

Multi-Objective Optimization Through Pareto Minimal Correction Subsets*

Miguel Terra-Neves, Inês Lynce and Vasco Manquinho

INESC-ID / Instituto Superior Técnico, Universidade de Lisboa, Portugal

{neves,ines,vmm}@sat.inesc-id.pt

Abstract

A Minimal Correction Subset (MCS) of an unsatisfiable constraint set is a minimal subset of constraints that, if removed, makes the constraint set satisfiable. MCSs enjoy a wide range of applications, such as finding approximate solutions to constrained optimization problems. However, existing work on applying MCS enumeration to optimization problems focuses on the single-objective case. In this work, Pareto Minimal Correction Subsets (Pareto-MCSs) are proposed for approximating the Pareto-optimal solution set of multi-objective constrained optimization problems. We formalize and prove an equivalence relationship between Pareto-optimal solutions and Pareto-MCSs. Moreover, Pareto-MCSs and MCSs can be connected in such a way that existing state-of-the-art MCS enumeration algorithms can be used to enumerate Pareto-MCSs. Finally, experimental results on the multi-objective virtual machine consolidation problem show that the Pareto-MCS approach is competitive with state-of-the-art algorithms.

1 Introduction

Given an unsatisfiable set of constraints F , a Minimal Correction Subset (MCS) is a minimal subset of constraints C such that, if all constraints in C are removed from F , then F becomes satisfiable. MCSs enjoy a wide range of applications, such as analysis of over-constrained systems [Felfernig *et al.*, 2012; Junker, 2004], minimal model computation [Ben-Eliyahu and Dechter, 1996], interactive constraint satisfaction [O’Callaghan *et al.*, 2005], among others [Mencía *et al.*, 2015]. Many efficient MCS enumeration algorithms have been recently proposed [Bacchus *et al.*, 2014; Felfernig *et al.*, 2012; Marques-Silva *et al.*, 2013], and MCS-based approximation algorithms are able to find good quality approximations of optimal solutions efficiently [Mencía *et al.*, 2015].

*This paper is an abridged version of the paper entitled “Introducing Pareto Minimal Correction Subsets” that won a best student paper award at the 20th International Conference on Theory and Applications of Satisfiability Testing [Terra-Neves *et al.*, 2017]. For the interested reader, proofs of theoretical results presented throughout the paper are available in the full version.

However, the usage of MCSs has focused only on approximating single-objective problems.

In many scenarios, such as scheduling [Iturriaga *et al.*, 2017] or green computing [Zheng *et al.*, 2016], a decision maker may need to optimize multiple conflicting objectives [Ulungu and Teghem, 1994]. In this case, multiple optimal solutions may exist, referred to as Pareto-optimal solutions [Pareto, 1906], each of them favoring certain objectives at the expense of others. There is a wide plethora of stochastic algorithms that try to approximate the Pareto front [Deb *et al.*, 2000; Xu and Fortes, 2010]. Such approaches are non-deterministic and are known to be parameter sensitive, resulting in each different problem requiring a distinct configuration of the algorithm to achieve a competitive performance. Moreover, stochastic algorithms are known to struggle as instances become more tightly constrained. On the other hand, it is widely accepted that constraint-based methods, compared to other approaches, usually thrive in tightly constrained problems, and some exact algorithms exist for computing the set of Pareto-optimal solutions of multi-objective problems [Rayside *et al.*, 2009]. However, such solutions are impractical for large scale instances.

The main contributions of this paper are: (1) a first definition of Multi-MCSs and Pareto-MCSs, an extension of MCSs to constrained multi-objective optimization problems; (2) a proof of an equivalence relationship between Pareto-MCSs and Pareto-optimal solutions; (3) a proof of a relationship between Multi-MCSs and MCSs that allows one to use off-the-shelf MCS enumerators as Multi-MCS enumerators; (4) an extensive experimental evaluation on virtual machine consolidation instances from the Google Cluster Data project which clearly shows the suitability of the Pareto-MCS approach for finding good quality approximations of the Pareto front.

