Proceedings, Fifteenth International Conference on
Principles of Knowledge Representation and Reasoning (KR 2016)

Representative Solutions for Multi-Objective Constraint Optimization Problems

Nicolas Schwind
Transdisciplinary Research Integration Center
National Institute of Advanced Industrial
Science and Technology, Tokyo, Japan
(email: schwind@nii.ac.jp)

Tenda Okimoto
Faculty of Maritime Sciences
Kobe University, Kobe, Japan
(email: tenda@maritime.kobe-u.ac.jp)

Maxime Clement
The Graduate University for Advanced Studies
Tokyo, Japan (email: Maxime-clement@nii.ac.jp)

Katsumi Inoue
National Institute of Informatics
The Graduate University for Advanced Studies
Tokyo, Japan (email: inoue@nii.ac.jp)

Abstract

Solving a multi-objective constraint optimization problem (MO-COP) typically consists in computing all Pareto optimal solutions, which are exponentially many in the general case. This causes two problems: time complexity and lack of decisiveness. We present an approach which, given a number k of desired solutions, selects k Pareto optimal solutions that are representative of the Pareto front. We analyze the computational complexity of the underlying computational problem and provide exact and approximation procedures.

Introduction

For many computational problems involving an exponentially large number of solutions, computing a subset of alternatives may be of interest for decision-making (Hebrard et al. 2005). We address the issue on multi-objective constraint optimization problems (MO-COPs), which consist in finding assignments of variables to values which satisfy some constraints and optimize several objectives simultaneously. Typically, assignments are preferred based on the notion of Pareto optimality, thus most of standards algorithms addressing MO-COPs (Marinescu 2009) solve a given MO-COP by computing all Pareto optimal solutions. The main issues are time complexity and lack of decisiveness. Schwind et al. (2014) have proposed an approach to addess these issues. It consists in computing a restricted set of Pareto optimal solutions given some preferences among the different objectives expressed as a "weight vector." However, expressing quantitative relative importance between preferences may be hard for a user, e.g., in a context of product configuration (Sabin and Weigel 1998). This calls for an appropriate function which filters available alternatives.

A class of filtering functions called *diversities* has been introduced for Constraint Satisfaction Problems (CSP) (Hebrard et al. 2005). They are used to extract a small subset of solutions which are pairwise "distant." We argue that this notion is not desirable for MO-COPs, as the selected solutions are intuitively not "representative" of the Pareto front. To fill the gap, we take our inspiration from the problem of facility location (Ilhan and Pinar 2001) and introduce a new filtering function for MO-COPs called *representativity*.

Preliminaries

We consider a given fixed number m of *objectives*. A multi-objective constraint optimization problem (MO-COP) is a tuple $\langle \mathcal{X}, \mathcal{D}, \mathcal{C}^h, \mathcal{C}^s \rangle$, where: $\mathcal{X} = \{x_1, \dots, x_n\}$ is a set of variables; $\mathcal{D} = \{D_1, \dots, D_n\}$ is a multiset of non-empty domains for the variables; $\mathcal{C}^h$ is a finite set of *hard* constraints, i.e., for each $C^h_j \in \mathcal{C}^h$, $C^h_j \subseteq D_{i_1} \times \cdots \times D_{i_j}$ for some $\{D_{i_1}, \dots, D_{i_j}\} \subseteq \mathcal{D}$; and $\mathcal{C}^s$ is a finite set of *soft*, polyadic constraints, i.e., each $C^s_j \in \mathcal{C}^s$ is a mapping from $\{D_{i_1}, \dots, D_{i_j}\}$ to $\mathbb{N}^m$, for some $\{D_{i_1}, \dots, D_{i_j}\} \subseteq \mathcal{D}$. Each constraint from $\mathcal{C}^h \cup \mathcal{C}^s$ involves a set of variables $\mathcal{X}_j \subseteq \mathcal{X}$ called its *scope*.

