A RankMOEA to Approximate the Pareto Front of a Dynamic Principal-Agent Model

Juan Arturo Herrera-Ortiz
jherrera@uxmcc2.iimas.unam.mx
Katya Rodriguez-Vazquez
katya@uxdea4.iimas.unam.mx
IIMAS-UNAM
Circuito Escolar, CU, 04510
Coyoacan, Mexico City, MEXICO

Itza T.Q. Curiel Cabral
itza.curiel@gmail.com
Sonia Di Giannatale Menegalli
sonia.digiannatale@cide.edu
CIDE
Mexico-Toluca 3655, Lomas de Santa Fe 01210
Cuajimalpa, Mexico City, MEXICO

ABSTRACT

In this paper, a new Multi-Objective Evolutionary Algorithm (MOEA) named RankMOEA is proposed. Innovative niching and ranking-mutation procedures which avoid the need of parameters definition are involved; such procedures outperform traditional diversity-preservation mechanisms under spread-hardness situations. RankMOEA performance is compared with those of other state of the art MOEAs: MOGA, NSGA-II and SPEA2, showing remarkable improvements. RankMOEA is also applied to approximate the Pareto Front of a Dynamic Principal-Agent model with Discrete Actions posed in a Multi-Objective Optimization framework allowing to consider more powerful assumptions than those used in the traditional single-objective optimization approach. Within this new framework a set of feasible contracts is described, while others similar studies only focus on one single contract. The results achieved with RankMOEA show better spread and minor error than those obtained by already mentioned MOEAs, allowing to perform better economic analysis in the contracts trade-off surface.

Categories and Subject Descriptors

I.2.8 [**Artificial Intelligence**]: Problem Solving, Control Methods, and Search

General Terms

Algorithms, Design, Performance

Keywords

Multi-objective optimization, Evolutionary Algorithms

1. INTRODUCTION

Innumerable situations in real world involve in a natural way problems with multiple objectives to be optimized; Multi-Objective Optimization (MO) studies such kind of problems. Since MO implies to optimize conflicting objectives subject to certain constraints, most of the time it is impossible to determine a unique solution. Thus, Multi-Objective Optimization Problems (MOPs) are characterized by a set of alternative solutions that must be considered as equivalents given the lack of information about relevance of one objective with regard to the others. In MO a space for decision variables and a space for their objective functions evaluation are considered. In real valued functions, those two spaces are related by a mapping $F : \Re^n \rightarrow \Re^k$. It is assumed that a solution to the MOP can be defined in terms of a decision vector $\overrightarrow{x} = [x_1, x_2, \ldots, x_n]^T$ in the decision space $\Re^n$. The set of imposed constraints on $F(\overrightarrow{x}) = [f_1(\overrightarrow{x}), f_2(\overrightarrow{x}), \ldots f_k(\overrightarrow{x})]$ define a feasible region $\Omega \subseteq \Re^n$ in the decision space along with its corresponding image $\Lambda \subseteq \Re^k$ on the objective space. By definition, there is a possibly infinite set of optimal solutions which are found at the frontier of Λ and are called the Pareto Optimal Frontier (PF^*), while their corresponding decision variable values in Ω are called the Pareto Optimal Set (PS). A solution $\overrightarrow{x}$ in PS is Pareto optimal (also called nondominated with respect to Ω), which means that there is no other solution $\overrightarrow{y} \in \Omega$ for which $F(\overrightarrow{y})$ dominates $F(\overrightarrow{x})$ (denoted by $F(\overrightarrow{y}) \prec F(\overrightarrow{x})$). $F(\overrightarrow{y})$ is said to dominate$F(\overrightarrow{x})$ if and only if $\overrightarrow{y}$ improves any objective to optimize with respect to $\overrightarrow{x}$ without inducing some simultaneous deterioration in at least another objective, e.g. assuming only minimization$F(\overrightarrow{y})$ is partially less than $F(\overrightarrow{x})$, i.e., $\forall f_i \in F : f_i(\overrightarrow{y}) \leq f_i(\overrightarrow{x}) \wedge \exists f_i \in F : f_i(\overrightarrow{y}) < f_i(\overrightarrow{x})$.

Evolutionary Algorithms (EAs) have shown to be a promising approach to deal with real MOPs [4]. They usually do not guarantee to identify optimal trade-offs, but to find good assessments, i.e., the set of solutions (Pareto Frontier approximation) whose objective vectors are not too far away from the optimal objective vectors. One of the aspects in EAs is its exploration mechanism (mutation), which should be carefully adjusted in order to achieve good performance. Several heuristics have been developed in an attempt to solve this problem in single-objective problems, although no one is successful for every case given that some are too general or require an extremely fine-grained level of detail. A new rank mutation concept [2] has shown to overcome the aforesaid heuristics weaknesses by applying mutation such that different individuals are characterized by different mutation rates, each individual being assigned a particular mutation rate ac-

GECCO'11, July 12–16, 2011, Dublin, Ireland.

cording to some rule. The rule will work where the mutation rate is related to fitness rank [1]. A new EA which extends this concept to the MO framework and involves and innovative niching, called RankMOEA, is designed and tested in this paper with some already well known MOEAs, showing a good performance over spread-harness situations. In addition, RankMOEA is used to numerically approximate the PF^* of a Dynamic Principal-Agent model with Discrete Actions, which is analyzed from a MO framework.

The remainder of this paper is organized as follows. In Section 2, a detailed description of RankMOEA is given. Section 3 presents some spread-hardness situations that are used to compare the performance of RankMOEA with some well-known MOEAs. The Dynamic Principal-Agent model with Discrete Actions is presented and explained in Section 4. In Section 5, experimental results of applying RankMOEA to the previous Principal-Agent model are reported and discussed. Finally, conclusions are drawn in Section 6.