The paper is organized as follows. Section 2 introduces basic definitions used in the remainder of the paper. In section 3, a definition of Multi-MCSs and Pareto-MCSs is proposed and some of their properties are described. Experimental results showing the effectiveness of the Pareto-MCS approach are presented in section 4. Section 5 concludes the paper.

2 Preliminaries

In this section, we introduce the necessary definitions and notations that will be used throughout the rest of the paper. We start by describing Weighted Boolean Optimization (WBO)

and its multi-objective version. Next, Minimal Correction Subsets (MCSs) are defined and we review how these can be used to find solutions for WBO problems.

2.1 Weighted Boolean Optimization

Let $X = \{x_1, x_2, \dots, x_n\}$ be a set of n Boolean variables. A literal is either a variable x_i or its complement $\neg x_i$. Given an integer $k \in \mathbb{N}$, a set of m literals l_1, l_2, \dots, l_m and their respective coefficients $\omega_1, \omega_2, \dots, \omega_m \in \mathbb{N}$, a Pseudo-Boolean (PB) constraint is a linear inequality with the form:

$$\sum \omega_i \cdot l_i \bowtie k, \quad \bowtie \in \{\leq, \geq, =\}. \quad (1)$$

A weighted PB constraint is a pair (c, ω) , where c is a PB constraint and $\omega \in \mathbb{N}^+$ is the cost of not satisfying c .

Given a formula $\mathbb{F} = (F_H, F_S)$, where F_H and F_S denote sets of hard and soft weighted PB constraints respectively, the WBO problem [Manquinho *et al.*, 2009] consists of finding a complete assignment $\alpha : X \rightarrow \{0, 1\}$ that satisfies all constraints in F_H , denoted $\alpha(F_H) = 1$, and minimizes the sum of the weights of the constraints in F_S that are not satisfied. If F_H is unsatisfiable, we say that $\alpha(F_H) = 0$ for any assignment α . Analogously, given a PB constraint c , $\alpha(c) = 1$ ($\alpha(c) = 0$) denotes that α satisfies (does not satisfy) c .

Example 2.1. Let $F_H = \{(x_1 + x_2 + x_3 \leq 2)\}$ be the set of hard PB constraints and $F_S = \{(x_1 \geq 1, 4), (x_2 \geq 1, 2), (x_3 \geq 1, 3)\}$ the set of weighted soft PB constraints of a WBO instance. $\alpha_1 = \{(x_1, 1), (x_2, 0), (x_3, 1)\}$ is an optimal assignment that does not satisfy $(x_2 \geq 1)$, having a cost of 2. $\alpha_2 = \{(x_1, 0), (x_2, 1), (x_3, 1)\}$ is not optimal because it does not satisfy $(x_1 \geq 1)$, which has a weight of 4. $\alpha_3 = \{(x_1, 1), (x_2, 1), (x_3, 1)\}$ is an invalid assignment because it does not satisfy F_H .

A Multi-Objective Weighted Boolean Optimization (MOWBO) instance is composed of a set F_H of hard PB constraints and a set $O_S = \{F_{S1}, F_{S2}, \dots, F_{S1}\}$ of soft weighted PB constraint sets. Given an assignment α and a set $F_{Si} \in O_S$, let $w(F_{Si}, \alpha)$ denote the sum of the weights of the PB constraints in F_{Si} unsatisfied by α , i.e.,

$$w(F_{Si}, \alpha) = \sum_{(c, \omega) \in F_{Si}, \alpha(c)=0} \omega. \quad (2)$$

Definition 2.1. Let $\mathbb{W} = (F_H, O_S)$ be a MOWBO instance, with $O_S = \{F_{S1}, F_{S2}, \dots, F_{S1}\}$. Let $\alpha, \alpha' : X \rightarrow \{0, 1\}$ be two complete assignments such that $\alpha \neq \alpha'$ and $\alpha(F_H) = \alpha'(F_H) = 1$. We say that α dominates α' ($\alpha \prec \alpha'$) if, and only if, $\forall F_{Si} \in O_S w(F_{Si}, \alpha) \leq w(F_{Si}, \alpha')$ and $\exists F_{Si} \in O_S w(F_{Si}, \alpha) < w(F_{Si}, \alpha')$.

