The Multi-Objective Constrained Assignment Problem *

Mark P. Kleeman
Air Force Institute of Technology
Dept. of Electrical and Computer Engineering
Graduate School of Engineering and Mgmt
WPAFB (Dayton), Ohio, 45433, USA
Mark.Kleeman@afit.edu

Gary B. Lamont
Air Force Institute of Technology
Dept. of Electrical and Computer Engineering
Graduate School of Engineering and Mgmt
WPAFB (Dayton), Ohio, 45433, USA
Gary.Lamont@afit.edu

Categories and Subject Descriptors

J.1 [**Computer Applications**]: Administrative Data Processing—*Military*

General Terms

Algorithms, Performance

Keywords

MOEA, scheduling, assignment problems

1. INTRODUCTION

There has been much research done in the area of assignment problems. These problems span many different disciplines and researchers from an array of different fields are exploring ways to solve them. As such, the assignment problem has been modified over the years to meet specific needs in different application areas. But how do the mathematical formulations of the various assignment problems compare with one another? By knowing the similarities and differences of these problems, a researcher can best fit a particular model to the problem he is attempting to solve. In [2], some of the basic differences between the various models are outlined and a baseline model that researchers can use as a starting point to build their specific model is presented. This abstract briefly describes the constrained assignment problem (CAP) and a specific problem called the Airman Assignment Problem (AAP) and then discusses the results of experiments using the NSGA-II.

2. ASSIGNMENT PROBLEMS

There has been much research done in the realm of "assignment problems". There is the assignment problem, the generalized assignment problem, the channel assignment problem, the frequency assignment problem, the pin assignment problem, the layer assignment problem, the traffic assignment problem, the quadratic assignment problem, the job assignment problem, the sailor assignment problem, the airman assignment problem, etc. Scheduling problems, such as the nurse scheduling problem, also have many of the same characteristics as some of the assignment problems.

Many researchers have applied numerous approaches to address specific assignment problems in their fields. Are each of the problems really that different? In [2], some of the various assignment problems are compared and classified based on their similarities and differences.

The constrained assignment problem is a generalization of many specific assignment problems, including the sailor assignment problem and the nurse scheduling problem. Most assignment problems have to deal with outside constraints on the problem. And with the exception of the quadratic assignment problem and the stable marriage problem, the problems typically don't have a one-to-one matching. Essentially, many assignment problems can be generalized to the following form:

$$\min \quad \sum_{i=1}^{N}\sum_{j=1}^{M} x_{i,j}y_{i,j} \tag{1}$$

$$\text{s.t.} \quad x_{i,j} \quad \text{and} \quad y_{i,j} \quad \text{are constrained}$$

where $x_{i,j}$ and $y_{i,j}$ are costs related to the assignment of an entity to a job. Not all assignments may be valid. Invalid assignments are classified as hard constraints. Assignments that are allowed, but are penalized, are called soft constraints. For example, a job that requires a certain skill can eliminate people from occupying that position. Hard constraints can be viewed as boolean types.

A soft constraint can be a trait, skill, or physical quality that is beneficial to fulfill the duties of a particular job, but they are not required. Unlike hard constraints, soft constraints can vary in significance. So different job skills may have varying penalties associated with them. For example, a person that has the job skill specified may get a high score for that constraint, while someone who has a job skill closely related to the job may be penalized with a higher score. But someone who is overqualified for the job may be penalized with a much higher score.

A typical CAP has a mixture of hard and soft constraints. These constraints can vary depending upon the job, so a hard constraint for one job can be a soft constraint for an-

*The views expressed in this article are those of the authors and do not reflect the official policy of the United States Air Force, Department of Defense, or the United States Government.

GECCO'06, July 8–12, 2006, Seattle, Washington, USA.
ACM 1-59593-186-4/06/0007.

other job. The CAP is quite common. It can be used in any problem where you are assigning personnel to jobs.

Airman Assignment Problem

For our research example, we used the United States Air Force assignment system as our model, which we call the Airmen Assignment Problem (AAP). Mathematically, the problem can be modelled as follows:

$$\min \quad \sum_{i=1}^{N} \sum_{j=1}^{M} F_{i,j} h_{i,j} a_{i,j} \tag{2}$$

where $F_{i,j}$ denotes the fitness (including any associated penalties from soft constraints) of assigning airman i to job j, H is a hard constraint matrix such that $h_{i,j} = 1$ when Airman i meets all hard constraints for job j and $h_{i,j} = 0$ otherwise, and A is an assignment matrix such that $a_{i,j} = 1$ when Airman i is assigned to job j and $a_{i,j} = 0$ otherwise.

