Representativity versus Diversity: Focusing on Specific Solutions in Multi-Objective Contraint Optimization Problems

Nicolas Schwind[*1*2] Tenda Okimoto[*3*1] Maxime Clément[*4*2] Katsumi Inoue[*2*4]

[*1] Transdisciplinary Research Integration Center [*2] National Institute of Informatics

[*3] Faculty of Maritime Sciences, Kobe University

[*4] The Graduate University for Advanced Studies

Solving a multi-objective constraint optimization problem (MO-COP) typically consists in computing the set of all Pareto optimal solutions, which is exponentially large in the general case. Besides the time complexity concern, the other main drawback of this process is its lack of decisiveness, leading to thousands of solutions even for problems with a simple structure. Taking our inspiration from the well-known vertex center problem in discrete location theory, we present a procedure which, given a number k of desired solutions, returns k Pareto optimal solutions which are well-distributed and representative of the Pareto front. Compared to previous approaches, we show that our procedure exhibits a desirable behavior in the context of MO-COPs. We analyze the computational complexity of the underlying computational problem and provide both exact and approximation procedures.

1. Introduction

We address the issue on multi-objective constraint optimization problems (MO-COPs), which is the problem to find assignments of variables to values which satisfy some constraints and optimize several objectives simultaneously. Typically, the notion of "optimality" for such assignments is based on the notion of Pareto optimality, thus most of standards algorithms addressing MO-COPs [10, 11, 7] solve a given MO-COP by computing the whole set of Pareto optimal solutions. But this causes two problems: time complexity and lack of decisiveness.

A recent approach for MO-COPs has been proposed to address the decisiveness issue [13]. It consists in computing a restricted set of Pareto optimal solutions given some preferences among the different objectives, expressed as a weighted vector. However, expressing quantitative relative importance between preferences may be cumbersome for the end user, e.g., in a context of product configuration [12, 6]. Indeed, not more than 7 alternatives can be pairwise compared by a user [5]. This calls for an appropriate function which filters available alternatives. A class of filtering functions called *diversities* has been introduced in [3, 1]. A diversity extracts a small subset of solutions which are pairwise "distant." In this paper, we argue that this notion is not desirable for MO-COPs, in the sense that reasonable trade-off solutions may be missed. Taking our inspiration from facility location [8, 2, 4], we introduce a new filtering function for MO-COPs called *representativity.* Unlike diverse solutions, representative solutions provide implicit information about the shape of the Pareto front.

2. Preliminaries on MO-COPs

We consider m elements $o_1, \ldots, o_m$ called *objectives.* A multi-objective constraint optimization problem (MO-COP) is a tuple $\langle \mathcal{X}, \mathcal{D}, \mathcal{C} \rangle$, where: $\mathcal{X} = \{x_1, \ldots, x_n\}$ is a set of variables; $\mathcal{D} = \{D_1, \ldots, D_n\}$ is a multiset of non-empty domains for the variables; $\mathcal{C}$ is a finite set of polyadic constraints (i.e., each $C_j \in \mathcal{C}$ is a mapping from some specific set of domains $\mathcal{D}_j \subseteq \mathcal{D}$ to $\mathbb{N}^m \cup \{\infty\}$). Each constraint C_j involves a set of variables $\mathcal{X}_j \subseteq \mathcal{X}$ called the *scope* of C_j.

Let P be an MO-COP $\langle \mathcal{X}, \mathcal{D}, \mathcal{C} \rangle$. An *assignment* A of P associates each variable $x_i \in X$ with a value $\alpha_i \in D_i$. An assignment A is *forbidden* by a constraint $C_j \in \mathcal{C}$ if $C_j(A(x_{i_1}, \ldots, x_{i_j})) = \infty$, where $\{x_{i_1}, \ldots, x_{i_j}\}$ is the scope of C_j. An assignment of P is a *solution* if no constraint from $\mathcal{C}$ forbids it. $Sols(P)$ denotes the set of all solutions of P. The *cost vector* of a solution $S \in Sols(P)$ is the m-vector denoted by $V(S) = (V^1(S), \ldots, V^m(S))$ defined for each $k \in \{1, \ldots, m\}$ as $V^k(S) = \sum_{C_j \in \mathcal{C}} C_j(S(x_{i_1}, \ldots, x_{i_j}))$, where $\{x_{i_1}, \ldots, x_{i_j}\}$ is the scope of C_j.