Definition 2.2. Let $\mathbb{W} = (F_H, O_S)$ be a MOWBO instance and $\alpha : X \rightarrow \{0, 1\}$ a complete assignment. α is Pareto-optimal if, and only if, $\alpha(F_H) = 1$ and no other complete assignment α' exists such that $\alpha'(F_H) = 1$ and $\alpha' \prec \alpha$.

In MOWBO, the goal is to find the set of Pareto-optimal [Pareto, 1906] solutions, i.e., the Pareto front.

Example 2.2. Let $F_H = \{(x_1 + x_2 \leq 1)\}$ be the set of PB constraints and $O_S = \{F_{S1}, F_{S2}\}$, where $F_{S1} = \{(\neg x_1, 1), (x_2, 2)\}$ and $F_{S2} = \{(x_1, 1)\}$, the set

x_1	x_2	$\{(\neg x_1, 1), (x_2, 2)\}$	$(x_1, 1)$
0	0	2	1
0	1	0	1
1	0	3	0
1	1	-	-

Table 1: Possible assignments and respective costs for example 2.2.

of soft constraint sets of a MOWBO instance. Table 1 shows the costs for each possible assignment. Note that $\{(x_1, 1), (x_2, 1)\}$ violates F_H . Hence, it is not a valid assignment. $\{(x_1, 0), (x_2, 0)\}$ is not Pareto-optimal because it is dominated by $\{(x_1, 0), (x_2, 1)\}$. However, $\{(x_1, 0), (x_2, 1)\}$ and $\{(x_1, 1), (x_2, 0)\}$ are Pareto-optimal solutions since they are not dominated by any other assignment that satisfies F_H .

2.2 Minimal Correction Subsets

Given an unsatisfiable set of PB constraints F , an MCS is a minimal subset $C \subseteq F$ such that $F \setminus C$ is satisfiable.

Definition 2.3. Let F be an unsatisfiable set of constraints. A subset $C \subseteq F$ is an MCS of F if, and only if, $F \setminus C$ is satisfiable and $(F \setminus C) \cup \{c\}$ is unsatisfiable for all $c \in C$.

Example 2.3. Consider the unsatisfiable set of PB constraints $F = \{(x_1 + x_2 = 1), (x_1 \geq 1), (x_2 \geq 1)\}$. F has three MCSs $C_1 = \{(x_1 \geq 1)\}$, $C_2 = \{(x_2 \geq 1)\}$ and $C_3 = \{(x_1 + x_2 = 1)\}$.

Several algorithms exist for finding MCSs [Bailey and Stuckey, 2005; Felfernig *et al.*, 2012; Marques-Silva *et al.*, 2013; Mencia *et al.*, 2015]. The following definition extends the notion of MCS to WBO instances. For simplicity, we assume that the set of hard PB constraints F_H of a WBO instance is always satisfiable, but this can be checked using a single call to a satisfiability solver.

Definition 2.4. Let $\mathbb{F} = (F_H, F_S)$ be a WBO instance, where F_H and F_S denote the hard and soft PB constraint sets respectively. A subset $C \subseteq F_S$ is an MCS of \mathbb{F} if, and only if, $F_H \cup (F_S \setminus C)$ is satisfiable and $F_H \cup (F_S \setminus C) \cup \{c\}$ is unsatisfiable for all $c \in C$.

An MCS provides an approximation to a WBO optimal solution, and in some problems is faster to compute [Ignatiev *et al.*, 2014]. Actually, the WBO problem can be reduced to finding the MCS $C \subseteq F_S$ that minimizes the sum of the weights of its constraints [Birnbaum and Lozinskii, 2003].

Example 2.4. Consider the WBO instance given by $F_H = \{(x_1 + x_2 = 1)\}$ and $F_S = \{(x_1 \geq 1, 1), (x_2 \geq 1, 2), (\neg x_1 \geq 1, 4), (\neg x_2 \geq 1, 6)\}$. It has two MCSs $C_1 = \{(x_1 \geq 1, 1), (\neg x_2 \geq 1, 6)\}$ and $C_2 = \{(x_2 \geq 1, 2), (\neg x_1 \geq 1, 4)\}$. The sum of the weights of the constraints in C_1 and C_2 is 7 and 6 respectively. C_2 is the MCS that minimizes the sum of the weights of its constraints. Therefore, any assignment that satisfies $\{(x_1 + x_2 = 1), (x_1 \geq 1), (\neg x_2 \geq 1)\}$ is an optimal solution of the WBO instance.