2. A NEW MULTI-OBJECTIVE EVOLUTIONARY ALGORITHM

2.1 Some Issues about Evolutionary MO

When using heuristics in MO, attaining good convergence to PF^* and maintaining Pareto Front approximation (outcome of an approximation algorithm, PF^*_{known}) distribution as diverse as possible are two desirable goals [19]. EAs have shown to be a good heuristic in order to achieve such goals. The main reason is that population approach is well suited to find multiple solutions, besides niche-preservation methods can be exploited to find diverse solutions at the same time that implicit parallelism helps. Several MOEAs with good performance [4] have been proposed through recent years focusing in some of the following elements: fitness assignment, diversity-preservation mechanism and elitism.

- Fitness assignment has been conceived either as aggregation based, or as objective based, or as Pareto dominance based [4]. Aggregation based fitness assignment combines objectives to be optimized in one single linear or nonlinear parameterized function. Objective based fitness assignment chooses the most suitable sequence of the objective(s) to optimize during selection process, i.e., only one subset of objectives is optimized at the time [9, 15]. Finally, Pareto dominance based fitness assignment ranks solutions according to the Pareto dominance concept using distinctive rules [8, 21, 12, 6, 20] .

- Diversity-preservation mechanisms impulse divergence in tangential direction to the promising regions discovered by the MOEA, this through probability selection bias towards less conglomerated regions. The most common mechanisms to maintain diversity are: fitness sharing [10], mating restriction, reinitialization [8], clustering [21], grid mapping [12], crowding [6] and truncation [21]. Nevertheless, most of such mechanisms require parameters specification.

- Elitism is extended to the concept of an offline population $PF^*_{known}(t)$, which stores all nondominated solutions found up to epoch t, $PF^*_{known}(t)$ can be interactive or isolated, which means that can either cooperate in the evolutionary process by selecting new parents from it, or act only as a repository unit to store solutions.

A new MOEA which extends Pareto dominance-based fitness assignment strength to also promote diversity, avoids the need for parameter specification, outperforms some well-known diversity-preservation mechanisms and uses elitism with interactive population is proposed in the next subsection.

2.2 RankMOEA

Three premises were considered to design the new proposed algorithm called RankMOEA:

- Since Pareto dominance rules sort candidate solutions in a certain order according to their proximity to the frontier of Λ, some advantage can be taken from such arrangement by intensifying exploration in candidate solutions far from the frontier and reducing exploration in candidate solutions close to the frontier. This assuming that the first type of solutions does not have much information about PF^* structure and need more effort to achieve a good performance.

- The structure of the search is defined by Ω and not by Λ, thus, diversity preservation mechanisms should work better in Ω if they are compliant with Ω structure. Hence, exploitation of the information could be successful by mating nearby candidate solutions in Ω since such process is less disruptive.

- In most of the cases, after a certain threshold in the evolutionary process of MOEAs, the number of nondominated solutions grows rapidly, thus reduced mutation in solutions closer to the frontier of Λ that are less conglomerated in Ω should improve performance by controlling exploration and preserving the emphasized exploitation in such regions.

RankMOEA extends the rank mutation presented in [1, 2] to the MO framework, overcoming the mutation fine tuning drawbacks and promoting a controlled diversity according to Pareto dominance and the degree of conglomerate. RankMOEA is described in Algorithm 1. First, a set of m individuals are initialized as the early population $\overrightarrow{x} \in P(1)$ and evaluated in the set of k objectives to be optimized (lines 1 to 3). Then, Goldberg's ranking [10] is used to sort nondominated and dominated individuals (line 4):

$$rank_g\left(\overrightarrow{x}, P(t)\right) =$$

$$\begin{cases} 1 & \text{iff} \nexists \overrightarrow{y} \in P(t) \,|\, F\left(\overrightarrow{y}\right) \prec F\left(\overrightarrow{x}\right) \\ \max_{\overrightarrow{y} \in P(t) | F(\overrightarrow{y}) \prec F(\overrightarrow{x})} \left[rank_g\left(\overrightarrow{y}, P(t)\right)\right] + 1 & \text{otherwise} \end{cases} \tag{1}$$

By definition, nondominated individuals in $P(t)$ have ranking value of 1, thus individuals closer to such nondominated inidividuals in Λ have lower ranking values. Goldberg's ranking was preferred since it allows smoother ranking landscapes of Pareto domination. RankMOEA uses an interactive online file $PF^*_{known}(t)$ to store continuously its approximation to PF^* (line 5). During the evolution process only $m/2$ of the parents $P(t)$ are chosen (line 7) by the selection procedure.

Mates of the $m/2$ parents are chosen using a minimum spanning tree niching which works over the phenotypic space

Algorithm 1 RankMOEA

1 $t \leftarrow 1$
2 **random initialization** of each individual $\vec{x} \in P(t)$
3 $\forall \vec{x} \in P(t)$ evaluate $f_i(\vec{x}) \, \forall f_i \in F$
4 $ranking_{\vec{x} \in P(t)} \leftarrow rank_g(\vec{x}, P(t)) \, \forall \vec{x} \in P(t)$
5 $PF^*_{known}(t) \leftarrow nondominated(P(t))$
6 **do while** $t <$ stop criterion
7 $\quad P'(t) \leftarrow selection(P(t))$
8 $\quad P''(t) \leftarrow mst\_niching\left(P'(t), P(t)\right)$
9 $\quad P'''(t) \leftarrow rank\_mutation\left(P''(t)\right)$
10 $\quad \forall \vec{x} \in P'''(t)$ evaluate $f_i(\vec{x}) \, \forall f_i \in F$
11 $\quad ranking_{\vec{x} \in P'''(t)} \leftarrow rank_g\left(\vec{x}, P'''(t)\right) \forall \vec{x} \in P'''(t)$
12 $\quad PF^*_{known}(t+1) \leftarrow nondominated\left(\left\{PF^*_{known}(t) \cup P'''(t)\right\}\right)$
13 $\quad$ **if** $\left|PF^*_{known}(t+1)\right| > m$
14 $\quad\quad PF^*_{known}(t+1) \leftarrow truncation\left(PF^*_{known}(t+1), m\right)$
15 $\quad\quad P(t+1) \leftarrow PF^*_{known}(t+1)$
16 $\quad$ **else**
17 $\quad\quad$ **sort** $\left\{P'''(t) - PF^*_{known}(t+1)\right\}$ by $ranking_{\vec{x} \in \left\{P'''(t) - PF^*_{known}(t+1)\right\}}$ **ascendant**
18 $\quad\quad P(t+1) \leftarrow \left\{PF^*_{known}(t+1) \cup \left\{P'''(t)\right\}[1 : m - |PF^*_{known}(t+1)|]\right\}$
19 $\quad$ **end if**
20 $\quad t \leftarrow t + 1$
21 **end do**

