

Parameter Tuning for Self-optimizing Software at Scale

Dmytro Pukhkaiev

Software Technology Group, Technische Universität
Dresden, Germany
dmytro.pukhkaiev@tu-dresden.de

Uwe Aßmann

Software Technology Group, Technische Universität
Dresden, Germany
uwe.assmann@tu-dresden.de

ABSTRACT

Efficiency of self-optimizing systems is heavily dependent on their optimization strategies, e.g., choosing exact or approximate solver. A choice of such a strategy, in turn, is influenced by numerous factors, such as re-optimization time, size of the problem, optimality constraints, etc. Exact solvers are domain-independent and can guarantee optimality but suffer from scaling, while approximate solvers offer a “good-enough” solution in exchange for a lack of generality and parameter-dependence. In this paper we discuss the trade-offs between exact and approximate optimizers for solving a quality-based software selection and hardware mapping problem from the scalability perspective. We show that even a simple heuristic can compete with thoroughly developed exact solvers under condition of an effective parameter tuning. Moreover, we discuss robustness of the obtained algorithm’s configuration. Last but not least, we present a software product line for parameter tuning, which comprise the main features of this process and can serve as a platform for further research in the area of parameter tuning.

CCS CONCEPTS

• **Software and its engineering** → **Search-based software engineering**; *Software product lines*; • **Theory of computation** → *Simulated annealing*; • **General and reference** → Empirical studies; • **Applied computing** → Multi-criterion optimization and decision-making.

KEYWORDS

parameter tuning, software product lines, search-based software engineering, simulated annealing, optimization, active learning

1 INTRODUCTION

Combinatorial optimization problems are repeatedly occurring in all spheres of life. Navigation and route planning for autonomic vehicles [30], planning and scheduling of production tasks [25] and sport events [21], recommender systems [20] and many more. Naturally, the decades-long attention to these topic brought numerous approaches and algorithms for solving one or another problem from this list.

As a running example for this paper we use a problem of quality-based software selection and hardware mapping [15]. It is an important problem occurring in the field of self-optimizing software [14]. In short, it requires a satisfaction of functional and nonfunctional requirements of the users by serving their requests at a certain point in time. A request is specified in a form of a contract that requires a variant of a software component providing a certain functionality. The software components themselves can require other software components and a hardware instance to be deployed on. The solution is an optimal (w.r.t. a quality measure, e.g., overall energy consumption) mapping of software variants onto the available hardware, whilst satisfying all contracts.

This problem was used as a case of the Transformation Tool Contest 2018¹ (TTC), where several promising solutions were presented. The results of TTC’18 are very expressive as the best two approaches are diametrically different by their nature: a heuristic-based ant-colony optimization approach [17] and an exact integer linear programming-based (ILP) approach [16]. The result continues a decades-long competition between exact and approximate solvers [24]. The general claim is that “good enough” heuristics have much better scaling factor than exact solvers. However, the recent advancements in computational resources made exact solvers very attractive for tackling optimization problems. For example, in [16] the authors showed the superiority of an ILP-based solution to the heuristic [17] in terms of scalability; thus, arguing the utilization of heuristics.

In this paper we aim to back up the usage of approximate solvers for combinatorial optimization problems. The main problem of such approaches lies in costly and often not obvious parameter optimization (tuning) [5]. Choosing an appropriate parameter configuration can significantly improve performance of an algorithm. Therefore, we present BRISE 2, a software product line (SPL) for parameter tuning, which effectively reduces the effort spent on optimizing the algorithm’s parameters. Its idea was inspired by a successful application of a similar approach in the energy benchmarking area [26].

To show the power of the automatic parameter tuning we introduce a simple heuristic based on a simulated annealing algorithm [22] that randomly varies software and hardware components. We pick an initial configuration according to the established guidance, which is already a manual parameter tuning. BRISE 2 was able to further improve the solution’s quality by 11% (for larger problem sizes). We also show how and under which conditions a combination of parameter tuning with such a heuristic can be beneficial for solving the quality-based software selection and hardware mapping in comparison to the exact ILP-based solution [16].

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).

MoST-Rec’19, 2019, Beijing, China

© 2019 Copyright held by the owner/author(s).

¹<https://www.transformation-tool-contest.eu/2018/>

2 RESEARCH QUESTIONS

The goal of this paper is to motivate the utilization of heuristics for the problem of quality-based software selection and hardware mapping. The research objective is to determine the conditions under which a utilization of a heuristic-based solution can be beneficial. To reach the research objective we need to answer the following research questions.

- RQ1.** What is the effect of parameter tuning on the solution quality of the approximate algorithm?
- RQ2.** Is it possible for a tuned heuristic to provide a comparable solution to the ILP-based solver?
- RQ3.** What is the scalability of the heuristic in comparison to ILP-based solver?

3 QUALITY-BASED SOFTWARE SELECTION AND HARDWARE MAPPING AND ITS SOLVERS

3.1 Problem definition

Let us discuss the problem domain of the paper: the problem of quality-based software selection and hardware mapping [15]. This optimization problem is twofold. It presumes finding an optimal serving of incoming requests by the means of mapping the corresponding software components (SWC) onto available resources, hardware components (HWC) w.r.t. energy-efficiency.

Assume a scenario where the user requires a video service with a specified non-functional properties (NFP) such as, e.g., a minimal resolution (1440 p) and a frame rate (60 fps). The required functionality can be provided by a core SWC: VIDEOPLAYER. Each SWC can, in turn, have a set of different *implementations* such as, e.g. VLC (Video LAN Client) and QT (Quicktime). Moreover, a SWC can require other SWCs, e.g., a VIDEOPLAYER SWC requires a DECODER and DATA PROVIDER. A HWC is a hierarchical structure, where a component *server* comprises several sub-components: *CPUs, RAM, disk, network*.

The optimal solution consists of a selection of the best SWC implementations and their best mapping to the available HWCs for a given set of requests.

