Multiobjective Evolutionary Algorithms for Portfolio Management: A comprehensive literature review

K. Metaxiotis [1], K. Liagkouras *

Decision Support Systems Laboratory, Department of Informatics, University of Piraeus, 80, Karaoli & Dimitriou Str., 18534 Piraeus, Greece

ABSTRACT

In this paper we provide a review of the current state of research on Portfolio Management with the support of Multiobjective Evolutionary Algorithms (MOEAs). Second we present a methodological framework for conducting a comprehensive literature review on the Multiobjective Evolutionary Algorithms (MOEAs) for the Portfolio Management. Third, we use this framework to gain an understanding of the current state of the MOEAs for the Portfolio Management research field and fourth, based on the literature review, we identify areas of concern with regard to MOEAs for the Portfolio Management research field.

© 2012 Elsevier Ltd. All rights reserved.

1. Introduction

Multiobjective optimization (MO) is the problem of simultaneously optimizing two or more conflicting objectives subject to certain constraints. Many real-world problems involve simultaneous optimization of several often conflicting objectives. The portfolio optimization problem is a characteristic example of this category of problems. According to Markowitz's (1952) mean–variance theory (MV) an investor attempts to maximize portfolio expected return for a given amount of portfolio risk or minimize portfolio risk for a given level of expected return. The MV theory was criticized for unrealistic assumptions. Responding to the critics, researchers incorporated to the portfolio model some real world constraints like bounds on holdings, cardinality, minimum transaction lots and sector capitalization constraints.

Moreover they proposed new risk measures such as the value-at-risk (VaR), the lower partial moments (LPM) of nth order, the conditional value-at-risk (CVaR) (kibzun & Kuznetsov, 2006), and the conditional drawdown-at-risk (CDaR) (Balzer, 1994). However, these additional constraints and risk measures made the Portfolio Selection problem difficult to be solved with exact methods.

Evolutionary algorithms (EAs) have become the method of choice for optimization problems that are too complex to be solved using deterministic techniques. EAs are well suited to multiobjective optimization problems (MOP) as they are inspired by the biological processes which are inherently multiobjective.

Thanks to Multiobjective Evolutionary Algorithms (MOEAs) techniques the classical portfolio model can be extended to handle two or more conflicting objectives subject to various realistic constraints.

Because the various objectives functions in the portfolio selection problem are usually in conflict with each other, each time that we attempt to optimize further an objective other objectives suffer as a result. Therefore, the objective in MOEAs is to find the Pareto front of efficient solutions that provide a tradeoff between the various objectives.

During the last decade MOEAs for the Portfolio Management have attracted attention from both academics and practitioners and we feel that now is a good time to study how the MOEAs for the Portfolio Management field has evolved and what its present state is. The purpose of this paper is fourfold. The first objective is to present the current state of research in Portfolio Management with the support of MOEAs by providing a brief review of the available literature in the field. The second objective is to develop a methodological framework for conducting a comprehensive literature study based on the papers published in MOEAs for the Portfolio Management over a long time span across various disciplines. The third objective is to use this framework to gain an understanding of the current state of the MOEAs for the Portfolio Management research field. The fourth objective is, based on the literature review, to identify potential areas of concern in regard to MOEAs for the Portfolio Management.

The paper is organized as follows. In Section 2 we present the most well known techniques for multiobjective optimization with the use of MOEAs. In Section 3 we provide a brief review of the available literature in Portfolio Management with the support of MOEAs. In Section 4 the methodological framework for carrying out the literature study is presented and additionally the findings of the review are analyzed. Finally, in Section 5 we identify some possible paths for future research in MOEAs for the Portfolio Management.

* Corresponding author. Tel.: +30 6948612347.
E-mail addresses: kmetax@unipi.gr (K. Metaxiotis), kliagk@unipi.gr (K. Liagkouras).
[1] Tel.: +30 2104142378.

2. Multi-objective evolutionary algorithms (MOEAs): a brief review

The idea of using techniques based on the emulation of the mechanism of natural selection to solve problems can be traced as long back as the 1930s (Cannon, 1932). Since 1930s, this research field did not see major developments for almost three decades. However, in 1960s three studies (Fogel, 1966; Holland, 1962; Schwefel, 1965) set the foundations of what is nowadays denominated "evolutionary algorithms" (EAs).

EAs are population based stochastic optimization heuristics inspired by Darwin's Evolution Theory. An EA searches through a solution space in parallel by evaluating a set (population) of possible solutions (individuals). An EA starts with a random initial population. Then the 'fitness' of each individual is determined by evaluating the objective function. After the best individuals are selected, new individuals for the next generation are created. The new individuals are generated by altering the individuals through random mutation and by mixing the decision variables of multiple parents through crossover. Then the generational cycle repeats until a breaking criterion is fulfilled. EAs have been applied successfully to a wide range of problems such as engineering, biology, genetics, finance, etc. According to the literature EAs have been proved very effective for single-optimization (Fogel, 1999; Goldberg, 1989; Schwefel, 1981).

