

Collective mind: Towards practical and collaborative auto-tuning

Grigori Fursin ^{a,*}, Renato Miceli ^b, Anton Lokhmotov ^c, Michael Gerndt ^d, Marc Baboulin ^a, Allen D. Malony ^e, Zbigniew Chamski ^f, Diego Novillo ^g and Davide Del Vento ^h

^a *Inria and University of Paris-Sud, Orsay, France*

^b *University of Rennes 1, Rennes, France and ICHEC, Dublin, Ireland*

^c *ARM, Cambridge, UK*

^d *Technical University of Munich, Munich, Germany*

^e *University of Oregon, Eugene, OR, USA*

^f *Infrasoft IT Solutions, Plock, Poland*

^g *Google Inc., Toronto, Canada*

^h *National Center for Atmospheric Research, Boulder, CO, USA*

Abstract. Empirical auto-tuning and machine learning techniques have been showing high potential to improve execution time, power consumption, code size, reliability and other important metrics of various applications for more than two decades. However, they are still far from widespread production use due to lack of native support for auto-tuning in an ever changing and complex software and hardware stack, large and multi-dimensional optimization spaces, excessively long exploration times, and lack of unified mechanisms for preserving and sharing of optimization knowledge and research material.

We present a possible collaborative approach to solve above problems using Collective Mind knowledge management system. In contrast with previous cTuning framework, this modular infrastructure allows to preserve and share through the Internet the whole auto-tuning setups with all related artifacts and their software and hardware dependencies besides just performance data. It also allows to gradually structure, systematize and describe all available research material including tools, benchmarks, data sets, search strategies and machine learning models. Researchers can take advantage of shared components and data with extensible meta-description to quickly and collaboratively validate and improve existing auto-tuning and benchmarking techniques or prototype new ones. The community can now gradually learn and improve complex behavior of all existing computer systems while exposing behavior anomalies or model mispredictions to an interdisciplinary community in a reproducible way for further analysis. We present several practical, collaborative and model-driven auto-tuning scenarios. We also decided to release all material at *c-mind.org/repo* to set up an example for a collaborative and reproducible research as well as our new publication model in computer engineering where experimental results are continuously shared and validated by the community.

Keywords: High performance computing, systematic auto-tuning, systematic benchmarking, big data driven optimization, modeling of computer behavior, performance prediction, collaborative knowledge management, public repository of knowledge, NoSQL repository, code and data sharing, specification sharing, collaborative experimentation, machine learning, data mining, multi-objective optimization, model driven optimization, agile development, plugin-based tuning, performance regression buildbot, open access publication model, reproducible research

1. Introduction and related work

Computer systems' users are always eager to have faster, smaller, cheaper, more reliable and power efficient computer systems either to improve their every day tasks and quality of life or to continue innovation in science and technology. However, designing and optimizing such systems is becoming excessively time

consuming, costly and error prone due to an enormous number of available design and optimization choices and complex interactions between all software and hardware components. Furthermore, multiple characteristics have to be carefully balanced including execution time, code size, compilation time, power consumption and reliability using a growing number of incompatible tools and techniques with many ad-hoc, intuition based heuristics.

* Corresponding author. E-mail: grigori.fursin@inria.fr.

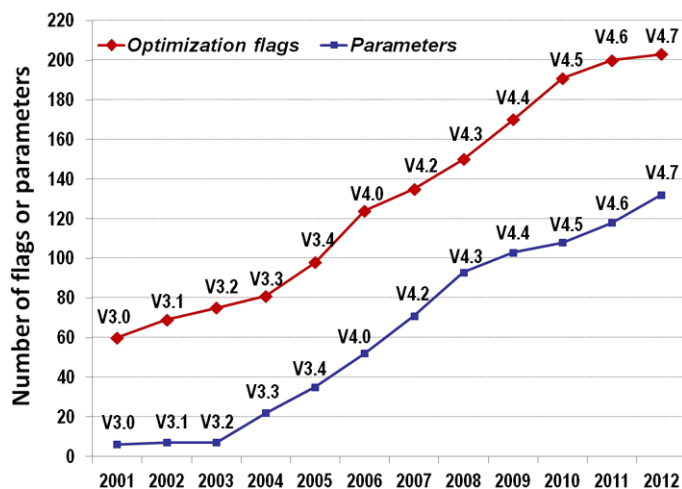


Fig. 1. Rising number of optimization dimensions in GCC in the past 12 years (Boolean or parametric flags). Obtained by automatically parsing GCC manual pages, therefore small variation is possible (script was kindly shared by Yuriy Kashnikov). (Colors are visible in the online version of the article; <http://dx.doi.org/10.3233/SPR-140396>.)

At the same time, development methodology for computer systems has hardly changed in the past decades: hardware is first designed and then the compiler is tuned for the new architecture using some ad-hoc benchmarks and heuristics. As a result, peak performance of the new systems is often achieved only for a few previously optimized and not necessarily representative benchmarks while leaving most of the real user applications severely underperforming. Therefore, users are often forced to resort to a tedious and often non-systematic optimization of their programs for each new architecture. This, in turn, leads to an enormous waste of time, expensive computing resources and energy, dramatically increases development costs and time-to-market for new products and slows down innovation [4,22,65,74].

Various off-line and on-line auto-tuning techniques together with run-time adaptation and split compilation have been introduced during the past two decades to address some of the above problems and help users automatically improve performance, power consumption and other characteristics of their applications. These approaches treat rapidly evolving computer system as a black box and explore program and architecture design and optimization spaces empirically [1,8,21,27,38,44,48,49,51,53,59,64,70,73,77,78].

Since empirical auto-tuning was conceptually simple and did not require deep user knowledge about programs and computer systems, it quickly gained popularity. At the same time, users immediately faced a fundamental problem: a continuously growing number of available design and optimization choices makes it im-

possible to exhaustively explore the whole optimization space. For example, Fig. 1 shows a continuously rising number of available boolean and parametric optimizations available in a popular, production, open-source compiler GCC used in practically all Linux and Android based systems. Furthermore, there is no more single combination of flags such as `-O3` or `-Ofast` that could deliver the best execution time across all user programs. Figure 2 demonstrates 79 distinct combinations of optimizations for GCC 4.7.2 that improve execution time across 285 benchmarks with just one data set over `-O3` on Intel Xeon E5520 based platform after 5000 explored solutions using traditional iterative compilation [37] (random selection of compiler optimization flags and parameters).

Optimization space explodes even further when considering heterogeneous architectures, multiple data sets and fine grain program transformations and parameters including tiling, unrolling, inlining, padding, prefetching, number of threads, processor frequency and MPI communication [10,32,37,38,55,60]. For example, Fig. 3 shows execution time of a matrix–matrix multiplication kernel for square matrices on CPU (Intel E6600) and GPU (NVIDIA 8600 GTS), depending on their size. It motivates the need for adaptive scheduling since it may be beneficial either to run kernel on CPU or GPU depending on data set parameters (to amortize the cost of data transfers to GPU). However, as we show in [6,46,75], the final decision tree is architecture and kernel dependent and requires both off-line kernel cloning and some on-line, automatic and ad-hoc modeling of application behavior.

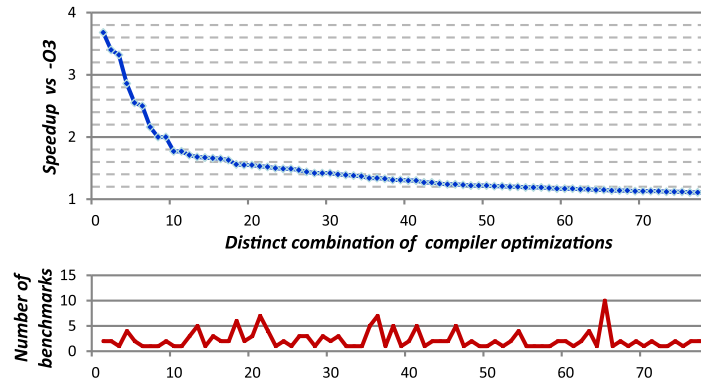


Fig. 2. Number of distinct combinations of compiler optimizations for GCC 4.7.2 with a maximum achievable execution time speedup over -O3 optimization level on Intel Xeon E5520 platform across 285 shared Collective Mind benchmarks after 5000 random iterations (top graph) together with a number of benchmarks where these combinations achieve more than 10% speedup (bottom graph). (Colors are visible in the online version of the article; <http://dx.doi.org/10.3233/SPR-140396>.)

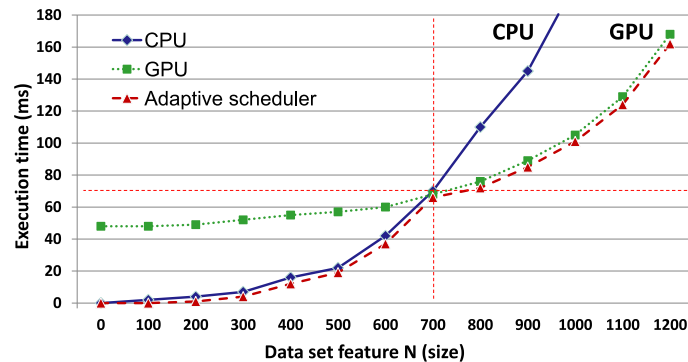


Fig. 3. Execution time of a matrix–matrix multiply kernel when executed on CPU (Intel E6600) and on GPU (NVIDIA 8600 GTS) depending on size N of square matrix as a motivation for online tuning and adaptive scheduling on heterogeneous architectures [46]. (Colors are visible in the online version of the article; <http://dx.doi.org/10.3233/SPR-140396>.)

Machine learning techniques (predictive modeling and classification) have been gradually introduced during the past decade as an attempt to address the above problems [2,14,23,46,57,62,67,71,72,81]. These techniques can help speed up program and architecture analysis, optimization and co-design by narrowing down regions in large optimization spaces with the most likely highest speedup. They usually use prior training similar to Fig. 2 in case of compiler tuning and predict optimizations for previously unseen programs based on some code, data set and system features.

During the MILEPOST project in 2006–2009, we made the first practical attempt to move auto-tuning and machine learning to production compilers including GCC by combining a plugin-based compiler framework [45] and a public repository of experimental results (cTuning.org). This approach allowed to substitute and automatically learn default compiler

optimization heuristics by crowdsourcing auto-tuning (processing a large amount of performance statistics collected from many users to classify application and build predictive models) [31,37,61]. However, this project exposed even more fundamental challenges including:

- Lack of common, large and diverse benchmarks and data sets needed to build statistically meaningful predictive models;
- Lack of common experimental methodology and unified ways to preserve, systematize and share our growing optimization knowledge and research material including benchmarks, data sets, tools, tuning plugins, predictive models and optimization results;
- Problem with continuously changing, “black box” and complex software and hardware stack with many hardwired and hidden optimization choices

and heuristics not well suited for auto-tuning and machine learning;

- Difficulty to reproduce performance results from the cTuning.org database submitted by users due to a lack of full software and hardware dependencies;
- Difficulty to validate related auto-tuning and machine learning techniques from existing publications due to a lack of culture of sharing research artifacts with full experiment specifications along with publications in computer engineering.