3.2 Exact and approximate solvers

The classic solution to the aforementioned problem is: Multi-Quality Auto-Tuning (MQuAT) [14]. It is an approach for model-driven, component-based development of self-adaptive systems. The key idea is a utilization of model transformation from a problem specification to a reasoning logic. The problem specification is described with a domain specific language, while the reasoning logic is represented in a form of an ILP.

MQuAT is an example of a two-phased optimization approach with two distinct parts: domain- and solver-specific. In this case the problem modeling is done by a domain expert in domain-specific part, whilst after the transformation the described model is domain independent and can be solved by a standard constraint solver.

Heuristic-based solvers are the algorithms that are situated solely in the domain-specific part. They are enriched with domain knowledge and therefore are able to effectively find a valid solution. However, they do not guarantee optimality and in the long run are often losing to the exact solvers in terms of solution's quality.

3.3 Metaheuristic-based solver

In this paper we present a simple heuristic based on the simulated annealing meta-heuristic that solves the stated problem. It is developed with help of OptaPlanner², a lightweight, embeddable planning engine. The choice of an underlying meta-heuristic is very important for the resulting quality of the solver. However, in this paper we want to show that even a straightforward heuristic can deliver solutions comparable and even outperform thoroughly developed ILP-based solvers. Thus, we leave a question of heuristic selection to specialized works [2, 6, 31]. First, let us discuss the main concepts of the approach.

Figure 1 depicts a class diagram of the heuristic. COMPONENTASSIGNMENT is a placeholder for interchangeable variables, which in context of our case study are HWCs and SWCs. Each COMPONENTASSIGNMENT contains exactly one SWC-HWC pair and the REQUEST it corresponds to. A list of COMPONENTASSIGNMENTS forms an ALLOCATION, the optimization of which is our final goal. Note, there are several unintuitive design decisions that were made in order to make the problem model compatible with OptaPlanner. For example, from the problem description we know that different implementations of the same SWC can have an arbitrary number of requested components. Therefore, the overall number of COMPONENTASSIGNMENTS can vary depending on the current choice of the solution. OptaPlanner, however, needs to create all planning entities in advance to manipulate them during the search. Hence, we created the maximum number of COMPONENTASSIGNMENTS and enable or disable them depending on component requirements of a current solution.

Each iteration of any local search (LS) is based on two main notions: *move* and *score calculation*. Move is a change of a planning variable that provides a new solution, i.e., explores the search space. Score calculation is necessary to assess the quality of the obtained solution and, thus, helps to decide which move to take.

There are three types of moves in our MH: HWC-CHANGE MOVE, HWC-SWAP MOVE and SWC-CHANGE MOVE. Change move means a random substitution of an assigned HWC or SWC by a different one, while swap move switches an HWC of one COMPONENTASSIGNMENT by the content of the another one and vice versa. An SWC-CHANGE MOVE is a bit more complicated and deserves additional details. After changing a random SWC at a random level of hierarchy it causes a modification of requirements for the dependent COMPONENTASSIGNMENTS, removal of the obsolete assignments and addition of the new ones. Note, that the addition of the new assignments implies not a creation of new SWC-HWC pair, but unblocking an existent COMPONENTASSIGNMENT, thus, allowing to utilize them at the current stage of the LS. Figure 2 shows a class diagram of possible moves.

After the move is performed, the score calculation takes place. We iterate through all COMPONENTASSIGNMENTS and check whether

²<https://www.optaplanner.org/>

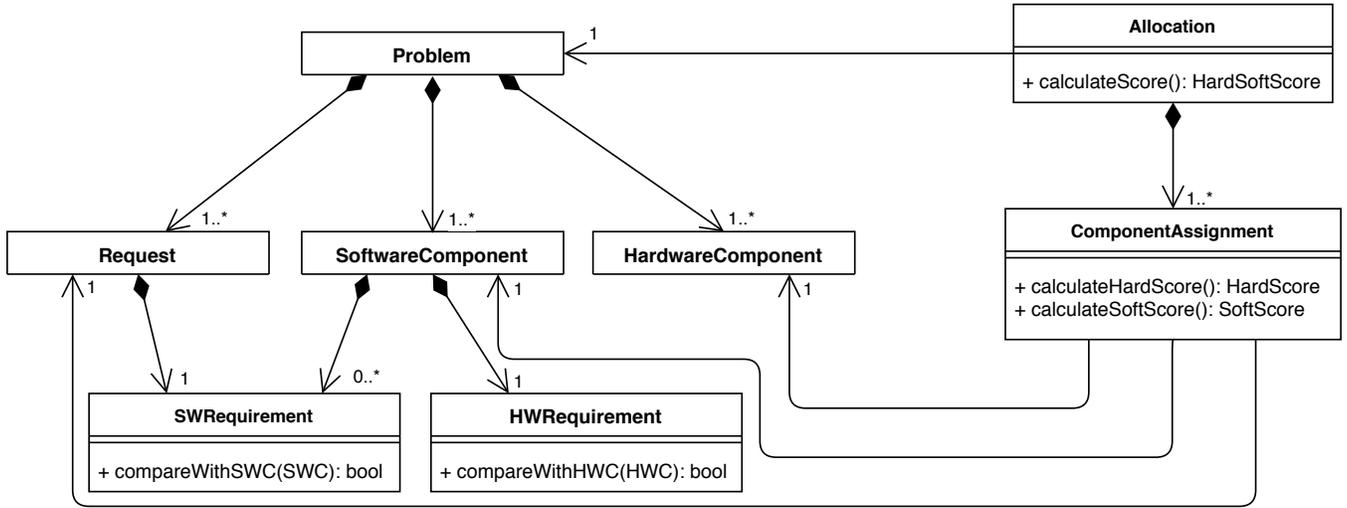


Figure 1: Class diagram of our metaheuristic-based solver

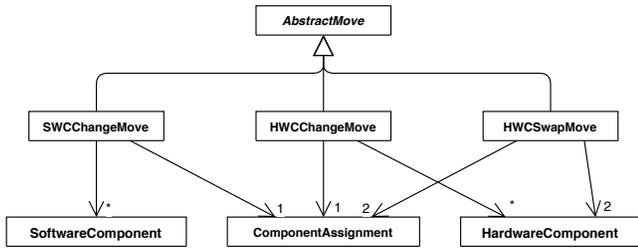


Figure 2: Class diagram of the possible moves

the underlying SWCs and HWCs are complying with the stated requirements. We utilize a so-called “hardsoft” score³ in our MH. It allows to decouple the objective function (energy consumption) from other requirements while comparing two solutions. If any constraint, such as missing or unsatisfying qualities of an SWC or HWC, is violated, the hard score is increasing. The soft score summarizes the energy consumption of all HWCs executing picked SWCs. The optimal score looks as follows:

$$0 \text{ hard} / \min(\text{energy_consumption}) \text{ soft}. \quad (1)$$

Note, that the score 0 hard / 1000 soft is better than 1 hard / 10 soft as the hard score has a higher priority than the soft. By using such a scoring strategy we can ensure the knowledge whether the obtained solution is valid (serves all requests).

The employed simulated annealing algorithm, which provides a basis for our heuristic, has a strong dependence from its parameters. Picking the right parameter configuration can make the algorithm highly effective, while a poor choice can prevent search space exploration at all.

Our heuristic has the following input parameters:

- *subComponentUnassignedFactor*: Factor of influence for an unsatisfying HWC sub-component. The number of unsatisfying sub-components is multiplied by the chosen factor and added to the overall hard score. Possible values: [1, 2, 3, 5, 10, 100, 1000, 10000];
- *softwareComponentUnassignedFactor*: Factor of an unsatisfying SWC. The number of unsatisfying SWCs is multiplied by the chosen factor and added to the overall hard score. Possible values: [1, 2, 3, 5, 10, 100, 1000, 10000];
- *hardScoreStartingTemperature*: Maximal accepted hard score degradation of a new solution in comparison to the current solution. Is relative to the hard score of the initial solution. Possible values: [1, 2, 3, 5, 10, 20, 30, 50, 75, 100];
- *softScoreStartingTemperature*: Maximal accepted soft score degradation of a new solution in comparison to the current solution. Possible values: [1, 2, 3, 5, 10, 20, 30, 50, 75, 100];
- *neighborhoodSize*: Size of the neighborhood, i.e., how many neighboring solutions are being compared each iteration. Possible values: [1, 2, 5, 10, 20, 30, 40, 50].

The total number of configurations needed to be measured to find an optimal parameter combination equals 51200 configurations for a single problem size. Assume a testing time of just 10 seconds (in the TTC case it was 15 minutes) and a single-threaded execution the minimum required time effort equals to almost 6 days of continuous execution. In this paper we present an approach that allows to find a near-optimal configuration for the same problem size in 11 minutes, thus decreasing this effort by 99.9%. But first we need to discuss state-of-the-art approaches of parameter tuning.

4 PARAMETER TUNING

4.1 Background

Parameter tuning is a well-known problem arising in different areas, such as *machine learning* [12], where a selection of an optimal configuration can significantly influence quality of the production

³https://docs.optaplanner.org/7.21.0.Final/optaplanner-docs/html_single/index.html#hardSoftScore

phase, *search-based software engineering* [3], where an effective parameter configuration can dramatically speed-up search for an optimal solution and *benchmarking process* [26], where it can minimize execution time, energy consumption or maximize quality of an algorithm.

Approaches that aim to reduce the number of measurements needed to identify an optimal configuration can be divided into three major groups: factorial designs [9, 11], combinatorial optimization approaches [7, 32] and active learning approaches [1, 10, 26].

Fractional factorial design is an identification and removal of parameters that do not have much influence on the resulting quality [9]. It is similar to feature selection in machine learning, and even though it was originally developed for another goal (reduction the measurements' number), fractional factorial design is applicable in this area [29]. Being able to drastically reduce the search space in case of a high parameter number, these approaches lack flexibility while dealing with remaining parameter combinations, giving no guidance on their exploration. Thus, factorial designs should be complemented with some other strategy.

General purpose combinatorial optimization algorithms such as evolutionary algorithms or meta-heuristics are widely used in parameter tuning [7, 18, 32]. However, due to the nature of combinatorial optimization itself, they are suitable to only a limited subset of experiments [26]. The optimization problem should have a huge search space, while time of identifying a resulting value for a single configuration should be as small as possible. Such a scenario allows to compare a vast number of the neighbors to find a solution for the next step. Being applied to a long-running target algorithm, these approaches are unforgivably wasteful.

Active learning approaches such as [1, 8, 10, 13, 26] iteratively sample the search space and construct the model which is used to predict the next configuration to be measured, thus, keeping the number of measured configurations at each point of time as small as possible. These approaches have received the highest attraction during the recent years. We refer the interested reader to the respective papers for details and describe a generalized approach in this paper.

4.2 Generalization and Parameter Tuning SPL

A typical active learning approach looks as follows. After getting a search space from the user, the algorithm starts with *sampling of configurations*, this sampling can be random [1, 10, 13, 26] or adhere to some specific strategy [26]. After each measured configuration the *model* of the search space based on the measured configurations is created. It can be Bayesian [8, 10, 13], direct-acyclic-graph-based [8], linear-regression-based [26], etc. The process of building the model starts with a very small data set. Thus, it can be very inaccurate and one needs a metric to *validate* it [26]. It is based on a comparison of possible results for the task, specified by a user and of predicted results by the model. After the model is successfully validated, sampling is dropped and the model leads the search into the promising areas of the search space. Each newly found *solution candidate* should also be *validated* to determine the algorithm's termination [1, 13, 26], otherwise, a user-defined timeout is another popular stopping criteria [10].

Unlike its predecessor BRISE 1 [26], BRISE 2 uses this generalized flow as a basis, has multiple variants for each previously described feature, is easily extensible; and thus, is a full-fledged software product line for parameter tuning. Figure 3 depicts a high level architecture of the SPL.

