



Discrete Optimization

Automatically improving the anytime behaviour of optimisation algorithms



Manuel López-Ibáñez*, Thomas Stützle

IRIDIA, Université Libre de Bruxelles (ULB), CP 194/6, Av. F. Roosevelt 50, B-1050 Brussels, Belgium

ARTICLE INFO

Article history:

Received 16 May 2012

Accepted 17 October 2013

Available online 28 October 2013

Keywords:

Metaheuristics

Anytime algorithms

Automatic configuration

Offline tuning

ABSTRACT

Optimisation algorithms with good anytime behaviour try to return as high-quality solutions as possible independently of the computation time allowed. Designing algorithms with good anytime behaviour is a difficult task, because performance is often evaluated subjectively, by plotting the trade-off curve between computation time and solution quality. Yet, the trade-off curve may be modelled also as a set of mutually nondominated, bi-objective points. Using this model, we propose to combine an automatic configuration tool and the hypervolume measure, which assigns a single quality measure to a nondominated set. This allows us to improve the anytime behaviour of optimisation algorithms by means of automatically finding algorithmic configurations that produce the best nondominated sets. Moreover, the recently proposed weighted hypervolume measure is used here to incorporate the decision-maker's preferences into the automatic tuning procedure. We report on the improvements reached when applying the proposed method to two relevant scenarios: (i) the design of parameter variation strategies for MAX-MIN Ant System and (ii) the tuning of the anytime behaviour of SCIP, an open-source mixed integer programming solver with more than 200 parameters.

© 2013 Elsevier B.V. All rights reserved.

1. Introduction

Many optimisation algorithms are designed without a specific termination criterion, and generate a sequence of feasible solutions that are increasingly better approximations of the optimal solution. However, the performance of an algorithm is often crucially determined by the choice of the termination criterion and the parameters of the algorithm. If the parameter settings of an algorithm result in fast convergence to good solutions, this may prevent the algorithm from adequately exploring the search space to find better solutions if given ample time. On the other hand, parameter settings that give higher exploration capabilities may produce poor results if the termination criterion is too short. Hence, there is a trade-off between solution quality and the runtime of the algorithm that can be adjusted by appropriately setting the parameters of the algorithm.

In many practical scenarios, an optimisation algorithm may be terminated at an arbitrary time, and, upon termination, the algorithm returns the best solution found since the start of the run. In such scenarios, the termination criterion is not known in advance, and, hence, the algorithm should produce as high quality solutions as possible at any moment of its run time. Algorithms

that show a better trade-off between solution quality and runtime are said to have a better *anytime behaviour* (Zilberstein, 1996).

There are two classical views when analysing the anytime behaviour (Hoos & Stützle, 2005). One view defines a number of termination criteria and analyses the quality achieved by the algorithm at each termination criterion. In this quality-over-time view, the anytime behaviour can be analysed as a series of plots of time-dependent solution quality distributions. A different view defines a number of target quality values and analyses the time required by the algorithm to reach each target. In this time-over-quality view, algorithms are often analysed in terms of a series of qualified runtime distributions.

In this paper, we consider a third view that does not favour time over quality or vice versa. Instead, this third view models the performance profile of an algorithm as a nondominated set in a multi-objective space. An algorithm has better anytime behaviour when it produces better nondominated sets, where “better” means better in terms of Pareto optimality. Surprisingly, this third view has received little attention (Hoos & Stützle, 2005; Chiarandini, 2005; den Besten, 2004), despite the important advances in theory and practice achieved in performance assessment of multi-objective optimisers in the last decade. Essentially, this model allows us to apply the same unary quality measures used in multi-objective optimisation to assign a single numerical value to the anytime behaviour of an algorithm's run. In this paper, we use the hypervolume measure as the unary quality measure for this purpose. The main reason is that the hypervolume is the quality measure with

* Corresponding author. Tel.: +32 (0) 2650 2745.

E-mail addresses: manuel.lopez-ibanez@ulb.ac.be (M. López-Ibáñez), stuetzle@ulb.ac.be (T. Stützle).

the highest discriminatory power among the known unary quality measures (Zitzler, Thiele, Laumanns, Fonseca, & Grunert da Fonseca, 2003). In addition, recent work has made possible to describe user preferences in terms of a weighted hypervolume measure (Auger, Bader, Brockhoff, & Zitzler, 2009), and, hence, our proposal allows incorporating user preferences when analysing the anytime behaviour of an algorithm. Moreover, as shown in this paper, evaluating the anytime behaviour of an algorithm in terms of the hypervolume allows applying automatic algorithm configuration methods to find parameter settings of an algorithm that optimise the trade-off between quality and time.

Recent advances in automatic configuration of algorithms (also called offline parameter tuning) have shown that such methods can save a significant amount of human effort and improve the performance of optimisation algorithms, when designing and evaluating new algorithms and when tuning existing algorithms to specific problems (Birattari, 2009; Bartz-Beielstein, 2006; Hutter, Hoos, Leyton-Brown, & Stützle, 2009; Eiben & Smit, 2011; Hutter, Hoos, & Leyton-Brown, 2011; Hoos, 2012). Our proposal here is to combine automatic configuration with the use of the hypervolume as a surrogate measure of anytime behaviour in order to enable the automatic configuration of algorithms in terms of anytime behaviour.

In the scenario described above, where the algorithm does not know its termination criterion in advance, techniques such as parameter adaptation are often applied to improve the anytime behaviour of the algorithm (Eiben, Michalewicz, Schoenauer, & Smith, 2007; Aine, Kumar, & Chakrabarti, 2009; Stützle et al., 2012). However, designing such parameter adaptation strategies is an arduous task, and they usually add new parameters to the algorithm that need to be tuned. The method proposed in this paper will help algorithm designers to compare and fine-tune such parameter adaptation strategies to find the settings that improve the anytime behaviour of the algorithm on the problem at hand. Indeed, the first case study reported here derives from our own efforts on designing parameter adaptation strategies for ant colony optimisation algorithms. This experience motivated us to develop the method proposed here, since the classical trial-and-error approach for designing such strategies proved extremely time-consuming.

The second case study reported here deals with a different scenario, in particular, a general purpose black-box solver (SCIP (Achterberg, 2009)) with a large number of parameters. The default parameter settings of such solvers are tuned for solving problem instances to optimality as fast as possible. However, in some practical scenarios, users may not want to wait until a problem instance is solved to optimality, and may decide to stop the solver at an arbitrary time. Using our method for fine-tuning the parameters of the solver with respect to anytime behaviour, users can improve the quality of the solutions found when the solver is stopped before reaching optimality, without knowing in advance the particular termination criterion.

The outline of the paper is as follows. Section 2 provides a background on automatic algorithm configuration, summarises the state of the art and describes the automatic configuration method (*irace*) used throughout this paper. Section 3 introduces the two classical views of the analysis of anytime algorithms and the less-explored multi-objective view. In Section 4, we describe our proposal in detail. We explain the benefits of using the hypervolume to evaluate the anytime behaviour of an algorithm in the context of an automatic configuration method. We discuss the choice of reference point and how to combine *irace* with the hypervolume measure. An additional section summarises related work and highlights the differences with our proposed approach. Section 5 describes our first case study, where we apply this proposal to the design of parameter adaptation strategies for MMAS. Section 6

discusses how our proposal enables a decision maker to incorporate preferences regarding the anytime behaviour of an algorithm to the automatic configuration procedure. A second case study is considered in Section 7, where we tune the anytime behaviour of SCIP. Finally, Section 8 provides a summary of our results and discusses possible extensions of the present work.

2. Preliminaries: automatic algorithm configuration

This section is a brief introduction to automatic algorithm configuration. We define the algorithm configuration problem, give an overview on the state of the art of automatic configuration methods, and describe *irace*, the automatic configuration method used throughout this paper. A more detailed and formal introduction is available from the literature referenced here and in the extended version of the paper (López-Ibáñez & Stützle, 2012a).

2.1. The algorithm configuration problem

Most algorithms for computationally hard optimisation problems have a number of parameters that need to be set. As an example, ACO algorithms (Dorigo & Stützle, 2004) often require the user to specify not only numerical parameters like the evaporation factor and the number of ants, but also components like the type of heuristic information and update method. Another example is mixed-integer programming solvers, such as SCIP (Achterberg, 2009), which often have a large number of configurable parameters affecting the main algorithm used internally, e.g., selecting among different branching strategies. The process of designing complex algorithms from a framework of algorithm components can be seen as an algorithm configuration problem (KhudaBukhsh, Xu, Hoos, & Leyton-Brown, 2009; Montes de Oca, Stützle, Birattari, & Dorigo, 2009; López-Ibáñez & Stützle, 2012c).

Given a parametrised algorithm, where each parameter may take different values (settings), a configuration of the algorithm is a unique assignment of values to parameters. When considering a problem to be solved by this parametrised algorithm, the goal of automatic configuration is to find the configuration that minimises a particular cost function over the set of possible instances of the problem. The cost function assigns a value to each configuration when applied to a single problem instance. In the case of stochastic algorithms, this cost measure is a random variable. Since most algorithms and problems of practical interest are sufficiently complex to preclude an analytical approach, the configuration of such algorithms follows an experimental approach (Birattari, 2009; Bartz-Beielstein, 2006).

2.2. Automatic configuration methods

The traditional approach to algorithm configuration consists of ad hoc experiments testing relatively few configurations. The use of experimental design techniques (Coy, Golden, Runger, & Wasil, 2001; Adenso-Díaz & Laguna, 2006) began a trend in which the task of finding the most promising configurations to be tested is performed automatically. The natural evolution of this trend has been to tackle algorithm configuration as an optimisation problem (Nannen & Eiben, 2006; Ansótegui, Sellmann, & Tierney, 2009; Hutter et al., 2009; Bartz-Beielstein, 2006; Bartz-Beielstein, Lasarczyk, & Preuss, 2010; Hutter et al., 2011). It is becoming widely accepted that automatic configuration methods may save substantial human effort during the empirical analysis and design of optimisation algorithms, and, at the same time, lead to better algorithms (Hoos, 2012; Eiben & Smit, 2011; Bartz-Beielstein, 2006; Birattari, 2009).

2.3. Iterated racing (*irace*)

Iterated F-race (I/F-Race) (Balaprakash, Birattari, & Stützle, 2007; Birattari, Yuan, Balaprakash, & Stützle, 2010) is a method for automatic configuration that consists of three steps: (1) sampling new configurations according to a probability distribution, (2) selecting the best configurations from the newly sampled ones by means of F-Race, and (3) updating the probability distribution in order to bias the sampling towards the best configurations. These three steps are repeated until a termination criterion is met, usually a predefined budget of runs of the algorithm being tuned. F-Race (Birattari, Stützle, Paquete, & Varrentapp, 2002) is a racing procedure (Maron & Moore, 1997) for the selection of the best among a given set of algorithm configurations, by means of the non-parametric Friedman's two-way analysis of variance by ranks, and its associated post-hoc test (Conover, 1999).

The *irace* software (López-Ibáñez, Dubois-Lacoste, & Stützle, 2011) implements a general *iterated racing* procedure, which includes I/F-Race as a special case. There are some notable differences between *irace* and the original description of I/F-Race, such as the use of truncated normal distribution for sampling numerical parameters, a restart mechanism for avoiding premature convergence, and other features described in the *irace* documentation (López-Ibáñez et al., 2011). For an outline of the iterated racing algorithm, we refer the reader to the extended version of the paper (López-Ibáñez & Stützle, 2012a).