As a result, we spent a considerable amount of our “research” time on re-engineering existing tools or developing new ones to support auto-tuning and learning. At the same time, we were trying to somehow assemble large and diverse experimental sets to make our research and experimentation on machine learning and data mining statistically meaningful. We spent even more time when struggling to reproduce existing machine learning-based optimization techniques from numerous publications.

Worse, when we were ready to deliver auto-tuning solutions at the end of such tedious developments, experimentation and validation, we were already receiving new versions of compilers, third-party tools, libraries, operating systems and architectures. As a consequence, our developments and results were already potentially outdated even before being released while optimization problems considerably evolved.

We believe that these are major reasons why so many promising research techniques, tools and data sets for auto-tuning and machine learning in computer engineering have a life span of a PhD project, grant funding or publication preparation, and often vanish shortly after. Furthermore, we witness diminishing attractiveness of computer engineering often seen by students as “hacking” rather than systematic science. Some of the recent long-term research visions acknowledge these problems for computer engineering and many research groups search for “holy grail” auto-tuning solutions but no widely adopted solution has been found yet [22,74].

In this paper, we describe the first, to our knowledge, alternative, orthogonal, community-based and big-data driven approach to address above problems. It may help make auto-tuning a mainstream technology based on our practical experience in the MILEPOST, cTuning and Auto-tune projects, industrial usage of our frameworks and community feedback. Our main contribution is a collaborative knowledge management framework for computer engineering called Collective

Mind (or cM for short) that brings interdisciplinary researchers and developers together to organize, systematize, share and validate already available or new tools, techniques and data in a unified format with gradually exposed actions and meta-information required for auto-tuning and learning (optimization choices, features and tuning characteristics).

Our approach should allow to collaboratively prototype, evaluate and improve various auto-tuning techniques while reusing all shared artifacts just like LEGO™ pieces and applying machine learning and data mining techniques to find meaningful relations between all shared material. It can also help crowdsource long tuning and learning process including classification and model building among many participants while using Collective Mind as a performance tracking buildbot. At the same time, any unexpected program behavior or model mispredictions can now be exposed to the community through unified cM web-services for collaborative analysis, explanation and solving. This, in turn, enables reproducibility of experimental results naturally and as a side effect rather than being enforced – interdisciplinary community needs to gradually find and add missing software and hardware dependencies to the Collective Mind (fixing processor frequency, pinning code to specific cores to avoid contentions) or improve analysis and predictive models (statistical normality tests for multiple experiments) whenever abnormal behavior is detected.

We hope that our approach will eventually help the community collaboratively evaluate and derive the most effective auto-tuning and learning strategies. It should also eventually help the community collaboratively learn complex behavior of all existing computer systems using top-down methodology originating from physics. At the same time, continuously collected and systematized knowledge (“big data”) should help us make more scientifically motivated advice about how to improve design and optimization of the future computer systems (particularly on our way towards extreme scale computing). Finally, we believe that it can naturally make computer engineering a systematic science while supporting Vinton G. Cerf’s recent vision [15].

This paper is organized as follows: the current section provides motivation for our approach and related work. It is followed by Section 2 presenting possible solution to collaboratively systematize and unify our knowledge about program optimization and auto-tuning using public Collective Mind framework and repository. Section 3 presents mathematical formaliza-

tion of auto-tuning techniques. Section 4 demonstrates how our collaborative approach can be combined with several existing plugin-based auto-tuning infrastructures including MILEPOST GCC [31], OpenME [30] and Periscope Tuning Framework (PTF) [60] to start systematizing and making practical various auto-tuning scenarios from our industrial partners including continuous benchmarking and comparison of compilers, validation of new hardware designs, crowdsourcing of program optimization using commodity mobile phones and tablets, automatic modeling of application behavior, model driven optimization and adaptive scheduling. It is followed by a section on reproducibility of experimental results in our approach together with a new publication model proposal where all research material is continuously shared and validated by the community. The last section includes conclusions and future work directions.

2. Cleaning up research and experimental mess

Based on our long experience with auto-tuning and machine learning in both academia and industry, we now strongly believe that the missing piece of the puzzle to make these techniques practical is to enable sharing systematization, and reuse of all available optimization knowledge and experience from the community. However, our first attempt to crowdsource auto-tuning and machine learning using cTuning

plugin-based framework and MySQL-based repository (cTuning) [29,31] suffered from many orthogonal engineering issues. For example, we had to spend considerable effort to develop and continuously update ad-hoc research and experimental scenarios using many hardwired scripts and tools while being able to expose only a few dimensions, monitor a few characteristics and extract a few features. At the same time, we struggled with collection, processing and storing of a growing amount of experimental data in many different formats as conceptually shown in Fig. 4(a). Furthermore, adding a new version of a compiler or comparing multiple compilers at the same time required complex manipulations with numerous environment variables.

Eventually, all these problems motivated the development of a modular Collective Mind framework and NoSQL heterogeneous repository [18,30] to unify and preserve the whole experimental setups with all related artifacts and dependencies. First of all, to avoid invoking ad-hoc tools directly, we introduced cM modules which serve as wrappers around them to be able to transparently set up all necessary environment variables and validate all software and hardware dependencies before eventually calling these tools. Such an approach allows easy co-existence of multiple versions of tools and libraries while protecting experimental setups from continuous changes in the system. Furthermore, cM modules can now transparently monitor and unify all information flow in the system. For example, we currently monitor tools' command line together

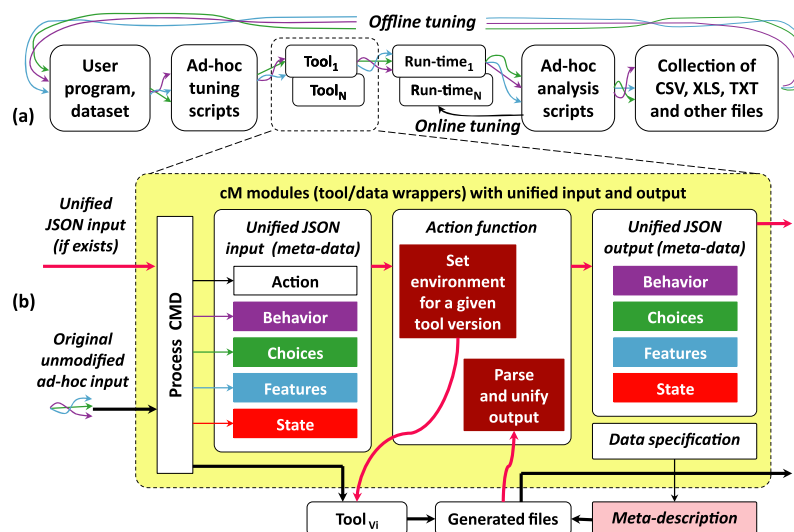


Fig. 4. Converting (a) continuously evolving, ad-hoc, hardwired and difficult to maintain experimental setups to (b) interconnected cM modules (tool wrappers) with unified, dictionary-based inputs and outputs, data meta-description, and gradually exposed characteristics, tuning choices, features and a system state. (Colors are visible in the online version of the article; <http://dx.doi.org/10.3233/SPR-140396>.)

with their input and output files to expose measured characteristics (behavior of computer systems), optimization and tuning choices, program, data set and architecture features, and a system state used in all our existing auto-tuning and machine learning scenarios as conceptually shown in Fig. 4(b).

Since researchers are often eager to quickly prototype their research ideas rather than sink in low-language implementations, complex APIs and data structures that may change over time, we decided to use a researcher friendly and portable Python language as the main language in Collective Mind (though we also provide possibility to use any other language for writing modules through an OpenME interface described later in this paper). Therefore, it is possible to run minimal cM on practically any Linux and Windows computer supporting Python. An additional benefit of using Python is a growing collection of useful packages for data management, mining and machine learning.

We also decided to switch from traditional TXT, CSV and XML formats used in the first cTuning framework to a schema-free JSON data format [63] for all module inputs, outputs and meta-description. JSON is a popular, human readable and open standard format that represent data objects as attribute–value pairs. It is now backed up by many companies, supported by most of the recent languages and powerful search engines [24], and can be immediately used for web services and P2P communication during collaborative research and experimentation. Only when the format of data becomes stable or a research technique is validated, the community can provide data specification as will be described later in this paper.

At the same time, we noticed that we can apply exactly the same concept of cM modules to systematize and describe any research and development material (code and data) while making sure that it can be easily found, reused and exposed to the Web. Researchers and developers can now categorize any collections of their files and directories by assigning an existing or adding a new cM module and moving their material to a new directory with a unique ID (UID) and an optional alias. Thus we can now abstract an access to highly heterogeneous and evolving material by gradually adding possible data actions and meta-description required for user’s research and development. For example, all cM modules have common actions to manage their data in a unified way similar to any repository such as *add*, *list*, *view*, *copy*, *move* and *search*. In addition, module *code.source* abstracts access to programs and has an

individual action *build* to compile a given program. In fact, all current cM functionality is implemented as interconnected modules including *kernel*, *core* and *repo* that provide main low-level cM functions documented at *c-mind.org/doxygen*.

In contrast with using SQL-based databases, our approach can help systematize, preserve and describe any heterogeneous code and data on any native file system without any need for specialized databases, pre-defined data schema and complex table restructuring as conceptually shown in Fig. 5. Since cM modules also have their own UOA (UID or alias), it is now possible to easily reference and find any local user material similar to DOI by a unified Collective ID (CID) of the following format:

$$\langle \text{cM module UOA} \rangle : \langle \text{cM data UOA} \rangle.$$

In addition, cM provides an option to transparently index meta-description of all artifacts using a third-party, open-source and JSON-based Elastic-Search framework implemented on Hadoop engine [24] to enable fast and powerful search queries for such a schema-free repository. Note that other similar NoSQL databases including MongoDB and CouchDB or even SQL-based repositories can be easily connected to cM to cache data or speed up queries, if necessary.