BRISE 2 adheres to the client-server architecture and consists of three main components. *Main Node* is a core component of the framework, it is responsible for the whole flow of the application except the measurement process. *Worker* is a lightweight component, containing only the logic of the algorithm to be tuned (or its wrapper), which should be specified by the user. It gets a parameter combination to be tested as an input and outputs the quality metric to be later evaluated in the *Main Node*. The communication between the *Main Node* and *Workers* is performed via the *Worker Service*, which also manages liveness of the *Workers* and is responsible for the distribution of tasks. There are several additional components like *Front End* to visualize the parameter tuning process and *Benchmark* that automates the execution of multiple experiments and produces graphical reports out of these executions, but we leave them out of scope of this paper.

Now let us discuss the *Main Node* in more details. *Experiment* is a core entity. It stores the search space and measured *configurations*, parameters to be tuned and the configuration of the SPL itself.

Selection algorithm is used to cover the search space when there is no information on its structure or the model is unreliable. Out of the box we provide two selection algorithms: Sobol sampling [28] and Fedorov's exchange algorithm [11], but the user can extend the SPL with their own approaches. Each selection algorithm has to provide functionality on how to get the new configuration from the search space and how to disable a point after measuring it.

Model is being continuously built after each measurement phase iteration, it predicts new configurations to be measured and performs self-validation. Regression model is complemented by a Bayesian optimization model inspired by [10].

Repeater is a feature that automatically decides on the number of repetitions needed to obtain the required accuracy for each configuration. The basic idea is to decrease the number of repetitions for non-promising configurations, while measuring the important ones with higher preciseness. We have implemented several variants of repetition strategies. Quantity-based is a simplistic strategy, where the user specifies the number of repetitions to be performed for each configuration. Student-based repetition reduction is a strategy, that takes into account preciseness of measurements in terms of the relative measurement error. It is also possible to use knowledge available in the model to further reduce repetitions. Model-aware Student repeater does it manipulating the affordable measurement error. E.g., for those configurations which are considered not promising, the accuracy requirements are relaxed.

Stop condition is another mandatory feature which validates a solution received from the model and decides, whether to perform subsequent measurements or to stop the experiment. The overall Stop Condition can be combined from several stop criteria such as exceeding a specific number of measured configurations (quantity-based), percentage of the search space (adaptive) or time limit (time-based) specified by the user. The lack of the result's improvement

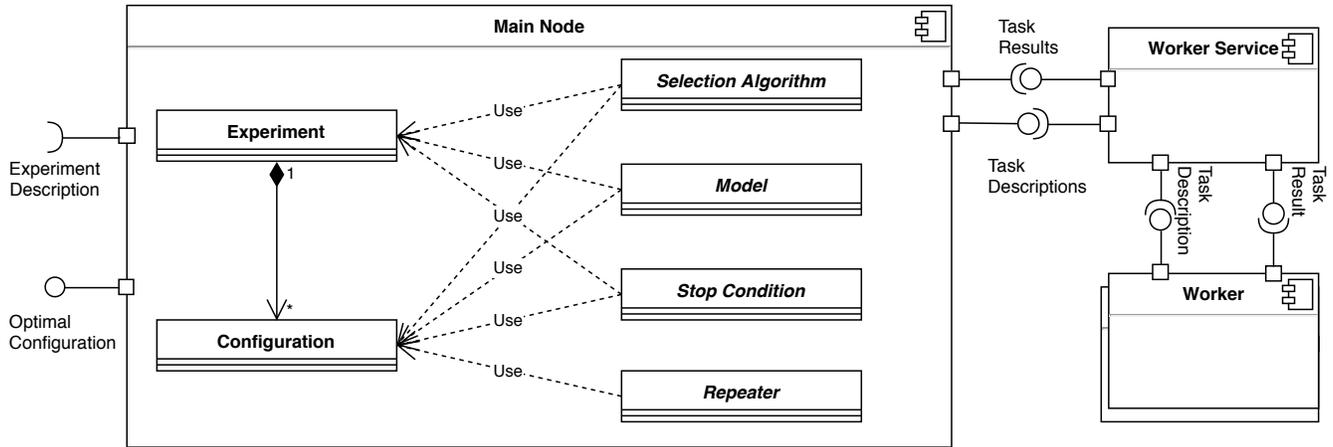


Figure 3: High-level architecture of the SPL

since finding the last new solution (improvement-based) or a guaranty of finding a configuration better than the starting one (guaranteed) are also possible stopping criteria. Figure 4 summarizes the features of the SPL and their variants.

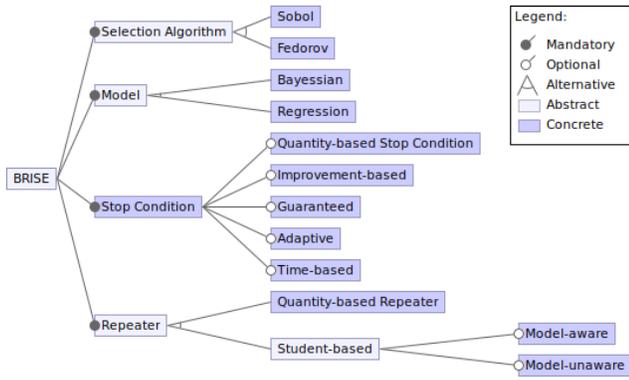


Figure 4: Feature tree of the SPL

5 EVALUATION

The goal of this paper is to motivate the usage of approximate solvers in combination with automatic parameter tuning based on a specific case study (quality-based software selection and hardware mapping). Therefore, we leave a full-fledged evaluation of the SPL itself out of scope of this paper, but discuss the effect it makes on the metaheuristic-based solver and compare both variants with the ILP-based solver.

We used the same ILP solver as the authors of the approach [15], GLPK 4.65. Note, that one can use a much more efficient commercial solver which can be also tuned for performance [19] like the MH-based solver. However, as will be shown later, the main drawback of ILP-based solvers is their costly generation time. Thus, the choice of the ILP solver does not impact our findings.