"Solving" an MO-COP P typically consists in providing the *Pareto optimal* solutions from $Sols(P)$. 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 $\forall k \in \{1, \ldots, m\}$, $V_1^k \leq V_2^k$. The preordering $\preceq_{Par}$ over $Sols(P)$, called the Pareto dominance relation, is defined $\forall S, S' \in Sols(P)$ as $S \preceq_{Par} S'$ iff $V(S) \preceq_m V(S')$; we say that S Pareto dominates S'. A *Pareto optimal solution* of P is a solution $S \in Sols(P)$ which is not strictly Pareto dominated by any other solution S'. $\mathcal{S}_{Par}(P)$ denotes the set of Pareto optimal solutions, and $\mathsf{PF}(P) = \bigcup_{S \in \mathcal{S}_{Par}(P)} V(S)$ is called the *Pareto front.*

Example 1. *Consider the bi-objective MO-COP* $P_\star = \langle \mathcal{X}_\star, \mathcal{D}_\star, \mathcal{C}_\star \rangle$ *such that* $\mathcal{X}_\star = \{x_1, x_2, x_3\}$, $D_i = \{a, b\}$ $\forall D_i \in D_\star$ *and* $\mathcal{C}_\star = \{C_1, C_2, C_3\}$ *defined as follows:*

(x_1, x_2)	C_1	(x_2, x_3)	C_2	(x_1, x_2, x_3)	C_3
(a, a)	$(1, 5)$	(a, a)	$(0, 4)$	(b, b, a)	∞
(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 cost vectors of the rest of (Pareto optimal) solutions of* $P_\star$ *are as follows:* $V((aab)) = (2, 8)$*;* $V((aba)) = (3, 6)$*;* $V((abb)) = (9, 2)$*;* $V((baa)) = (6, 4)$*;* $V((bab)) = (7, 3)$*;* $V((bbb)) = (10, 1)$*. The Pareto front of* $P_\star$ *can be seen in Figure 1.*

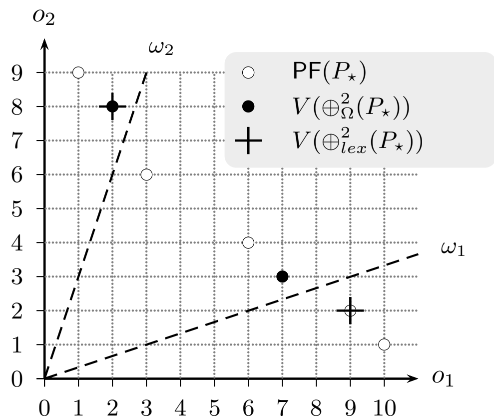

Figure 1: The Pareto front of $P_\star$ and results for exact and approximation approaches for $k = 2$.

An MO-COP operator $\oplus$ associates with an MO-COP P a subset of Pareto solutions from $Sols(P)$ [13]. Typically, the "quality" of a solution is evaluated through its cost vector only, thus we focus on operators $\oplus$ which satisfy $S, S' \in \oplus(P) \wedge S \neq S' \implies V(S) \neq V(S')$.

3. Representative Solutions

It can be easily seen that in general, the size of the Pareto front is exponential in the size of the problem: consider a simple bi-objective problem P, i.e., $m = 2$ with n binary variables (i.e., for instance each $D_i = \{a, b\}$) and n unary constraints C_j of scope $\{x_j\}$ defined as $C_j(a) = (0, 2^m)$ and $C_j(b) = (2^m, 0)$, then one can easily verify that all 2^n assignments are Pareto optimal solutions of P, thus $|\oplus(P)| = 2^n$. In Example 1, all solutions of $P_\star$ are Pareto optimal. Therefore, providing all Pareto optimal solutions implies a lack of decisiveness. In product configuration [12, 6] where a user (e.g., a customer) is in charge of selecting her preferred choice among the given alternatives, not more than 7 options can be rationally handled by an end-user customer [5]. We thus 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 $\oplus^k_\Gamma$, *is defined as*

$$\oplus^k_\Gamma(P) = \gamma(\arg\min\{\Gamma(\mathcal{S}) \mid \mathcal{S} \subseteq \mathcal{S}_{Par}(P), |\mathcal{S}| = k\}),$$

where γ is any choice function.

A class of filtering functions, called *diversities*, has been introduced in the framework of Constraint Satisfaction Problems (CSPs) [3] and logic programs [1]. 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 $\mathbb{R}$). 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 $\delta_1(V_1, V_2) = \sum_{k=1}^m |V_1^k - V_2^k|$.

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

Definition 2 (Diversity [3]). *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*

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

Note that a diversity is a function to be maximized.