3 Pareto Minimal Correction Subsets

This section introduces the novel concept of Multi-MCSs and explains how they can be used to approximate the Pareto front of a MOWBO instance. First, Multi-MCSs and Pareto-MCSs are defined. Next, we describe some properties of Multi-MCSs and Pareto-MCSs.

3.1 Multi and Pareto Minimal Correction Subsets

The definition of Multi-MCS, an extension of MCSs to MOWBO formulas, builds upon the concept of MCS for WBO formulas (see section 2.2).

Definition 3.1. Let $\mathbb{W} = (F_H, O_S)$ be a MOWBO instance, with $O_S = \{F_{S1}, F_{S2}, \dots, F_{Sl}\}$. Let $\mathbb{C} = (C_1, C_2, \dots, C_l)$ be a tuple of sets such that $C_i \subseteq F_{Si}$, $1 \leq i \leq l$. \mathbb{C} is a Multi-MCS of \mathbb{W} if, and only if, $F_H \cup \bigcup_{i=1}^l (F_{Si} \setminus C_i)$ is satisfiable and $F_H \cup \bigcup_{i=1}^l (F_{Si} \setminus C_i) \cup \{c\}$ is unsatisfiable for all $c \in \bigcup_{i=1}^l C_i$.

The dominance relation between two MOWBO solutions in definition 2.1 can also be extended to pairs of Multi-MCSs.

Definition 3.2. Let $\mathbb{W} = (F_H, O_S)$ be a MOWBO instance. Let $\mathbb{C} = (C_1, C_2, \dots, C_l)$ and $\mathbb{C}' = (C'_1, C'_2, \dots, C'_l)$ be two Multi-MCSs of \mathbb{W} . We say that \mathbb{C} dominates \mathbb{C}' ($\mathbb{C} \prec \mathbb{C}'$) if, and only if, $\forall 1 \leq i \leq l \sum_{(c, \omega) \in C_i} \omega \leq \sum_{(c', \omega') \in C'_i} \omega'$ and $\exists 1 \leq j \leq l \sum_{(c, \omega) \in C_j} \omega < \sum_{(c', \omega') \in C'_j} \omega'$.

Considering the definition of dominance for Multi-MCSs, the concept of Pareto-MCS can be formalized.

Definition 3.3. Let $\mathbb{W} = (F_H, O_S)$ be a MOWBO instance and \mathbb{C} a Multi-MCS of \mathbb{W} . \mathbb{C} is a Pareto-MCS if, and only if, no other Multi-MCS \mathbb{C}' exists such that $\mathbb{C}' \prec \mathbb{C}$.

Example 3.1. Recall the MOWBO instance $\mathbb{W} = (F_H, O_S)$ from example 2.2, with $F_H = \{(x_1 + x_2 \leq 1)\}$ and $O_S = \{F_{S1}, F_{S2}\}$, where $F_{S1} = \{(-x_1, 1), (x_2, 2)\}$ and $F_{S2} = \{(x_1, 1)\}$. \mathbb{W} has two Multi-MCSs $\mathbb{C}_1 = (\{\}, \{(x_1, 1)\})$ and $\mathbb{C}_2 = (\{(-x_1, 1), (x_2, 2)\}, \{\})$, which are also Pareto-MCSs.

3.2 Properties of Multi and Pareto Minimal Correction Subsets

Recall from section 2.2 that the WBO problem can be reduced to finding an MCS that minimizes the sum of the weights of its soft constraints. Therefore, it is expected that a MOWBO instance can be reduced to finding the set of its Pareto-MCSs. The following results reveal an equivalence relationship between Pareto-MCSs and Pareto-optimal solutions.