There exist numerous recommendations for picking a default configuration of simulated annealing [4, 23, 27], which are sometimes contradictory, e.g., in [23] the authors suggest high starting temperature, while in [27] the authors motivate usage of smaller temperature values. We decided to take a maximum possible initial temperature, so that the algorithm starts with global exploration of the search space and resides in local regions after getting familiar with its structure. Regarding the other parameters we have decided to start with a balanced influence of picking wrong software or hardware components' on the solution's score, 1 point for an unsatisfying HWC sub-component, 5 points for an unsatisfying SWC. We have also picked the largest allowed neighborhood size and pursue the steepest descent strategy. Thus, the default (manually tuned) configuration looks as follows:

- *subComponentUnassignedFactor*: 1;
- *softwareComponentUnassignedFactor*: 5;
- *hardScoreStartingTemperature*: 100;
- *softScoreStartingTemperature*: 100;
- *neighborhoodSize*: 50.

We configured our parameter tuning experiment in the following way. We picked Sobol selection algorithm, Bayesian model, improvement-based stop condition, which fires after 50 subsequent configurations without improvement and quantity-based repeater with 2 repetition for each configuration. For parameter tuning we decreased the timeout for running a single configuration to 10 seconds. This reduction has sped up the tuning (the overall process of parameter tuning for all use cases took 2.34 hours), but may have influenced the result's quality as some good configurations could have been cut out by the model. After the parameter tuning stage we returned to the maximum solving time of 15 minutes as it was in the TTC case. The measurements were performed on an Intel Core i7-8700 CPU machine with 64G of memory using Fedora Server 29 with GLPK 4.65 and Oracle Java 1.8.0 201.

Table 1 shows the results for the scenarios introduced in [15], using the ILP-based solver (GLPK 4.65, first value in each column), manually tuned MH-solver (second value in each column) and automatically tuned MH-solver (bold value). We measure the validity

and quality of the result, i.e., both hard and soft score, for a corresponding problem size for each approach. Note, that solution quality is taken as a relation to the optimal solution, which was obtained by the ILP-based solver without time constraints. Moreover, we compare the time to transform the problem into a suitable format ready to be solved by the ILP and MH solvers, respectively. We also track the timestamp of the first valid solution obtained by each solver and the last improvement of the solution. Note, that the ILP solver guarantees the optimality of the obtained solution and stops after it finds it, while MH solver always runs until the timeout.

Scenarios are separated into four main groups based on the number of requests: small (1 request), medium (10-15 requests), large (20 requests) and huge (50 requests). Moreover, each group has three variations: basic, increased number of hardware resources (*much hardware*) and increased depth of the software tree (*complex software*). The resulting numbers of available implementations and total number of resources are presented in Table 1. To get a deeper understanding of problem scenarios, the interested reader is referred to the TTC'18 case [15].

From Table 1 we can see that both ILP- and MH-based solvers are unable to solve problems 6, 9 and 12, i.e., the problems with complex SWC structure. In terms of the MH-based solver it is determined by the move strategy of the heuristic. A move changes a single SWC at a random position in the dependency tree; thus, increasing the hard score. As a complex SWC structure results in frequent SWC-moves, the temperature should be very high; henceforth, the heuristic behaves like a random search.

As can be seen from Table 1, for the smaller problem sizes 0-4 and 7 the utilization of ILP-based solver is preferential. Transformation time for these problems is very small (less than 2 seconds), whilst solving time is even shorter. Though the MH solver has almost negligible transformation time (< 0.1 s), it needs more time to find a valid solution and even more time for an optimal (problem sizes 4).

The problem size 5 is a border case, which shows the scalability behavior of both approaches. Here, the transformation time of ILP-based solver exceeds the first solution found timestamp of the MH-solver; thus, delivering the first valid solution later and is not able to deliver an optimal solution until the timeout. Although delivering a better final solution, the ILP-based solver loses to the MH-based in terms of getting the valid result. With problem size growth (problem sizes 8, 10, 11) we can see a development of this trend. Scenario 8 is the largest problem that can be transformed in the given time limit, but no solution is found, while for scenarios 10 and 11 the problem cannot be even transformed. On contrary, MH-based solver delivers valid solutions for all these scenarios.

To illustrate this scalability trend more vividly, we stabilized the number of requests, SWCs and their structure (4 requests each requesting a SWC chain of length 4) and scaled the number of HWCs. Figure 5 shows the improvement of solutions' quality found by MH- and ILP-based solvers in comparison to an optimal result for a selected problem size (2048 HWCs). Note, that solution's validity is normalized here, 1 means a valid solution, while 0 is the starting solution's validity. Solution's quality is taken in relation to the optimal solution. The blue line shows the change of validity of the MH solution, it instantly improves from the initial solution to an almost valid, with only several unsatisfied constraints. Being

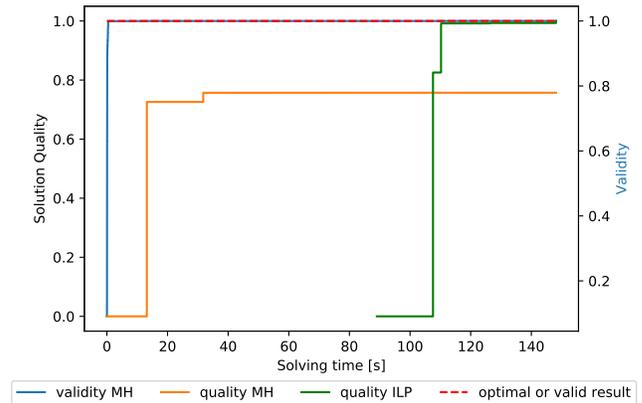


Figure 5: Solution quality improvement of MH- and ILP-based solvers

“almost” valid is certainly not enough and heuristic needs about 15 seconds to find the first valid solution (see a leap of the orange line). The orange line represents the quality of the MH result, until the valid solution is found we treat the quality as zero. After the solution is found we compare its quality with an optimal result, which is about 70%. During the optimization process heuristic finds one better solution and resides with the result of 75% comparing to the optimum.