Let P be an MO-COP $\langle \mathcal{X}, \mathcal{D}, \mathcal{C}^h, \mathcal{C}^s \rangle$. An *assignment* A of P associates each $x_i \in X$ with a value from D_i; A is a *solution* of P if there is no $C^h_j \in \mathcal{C}^h$ such that $(A(x_{i_1}), \dots, A(x_{i_j})) \in C^h_j$, where $\{x_{i_1}, \dots, x_{i_j}\}$ is the scope of C^h_j. $Sols(P)$ denotes the set of solutions of P. Given an m-vector U, we denote by U^k or $U(k)$ its k^{th} component, with $k \in \{1, \dots, m\}$. The *cost vector* of A is the m-vector denoted by $V(A)$ defined for each $k \in \{1, \dots, m\}$ as $V(A)^k = \sum_{C^s_j \in \mathcal{C}^s} C^s_j(A(x_{i_1}), \dots, A(x_{i_j}))(k)$, where for each C^s_j, $\{x_{i_1}, \dots, x_{i_j}\}$ is the scope of C^s_j. When $\mathcal{S}$ is a set of solutions, $V(\mathcal{S})$ denotes the set $\{V(A) \mid A \in \mathcal{S}\}$.

Let $\preceq_m$ be the product ordering over $\mathbb{N}^m$, i.e., $\forall V_1, V_2 \in \mathbb{N}^m$, $V_1 \preceq_m V_2$ iff $\forall k \in \{1, \dots, m\}$, $V^k_1 \leq V^k_2$. The preordering $\preceq_{Par}$ over $Sols(P)$, called the Pareto dominance relation, is defined $\forall A, A' \in Sols(P)$ as $A \preceq_{Par} A'$ iff $V(A) \preceq_m V(A')$; we say that A Pareto dominates A'. A *Pareto optimal solution* of P is a solution $S \in Sols(P)$ which is not strictly Pareto dominated by an other solution $S' \in Sols(P)$. $\mathcal{S}_{Par}(P)$ denotes the set of Pareto optimal solutions of P, and $\mathsf{PF}(P) = \bigcup_{S \in \mathcal{S}_{Par}(P)} V(S)$ is called the *Pareto front* of P.

Example 1. *Let $P_\star = \langle \mathcal{X}_\star, \mathcal{D}_\star, \mathcal{C}^h_\star, \mathcal{C}^s_\star \rangle$ be the bi-objective MO-COP defined as $\mathcal{X}_\star = \{x_1, x_2, x_3\}$, $D_i = \{a, b\}\ \forall D_i \in D_\star$, $\mathcal{C}^h_\star = \{\{(b, b, a)\}\}$ and $\mathcal{C}^s_\star = \{C^s_1, C^s_2\}$, defined as:*

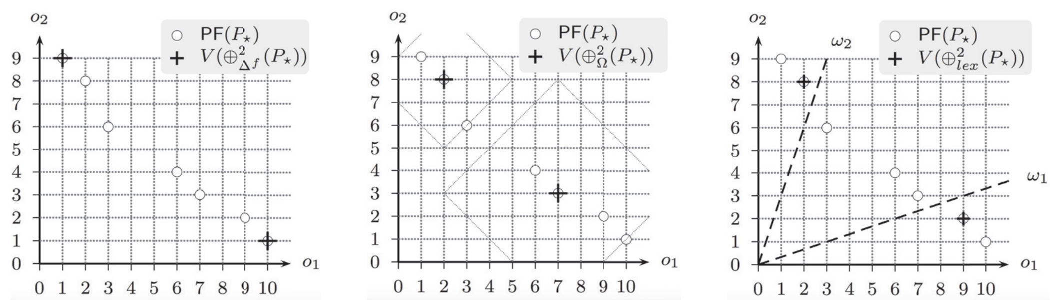

(a) The Pareto front of $P_\star$ and the most diverse solutions ($k = 2$).

(b) An optimal set of representative solutions for $P_\star$ ($k = 2$).

(c) The solutions of $P_\star$ from the approximation approach ($k = 2$).

Figure 1: The Pareto front of $P_\star$ and results for exact and approximation approaches (for $k = 2$).