The last years there has been an increasing interest to explore the application of evolutionary algorithms for the solution of multiobjective optimization problems (MOPs).

Multiobjective optimization (MO) is the problem of maximizing/minimizing a set of conflicting objectives functions subject to a set of constraints. In Multiobjective optimization there is not a single solution that maximizes/minimizes each objective to its fullest (Gandibleux & Ehrgott, 2005). This happens because the various objectives functions in the problem are usually in conflict with each other. Therefore, the objective in MO is to find the Pareto front of efficient solutions that provide a tradeoff between the various objectives (Zitzler et al., 2000, 2002). EAs are naturally suited for extension into the multiobjective optimization problems domain, thanks to the population based search strategy and the simple selection strategy. The ability of EAs to deal simultaneously with a set of possible solutions (population) allows to find several members of the Pareto optimal set in a single run of the algorithm, instead of having to perform a series of separate runs as in the case of the traditional mathematical programming techniques. Moreover, EAs are less susceptible to the shape or continuity of the Pareto front (they can easily deal with discontinuous and concave Pareto fronts), whereas these two issues are known problems with mathematical programming techniques (Coello Coello, Van Veldhuizen, & Lamont, 2002; Coello Coello & Toscano Pulido, 2001). During multiobjective optimization two goals are to be reached. On the one hand the solutions should be as close to the global Pareto-optimal front as possible and on the other hand the solutions should also cover the whole Pareto-front. The first goal is often achieved through elitism by replacing random individuals with individuals on the Pareto front. The second goal can be achieved by punishing individuals that are too close together (Fitness Sharing).

According to Coello Coello (2006), the traditional evolutionary algorithms cannot efficiently deal with multi-objective optimization problems for two (2) reasons:

(1) Due to stochastic noise, evolutionary algorithms tend to converge to a single solution if run for a sufficiently large number of iterations. Thus, it is necessary to block the selection mechanism so that different solutions (non-dominated) are preserved in the population of an evolutionary algorithm.

(2) All non-dominated solutions should be sampled at the same rate during the selection stage, since all non-dominated solutions are equally good among themselves.

Over the past years researchers developed several approaches for the solution of multi-objective optimization problems with the use of EAs. The first implementation of a multi-objective evolutionary algorithm (MOEA) dates back to the mid-1980s (Schaffer, 1984, 1985). Since then, a considerable amount of research has been done in this area, now known as evolutionary multiobjective optimization. A brief review of the most well-known MOEAs is presented below.

2.1. Vector Evaluation Genetic Algorithm (VEGA)

Schaffer (1984, 1985) and Schaffer and Grefenstette (1985) introduce the Vector Evaluation Genetic Algorithm (VEGA) in the mid 1980s, which was the first implementation of a multi-objective evolutionary algorithm (MOEA). In reality VEGA was not nothing more than a simple genetic algorithm with a modified selection mechanism. A weakness of VEGA approach according to Schaffer is that the generated solutions are non-dominated in a local sense, because their non-dominance is limited to the current population. Schaffer highlighted another problem that in genetics is known as "speciation" which is the evolution of "species" within the population which excel in one dimension of performance, without considering at the other dimensions. Thus VEGA approach tend to reject individuals with acceptable performance, perhaps above average, but not outstanding for any of the objective functions. Schaffer called this tendency of VEGA to reject individuals with acceptable performance for any objective function as "middling", and it is the most serious shortcoming of VEGA method.

2.2. Niched Pareto Genetic Algorithm (NPGA)

NPGA developed by Horn, Nafpliotis, and Goldberg (1994). NPGA extends the traditional GA to multiple objectives through the use of Pareto domination ranking and fitness sharing (or niching). The main advantages of this method is its adaptability to a wide range of multiobjective optimization problems and its ability to search non-linear and discontinuous search spaces without relying on the need for continuous first and second derivatives.

2.3. Niched Pareto Genetic Algorithm II (NPGA-II)

Erickson, Mayer, and Horn (2001) proposed the NPGA-II. This algorithm relies on the traditional Pareto ranking approach, but it keeps its tournament selection scheme. Ties are solved through fitness sharing as in its predecessor. However, the niche count of the NPGA-II is computed using individuals from the next partially filled generation using a technique called "continuously updated fitness sharing".

2.4. Nondominated Sorting Genetic Algorithm (NSGA)

Nondominated Sorting Genetic Algorithm (NSGA) proposed by Srinivas and Deb (1994) is based on Goldberg (1989). Goldberg suggested a nondominated sorting procedure to overcome the weakness of VEGA to bias towards some Pareto-optimal solutions. Schaffer himself the creator of VEGA had realized VEGA's bias towards some Pareto-optimal solutions. NSGA eliminates the bias in VEGA and thereby distributes the population over the entire Pareto-optimal regions.