We claim here that a diversity Δ^f is not be appropriate for MO-COPs. First, the aggregation function f is arbitrarily chosen, e.g., the standard max and summation functions [3]. But the consequence of such a choice for f is unclear; obviously enough, different aggregation functions induce different filtering functions. Second, a diversity-based MO-COP operator is not "representative" of the Pareto front; indeed, it can be seen from Definition 2 that a diversity associates with a subset $\mathcal{S}$ of Pareto optimal solutions a value which is independent from the shape of the Pareto front. Third, such operators can output "extreme" solutions, i.e., which do not provide a reasonable trade-off among the objectives:

Example 1 (continued). *For any aggregation function f, we get that $V(\oplus^2_{\Delta^f}(P_\star)) = \{(10, 1), (1, 9)\}$. One can see from Figure 1 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.*

Such undesirable behavior also occurs for instances with a higher number of objectives. As a concrete example, consider the case where the Pareto front represents the outcomes of Pareto optimal second-hand cars available for sale, where the two objectives to be minimized respectively represent the price and the age of a car. Without any prior knowledge of the possible alternatives a customer may ask for two options. Any diversity-based filtering MO-COP operator $\otimes^2_{\Delta^f}$ will recommend only the newest and the oldest available cars, though better trade-offs are available.

We introduce a more parsimonious filtering function which is directly inspired from the well-known discrete p-center problem in discrete location theory [8, 2, 4]. 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*

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

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

A representativity is a function to be minimized. Note that compared to a diversity, a representability does not depend only on a given a set $\mathcal{S}$ of Pareto optimal solutions, but also on the structure 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). In both cases, the radius is equal to* 5.

It can be seen that a representativity-based MO-COP operator offers better trade-offs than a diversity-based one. No Pareto solution is left apart: intuitively, for each one of them there is a solution among the proposed ones which is not "too far" from it.

4. Computational Complexity of $\oplus_\Omega^k$

We assume that the reader is familiar with the complexity class NP (see [9] for more details). Higher complexity classes are defined using oracles. In particular, $\Sigma_2^P = \mathsf{NP}^{\mathsf{NP}}$ 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.

We investigate the computational complexity of two decision problems inherent in our representativity-based MO-COP operator $\oplus_\Omega^k$. We assume that k is bounded by a polynomial in the size of the input [3, 1]. 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)**.**

- ***Input:*** *An MO-COP P, its Pareto front* $\mathsf{PF}(P)$ *and two integers α, k.*
- ***Question:*** *Does there exist $\mathcal{S} \subseteq \mathcal{S}_{Par}(P)$ such that $|\mathcal{S}| = k$ and $\Omega(\mathcal{S}) \leq \alpha$?*

Definition 5 (DP2)**.**

- ***Input:*** *An MO-COP P and two integers α, k.*
- ***Question:*** *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 [3] 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 our second problem: DP2 lies in the second level of the polynomial hierarchy, whereas the counterpart diversity problem is NP-complete [3]:

Proposition 2. DP2 *is* Σ_2^{P}*-complete.* Σ_2^{P}*-hardness holds even when $k = 1$.*

Let us provide a brief description of a procedure for computing $\oplus_\Gamma^k(P)$, i.e., the optimization problem corresponding to DP2. This so-called "exact procedure" consists in three phases. We first use any algorithm to compute the Pareto front [10, 11, 7]; this can be made for instance using a standard branch-and-bound technique. Second, we randomly generate a k-subset $\mathcal{S}$ of $\mathcal{S}_{Par}(P)$ and define α_{up} to be the radius of $\mathcal{S}$, i.e., $\alpha_{up} = \Omega(\mathcal{S})$. Third, we adapt an efficient encoding of the vertex p-center problem proposed in [4] into our MO-COP framework. Such an encoding is parameterized by an integer α (a specific radius) and serves as an NP-oracle the decision problem DP1. The result with the minimum radius α_{min}, i.e., $\oplus_\Omega^k(P)$, is found by calling the oracle iteratively: we adjust the radius α at each call using a dichotomic search between 0 (the utopic value) and α_{up}, which serves as an initial upper bound.

5. Approximation Operator $\oplus_{lex}^k$

Despite the Σ_2^P-hardness of DP2, we introduce in this section an MO-COP operator approximating $\oplus_\Omega^k$ (in the sense of the associated radius) which can be computed much more efficiently, especially for a high number of objectives where the exact approach does not work anymore. This approximation operator, denoted $\oplus_{lex}^k$, intuitively "targets" k specific areas of the Pareto front without actually computing it explicitely. The notion of *normalized weighted vector* ("weight vector" for short) [14] 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. We take advantage of a direct adaptation of our representativity measure (cf. Definition 3) to weighted 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*

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

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 $\odot_m^k$ is defined as*

$$\odot_m^k = \gamma(\arg\min\{\Omega^{\mathcal{W}_m}(\mathcal{W}') \mid \mathcal{W}' \subset \mathcal{W}_m, |\mathcal{W}'| = k\}),$$

where γ is any choice function.