(x_1, x_2)	C_1^s	(x_2, x_3)	C_2^s
(a, a)	$(1, 5)$	(a, a)	$(0, 4)$
(a, b)	$(2, 1)$	(a, b)	$(1, 3)$
(b, a)	$(6, 0)$	(b, a)	$(1, 5)$
(b, b)	$(3, 0)$	(b, b)	$(7, 1)$

Here, $\mathcal{S}_{Par}(P_\star) = Sols(P_\star) = \{a, b\}^3 \setminus \{(b, b, a)\}$. *For instance* $V((aaa)) = (1, 5) + (0, 4) = (1, 9)$. *The Pareto front of* $P_\star$ *can be seen in Figure 1(a).*

An MO-COP operator $\oplus$ associates with an MO-COP P a subset of Pareto optimal solutions from $Sols(P)$ (Schwind et al. 2014). Here, a solution is evaluated through its cost vector only, thus we focus on operators $\oplus$ which satisfy $A, A' \in \oplus(P) \wedge A \neq A' \implies V(A) \neq V(A')$.

Representative Solutions

In general, the size of the Pareto front is exponential in the size of the MO-COP, so providing all Pareto optimal solutions lacks decisiveness. So we now introduce a general class of "filtering" MO-COP operators, which given an integer k associate with every MO-COP P a subset of its Pareto optimal solutions such that $|\oplus^k(P)| = k$. These operators can be characterized by a *filtering function* Γ associating a set of assignments with a number to be optimized:[1]

Definition 1 (Filtering MO-COP operator)**.** *Let k be an integer and Γ be a filtering function. The* $\langle\Gamma, k\rangle$-filtering MO-COP operator $\oplus_\Gamma^k$*, is defined as*

$$\oplus_\Gamma^k(P) = \gamma(\text{argmin}\{\Gamma(\mathcal{S}) \mid \mathcal{S} \subseteq \mathcal{S}_{Par}(P), |\mathcal{S}| = k\}),$$

where γ is any choice function.

Diversities (Hebrard et al. 2005) constitute a specific class of filtering functions. A diversity is characterized by a distance δ between assignments and an aggregation function f (i.e., a function from $\mathbb{R} \times \cdots \times \mathbb{R}$ to $\mathbb{R}$, e.g., the max and summation functions). Without loss of generality, we consider the standard distance $\delta = \delta_1$, i.e., the L_1-norm defined for all $V_1, V_2 \in \mathbb{R}^m$ as $\delta_1(V_1, V_2) = \sum_{k=1}^m |V_1^k - V_2^k|$.

Definition 2 (Diversity (Hebrard et al. 2005))**.** *Given an aggregation function f, the* diversity Δ^f *is the filtering function defined for every set* $\mathcal{S} \subseteq \mathcal{S}_{Par}(P)$ *as*

$$\Delta^f(\mathcal{S}) = f\{\delta_1(V(A), V(A')) \mid A, A' \in \mathcal{S}\}.$$

[1]The definition is given in the case where Γ is to be minimized. The maximization counterpart is defined similarly.

Note that a diversity is a function to be maximized.

We claim that a diversity Δ^f is not be appropriate for MO-COPs. First, the choice of the aggregation function f is arbitrary; obviously enough, different aggregation functions induce different filtering functions. Second, diversity-based MO-COP operators may output only "unbalanced" solutions, i.e., they may not provide alternatives with parsimonious trade-offs among the objectives. Third, such operators are not "representative" of the Pareto front, as shown in our running example:

Example 1 (**continued**)**.** *For any aggregation function f, we get that* $V(\oplus_{\Delta^f}^2(P_\star)) = \{(10, 1), (1, 9)\}$. *One can see from Figure 1(a) that these vectors are the most "unbalanced" ones from the Pareto front, and that whichever the vectors lying "between" these two, the result remains unchanged.*

As a concrete example, consider a Pareto front representing the outcomes of second-hand cars available for sale, where the two objectives to be minimized represent the price and the age of a car. Without any prior knowledge of the possible alternatives a customer may ask for two options; then any diversity-based filtering MO-COP operator $\otimes_{\Delta^f}^2$ will recommend only the newest and the oldest available cars despite the presence of balanced trade-off solutions.