Proposition 1. Let $\mathbb{W} = (F_H, O_S)$ be a MOWBO instance, with $O_S = \{F_{S1}, F_{S2}, \dots, F_{Sl}\}$, and α a Pareto-optimal solution of \mathbb{W} . Let $\mathbb{C} = (C_1, C_2, \dots, C_l)$, where $C_i = \{(c, \omega) : (c, \omega) \in F_{Si} \wedge \alpha(c) = 0\}$ for all $1 \leq i \leq l$. Then, \mathbb{C} is a Pareto-MCS of \mathbb{W} .

Proposition 2. Let $\mathbb{W} = (F_H, O_S)$ be a MOWBO instance, with $O_S = \{F_{S1}, F_{S2}, \dots, F_{Sl}\}$, and $\mathbb{C} = (C_1, C_2, \dots, C_l)$ a Pareto-MCS of \mathbb{W} . Then, any complete assignment α that satisfies $F_H \cup \bigcup_{i=1}^l (F_{Si} \setminus C_i)$ is a Pareto-optimal solution.

Propositions 1 and 2 show that for each Pareto-MCS there is at least one Pareto-optimal solution and that each Pareto-optimal solution has an associated Pareto-MCS. Therefore, MOWBO can be reduced to Pareto-MCS enumeration in the same way that WBO can be reduced to MCS enumeration.

Proposition 3. Let $\mathbb{W} = (F_H, O_S)$ be a MOWBO instance, with $O_S = \{F_{S1}, F_{S2}, \dots, F_{Sl}\}$, and $\mathbb{C} = (C_1, C_2, \dots, C_l)$ a Multi-MCS of \mathbb{W} . Then, $C = \bigcup_{i=1}^l C_i$ is an MCS of the WBO instance $\mathbb{F} = (F_H, \bigcup_{i=1}^l F_{Si})$.

Proposition 3 implies that Pareto-MCS enumeration of a MOWBO instance $\mathbb{W} = (F_H, O_S)$ can be reduced to enumerating MCSs of the WBO instance $\mathbb{F} = (F_H, \bigcup_{F_S \in O_S} F_S)$ as follows: (1) build \mathbb{F} from \mathbb{W} ; (2) enumerate MCSs C of \mathbb{F} using an off-the-shelf MCS enumerator and convert them to Multi-MCSs $\mathbb{C} = (C_1, C_2, \dots, C_l)$ of \mathbb{W} , where $C_i = \{(c, \omega) : (c, \omega) \in C \cap F_{Si}\}$ for all $1 \leq i \leq l$; (3) filter out Multi-MCSs dominated by any other Multi-MCS using non dominated sorting [Deb *et al.*, 2000]. Therefore, efficient state-of-the-art algorithms for MCS enumeration can be used to enumerate Pareto-MCSs right off-the-shelf.

4 Experimental Results

This section evaluates the performance of the Pareto-MCS based approach on instances of the Virtual Machine Consolidation (VMC) problem. We refer to our approach for enumerating Pareto-MCSs as PCLD, since the CLD [Marques-Silva *et al.*, 2013] algorithm was used to compute MCSs. The performance of PCLD is compared with that of the state-of-the-art evolutionary algorithms for VMC. PCLD was built on top of Sat4j-PB [Le Berre and Parrain, 2010] (version 2.3.4), and the evolutionary algorithms were implemented on top of the MOEA Framework¹ (version 2.9.1).