The ILP-based solver finds the optimal solution very fast (in about 60 seconds). However, the transformation time, which implies a construction of the ILP provides a 90 second overhead resulting in overall 150 seconds of execution. The first valid solution is already better than the one found by the heuristic, but is provided only after 110 seconds. Therefore, we can observe a window of 95 seconds, when the utilization of heuristic is beneficial as the ILP-based solver is performing costly model transformation (RQ2).

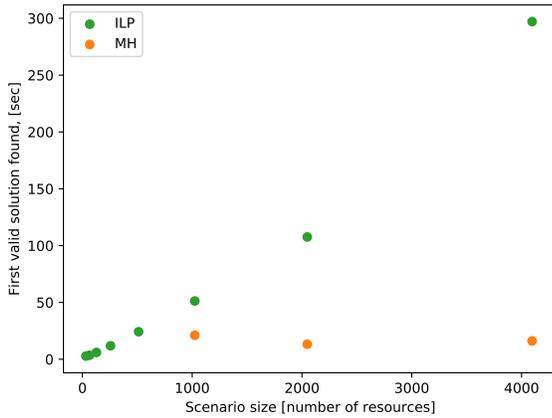
Figure 6 depicts the window between the first solution found by MH- and ILP-based solvers respectively and its change with the problem size. We can see that the size of the window grows exponentially with more HWC being added, thus, making the MH-based solver more and more viable (RQ3). In Figure 6 we can also see that the MH-based solver starts finding a solution starting from 1024 HWCs and is unable to solve smaller problems. It is a drawback of all parameter-dependent approaches, a single parameter setting cannot fit all problem sizes.

The reason for such an inefficient scalability of the ILP-based solver's transformation time is imposed by its nature. In order to form a linear program one needs to list all the possible combinations of HWCs and SWCs in the objective function. Additionally, to form the ILP constraints one needs to evaluate each mapping. MH-based Solver, on contrary, performs the evaluation of HWC-SWC mappings only during the score calculation process, i.e., at runtime. The transformation in this case is only a single iteration through all components and thus, scales linearly.

Table 1 shows us not only a comparison of MH- and ILP-based solvers, but the effect of parameter tuning and robustness of a single

Table 1: Comparison of ILP-based approach / manually tuned MH / automatically tuned MH

Scenario	Impl's	Compute resources	Valid?	Quality	Transformation time (s)	First valid solution found (s)	Last valid solution found (s)
0 trivial	1	1	✓/✓/✓	1.0 / 1.0 / 1.0	0.01 / 0.01 / 0.01	0.01 / 0.01 / 0.01	0.01 / 0.01 / 0.01
1 small	6	5	✓/✓/✓	1.0 / 1.0 / 1.0	0.02 / 0.01 / 0.01	0.02 / 4.21 / 4.21	0.02 / 4.21 / 4.21
2 small, much hardware	6	15	✓/✓/✓	1.0 / 1.0 / 1.0	0.03 / 0.01 / 0.01	0.03 / 0.11 / 0.11	0.03 / 0.11 / 0.11
3 small, complex software	62	47	✓/✓/✓	1.0 / 1.0 / 1.0	0.17 / 0.01 / 0.01	0.19 / 3.45 / 3.45	0.19 / 890 / 890
4 medium	30	68	✓/✓/✓	1.0 / 0.84 / 0.84	1.59 / 0.01 / 0.01	2.38 / 3.38 / 3.38	3.3 / 881 / 881
5 medium, much hardware	30	225	✓/✓/✓	1.0 / 0.94 / 0.94	5.12 / 0.01 / 0.01	5.72 / 1.62 / 1.62	33.5 / 464 / 464
6 medium, complex software	155	465	✗/✗/✗	✗/✗/✗	27.8 / 0.01 / 0.01	✗/✗/✗	✗/✗/✗
7 large	60	90	✓/✓/✓	0.88 / 0.61 / 0.61	7.75 / 0.01 / 0.01	13.97 / 338 / 338	804 / 338 / 338
8 large, much hardware	60	300	✗/✓/✓	✗/ 0.78 / 0.83	23.55 / 0.01 / 0.01	✗/ 536 / 94.5	✗/ 840 / 240
9 large, complex software	310	930	✗/✗/✗	✗/✗/✗	✗/ 0.02 / 0.02	✗/✗/✗	✗/✗/✗
10 huge	150	225	✗/✓/✓	✗/ 0.43 / 0.51	✗/ 0.02 / 0.02	✗/ 452 / 705	✗/ 743 / 705
11 huge, much hardware	150	750	✗/✓/✓	✗/ 0.38 / 0.43	✗/ 0.02 / 0.02	✗/ 15.28 / 895	✗/ 285 / 899
12 huge, complex software	620	2325	✗/✗/✗	✗/✗/✗	✗/ 0.05 / 0.05	✗/✗/✗	✗/✗/✗

**Figure 6: Execution time to find the first valid solution by MH and ILP-based solvers with changing problem size**

configuration. Manually obtained initial configuration is a good example for the robustness. From Table 1 we can see that the manually tuned configuration can solve all the same problem instances as the automatically tuned solver. Moreover, for the smaller scenarios 1–5 and 7 BRISE 2 was not able to find a better configuration. For the scenarios 8, 10 and 11 automatic parameter tuning can provide improved solutions, 11% better (RQ1). We assume that this trend remains with scaling. The full results as well as the SPL itself are available online⁴.

6 THREATS TO VALIDITY

The highest threat of internal validity is an implementation of the heuristic. In general, a thorough implementation of a heuristic-based solver as well as choice of a different meta-heuristic as a basis has a major influence on the quality of the resulting solution. While

developing the MH-based solver, we have developed several strategies of moves and score calculation with each subsequent iteration being more efficient. Investing more effort in solver engineering could have improved the resulting solution to a higher extent than our parameter tuning approach. However, parameter tuning can be applied on top of numerous optimization solver, thus, allowing to focus exactly on development of the solver.

