Local search for mixed-integer nonlinear optimization: methodology and applications

Thierry Benoist[1], Bertrand Estellon[2],
Frédéric Gardi[1], Antoine Jeanjean[1], Karim Nouioua[2]

[1] Bouygues e-lab, 40 rue Washington, 75008 Paris, France

[2] Laboratoire d'Informatique Fondamentale – CNRS UMR 6166,
Université Aix-Marseille II, Parc Scientifique et Technologique de Luminy,
case 901, 163 avenue de Luminy, 13288 Marseille cedex 9, France

fgardi@bouygues.com

Keywords : Mixed-integer nonlinear optimization, local search, engineering methodology, very large-scale problems, industrial applications.

Local search [1] is a practical solution approach for tackling large-scale discrete optimization problems, particularly whose arising in real-life applications. The authors have designed a methodology [5] for designing and engineering high-performance local-search heuristics in such contexts. This methodology has been shown to be successful for solving various industrial problems; several of these applications have been awarded on the occasion of OR competitions (1er Junior/Senior Prize at ROADEF 2005 Challenge [4], 2nd Senior Prize at ROADEF 2007 Challenge [5]). Recently, this methodology was extended to deal with mixed-integer optimization problems (linear as well as nonlinear). Note that local search is rarely used for solving mixed-integer optimization problems. A classical way to address these problems in practice is to use decomposition approaches (empirical or mathematical). Here we present the main ingredients of our methodology and its application for solving three industrial problems with high economic stakes (but short running times): rich inventory routing [2], resource scheduling for mass transportation [8], nuclear maintenance planning [6]. This methodology seems to be suited for tackling large-scale unit commitment problems arising in energy management, especially in operational contexts (daily or hourly use, very short running times, on-line replanification).

The first particularity of our local search is to be *pure* and *direct*. Indeed, *no decomposition* is done; the problem is tackled frontally. The search space explored by our algorithm is close to the original solution space. In particular, the combinatorial and continuous parts of the problem are treated together: combinatorial and continuous decisions can be simultaneously modified by a move during the search. By avoiding decompositions or reductions, no solution is lost and the probability to find good-quality ones is increased. Then, *no hybridization* is done: no particular metaheuristic is used, no tree-search technique is used. In this way, we avoid complex parameter tuning and simplify the architecture of the resulting software. Then, the second specificity of our local search is to be *highly randomized*, in order to avoid bias while exploring the search space. Such a diversification of the search is obtained by exploring in a first-improvement fashion a large variety of randomized neighborhoods. The union of these (small) randomized neighborhoods induces in effect a very large neighborhood, allowing to converge in practice toward high-quality local optima despite hard constraints. Finally, its third specificity is to be *very aggressive: millions of feasible solutions are visited within the time limit*. Indeed, randomized local search is a non deterministic, incomplete exploration of the search space. Therefore, exploring a huge number of (feasible) solutions during the allocated time augments the probability to find good-quality solutions.

Our local-search heuristics are composed of three layers: general heuristic, moves, evaluation machinery. The evaluation machinery forms the engine of the local search; it computes the impacts of moves on constraints and objectives during the search. The time spent to engineer each layer during the project follows the following distribution: 10 % on general heuristic, 20 % on moves, 70 % on evaluation machinery. In summary, our work was focused on: *designing a large variety of randomized moves* allowing an effective exploration of the search space, and *speeding up the evaluation of these moves*. In a mixed-integer optimization context, these two points are declined as follows. First, the moves are designed in order to treat together the combinatorial and continuous dimensions of the problem. For this, discrete and continuous decisions are simultaneously modified by the moves during the search. Then, a difficulty arises: recovering the feasibility of the continuous part of the solution for evaluating the move. Roughly speaking, it imposes to be able to solve the continuous subproblem, which is generally very time-consuming. That is why the second point is concentrated on implementing an *incremental randomized combinatorial algorithm* for solving approximately but very efficiently the continuous subproblem (even if it is polynomial-time solvable). Numerical experiments show that such approximate algorithms are more than *1000 times faster* than state-of-the-art *exact* algorithms (in particular, linear programming approaches), while providing near-optimal solutions.

As first application, a new practical solution approach based on randomized local search is presented for tackling a real-life inventory routing problem [2]. This problem was posed in 2008 to Bouygues e-lab by a major French industrial actor. Inventory routing refers to the optimization of transportation costs for the replenishment of customers' inventories: based on consumption forecasts, the vendor organizes delivery routes. The instances contain several hundreds customers, while the running time is limited to 5 minutes. Our model takes into account pickups, time windows, drivers' safety regulations, orders and many other real-life constraints. This generalization of the vehicle routing problem was often handled in two stages in the past: inventory first, routing second. On the contrary, a characteristic of our local-search approach is the absence of decomposition, made possible by the use of a fast incremental greedy heuristic for maximizing flows in a directed acyclic graph. Moreover, thanks to a large variety of randomized neighborhoods, a simple first-improvement descent is used instead of tuned, complex metaheuristics. An extensive computational study shows that our solution is effective, efficient and robust, providing long-term savings exceeding 20 % on average compared to solutions built by expert planners or even a classical urgency-based constructive algorithm. Confirming the promised gains in operations, the resulting decision support system is progressively deployed worldwide.