3. Preliminaries: anytime algorithms

This section is an introduction to the analysis of stochastic optimisation algorithms and, in particular, anytime algorithms. We describe the two classical views, and the lesser studied multi-objective view. The concepts of bivariate runtime distributions (RTD), performance profiles, and the two classical views are described in textbooks (Hoos & Stützle, 2005). The extended version of the paper provides a more detailed introduction to these concepts (López-Ibáñez & Stützle, 2012a).

3.1. Classical views of the analysis of runtime distributions

Dean and Boddy (1988) describe an anytime algorithm as one that, first, may be interrupted at any moment and return a solution and, second, it keeps steadily improving its solution until interrupted, eventually finding the optimal. Most metaheuristics and other optimisation algorithms satisfy this condition, and, hence, they are anytime optimisation algorithms. A concept of anytime behaviour more useful in the context of metaheuristics was introduced by Zilberstein (1996), who highlights that algorithms with *good anytime behaviour* return as high-quality solutions as possible at any moment of their execution. A single run of an anytime algorithm generates a sequence of solutions that are increasingly better approximations of the optimal solution. Hence, in the context of anytime algorithms, an algorithm run is often described as a performance profile:

Definition 1 (*Performance profile* (Zilberstein, 1996)). Let us consider a single run r of an anytime optimisation algorithm A on a problem instance π , and record the computational effort (t_i , measured, for example as CPU-time in seconds) and the solution quality (q_i , measured, for example, as relative percentage deviation from the optimal solution quality), whenever a new best-so-far solution is found during the run of the algorithm. The set $P_r = \{(t_1, q_1), (t_2, q_2), \dots, c\}$ is called the *performance profile* of run r , where (t_i, q_i) are sampled with probability $rd(t, q)$, and $t_i < t_j \wedge q_i > q_j, \forall i < j$, where $rd(t, q)$ is called the bivariate run-time distribution (RTD) (Hoos & Stützle, 2005).

Fig. 1 gives an example of performance profiles for three runs A, B, and C. The above definition of performance profile does not favour quality over time, or viceversa. It would be equivalent to plot quality or time on either axis. Nonetheless, performance profile plots traditionally place time on the x-axis and quality on the y-axis, and we follow this custom here. A problem arises, however, when one wants to aggregate the information from several performance profiles in order to analyse the behaviour of an algorithm.

There are two classical views on how to summarise performance profiles over multiple runs (Hoos & Stützle, 2005). The solution-quality-over-time (SQT) view examines solution quality distributions (SQDs) over fixed run-times. However, instead of plotting the SQDs, it is far more common to aggregate the performance profiles over fixed run-times and examine SQT curves. This SQT view is the most popular in combinatorial optimisation, where it is common to compare the anytime behaviour of optimisation algorithm by visually inspecting mean SQT curves (Hoos & Stützle, 2005; Wah & Chen, 2000; Loudni & Boizumault, 2008; Stützle et al., 2012).

A second classical view is to aggregate performance profiles over fixed quality targets. Although plots of mean time over quality are possible (Hoos & Stützle, 2005), the most common approach is to examine the *qualified run time distribution* (QRTD), which is the probability of attaining fixed quality-targets over time. In an analysis based on QRTDs one can state that an algorithm is faster by some factor than another at reaching a particular quality target.

3.2. A multi-objective view of runtime distributions

A third alternative is to not favour either of the classical views, but instead use techniques from multi-objective optimisation (den Besten, 2004; Chiarandini, 2005). Thus, we do not aggregate the performance profiles over fixed run-times or fixed quality-targets. Instead, we describe the performance profiles generated by an optimisation algorithm as random nondominated sets. First, let us introduce some basic definitions borrowed from multi-objective optimisation, but adapted to the analysis of RTDs.

Definition 2 (*Weak dominance*). Given two vectors $(t, q), (t', q') \in \mathbb{R}_0 \times \mathbb{R}_0$, we say that (t, q) *weakly dominates* (t', q') , and denote it by $(t, q) \leq (t', q')$ iff $t \leq t' \wedge q \leq q'$.

Definition 3 (*Nondominated set*). A set $X = \{(t_1, q_1), (t_2, q_2), \dots, c\}$ is called *nondominated* iff $\nexists (t_i, q_i), (t_j, q_j) \in X, i \neq j$ such that $(t_i, q_i) \leq (t_j, q_j)$.

According to the definition of performance profile above (Definition 1), the elements of a performance profile, must also be mutually nondominated. Therefore, we can analyse performance profiles as nondominated sets, using the same techniques as in multi-objective optimisation.

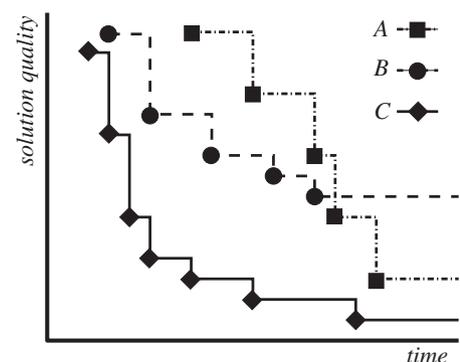


Fig. 1. Performance profiles of three algorithms A, B, and C (or three independent runs of the same stochastic algorithm).

The multi-objective view considers the bivariate RTD without aggregation, which is a well-known approach to the analysis of optimisation algorithms (Hoos & Stützle, 2005), but much less studied due to the inherent difficulty of analysing a bivariate distribution. Chiarandini (Chiarandini, 2005) considered such a view for analysing the anytime behaviour of several metaheuristics for graph colouring, concretely using the attainment function (Grunert da Fonseca & Fonseca, 2010).¹ For the purposes of automatic configuration, the available Kolmogorov-Smirnov test based on the attainment function (Knowles, Thiele, & Zitzler, 2006) is not enough, and graphical exploration methods (Knowles et al., 2006; Knowles, 2005; López-Ibáñez, Paquete, & Stützle, 2010) require substantial human interaction.

4. Our proposal: automatic configuration of anytime algorithms by means of the hypervolume measure

As mentioned above, directly using the attainment function as the basis of an automatic configuration tool for anytime algorithms seems difficult. Instead, we identify the hypervolume measure as the best available choice for this task. The main reasons are its high discriminatory power, being Pareto-compliant, and the possibility of incorporating user preferences into the automatic configuration process. There is also substantial and ongoing research on the theoretical properties of the hypervolume (Auger, Bader, Brockhoff, & Zitzler, 2012).

4.1. Hypervolume measure of performance profiles

As a first step, let us define the classical Pareto-dominance relation on performance profiles:

Definition 4 (Better in terms of Pareto-optimality, \triangleleft). Given two performance profiles P_i and P_j , which are the result of running two algorithms A_i and A_j on the same problem instance π , we say that P_i is better, in terms of Pareto-optimality, than P_j ($P_i \triangleleft P_j$) iff $P_i \neq P_j$, and $\forall (t_j, q_j) \in P_j, \exists (t_i, q_i) \in P_i$, such that $(t_i, q_i) \leq (t_j, q_j)$.

It is often the case, however, that neither performance profile is better than the other, i.e., they are *incomparable*. These relations are independent of the classical views of fixed-quality (first view) or fixed-runtimes (second view) as described above, and only make sense in the third view that considers both quality and runtime in terms of Pareto-optimality.

In the following, we assume that, without any a priori information about the actual termination criterion of the algorithm or the preferred trade-off between solution-quality and computation time, a performance profile that is better than another in terms of Pareto-optimality is also better in terms of anytime behaviour, in the sense that the former is always preferred to the latter.

In order to apply an automatic configuration tool from the literature for improving the performance profiles produced by an algorithm, we would desire a unary quality measure that unequivocally indicates whether a performance profile is better than another. Unfortunately, a well-known result from multi-objective optimisation states that no unary quality measure (or finite combination thereof) can indicate whether a performance profile is better, as defined above, than another (Zitzler et al., 2003). The most powerful of the unary quality indicators can at most indicate that a performance profile is not worse than (better than or incomparable to) another. The hypervolume measure (Zitzler & Thiele, 1999) is the

only unary quality measure known to have such discriminatory power (Zitzler et al., 2003).

The hypervolume can be defined as the measure of the region that is simultaneously weakly dominated by any point in a non-dominated set and bounded above by a reference point that is strictly dominated by all points in the set. In the context of performance profiles of single-objective optimisation algorithms, this region is the area contained within the orthogonal polygon defined by the elements of a performance profile P and an arbitrary reference point (t_r, q_r) , such that $(t_i, q_i) < (t_r, q_r), \forall (t_i, q_i) \in P$.

In this paper, we only consider single-objective optimisation algorithms. Nonetheless, it is trivial to extend the above discussion to multi-objective algorithms, where there is more than one measure of solution quality. In fact, we have applied the method proposed here to automatically improve the anytime behaviour of multi-objective evolutionary algorithms (Radulescu, López-Ibáñez, & Stützle, 2013).

4.1.1. The choice of reference point

In the context of anytime algorithms, the choice of reference point is application specific. The bivariate runtime distribution is defined without any limits on how bad the solution quality might be or how long it may take to generate any solution. In practice, however, large deviations from the optimal quality may be of little interest (no matter how fast they can be generated) and algorithms need to be stopped at some point (cut-off time). These limitations are not specific to any of the three views discussed above. A default approach is to consider a cut-off quality that corresponds to the worst solution quality found by any run of the algorithms under analysis, and a cut-off time that is slightly larger than the maximum time that would be reasonable for a single run of the algorithm. The reference point would then be defined as a factor larger than the cut-off quality and cut-off time.

How much larger this factor should be is an open question in multi-objective optimisation. There are some theoretical results on how the choice of reference point affects the distribution of elements within a nondominated set that maximises the hypervolume (Auger et al., 2012). However, it is not clear how these results extend to the relations between nondominated sets. In any case, the only effect of the reference point is to bias the preference between incomparable performance profiles. In that sense, our suggestion (and the usual practice in multi-objective optimisation) is to define the reference point in a consistent manner and bias this preference by other means, such as the weighted hypervolume (Section 6).

4.2. Automatic configuration of anytime algorithms

The use of the hypervolume to compare performance profiles in terms of Pareto-optimality has the additional benefit of providing a unary scalar measure to evaluate anytime behaviour. Integrating such a measure in most automatic configuration methods should be straightforward. Here, we discuss the practical aspects of the integration of the hypervolume in *irace* for improving the anytime behaviour of single-objective optimisation algorithms.

Our procedure requires to specify a maximum cut-off time for the algorithm being tuned by *irace*. As discussed above, this cut-off time is necessary because anytime algorithms may in principle run forever. The cut-off time could be dynamic or different for each run, however, for the sake of simplicity, we do not explore these possibilities here. Each run of the algorithm must produce a performance profile as defined in Definition 1.