Another threat is the maximal time we allowed for each configuration during tuning. We have picked 10 seconds as this cap, while in the TTC case (production in some sense) the maximal allowed time is 15 minutes. Such a reduction could have resulted in a wrongly picked optimal configuration as some configuration may have a steadier improvement speed, but ending up in a global optimum in the end.

One more threat comes from the area of exact solvers. The main drawbacks of exact solvers are high generation time and inefficient scaling with problem size growth. Utilization of abstraction on some problem entity, e.g., packing of similar hardware components in clusters may decrease both generation and solving time, while decreasing solution quality to some extent.

A threat to external validity lies in utilizing a single use case as a case study in this paper. We have still yet to define the limits of successful application of parameter tuning for combinatorial optimization.

7 CONCLUSION AND FUTURE WORK

Utilization of heuristic-based approaches for combinatorial optimization problems has its benefits and drawbacks. Active learning approaches to parameter tuning can aid the researcher to focus on improving the solver itself, while leaving a costly and often not obvious choice of parameters to the tuner. In this paper we present a simple metaheuristic-based solver to quality-based software selection and hardware mapping problem that enhanced with a parameter tuning software product line is able to outperform the state-of-the-art exact approaches in terms of scaling and time of finding the first valid solution, while the same solver being 11% less efficient without tuning.

⁴<https://github.com/dpukhkaiev/BRISE2/tree/MoST-Rec-19>

Our future work goes in several directions. We will continue extending the software product line with new features, such as, outlier detection, multi-objectiveness and adaptive task interruption strategies. We will try to push the limits of parameter tuning to combinatorial optimization problems themselves by decomposing variant selection and resource allocation into separate stages, thus, trading of exploration (variant selection) and exploitation (resource allocation). Exploration is a natural task of the *model*, while exploitation can be treated as a separate optimization problem.

Another promising research direction lies in abstraction of a problem formulation of an exact solver. Here, we aim to increase the performance of the approaches, whilst trading-off optimality.

ACKNOWLEDGMENTS

This work is supported (in part) by the German Research Foundation (DFG) within the Collaborative Research Center SFB 912-HAEC and BMBF project "Software Campus" (grant no. 01IS17044).