The second application addressed is a real-life earthmoving optimization problem [8]. The problem takes as input the optimal solution of a mass transportation problem which is one of the first operation research problems, known as Monge-Kantorovich problem (introduced by Monge in 1781 in his famous "Mémoire sur la théorie des déblais et des remblais"). The earthmoving optimization problem addressed here consists in scheduling a set of resources traveling between blocks located on a linear axis, while ensuring the transportation of earth quantities from sources to destinations. This problem was posed in 2009 to Bouygues e-lab by DTP Terrassement, a subsidiary of Bouygues Construction. The resulting software, based on randomized local search, is now exploited for optimizing the provisional schedule of large linear construction sites (highways and high-speed railways) over several years, in less than 1 minute of running time.

The third application was posed by Électricité de France (EDF) as subject of the ROADEF/EURO Challenge 2010 [9]. It is focused on the medium-term (5 years) management of the EDF French thermal power park, and especially of nuclear plants which have to be repeatedly shut down for refueling and maintenance. The resulting optimization problem is a very large-scale mixed-integer nonlinear problem (involving more than one billion decision variables), whereas the running time is restricted to 1 hour on a standard computer. The solution approach that we have implemented during the competition [6] is a randomized local search, which follows the methodology presented above. The results obtained by our algorithm are among the best ones of the competition. Benchmarks are divided into three categories A, B, X containing each one 5 instances. Our algorithm was ranked 1st on instances A and B (among 44 teams engaged, 16 finalists), before falling to the 8th place due to a late-working-hours bug appearing on some instances X (note that only 4 teams among the 16 finalists have been able to provide all the solutions to instances X). Once corrected, our algorithm provides state-of-the-art results on instances X in conditions similar to ones of the Challenge. The results on the realistic instances B, as computed by the Challenge's organizers[3], show an average gap greater than 1 % (resp. 10 %) between our solutions and the ones of the team ranked 3rd (resp. 6th). Such gaps are important because corresponding from dozens to hundreds million euros of savings [7]. One can observe that the majority of approaches proposed by the other competitors corresponds to MIP/CP-based decomposition heuristics.

References

1. E. Aarts, J.K. Lenstra, eds (1997). *Local Search in Combinatorial Optimization.* Wiley-Interscience Series in Discrete Mathematics and Optimization, John Wiley

[3] `http://challenge.roadef.org/2010/files/talks/finalResultsEDF.pdf`

& Sons, Chichester, England, UK.
2. T. Benoist, B. Estellon, F. Gardi, A. Jeanjean (2010). Randomized local search for real-life inventory routing. *Transportation Science* (to appear).
3. B. Estellon, F. Gardi, K. Nouioua (2006). Large neighborhood improvements for solving car sequencing problems. *RAIRO Operations Research* 40(4), pp. 355–379.
4. B. Estellon, F. Gardi, K. Nouioua (2008). Two local search approaches for solving real-life car sequencing problems. *European Journal of Operational Research* 191(3), pp. 928–944.
5. B. Estellon, F. Gardi, K. Nouioua (2009). High-performance local search for task scheduling with human resource allocation. In *Proceedings of SLS 2009, the 2nd International Workshop on Engineering Stochastic Local Search Algorithms* (T. Stützle, M. Birattari, H.H. Hoos, eds.), *Lecture Notes in Computer Science* 5752, pp. 1–15. Springer, Berlin, Germany.
6. F. Gardi, K. Nouioua (2011). Local search for mixed-integer nonlinear optimization: a methodology and an application. In *Proceedings of EvoCOP 2011, the 11th European Conference on Evolutionary Computation in Combinatorial Optimisation* (P. Merz, J.-K. Hao, eds.), *Lecture Notes in Computer Science* 6622, pp. 167–178. Springer, Berlin, Germany.
7. A. Gorge (2008). Planification des arrêts pour rechargement des centrales nucléaires. In *JOR 2008, la 3ème Journée de Recherche Opérationnelle et Optimisation dans les Réseaux*, Paris, France (oral communication).
8. A. Jeanjean (2010). Resource scheduling optimization in mass transportation problems. In *Abstracts of PMS 2011, the 12th International Workshop on Project Management and Scheduling* (V. T'Kindt, ed.). Tours, France.
9. M. Porcheron, A. Gorge, O. Juan, T. Simovic, G. Dereu (2010). ROADEF/EURO 2010: a large-scale energy management problem with varied constraints. EDF R&D, Clamart, France (27 pages).
`http://challenge.roadef.org/2010/index.en.htm`