(line 8). This mechanism builds a minimum spanning tree in Ω including all individuals in $P(t)$, distance in Ω is computed normalizing every phenotypic feature which allows to handle incommensurable variables. In this approach, niches are not isolated elements, moreover they are elements partially coupled by the tree structure (see left side of Figure 1). Since each individual in the minimum spanning tree can be connected with more than one individual, every $\vec{x} \in P(t)$ is weighted with

$$mst(\vec{x}) = \frac{1}{rank_g(\vec{y}, P(t), F) - \left[mst\_arity(\vec{x})\right]^{-1} + 1} \quad (2)$$

where $mst\_arity(\vec{x})$ counts the number of individuals connected to $\vec{x}$ in the minimum spanning tree. So individuals with lower Goldberg ranking value and lower arity (conglomerate measure) will accomplish a higher value of $mst(\vec{x})$, a hierarchical preference of ranking over arity is denoted. In order to select the mates of the $m/2$ parents, a stochastic selection process (e.g. stochastic universal selection) can be used with $mst(\vec{x})$ as the desirability of selection, including all the neighbors of the parent in the minimum spanning tree, then parents will be mated with less conglomerated individuals that are closer to PF^*. It is important to observe that there is no need to define a proximity value. This procedure can be performed using Chazelle's algorithm [3] which is based on the soft heap, the most asymptotically efficient known structure to find the minimum spanning tree. Its running time is $O(\lambda\alpha(\lambda, m))$, where λ is the number of edges and α is the classical functional inverse of the Ackermann function. The function α grows extremely slowly, so that for all practical purposes it may be considered a constant no greater than 4; thus Chazelle's algorithm takes very close to linear time.

The proposed rank mutation considers pre-order, the intrinsic inconvenient of MOPs. Rank mutation (line 9) consists of the definition of a mutation rate range and the assign-

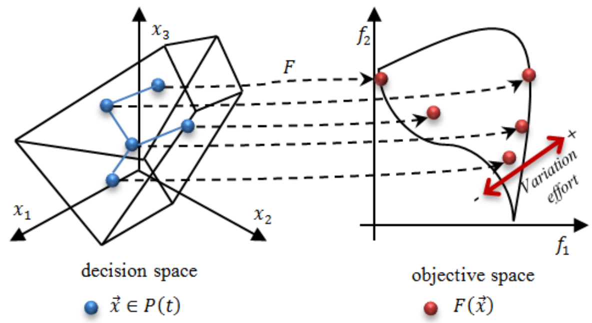

Figure 1: Minimum spanning tree niching and ranking mutation.

ment of a uniformly distributed mutation rate to individuals according to their inherited $mst(\vec{x})$ value, i.e. individuals with lower Goldberg ranking value and lower arity will get a lower mutation rate and individuals with higher Goldberg ranking value and higher arity will get a higher mutation rate, denoting tight exploration in the neighborhood of individuals closer to PF^* and widespread exploration in the neighborhood of individuals farther from PF^* (see right side of Figure 1). When the entire population is nondominated, tight exploration is performed in the neighborhood of individuals with lower arity and widespread exploration in the neighborhood of individuals with higher arity. The mutation rate range will be specified by a minimum and maximum mutation rates, p_{min} and p_{max} respectively, and divided into m steps to generate the deterministic rule of choosing the mutation rate. So the mutation rate of the i-esim individual in $P''(t)$ sorted in descendant order according to $mst(\vec{x})$ is $p_{min} + i \cdot (p_{max} - p_{min}) / (m-1)$. According to [2] a natural range to cover any eventuality is $p_{min} = 0$ and $p_{max} = 1 - 1/l$, where l is individual length, however if there is knowledge of the vicinity of the optimum, and the population is in the vicinity, then a lower p_{max} may be appropriate. Mutation range remains fixed during entire evolution. Since mutation only requires to sort individuals in order to assign the mutation rate, this step has a complexity of $O(m \log m)$.

Thereafter, RankMOEA evaluates the offspring in the set of k objectives and ranks them with Goldberg's ranking (lines 10 and 11). $PF^*_{known}(t)$ is updated with new nondominated solutions (line 12). Finally, if $PF^*_{known}(t)$ size is larger than m, a truncation process proposed in [21] is used to reduce its size (lines 14 and 15) and the $P(t+1)$ is constituted by such reduction; else $P(t+1)$ is constituted by $PF^*_{known}(t)$ and a controlled insertion using ranking based selection of the best offspring that were not already included in $PF^*_{known}(t)$ (lines 17 and 18).

The speed performance of RankMOEA is ruled by the mutation process, therefore its computational complexity can be calculated as $O(m \log m)$, which makes it a fast algorithm, worthy to compete with other state of the art MOEAs.

3. TESTING RANKMOEA

In the following tests two well-known MOEAs (NSGA-II [6] and SPEA2 [20]) were used to compare the performance of RankMOEA. The three MOEAs used binary-coded chromosomes, one point crossover and bit-wise mutation. RankMOEA was tested using the stochastic universal sampling, while NSGA-II and SPEA2 were tested using their tournament selection operator. NSGA-II's crowding, SPEA2's k-nearest neighbor and RankMOEA's minimum spanning tree niching were implemented in the phenotypic space. The mating rates used for NSGA-II, SPEA2 and RankMOEA were: 70%, 80% and 90%. The mutation rates used for NSGA-II and SPEA2 were 1%, 2%, 3%, 4%, 5% and 6%, whereas for RankMOEA p_{min} was set to 0% and p_{max} to 6%. A precision of 0.001 was set for each variable in the phenotype. The three algorithms were run 30 times with each mating-mutation configuration, the average behavior of each configuration was assessed using a version of $G-metric$ [14] to work in Ω, a n-ary quality indicator that ranks PF^*_{known}s based on the their attained dispersion and convergence.

