Unconventional Pivoting Rules for Local Search

Matthieu Basseur[1] and Adrien Goëffon[1]

Laboratoire d'Étude et de Recherche en Informatique d'Angers, 2 bd Lavoisier, 59045 Angers Cedex 01
matthieu.basseur@univ-angers.fr, adrien.goeffon@univ-angers.fr

1 Towards Worst Improvement

Solving combinatorial optimization problems using local search techniques consists in transforming an initial configuration by application of local moves chosen from a neighborhood structure. During the last decades, a large panel of neighborhood searches have been designed, like hill-climbing, tabu search, simulated annealing or iterated local search. These metaheuristics mainly differ in the move selection heuristic employed. Hill-climbings are simple local search techniques and are widely used as basic elements of more sophisticated metaheuristics.

Two commonly-used climbing move strategies, first and best improvement, constitute the widely used pivoting rules. These rules define how to select a better neighbor from a not locally optimal configuration [4, 2]. More precisely, the *best improvement* strategy consists in selecting, at each iteration, a neighbor which achieves the best fitness. This implies to generate and evaluate the whole neighborhood at each step of the search. The *first improvement* strategy accepts the first evaluated neighbor which satisfies the moving condition and avoids the systematic generation of the entire neighborhood. The *worst improvement* strategy, can easily be considered but is never envisaged for the design of local search algorithms. We focus here on determining the actual efficiency of this underestimated move strategy.

In a previous study [1], we showed that first improvement tends to outperform best improvement on many landscapes derived from several combinatorial problems. An analysis of landscapes was also provided, indicating that first improvement is mostly efficient to explore most landscapes, exepting smooth ones. Intuitively, since best improvement systematically choose the highest neighbor, it should conduce the search towards the nearest peak (local optima) of the landscape. On the contrary, first improvement often performs reduced improvements, which tends to drive progressively the search toward higher areas, where the potential local optima are also higher. Then, it seems interesting to determine if choosing the worst improving neighbor at each step of the search increases the possibility to reach higher areas of landscapes by avoiding the climbing of small steepest peaks.

2 Experiments on NK landscapes

The NK family of landscapes has been proposed in [3] in order to generate artificial combinatorial landscapes with tunable shape properties: size and ruggedness. NK landscapes use a basic search space, with binary strings as configurations and bit-flip as neighborhood. Characteristics of an NK landscape are determined by two parameters N and K. N refers to the length of binary string configurations. $K \in \{0, \ldots, N-1\}$ specifies the level of variables interdependency, which directly affects the ruggedness of the landscape. By increasing the value of K from 0 to $N-1$, NK landscapes can be tuned from smooth to rugged.

Experimental process aim at comparing 3 hill-climbing versions (best, first and worst improvement), as well as several approximated worst improvements, which consist in selecting the worst neighbor among the k first improving ones. These 7 climber variants will be respectively denoted as B, F, W, W_k ($k \in \{2, 4, 8, 16\}$). In our experiments, we have considered many NK landscapes parameterizations. Results for $N = 1024$ and $K \in \{4, 6, 8, 10, 12\}$[1] are given in table 1.

100 random configurations were generated for each NK landscape, which will be used as starting points for hill-climbings. For each NK landscape, 100 executions of the 7 hill-climbing versions were performed. Searches are stopped when a local optimum is reached. Since we focus on the quality

[1] In a previous study, we reported that first improvement performs better than best improvement on rugged NK landscape instances, where $K \geqslant 4$.

of the local optima reached by hill-climbings variants, we report the average fitness of the 1,000 local optima resulting from the corresponding searches. Comparison between best, first and worst improvement show that worst improvement is clearly more efficient that first improvement each time first improvement outperforms best improvement.

However, the worst improvement strategy is time-consuming for two reasons: the number of steps needed to reach a local optimum is increase, and the whole neighborhood has to be generated for ensuring the selection of the worst neighbor. This leads us to consider alterative hill-climbing variants which approximate worst improvement. Comparative results of the 7 hill-climbings shows that the quality of the local optima obtained is negatively correlated with the average quality of the selected improving solutions. Globally, $W \succ W_{16} \succ W_8 \succ W_4 \succ W_2 \succ F \succ B$. Worst improvement approximation (W_k) is particularly interesting since it clearly requires less solution evaluations than complete worst improvement. Setting k consists in determining the best compromise between hill-climbing efficiency and computational costs. This is emphasized in figure 1, which report the average fitness evolution w.r.t. the number of solutions evaluated on a 1024_6 instance, focusing on the first 400000 evaluations, which is enough to terminate all local searches, expect the (complete) worst improvement strategy. To summarize this work, experiment showed that choosing the worst improving neighbor often leads to attain better local optima. Morevoer by slightly modifying the worst improvement strategy, one can design efficient hill-climbings which outperform first and best improvement in terms of tradeoff between quality and computational effort.

N_K	B	F	W_2	W_4	W_8	W_{16}	W
1024_4	.7232	.7238	.7253	.7270	.7286	.7291	**.7298**
	302k	12k	25k	54k	112k	223k	2336k
1024_6	.7223	.7250	.7280	.7310	.7330	.7343	**.7353**
	251k	13k	29k	65k	145k	311k	4151k
1024_8	.7176	.7215	.7251	.7285	.7306	.7316	**.7330**
	214k	13k	31k	74k	172k	382k	5837k
1024_10	.7121	.7165	.7204	.7235	.7255	.7265	**.7270**
	187k	14k	33k	80k	189k	439k	7325k
1024_12	.7064	.7107	.7150	.7178	.7197	.7206	**.7210**
	166k	14k	34k	84k	206k	487k	8544k

Table 1. Comparison of best, first, worst based pivoting rules (average fitness of local optima reached from 1,000 hill-climbings distributed on 10 random instances).

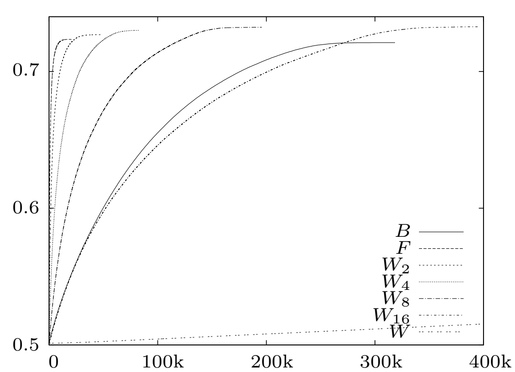

Fig. 1. Average fitness evolution on a 1024_6 NK instance, w.r.t. the number of evaluations (focus on the first 400,000 evaluations).

References

1. Matthieu Basseur and Adrien Goëffon. Hill-climbing strategies on various landscapes: an empirical comparison. In *Genetic and Evolutionary Computation Conference*, pages 479–486. ACM, 2013.
2. Holger Hoos and Thomas Stützle. Stochastic Local Search: Foundations & Applications. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2004.
3. Stuart A. Kauffman. *The Origins of Order: Self-Organization and Selection in Evolution.* Oxford University Press, USA, 1 edition, 06 1993.
4. Mihalis Yannakakis. The analysis of local search problems and their heuristics. In *Proceedings of the seventh annual symposium on Theoretical aspects of computer science*, STACS 90, pages 298–311, New York, NY, USA, 1990. Springer-Verlag New York, Inc.