Any cM module with a given action can be executed in a unified way using JSON format as both input and output either through cM command line front-end

```
cm <module UOA><action>@input.json
```

or using one Python function from a cM kernel module

```
r = cm_kernel.access({'cm_run_module_uoa' :
                      <module UOA>,
                      'cm_action' : <action>,
                      action_parameters})
```

or as a web service when running internal cM web server (also implemented as cM *web.server* module) using the following URL

```
http://localhost:3333?cm_web_module_uoa=
<module UOA>&cm_web_action=<action>...
```

For example, a user can list all available programs in the system using *cm code.source list* and then compile a

Category	cM module	Module actions	All data	Meta description
Third-party tools, libraries	package	common*, install	<input type="checkbox"/> <input type="checkbox"/> ... <input type="checkbox"/>	<input checked="" type="checkbox"/> <input checked="" type="checkbox"/> ... <input checked="" type="checkbox"/>
High-level algorithms	algorithm	common*, transform	<input type="checkbox"/> <input type="checkbox"/> ... <input type="checkbox"/>	<input checked="" type="checkbox"/> <input checked="" type="checkbox"/> ... <input checked="" type="checkbox"/>
Applications, benchmarks, kernels	code.source	common*, build	<input type="checkbox"/> <input type="checkbox"/> ... <input type="checkbox"/>	<input checked="" type="checkbox"/> <input checked="" type="checkbox"/> ... <input checked="" type="checkbox"/>
Multiple compilers	ctuning.compiler	common*, compile_program	<input type="checkbox"/> <input type="checkbox"/> ... <input type="checkbox"/>	<input checked="" type="checkbox"/> <input checked="" type="checkbox"/> ... <input checked="" type="checkbox"/>
Predictive models	math.model	common*, build, predict, fit, detect_representative_points	<input type="checkbox"/> <input type="checkbox"/> ... <input type="checkbox"/>	<input checked="" type="checkbox"/> <input checked="" type="checkbox"/> ... <input checked="" type="checkbox"/>
Binaries and libraries	code	common*, run	<input type="checkbox"/> <input type="checkbox"/> ... <input type="checkbox"/>	<input checked="" type="checkbox"/> <input checked="" type="checkbox"/> ... <input checked="" type="checkbox"/>
Program data sets	dataset	common*, create	<input type="checkbox"/> <input type="checkbox"/> ... <input type="checkbox"/>	<input checked="" type="checkbox"/> <input checked="" type="checkbox"/> ... <input checked="" type="checkbox"/>
Operating Systems	os	common*, detect_host_family	<input type="checkbox"/> <input type="checkbox"/> ... <input type="checkbox"/>	<input checked="" type="checkbox"/> <input checked="" type="checkbox"/> ... <input checked="" type="checkbox"/>
Processors	processor	common*, detect_host_processor	<input type="checkbox"/> <input type="checkbox"/> ... <input type="checkbox"/>	<input checked="" type="checkbox"/> <input checked="" type="checkbox"/> ... <input checked="" type="checkbox"/>
Statistical and data mining functions	math.statistics.r	common*, analyze	<input type="checkbox"/> <input type="checkbox"/> ... <input type="checkbox"/>	<input checked="" type="checkbox"/> <input checked="" type="checkbox"/> ... <input checked="" type="checkbox"/>

* add, list, view, copy, move, search

cM repository directory structure: `.cmr: / module UOA (UID or alias)` / data UOA / `.cm / data.json`

Fig. 5. Gradually categorizing all available user artifacts using cM modules while making them searchable through meta-description and reusable through unified cM module actions. All material from this paper is shared through Collective Mind online live repository at c-mind.org/browse and c-mind.org/github-code-source. (Colors are visible in the online version of the article; <http://dx.doi.org/10.3233/SPR-140396>.)

given program using

```
cm code.source build
work_dir_data_uoa
  =benchmark-cbench-security-blowfish
build_target_os_uoa=windows-generic-64
```

If some parameters or dependencies are missing, the cM module should be implemented as such to inform users about how to fix these problems.

In order to simplify validation and reuse of shared experimental setups, we provide an option to keep tools inside a cM repository also together with a unified installation mechanism that resolves all software dependencies. Such packages including third-party tools and libraries can now be installed into a different entry in a cM repository with a unique IDs abstracted by cM *code* module. At the same time, OS-dependent script is automatically generated for each version of a package to set up appropriate environment including all paths. This script is automatically called before executing a given tool version inside an associated cM module as shown in Fig. 4.

In spite of its relative simplicity, the Collective Mind approach helped us to gradually clean up and systematize our material that can now be easily searched, shared, reused or exposed to the web. It also helps substitute all ad-hoc and hardwired experimental se-

tups with interconnected and unified modules and data that can be protected from continuous changes in computer systems and easily shared among workgroups. Users only need to categorize new material, move related files to a special directory of format `.cmr/(module UOA)/(data UOA)` (where `.cmr` is an acronym for Collective Mind Repository) to be automatically discovered and indexed by cM, and provide some meta-information in JSON format depending on research scenarios.

In contrast with public web-based sharing services, we provide an open-source, technology-neutral, agile, customizable, and portable knowledge management system which allows both private and public systematization of research and experimentation. To initiate and demonstrate gradual and collaborative systematization of a research material for auto-tuning and machine learning, we decided to release all related code and data at c-mind.org/browse to discuss, validate and rank shared artifacts while extending their meta-description and abstract actions with the help of the community. We described and shared multiple benchmarks, kernels and real applications using cM *code.source* module, various data sets using cM *dataset* module, various parameterized classification algorithms and predictive models using cM *math.model* module, and many others.

We also shared packages with exposed dependencies and installation scripts for many popular tools and libraries in our public cM repository at c-mind.org/repo

including GCC, LLVM, ICC, Microsoft Visual Studio compilers, PGI compilers, Open64/PathScale compilers, ROSE source-to-source compilers, Oracle JDK, VTune, NVIDIA GPU toolkit, perf, gprof, GMP, MPFR, MPC, PPL, LAPACK and many others. We hope that this will ease the burden of the community to continuously (re-)implement some ad-hoc and often unreleased experimental scenarios. In the next sections, we will show how this approach can be used to systematize and formalize auto-tuning.

3. Formalizing auto-tuning and predictive modeling

Almost all research on auto-tuning can be formalized as finding a function of a behavior of a given user program B running on a given computer system with a given data set, selected design and optimization choices including program transformations and architecture configuration \mathbf{c} , and a system state \mathbf{s} ([1,28,31,78]):

$$\mathbf{b} = B(\mathbf{c}, \mathbf{s}).$$

For example, in our current and past research and experimentation, \mathbf{b} is a behavior vector that includes execution time, power consumption, accuracy, compilation time, code size, device cost, and other important characteristics; \mathbf{c} represents the available design and optimization choices including algorithm selection, the compiler and its optimizations, number of threads, scheduling, affinity, processor ISA, cache sizes, memory and interconnect bandwidth, etc.; and finally \mathbf{s} represents a state of the system including processor frequency and cache or network contentions.

Knowing and minimizing this function is of a particular importance to our industrial partners when designing, validating and optimizing the next generation of hardware and software including compilers for a broad range of customers' applications, data sets and requirements (constraints), since it can help reduce time to market for the new systems and increase return on investment (ROI). However, the fundamental problem is that *this function B is highly non-linear with a multi-dimensional discrete and continuous space of choices* [28,38] *which is rarely possible to model analytically or evaluate empirically* using exhaustive search unless really small kernels and libraries are used with just one or a few program transformations [1,78].

This problem motivated research on automatic and empirical modeling of an associated function P that can quickly predict better design and optimization choices for a given computer system \mathbf{c} based on some features (properties) of an end-users' program, data set and a given hardware \mathbf{f} , and a current state of a computer system \mathbf{s} :

$$\mathbf{c} = P(\mathbf{f}, \mathbf{s}).$$

For example, in our research on machine-learning based optimization, vector \mathbf{f} includes semantic or static program features [2,31,62,72], data set features and hardware counters [14,46], system configuration, and run-time environment parameters among many others. However, when trying to implement practical and industrial scenarios in cTuning framework, we spent most of our time on engineering issues trying to expose characteristics, choices, features and system state using numerous, "black box" and not necessarily documented tools. Furthermore, when colleagues with a machine learning background were trying to help us improve optimization predictions, they were often quickly demotivated when trying to understand our terminology and problems.

The Collective Mind approach helped our colleagues solve this problem by formalizing the problem and gradually exposing characteristics \mathbf{b} , choices \mathbf{c} , system state \mathbf{s} and features \mathbf{f} (meta information) in experimental setups using JSON format as shown in the following real example:

```
{ "characteristics": {
  "execution_times": [ "10.3", "10.1", "13.3" ],
  "code_size": "131938", ... },
  "choices": {
    "os": "linux", "os_version": "2.6.32-5-amd64",
    "compiler": "gcc", "compiler_version": "4.6.3",
    "compiler_flags": "-O3 -fno-if-conversion",
    "platform": {
      "processor": "intel_xeon_e5520", "l2": "8192",
      "memory": "24" ... }, ... },
  "features": {
    "semantic_features": { "number_of_bb": "24", ... },
    "hardware_counters": { "cpi": "1.4" ... }, ... }
  "state": {
    "frequency": "2.27", ... }
}
```

Furthermore, we can easily convert JSON hierarchical data into a *flat vector format* to apply above mathematical formalization of auto-tuning and learning problem while making it easily understandable

to an interdisciplinary community particularly with a background in mathematics and physics. In our flat format, a *flat key* can reference any key in a complex JSON hierarchy as one string. Such *flat key* always starts with # followed by #*key* if it is a dictionary key or @*position_in_a_list* if it is a value in a list. For example, flat key for the second execution time “10.1” in one of the previous examples of information flow can be referenced as “##characteristics#execution_time@1”. Finally, users can gradually provide the following cM data specification for the flat keys in information flow to fully automate program optimization and learning (kept together with a given cM module):

```
"flattened_json_key": {
  "type": "text" | "integer" | "float" | "dict" | "list" | "uid",
  "characteristic": "yes" | "no",
  "feature": "yes" | "no",
  "state": "yes" | "no",
  "has_choice": "yes" | "no",
  "choices": [list of strings if categorical choice],
  "explore_start": "start number if numerical range",
  "explore_stop": "stop number if numerical range",
  "explore_step": "step if numerical range",
  "can_be_omitted": "yes" | "no",
  "default_value": "string"
  ...
}
```

Of course, such format may have some limitations, but it supports well our current research and experimentation on auto-tuning and will be extended only when needed. Furthermore, such implementation allowed us and our colleagues to collaboratively prototype, validate and improve various auto-tuning and learning scenarios simply by chaining available cM modules similar to components and filters in electronics (cM experimental pipelines) and reusing all shared artifacts. For example, we converted our ad-hoc build and run scripts from cTuning framework to a unified cM pipeline consisting of chained cM modules as shown in Fig. 6. This pipeline (*ctuning.pipeline.build_and_run*) is implemented and executed as any other cM module to help researchers simplify the following operations during experimentation:

- Source to source program transformation and instrumentation (if required). For example, we added support for PLUTO polyhedral compiler to enable automatic restructuring and parallelization of loops [12].
- Compilation and execution of any shared programs (real applications, benchmarks and ker-

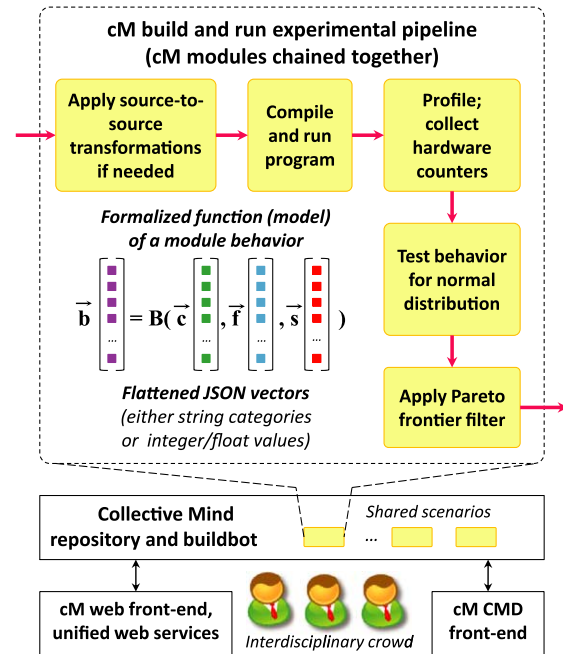


Fig. 6. Unified build and run cM pipeline implemented as chained cM modules. (Colors are visible in the online version of the article; <http://dx.doi.org/10.3233/SPR-140396>.)

nels) using *code.source* module. Meta-description of these programs includes information about how to build and execute them. Code can be executed with any shared data set as an input (*dataset* module). The community can gradually share more data sets together with the unified descriptions of their features such as dimensions of images, sizes of matrices and so on.

- Testing of measured characteristics from repeated executions for normal distribution [25] using shared cM module *ctuning.filter.variation* to be able to expose unusual behavior in a reproducible way to the community for further analysis. Unexpected behavior often means that some feature is missing in the experimental pipeline such as frequency or cache contention that can be gradually added by the community to separate executions with different contexts as described further in Section 5.
- Applying Pareto frontier filter [31,44,50] to leave only optimal solutions during multi-objective optimization when multiple characteristics have to be balanced at the same time such as execution time vs code size vs power consumption vs compilation time. This, in turn, can help to avoid collecting large amounts of off-line and often un-

essary experimental data that can easily saturate repositories and make data analysis too time consuming or even impossible (as happened several times with a public cTuning repository).

In the next section we show how we can reuse and customize this pipeline (demonstrated online at c-mind.org/ctuning-pipeline) to systematize and run some existing auto-tuning scenarios from our industrial partners.

4. Systematizing auto-tuning and learning scenarios

Unified cM build and run pipeline combined with mathematical formalization allows researchers and engineers to focus their effort on implementing and extending universal auto-tuning and learning scenarios rather than hardwiring them to specific systems, compilers, optimizations or tuned characteristics. This, in turn, allows to distribute long tuning process across multiple users while potentially solving an old and well-known problem of using a few possibly non-

representative benchmarks and a limited number of architectures when developing and validating new optimization techniques.

Furthermore, it is now possible to take advantage of mature interdisciplinary methodologies from other sciences such as physics and biology to analyze and learn the behavior of complex systems. Therefore, cM uses a top-down methodology to decompose software and hardware into simple sub-components to be able to start learning and tuning of a global, coarse-grain program behavior with respect to exposed coarse-grain tuning choices and features. Later, depending on user requirements, time budget and expected return on investment during optimization, the community can extend components to cover finer-grain tuning choices and behavior as conceptually shown in Fig. 7. Note, that when analyzing a application at a finer-grain levels such as code regions, we consider them as interacting cM components with their own vectors of tuning choices, characteristics, features and internal states. In doing so, we can analyze and learn their behavior using methodologies from quantum mechanics or agent-based modeling [69].

Gradually extend cM build and run pipeline module		Gradually expose characteristics	Gradually expose design and optimization choices, features
Select algorithm		(time) productivity, variable-accuracy, complexity ...	Language, MPI, OpenMP, TBB, MapReduce ...
Analyze and transform program	Process	time; memory usage; code size ...	transformation ordering; polyhedral transformations; transformation parameters; instruction ordering; MPI parameters; number of threads;
	Function		
	Kernel		
	Loop		
	Instruction		
Build program		time ...	compiler flags; pragmas ...
Run code	Run-time environment	time; power consumption ...	pinning/scheduling ...
	System	cost; size ...	CPU/GPU; frequency; memory ...
	Data set	size; values; description ...	precision ...
	Run-time analysis	time; precision ...	hardware counters; power meters ...
	Run-time state	processor state; cache state ...	helper threads; hardware counters ...
Analyze profile	Statistical analysis	time; size ...	instrumentation; profiling ...
Model behavior		size; precision	model type ...

Fig. 7. Gradual and collaborative top-down decomposition of computer system software and hardware using cM modules (wrappers) similar to methodology in physics. First, coarse-grain design and optimization choices and features are exposed and tuned, and later more fine-grain choices are exposed depending on the available tuning time budget and expected return on investment. (Colors are visible in the online version of the article; <http://dx.doi.org/10.3233/SPR-140396>.)

As the first practical usage scenario, we developed a universal and customizable design and optimization space exploration as cM module *ctuning.scenario.exploration* on top of *ctuning.pipeline.build_and_run_program* module to substitute most ad-hoc tuning scripts and frameworks from our past research. This scenario can be executed from the command line as any other cM module thus enabling relatively easy integration with third-party tools including compiler regression buildbots or Eclipse-based framework. However, the most user friendly way is to run scenarios is through cM web interface as demonstrated at c-mind.org/ctuning-exploration (note that we plan to improve the usability of this interface with dynamic HTML, JavaScript and Ajax technology [39] while hiding unnecessary information from users and avoiding costly page refreshes). In such a way, cM will query all chained modules for this scenario to automatically visualize all available tuning choices, characteristics, features and system states. cM will also preset all default values (if provided by specification) while allowing a user to select which choices to explore, characteristics to measure, search strategy to use, and statistical analysis for experimental results to apply.

We currently implemented and shared *uniform random and exhaustive* exploration strategies. We also plan to add adaptive, probabilistic and hill climbing sampling from our past research [27,37] or let users develop and share any other universal strategy which is not hardwired to any specific tool but can explore any available choices exposed by the scenario.

Next, we present several practical and industrial auto-tuning scenarios using above customized exploration module.

4.1. Systematizing compiler benchmarking

Validating new architecture designs across multiple benchmarks, tuning optimization heuristics of multiple versions of compilers, or tuning compiler flags for a customer application is a tedious, time consuming and often ad-hoc process that is far from being solved. In fact, it becomes even tougher with time due to the ever rising number of available optimizations (Fig. 1) and many strict requirements placed on compilers such as generating fast and small code for all possible existing architectures within a reasonable amount of time.

Collective Mind framework helps unify and distribute this process among many machines as a performance tracking buildbot. For this purpose, we customized universal cM exploration module for com-

piler flag tuning as a new *ctuning.scenario.compiler.optimizations* module. Users just need to choose a compiler version and related description of flags (example is available at c-mind.org/ctuning-compiler-desc) as an input and select either to explore a compiler flag optimization space for a given program or distribute tuning of a default compiler optimization heuristic across many machines using a set of shared benchmarks. Note, that it is possible to use Collective Mind not only on desktop machines, servers, data centers and cloud services but also on bare metal hardware or Android-based mobile devices (either through SSH or using a special *Collective Mind Node* application available in Google Play Store [20] to help deploy and crowdsource experiments on mobile phones and tablets while aggregating results in web-based cM repositories).

To demonstrate this scenario, we optimized a real image corner detection program on a commodity Samsung Galaxy Series mobile phone with ARMv6 830 MHz processor using Sourcery GCC v4.7.2 with randomly generated combinations of compiler flags of format *-O3 -f(no-)optimization_flag -parameter param=random_number_from_range*, LLVM v3.2 with *-O3* flag, and a chained Pareto frontier filter (cM module *ctuning.filter.frontier*) for multi-objective optimization (balancing execution time, code size and compilation time).

Experimental results during such exploration (cM module output) are continuously recorded in a repository in a unified flat vector format making it possible to immediately take advantage of numerous and powerful public web services for visualization, data mining and analytics (for example from Google, Microsoft, Oracle and IBM) or available as packages for Python, Weka, MATLAB, SciLab and R. For example, Fig. 8 shows 2D visualization of these experimental results using public Google Web Services integrated with cM. Such interactive graphs are particularly useful when working in workgroups or for interactive publications (as demonstrated at c-mind.org/interactive-graph-demo).

Note, that we always suggest to run optimized code several times to check variation and test distribution for normality as we used to do in physics and electronics. If such a test fails or the variation of any characteristic dimension is more than some threshold (currently set as 2%), we do not skip such case but record it as *suspicious* including all inputs and outputs for further validation and analysis by the community as described in Section 5. At the same time, using a Pareto frontier filter allows users to easily select the most appropri-

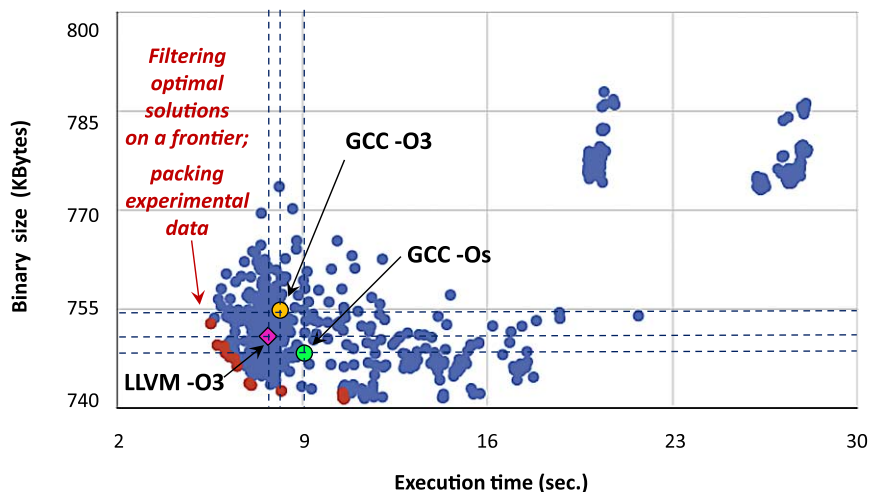


Fig. 8. Compiler flag auto-tuning to improve execution time and code size of a shared image corner detection program with a fixed data set on Samsung Galaxy Series mobile phone using cM for Android. Highlighted points represent frontier of optimal solutions as well as GCC with -O3 and -Os optimization flags versus LLVM with -O3 flag (c-mind.org/interactive-graph-demo). (Colors are visible in the online version of the article; <http://dx.doi.org/10.3233/SPR-140396>.)

ate solution depending on the further intended usage of their applications, i.e. the fastest variant if used for HPC systems, the smallest variant if used for embedded devices with very limited resources, such as credit card chips or the future “Internet of Things” devices, or balanced for both speed and size when used in mobile phones and tablets.