We introduce now a parsimonious filtering function which is directly inspired from the well-known discrete p-center problem in discrete location theory (Ilhan and Pinar 2001). This problem consists of locating p facilities in a network and assigning clients to them so as to minimize the maximum distance between any client and the facility she is assigned to. We adapt the notion to MO-COPs:

Definition 3 (Representativity)**.** *The* representativity Ω *is the filtering function defined for every* $\mathcal{S} \subseteq \mathcal{S}_{Par}(P)$ *as*

$$\Omega(\mathcal{S}) = \max_{A \in \mathcal{S}_{Par}(P)} \min_{A' \in \mathcal{S}} \delta_1(V(A), V(A')).$$

$\Omega(\mathcal{S})$ *is called the* radius *of* $\mathcal{S}$.

The rationale behind a representativity Ω is to extract from an MO-COP P a subset $\mathcal{S}_k$ of Pareto optimal solutions such that $|\mathcal{S}_k| = k$ so as to *minimize* $\Omega(\mathcal{S})$, the radius of $\mathcal{S}_k$, that is the maximum distance between any Pareto optimal solution from $\mathsf{PF}(P)$ and the Pareto optimal solution from $\mathcal{S}_k$ which is the closest. So a representativity is a filtering function to be minimized. From now on, we focus on the representativity-based filtering MO-COP operators $\oplus^k_\Omega$. Note that these operators, unlike diversity-based ones, are more "sensitive" to the shape of the Pareto front:

Example 1 (continued). *We have* $V(\oplus^2_\Omega(P_\star)) = \{(7,3),(2,8)\}$ *or* $V(\oplus^2_\Omega(P_\star)) = \{(9,2),(2,8)\}$ *(only the former case is depicted in Figure 1(b)). In both cases, the radius is equal to* 5.

Using a representativity-based MO-COP operator provides us with the guarantee that no Pareto optimal solution is left apart: for each one of them there is a representative solution that is not "too far" from it.

Computational Complexity of $\oplus^k_\Omega$

We assume that the reader is familiar with the complexity class NP. Higher complexity classes are defined using oracles. In particular, $\Sigma^P_2 = \mathsf{NP}^{\mathsf{NP}}$ corresponds to the class of decision problems that are solved in non-deterministic polynomial time by deterministic Turing machines using an oracle for NP in polynomial time (Papadimitriou 1994).

We investigate the computational complexity of two decision problems inherent in our representativity-based MO-COP operator $\oplus^k_\Omega$. We assume that k is bounded by a polynomial in the size of the input (Hebrard et al. 2005). The first problem DP1 considers that the Pareto front is given in input (for instance, when it is computed as a preprocessing). The second problem DP2 does not require any prior processing step on the input MO-COP:

Definition 4 (DP1)**.** *Given an MO-COP P, its Pareto front* $\mathsf{PF}(P)$ *and two integers α, k, does there exist $\mathcal{S} \subseteq \mathcal{S}_{Par}(P)$ such that $|\mathcal{S}| = k$ and $\Omega(\mathcal{S}) \leq \alpha$?*

Definition 5 (DP2)**.** *Given an MO-COP P and two integers α, k, does there exist $\mathcal{S} \subseteq \mathcal{S}_{Par}(P)$ such that $|\mathcal{S}| = k$ and $\Omega(\mathcal{S}) \leq \alpha$?*

Proposition 1. DP1 *is* NP*-complete.*

Without stating it formally, we claim that the corresponding diversity problem (Hebrard et al. 2005) in the case where the Pareto front is given in input, is also NP-hard, even for aggregation functions f computable in polynomial time. Thus for DP1, considering a representativity measure instead of a diversity does not result in a complexity shift. This does hold anymore for DP2 which lies in the second level of the polynomial hierarchy, whereas the counterpart diversity problem is NP-complete (Hebrard et al. 2005):

Proposition 2. DP2 *is* Σ^{P}_2*-complete.* Σ^{P}_2*-hardness holds even when* $k = 1$.

