Neighbourhood Selection in Variable Neighborhood Search

Stefan Voss

Chair and Director of the Institute of Information Systems,
Faculty of Economics and Social Sciences, University of Hamburg
Von-Melle-Park 5, D-20146 Hamburg, Germany

Abstract: Variable neighborhood search (VNS) with its numerous successful applications is based on a simple principle: systematic changes of neighborhoods within the search, both in the descent to local minima and in the escape from them. Many extensions have been made, mainly to be able to solve large problem instances. Nevertheless, a main driver behind VNS is to keep the simplicity of the basic scheme. Variable neighborhood descend (VND) belongs to the family of methods within VNS. The idea behind VND is to systematically switch between different neighborhoods where a local search is comprehensively performed within one neighborhood until no further improvements are possible. Given the found local optimum, VND continues with a local search within the next neighborhood. If an improved solution is found one may resort to the first neighborhood; otherwise the search continues with the next neighborhood, etc.

Designing these neighborhoods and applying them in a meaningful fashion is not an easy task. Moreover, incorporating these neighborhoods, e.g., into VNS may still need a considerable degree of ingenuity and may be confronted with the following research questions (among others):

- Which neighborhoods can be designed?
- How should they be explored and how efficient are they?
- How structurally different are they?
- In which order should they be applied?

While all of these questions seem to deserve in-depth study, in this paper we focus especially the latter concern. Assume that we are given an optimization problem that is intended to be solved by applying the VNS scheme, how many and which types of neighborhoods should be investigated and what could be appropriate selection criteria to apply these neighborhoods. More specifically, does it pay to "look ahead" when attempting to switch from one neighborhood to another? Is it reasonable to apply "nested neighborhood structures" (like first 2-optimal exchanges, then 3-optimal exchanges, then 4-optimal etc.) or should neighborhood structures be considerably different from each other? Often different neighborhoods are proposed allowing for ideas regarding sequential and nested changes. We strive to provide insights on these questions on different optimization problems.