3.1 Spread-hardness Test

The first test was designed to examine the robustness of the diversity-preservation mechanisms of NSGA-II, SPEA2 and RankMOEA by finding a good diversity of the solutions in Ω. A function with three reference points $\overrightarrow{z_i} : i = 1, 2, 3$ in a bidimensional Ω was defined; the idea is to minimize the distance to such reference points, i.e., $\min F(\overrightarrow{x}) = [f_1(\overrightarrow{x}), f_2(\overrightarrow{x}), f_3(\overrightarrow{x})]$ with $f_i(\overrightarrow{x}) = \| \overrightarrow{z_i} - \overrightarrow{x} \|_2$, where $\overrightarrow{z_1} = (2, 1)$, $\overrightarrow{z_2} = (3, 5)$, $\overrightarrow{z_3} = (4, 1)$ and subject to $x_j \in [0, 6] : j = 1, 2$. It is clear that PS is formed by all the points located within the triangle constituted by the three reference points. It is expected that a diversity-preservation mechanism with good performance and compliant with Ω structure will achieve a quasi-uniform spread. A population and $PF^*_{known}(t)$ size of 50 individuals with 10,000 objective function evaluations were considered. Figure 2 shows the best approximation to PS achieved by the best run of the best mating-mutation configuration of each MOEA according to the average of $G-metric$ over the 30 runs. The best distribution is achieved by RankMOEA, followed by SPEA2 and finally by NSGA-II.

3.2 Complicated PS Test

The second test was performed using UF4 problem from CEC'09 contest [18], a MOP with complicated PS which demonstrated to be a very hard problem even for the best algorithms that participated in MO contest in CEC'09. Figure 3 shows the best PF^*_{known} achieved by the best run of the best mating-mutation configuration of each MOEA according to the average of $G-metric$ over the 30 runs. A population and $PF^*_{known}(t)$ size of 100 individuals with 300,000 objective function evaluations were considered. RankMOEA shows the best spread and the lowest convergence error, followed by NSGA-II with worse spread, and finally SPEA2 with the highest convergence error.

4. DYNAMIC PRINCIPAL-AGENT MODEL WITH DISCRETE ACTIONS

The Principal-Agent problem is a political science and economics well known problem, which analyzes a situation of asymmetric information where a risk neutral Principal delegates tasks to a risk averse Agent. Asymmetric information arises because the Principal cannot observe the effort level that the Agent chooses, and monitoring the Agent is too costly for the Principal. Moreover, the existence of uncertainty in the production process makes the design of the Agent's compensation plan a non-trivial problem.

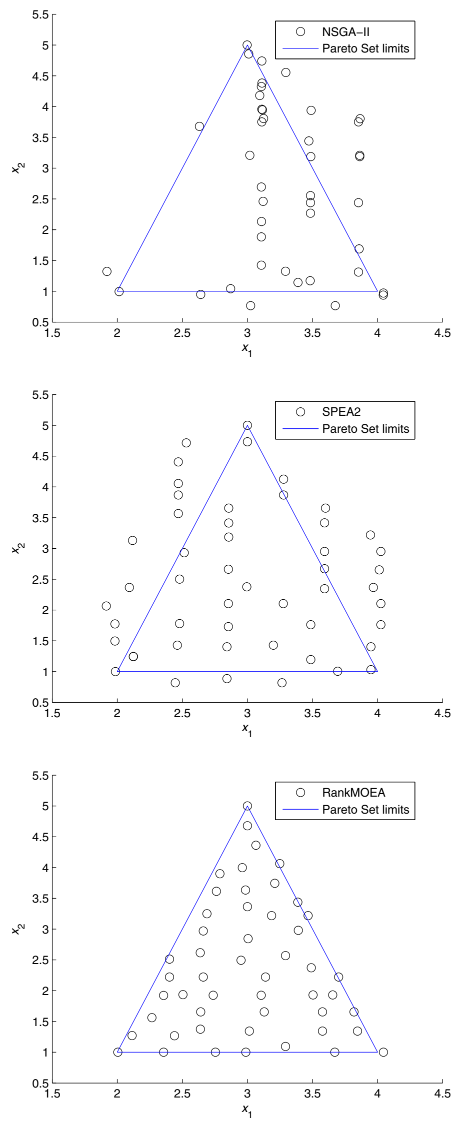

Figure 2: Best approximation of PS achived by NSGA-II, SPEA2 and RankMOEA, minimizing the distance to reference points in a bidimensional Ω.

A dynamic problem can be modeled when the Principal-Agent relationship is recurrent, i.e., the relationship goes on for infinite periods [16, 17]. In such context, the Agent's compensation plan has two components: present and future compensation. Both components of the Agent's compensation plan aim to link the Agent's wealth with the Principal's wealth. In the Dynamic Principal-Agent problem, the Principal maximizes his discounted expected utility subject to two fundamental constraints: the participation constraint (i.e. the contractual relationship should be accepted by the Agent), and the incentive compatibility constraint (i.e. the level of effort implemented by the Principal in every period should be chosen by the Agent given the unobservability of his effort).

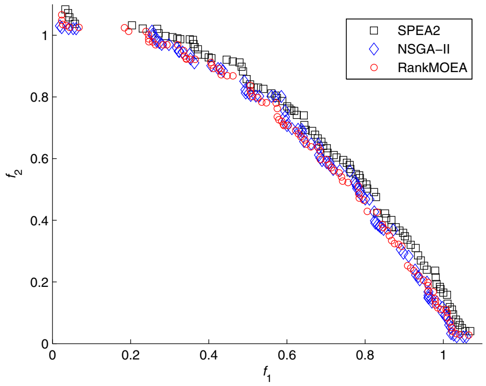

Figure 3: Best PF^*_{known} achived by NSGA-II, SPEA2 and RankMOEA in UF4-CEC'09.