Since Collective Mind framework also enables co-existence of multiple versions of different compilers, checks output of programs for correct execution during optimization, and supports multiple shared benchmarks and data sets, it can be easily used as a distributed and public buildbot for rigorous performance tracking and simultaneous tuning of compilers (as shown in Fig. 2) while taking advantage of a growing number of shared benchmarks and data sets [19]. Longer term, we expect that such an approach will help the community fully automate compiler tuning for new architectures or even validate new processor designs for errors. It can also help derive a realistic, diverse and representative training set of benchmarks and data sets [55] to systematize and speed up training for machine learning based optimization prediction for previously unseen programs and architectures [31].

To continue this collaborative effort, we shared the description of all parametric and boolean (on or off) compiler flags in JSON format as “choices” for a number of popular compilers including GCC, LLVM, Open64, PathScale, PGI and ICC under *ctuning.compiler* module. We also implemented and shared several off-the-shelf classification and predictive mod-

els including KNN and SVM from our past research [2, 14,31] using *math.model* module to be able to automatically predict better compiler optimization using semantic and dynamic program features. Finally, we started implementing standard complexity reduction and differential analysis techniques [47,66] in cM to iteratively isolate unusual program behavior [36] or to find minimal set of representative benchmarks, data sets and correlating features [30,55]. Users can now collaboratively analyze unexpected program behavior, improve predictive models, find best tuning strategies and collect minimal set of influential optimizations, representative features, most accurate models, benchmarks and data sets.

4.2. Systematizing modeling of application behavior to focus optimizations

Since programs may potentially have an infinite number of data sets while auto-tuning is already time consuming, it is usually performed for one or a few and not necessarily representative data sets. Collective Mind framework can help to systematize and automate modeling of a behavior of a given application across multiple data sets to suggest where to focus further tuning (adaptive sampling and online learning) [2, 9,37]. We just needed to customize previously introduced auto-tuning pipeline to explore data set parameters (already exposed through *dataset* module) and model program behavior at the same time using either off-the-shelf predictive models including linear re-

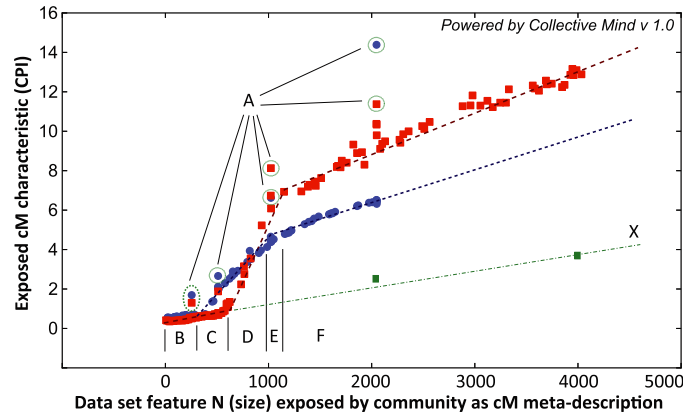


Fig. 9. Online learning (predictive modeling) of a CPI behavior of a shared LU-decomposition benchmark on 2 different platforms (Intel Core2 shown in red vs Intel i5 shown in blue) vs vector size N (data set feature). (Colors are visible in the online version of the article; <http://dx.doi.org/10.3233/SPR-140396>.)

gression, Support Vector Machines (SVM), Multivariate Adaptive Regression Splines (MARS), and neural networks available for R language and abstracted by cM module *math.model.r*, or shared user-defined hybrid models specific for a given application.

For example, Fig. 9 demonstrates how such exploration and online learning is performed using cM together with shared LU-decomposition benchmark versus size of input vector (N), measured CPI characteristic, and 2 Intel-based platforms (Intel Core2 Centrino T7500 Merom 2.2 GHz L1 = 32 KB 8-way set-associative, L2 = 4 MB 16-way set associative – red dots vs. Intel Core i5 2540M 2.6 GHz Sandy Bridge L1 = 32 KB 8-way set associative, L2 = 256 KB 8-way set associative, L3 = 3 MB 12-way set associative – blue dots).

In the beginning, cM does not have any knowledge about behavior of this (or any other) benchmark, so it simply observes and stores available characteristics along with the data set features. At each exploration (sampling) step, cM processes all historical observations using various available or shared predictive models such as SVM or MARS in order to find correlations between data set features and characteristics. At the same time it attempts to minimize Root-Mean-Square Deviation (RMSE) between predicted and measured values for all available points. Even if RMSE is relatively low, cM can continue exploring and observing behavior in order to detect discrepancies (failed predictions).

Interestingly, in our example, practically no off-the-shelf model could detect the A outliers (singularities) which appear due to cache alignment problems. However, having mathematical formalization helps inter-

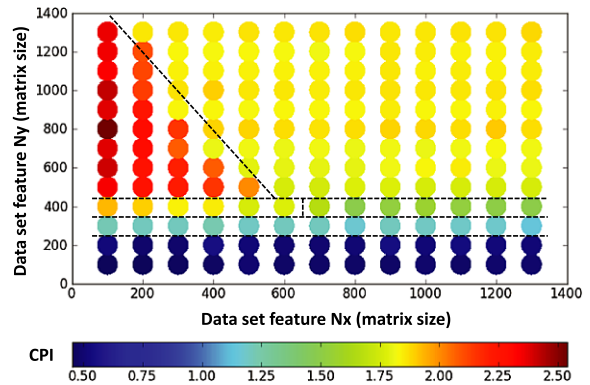


Fig. 10. CPI behavior of a matrix–matrix multiply benchmark on Intel i5 platform vs matrix size. Hyperplanes separate areas with similar behavior found using multivariate adaptive regression splines (MARS). (Colors are visible in the online version of the article; <http://dx.doi.org/10.3233/SPR-140396>.)

disciplinary community to find and share better models that minimized RMSE and model size at the same time. In the presented case, our colleagues from machine learning department managed to fit and share a hybrid, parameterized, rule-based model that first validates cases where data set size is a power of 2, otherwise it uses linear models as functions of a data set and cache size. This model resembles reversed analytical roofline model [79] though is continuously and empirically refined to capture even fine-grain effects. In contrast, standard MARS model managed to predict the behavior of a matrix–matrix multiplication kernel for different matrix sizes as shown in Fig. 10.

Such models can help focus auto-tuning on areas with distinct behavior as described in [9,27,37]. For example presented in Fig. 9, outlier points A can be

optimized using array padding; area B can profit from parallelization and traditional compiler optimizations targeting ILP; areas C–E can benefit from loop tiling; points A saturate memory bus and can also benefit from reduced processor frequency to save energy. Such optimizations can be performed automatically if exposed through cM or provided by the community as shared advices using *ctuning.advice* module.

In the end, multiple customizable models can be shared as parameterized cM modules along with applications thus allowing the community to continuously refine them or even reuse them for similar classes of applications. Finally, such predictive models can be used for effective and online compaction of experiments while avoiding collection of a large amount of data (known in other fields as a “big data” problem) and leaving only representative or unexpected behavior. It can, in turn, minimize communications between cM nodes while making Collective Mind a giant and distributed learning and decision making network to some extent similar to the brain [30].

4.3. Enabling fine-grain auto-tuning through plugins

After learning and tuning coarse-grain behavior, we gradually move to finer-grain levels including selected code regions, loop transformations, MPI parameters and so on, as shown in Fig. 7. However, in our past research, it required messy instrumentation of appli-

cations and development of complex source-to-source transformation tools and pragma-based languages.

As an alternative and simpler solution, we developed an event-based plugin framework (Interactive Compilation Interface and was recently substituted by a new and universal OpenME plugin-based framework connected to cM) to expose tuning choices and semantic program features from production compilers such as GCC and LLVM through external plugins [29,33, 45]. This plugin-based tuning technique helped us to start unifying, cleaning up and converting rigid compilers into powerful and flexible research toolsets. Such an approach also helped companies and end-users to develop their own plugins with customized optimization and tuning scenarios without rebuilding compilers and instrumenting applications thus keeping them clean and portable.

This framework also allowed to easily expose multiple semantic code features to automatically learn and improve all optimization and tuning decisions using standard machine learning techniques as conceptually shown in Fig. 11. This plugin-based tuning technique was successfully used in the MILEPOST project to automate online learning and tuning of the default optimization heuristic of GCC for new reconfigurable processors from ARC during software and hardware co-design [31]. The plugin framework was eventually added to mainline GCC since version 4.6. We are gradually adding to cM support for plugin-based selection and ordering of passes or tuning and learning of

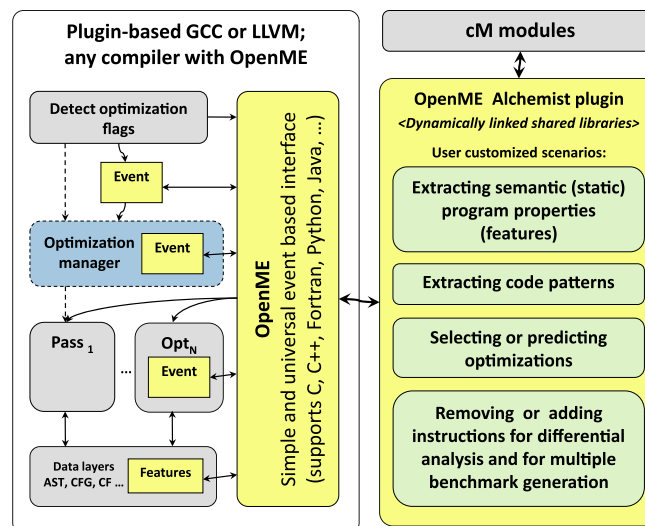


Fig. 11. Conceptual structure of compilers supporting plugin and event based interface to enable fine-grain tuning and learning of their internal and often hidden heuristics through external plugins [45]. (Colors are visible in the online version of the article; <http://dx.doi.org/10.3233/SPR-140396>.)

