

Towards a novel optimisation approach with simultaneous knowledge acquisition for distributed computing environments

Siyu Yang, Antonis Kokossis and Patrick Linke

Centre for Process and Information Systems Engineering, University of Surrey,
Guildford, GU2 7XH, U.K.

Current optimisation methods are of limited use for decision-support in complex systems due to two main short-comings. Firstly, they require long computational times to identify optimal solutions for large-scale problems and the algorithms are not easily parallelised for use in large-scale distributed computing environments. Such environments become increasingly available with the advent of Grid Technologies and new generations of optimisation methods are required that can exploit the vast available distributed computing resources effectively. Secondly, the results obtained from optimisation runs are often difficult to interpret by the user in the context of the decisions to be taken. This is particularly true for stochastic optimisation methods, where important solution features are often blurred by features not strongly impacting on the systems performance. An optimisation approach that could provide the user with optimal solutions alongside insights into the importance of individual solution features would be very attractive for decision-support.

This paper reports on ongoing developments of a novel optimisation scheme that is amenable for highly distributed computing environments and enables data analysis and knowledge acquisition in the course of the optimisation. The scheme is built upon the strategy of Simulated Annealing: transitions are considered towards improving as well as dis-improving states with the probability of significant uphill excursions being reduced from high towards low system temperature. Hence, the system is gradually forced towards states of better performance; however, on the path towards better performance the system is

allowed to attain inferior states. At each temperature, a number of random, reversible state transitions (homogeneous Markov chains) are performed to equilibrate the system. In contrast to Simulated Annealing, which is a sequential optimisation approach that gradually transforms an initial solution into an optimal solution through a number of Markov processes at reducing temperatures, the proposed optimisation schemes consists of a number of “solution pools”, each of which is associated with a system temperature. The solutions in these pools are generated by performing constant temperature Markov processes on existing solutions in these pools and dynamically assigned on the basis of temperature and solution quality dependent acceptance criteria. All pools are stored in a central database. The Markov processes can be completed quickly in distributed computing environments. As new solutions are constantly generated during the optimisation run, solutions are migrated between the pools so as to ensure all solutions in a given pool are acceptable. The highest temperature pool contains a wide distribution of solutions, whereas the lowest temperature pool contains only the best solutions with a low distribution of solution quality. The solution database contains all the information acquired during the optimisation as well as information on the evolution of solutions from low to high temperature pools. Data mining techniques can be used on the solution database to extract solution features that are essential to the systems performance.

The paper will focus on the general concepts of the new optimisation scheme as well as on implementation issues including controls of the algorithm. We will also present applications to a number of global optimisation test problems to illustrate the performance of the approach.