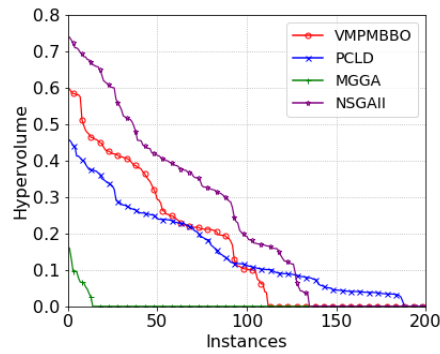

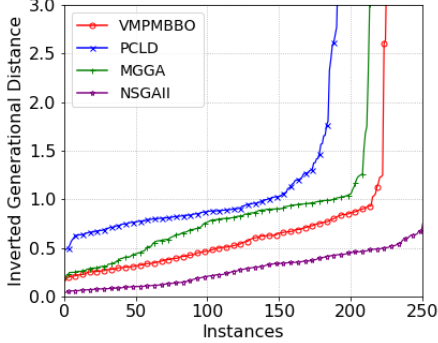
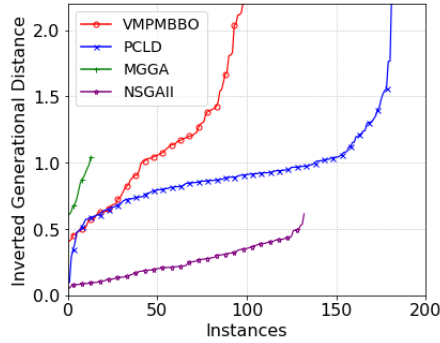
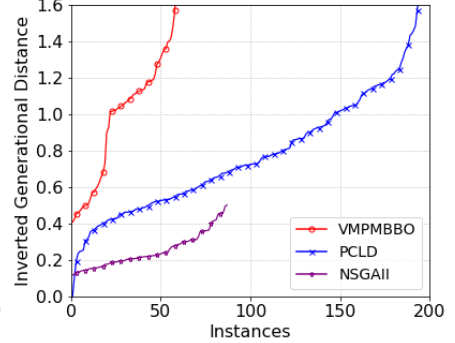
In VMC, we have several servers with fixed resource capacities and Virtual Machines (VMs) with requirements of those same resources. Each VM must be placed in some server, but server capacities cannot be exceeded and some VMs cannot be placed in the same server. There exists an initial placement, i.e., a VM can be associated with an initial server, incurring a migration cost if the VM is placed in a different one. A migration budget constraint can be used to enforce an upper limit on the migration costs, and is specified as a percentile bp of the total memory capacity of the servers. The goal is to find a placement for all VMs that satisfies the constraints and simultaneously minimizes (1) energy consumption of servers, (2) migration costs and (3) resource wastage. The latter is a measure of the imbalance of server resource usage. A detailed description of the VMC problem can be found in the literature [Terra-Neves *et al.*, 2017; Zheng *et al.*, 2016]. The evaluation is performed on publicly available benchmarks², based on subsets of workload traces randomly selected from the Google Cluster Data project³.

In the VMC problem instances, it is impractical to find the full set of Pareto-optimal solutions within a reasonable amount of time. Therefore, the evaluation process considers approximations of the Pareto-optimal solution set that each algorithm is able to produce within the time limit of 1800 seconds. The Hypervolume (HV) quality indicator [Zitzler and Thiele, 1999] provides a combined measure of convergence and diversity. Larger values of HV mean that the solution set is composed of solutions of better quality and/or diversity. The Inverted Generational Distance (IGD) indicator [Zhang and Li, 2007] measures the average Euclidean distance, in the cost space, between the Pareto optimal solutions and the approximation returned by the algorithm, and a smaller value

¹<http://moaframework.org/>

²<http://sat.inesc-id.pt/dome/>

³<http://code.google.com/p/googleclusterdata/>


 Figure 1: HV distributions ($bp = 100\%$).

 Figure 3: HV distributions ($bp = 5\%$).

 Figure 5: HV distributions ($bp = 1\%$).

 Figure 2: IGD distributions ($bp = 100\%$).

 Figure 4: IGD distributions ($bp = 5\%$).

 Figure 6: IGD distributions ($bp = 1\%$).

is preferred. Since the Pareto front is unknown, we use instead the combination of the solutions produced by all algorithms evaluated. The evaluation was conducted on an AMD Opteron 6376 (2.3 GHz) running Debian jessie and each algorithm was executed with a memory limit of 4 GB. Evolutionary algorithms were executed with 10 different seeds for each instance, and the analysis is performed using the median values over all executions.

The performance of PCLD is compared with that of the evolutionary algorithms VMPMBBO [Zheng *et al.*, 2016], MGGA [Xu and Fortes, 2010] and NSGAI [Deb *et al.*, 2000]. MGGA was adapted to consider migration costs instead of thermal dissipation and configured as suggested by Xu and Fortes [2010]. We note that VMPMBBO was originally designed for optimizing energy consumption and resource wastage and was run with the configuration suggested by Zheng *et al.* [2016]. The only difference is that we use 2 additional subsystems to account for migration costs. NSGAI is a general purpose genetic algorithm and was fine tuned to solve the VMC problem.

