the first few seconds. This example shows that OpenTuner can be used for more types of search problems than just program autotuning.

4.7 Mario

To demonstrate OpenTuner's versatility, we used the system to learn to play Super Mario Bros., a Nintendo Entertainment System game in which the player, as Mario, runs to the right, jumping over pits and onto enemies' heads to reach the flagpole at the end of each level. We use OpenTuner to search for a sequence of button presses that completes the game's first level. The button presses are written to a file and played back in the NES emulator FCEUX. The emulator is configured to evaluate candidate input sequences faster than real time. OpenTuner evaluates 96 input sequences in parallel in different instances of the emulator. A Lua script running in the emulator detects when Mario dies or reaches the flagpole by monitoring a location in the emulated memory. At that time it reports the number of pixels Mario has moved to the right, which is the objective function OpenTuner tries to maximize. The game is deterministic, so only one trial is necessary to evaluate this function. OpenTuner does not interpret the game's graphics or read the emulated memory; the objective function's value is the only information OpenTuner receives. Horizontal movement is controlled by many `EnumParameters` choosing

Figure 13: **Mario**: Pixels Mario moves to the right (higher is better) as a function of autotuning time. Reaching pixel 3161 (the flagpole) completes the level. Aggregated performance of 30 runs of OpenTuner, error bars indicate median, 20th, and 80th percentiles.

to run or walk left or right or to not move, each paired with an `IntParameter` specifying the duration of the movement in frames. As the goal is to move right, the search space is given a 3 to 1 right-moving bias. Jumping is controlled by `IntParameters` specifying on which frames to jump and how long to hold the jump button.

Figure 13 shows the performance of the Mario example. OpenTuner rapidly reaches 1500 pixels because many initial (mostly random) configurations will make it that far in the level. The bottlenecks evident in the figure correspond to pits which Mario must jump over, which requires coordinating all four kinds of parameter: Mario must jump on the correct frame for a long enough duration, having built up speed from walking or running right for a long enough duration, and then must not return left into the pit. Performance is aided by a “ratcheting” effect: due to the NES’s memory limits, portions of the level that have scrolled off the screen are quickly overwritten with new portions, so the game only allows Mario to move left one screen (up to 256 pixels) from his furthest-right position. The final sequences complete the level after between 3500 and 5000 input actions.

5. CONCLUSIONS

We have shown OpenTuner, a new framework for building domain-specific multi-objective program autotuners. OpenTuner supports fully customizable configuration representations and an extensible technique representation to allow for domain-specific techniques. OpenTuner introduces the concept of ensembles of search techniques in autotuning, which allow many search techniques to work together to find an optimal solution and provides a more robust search than a single technique alone.

While implementing these seven autotuners in OpenTuner, the biggest lesson we learned reinforced a core message of this paper of the need for domain-specific representations and domain-specific search techniques in autotuning. As an example, the initial version of the PetaBricks autotuner we implemented just used a point in high dimensional space as the configuration representation. This generic mapping of the search space did not work at all. It produced final configurations an order of magnitude

slower than the results presented from our autotuner that uses selector parameter types. Similarly, Halide’s search space strongly motivates domain specific techniques that make large coordinated jumps, for example, swapping scheduling operators on x and y across all functions in the program. We were able to add domain-specific representations and techniques to OpenTuner at a fraction of the time and effort of building a fully custom system for that project. OpenTuner was able to seamlessly integrate the techniques with its ensemble approach.

OpenTuner is free and open source⁴ and as the community adds more techniques and representations to this flexible framework, there will be less of a need to create a new representation or techniques for each project and we hope that the system will work out-of-the-box for most creators of autotuners. OpenTuner pushes the state of the art forward in program autotuning in a way that can easily be adopted by other projects. We hope that OpenTuner will be an enabling technology that will allow the expanded use of program autotuning both to more domains and by expanding the role of program autotuning in existing domains.