The contractual arrangement between the Principal and the Agent affects how the economic surplus is divided and the sheer magnitude of such surplus. Hence, characterizing a Pareto Front where the Principal and the Agent have different levels of bargaining power is an interesting exercise to shed light into how the economic surplus is affected by those different contractual arrangements.

4.1 A Multi-Objective Framework

Given that in MO, each Pareto optimal solution represents a different compromise among objectives, finding different Pareto optimal solutions implies finding the structure of the trade-off surface involved in the problem. Thus, since the Principal-Agent model represents a situation of conflict of interests, modeling it as a MOP allows:

- to consider diverse contractual arrangements between the Principal and the Agent in which they have dissimilar levels of bargaining power,
- to achieve a better insight on how the creation of economic surplus is affected by the diverse contractual arrangements, and
- to obtain a better idea of how the conflict of interest and asymmetry of information between the Principal and the Agent affect the creation of economy surplus.

4.2 Mathematical Model

In order to reconceive the Dynamic Principal-Agent model in a MO framework, two objectives are considered: maximize the Principal's discounted expected utility U, and maximize the Agent's discounted expected utility V. The Dynamic Principal-Agent model is about choosing an action plan, a compensation plan for each level of production and a future utility plan such that U and V are simultaneously maximized. The Dynamic Principal-Agent model with Discrete Actions can be stated as:

$$\max_{a(V),w(y,V),\widetilde{V}(y,V)} \{U, V\} \tag{3}$$

$$U_t = \sum_{y\in Y} f\left(y; a\left(V\right)\right) u_t\left(y, w\left(y, \overline{V}\right)\right) + \beta U_{t+1}\left(\widetilde{V}\left(y, \overline{V}\right)\right) \tag{4}$$

$$V_t = \sum_{y\in Y} f\left(y; a\left(V\right)\right) v_t\left(a\left(V\right), w\left(y, \overline{V}\right)\right) + \beta \widetilde{V}_{t+1}\left(y, \overline{V}\right) \tag{5}$$

where $f\left(y; a\left(V\right)\right)$ is the probability function that associates action $a(V)$ and output y, $u_t\left(y, w\left(y, \overline{V}\right)\right)$ is the Principal's expected utility at time t, $v_t\left(a\left(V\right), w\left(y, \overline{V}\right)\right)$ is the Agent's expected utility at time t, $w\left(y, \overline{V}\right)$ is the current compensation, $\overline{V}$ is the Agent's reservation utility, $\widetilde{V}\left(y, \overline{V}\right)$ is the state variables for tomorrow on, and β is the discount factor. This Dynamic Principal-Agent model with Discrete Actions can be represented as a Bellman equation given [16] methodology. This model is subject to the participation constraint,

$$\sum_{y\in Y} f\left(y; a\left(V\right)\right) v\left(a\left(V\right), w\left(y, \overline{V}\right)\right) + \beta \widetilde{V}\left(y, \overline{V}\right) = \overline{V} \tag{6}$$

to the fact that the actions are in the space of feasible actions,

$$a\left(\overline{V}\right) \in A \tag{7}$$

where A is the action set, to the inability of temporal borrow,

$$0 \leq w\left(y, \overline{V}\right) \leq y \; \forall y \in Y \tag{8}$$

where Y is the output set, and to the fact that the future compensation is in the feasible space.

$$\overline{V}\left(y, V\right) \in V \; \forall y \in Y \tag{9}$$

5. APPROXIMATING THE PRINCIPAL-AGENT MODEL SOLUTION

Given the difficulty of obtaining analytical results with Dynamic Principal-Agent models, close solution methods are not applicable, hence it is common in the literature to numerically approximate the optimal contracts, see e.g. Wang [17] and Fernandes & Phelan [7].

5.1 Finding the Optimal Contract: a Numerical Example

The same functional forms and parameters of Wang [17] are used. In particular, the Agent's utility function is assumed to be exponential because the Agent is risk averse, i.e. $v_t\left(a\left(V\right), w\left(y, \overline{V}\right)\right) = e^{\gamma\left(a(V) - aw\left(y, \overline{V}\right)\right)}$ where $\gamma > 0$ is the coefficient of absolute risk aversion and $\alpha > 0$ measures the relative cost for the Agent of exercising an unit of effort. On the other hand, the Principal's utility is $u_t\left(y, w\left(y, \overline{V}\right)\right) = y - w\left(y, \overline{V}\right)$ because risk neutral is assumed.

For the standard model $\gamma = \alpha = 1$ and two feasible action levels $A = a_L = 0.1, a_H = 0.2$ are assumed, i.e., the Agent can choose either to shirk or to work. Hence, $a_H > a_L$ indicates that shirking is less costly than working. Also, it is assumed that there are two levels of output: low or high $Y = y_L = 0.4, y_H = 0.8$, and the probability function that associates effort and output is defined as: $f\left(y_L; a_L\right) = f\left(y_H; a_H\right) = 2/3$ and $f\left(y_L; a_H\right) = f\left(y_H; a_L\right) = 1/3$. These probabilities capture the idea that the more diligently the Agent works, the greater the likelihood of the realization of the high output level. Finally, the Principal and the Agent's common discount factor was set to $\beta = 0.96$.