A consequence of Proposition 2 is that unless the polynomial hierarchy collapses at the first level, DP2 cannot be solved using a standard NP problem solver. However, Proposition 1 shows that if the Pareto front is part of the input (thus computed in a first step), one can take advantage of an NP-oracle to compute $\oplus^k_\Omega(P)$. The (so-called "exact") procedure consists in three phases. First, we compute all Pareto optimal solutions using a standard branch-and-bound technique (Marinescu 2009). Second, we randomly generate a k-subset $\mathcal{S}$ of $\mathcal{S}_{Par}(P)$ and set $\alpha_{up} = \Omega(\mathcal{S})$. Third, we adapt an efficient encoding of the vertex p-center problem proposed in (Ilhan and Pinar 2001) into our MO-COP framework. Such an encoding is parameterized by an integer α (a specific radius) and serves as an NP solver for the decision problem DP1. The result with the minimum radius α_{min}, i.e., $\oplus^k_\Omega(P)$, is found by calling the oracle iteratively: we adjust the radius α at each call using a dichotomic search between 0 and α_{up}, the initial upper bound.

Approximation Operator $\oplus^k_{lex}$

We now introduce an MO-COP operator approximating $\oplus^k_\Omega$ (in the sense of the associated radius). Despite the Σ^P_2-hardness of DP2, our approximating MO-COP operator, denoted $\oplus^k_{lex}$, can be computed much more efficiently, especially for a high number of objectives where the "exact" approach does not work anymore. $\oplus^k_{lex}$, intuitively "targets" k specific areas of the Pareto front without computing it explicitly. The notion of "weight vector" (Torra 1997) is at the core of the idea: it is an m-vector $\omega = (\omega^1, \ldots, \omega^m)$ such that $\omega \in]0,1]^m$ and $\sum_{k=1}^m \omega^k = 1$. The set $\mathcal{W}_m$ denotes the set of all m-weight vectors. For instance, $(0.3, 0.6, 0.1)$ and $(0.45, 0.45, 0.1)$ are both 3-weight vectors. We take advantage of a direct adaptation of our representativity measure (cf. Definition 3) to weight vectors, defined as follows:

Definition 6 ($\mathcal{W}_m$-representativity)**.** *The $\mathcal{W}_m$-representativity $\Omega^{\mathcal{W}_m}$ is the mapping from $2^{\mathcal{W}_m}$ to $\mathbb{R}^+$ defined for every set of m-weight vectors $\mathcal{W}' \subseteq \mathcal{W}_m$ as*

$$\Omega^{\mathcal{W}_m}(\mathcal{W}') = \max_{\omega \in \mathcal{W}_m} \min_{\omega' \in \mathcal{W}'} \delta_1(\omega, \omega').$$

$\Omega^{\mathcal{W}_m}(\mathcal{W}')$ *is called the* radius *of* $\mathcal{W}'$.

The induced $\mathcal{W}_m$-representativity-based filtering function is given as follows:

Definition 7 (Weight vector filtering)**.** *Given two integers k, m, the weight vector filtering $\odot^k_m$ is defined as*

$$\odot^k_m = \gamma(\mathrm{argmin}\{\Omega^{\mathcal{W}_m}(\mathcal{W}') \mid \mathcal{W}' \subset \mathcal{W}_m, |\mathcal{W}'| = k\}),$$

where γ is any choice function.