Note that the result of a weight vector filtering is completely characterized by k and m, i.e., it is independent from any MO-COP instance. As a consequence $\odot_m^k$ can be assumed to be computed as a preprocessing step.

Example 1 (continued)**.** *Consider again the bi-objective MO-COP $P_\star$, 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 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 set $\{(V_1, V_2) \mid \omega_1^i.V_1 = \omega_2^i.V_2\}$.*

The MO-COP operator $\oplus_{lex}^k$ we are going to introduce uses the set $\odot_m^k$. Given an m-vector V, $V_>$ denotes the vector composed of each element of V rearranged in a non-decreasing order. Additionally, given two m-vectors U, V, $U.V$ denotes the vector Z defined for each $i \in \{1, \ldots, m\}$ as $Z^i = U^i.V^i$. Given an m-weight vector ω, let $\preceq_\omega^{lex}$ be the total preordering over $Sols(P)$ defined as follows. First, the strict part of $\preceq_\omega^{lex}$ is defined for all solutions $S, S' \in Sols(P)$ as $S \prec_\omega^{lex} S'$ iff $(V(S).\omega)_>$ lexically precedes $(V(S').\omega)_>$, i.e., iff there is $i \in \{1, \ldots, m\}$ such that $(V(S).\omega)_>^i < (V(S').\omega)_>^i$ and if $i > 1$, then for every $j \in \{1, \ldots, i-1\}$, $(V(S).\omega)_>^j = (V(S').\omega)_>^j$. Then $\preceq_\omega^{lex}$ is defined for all solutions $S, S' \in Sols(P)$ as $S \preceq_\omega^{lex} S'$ iff $V(S.\omega)_> = V(S'.\omega)_>$ or $S \prec_\omega^{lex} S'$. For instance, if $R(S) = (2, 6, 3, 5)$, $R(S') = (6, 1, 5, 4)$ and $\omega = (0.1, 0.5, 0.2, 0.2)$, then $S' \prec_\omega^{lex} S$ since the vector $(0.1*2, 0.5*6, 0.2*3, 0.2*5)_> = (1/5, 3, 3/5, 1)_> = (3, 1, 3/5, 1/5)$ lexically precedes the vector $(0.1*6, 0.5*1, 0.2*5, 0.2*5)_> = (3/10, 1/2, 1, 4/5)_> = (1, 4/5, 1/2, 3/10)$.

To define our MO-COP operator $\oplus_{lex}^k$, we take advantage of the so-called ω-weighted egalitarian operator $\oplus_\omega$ recently introduced in [13]:

Definition 8 (Weighted egalitarian operators)**.** *Let ω be a weight vector. The* ω-weighted egalitarian operator *$\oplus_\omega$ is defined for every MO-COP P as*

$$\oplus_\omega(P) = \min(Sols(P), \preceq_\omega^{lex}).$$

These operators exhibits a number of interesting properties: (i) they return a subset of Pareto optimal solutions; (ii) they have a high decisiveness power: it has been shown experimentally that only 1 solution is returned in most cases [13]; (iii) any Pareto optimal solution can be reached by adjusting the weight vector (even for non-convex instances); (iv) graphically speaking, the resulting Pareto optimal solution is the "closest" to some utopia point which is located on the line specified by ω w.r.t. some specific class of weighted norms. We are ready to define our approximation operator $\oplus^k_{lex}$, which is characterized by $\odot^k_m$ and weighted egalitarian operators $\oplus_\omega$:

Definition 9 ($\oplus^k_{lex}$)**.** *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 $V(\oplus^{lex}_{\omega_1}(P_\star)) = \{(9,2)\}$ (cf. Figure 1). Intuitively, the vector $\omega_1.V(S) = (0.25 * 9, 0.75 * 2) = (2.25, 1.5)$ is the most "balanced" one among $\mathsf{PF}(P)$. Similarly, for $\omega_2 = (0.75, 0.25)$, we have $V(\oplus^{lex}_{\omega_2}(P_\star)) = \{(2,8)\}$. Therefore, we get that $V(\oplus^k_{lex}) = \{(9,2),(2,8)\}$.*

6. Empirical Results

We empirically evaluated the CPU time of computing both $\oplus^k_\Omega$ and $\oplus^k_{lex}$ on MO-COP instances with 15 binary variables for $k = 7$. We carried out all experiments on one core running at 2.3GHz with 4GB RAM. We considered MO-COP instances $\langle \mathcal{X}, \mathcal{D}, \mathcal{C} \rangle$ randomly generated as complete constraint graphs such that each constraint $C^{ij} \in \mathcal{C}$ associates a vector V from $\mathbb{N}^m$, with each V^k being a random value ranging over $\{0, \ldots, 100\}$.