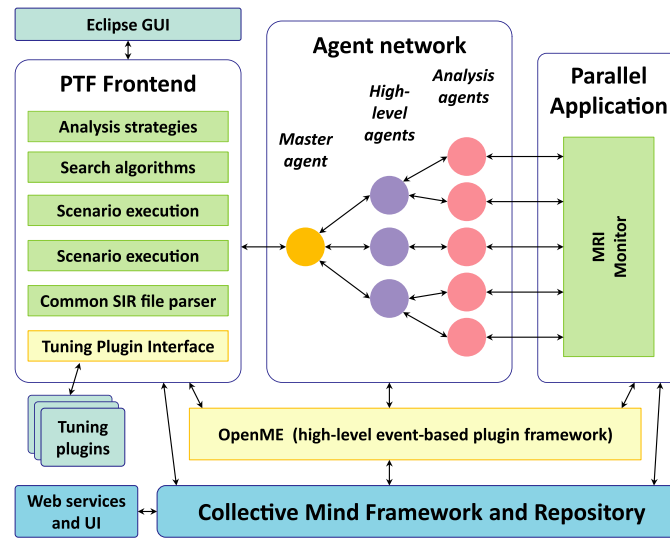


Fig. 12. PTF plugin-based application online tuning framework. (Colors are visible in the online version of the article; <http://dx.doi.org/10.3233/SPR-140396>.)

internal compiler decisions while aggregating semantic program features in a unified format using *ctuning.scenario.program.features.milepost* module.

Plugin-based static compilers can help users automatically or interactively tune a given application with a given data set for a given architecture. However, different data sets or run-time system state often require different optimizations and tuning parameters that should be dynamically selected during execution. Therefore, Periscope Tuning Framework (PTF) [60] was designed to enable and automate online tuning of parallel applications using external plugins with integrated tuning strategies. Users need to instrument application to expose required tuning parameters and measured characteristics for a given application. At the same time, tuning space can be considerably reduced inside such plugins per given application using previous compiler analysis or expert knowledge about typical performance bottlenecks and ways to detect and improve them as conceptually shown in Fig. 12.

Once the online tuning process is finished, PTF generates a report with the recommended tuning actions which can be integrated either manually or automatically into the application for further production runs. Currently, PTF includes plugins to tune execution time of high-level parallel kernels for GPGPUs, balance energy consumption via CPU frequency scaling, optimize MPI runtime parameters among many other scenarios in development.

Collective Mind can help PTF distribute tuning of shared benchmarks and data sets among many users,

aggregate results in a common repository, apply data mining and machine learning plugins to prune tuning spaces, and automate prediction of optimal tuning parameters. PTF and cM can also complement each other in terms of tuning coverage since cM currently focuses on global, high-level, machine-learning guided optimizations and compiler tuning while PTF currently focuses on finer-grain online application tuning. In our future work we plan to connect PTF and cM together using cM OpenME interface.

4.4. Systematizing split compilation and adaptive scheduling

Many current online auto-tuning techniques have a limitation – they usually do not support arbitrary online code restructuring unless complex just-in-time (JIT) compilers are used. As a possible solution to this problem, we introduced split compilation to statically enable dynamic optimizations and adaptation by cloning hot functions or kernels during compilation and providing run-time selection mechanism depending on data set features, target architecture features and a system state [34,46,55,58]. However, since this approach still requires a long and off-line training phase, we can now use Collective Mind infrastructure to systematize off-line tuning and learning of a program behavior across many data sets and computer systems as conceptually shown in Fig. 13.

Now, users can take advantage of continuously collected knowledge about program behavior and opti-

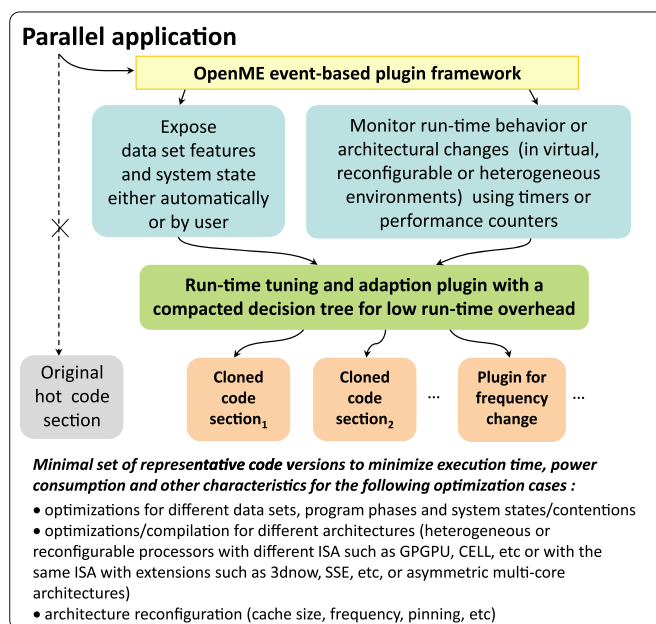


Fig. 13. Making run-time adaptation and tuning practical using static multi-versioning, features exposed by users or automatically detected, and predictive modeling (decision trees) while avoiding complex dynamic recompilation frameworks. (Colors are visible in the online version of the article; <http://dx.doi.org/10.3233/SPR-140396>.)

mization in the repository to derive a minimal set of representative optimizations or tuning parameters covering application behavior across as many data sets and architectures as possible [55]. Furthermore, it is now possible to reuse machine learning techniques from cM to automatically derive small and fast decision trees needed for realistic cases shown in Figs 3, 9 and 10. Such decision trees can now be integrated with the application through OpenME or PTF plugins to dynamically select appropriate clones and automatically adapt for heterogeneous architectures particularly in supercomputers and data centers, or even execute some external tools to reconfigure architecture (change frequency, for example) based on exposed features to minimize execution time, power consumption and other user objectives. These data set, program and architecture feature can also be exposed through plugins either automatically using OpenME-based compilers or manually through application annotation and instrumentation.

OpenME was designed especially to be very easy to use for researchers and provide a simple connection between Python-based Collective Mind and other modules or plugins written in other languages including C, C++, Fortran and Java. It has only two functions to initialize an event with an arbitrary string name, and to call it with a *void type* argument that will be handled

by a user plugin and can range from a simple integer to a cM JSON dictionary. However, since such implementation of OpenME can be relatively slow, we use fast Periscope Tuning Framework for fine-grain tuning. Possible example of such implementation for predictive scheduling of matrix multiply using OpenME interface and several clones for heterogeneous architectures [46] is presented in Fig. 14.

Our static function cloning approach with dynamic adaptation was recently added to mainline GCC since version 4.8. We hope that together with OpenME, PTF and cM, it will help systematize research on split compilation while focusing on finding and exposing the most appropriate features to improve run-time adaptation decisions [55] using recent advances in machine learning, data mining and decision making [11,26,43, 52].

4.5. Automating benchmark generation and differential analysis

Our past research on machine learning to speed up auto-tuning suffered from yet another well-known problem: lack of large and diverse benchmarks. Though Collective Mind helps share multiple programs and data sets including ones from [16,31,32], it may still not be enough to cover all possible pro-

```

2mm.c / 2mm.c
...
#ifdef OPENME
#include <openme.h>
#endif
...

int main(void) {
...
#ifdef OPENME
openme_init(NULL,NULL,NULL,0);
openme_callback("PROGRAM_START", NULL);
#endif
...

#ifdef OPENME
openme_callback("SELECT_KERNEL", &adapt);
#endif
}

#ifdef OPENME
openme_callback("KERNEL_START", NULL);
#endif

if (adaptive_select==0) mm2_cpu(A, B, C, D, E);
elif (adaptive_select==1) cl_launch_kernel(A,B,C,D,E);
elif (adaptive_select==2) mm2Cuda(A, B, C, D, E,
                                E_outputFromGpu);

#ifdef OPENME
openme_callback("KERNEL_END", NULL);
#endif
...

#ifdef OPENME
openme_callback("PROGRAM_END", NULL);
#endif
}

```

Fig. 14. Example of predictive scheduling of matrix–matrix multiplication kernel for heterogeneous architectures using OpenME interface and statically generated kernel clones with different algorithm implementations and optimizations to find the winning one at run-time. (Colors are visible in the online version of the article; <http://dx.doi.org/10.3233/SPR-140396>.)

gram behavior and features. One possibility is to generate many synthetic benchmarks and data sets but it always result in explosion in tuning and training times. Instead, we propose to use Alchemist plugin [30] together with plugin-enabled compilers such as GCC to use existing benchmarks, kernels and even data sets as templates and randomly modify them by removing, modifying or adding various instructions, basic blocks, loops and so on. Naturally, we can ignore crashing variants of the code and continue evolving only the working ones.

We can use such an approach not only to gradually extend realistic training sets, but also to iteratively identify various behavior anomalies or detect missing code feature to explain unexpected behavior similar to differential analysis from electronics [47,66]. For example, we are adding support to Alchemist plugin to iteratively scalarize memory accesses to characterize code and data set as CPU or memory bound [28,36]. Its prototype was used to obtain line X in Fig. 9 showing ideal code behavior when all floating point memory accesses are NOPed. Additionally, we use Alchemist plugin to unify extraction of code structure, patterns and other features to collaboratively improve prediction during software/hardware co-design [23].

5. Enabling and improving reproducibility of experimental results

Since cM allows to implement, preserve and share the whole experimental setup, it can also be used for reproducible research and experimentation. For example, unified module invocation in cM makes it pos-

sible to reproduce (replay) any experiment by saving JSON input for a given module and an action, and comparing JSON output. At the same time, since execution time and other characteristics often vary, we developed and shared cM module that applies Shapiro–Wilk test from R to test monitored characteristic for normality. However, in contrast with current experimental methodologies where results not passing such test are simply skipped, we record them in a reproducible way to find and explain missing features in the system. For example, when analyzing multiple executions of image corner detection benchmark on a smart phone shown in Fig. 8, we noticed an occasional $4\times$ difference in execution times as shown in Fig. 15. Simple analysis showed that our phone was often in the low power state at the beginning of experiments and then gradually switched to the high-frequency state ($4\times$ difference in frequency). Though relatively obvious, this information allowed us to add CPU frequency to the build and run pipeline using *cpufreq* module and thus separate such experiments. Therefore, Collective Mind research methodology can gradually improve reproducibility as a side effect and with the help of the community rather than trying to somehow enforce it from the start.