Example 1 (continued)**.** *In $P_\star$, we have $m = 2$. Let $k = 2$. We have that $\odot^2_2 = \{\omega_1, \omega_2\}$, with $\omega_1 = (0.25, 0.75)$, $\omega_2 = (0.75, 0.25)$. Figure 1(c) graphically depicts two dashed lines associated respectively with ω_1 and ω_2, where for each $\omega_i \in \{\omega_1, \omega_2\}$, the line associated with ω_i is the line characterized by the set $\{(V_1, V_2) \mid \omega^i_1.V_1 = \omega^i_2.V_2\}$.*

Note that a weight vector filtering is solely characterized by k and m, so one can assume that $\odot^k_m$ is computed as a preprocessing step: this can be done analytically, or by finely discretizing the set $\mathcal{W}_m$ and take advantage of the exact procedure described in the previous section.

The MO-COP operator $\oplus^k_{lex}$ we are going to introduce (in Definition 8 below) uses the set $\odot^k_m$. Intuitively and graphically speaking, the weight vectors from $\odot^k_m$ are used to "target" different areas of the search space in the most distributed fashion. Indeed, the operator $\oplus^k_{lex}$ will associate with an MO-COP P and each weight vector ω_i from $\odot^k_m$ a Pareto optimal solution of P which is "close" to ω_i,

Table 1: CPU time and value of radius obtained from MO-COPs with $n = 16$ and $k = 7$.

m	$\lvert\mathsf{PF}(P)\rvert$	CPU time (in sec.)		radius	
		$\oplus^7_\Omega(P)$ (PF / opt)	$\oplus^7_{lex}(P)$	$\oplus^7_\Omega(P)$	$\oplus^7_{lex}(P)$
2	15	2.34 / 0.26	2.80	156	245
3	98	2.97 / 0.62	3.01	248	498
4	467	4.70 / 7.40	3.30	417	802
5	1301	11.59 / 13.5	3.72	752	1280
6	2709	28.79 / 540	4.04	966	1488
7	4903	73.38 / time out	4.34	time out	1702
8	8210	139.2 / time out	4.65	time out	1934
9	14820	152.1 / time out	4.65	time out	2196

in attempt to return a set of k representative Pareto optimal solutions. For this purpose, we take advantage of *ω-weighted egalitarian operators* $\oplus_\omega$ introduced in (Schwind et al. 2014): for each weight vector ω, $\oplus_\omega$ is defined for every MO-COP P as $\oplus_\omega(P) = \min(Sols(P), \preceq^{lex}_\omega)$, where $\preceq^{lex}_\omega$ is the total preorder over $Sols(P)$ such that $\forall A, A' \in Sols(P)$, $A \preceq^{lex}_\omega A'$ if and only if $V(A)_{<,\omega} \leq_{lex} V(A')_{<,\omega}$, where $\leq_{lex}$ is the lexicographic ordering induced by the natural ordering, and $V(A)_{<,\omega}$ is the vector of numbers $\alpha^1, \ldots, \alpha^m$ obtained by sorting in a non-increasing order the vector $(V(A)^1/\omega(1), \ldots, V(A)^m/\omega(m))$ (see (Schwind et al. 2014) for more details). An ω-weighted egalitarian operator exhibits a number of interesting properties. In particular, it returns the Pareto optimal solutions that are the "closest" to some utopia point located on the line specified by ω w.r.t. some a specific weighted norm. Hence, one can use this operator to efficiently compute a Pareto optimal solution which "targets" a specific area of the Pareto front specified by ω. Our approximation MO-COP operator $\oplus^k_{lex}$ is characterized by $\odot^k_m$ and the ω_i-weighted egalitarian operators $\oplus_{\omega_i}$ for all $\omega_i \in \odot^k_m$:

Definition 8 ($\oplus^k_{lex}$)**.** *Given an integer k, the MO-COP operator $\oplus^k_{lex}$ is defined as*

$$\oplus^k_{lex}(P) = \{\gamma(\oplus_{\omega^i}(P)) \mid \omega^i \in \odot^k_m\},$$

where γ is any choice function.