Table 1 gives the results of computing both $\oplus^7_\Omega$ and $\oplus^7_{lex}$, varying the number of objectives m from 2 to 8. For each m, all values represent an average on 100 MO-COPs instances, and we give the CPU time and the radius of the associated set of 7 Pareto optimal solutions. We fixed a time out of 3,600 seconds for computing the Pareto front (i.e., the first phase of the exact procedure) and a time out of 50 seconds for each call of the NP-oracle used to find an optimal set $\mathcal{S}$ of 7 Pareto solutions. For example, MO-COP instances with $m = 4$ had (in average) a Pareto front of size 388; computing the Pareto front (first phase of the exact procedure for $\oplus^7_\Omega$) took 4.70 seconds, while the phase searching for an optimal set $\mathcal{S}$ of 7 Pareto solutions took 190 seconds, thus 194.70 seconds in total; the radius of the average $\mathcal{S}$ was 918; and it was proved that no set of 7 Pareto solutions exists with a radius of 896. In comparison, for the same set of instances computing $\mathcal{S} = \oplus^7_{lex}(P)$ took 3.30 seconds and the radius of $\mathcal{S}$ was 1835 in average. As to $\oplus_{lex}$, computing the set of weight vectors $\odot^7_m$ was made for each m by discretizing the set Ω_7 with a precision of two digits per vector component. Since $\odot^7_m$ is independent from a specific MO-COP, it was computed offline as an upstream step.

Table 1 shows the impact of the increasing number of objectives on the computational time of the exact procedure for $\oplus^7_\Omega(P)$: for MO-COPs with 15 variables the time out is reached from 4 objectives and becomes sub-optimal. From 7 objectives and above the exact procedure becomes infeasible as it requires to solve iteratively to solve an NP-hard problem on an exponential input with thousands of Pareto optimal solutions. In comparison, computing $\oplus^7_{lex}(P)$ is much faster and offers an acceptable alternative for problems with a high number of objectives.

Table 1: CPU time and value of radius obtained from MO-COPs with $n = 15$ and $p = 7$.

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	16	2.34 / 0.26	2.80	156 (157)	457
3	92	2.97 / 0.62	3.01	494 (493)	975
4	388	4.70 / 190	3.30	918 (896)	1835
5	1346	11.59 / 583	3.72	1401 (1162)	2705
6	3135	28.79 / 681	4.04	1803 (1396)	3208
7	5278	73.38 / time out	4.34	time out	4058
8	8044	139.22 / time out	4.65	time out	4064

7. Summary

We adapted a well-known concept in discrete location theory to multi-objective constraint optimization problems (MO-COPs). We characterized a well-behaved, restricted subset of Pareto optimal solutions from any MO-COP. The notion is of interest in a decision-making context where no a priori preferences among the different objectives may be available. We investigated the computational complexity of some associated decision problems and provided exact and approximation procedures to compute such sets.

References

[1] T. Eiter, E. Erdem, H. Erdogan, and M. Fink. Finding similar/diverse solutions in answer set programming. *TPLP*, 13(3):303–359, 2013.

[2] Jonathan Halpern and Oded Maimon. Algorithms for the m-center problems: A survey. *European Journal of Operational Research*, 10(1):90 – 99, 1982.

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

[4] T. Ilhan and M. C. Pinar. An efficient exact algorithm for the vertex p-center problem. Technical report, Bilkent University Technical Report, Department of Industrial Engineering 06533 Ankara Turkey, 2001.

[5] S. S. Iyengar and M. R. Lepper. When choice is demotivating: Can one desire too much of a good thing? *Journal of personality and social psychology*, 79(6):995–1006, 2000.

[6] D. Mailharro. A classification and constraint-based framework for configuration. *Artificial Intelligence for Engineering Design, Analysis and Manufacturing*, 12:383–397, 1998.

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

[8] E. Minieka. The m-center problem. *SIAM Review*, 12:138–139, 1970.

[9] C. M. Papadimitriou. *Computational complexity*. Addison-Wesley, Reading, Massachusetts, 1994.

[10] E. Rollon and J. Larrosa. Bucket elimination for multiobjective optimization problems. *Journal of Heuristics*, 12(4-5):307–328, 2006.

[11] E. Rollon and Javier Larrosa. Multi-objective Russian doll search. In *AAI'07*, pages 249–254, 2007.

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

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

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