REFERENCES

- [1] Omid Alipourfard, Hongqiang Harry Liu, Jianshu Chen, Shivaram Venkataraman, Minlan Yu, and Ming Zhang. 2017. Cherrypick: Adaptively Unearthing the Best Cloud Configurations for Big Data Analytics. In *Proceedings of the 14th USENIX Conference on Networked Systems Design and Implementation (NSDI'17)*. USENIX Association, Berkeley, CA, USA, 469–482. <http://dl.acm.org/citation.cfm?id=3154630.3154669>
- [2] M. Antosiewicz, G. Koloch, and B. Kamiński. 2013. Choice of best possible metaheuristic algorithm for the travelling salesman problem with limited computational time: quality, uncertainty and speed. *Journal of Theoretical and Applied Computer Science* Vol. 7, nr 1 (2013), 46–55.
- [3] Andrea Arcuri and Gordon Fraser. 2011. On Parameter Tuning in Search Based Software Engineering. In *Search Based Software Engineering*, Myra B. Cohen and Mel Ó Cinnéide (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 33–47.
- [4] Walid Ben-Ameur. 2004. Computing the Initial Temperature of Simulated Annealing. *Computational Optimization and Applications* 29, 3 (01 Dec 2004), 369–385. <https://doi.org/10.1023/B:COAP.0000044187.23143.bd>
- [5] Mauro Birattari. 2009. *Tuning Metaheuristics: A Machine Learning Perspective* (1st ed. 2005. 2nd printing ed.). Springer Publishing Company, Incorporated.
- [6] Ilhem Boussaid, Julien Lepagnot, and Patrick Siarry. 2013. A survey on optimization metaheuristics. *Information Sciences* 237 (2013), 82 – 117. <https://doi.org/10.1016/j.ins.2013.02.041>
- [7] Michel Camilleri, Filippo Neri, and MICHALIS Papoutsidakis. 2014. An algorithmic approach to parameter selection in machine learning using meta-optimization techniques. *WSEAS Transactions on systems* 13 (2014), 202 – 213.
- [8] Valentin Dalibard, Michael Schaarschmidt, and Eiko Yoneki. 2017. BOAT: Building Auto-Tuners with Structured Bayesian Optimization. In *Proceedings of the 26th International Conference on World Wide Web (WWW '17)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, Switzerland, 479–488. <https://doi.org/10.1145/3038912.3052662>
- [9] Angela Dean, Max Morris, John Stufken, and Derek Bingham. 2015. *Handbook of design and analysis of experiments*. Vol. 7. CRC Press.
- [10] Stefan Falkner, Aaron Klein, and Frank Hutter. 2018. BOHB: Robust and Efficient Hyperparameter Optimization at Scale. In *Proceedings of the 35th International Conference on Machine Learning (Proceedings of Machine Learning Research)*, Jennifer Dy and Andreas Krause (Eds.), Vol. 80. PMLR, Stockholmsmässan, Stockholm Sweden, 1437–1446. <http://proceedings.mlr.press/v80/falkner18a.html>
- [11] Valerii Vadimovich Fedorov. 2013. *Theory of optimal experiments*. Elsevier.
- [12] Matthias Feurer and Frank Hutter. 2019. *Hyperparameter Optimization*. Springer International Publishing, Cham, 3–33. https://doi.org/10.1007/978-3-030-05318-5_1
- [13] Daniel Golovin, Benjamin Solnik, Subhodeep Moitra, Greg Kochanski, John Karro, and D. Sculley. 2017. Google Vizier: A Service for Black-Box Optimization. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '17)*. ACM, New York, NY, USA, 1487–1495. <https://doi.org/10.1145/3097983.3098043>
- [14] Sebastian Götz, Thomas Kühn, Christian Piechnick, Georg Püschel, and Uwe Aßmann. 2014. A Models@run.time Approach for Multi-objective Self-optimizing Software. In *Adaptive and Intelligent Systems*, Abdelhamid Bouchachia (Ed.). Springer International Publishing, Cham, 100–109.
- [15] Sebastian Götz, Johannes Mey, René Schöne, and Uwe Aßmann. 2018. Quality-based Software-Selection and Hardware-Mapping as Model Transformation Problem. In *Proceedings of the Transformation Tool Contest, a part of the Software Technologies: Applications and Foundations (STAF)*, Antonio Garcia-Dominguez, Georg Hinkel, and Filip Krikava (Eds.). CEUR-WS.org.
- [16] Sebastian Götz, Johannes Mey, René Schöne, and Uwe Aßmann. 2018. A JastAdd and ILP-Based Solution to the Software-Selection and Hardware-Mapping-Problem at the TTC 2018. In *Proceedings of the Transformation Tool Contest, a part of the Software Technologies: Applications and Foundations (STAF)*, Antonio Garcia-Dominguez, Georg Hinkel, and Filip Krikava (Eds.). CEUR-WS.org.
- [17] Samaneh HoseinDoost, Meysam Karimi, Shekoufeh Kollahdouz-Rahimi, and Bahman Zamani. 2018. Solving the Quality-based Software-Selection and Hardware-Mapping Problem with ACO. In *Proceedings of the Transformation Tool Contest, a part of the Software Technologies: Applications and Foundations (STAF)*, Antonio Garcia-Dominguez, Georg Hinkel, and Filip Krikava (Eds.). CEUR-WS.org.
- [18] Hui-Ling Huang and Fang-Lin Chang. 2007. ESVM: Evolutionary support vector machine for automatic feature selection and classification of microarray data. *Biosystems* 90, 2 (2007), 516 – 528. <https://doi.org/10.1016/j.biosystems.2006.12.003>
- [19] Frank Hutter, Holger H. Hoos, and Kevin Leyton-Brown. 2010. Automated Configuration of Mixed Integer Programming Solvers. In *Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems*, Andrea Lodi, Michela Milano, and Paolo Toth (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 186–202.
- [20] Michael Jugovac, Dietmar Jannach, and Lukas Lerche. 2017. Efficient optimization of multiple recommendation quality factors according to individual user tendencies. *Expert Systems with Applications* 81 (2017), 321 – 331. <https://doi.org/10.1016/j.eswa.2017.03.055>
- [21] Graham Kendall, Sigrid Knust, Celso C. Ribeiro, and Sebastián Urrutia. 2010. Scheduling in sports: An annotated bibliography. *Computers and Operations Research* 37, 1 (2010), 1 – 19. <https://doi.org/10.1016/j.cor.2009.05.013>
- [22] A Khachaturyan, S Semenovskaya, and B Vainstein. 1979. Statistical-thermodynamic approach to determination of structure amplitude phases. *Sov. Phys. Crystallography* 24, 5 (1979), 519–524.
- [23] S. Kirkpatrick, C. D. Gelatt, Jr., and M. P. Vecchi. 1987. Readings in Computer Vision: Issues, Problems, Principles, and Paradigms. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, Chapter Optimization by Simulated Annealing, 606–615. <http://dl.acm.org/citation.cfm?id=33517.33566>
- [24] Robert E. Kuenne and Richard M. Soland. 1972. Exact and approximate solutions to the multisource weber problem. *Mathematical Programming* 3, 1 (01 Dec 1972), 193–209. <https://doi.org/10.1007/BF01584989>
- [25] Michael Pinedo. 2012. *Scheduling*. Springer.
- [26] Dmytro Pukhkaiev and Sebastian Götz. 2018. BRISE: Energy-efficient Benchmark Reduction. In *Proceedings of the 6th International Workshop on Green and Sustainable Software (GREENS '18)*. ACM, New York, NY, USA, 23–30. <https://doi.org/10.1145/3194078.3194082>
- [27] H. Shakouri G., K. Shojaaee, and M. Behnam T. 2009. Investigation on the choice of the initial temperature in the Simulated Annealing: A mushy state SA for TSP. In *2009 17th Mediterranean Conference on Control and Automation*. 1050–1055. <https://doi.org/10.1109/MED.2009.5164685>
- [28] IM Sobol and Yu L Levitan. 1999. A pseudo-random number generator for personal computers. *Computers & Mathematics with Applications* 37, 4-5 (1999), 33–40.
- [29] Xiaochuan Tang, Yuanshun Dai, Peng Sun, and Sa Meng. 2018. Interaction-based feature selection using Factorial Design. *Neurocomputing* 281 (2018), 47 – 54. <https://doi.org/10.1016/j.neucom.2017.11.058>
- [30] Nianfeng Wan, Ardalan Vahidi, and Andre Luckow. 2016. Optimal speed advisory for connected vehicles in arterial roads and the impact on mixed traffic. *Transportation Research Part C: Emerging Technologies* 69 (2016), 548 – 563. <https://doi.org/10.1016/j.trc.2016.01.011>
- [31] Thomas Wortmann, Christoph Waibel, Giacomo Nannicini, Ralph Evins, Thomas Schropfer, and Jan Carmeliet. 2017. Are Genetic Algorithms Really the Best Choice for Building Energy Optimization?. In *Proceedings of the Symposium on Simulation for Architecture and Urban Design (SIMAUD '17)*. Society for Computer Simulation International, San Diego, CA, USA, Article 6, 8 pages. <http://dl.acm.org/citation.cfm?id=3289787.3289793>
- [32] Qian Yi, Congbo Li, Ying Tang, and Kingzheng Chen. 2015. Multi-objective parameter optimization of CNC machining for low carbon manufacturing. *Journal of Cleaner Production* 95 (2015), 256 – 264. <https://doi.org/10.1016/j.jclepro.2015.02.076>