Incremental Refinement of Solutions for Multiple Objective Optimization Problems

C. Mariano-Romero
Mexican Institute of Water Technology
Paseo Cuauhnahuac 8532, Progreso
Jiutepec, Morelos, Mexico
cmariano@tlaloc.imta.mx

V. Alcocer-Yamanaka
Mexican Institute for Water Technology
Paseo Cuauhnahuac 8532, Progreso
Jiutepec, Morelos, Mexico
valcocer@tlaloc.imta.mx

E. F. Morales
INAOE
Luis Enrique Erro No. 1, Santa Maria Tonantzintla
Puebla, 72840, Mexico
emorales@inaoep.mx

ABSTRACT

MDQL is an algorithm, based on reinforcement learning, for solving multiple objective optimization problems, that has been tested on several applications with promising results [1]. *MDQL* discretizes the decision variables into a set of states, each associated with actions to move agents to contiguous states. A group of agents explore this state space and are able to find Pareto sets applying a distributed reinforcement learning algorithm. The precision of the Pareto solutions depends on the chosen granularity of the states. A finer granularity on the states creates more precise solutions but at the expense of a larger search space, and consequently the need for more computational resources. An important improvement is presented. The new algorithm, called *IMDQL*, starts with a coarse granularity to find an initial Pareto set. A vicinity for each of the Pareto solutions in refined and a new Pareto set is founded in this refined state space. This process continuous until there is no more improvement within a small threshold value. It is shown that *IMDQL* not only improves the solutions found by *MDQL*, but also converges faster and is capable to approximate dynamic Pareto fronts.

The main consideration in the application of *IMDQL* to dynamic environments is that the agents in the algorithm start from the Pareto solutions obtained. Agents start with a deterministic environment constructed with fixed values for the value functions for the first dynamic parameters; when convergence is reached and a Pareto set is obtained, a new cycle is started, changing to the next value for the dynamic parameters. Agents start searching (adapting solutions) from the existing environments which correspond to the Pareto solutions obtained for the previous value for the dynamic parameters. Searching for new solutions, from the last Pareto set, given the new values for the dynamic parameters, significantly reduces the convergence time.

IMDQL is tested on a real water distribution network design involving water-reusing treatment plants and different contaminants concentrations [1]. In this problem, the concentration of contaminants can change over time so the search for optimal solutions becomes a continuous process. It is shown that *IMDQL* improves on the solutions found by *MDQL* with fixed concentrations and, that due to its incremental nature, it is able to adequately adjust its Pareto set solutions with dynamic changes in the contaminants concentrations as long as they are within the vicinity of the previous Pareto set. This is, to our knowledge, the first multi-objective optimization algorithm that is able to dynamically adjust the Pareto set with changing conditions and that can adjust the accuracy of its solutions.

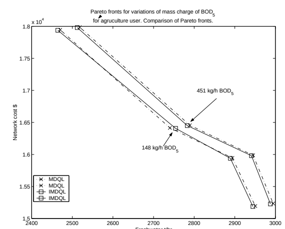

Figure 1: Agriculture user

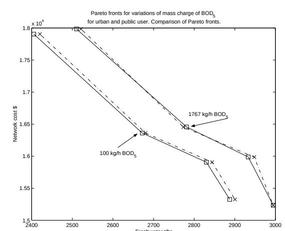

Figure 2: Urban and Public user

The solutions were also compared against *MDQL* that was compared against a reduced gradient method using a weighted combination of the two objective functions [1], see Figures 1 and 2 for comparisons of Pareto fronts obtained with *MDQL*, $\times$, and *IMDQL*, $\square$.

1. REFERENCES

[1] C. Mariano, V. Alcocer, and E. Morales. Multiobjective water pinch analysis of the cuernavaca, city water distribution network. *European Journal on Operational Research*, in press 2007.

GECCO'07, July 7–11, 2007, London, England, United Kingdom.
ACM 978-1-59593-697-4/07/0007.