The numerical solution of the Bellman equation is

$$\{U, V, \widehat{a}(V), w(y_H, \overline{V}), w(y_L, \overline{V}), \widetilde{V}(y_H, \overline{V}), \widetilde{V}(y_L, \overline{V})\}$$

where $\widehat{a}(V)$ is the optimal action. Given a finite horizon, the chromosome of the individuals in the population is characterized by 3 substrings of length N, where N is the number of periods of time an individual lives, i.e. the length of each chromosome is $3N$. The first substring indicates the history of actions of the individual, the second and third one show the history of compensations conditional on a high or low output level respectively. Therefore, the phenotype of an individual is defined as:

$$\begin{bmatrix} a_1(V), a_2(V), \dots a_N(V), \\ w_1(y_H, \overline{V}), w_2(y_H, \overline{V}), \dots w_N(y_H, \overline{V}), \\ w_1(y_L, \overline{V}), w_2(y_L, \overline{V}), \dots w_N(y_L, \overline{V}) \end{bmatrix}$$

In order to compute U and V a backward induction must be used [11]. The number of periods in the Agent's life-span was set $N = 70$.

5.2 Experimental Results

In this test, besides NSGA-II, SPEA2 and RankMOEA, a well-known MOEA called MOGA [8] was included as an inferior bound. The same algorithmic specifications described in Section 3 were used and extended to MOGA, which was tested using the stochastic universal sampling. A precision of 0.0001 was required for each variable in the phenotype, thus binary chromosomes of 1820 bits were used, a population and $PF^*_{known}(t)$ size of 200 individuals with 100,000 objective function evaluations were considered. The four algorithms were run 50 times with different mating and mutation rates combination, the average behavior of each configuration was computed using a set of MO quality indicators: I_ε [22, 13], Distribution of the Found Pareto Front ($DFPF$) [19] and $G - metric$, this due to collectively, they fulfill the following evaluation issues in MO:

- Convergence to PF^*, how close PF^*_{known} is from PF^*. This can be difficult for highly discontinuous landscapes where a search method can be trapped in local minima. If in addition, the considered objectives have complex interactions; this can yield difficult optimization problems.

- Sample representatively PF^*. This entails a diversification of the set of solutions along the entire PF^* and can be defined as: Uniformity, how appropriate PF^*_{known} distribution is, meaning the relative distance among solutions; most of the time a homogeneous dispersion is ideal; and as Spread, how appropriate PF^*_{known} extension is; a wider PF^*_{known} involves more options.

The mating rates used for MOGA, NSGA-II, SPEA2 and RankMOEA were: 40%, 50%, 60%, 70%, 80% and 90%. The mutation rates used for MOGA, NSGA-II and SPEA2 were 1%, 2%, 3%, 4%, 5%, 6%, 7% and 8%, whereas for RankMOEA p_{min} was set to 0% and p_{max} to 8%. For the four algorithms, constraints were handled with the idea of superiority of feasible points proposed in [5].

In Evolutionary MO, how to evaluate PF^*_{known} quality that different MOEAs generate is still an open problem. Detailed analysis in each of the used quality indicators is possible; however discrimination among their measurements constitutes itself a MOP. So we propose to combine the three MO quality indicators in order to discriminate among PF^*_{known}s quality. Since linear combination of the quality indicators could smooth differences and hide trade-offs when there is no an absolute winner-MOEA in all indicators, we propose a lexicographic combination of their preference according to expectations of the user, hence providing a suitable framework of analysis from the point of view of the user.

First, for each $MOEA_i \in \Gamma$ to be compared, we used its 50 runs to calculate the mean and variance of every indicator, the reference set used to execute this task was formed as the total Pareto Front taking in count all the runs from all MOEAs. Then we ranked MO quality indicators using the statistical measure ε' presented in [2]:

$$\varepsilon'_{i,j} = 10\left(\frac{\mu_{I,MOEA_i} - \mu_{I,MOEA_j}}{\sqrt{\sigma^2_{I,MOEA_i} + \sigma^2_{I,MOEA_j}}}\right) \quad (10)$$

where the numerator is the difference in average of indicator I between $MOEA_i$ and $MOEA_j$, and the denominator is the variance difference. If we assume a normal distribution then, $\varepsilon'_{i,j} > 2$ corresponds to a 95% confidence interval which will take to mean that it is statistically significant than the $MOEA_j$ is leading to better I values than the $MOEA_i$. Next, we ranked every indicator I of every MOEA using:

$$rank_I(MOEA_i) = \begin{cases} 0 & \text{if} \nexists MOEA_j \in \Gamma | \varepsilon'_{j,i} > 2 \\ \max_{MOEA_j \in \Gamma | \varepsilon'_{j,i} > 2} [rank_I(MOEA_j)] + 1 & \text{otherwise} \end{cases} \quad (11)$$

With this ranking, a landscape of dominance of every quality indicator is achieved, differentiating which MOEA outperforms others MOEAs with regard to a specific quality indicator with 95% of statistic confidence. Finally, we lexicographically combined the ranking of the three quality indicators using the following order: 1) I_ε, 2) $DFPF$ and 3) $G - metric$. It is easy to see that the MOEA with the best performance according to the defined order of the MO quality indicators will have the highest ranking value. The performance of the four MOEAs in the Dynamic Principal-Agent model with Discrete Actions using the comparison process previously described is plotted in Figure 4.

MOEA's configuration with performance ranking values lower than 2 did not achieved feasible solutions by violating some constraints. MOGA shows a bad performance, since only four configurations achieved feasible solutions but with very low ranking value, i.e. poor convergence and spread, besides mutation percentage seems to affect the performance of MOGA in an erratic way.

Mutation percentage seems to have an important role in the performance of NSGA-II and SPEA2, since lower mutations rates allow to achieve a better performance. In NSGA-II higher values of mating rate seem to offer a better PF^*_{known}, while in SPEA2, medium values of mating rate subjugated to a low mutation rate is clearly a key to achieve better PF^*_{known}s. Both algorithms have analogous average behavior over all the combinations of mutation and mating rates. About performance of RankMOEA, lower values of mating rates allow to achieve a better performance, even better than those obtained by NSGA-II and SPEA2. A remarkable fact of RankMOEA's behavior contrary to the other three MOEAs, is that it always achieves feasible solutions. Even worst approximations of RankMOEA are comparable to best approximations of NSGA-II and SPEA2.