Within a single race in *irace*, a set of algorithm configurations are evaluated on a sequence of training instances. After evaluating an instance, some configurations may be discarded. Let Θ_π denote the configurations that have *not* been discarded *before* evaluating

¹ For brevity, we do not explain the attainment function approach here and its application to the analysis of anytime algorithms. A complete description is provided in the extended version of the paper (López-Ibáñez & Stützle, 2012a).

instance π . Let $P_{\theta,\pi}$ denote the performance profile generated by running configuration θ on instance π . First, we normalise the performance profiles $P_{\theta,\pi}$, for each $\theta \in \Theta_\pi$ to the range $[0.0, 0.9]$. Thus, after normalisation we obtain for each θ a new performance profile:

$$P'_{\theta,\pi} = \{(t'_i, q'_i) | \forall (t_i, q_i) \in P_{\theta,\pi}\}$$

where

$$\begin{cases} t'_i = 0.9 \cdot (t_i - t_{\min}) / (t_{\max} - t_{\min}) \\ q'_i = 0.9 \cdot (q_i - q_{\min}) / (q_{\max} - q_{\min}) \end{cases}$$

where t_{\min} is usually zero, t_{\max} is the cut-off time, $q_{\min} = \min\{q_i | \forall (t_i, q_i) \in P_{\theta,\pi}, \forall \theta \in \Theta_\pi\}$ and q_{\max} is defined similarly for the maximum solution quality found after running all configurations in Θ_π on instance π .

Finally, we compute the hypervolume $h\nu(\theta, \pi)$ of each $P'_{\theta,\pi}$ using $(1.0, 1.0)$ as the reference point. In our proposal, this $h\nu(\theta, \pi)$ value becomes the cost measure used by *irace*.

Neither quality values, nor hypervolume values from different instances are directly compared because normalisation is done within each instance π , as defined above, and the application of the F-test within *irace* transforms the hypervolume values into ranks per instance. Hence, quality values or hypervolume values may have different ranges on each instance without introducing a bias. This approach also allows for instance-dependent cut-off times, although we do not explore this possibility in this paper.

4.3. Related work on automatic configuration of anytime algorithms

There is substantial work on automatic configuration for decision problems and single-objective optimisation problems. We refer to recent overviews (Hoos, 2012; Eiben & Smit, 2011) and books (Bartz-Beielstein, 2006; Birattari, 2009) for a complete bibliography. By comparison, there are relatively few works on tuning multi-objective optimisation algorithms. Wessing, Beume, Rudolph, and Naujoks (2010) automatically tuned the variation operator of a multi-objective evolutionary algorithm applied to a single problem instance. Simultaneously, López-Ibáñez and Stützle (2010, 2012c) automatically instantiated new designs of multi-objective ant colony optimisation (MOACO) algorithms for the bi-objective travelling salesman problem from a framework of MOACO algorithmic components. More recently, Dubois-Lacoste, López-Ibáñez, and Stützle (2011) applied this latter approach to outperform the state of the art in several bi-objective permutation flow-shop problems. These works share with our proposal the use of unary quality measures, such as the hypervolume, as the cost function used by the automatic configuration method.

On the other hand, our proposal should not be confused with parameter tuning as a multi-objective problem (Dréo, 2009), where the aim is to produce a set of parameter configurations that are mutually nondominated with respect to multiple criteria. In this paper, our aim is to produce a single parameter configuration that generates an anytime behaviour that is as good as possible.

There have been some recent attempts at tackling the problem of tuning anytime algorithms. As mentioned above, Chiarandini (2005) used the attainment function to analyse the anytime behaviour of several metaheuristics for graph colouring. However, it is far from obvious how to effectively use the attainment function in an automatic configuration method. The proposal closest to ours is by den Besten (2004), who combined racing and a performance measure based on the binary ϵ -indicator. The use of a binary measure involves computing a matrix of ϵ -measure values, comparing each alternative with the rest, and transforming it into ranks. More recently, Branke and Elomari (2011) combined a meta-level evolutionary algorithm and an ad hoc ranking procedure for tuning the

mutation rate of a lower-level algorithm for multiple termination criteria in a single tuner run. Their ranking method is not based on any multi-objective quality measure. Instead, it ranks each configuration with respect to the number of discrete time steps in which the configuration was better than other configurations. In that sense, it is an example of the classical fixed-runtimes view (what we call first view above).

5. Case Study: design of parameter variation strategies for MAX-MIN Ant System on the TSP

Many anytime algorithms use parameter adaptation strategies (Eiben et al., 2007; Aine et al., 2009), that is, the variation of parameter settings while solving a problem instance, to adapt the parameters to different phases of the search, and to balance exploration of the search space and exploitation of the best solutions found.

Designing and comparing parameter adaptation strategies is, however, an arduous and complex task. Traditionally, the analysis is performed in terms of one (or both) classical views, that is, either measuring solution quality over fixed-runtimes (Aine et al., 2009), or runtime (CPU-time or function evaluations) over fixed quality-targets (Auger & Hansen, 2005).

In previous work (Stützle et al., 2012; Maur, López-Ibáñez, & Stützle, 2010), we studied parameter adaptation strategies for ant colony optimisation (ACO) algorithms using the classical solution quality over fixed-runtimes view. In particular, we studied the anytime behaviour of MAX-MIN Ant System (MMAS) on the travelling salesman problem (TSP) by experimenting with various static parameter settings and parameter variation strategies. As a result, we identified parameter configurations that are significantly better in terms of anytime behaviour than the default settings of MMAS (Maur et al., 2010; Stützle et al., 2012). Our analysis relied on visually comparing the mean SQT curves of various strategies that were deemed interesting. Needless to say, this was a human intensive task that required many iterations of experimentation and analysis. We roughly estimate that the overall effort for obtaining the best configurations was close to one person-year.

This effort could have been significantly reduced by using an automatic algorithm configuration tool for improving the anytime behaviour. This case study was our main motivation for developing the method proposed in this paper. In this section, we describe the case study in detail, we examine the setup required for applying automatic algorithm configuration, and we compare the parameter adaptation strategies identified as the best by the automatic configuration procedure versus the ones identified in our previous work.

5.1. MAX-MIN Ant System

MMAS is an ACO algorithm that incorporates an aggressive pheromone update procedure and mechanisms to avoid search stagnation. When applying MMAS to the TSP, each ant starts at a randomly chosen initial city, and constructs a tour by randomly choosing at each step the city to visit next according to a probability defined by pheromone trails and heuristic information. In particular, the probability that ant k chooses a successor city j when being at city i is given by

$$p_{ij} = \frac{[\tau_{ij}]^\alpha \cdot [\eta_{ij}]^\beta}{\sum_{h \in N^k} [\tau_{ih}]^\alpha \cdot [\eta_{ih}]^\beta} \quad \text{if } j \in N^k, \quad (\text{otherwise, } p_{ij} = 0) \quad (1)$$

where τ_{ij} is the pheromone trail strength associated to edge (i, j) , η_{ij} is the corresponding heuristic information; α and β are two parameters that influence the weight given to pheromone and heuristic information, respectively; N^k is the feasible neighbourhood, that is, a candidate list of cities not yet visited in the partial tour of ant k .

Following previous work (Stützle, 1997), we also incorporate the *pseudo-random action choice rule* of ACS (Dorigo & Gambardella, 1997), which allows for a greedier solution construction. With a probability q_0 , an ant chooses next a city $j \in N^k$ such that $j = \arg \max_{h \in N^k} \{[\tau_{ih}]^\alpha \cdot [\eta_{ih}]^\beta\}$; otherwise, the ant performs the probabilistic selection based on Eq. (1). A value of $q_0 = 0$ reverts back to the original MMAS.

The pheromone update of MMAS updates all pheromone trails as

$$\tau_{ij} \leftarrow \max \left\{ \tau_{\min}, \min \left\{ \tau_{\max}, (1 - \rho) \cdot \tau_{ij} + \Delta\tau_{ij}^{\text{best}} \right\} \right\}, \quad (2)$$

where ρ , $0 < \rho \leq 1$, is a parameter called evaporation rate and $\Delta\tau_{ij}^{\text{best}} = 1/f(s^{\text{best}})$ if edge $(i, j) \in s^{\text{best}}$, ($\Delta\tau_{ij}^{\text{best}} = 0$, otherwise), where $f(s)$ is the tour length of solution s , and s^{best} is either the iteration-best solution, the best-so-far solution or the best solution since a re-initialisation of the pheromone trails (restart-best). In MMAS, these solutions are chosen alternately (Stützle & Hoos, 2000).

Finally, solutions constructed by the ants may be further improved by the application of a local search algorithm. In this paper, we will use MMAS with 2-opt local search, as was done in previous work (Maur et al., 2010; Stützle et al., 2012).

5.2. Parameter variation strategies in MMAS

Following our previous work (Stützle et al., 2012), we focus on two basic schemes for parameter variation in MMAS, which we call henceforth *delta* and *switch* strategies. During a single run of MMAS, the variation strategy called *delta* applied, for example, to parameter β increases the value of β at each iteration of the algorithm by a certain amount $\Delta\beta$, starting from the value β_{start} and stopping at the value β_{end} . If $\beta_{\text{start}} > \beta_{\text{end}}$, then the value of β is *decreased* at each iteration by $\Delta\beta$ instead of increased. Conversely, the variation strategy called *switch* changes, at iteration β_{switch} , the value of parameter β from the value β_{start} to the value β_{end} . An additional parameter β_{var} controls the variation strategy, which is either *delta*, *switch* or *none*, where *none* means that the parameter value of β stays constant throughout the run of the algorithm. As a result, we add to MMAS five additional parameters for varying β : β_{start} , β_{end} , $\Delta\beta$, β_{switch} and β_{var} .

We apply the same parameter variation strategies to parameters β , ρ , the number of ants (m), and q_0 . Hence, we add five additional static parameters for each parameter that is dynamically varied. Table 2 describes the domains of all parameters, and Table 1 describes the default values (Stützle & Hoos, 2000).

5.3. Automatic configuration of parameter adaptation strategies

We consider random uniformly generated instances of the symmetric TSP with 3000 cities (Johnson, McGeoch, Rego, & Glover, 2001). Our instances are available in the supplementary material page (López-Ibáñez & Stützle, 2012b). We generate 50 training (tuning) instances and 50 test instances.

We apply our proposed method for the automatic configuration of the anytime behaviour (Section 4.2). The automatic configuration tool is the implementation of I/F-Race provided by the *irace* software package (López-Ibáñez et al., 2011). As explained above, we incorporate the hypervolume measure to *irace* in order to evaluate the anytime behaviour of a single run of MMAS. We use

Table 1
Default settings of the parameters under study for MMAS.

Algorithm	Time _{CPU}	α	β	ρ	m	q_0
MMAS	500 s	1.0	2.0	0.2	25	0.0

Table 2
Parameter space for variation strategies of MMAS.