6. Conclusions and future work

This paper presents our novel, community-driven approach to make auto-tuning practical and move it to mainstream production environments. However, rather than searching for yet another “holy grail” auto-tuning technique, we propose to start preserving, sharing and

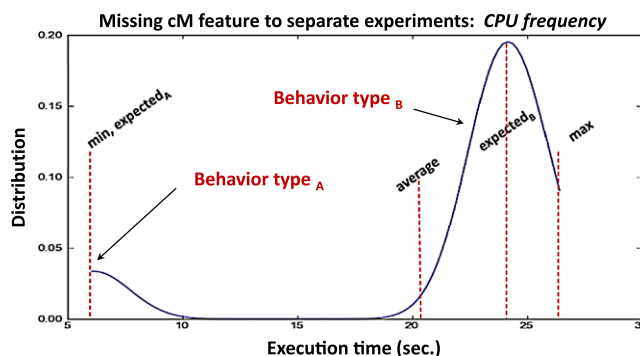


Fig. 15. Unexpected behavior helped to identify and add a missing feature to cM (processor frequency) as well as software dependency (cpufreq) that ensures reproducibility of experimental results. (Colors are visible in the online version of the article; <http://dx.doi.org/10.3233/SPR-140396>.)

reusing already available practical knowledge and experience about program optimization and hardware co-design using Collective Mind framework and repository. Such approach helps researchers and engineers quickly prototype and validate various auto-tuning and learning techniques as plugins connected into experimental pipelines while reusing all shared artifacts. Such pipelines can be distributed among many users to collaboratively learn, model and tune program behavior using standard top-down methodology from mature sciences such as physics by decomposing complex software into interconnected components while capturing first coarse-grain effects and later move to finer-grain levels. At the same time, any unexpected behavior and optimization mispredictions are exposed to the community in a reproducible way to be explained and improved. Therefore, we can collaboratively search for profitable optimizations, efficient auto-tuning strategies, truly representative benchmarks, and most accurate models to predict optimizations together with minimal set of relevant semantic and dynamic features.

Our future collaborative work includes exposing more tuning dimensions, characteristics and features using Collective Mind and Periscope tuning frameworks to eventually tune the whole computer system while extrapolating collected knowledge to build faster, more power efficient and reliable self-tuning computer systems. We are working with the community to gradually unify existing techniques and tools including pragma-based source-to-source transformations [41,80], plugin-based GCC and LLVM to expose and tune all internal optimization decisions [30,31]; polyhedral source-to-source transformation tools [12]; differential analysis to detect performance anomalies and CPU/memory bounds [28,36]; just-in-time compilation for Android Dalvik or Oracle JDK; algorithm-level tuning [3]; techniques to balance communication

and computation in numerical codes particularly for heterogeneous architectures [7,75]; Scalasca framework to automate analysis and modeling of scalability of HPC applications [13,40]; LIKWID for lightweight collection of hardware counters [76]; HPCC and HPCG benchmarks to collaboratively rank HPC systems [42,56]; benchmarks from GCC and LLVM, TAU performance tuning framework [68]; and all recent Periscope application tuning plugins [10,60].

At the same time we plan to use collected and unified knowledge to improve our past techniques on decomposition of complex programs into interconnected kernels, predictive modeling of program behavior, and run-time tuning and adaptation [5,9,17,34,46,54,55,69,82]. Finally, we are extending Collective Mind to assist recent initiatives on reproducible research and new publication models in computer engineering where all experimental results and related research artifacts with all dependencies are continuously shared along with publications to be validated and improved by the community [35].

Acknowledgements

We would like to thank David Kuck, David Wong, Abdul Memon, Yuriy Kashnikov, Michael Pankov, Siegfried Benkner, David Padua, Olivier Zendra, Arnaud Legrand, Sascha Hunold, Jack Davidson, Bruce Childers, Alex K. Jones, Daniel Mosse, Jean-Luc Gaudiot, Christian Poli, Egor Pasko, Francois Bodin, Thomas Fahringer, I-hsin Chung, Paul Hovland, Prasanna Balaprakash, Stefan Wild, Hal Finkel, Bernd Mohr, Arutyun Avetisyan, Anton Korzh, Andrey Slepudin, Vladimir Voevodin, Christophe Guillon, Christian Bertin, Antoine Moynault, Francois De-Ferriere, Reiji Suda, Takahiro Katagiri, Weichung

Wang, Chengyong Wu, Marisa Gil, Lasse Natvig, Nacho Navarro, Vittorio Zaccaria, Chris Fensch, Ayal Zaks, Pooyan Dadvand, Paolo Faraboschi, Ana Lucia Varbanescu, Cosmin Oancea, Petar Radojkovic, Christian Collberg, Olivier Temam and Panos Tsarchopoulos, as well as cTuning, Collective Mind, Periscope and HiPEAC communities for interesting related discussions and feedback. We are also very grateful to anonymous reviewers for their insightful comments.

References

- [1] B. Aarts et al., OCEANS: Optimizing compilers for embedded applications, in: *Proc. Euro-Par 97*, Lecture Notes in Computer Science, Vol. 1300, 1997, pp. 1351–1356.
- [2] 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: *Proceedings of the International Symposium on Code Generation and Optimization (CGO)*, 2006.
- [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: *Proceedings of the 2009 ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI’09*, ACM, New York, NY, USA, 2009, pp. 38–49.
- [4] K. Asanovic et al., The landscape of parallel computing research: a view from Berkeley, Technical Report UCB/EECS-2006-183, Electrical Engineering and Computer Sciences, University of California at Berkeley, December 2006.
- [5] C. Augonnet, S. Thibault, R. Namyst and P.-A. Wacrenier, Starpu: A unified platform for task scheduling on heterogeneous multicore architectures, *Concurr. Comput.: Pract. Exper.* **23**(2) (2011), 187–198.
- [6] M. Baboulin, D. Becker and J. Dongarra, A parallel tiled solver for dense symmetric indefinite systems on multicore architectures, in: *Parallel Distributed Processing Symposium (IPDPS), 2012 IEEE 26th International*, May 2012, pp. 14–24.
- [7] M. Baboulin, S. Donfack, J. Dongarra, L. Grigori, A. Rémy and S. Tomov, A class of communication-avoiding algorithms for solving general dense linear systems on CPU/GPU parallel machines, in: *ICCS*, 2012, pp. 17–26.
- [8] D.H. Bailey et al., Peri auto-tuning, *Journal of Physics: Conference Series (SciDAC 2008)* **125** (2008), 6 pp.
- [9] P. Balaprakash, S.M. Wild and P.D. Hovland, Can search algorithms save large-scale automatic performance tuning?, *Proceedia Computer Science* **4** (2011), 2136–2145 (Proceedings of the International Conference on Computational Science, ICCS 2011).
- [10] S. Benedict, V. Petkov and M. Gerndt, Periscope: An online-based distributed performance analysis tool, in: *Tools for High Performance Computing 2009*, 2010, pp. 1–16.
- [11] C.M. Bishop, *Pattern Recognition and Machine Learning (Information Science and Statistics)*, 1st edn, Springer, 2006. (Corr. 2nd printing 2011 edition, October 2007.)
- [12] U. Bondhugula, A. Hartono, J. Ramanujam and P. Sadayappan, A practical automatic polyhedral program optimization system, in: *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, June 2008.
- [13] A. Calotoiu, T. Hoefler, M. Poke and F. Wolf, Using automated performance modeling to find scalability bugs in complex codes, in: *SC*, 2013, p. 45.
- [14] J. Cavazos, G. Fursin, F. Agakov, E. Bonilla, M. O’Boyle and O. Temam, Rapidly selecting good compiler optimizations using performance counters, in: *Proceedings of the International Symposium on Code Generation and Optimization (CGO)*, March 2007.
- [15] V.G. Cerf, Where is the science in computer science?, *Communications of the ACM* **55**(10) (2012), 5.
- [16] Y. Chen, L. Eeckhout, G. Fursin, L. Peng, O. Temam and C. Wu, Evaluating iterative optimization across 1000 data sets, in: *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, 2010.
- [17] Y. Chen, Y. Huang, L. Eeckhout, G. Fursin, L. Peng, O. Temam and C. Wu, Evaluating iterative optimization across 1000 data sets, in: *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, June 2010.
- [18] Collective Mind: open-source plugin-based infrastructure and repository for systematic and collaborative research, experimentation and management of large scientific data, available at: <http://c-mind.org>.
- [19] Collective Mind Live Repo: public repository of knowledge about design and optimization of computer systems, available at: <http://c-mind.org/repo>.
- [20] Collective Mind Node: deploying experiments on Android-based mobile computer systems, available at: https://play.google.com/store/apps/details?id=com.collective_mind.node.
- [21] K.D. Cooper, P.J. Schielke and D. Subramanian, Optimizing for reduced code space using genetic algorithms, in: *Proceedings of the Conference on Languages, Compilers, and Tools for Embedded Systems (LCTES)*, 1999, pp. 1–9.
- [22] J. Dongarra et al., The International Exascale Software Project roadmap, *Int. J. High Perform. Comput. Appl.* **25**(1) (2011), 3–60.
- [23] C. Dubach, T.M. Jones, E.V. Bonilla, G. Fursin and M.F.P. O’Boyle, Portable compiler optimization across embedded programs and microarchitectures using machine learning, in: *Proceedings of the IEEE/ACM International Symposium on Microarchitecture (MICRO)*, December 2009.
- [24] ElasticSearch: open source distributed real time search and analytics, available at: <http://www.elasticsearch.org>.
- [25] T.W. Epps and L.B. Pulley, A test for normality based on the empirical characteristic function, *Biometrika* **70**(3) (1983), 723–726.
- [26] D. Ferrucci et al., Building Watson: An overview of the DeepQA project, *AI Magazine* **31**(3) (2010), 59–79.
- [27] B. Franke, M.F.P. O’Boyle, J. Thomson and G. Fursin, Probabilistic source-level optimisation of embedded programs, in: *Proceedings of the Conference on Languages, Compilers, and Tools for Embedded Systems (LCTES)*, 2005.
- [28] G. Fursin, Iterative compilation and performance prediction for numerical applications, PhD thesis, University of Edinburgh, United Kingdom, 2004.
- [29] G. Fursin, Collective tuning initiative: automating and accelerating development and optimization of computing systems, in: *Proceedings of the GCC Developers’ Summit*, June 2009.