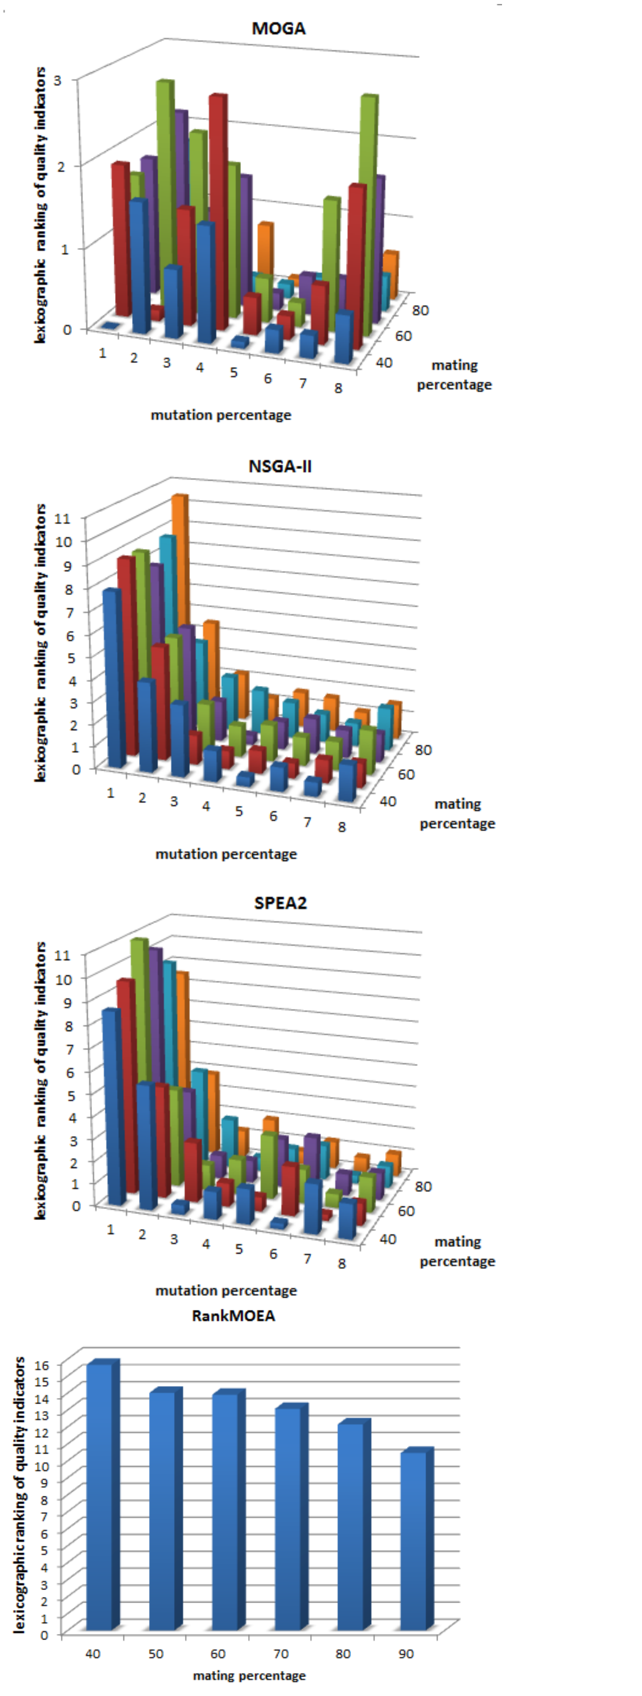

Figure 4: MOGA, NSGA-II, SPEA and RankMOEA's performance in Principal-Agent model.

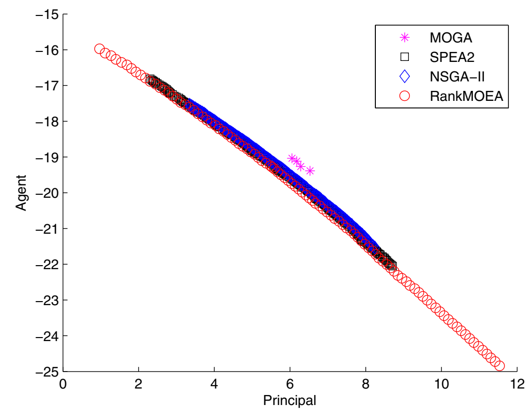

Figure 5: Best PF^*_{known} achieved by the best configuration of every MOEA tested in the Dynamic Principal-Agent model with Discrete Actions.

In order to have a better idea of MOEAs performance, the best PF^*_{known} achieved by every MOEA with its best mutation-mating configuration is plotted in Figure 5. Proposed comparison process success is confirmed by the correct classification of the quality of the achieved MOEA's outcomes. RankMOEA clearly enhances the convergence and spread achieved by MOGA, NSGA-II and SPEA2.

5.3 Analysis of the Achieved Approximation of the Pareto Front

As a result, a concave PF^* is numerically approximated, which is consequence of the information asymmetry between the Principal and the Agent (Figure 5). As contracts vary in the trade-off surface towards those that are more advantageous to the Agent, it is observed the prevalence of compensation plans in which the Principal assumes most of the risk of the productive activity. When the Principal and the Agent are more patient, both obtain higher values of their discounted expected utilities, which generates a higher level of economic surplus. The Agent faces lower variability in future compensation when it is costlier for him to exert an additional effort unit. In Figure 6 the current compensation schedules of the most advantageous contract for the Principal (PC) and the most advantageous contract for the Agent (AC) can be observed over the periods of time. Low and high salaries of AC are higher than those of PC, moreover, in most of the cases the low level of the salary for AC is higher than the high level of the salary for PC. Note that the low salary schedules of these two contracts do not vary, i.e., both PC and AC provide incentives to the Agent through variability in the high levels of salary.

