

Systematic Integration of Parameterized Local Search Techniques in Evolutionary Algorithms

Neal K. Bambha¹, Shuvra S. Bhattacharyya¹,
Jürgen Teich², and Eckart Zitzler³

¹ Department of Electrical and Computer Engineering, and
Institute for Advanced Computer Studies,
University of Maryland,
College Park, MD USA

{nbambha, ssb}@eng.umd.edu

² Computer Science Institute,
Friedrich-Alexander University,
Erlangen-Nuremberg

³ Department of Computer Engineering
Swiss Federal Institute of Technology

Abstract. Application-specific, parameterized local search algorithms (PLSAs), in which optimization accuracy can be traded off with runtime, arise naturally in many optimization contexts. We introduce a novel approach, called simulated heating, for systematically integrating parameterized local search into evolutionary algorithms (EAs). Using the framework of simulated heating, we investigate both static and dynamic strategies for systematically managing the trade-off between PLSA accuracy and optimization effort. Our goal is to achieve maximum solution quality within a fixed optimization time budget. We show that the simulated heating technique better utilizes the given optimization time resources than standard hybrid methods that employ fixed parameters, and that the technique is less sensitive to these parameter settings. We demonstrate our techniques on the well-known binary knapsack problem and two problems in electronic design automation. We compare our results to the standard hybrid methods, and show quantitatively that careful management of this trade-off is necessary to achieve the full potential of an EA/PLSA combination.

1 Introduction

For many optimization problems, efficient algorithms exist for refining arbitrary points in the search space into better solutions. Such algorithms are called *local search algorithms* because they define neighborhoods, typically based on initial “coarse” solutions, in which to search for optima. Many of these algorithms are parameterizable in nature. Based on the values of one or more algorithm parameters, such a *parameterized local search algorithm (PLSA)* can trade off time or space complexity for optimization accuracy.

In this poster, we describe a technique called *simulated heating* [1], which systematically incorporates parameterized local search into the framework of global search. The idea is to increase the time allotted to each PLSA invocation during the optimization process—low accuracy of the PLSA at the beginning and high accuracy at the end. This is in contrast to most existing hybrid techniques, which consider a fixed local search function, usually operating at the highest accuracy.

We demonstrate our techniques on three applications—the well-known binary knapsack problem, a memory cost optimization problem in embedded systems, and a voltage scaling optimization problem for multiprocessor systems.

A simulated heating optimization begins with a low cost and accuracy for the PLSA. At certain points in time during the optimization the cost and accuracy are increased. The goal is to focus on the global search at the beginning and to find promising regions of the search space first; for this phase, the PLSA runs with low accuracy, which in turn allows a greater number of optimization steps of the global search. Afterward, more time is spent by the PLSA in order to improve the solutions found or to assess them more accurately. As a consequence, fewer global search operations are possible during this phase of optimization. Since the accuracy is systematically increased during the process, we use the term *simulated heating* for this approach by analogy to simulated annealing where the ‘temperature’ is continuously decreased according to a given cooling scheme.

Within the context of simulated heating optimization, we consider both static and dynamic strategies for systematically increasing the PLSA accuracy and the corresponding optimization effort. In the static optimization, the PLSA is varied at regular intervals. In the dynamic optimization, the PLSA is varied according to the overall progress of the optimization and the time remaining.

In most heuristic optimization techniques, there are some parameters that must be set by the user. In many cases, there are no clear guidelines on how to set these parameters. Moreover, the optimal parameters are often dependent on the exact problem specification. We show that the simulated heating technique, while still requiring parameters to be set by the user, is less sensitive to the parameter settings.

Our results for the three different applications show that, in the context of a fixed optimization time budget, simulated heating better utilizes the time resources and outperforms the standard fixed parameter hybrid methods. We show that careful management of this trade-off is necessary to achieve the full potential of an EA/PLSA combination, and to develop an efficient strategy for achieving this trade-off management.

References

1. N. K. Bambha, S. S. Bhattacharyya, J. Teich, and E. Zitzler, “Systematic integration of parameterized local search into evolutionary algorithms,” *IEEE Transactions on Evolutionary Algorithms*, April 2004, to appear.