Parameter	Domain	Constraint
$m_{\text{var}}, \beta_{\text{var}}, \rho_{\text{var}}, q_{0\text{var}}$	{delta, switch, none}	
m	[1, 100]	if var = none
β	[0, 20]	
ρ	[0.01, 1.0]	
q_0	[0.0, 1.0]	
Δm	{0.01, 0.05, 0.1, 0.25, 0.5, 1, 2, 5}	if var = delta
$\Delta\beta$	{0.01, 0.02, 0.05, 0.1, 0.2, 0.5, 1.0}	
$\Delta\rho$	{0.001, 0.002, 0.005, 0.01}	
Δq_0	{0.0001, 0.0002, 0.0005, 0.001, 0.002, 0.005}	
$m_{\text{switch}}, \beta_{\text{switch}}, \rho_{\text{switch}}, q_{0\text{switch}}$	[1, 500]	if var = switch
$m_{\text{start}}, m_{\text{end}}, \beta_{\text{start}}, \beta_{\text{end}}, \rho_{\text{start}}, \rho_{\text{end}}, q_{0\text{start}}, q_{0\text{end}}$	1, [1, 500], [0, 20], [0, 5], [0.001, 1.0], [0.0, 1.0]	if var \in {delta, switch}

a publicly available implementation of the hypervolume measure (Fonseca, Paquete, & López-Ibáñez, 2006), and the MMAS implementation is based on the ACOTSP software (Stützle, 2002). All experiments are run on Intel Xeon E5410 CPUs (2.33 GHz, 2×6MB L2 cache) running Cluster Rocks Linux version 6/CentOS 6.3, 64 bits. Each individual run of the algorithms being tuned uses just one core.

In a first step, we tune separately the variation strategy of one dynamic parameter at a time, while other parameters are fixed to their default values as given in Table 1. That is, we perform one run of *irace* for each parameter $\{m, \beta, q_0, \rho\}$. For example, in the run of *irace* that tunes the variation strategy of β , the parameters tuned are $\beta_{\text{var}}, \beta, \Delta\beta, \beta_{\text{switch}}, \beta_{\text{start}}, \beta_{\text{end}}$, whereas the other parameters (m, q_0, ρ) are fixed to their default values (Table 1) and their variation strategies ($param_{\text{var}}$) are set to none, and, hence, their corresponding variation parameters ($\Delta param, param_{\text{switch}}, param_{\text{start}}, param_{\text{end}}$) are not considered (see Table 2). We give each run of *irace* a budget of 1000 runs of MMAS. Each run of MMAS is stopped after Time_{CPU} seconds (Table 1).

In a second step, we automatically configure all parameter variation strategies at the same time, that is, we configure 24 parameters instead of six. Since the parameter space is much larger now, we assign a larger tuning budget to this run of *irace*, specifically 10000 runs of MMAS.

After each run of *irace* finishes, we apply the resulting parameter configurations (Table 4) to the test instances. In addition, we also run, on the test instances, the *default* parameter configuration of MMAS (Table 1) without any variation strategy and several variation strategies previously found by manual ad hoc experimentation (Table 3) (Stützle et al., 2012). We present these results in the following sections.

5.4. Analysis of the results

5.4.1. Automatic configuration vs. manual configuration

Our goal is to improve the anytime behaviour of MMAS over the whole set of test instances. Hence, we analyse the overall results by plotting the average solution quality over time (SQT) for all test instances at once. However, as explained above, our proposed approach *does not rely* on these SQT curves for improving the anytime behaviour, and it does not favour solution-quality over

Table 3

Parameter configurations found by a human expert when varying one parameter at a time in MMAS (Stützle et al., 2012).

Configuration	Parameter settings
manual var ants	$m_{var} = \text{delta}$, $m_{start} = 1$, $m_{end} = 25$, $\Delta m = 0.1$
manual var beta	$\beta_{var} = \text{switch}$, $\beta_{start} = 20$, $\beta_{end} = 3$, $\beta_{switch} = 50$
manual var rho	$\rho_{var} = \text{none}$, $\rho = 0.9$
manual var q0	$q_{0var} = \text{delta}$, $q_{0start} = 0.99$, $q_{0end} = 0$, $\Delta q_0 = 0.0005$

time or viceversa. We could also visualise our results in terms of qualified runtime distributions (Hoos & Stützle, 2005), or in terms of attainment surfaces (Grunert da Fonseca & Fonseca, 2010; Chiarandini, 2005). We chose SQT curves as one of the traditional means of visualising the anytime behaviour of an algorithm, and the most popular view in the literature on combinatorial optimisation algorithm (Hoos & Stützle, 2005).

For each algorithm configuration, we have the best solution quality found f_{rit} on run r on instance i at time t . We compute the relative percentage deviation (RPD) from the optimal solution for each instance as $RPD_{rit} = 100 \cdot f_{rit}/f_i^{opt}$, where f_i^{opt} is the optimal tour length of instance i . Then, we compute the mean RPD over all 50 instances and over all 15 independent runs of each algorithm as $RPD_t = \frac{1}{50 \cdot 15} \cdot \sum_{i=1}^{50} \sum_{r=1}^{15} RPD_{rit}$. Each line in the plots in Fig. 2 corresponds to the RPD_t of one algorithm configuration, that is, aggregating solution-quality over time.

In the case of the hypervolume computation, we do not aggregate over time, but compute the hypervolume of the (nondominated) performance profile of each run on each instance by normalising both time and solution quality to the interval [0.0,0.9] and using (1.0,1.0) as the reference point. Then, for each algorithm configuration we compute its mean hypervolume over all its runs on all test instances, and we give this value in the legend of each plot.

The first important observation is how a larger hypervolume value matches a better anytime behaviour. The four plots in Fig. 2 show a large improvement in the anytime behaviour of the manually tuned configurations with respect to the default configuration of MMAS. Nonetheless, the automatically found configurations are able to match, and in most cases surpass, the manually tuned configurations in terms of hypervolume, despite the fact that the manually tuned configurations were found by extensive experimentation under the guidance of human expertise.

Fig. 3(a) compares the configuration obtained after automatically configuring all parameter variation strategies at once versus the best configurations obtained after automatically configuring the variation strategy of one parameter at a time. In our previous

study, the manual tuning and analysis of all parameter strategies at once was ruled out as infeasible, given the extremely large number of potential configurations and interactions among different parameters. Here, we see that automatically configuring all parameters at once leads to an additional improvement in anytime behaviour.

Figs. 2(a–d) and 3(a) show the mean hypervolume over all runs and all test instances. To assess whether the observed differences are statistically significant, we perform a statistical analysis of the results over the whole set of test instances. The analysis is based on the Friedman test for analysing non-parametric unrepeated complete block designs, and its associated post-test for multiple comparisons (Conover, 1999). First, we calculate the mean hypervolume of the 15 runs of each algorithm for each instance. Then, we perform a Friedman test using the instances as the blocking factor, and the different configurations of MMAS as the treatment factor. The null hypothesis is that the configurations have identical effect on the ranking according to the hypervolume within each instance. If the Friedman test rejects the null hypothesis given a significance level of $\alpha = 0.05$, we proceed to calculate the minimum difference between the sum of ranks of two configurations that is statistically significant (ΔR_α). In this manner, we identify which configurations are significantly different from the best ranked one, i.e., the one with the lowest sum of ranks.

Table 5 summarises the results of the statistical analysis. It shows the value of ΔR_α for $\alpha = 0.05$, the different configurations of MMAS sorted by increasing sum of ranks, and the difference between the sum of ranks of each configuration and the best configuration (ΔR). For each parameter considered, the ranking shown in Table 5 always ranks higher the configurations found automatically than their counterparts found by ad hoc experimentation (auto vs. manual, respectively). More importantly, it shows that the best ranked configuration is the one that automatically configured all parameters at once, and that the difference in ranks between this configuration and the rest is statistically significant.

5.4.2. Hypervolume vs. final quality

Here, we show that the use of the hypervolume as the tuning criterion is the key factor for improving the anytime behaviour. Fig. 3(b) shows four configurations of MMAS: the default configuration (default); the one resulting from automatically tuning all variation parameters (auto var ALL); a configuration obtained by tuning all variation parameters with respect to final quality, that is, the solution quality obtained at 500 seconds (auto var final); and a configuration obtained by tuning the classical MMAS parameters, without any variation, with respect to final quality (auto fix final). The plot shows that, independently of whether

Table 4

Parameter configurations found by irace for MMAS: (auto var param) was obtained by tuning, with respect to anytime behaviour, the variation strategy of parameter param, while the other parameters are set to their default settings; (auto var ALL) was obtained by tuning, with respect to anytime behaviour, all variation parameters at the same time; (auto fix final) was obtained by tuning the classical MMAS parameters, without any variation, with respect to final quality, that is, the solution quality obtained at 500 seconds; and (auto var final) was obtained by tuning all variation parameters with respect to final quality.

Configuration	Parameter settings
auto var ants	$m_{var} = \text{delta}$, $m_{delta} = 0.05$, $m_{start} = 1$, $m_{end} = 417$
auto var beta	$\beta_{var} = \text{delta}$, $\beta_{delta} = 0.05$, $\beta_{start} = 9$, $\beta_{end} = 4$
auto var q0	$q_{0var} = \text{switch}$, $q_{0switch} = 200$, $q_{0start} = 0.96$, $q_{0end} = 0.30$
auto var rho	$\rho_{var} = \text{delta}$, $\rho_{delta} = 0.001$, $\rho_{start} = 0.82$, $\rho_{end} = 0.84$
auto var ALL	$m_{var} = \text{delta}$, $m_{delta} = 1$, $m_{start} = 1$, $m_{end} = 384$, $\beta_{var} = \text{switch}$, $\beta_{switch} = 79$, $\beta_{start} = 5$, $\beta_{end} = 0$, $q_{0var} = \text{delta}$, $q_{0delta} = 0.002$, $q_{0start} = 0.87$, $q_{0end} = 0.57$, $\rho_{var} = \text{none}$, $\rho = 0.68$
auto fix final	$\beta = 5.9$, $\rho = 0.62$, $m = 84$, $q_0 = 0.099$
auto var final	$m_{var} = \text{switch}$, $m_{switch} = 50$, $m_{start} = 1$, $m_{end} = 317$, $\beta_{var} = \text{delta}$, $\beta_{delta} = 0.5$, $\beta_{start} = 8$, $\beta_{end} = 2$, $q_{0var} = \text{switch}$, $q_{0switch} = 139$, $q_{0start} = 0.6241$, $q_{0end} = 0.2725$, $\rho_{var} = \text{switch}$, $\rho_{switch} = 493$, $\rho_{start} = 0.338$, $\rho_{end} = 0.7495$