Figures 1 and 2 show the HV and IGD distributions, respectively, for all algorithms when $bp = 100\%$. A point (x, y) in the HV (IGD) distribution plot indicates that the given approach obtained an HV (IGD) equal to or greater (lower) than y for x instances. We can see that NSGAI has the best performance both in terms of HV and IGD, being able to find solution sets for more instances and of better quality. However, considering large migration costs is not realistic for an active data center where live migrations may result in a large performance deterioration of running applications. Therefore, considering a limited budget for migrations of VMs is more real-

istic. With migration budgets of 5% and 1%, the performance of the evolutionary algorithms degrades considerably, while PCLD remains robust, as we can see in figures 3, 4, 5 and 6. MGGA's performance degrades the most, since it is unable to find many feasible solutions. In fact, with $bp = 1\%$, MGGA fails to find feasible solutions for any instance. Moreover, PCLD is able to find feasible solutions for far more instances than the remaining algorithms.

5 Conclusion

This paper introduces the Pareto Minimal Correction Subset (Pareto-MCS) of a multi-objective constrained optimization problem. An equivalence relationship between Pareto-MCSs and Pareto-optimal solutions shows that Pareto-optimal solution enumeration can be reduced to Pareto-MCS enumeration. Additionally, we show that Pareto-MCS enumeration can be reduced to MCS enumeration, allowing the usage of off-the-shelf MCS enumeration algorithms, instead of developing entirely new algorithms for Pareto-MCS enumeration.

An experimental evaluation on instances of the Virtual Machine Consolidation (VMC) problem shows that Pareto-MCS enumeration is competitive with the state-of-the-art on a large set of problem instances. Not only is this new approach able to find solutions for more instances, but it is also able to find solution sets of higher quality in more constrained instances.

Acknowledgments

This work was supported by national funds through FCT with references UID/CEC/50021/2013, SFRH/BD/111471/2015 and CMU/AIR/0022/2017.