6. ACKNOWLEDGEMENTS

We would like to thank Clarice Aiello for contributing the Unitary benchmark program. We gratefully acknowledge Connelly Barnes for helpful discussions and bug fixes related to autotuning the Halide project.

This work is partially supported by DOE award DE-SC0005288 and DOD DARPA award HR0011-10-9-0009. This research used resources of the National Energy Research Scientific Computing Center, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

7. REFERENCES

- [1] F. Agakov, E. Bonilla, J. Cavazos, B. Franke, G. Fursin, M. F. P. O’boyle, J. Thomson, M. Toussaint, and C. K. I. Williams, “Using machine learning to focus iterative optimization,” in *CGO’06*, 2006, pp. 295–305.
- [2] L. Almagor, K. D. Cooper, A. Grosul, T. J. Harvey, S. W. Reeves, D. Subramanian, L. Torczon, and T. Waterman, “Finding effective compilation sequences,” in *LCTES’04*, 2004, pp. 231–239.
- [3] 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 *PLDI*, Dublin, Ireland, Jun 2009.
- [4] J. Ansel, M. Pacula, S. Amarasinghe, and U.-M. O’Reilly, “An efficient evolutionary algorithm for solving bottom up problems,” in *Annual Conference on Genetic and Evolutionary Computation*, Dublin, Ireland, July 2011.
- [5] W. Baek and T. Chilimbi, “Green: A framework for supporting energy-conscious programming using controlled approximation,” in *PLDI*, June 2010.
- [6] V. Bhat, M. Parashar, . Hua Liu, M. Khandekar, N. Kandasamy, and S. Abdelwahed, “Enabling self-managing applications using model-based online