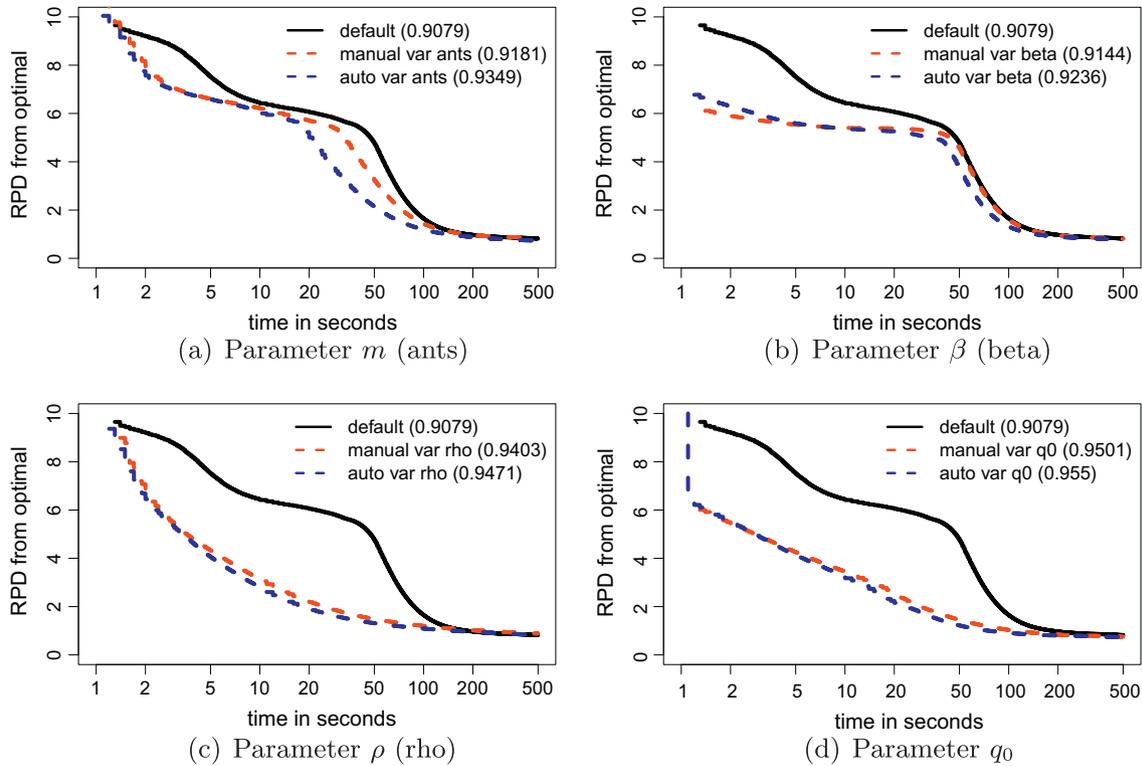


Fig. 2. Anytime behaviour of manually tuned vs. automatically tuned configurations of MMAS. The number in the legend is the mean hypervolume of each configuration over all runs. For the meaning of the labels, see the caption of Table 4.

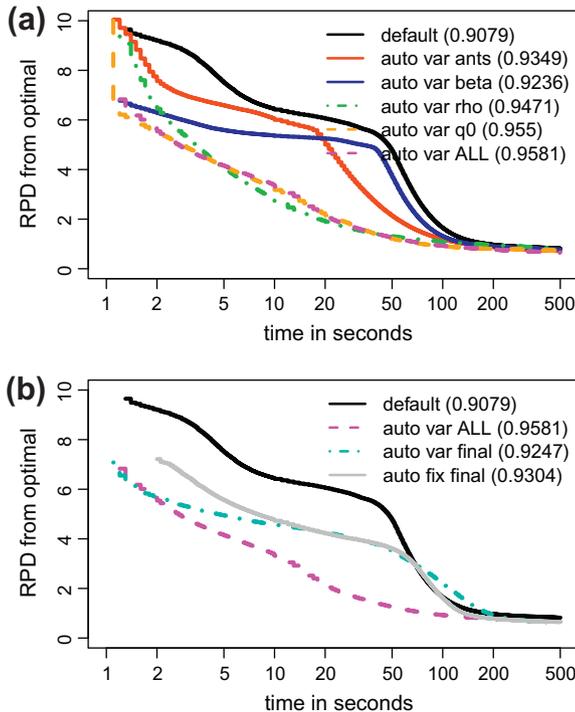


Fig. 3. Anytime behaviour of automatically tuned configurations of MMAS vs. the default configuration. The number in the legend is the mean hypervolume of each configuration over all runs. For the meaning of the labels, see the caption of Table 4.

MMAS uses parameter variation or not, the results not tuned with respect to the hypervolume have worse anytime behaviour.

A possible concern of tuning for anytime behaviour is a significant loss of final quality. Hence, we examine the final quality

Table 5

Various configurations of MMAS ordered according to the sum of ranks with respect to the hypervolume obtained over all test instances. The numbers in parenthesis are the differences of ranks relative to the best ranked configuration. ΔR_x is the statistically significant difference in ranks according to the post-hoc test for the Friedman-test with $\alpha = 0.05$. All configurations are statistically significantly worse than the best one (auto var ALL). For the meaning of the labels, see the caption of Table 4.

ΔR_x	Configurations (ΔR)
9.87	auto var ALL (0) , auto var q0 (34), manual q0 (94), auto var rho (125), manual rho (189), auto var ants (236), auto var beta (294), manual ants (345), manual beta (380), default (433)

achieved by these four variants of MMAS in Fig. 4. According to the boxplots, there is an important improvement in the final quality achieved in comparison with the default configuration of MMAS, even for the configuration tuned for anytime behaviour. The boxplot does not show a large difference between the three automatically configured variants. Nonetheless, the Friedman test indicates that the final quality obtained by the variant tuned for anytime behaviour is statistically worse than the variants tuned for final quality (Table 6). In order to assess the loss of final quality, we compute the 95% confidence interval on the mean difference in final quality between the configuration tuned for anytime behaviour and the best ranked configuration, which is [0.0244, 0.0480], measured in RPD.² Although the loss in final quality when tuning for anytime behaviour is small in this case, an anytime algorithm should aim to match the best possible final quality in the ideal case.

² The confidence interval is computed using the Welch's *t* statistic for two paired samples, which assumes that the samples follow a normal distribution. Nonetheless, for large sample sizes, as used here, the method is robust against deviations from normality.

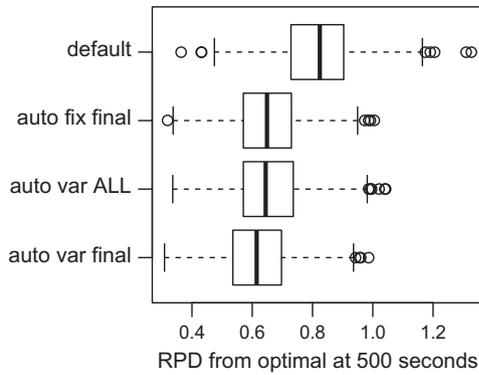


Fig. 4. Final quality achieved by several variants of MMAS. For the meaning of the labels, see the caption of Table 4.

6. Articulation of preferences in automatic configuration of anytime algorithms

The use of the hypervolume for automatic tuning of anytime algorithms has an additional advantage compared to other unary measures, that is, the possibility of specifying the decision-maker's preferences. A recent proposal extends the hypervolume indicator by a weight function over the objective space (Zitzler, Brockhoff, & Thiele, 2007; Auger et al., 2009). A weight function that assigns a larger value to a certain region of the objective space will bias the hypervolume indicator to favour nondominated sets that dominate that region. We show here that this formulation can straightforwardly be used to introduce a bias in the anytime behaviour produced by automatic configuration.

As an example, let us assume that the decision maker's preference is to obtain as good final solution quality as possible, while still giving some minor importance to achieving a good anytime behaviour. In other words, the decision maker prefers configurations that generate solution-quality curves that are better towards minimising the solution quality (in our case, the second objective). Zitzler et al. (2007) suggest to model this preference by considering the following weight function (adapted here to minimisation):

$$w^{qual}(z) = e^{20 \cdot (1-z_2)} / e^{20} \tag{3}$$

where $z = (z_1, z_2) \in Z$ is an objective vector, with z_1 representing time and z_2 representing solution quality, and $Z = [0, 1] \times [0, 1]$ represents the normalised bi-objective space of time \times quality.

The weighted hypervolume is computed as the integral of the weight function over the region dominated by a set of nondominated points and bounded above by a reference point. To give a rough idea of this integral when using the weighted function w^{qual} , Fig. 5(b) shows the value of the weighted hypervolume for each individual vector in the normalised objective space Z and with reference point (1,1). The plot shows that vectors with very small values of z_2 are assigned a high hypervolume, but vectors with values of z_2 larger than 0.2 are assigned a hypervolume close to zero. By comparison, Fig. 5(a) shows the value of the non-weighted hypervolume, which is symmetric around the diagonal, that is, without a preference for either objective.

As shown in Fig. 5(b), when using the weighted function w^{qual} , the gradient of the hypervolume values is very steep and most of the objective space has a hypervolume close to zero. We can make the gradient gentler by weighting also the z_1 component (corresponding to time), but then we have to increase the exponent associated to z_2 in order to keep a strong preference for low solution quality. This is done with the following weight function:

$$w^{xqual} = e^{10 \cdot z_1} / e^{10} + e^{100 \cdot (1-z_2)} / e^{100} \tag{4}$$

Table 6

Configurations of MMAS ordered according to the sum of ranks with respect to the final solution quality obtained after 500 seconds. The numbers in parenthesis are the difference of ranks relative to the best configuration. ΔR_z is the statistically significant difference in ranks according to the post-hoc test for the Friedman-test with $\alpha = 0.05$. Configurations that are not significantly different from the best one are indicated in bold face. For the meaning of the labels, see the caption of Table 4.

ΔR_z	Configurations (ΔR)
13.68	auto var final (0) , auto var ALL (48), auto fix final (52), default (132)

The weighted hypervolume using this weight function for each individual vector in the objective space Z is shown in Fig. 5(c). In this case, there is a gentler gradient of the hypervolume value than in Fig. 5(b). Moreover, the value of the hypervolume increases exponentially in the direction of decreasing z_2 (solution quality), while it stays roughly constant along z_1 (except for very high values of z_1).

We illustrate the differences between the original hypervolume and the two weighted variants above with an example. Fig. 6 shows five performance profiles (not aggregated over fixed runtime or over fixed quality-targets) in the normalised objective space Z . The plot shows the region $z_2 \in [0.0, 0.15]$, where we can see that the performance profiles are ordered according to the final quality achieved, with profile a being the best and profile e being the worst. The legend provides three numbers for each profile, which correspond to evaluating the profile with the classical hypervolume, the hypervolume weighted by w^{qual} and the hypervolume weighted by w^{xqual} , respectively. Table 7 gives the profiles in increasing order of preference according to each measure.

In this example, the classical hypervolume ranks profile a, which is the profile with the best final quality, worse than other three profiles. The weighted hypervolume functions increase the preference for profile a, and our proposed variant w^{xqual} gives it the highest rank.

Next, we test the effect of these two weighted hypervolume functions on the automatic configuration procedure. In particular, we carry out additional runs of *irace* using the weighted hypervolume variants described above, i.e., w^{qual} (Eq. (3)) and w^{xqual} (Eq. (4)). We run *irace* with the same setup as for tuning all parameter variations at once in Section 5.2, in particular, with a budget of 10000 runs of ACOTSP. These additional tuning runs produce two new configurations of MMAS, which we ran 25 times with different random seed on each test instance.

Fig. 7(a) plots the mean RPD over all runs of the resulting four configurations of MMAS: the default configuration (`default`); the one resulting from automatically tuning all variation parameters at once using the classical hypervolume (`auto var ALL`); the configuration obtained with the same tuning setup but using the weighted hypervolume with w^{qual} (`whv qual`); and the configuration obtained using the weighted hypervolume with w^{xqual} (`whv xqual`). In addition, the legend provides three numbers for each profile, which correspond to evaluating the results with the classical hypervolume, the hypervolume weighted by w^{qual} and the hypervolume weighted by w^{xqual} , each of them averaged over all runs.