References

- [Bacchus *et al.*, 2014] Fahiem Bacchus, Jessica Davies, Maria Tsimpoukelli, and George Katsirelos. Relaxation search: A simple way of managing optional clauses. In *28th Conference on Artificial Intelligence, AAAI*, pages 835–841, 2014.
- [Bailey and Stuckey, 2005] James Bailey and Peter J. Stuckey. Discovery of minimal unsatisfiable subsets of constraints using hitting set dualization. In *7th International Symposium on Practical Aspects of Declarative Languages, PADL*, pages 174–186, 2005.
- [Ben-Eliyahu and Dechter, 1996] Rachel Ben-Eliyahu and Rina Dechter. On computing minimal models. *Annals of Mathematics and Artificial Intelligence*, 18(1):3–27, 1996.
- [Birbaum and Lozinskii, 2003] Elazar Birbaum and Eliezer L. Lozinskii. Consistent subsets of inconsistent systems: structure and behaviour. *Journal of Experimental & Theoretical Artificial Intelligence*, 15(1):25–46, 2003.
- [Deb *et al.*, 2000] Kalyanmoy Deb, Samir Agrawal, Amrit Pratap, and T. Meyarivan. A fast elitist non-dominated sorting genetic algorithm for multi-objective optimisation: NSGA-II. In *6th International Conference on Parallel Problem Solving from Nature*, pages 849–858, 2000.
- [Felfernig *et al.*, 2012] Alexander Felfernig, Monika Schubert, and Christoph Zehentner. An efficient diagnosis algorithm for inconsistent constraint sets. *Artificial Intelligence for Engineering Design, Analysis and Manufacturing*, 26(1):53–62, 2012.
- [Ignatiev *et al.*, 2014] Alexey Ignatiev, Mikolás Janota, and João Marques-Silva. Towards efficient optimization in package management systems. In *36th International Conference on Software Engineering*, pages 745–755, 2014.
- [Iturriaga *et al.*, 2017] Santiago Iturriaga, Bernabé Dorronsoro, and Sergio Nesmachnow. Multiobjective Evolutionary Algorithms for Energy and Service Level Scheduling in a Federation of Distributed Datacenters. *International Transactions in Operational Research*, 24(1-2):199–228, 2017.
- [Junker, 2004] Ulrich Junker. QUICKXPLAIN: preferred explanations and relaxations for over-constrained problems. In *19th National Conference on Artificial Intelligence, 16th Conference on Innovative Applications of Artificial Intelligence*, pages 167–172, 2004.
- [Le Berre and Parrain, 2010] Daniel Le Berre and Anne Parrain. The sat4j library, release 2.2. *Journal on Satisfiability, Boolean Modeling and Computation*, 7(2-3):59–6, 2010.
- [Manquinho *et al.*, 2009] Vasco M. Manquinho, João P. Marques-Silva, and Jordi Planes. Algorithms for weighted boolean optimization. In *12th International Conference on Theory and Applications of Satisfiability Testing, SAT*, pages 495–508, 2009.
- [Marques-Silva *et al.*, 2013] João Marques-Silva, Federico Heras, Mikolás Janota, Alessandro Previti, and Anton Belov. On computing minimal correction subsets. In *23rd International Joint Conference on Artificial Intelligence, IJCAI*, pages 615–622, 2013.
- [Mencía *et al.*, 2015] Carlos Mencía, Alessandro Previti, and João Marques-Silva. Literal-based MCS extraction. In *24th International Joint Conference on Artificial Intelligence, IJCAI*, pages 1973–1979, 2015.
- [O’Callaghan *et al.*, 2005] Barry O’Callaghan, Barry O’Sullivan, and Eugene C. Freuder. Generating corrective explanations for interactive constraint satisfaction. In *11th International Conference on Principles and Practice of Constraint Programming*, pages 445–459, 2005.
- [Pareto, 1906] Vilfredo Pareto. *Manuale di economia politica*, volume 13. Societa Editrice, 1906.
- [Rayside *et al.*, 2009] Derek Rayside, H-Christian Estler, and Daniel Jackson. The guided improvement algorithm for exact, general-purpose, many-objective combinatorial optimization. Technical Report Technical Report MIT-CSAIL-TR-2009-033, MIT Massachusetts Institute of Technology, 2009.
- [Terra-Neves *et al.*, 2017] Miguel Terra-Neves, Inês Lynce, and Vasco Manquinho. Introducing Pareto Minimal Correction Subsets. In *Proceedings of the 20th International Conference on Theory and Applications of Satisfiability Testing*, pages 195–211. Springer, 2017.
- [Ulungu and Teghem, 1994] Ekunda Lukata Ulungu and Jacques Teghem. Multi-objective combinatorial optimization problems: A survey. *Journal of Multi-Criteria Decision Analysis*, 3(2):83–104, 1994.
- [Xu and Fortes, 2010] Jing Xu and José A. B. Fortes. Multi-objective virtual machine placement in virtualized data center environments. In *IEEE/ACM International Conference on Green Computing and Communications, GreenCom, & International Conference on Cyber, Physical and Social Computing, CPSCom*, pages 179–188, 2010.
- [Zhang and Li, 2007] Qingfu Zhang and Hui Li. MOEA/D: A Multiobjective Evolutionary Algorithm Based on Decomposition. *IEEE Transactions on Evolutionary Computation*, 11(6):712–731, 2007.
- [Zheng *et al.*, 2016] Qinghua Zheng, Rui Li, Xiuqi Li, Nazaraf Shah, Jianke Zhang, Feng Tian, Kuo-Ming Chao, and Jia Li. Virtual machine consolidated placement based on multi-objective biogeography-based optimization. *Future Generation Computer Systems*, 54:95–122, 2016.
- [Zitzler and Thiele, 1999] Eckart Zitzler and Lothar Thiele. Multiobjective evolutionary algorithms: a comparative case study and the strength pareto approach. *IEEE Transactions on Evolutionary Computation*, 3(4):257–271, 1999.