Example 1 (continued)**.** *For $\omega_1 = (0.25, 0.75)$, we have $V(\oplus_{\omega_1}(P_\star)) = \{(9,2)\}$ (cf. Figure 1(c)). Similarly, for $\omega_2 = (0.75, 0.25)$, we have $V(\oplus_{\omega_2}(P_\star)) = \{(2,8)\}$. Therefore, we get that $V(\oplus^2_{lex}) = \{(9,2),(2,8)\}$.*

Empirical Results

We empirically evaluated the efficiency and the radius obtained from computing both $\oplus^k_\Omega$ and $\oplus^k_{lex}$. We conducted experiments on graph-coloring problems, a well-known benchmark problem which has been adapted to MO-COPs (Fave et al. 2011). We considered MO-COP instances with 16 binary variables and set $k = 7$. We carried out all experiments on one core running at 2.67GHz with 12GB RAM. We varied the number of objectives m from 2 to 9. For each m, all values represent an average on 100 MO-COPs instances.

We used the latest version of CPLEX (version 12.6) to solve each NP-hard problem involved in the computation of $\oplus^7_\Omega$. We fixed a time out of 100 seconds for each call of the CPLEX solver. For example, MO-COP instances with $m = 6$ had (in average) a Pareto front of size 2709; computing the Pareto front took 28 seconds, and the exact operator $\oplus^7_\Omega$ and the approximation operator $\oplus^7_{lex}$ respectively took 568 seconds (in total) and 4 seconds. No instance from $m = 7$ objectives could be solved by the exact procedure $\oplus^7_\Omega$. The approximation procedure $\oplus^7_{lex}$ is not only much faster than just computing the Pareto front itself, but offers an acceptable alternative for problems with a high number of objectives. On the other hand, $\oplus^7_{lex}$ provides us with a reasonable approximation of the optimal radius.

Conclusion

We introduced the notion of representative solutions in multi-objective constraint optimization problems (MO-COPs), which are of interest in a decision-making context where no preference among the different objectives is available. We investigated the computational complexity of some associated decision problems and provided exact and approximation procedures to compute such representative sets.

In perspective, we will focus on local search algorithms to approximate representative sub-optimal solutions w.r.t. the Pareto dominance relation. Such techniques have been proposed for multi-objective optimization problems (MOOPs) (Paquete, Schiavinotto, and Stützle 2007) where the domain of variables is continuous, and more efficient to solve real-world problems. Based on these techniques, we plan to address more realistic MO-COP instances such as time tabling (Burke et al. 2012) which typically involve many objectives.

Acknowledgements

The authors would like to thank Sébastien Konieczny for his comments on the first draft of this paper.

References

Burke, E.; Marecek, J.; Parkes, A.; and Rudova, H. 2012. A branch-and-cut procedure for the udine course timetabling problem. *Annals of Operations Research* 194(1):71–87.

Fave, F. M. D.; Stranders, R.; Rogers, A.; and Jennings, N. R. 2011. Bounded decentralised coordination over multiple objectives. In *AAMAS'11*, 371–378.

Hebrard, E.; Hnich, B.; O'Sullivan, B.; and Walsh, T. 2005. Finding diverse and similar solutions in constraint programming. In *AAAI'05*, 372–377.

Ilhan, T., and Pinar, M. C. 2001. An efficient exact algorithm for the vertex p-center problem. Technical report, Bilkent University, Department of Industrial Engineering, Ankara, Turkey.

Marinescu, R. 2009. Exploiting problem decomposition in multi-objective constraint optimization. In *CP'09*, 592–607.

Papadimitriou, C. M. 1994. *Computational complexity*. Reading, Massachusetts: Addison-Wesley.

Paquete, L.; Schiavinotto, T.; and Stützle, T. 2007. On local optima in multiobjective combinatorial optimization problems. *Annals of Operations Research* 156(1):83–97.

Sabin, D., and Weigel, R. 1998. Product configuration frameworks - a survey. *IEEE Intelligent Systems* 13(4):42–49.

Schwind, N.; Okimoto, T.; Konieczny, S.; Wack, M.; and Inoue, K. 2014. Utilitarian and egalitarian solutions for multi-objective constraint optimization. In *ICTAI'14*, 170–177.

Torra, V. 1997. The weighted OWA operator. *International Journal of Intelligent Systems* 12(2):153–166.