Interestingly, the values reported in the legend of Fig. 7(a) indicate that the configuration tuned using w^{xqual} as the anytime criterion obtains a better hypervolume weighted by w^{qual} than the configuration tuned using w^{qual} as the anytime criterion. Taking into account Fig. 5(b and c), we can observe that both weight functions are strongly correlated and also that w^{qual} is "flatter" than w^{xqual} , that is, there are more plateau regions with almost the same value for different points. Our conjecture is that the strong correlation makes possible to tune for one weight function and maximise the other. At the same time, the relative flatness of w^{qual} makes it a

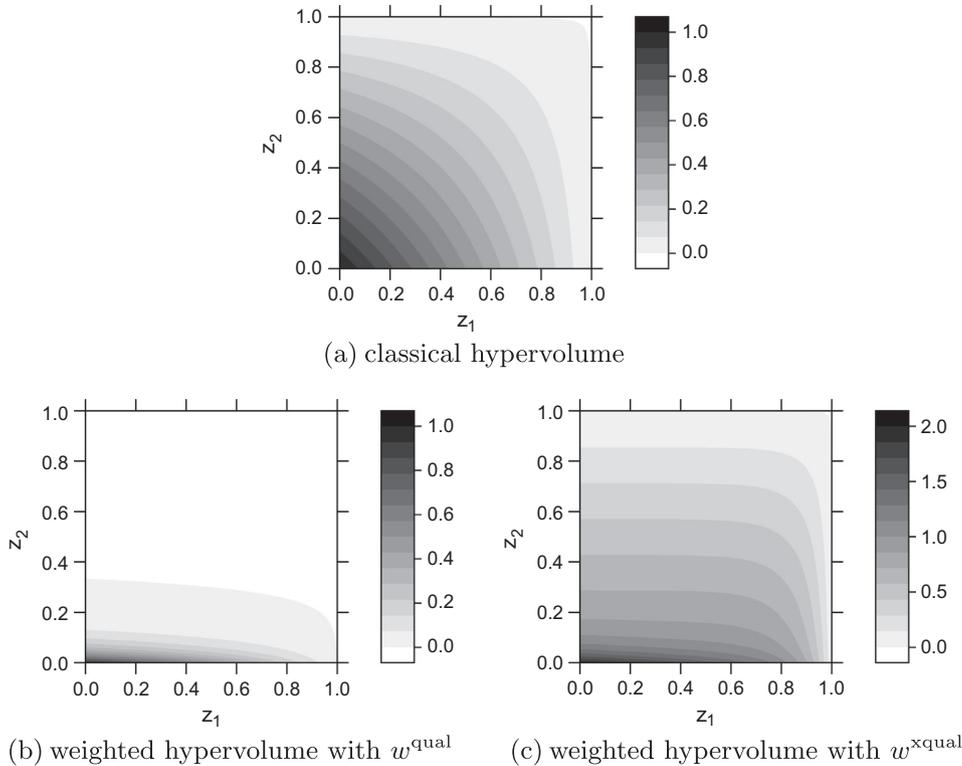


Fig. 5. Value of the classical hypervolume (a) and the weighted hypervolume when using weight functions w^{qual} (b) and w^{xqual} (c) on the normalised objective space. The grey level at each point gives the (weighted) hypervolume of that individual point, which is computed as the integral of the weight function over the region dominated by the point and bounded above by the reference point (1,1).

harder optimisation criterion for tuning than w^{xqual} . Additional experiments (López-Ibáñez & Stützle, 2012a) appear to confirm this conjecture.

In terms of final quality, the two configurations tuned with the weight functions (w^{qual} and w^{xqual}) are slightly better than the one tuned with the classical hypervolume, as indicated by the boxplots given in Fig. 7(b). Moreover, the configurations tuned with the weight functions obtain the lowest final quality in most instances. In fact, according to the Friedman test, these configurations are significantly better than configurations obtained by tuning for the classical hypervolume and for final quality (Table 8). The main conclusion of these experiments is that the weighted hypervolume allows us to set preferences on the trade-off between quality and time. For example, the weighted function w^{xqual} imposes a strong preference for good final quality.

7. Case study: automatic configuration of an anytime MIP solver

7.1. Experimental setup

In this second scenario, we apply our proposed approach to a very different problem with a large number of parameters. In particular, we tune 207 parameters of SCIP (Achterberg, 2009), a mixed integer programming (MIP) solver. The number of parameters is too large to be detailed here, but details can be found in the supplementary page (López-Ibáñez & Stützle, 2012b).

The benchmark set is composed of 2000 MIP-encoded instances (200 goods, 1000 bids) of the NP-hard winner determination problem for combinatorial auctions (Leyton-Brown, Pearson, & Shoham, 2000; Hutter et al., 2009). The benchmark set is split in two disjoint sets of 1000 instances each, one is used for training and the other for testing. In a combinatorial auction, bids are placed for subsets of goods. The goal in the winner determination problem is to find

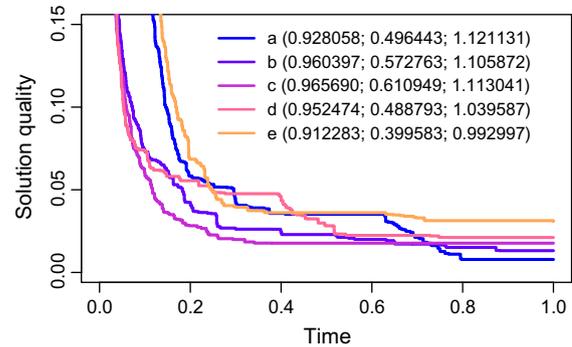


Fig. 6. For each performance profile, the legend shows the classical hypervolume, and the weighted hypervolume variants w^{qual} and w^{xqual} .

Table 7

Ranking of the performance profiles in Fig. 6 according to various preferences.

Preference	Ranking (best to worst)
Final quality	a b c d e
Hypervolume	c b d a e
w^{qual}	c b a d e
w^{xqual}	a c b d e

an assignment of goods to bids that maximises the total value of the winning bids.

For our experiments here, we use SCIP version 2.0.2 linked with the linear programming solver SoPlex 1.5.0. We set the maximum memory limit of SCIP to 350MB. During our experiments, we discovered that some parameter configurations produced an incorrect

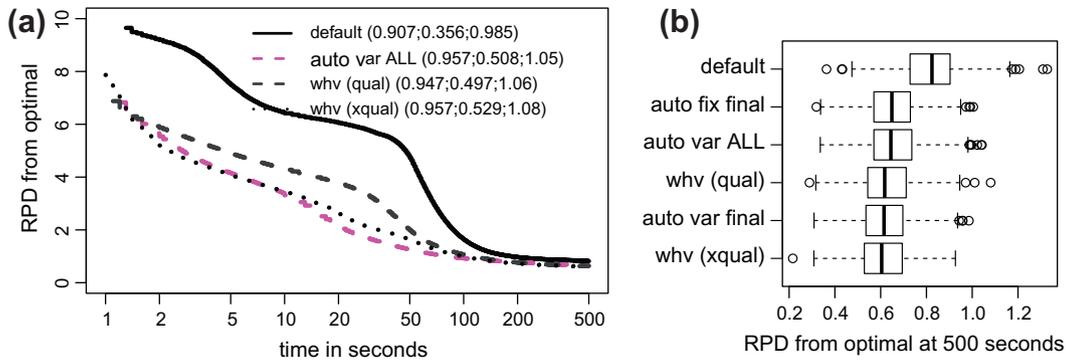


Fig. 7. Automatically tuned configurations of MMAS vs. the default configuration. (a) SQT curves, the legend shows the classical hypervolume, and the weighted hypervolume variants w^{qual} and w^{xqual} . (b) Final quality achieved. For the meaning of the labels, see the caption of Table 4.

behaviour of SCIP, and we assign those configurations the worst possible hypervolume. We give SCIP a time limit of 300 seconds, and we allow 5000 runs of SCIP for each run of irace. We carry out the tuning as before, that is, we combine irace with the hypervolume measure in order to improve the anytime behaviour of SCIP. We seed the automatic configuration procedure with the default configuration of SCIP.

For the purposes of comparison, we perform two additional tuning runs with two different objectives: (1) minimising the run-time to find the optimal solution and (2) maximising the final objective value obtained after 300 seconds. Thus, we obtain two additional configurations of SCIP, which we label as auto time and auto quality, respectively. We use these configurations to assess the potential loss of either run time or final solution quality, when tuning for improving the anytime behaviour. Finally, we run all configurations of SCIP obtained from the various tuning setups plus the default configuration one time on each test instance.

7.2. Analysis of SCIP configurations

As a first step, we graphically examine the solution quality over time. For each configuration, we compute the mean RPD over the 1000 test instances at each time step. Next, we plot the mean RPD over time (Fig. 8). The legend gives the mean hypervolume value corresponding to each configuration of SCIP. The plot uses a logarithmic scale for the x-axis (time), since the largest differences appear on the first half of the computation time limit.

The plot shows that the configuration tuned with the hypervolume (auto anytime) obtains a better anytime behaviour (and a higher hypervolume) than the rest. Moreover, both the configuration tuned for final quality and the one tuned for solving time show worse anytime behaviour (and lower hypervolume) than the default configuration of SCIP. The differences observed in the hypervolume values (and, hence, in the anytime behaviour) of each SCIP configuration are more evident in Fig. 9(a), which shows that the hypervolume values corresponding to auto anytime are much larger than those corresponding to the other configurations of SCIP.

Improving the anytime behaviour does not necessarily mean that instances are solved faster to optimality. Fig. 9(b) shows the time required by each configuration to solve each of the 1000 test instances. The best configurations of SCIP according to this criterion are the default configuration and the configuration tuned specifically for this criterion (auto time). This result is not surprising, since this is the most popular evaluation criterion in mixed-integer programming, and, hence, we presume that SCIP is by default tuned for it.

We also examine the potential loss of final quality. Fig. 9(c) shows the RPD from the optimal at the cut-off time of 300 seconds.

Table 8

Configurations of MMAS ordered according to the sum of ranks with respect to the final solution quality obtained. The numbers in parenthesis are the difference of ranks relative to the best configuration. ΔR_z is the statistically significant difference in ranks according to the post-hoc test for the Friedman-test with $\alpha = 0.05$. Configurations that are not significantly different from the best one are indicated in bold face. For the meaning of the labels, see the caption of Table 4.

ΔR_z	Configurations (ΔR)
23.16	whv (xqual) (0), auto var final (11) , whv (qual) (47), auto var ALL (94), auto fix final (99), default (199)

All configurations solve most of the instances to optimality (or very close to it). However, the configuration tuned for anytime (auto anytime) is the one that diverges most often from near-optimality. Hence, there is some loss of final quality when tuning using the hypervolume.

If we look at the solution quality up to a different cut-off time, the situation is certainly different. For example, if we consider solution quality up to 10 seconds (Fig. 9(d)), there is a large difference between the configurations. While the auto anytime configuration obtains an RPD value much lower than 10% in most cases, the RPD values of the default configuration are frequently larger than 10%.