⁴Available from <http://opentuner.org/>

- control strategies,” in *International Conference on Autonomic Computing*, Washington, DC, 2006.
- [7] C. Chan, J. Ansel, Y. L. Wong, S. Amarasinghe, and A. Edelman, “Autotuning multigrid with PetaBricks,” in *Supercomputing*, Portland, OR, Nov 2009.
 - [8] F. Chang and V. Karamcheti, “A framework for automatic adaptation of tunable distributed applications,” *Cluster Computing*, vol. 4, March 2001.
 - [9] M. Christen, O. Schenk, and H. Burkhart, “Patus: A code generation and autotuning framework for parallel iterative stencil computations on modern microarchitectures,” in *IPDPS*. IEEE, 2011.
 - [10] J. J. Dongarra, P. Luszczek, and A. Petit, “The LINPACK Benchmark: past, present and future,” *Concurrency and Computation: Practice and Experience*, vol. 15, no. 9, pp. 803–820, 2003.
 - [11] X. Fan, “Optimize your code: Matrix multiplication,” <https://tinyurl.com/kuvzbp9>, 2009.
 - [12] A. Fialho, L. Da Costa, M. Schoenauer, and M. Sebag, “Analyzing bandit-based adaptive operator selection mechanisms,” *Annals of Mathematics and Artificial Intelligence – Special Issue on Learning and Intelligent Optimization*, 2010.
 - [13] A. Fialho, R. Ros, M. Schoenauer, and M. Sebag, “Comparison-based adaptive strategy selection with bandits in differential evolution,” in *PPSN’10*, ser. LNCS, R. S. et al., Ed., vol. 6238. Springer, September 2010.
 - [14] M. Frigo and S. G. Johnson, “The design and implementation of FFTW3,” *IEEE*, vol. 93, no. 2, February 2005.
 - [15] G. Fursin, C. Miranda, O. Temam, M. Namolaru, E. Yom-Tov, A. Zaks, B. Mendelson, E. Bonilla, J. Thomson, H. Leather, C. Williams, M. O’Boyle, P. Barnard, E. Ashton, E. Courtois, and F. Bodin, “MILEPOST GCC: machine learning based research compiler,” in *Proceedings of the GCC Developers’ Summit*, Jul 2008.
 - [16] H. Hoffmann, J. Eastep, M. D. Santambrogio, J. E. Miller, and A. Agarwal, “Application heartbeats: a generic interface for specifying program performance and goals in autonomous computing environments,” in *ICAC*, New York, NY, 2010.
 - [17] H. Hoffmann, S. Misailovic, S. Sidirolglou, 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.
 - [18] H. Hoffmann, S. Sidirolglou, M. Carbin, S. Misailovic, A. Agarwal, and M. Rinard, “Power-aware computing with dynamic knobs,” in *ASPLOS*, 2011.
 - [19] H. Jordan, P. Thoman, J. J. Durillo, S. Pellegrini, P. Gschwandtner, T. Fahringer, and H. Moritsch, “A multi-objective auto-tuning framework for parallel codes,” in *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis*, ser. SC ’12, 2012.
 - [20] S. Kamil, C. Chan, L. Oliker, J. Shalf, and S. Williams, “An auto-tuning framework for parallel multicore stencil computations,” in *IPDPS’10*, 2010, pp. 1–12.
 - [21] S. A. Kamil, “Productive high performance parallel programming with auto-tuned domain-specific embedded languages,” Ph.D. dissertation, EECS Department, University of California, Berkeley, Jan 2013.
 - [22] G. Karsai, A. Ledeczi, J. Sztipanovits, G. Peceli, G. Simon, and T. Kovacs-hazy, “An approach to self-adaptive software based on supervisory control,” in *International Workshop in Self-adaptive software*, 2001.
 - [23] T. Murphy VII, “The first level of Super Mario Bros. is easy with lexicographic orderings and time travel,” April 2013.
 - [24] K. Nordkvist, “Solving TSP with a genetic algorithm in C++,” <https://tinyurl.com/lq3uqlh>, 2012.
 - [25] M. Pacula, J. Ansel, S. Amarasinghe, and U.-M. O’Reilly, “Hyperparameter tuning in bandit-based adaptive operator selection,” in *European Conference on the Applications of Evolutionary Computation*, Malaga, Spain, Apr 2012.
 - [26] E. Park, L.-N. Pouche, J. Cavazos, A. Cohen, and P. Sadayappan, “Predictive modeling in a polyhedral optimization space,” in *CGO’11*, April 2011, pp. 119–129.
 - [27] S. Pixel, “3D Basic Lessons: Writing a simple raytracer,” <https://tinyurl.com/lp8ncnw>, 2012.
 - [28] 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, 2004.
 - [29] J. Ragan-Kelley, A. Adams, S. Paris, M. Levoy, S. Amarasinghe, and F. Durand, “Decoupling algorithms from schedules for easy optimization of image processing pipelines,” *ACM Trans. Graph.*, vol. 31, no. 4, pp. 32:1–32:12, Jul. 2012.
 - [30] J. Ragan-Kelley, C. Barnes, A. Adams, S. Paris, F. Durand, and S. Amarasinghe, “Halide: a language and compiler for optimizing parallelism, locality, and recomputation in image processing pipelines,” in *Proceedings of the 34th ACM SIGPLAN conference on Programming language design and implementation*, ser. PLDI ’13. New York, NY, USA: ACM, 2013, pp. 519–530.
 - [31] C. Tapus, I.-H. Chung, and J. K. Hollingsworth, “Active harmony: Towards automated performance tuning,” in *In Proceedings from the Conference on High Performance Networking and Computing*, 2003.
 - [32] Top500, “Top 500 supercomputer sites,” <http://www.top500.org/>, 2010.
 - [33] R. Vuduc, J. W. Demmel, and K. A. Yelick, “OSKI: A library of automatically tuned sparse matrix kernels,” in *Scientific Discovery through Advanced Computing Conference*, San Francisco, CA, June 2005.
 - [34] R. C. Whaley and J. J. Dongarra, “Automatically tuned linear algebra software,” in *Supercomputing*, Washington, DC, 1998.
 - [35] S. Woo, M. Ohara, E. Torrie, J. Singh, and A. Gupta, “The SPLASH-2 programs: characterization and methodological considerations,” in *Symposium on Computer Architecture News*, June 1995.