10

Alternative Metaheuristics

> Nature is trying very hard to make us succeed, but nature does not depend on us. We are not the only experiment.
>
> R. Buckminster Fuller

10.1 Introduction

Evolutionary Algorithms (EAs) are not the only search techniques that have been used to deal with multiobjective optimization problems. In fact, as other search techniques (e.g., Tabu search and simulated annealing) have proved to have very good performance in many combinatorial (as well as other types of) optimization problems, it is only natural to think of extensions of such approaches to deal with multiple objectives.

The Operations Research (OR) and EA communities have shown a clear interest in pursuing these extensions. Since the multiobjective formulation of combinatorial optimization problems (e.g., the quadratic assignment problem) are known to be NP-complete [1465, 429], they present real challenges to researchers. Additionally, many real-world problems (e.g., scheduling) require efficient approaches that can at least approximate P_{true} and PF_{true} in a reasonable amount of time.

Any search technique such as those discussed in this chapter, can be extended to deal with multiple objectives in several ways, just as in the case of EAs (see Chapter 2). One could just aggregate the objective functions to form a single scalar value, or to use a target vector approach (defining ideal goals to be achieved by each objective and aggregating their differences with respect to the values obtained). However, dominance can also be checked locally (between two solutions generated by the algorithm) and then keep in an archive every nondominated solution generated over time, so that dominance can also be checked globally (i.e., with respect to this archive). Knowles and Corne [886] have argued that the use of a naive two-membered evolution strategy

(with an external archive) is sufficient to generate PF_{true} for relatively complex multiobjective optimization problems.

The issues are then of a different nature. For example: how to move from a certain state to another, or how to ensure that different portions of PF_{true} are being generated rather than only a certain fraction of it. Additionally, other issues such as diversity are an important concern with the heuristics of this chapter as well as with MOEAs.

This chapter is organized as follows. Section 10.2 discusses simulated annealing and the main proposals to extend it to problems with multiple objectives. Tabu search and scatter search as well as their corresponding multiobjective extensions are discussed in Section 10.3. The ant system (including the Ant-Q algorithm) is the subject of Section 10.4. Distributed reinforcement learning is analyzed in Section 10.5. Particle swarm optimization, differential evolution and artificial immune systems are discussed in Sections 10.6, 10.7 and 10.8, respectively.

Finally, Section 10.9 covers other promising heuristics that are good candidates for solving multiobjective optimization problems (i.e., cultural algorithms and cooperative search).

10.2 Simulated Annealing

As mentioned in Chapter 1 (Section 1.4), simulated annealing is a stochastic search algorithm based on the concept called "annealing". The annealing process consists of first raising the temperature of a solid to a point where its atoms can freely (i.e., randomly) move and then to lower the temperature, forcing the atoms to rearrange themselves into a lower energy state (i.e., a crystallization process). During this process the free energy of the solid is minimized (the crystalline state is the state of minimum energy of the system). The cooling schedule is vital in this process. If the solid is cooled too quickly, or if the initial temperature of the system is too low, it is not able to become a crystal and instead the solid arrives at an amorphous state with higher energy. In this case, the system reaches a local minimum (a higher energy state) instead of the global minimum (i.e., the minimal energy state) [407, 1292].

10.2.1 Basic Concepts

Nicholas C. Metropolis et al. [1095] proposed an algorithm to simulate the evolution of a solid in a heat bath until it reached its thermal equilibrium. The Monte Carlo method was used to simulate the process, which started from a certain thermodynamic state of the system, defined by a certain energy and temperature. Then, the state was slightly perturbed. If the change in energy produced by this perturbation was negative, the new configuration was accepted. If it was positive, it was accepted with a probability given by