The observations above are further confirmed by statistical analysis. We carry out four independent Friedman tests (as described in Section 5.4.1), one for each evaluation criterion shown in Fig. 9. The results of the four tests are reported in Table 9. As expected, the best configuration in terms of hypervolume is the one tuned for that criterion (auto anytime), which is significantly better than the rest by a large margin. The auto anytime configuration is also the clear winner in terms of the solution quality obtained if stopped after 10 seconds. Moreover, in terms of final

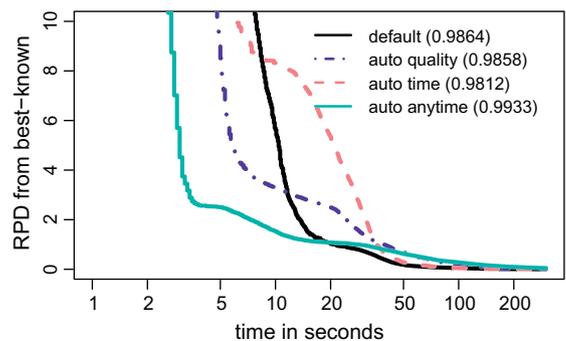


Fig. 8. Mean RPD over all test instances for different configurations of SCIP. The number in parentheses is the mean hypervolume corresponding to that configuration.

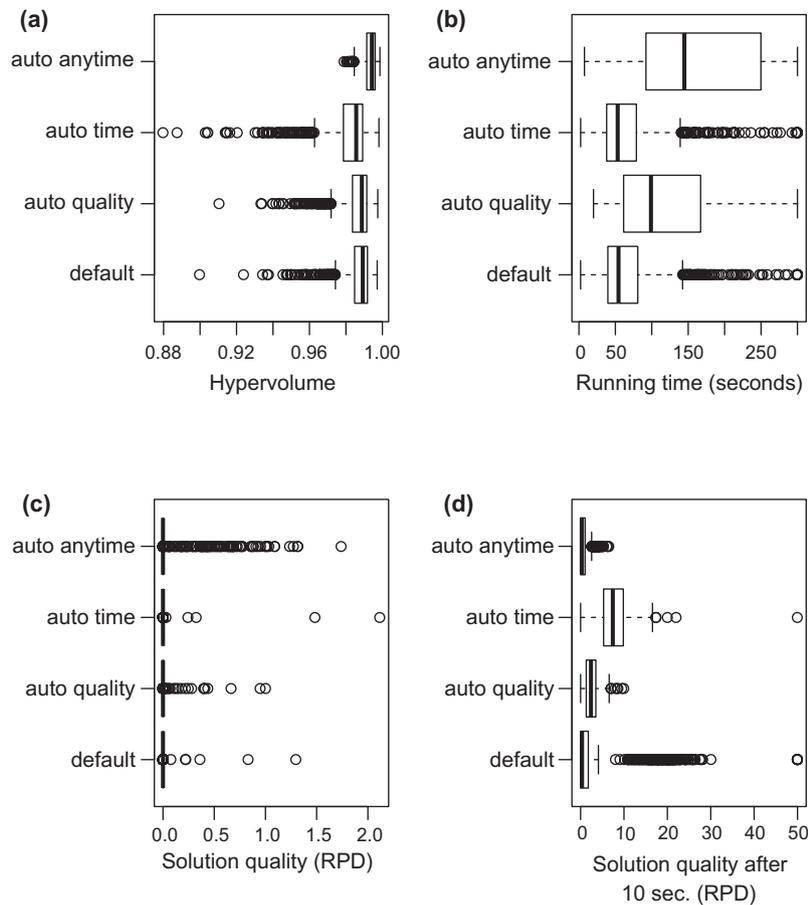


Fig. 9. Boxplots of the results obtained by four different parameter configurations of SCIP according to various evaluation criteria.

Table 9

Configurations of SCIP ordered according to the sum of ranks with respect to four different evaluation criteria. The numbers in parenthesis are the difference of ranks relative to the best ranked configuration. ΔR_x is the statistically significant difference in ranks according to the post-hoc test for the Friedman-test ($\alpha = 0.05$). Configurations that are not significantly different from the best one according to the Friedman test are in bold face.

ΔR_x	Configurations (ΔR)	Evaluation criterion
75.1	auto anytime (0), default (1183), auto quality (1490), auto time (2335)	Hypervolume
53.52	auto time (0), default (192), auto quality (1603), auto anytime (2353)	Time to best found
83.01	auto anytime (0), default (460), auto quality (1019.5), auto time (2024.5)	Quality after 10 seconds
∞	default (0), auto time (8.5), auto quality (19), auto anytime (44.5)	Final quality (300 seconds)

quality, the differences between the strategies are not statistically significant. The difference in the sum of ranks between *auto anytime* and *default* is only 44.5.

Finally, by using the weighted hypervolume as explained in Section 6, we are able to find a configuration of SCIP with good anytime behaviour and that ranks better than *default* according to final quality. However, the differences in ranks are still not statistically significant. Hence, for conciseness, we do not discuss the results of using the weighted hypervolume for tuning SCIP here, but we provide the results as supplementary material (López-Ibáñez & Stützle, 2012b). The results provided here are sufficient to conclude that the proposed method was able to find a configuration of SCIP that has better anytime behaviour than the default, without a significant loss of final quality.

8. Conclusions

In this paper, we have shown that the combination of *irace* and the hypervolume quality measure is effective at improving

the anytime behaviour of optimisation algorithms. We have presented two representative and challenging case studies. The first case study compared the results obtained automatically against those obtained by a human expert for the task of designing parameter variation strategies that show good anytime behaviour. Our results show that the automatic configuration method is able to match the anytime behaviour obtained by the parameter variation strategies designed by a human expert. Moreover, the automatic method allows exploring a much larger design space, potentially leading to configurations with better anytime behaviour. These are expected results when using automatic configuration tools for tuning with a fixed termination criterion. However, this is the first time that such results have been obtained when automatically designing anytime algorithms. In a follow-up work, we have applied the approach proposed here to improve the anytime behaviour of a state of the art optimiser for black-box continuous optimisation (López-Ibáñez, Liao, & Stützle, 2012). Our results there show that even for such state-of-the-art optimisers, the default parameter settings are not well-suited for scenarios where

the termination criterion is unknown in advance. Although the results presented here focus on single-objective optimisers, our approach is applicable to multi-objective optimisers as well. In another follow-up work, we have applied it to automatically configure the parameters of multi-objective evolutionary algorithms in order to improve their anytime behaviour (Radulescu et al., 2013).

In the second case study presented here, we apply our approach to an off-the-shelf optimisation solver, with a very large number of parameters. In this case, the optimisation solver is already tuned to solve problems to optimality as fast as possible. However, we show that if stopped before reaching optimality, the results may be very poor. Our proposed approach helps to tune such solvers in order to be more robust in case of earlier termination, without specifying in advance when the algorithm could be terminated. Our results show that important improvements can be obtained, specially for very early termination, without sacrificing much of the final quality. Moreover, comparing the configurations that produce better anytime behaviour versus those that produce better final quality (or shorter time to optimality) may lead to improvements in the solvers themselves.

The choice of the hypervolume measure also allows incorporating preference information into the automatic configuration process by means of the weighted hypervolume. We propose a weighted formulation that emphasises a good final quality but still takes into account the overall anytime behaviour of the algorithms. We show that, by adding such preferences, it is possible to effectively bias the configurations selected by the automatic configuration tool. This allows customising algorithms to very specific anytime scenarios, where an exact termination criterion is not known, but there is some a priori knowledge of what is expected.

An open question is how to extend the results to longer termination criteria than the ones that are feasible to test during automatic configuration. A problem that may arise is that configurations produce good results up until the tested termination criterion, but the performance becomes unsatisfactory for longer runs. Woodruff, Ritzinger, and Oppen (2011) have studied how to dynamically set a termination criterion. Survival analysis techniques may help to estimate the behaviour of the algorithms for longer runtime (Gagliolo & Legrand, 2010). These techniques could be incorporated into our approach in order to dynamically adjust the maximum cut-off time while tuning the anytime behaviour.

Finally, we are convinced that our approach contributes towards the final goal of designing algorithms that are more robust to different termination criteria and, hence, applicable to a wider range of scenarios, without sacrificing solution quality.

Acknowledgements

We are grateful to the SCIP developers, in particular, Timo Berthold, Stefan Heinz and Ambros Gleixner, for kindly providing information about SCIP parameters. The research leading to the results presented in this paper has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013)/ERC Grant Agreement n° 246939. This research and its results have also received funding from the COMEX project within the Interuniversity Attraction Poles Programme of the Belgian Science Policy Office and the Meta-X project, funded by the Scientific Research Directorate of the French Community of Belgium. Manuel López-Ibáñez and Thomas Stützle acknowledge support of the F.R.S.-FNRS of which they are a post-doctoral researcher and a senior research associate, respectively.