- [30] G. Fursin, Collective Mind: cleaning up the research and experimentation mess in computer engineering using crowdsourcing, big data and machine learning, *CoRR*, abs/1308.2410, 2013.
- [31] G. Fursin et al., MILEPOST GCC: Machine learning enabled self-tuning compiler, *International Journal of Parallel Programming* **39** (2011), 296–327.
- [32] G. Fursin, J. Cavazos, M. O’Boyle and O. Temam, MiDataSets: Creating the conditions for a more realistic evaluation of iterative optimization, in: *Proceedings of the International Conference on High Performance Embedded Architectures & Compilers (HiPEAC 2007)*, January 2007.
- [33] G. Fursin and A. Cohen, Building a practical iterative interactive compiler, in: *1st Workshop on Statistical and Machine Learning Approaches Applied to Architectures and Compilation (SMART’07)*, Colocated with HiPEAC 2007 Conference, January 2007.
- [34] G. Fursin, A. Cohen, M. O’Boyle and O. Temam, A practical method for quickly evaluating program optimizations, in: *Proceedings of the International Conference on High Performance Embedded Architectures & Compilers (HiPEAC 2005)*, November 2005, pp. 29–46.
- [35] G. Fursin and C. Dubach, Experience report: community-driven reviewing and validation of publications, in: *Proceedings of the 1st Workshop on Reproducible Research Methodologies and New Publication Models in Computer Engineering (ACM SIGPLAN TRUST’14)*, ACM, 2014.
- [36] G. Fursin, M. O’Boyle, O. Temam and G. Watts, Fast and accurate method for determining a lower bound on execution time, *Concurrency Comput.: Practice and Experience* **16**(2,3) (2004), 271–292.
- [37] G. Fursin and O. Temam, Collective optimization: A practical collaborative approach, *ACM Transactions on Architecture and Code Optimization (TACO)* **7**(4) (2010), 20:1–20:29.
- [38] G.G. Fursin, M.F.P. O’Boyle and P.M.W. Knijnenburg, Evaluating iterative compilation, in: *Proceedings of the Workshop on Languages and Compilers for Parallel Computers (LCPC)*, 2002, pp. 305–315.
- [39] J.J. Garrett, Ajax: A new approach to web applications, available at: <http://www.adaptivepath.com/ideas/ajax-new-approach-web-applications>.
- [40] M. Geimer, F. Wolf, B.J.N. Wylie, E. Ábrahám, D. Becker and B. Mohr, The Scalasca performance toolset architecture, *Concurr. Comput.: Pract. Exper.* **22**(6) (2010), 702–719.
- [41] A. Hartono, B. Norris and P. Sadayappan, Annotation-based empirical performance tuning using Orio, in: *IEEE International Symposium on Parallel Distributed Processing, 2009. IPDPS 2009*, 2009, pp. 1–11.
- [42] M.A. Heroux and J. Dongarra, Toward a new metric for ranking high performance computing systems, June 2013, available at: <http://0-www.osti.gov.iii-server.ualr.edu/scitech/servlets/purl/1089988>.
- [43] G.E. Hinton and S. Osindero, A fast learning algorithm for deep belief nets, *Neural Computation* **18** (2006), 2006, 1527–1554.
- [44] K. Hoste and L. Eeckhout, Cole: Compiler optimization level exploration, in: *Proceedings of the International Symposium on Code Generation and Optimization (CGO)*, 2008.
- [45] Y. Huang, L. Peng, C. Wu, Y. Kashnikov, J. Renneke and G. Fursin, Transforming GCC into a research-friendly environment: plugins for optimization tuning and reordering, function cloning and program instrumentation, in: *2nd International Workshop on GCC Research Opportunities (GROW)*, Colocated with HiPEAC’10 Conference, January 2010.
- [46] V. Jimenez, I. Gelado, L. Vilanova, M. Gil, G. Fursin and N. Navarro, Predictive runtime code scheduling for heterogeneous architectures, in: *Proceedings of the International Conference on High Performance Embedded Architectures & Compilers (HiPEAC 2009)*, January 2009.
- [47] Y. Jin, Fuzzy modeling of high-dimensional systems: complexity reduction and interpretability improvement, *IEEE Transactions on Fuzzy Systems* **8**(2) (2000), 212–221.
- [48] T. Kisuki, P.M.W. Knijnenburg and M.F.P. O’Boyle, Combined selection of tile sizes and unroll factors using iterative compilation, in: *Proceedings of the International Conference on Parallel Architectures and Compilation Techniques (PACT)*, 2000, pp. 237–246.
- [49] P. Kulkarni, W. Zhao, H. Moon, K. Cho, D. Whalley, J. Davidson, M. Bailey, Y. Paek and K. Gallivan, Finding effective optimization phase sequences, in: *Proceedings of the Conference on Languages, Compilers, and Tools for Embedded Systems (LCTES)*, 2003, pp. 12–23.
- [50] H.T. Kung, F. Luccio and F.P. Preparata, On finding the maxima of a set of vectors, *J. ACM* **22**(4) (1975), 469–476.
- [51] C. Lattner and V. Adve, LLVM: A compilation framework for lifelong program analysis & transformation, in: *Proceedings of the 2004 International Symposium on Code Generation and Optimization (CGO’04)*, Palo Alto, CA, March 2004.
- [52] Q. Le, M.A. Ranzato, R. Monga, M. Devin, K. Chen, G. Corrado, J. Dean and A. Ng, Building high-level features using large scale unsupervised learning, in: *International Conference in Machine Learning*, 2012.
- [53] J. Lu, H. Chen, P.-C. Yew and W.-C. Hsu, Design and implementation of a lightweight dynamic optimization system, *Journal of Instruction-Level Parallelism* **6** (2004), 24 pp.
- [54] C.-K. Luk, S. Hong and H. Kim, Qilin: exploiting parallelism on heterogeneous multiprocessors with adaptive mapping, in: *Proceedings of the 42nd Annual IEEE/ACM International Symposium on Microarchitecture, MICRO 42*, ACM, New York, NY, USA, 2009, pp. 45–55.
- [55] L. Luo, Y. Chen, C. Wu, S. Long and G. Fursin, Finding representative sets of optimizations for adaptive multiversioning applications, in: *3rd Workshop on Statistical and Machine Learning Approaches Applied to Architectures and Compilation (SMART’09)*, Colocated with HiPEAC’09 Conference, January 2009.
- [56] P.R. Luszczek, D.H. Bailey, J.J. Dongarra, J. Kepner, R.F. Lucas, R. Rabenseifner and D. Takahashi, The HPC challenge (HPCC) benchmark suite, in: *Proceedings of the 2006 ACM/IEEE Conference on Supercomputing, SC’06*, ACM, New York, NY, USA, 2006.
- [57] G. Marin and J. Mellor-Crummey, Cross-architecture performance predictions for scientific applications using parameterized models, *SIGMETRICS Perform. Eval. Rev.* **32**(1) (2004), 2–13.
- [58] J. Mars and R. Hundt, Scenario based optimization: A framework for statically enabling online optimizations, in: *CGO*, 2009, pp. 169–179.
- [59] F. Matteo and S. Johnson, FFTW: An adaptive software architecture for the FFT, in: *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing*, Vol. 3, Seattle, WA, May 1998, pp. 1381–1384.

- [60] R. Miceli et al., Autotune: A plugin-driven approach to the automatic tuning of parallel applications, in: *Proceedings of the 11th International Conference on Applied Parallel and Scientific Computing, PARA'12*, Springer-Verlag, Berlin/Heidelberg, 2013, pp. 328–342.
- [61] MILEPOST project archive (MachIne Learning for Embedded PrOgramS opTimization), available at: <http://cTuning.org/project-milepost>.
- [62] A. Monsifrot, F. Bodin and R. Quiniou, A machine learning approach to automatic production of compiler heuristics, in: *Proceedings of the International Conference on Artificial Intelligence: Methodology, Systems, Applications*, LNCS, Vol. 2443, 2002, pp. 41–50.
- [63] Online JSON introduction, available at: <http://www.json.org>.
- [64] Z. Pan and R. Eigenmann, Fast and effective orchestration of compiler optimizations for automatic performance tuning, in: *Proceedings of the International Symposium on Code Generation and Optimization (CGO)*, 2006, pp. 319–332.
- [65] PRACE: Partnership for Advanced Computing in Europe, available at: <http://www.prace-project.eu>.
- [66] H. Roubos and M. Setnes, Compact and transparent fuzzy models and classifiers through iterative complexity reduction, *IEEE Transactions on Fuzzy Systems* **9**(4) (2001), 516–524.
- [67] J. Shen, A.L. Varbanescu, H.J. Sips, M. Arntzen and D.G. Simons, Glinda: a framework for accelerating imbalanced applications on heterogeneous platforms, in: *Conf. Computing Frontiers*, 2013, p. 14.
- [68] S.S. Shende and A.D. Malony, The Tau parallel performance system, *Int. J. High Perform. Comput. Appl.* **20**(2) (2006), 287–311.
- [69] Y. Shoham and K. Leyton-Brown, *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*, Cambridge Univ. Press, New York, NY, USA, 2008.
- [70] B. Singer and M. Veloso, Learning to predict performance from formula modeling and training data, in: *Proceedings of the Conference on Machine Learning*, 2000.
- [71] M. Stephenson and S. Amarasinghe, Predicting unroll factors using supervised classification, in: *Proceedings of the International Symposium on Code Generation and Optimization (CGO)*, IEEE Computer Society, 2005.
- [72] M. Stephenson, S. Amarasinghe, M. Martin and U.-M. O'Reilly, Meta optimization: Improving compiler heuristics with machine learning, in: *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI'03)*, June 2003, pp. 77–90.
- [73] 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, Supercomputing'02*, IEEE Computer Society Press, Los Alamitos, CA, USA, 2002, pp. 1–11.
- [74] The HiPEAC vision on high-performance and embedded architecture and compilation (2012–2020), 2012, available at: <http://www.hipeac.net/roadmap>.
- [75] S. Tomov, J. Dongarra and M. Baboulin, Towards dense linear algebra for hybrid GPU accelerated manycore systems, *Parallel Comput.* **36**(5,6) (2010), 232–240.
- [76] J. Treibig, G. Hager and G. Wellein, Likwid: A lightweight performance-oriented tool suite for x86 multicore environments, *CoRR*, abs/1004.4431, 2010.
- [77] M.J. Voss and R. Eigenmann, ADAPT: Automated de-coupled adaptive program transformation, in: *Proceedings of International Conference on Parallel Processing*, 2000.
- [78] R. Whaley and J. Dongarra, Automatically tuned linear algebra software, in: *Proceedings of the Conference on High Performance Networking and Computing*, 1998.
- [79] S. Williams, A. Waterman and D. Patterson, Roofline: an insightful visual performance model for multicore architectures, *Commun. ACM* **52**(4) (2009), 65–76.
- [80] Q. Yi, K. Seymour, H. You, R. Vuduc and D. Quinlan, POET: Parameterized optimizations for empirical tuning, in: *Proceedings of the Workshop on Performance Optimization of High-Level Languages and Libraries (POHLL)*, Co-located with *IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, 2007.
- [81] M. Zhao, B.R. Childers and M.L. Soffa, A model-based framework: an approach for profit-driven optimization, in: *Third Annual IEEE/ACM International Conference on Code Generation and Optimization*, 2005, pp. 317–327.
- [82] S. Zuckerman, J. Suetterlein, R. Knauerhase and G.R. Gao, Using a “codelet” program execution model for exascale machines: position paper, in: *Proceedings of the 1st International Workshop on Adaptive Self-Tuning Computing Systems for the Exaflop Era, EXADAPT'11*, ACM, New York, NY, USA, 2011, pp. 64–69.



Hindawi

Submit your manuscripts at
<http://www.hindawi.com>