6. CONCLUSIONS

A RankMOEA is designated using minimum spanning tree niching and ranking-mutation procedure, the computational complexity of the new algorithm is $O(m \log m)$. The new diversity-preservation mechanism involved does not need extra parameters to work. RankMOEA outperforms traditional diversity-preservation mechanisms under spread-hardness situations, showing good spread and lower convergence error compared with other MOEAs. An alternative

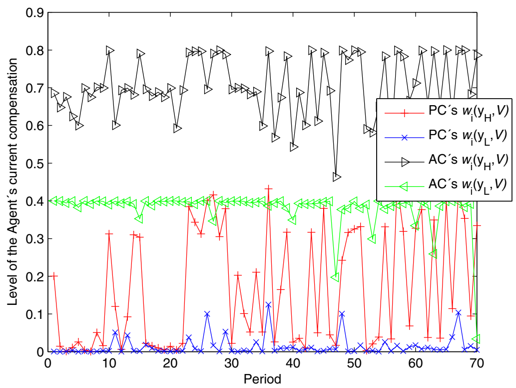

Figure 6: Agent's current compensation for PC and AC.

comparison process is proposed, where MOEAs performance is discriminated by using MO quality indicators statistical information and its lexicographical ranking combination.

In addition, RankMOEA is applied to approximate the Pareto Front of the Dynamic Principal-Agent model with Discrete Actions. The results achieved with RankMOEA show better spread and minor error than those obtained by some well-known MOEAs, allowing to perform better economic analysis by characterizing contracts in the trade-off surface. The achieved approximation of the Pareto Front allows to observe different compensation plans at different levels of bargaining power of the Agent and the Principal, and how the different contractual arrangements affect the generation of economic surplus.

7. REFERENCES

[1] J. Cervantes and C. Stephens. A rank proportional generic genetic algorithm. In *Proc. of the Conf. on Genetic and Evolutionary Computation (GECCO '08)*, volume 7, pages 71–74, 2008.

[2] J. Cervantes and C. Stephens. Limitations of existing mutation rate heuristics and how a rank ga overcomes them. *IEEE Trans. on Evolutionary Computation*, 13(2):369–397, 2009.

[3] B. Chazelle. A minimum spanning tree algorithm with inverse-ackermann type complexity. *J. ACM*, 47(6):1028–1047, 2000.

[4] C. Coello, G. B. Lamont, and D. Van Veldhuizen. *Evolutionary Algorithms for Solving Multi-Objective Problems*. Kluwer Academic Publishers, 2009.

[5] K. Deb. An efficient constraint handling method for genetic algorithms. *Comput. Methods Appl. Mech. Eng.*, 186:311–338, 2000.

[6] K. Deb, S. Agrawal, A. Pratab, and T. Meyarivan. A fast elitist non-dominated sorting genetic algorithm for multi-objective optimization: Nsga-ii. Technical Report KanGAL-200001, Indian Institute of Technology, Kanpur, India, 2000.

[7] A. Fernandes and C. Phelan. A recursive formulation for repeated agency with history dependence. *J. Economic Theory*, 91:223–247, 2000.

[8] C. Fonseca and P. Fleming. Genetic algorithms for multiobjective optimization: Formulation, discusión and generalization. In S. Forrest, editor, *Proc. 5th Int'l Conf. on Genetic Algorithms*, pages 416–423, San Mateo, California, 1993.

[9] M. Fourman. Compaction of symbolic layout using genetic algorithms. In *Genetic Algorithms and their Applications: Proc. the 1st Int'l Conf. on Genetic Algorithms*, pages 141–153. Lawrence Erlbaum, 1985.

[10] D. Goldberg. *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley Publishing Company, 1989.

[11] K. Judd. *Numerical methods in economics*. The MIT Press, 1998.

[12] J. Knowles and D. Corne. The pareto archived evolution strategy. a new baseline algorithm for multiobjective optimisation. In *1999 Congress on Evolutionary Computation*, pages 95–105, Washington, D.C., July 1999.

[13] J. Knowles, L. Thiele, and E. Zitzler. A tutorial on the performance assessment of stochastic multiobjective optimizers. Technical Report 214, Computer Engineering and Comunication Networks Lab (TIK), Swiss Federal Institute of Technology (ETH), Zurich, Switzerland, 2006.

[14] G. Lizárraga, A. Hernández, and S. Botello. G-metric: an m-ary quality indicator for the evaluation of non-dominated sets. In *Proc. the Genetic and Evolutionary Computation Conf. (GECCO 2007)*, volume 1, pages 665–672, 2007.

[15] D. Schaffer. Multiple objective optimization with vector-evaluated genetic algorithms. In *Int'l 1st Conf. on Genetic Algorithms and their Applications*, pages 93–100, 1985.

[16] S. Spear and S. Srivastava. On repeated moral hazard with discounting. *Review of Economic Studies*, 54:599–617, 1987.

[17] C. Wang. Incentives, ceo compensation, and shareholder wealth in a dynamic agency model. *J. Economic Theory*, 76:72–105, 1997.

[18] Q. Zhang, A. Zhou, S. Zhao, P. Suganthan, and W. Liu. Multiobjective optimization test instances for the cec 2009 special session and competition. Technical Report CES-887, University of Essex and Nanyang Technological University, 2008.

[19] E. Zitzler, K. Deb, and L. Thiele. Comparison of multiobjective evolutionary algorithms: Empirical results. *IEEE Trans. on Evolutionary Computation*, 8(2):173–195, June 2000.

[20] E. Zitzler, M. Laumanns, and L. Thiele. Spea2: Improving the strength pareto evolutionary algorithm. Technical Report 103, Computer Engineering and Networks Laboratory (TIK), Swiss Federal Institute of Technology (ETH) Zurich, May 2001.

[21] E. Zitzler and L. Thiele. An evolutionary algorithm for multiobjective optimization: The strength pareto approach. Technical report, Computer Engineering and Comunication Networks Lab (TIK), Swiss Federal Institute of Technology (ETH), Zurich, Switzerland, May 1998.

[22] E. Zitzler, L. Thiele, M. Laumanns, C. Fonseca, and V. Grunert. Performance assessment of multiobjective optimizers: An analysis and review. *IEEE Trans. on Evolutionary Computation*, 7(2):117–132, 2003.