References

- Achterberg, T. (2009). SCIP: Solving constraint integer programs. *Mathematical Programming Computation*, 10(1), 0 1–41.
- Adenso-Díaz, B., & Laguna, M. (2006). Fine-tuning of algorithms using fractional experimental design and local search. *Operations Research*, 54(1), 0 99–114.
- Aine, S., Kumar, R., & Chakrabarti, P. P. (2009). Adaptive parameter control of evolutionary algorithms to improve quality-time trade-off. *Applied Soft Computing*, 9(2), 0 527–540.
- Ansótegui, C., Sellmann, M., & Tierney, K. (2009). A gender-based genetic algorithm for the automatic configuration of algorithms. In I. P. Gent (Ed.), *Principles and practice of constraint programming, CP 2009*. LNCS (Vol. 5732, pp. 142–157). Springer.
- Auger, A., Bader, J., Brockhoff, D., & Zitzler, E. (2009). Investigating and exploiting the bias of the weighted hypervolume to articulate user preferences. In F. Rothlauf (Ed.), *Proceedings of GECCO 2009* (pp. 563–570). ACM Press.
- Auger, A., Bader, J., Brockhoff, D., & Zitzler, E. (2012). Hypervolume-based multiobjective optimization: Theoretical foundations and practical implications. *Theoretical Computer Science*, 425, 0 75–103.
- Auger, A., & Hansen, N. (2005). A restart CMA evolution strategy with increasing population size. In *Proceedings of CEC 2005* (pp. 1769–1776). IEEE Press.
- Balaprakash, P., Birattari, M., & Stützle, T. (2007). Improvement strategies for the F-race algorithm: Sampling design and iterative refinement. In T. Bartz-Beielstein et al. (Eds.), *Hybrid metaheuristics*. LNCS (Vol. 4771, pp. 108–122). Springer.
- Bartz-Beielstein, T. (2006). *Experimental research in evolutionary computation: The new experimentalism*. Springer.
- Bartz-Beielstein, T., Lasarczyk, C., & Preuss, M. (2010). The sequential parameter optimization toolbox. In T. Bartz-Beielstein et al. (Eds.), *Experimental methods for the analysis of optimization algorithms* (pp. 337–360). Springer.
- Birattari, M. (2009). *Tuning metaheuristics: A machine learning perspective*. *Studies in computational intelligence* (Vol. 197). Springer.
- Birattari, M., Stützle, T., Paquete, L., & Varrentapp, K. (2002). A racing algorithm for configuring metaheuristics 2002. In W. B. Langdon et al. (Eds.), *Proceedings of GECCO* (pp. 11–18). San Francisco, CA: Morgan Kaufmann Publishers.
- Birattari, M., Yuan, Z., Balaprakash, P., & Stützle, T. (2010). F-race and iterated F-race: An overview. In T. Bartz-Beielstein et al. (Eds.), *Experimental methods for the analysis of optimization algorithms* (pp. 311–336). Springer.
- Branke, J., & Elomari, J. (2011). Simultaneous tuning of metaheuristic parameters for various computing budgets. In N. Krasnogor & P. L. Lanzi (Eds.), *Proceedings of GECCO 2011* (pp. 263–264). ACM Press.
- Chiarandini, M. (2005). *Stochastic local search methods for highly constrained combinatorial optimisation problems*. PhD thesis, FG Intellektik, FB Informatik, TU Darmstadt, Germany.
- Conover, W. J. (1999). *Practical nonparametric statistics*. John Wiley & Sons.
- Coy, S. P., Golden, B. L., Runger, G. C., & Wasil, E. A. (2001). Using experimental design to find effective parameter settings for heuristics. *Journal of Heuristics*, 7(1), 77–97.
- Dean, T., & Boddy, M. S. (1988). An analysis of time-dependent planning. In *Proceedings of the 7th national conference on artificial intelligence. AAAI-88* (pp. 49–54). AAAI Press.
- den Besten, M. (2004). *Simple metaheuristics for scheduling*. PhD thesis, FB Informatik, TU Darmstadt, Germany.
- Dorigo, M., & Gambardella, L. M. (1997). Ant colony system: A cooperative learning approach to the traveling salesman problem. *IEEE Transactions on Evolutionary Computation*, 1(1), 0 53–66.
- Dorigo, M., & Stützle, T. (2004). *Ant colony optimization*. Cambridge, MA: MIT Press.
- Dréo, J. (2009). Using performance fronts for parameter setting of stochastic metaheuristics. In F. Rothlauf (Ed.), *GECCO (companion)* (pp. 2197–2200). ACM Press.
- Dubois-Lacoste, J., López-Ibáñez, M., & Stützle, T. (2011). Automatic configuration of state-of-the-art multi-objective optimizers using the TP + PLS framework. In N. Krasnogor & P. L. Lanzi (Eds.), *Proceedings of GECCO 2011* (pp. 2019–2026). ACM Press.
- Eiben, A. E., Michalewicz, Z., Schoenauer, M., & Smith, J. E. (2007). Parameter control in evolutionary algorithms. In F. Lobo et al. (Eds.), *Parameter setting in evolutionary algorithms* (pp. 19–46). Springer.
- Eiben, A. E., & Smit, S. K. (2011). Parameter tuning for configuring and analyzing evolutionary algorithms. *Swarm and Evolutionary Computation*, 10(1), 19–31.
- Fonseca, C. M., Paquete, L., & López-Ibáñez, M. (2006). An improved dimension-sweep algorithm for the hypervolume indicator. In *Proceedings of the 2006 congress on evolutionary computation (CEC 2006)* (pp. 1157–1163). Piscataway, NJ: IEEE Press.
- Gagliolo, M., & Legrand, C. (2010). Algorithm survival analysis. In T. Bartz-Beielstein et al. (Eds.), *Experimental methods for the analysis of optimization algorithms* (pp. 161–184). Springer.
- Grunert da Fonseca, V., & Fonseca, C. M. (2010). The attainment-function approach to stochastic multiobjective optimizer assessment and comparison. In T. Bartz-Beielstein et al. (Eds.), *Experimental methods for the analysis of optimization algorithms* (pp. 103–130). Springer.
- Hoos, H. H. (2012). Programming by optimization. *Communications of the ACM*, 55(2), 0 70–80.
- Hoos, H. H., & Stützle, T. (2005). *Stochastic local search—Foundations and applications*. San Francisco, CA: Morgan Kaufmann Publishers.

- Hutter, F., Hoos, H. H., & Leyton-Brown, K. (2011). Sequential model-based optimization for general algorithm configuration. In *Learning and intelligent optimization, 5th international conference, LION 5*. LNCS. Springer.
- Hutter, F., Hoos, H. H., Leyton-Brown, K., & Stützle, T. (2009). ParamILS: An automatic algorithm configuration framework. *Journal of Artificial Intelligence Research*, 36, 267–306.
- Johnson, D. S., McGeoch, L. A., Rego, C., & Glover, F. (2001). *8th DIMACS implementation challenge*. <http://www.research.att.com/dsj/chtsp/>.
- KhudaBukhsh, A. R., Xu, L., Hoos, H. H., & Leyton-Brown, K. (2009). SATenstein: Automatically building local search SAT solvers from components. In C. Boutilier (Ed.), *Proceedings of the twenty-first international joint conference on artificial intelligence (IJCAI-09)* (pp. 517–524). Menlo Park, CA: AAAI Press.
- Knowles, J. D. (2005). A summary-attainment-surface plotting method for visualizing the performance of stochastic multiobjective optimizers. In *Proceedings of the 5th international conference on intelligent systems design and applications* (pp. 552–557).
- Knowles, J. D., Thiele, L., & Zitzler, E. (2006). *A tutorial on the performance assessment of stochastic multiobjective optimizers*. TIK-Report 214, Computer Engineering and Networks Laboratory (TIK), Swiss Federal Institute of Technology (ETH), Zürich, Switzerland, February 2006. Revised version.
- Leyton-Brown, K., Pearson, M., & Shoham, Y. (2000). Towards a universal test suite for combinatorial auction algorithms. In A. Jhingran et al. (Eds.), *ACM conference on electronic commerce (EC-00)* (pp. 66–76). New York, NY: ACM Press.
- López-Ibáñez, M., & Stützle, T. (2012a). *Automatically improving the anytime behaviour of optimisation algorithms*. Technical Report TR/IRIDIA/2012-012, IRIDIA, Université Libre de Bruxelles, Belgium.
- López-Ibáñez, M., & Stützle, T. (2012b). *Automatically improving the anytime behaviour of optimisation algorithms: Supplementary material*. <http://iridia.ulb.ac.be/supp/IridiaSupp2012-011/>.
- López-Ibáñez, M., Dubois-Lacoste, J., Stützle, T., & Birattari M. (2011). *The irace package, iterated race for automatic algorithm configuration*. Technical Report TR/IRIDIA/2011-004, IRIDIA, Université Libre de Bruxelles, Belgium.
- López-Ibáñez, M., Liao, T., & Stützle, T. (2012). On the anytime behavior of IPOP-CMA-ES. In C. A. Coello Coello et al. (Eds.), *PPSN 2012, Part I*. LNCS (Vol. 7491, pp. 357–366). Springer.
- López-Ibáñez, M., Paquete, L., & Stützle, T. (2010). Exploratory analysis of stochastic local search algorithms in biobjective optimization. In T. Bartz-Beielstein et al. (Eds.), *Experimental methods for the analysis of optimization algorithms* (pp. 209–222). Springer.
- López-Ibáñez, M., & Stützle, T. (2010). Automatic configuration of multi-objective ACO algorithms. In M. Dorigo et al. (Eds.), *Swarm intelligence, 7th international conference, ANTS 2010*. LNCS (Vol. 6234, pp. 95–106). Springer.
- López-Ibáñez, M., & Stützle, T. (2012c). The automatic design of multi-objective ant colony optimization algorithms. *IEEE Transactions on Evolutionary Computation*, 16(6), 861–875.
- Loudni, S., & Boizumault, P. (2008). Combining VNS with constraint programming for solving anytime optimization problems. *European Journal of Operational Research*, 191, 705–735.
- Maron, O., & Moore, A. W. (1997). The racing algorithm: Model selection for lazy learners. *Artificial Intelligence Research*, 11(1–5), 193–225.
- Maur, M., López-Ibáñez, M., & Stützle, T. (2010). Pre-scheduled and adaptive parameter variation in MAX-MIN Ant System. In H. Ishibuchi et al. (Eds.), *Proceedings of the 2010 congress on evolutionary computation (CEC 2010)* (pp. 3823–3830). IEEE Press.
- Montes de Oca, M. A., Stützle, T., Birattari, M., & Dorigo, M. (2009). Frankenstein's PSO: A composite particle swarm optimization algorithm. *IEEE Transactions on Evolutionary Computation*, 13(5), 1120–1132.
- Nannen, V., & Eiben, A. E. (2006). A method for parameter calibration and relevance estimation in evolutionary algorithms. In M. Cattolico et al. (Eds.), *Proceedings of GECCO 2006* (pp. 183–190). ACM Press.
- Radulescu, A., López-Ibáñez, M., & Stützle, T. (2013). Automatically improving the anytime behaviour of multiobjective evolutionary algorithms. In R. C. Purshouse et al. (Eds.), *EMO 2013*. LNCS (Vol. 7811, pp. 825–840). Springer.
- Stützle, T. (1997). *MAX-MIN Ant System for the quadratic assignment problem*. Technical Report AIDA-97-4, FG Informatik, FB Informatik, TU Darmstadt, Germany, July 1997.
- Stützle T. (2002). ACOTSP: A software package of various ant colony optimization algorithms applied to the symmetric traveling salesman problem. <http://www.aco-metaheuristic.org/aco-code/>.
- Stützle, T., & Hoos, H. H. (2000). MAX-MIN ant system. *Future Generation Computer Systems*, 16(8), 889–914.
- Stützle, T., López-Ibáñez, M., Pellegrini, P., Maur, M., Montes de Oca, M. A., Birattari, M., et al. (2012). Parameter adaptation in ant colony optimization. In Y. Hamadi, E. Monfroy, & F. Saubion (Eds.), *Autonomous search* (pp. 191–215). Berlin, Germany: Springer.
- Wah, B. W., & Chen, Y. X. (2000). Optimal anytime constrained simulated annealing for constrained global optimization. In R. Dechter (Ed.), *Principles and practice of constraint programming, CP 2000*. LNCS (Vol. 1894, pp. 425–440). Springer.
- Wessing, S., Beume, N., Rudolph, G., & Naujoks, B. (2010). Parameter tuning boosts performance of variation operators in multiobjective optimization. In R. Schaefer et al. (Eds.), *Parallel problem solving from nature XI*. LNCS (Vol. 6238, pp. 728–737). Springer.
- Woodruff, D. L., Ritzinger, U., & Oppen, J. (2011). Research note: the point of diminishing returns in heuristic search. *International Journal of Metaheuristics*, 10(3), 0 222–231.
- Zilberstein, S. (1996). Using anytime algorithms in intelligent systems. *AI Magazine*, 17(3), 0 73–83.
- Zitzler, E., Brockhoff, D., & Thiele, L. (2007). The hypervolume indicator revisited: On the design of Pareto-compliant indicators via weighted integration. In S. Obayashi et al. (Eds.), *EMO 2007*. LNCS (Vol. 4403, pp. 862–876). Springer.
- Zitzler, E., & Thiele, L. (1999). Multiobjective evolutionary algorithms: A comparative case study and the strength Pareto evolutionary algorithm. *IEEE Transactions on Evolutionary Computation*, 3(4), 0 257–271.
- Zitzler, E., Thiele, L., Laumanns, M., Fonseca, C. M., & Grunert da Fonseca, V. (2003). Performance assessment of multiobjective optimizers: an analysis and review. *IEEE Transactions on Evolutionary Computation*, 7(2), 0 117–132.