The goal of the air force assignment system is to first meet the needs of the air force by filling all the necessary jobs with qualified individuals. The secondary goal is to try to put people into jobs that they want. Obviously, this is desirable in all organizations.

The problem has two objectives. The first objective is a measurement of how well each assignment meets the needs of the air force and at the same time satisfies the desires of the individual. This objective is met through the use of hard constraints and soft constraints. The hard constraints are the ones that must be satisfied for an assignment to be valid. The soft constraints are basically penalty functions. These penalties are applied when an assignment deviates from the ideal candidate or if the person is not given their ideal assignment. These penalties are weighted differently depending on their importance in the eyes of the decision maker. For our problem, we have 3 hard constraints (proper job training, rank, and security clearance) and 11 soft constraints, each with a user defined penalty function. The second objective function is the cost it takes to move each person from one duty location to another. Since we are optimizing two objectives this is actually a multiobjective constrained assignment problem (mCAP).

The goal is to minimize both objectives. For our initial runs of this problem, we generated "real-like" data that we deemed best fit the distribution for a typical assignment process in the air force. Air force assignments are decomposed into groups and then into subgroups. See [2] for more information on how problem instances were developed.

3. RESULTS AND ANALYSIS

The problem is initially integrated into the non-dominated sorting genetic algorithm - II (NSGA-II). NSGA-II is an MOEA designed to handle both real-valued and binary strings.

The main goal is to take a real world example of a CAP and find good solutions to the problem. Our secondary goal was to note how the algorithm explored the landscape. To do this, snapshots of the known Pareto front are saved at 100 generation intervals. A spacing algorithm and the two set coverage metric [1] determine the algorithm diversity and progression. The algorithm is run 30 times and the mean and standard deviation of each metric is determined. The NSGA-II is run with a standard "out of the box" configuration. We used a crossover probability of 0.9 and a mutation probability of 0.033.

For the AAP, an instance where 20 people are assigned to 30 jobs is used. This instance is chosen because it is a common size that small to mid-sized assignment team would have to deal with. As such, it is a typical problem that would be found in the real world. For more information on the design of experiments and metrics used, see [2].

The results of our experiments are discussed in more detail in [2]. Little change occurs after generation 700 and from generation 800 and up, the results are statistically similar. The spacing remained fairly constant throughout the run. This means that the solution vector maintained its spacing while converging on better solutions.

We also compared graphically the difference in the known Pareto fronts at generations 100 and 1000. These graphs were generated by finding the known Pareto front from all 30 runs. Figure 1 is a comparison with a population size of 200. It appears that the algorithm finds the interior points along the Pareto front first and then it expands outward.

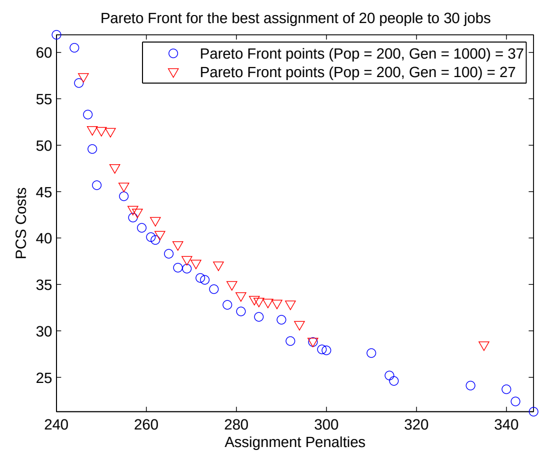

Figure 1: Comparison of PF known generated with population of 200 at generations 100 and 1000

Assignment problems are very common in the literature. They are quite common in the real-world as MOPs. In this abstract, we presented a generalized assignment problem called the constrained assignment problem. We then implemented a real-world, multi-objective example of this type of problem called the airman assignment problem. It was noted that MOEAs are well suited for these problems due to their landscape and complexity. We applied the NSGA-II algorithm to the problem and we found that it was able to find good solutions to the problem fairly quickly. Future research includes running more instances, including large instances similar to those used in army and navy scheduling, as well as comparing our results to other MOEAs.

4. REFERENCES

[1] C. A. Coello Coello, D. A. Van Veldhuizen, and G. B. Lamont. *Evolutionary Algorithms for Solving Multi-Objective Problems*. Kluwer Academic Publishers, New York, May 2002. ISBN 0-3064-6762-3.

[2] M. P. Kleeman and G. B. Lamont. The Multi-objective Constrained Assignment Problem. Technical Report AFIT-CSCE-2006-01, Air Force Institute of Technology, Wright-Patterson AFB, OH, USA, 2006.