

Tuning Tabu Search Strategies via Visual Diagnosis

Hoong Chuin LAU* Wee Chong WAN† Steven HALIM§

*School of Information Systems, Singapore Management University
hclau@smu.edu.sg

†The Logistics Institute Asia Pacific, National University of Singapore
tliwwc@nus.edu.sg

§ School of Computing, National University of Singapore
stevenha@comp.nus.edu.sg

1 Introduction

While designing metaheuristics can be straightforward, tuning the underlying *search parameters, configurations, and strategies* (collectively known as *search strategies*) to solve the problem well can be tricky. Since different problems or even instances of the same problem have their own search strategies that work well, some algorithm designers resort to trial-and-error through extensive experimentation. Others have adopted reactive or adaptive strategies, in which past knowledge or experience is used to set up the “adaptation” rules. From the industry standpoint, this process is not productive against a backdrop of tight development schedules.

Alternatively, human intelligence and machines can collaborate to shorten developmental time through the use of a well-designed visualization and interaction tool. Human plus computer collaboration has obtained considerable success in solving complex tasks such as CAD/CAM and combinatorial optimization problems. With the help of a visual diagnostic tool, an algorithm designer is able to examine search trajectories more systematically, steer the search and instantly see the impact of his action. This significantly reduces the time to design good search strategies.

Using visualization to assist optimization has been proposed in the seminal work of [5]. In this paper, we propose a human-guided scheme that collaborates with the tabu search to determine quickly an optimized set of adaptive rules. Unlike other works (e.g. [6]) which focus on problem-specific visualization, we emphasize the design of a *generic* tool called the *Visualizer for Metaheuristics Development Framework* (V-MDF), which is an extension on the work of [8]. Instead of relying on the specific problem domain information, V-MDF seeks to capture a pictorial view of the search trajectories and reports any anomalies to the human user. By visual inspection of these anomalies, the operator could determine with a high accuracy on the problems encountered in the search and consequently apply remedial strategies (such as tuning parameters, adjusting configurations, or deriving better adaptive rules). With V-MDF, the

Vienna, Austria, August 22–26, 2005

algorithm designer begins with some defined search strategies, and with the aid of the visualizer, observes the behavior of the search and dynamically changes the search strategies.

V-MDF differs from existing approaches for tuning metaheuristic which focus on the design of an efficient method for choosing the best parameter/configuration, e.g. [1,3]. Instead, we extend the idea of visualizing and analyzing the fitness landscape proposed in [4,9] to help users design better metaheuristics *on-the-fly*. This feature makes V-MDF especially useful for designing metaheuristics for new problems where search strategies have not been well-defined.

2 Tabu Search Concepts and Challenges

The strength of tabu search lies in its adaptive memory and intelligent exploration, which in turn lies in choosing a “good” set of parameters. For instance, the primary function of the adaptive memory is to prevent the search from revisiting explored solutions, and depends heavily on a good selection of tabu tenure. Unfortunately, an optimum tenure is hard to achieve and reactive methods such as [2] have been proposed to automatically and adaptively tune the value. However, the adaptive “rules” are often problem-specific and time-consuming to derive. The tuning process is further complicated with the usage of intelligent explorations (intensification or diversification) as these strategies are strongly dependent on the correct “timings” in which they are applied, which in turn introduce more parameters and rules. Recently, more complex intensifying and diversifying strategies have been proposed in the form of hybridization, in which tabu search is hybridized with other metaheuristics and/or with techniques such as linear programming and branch and bound. While such hybridization can further exploit the beneficial effect of intensification/diversification, it also adds another dimension of complexity to tune the strategies. In short, we see that the performance of tabu search is coupled tightly with its parameters, and it is the tuning of such a large number of parameters that can be a stumbling block to algorithm designers from devising complex search strategies. This leads us to the concept of a generic visualizer whose function is to allow the user to evolve his algorithm through visual observations of the search trajectories and fine tuning rules and parameters over a set of training instances. The user actions will be captured as guiding rules in a knowledge base such that the most promising parameter values and rules for the training sets will eventually form the elements of the tabu search algorithm for that problem.

3 Visual Diagnosis Tuning

We introduce some terminology. A *search trajectory* is stated as the path between the start and end of the search. Along this path, the search may encounter events (such as a new best solution found, solution cycling, etc). We define *incidents* as the occurrence of an event or a sequence of events which can be diagnosed visually. In response to certain incidents, one might decide to change the algorithm *strategy* such as to apply intensification/diversification or to adjust the search parameters. This new strategy might give rise to further incidents, and the process is

repeated. Each cycle can be seen as what we call a *{cause–action–outcome}* tuple which can be graphically presented for a user to evolve a good search strategy quickly.

The underlying graphical interface required to support visual diagnosis tuning is what we term in V-MDF as the *Distance Radar*. This tool displays the *degree of similarity (distance)* between solutions, measured by the number of local moves needed to translate one solution to the others, e.g: Hamming distance between the content of two solutions. The concept of distance has been proposed in [4,9]. For efficiency reason, distance computation should be done in linear time. For cases in which linear time is unattainable, approximation methods have been proposed in [4,9].

3.1 Distance Radar

The functionality of the Distance Radar is to display the incidents occurred in the search trajectory. These incidents either indicate the necessity for a remedial action or to display the outcome of an applied strategy. For example, the {passive searching – random restart – new best solution found} tuple signifies the effectiveness of diversification strategy, whereas {solution cycling – decrease tabu tenure – solution cycling} shows the ineffectiveness of decreasing tabu tenure. Distance Radar consists of a dual 2D graph in which X-axis represents elite solutions recorded in the search and Y-axis shows distance between current solution with each elite solution. Each graph is used to exhibit distance information from different perspective.

Radar A displays the sorted elite solutions by their objective value in descending order. The recency of the elite solution is denoted by the intensity of its colour (darkest as the most recent). Radar A displays only a visually manageable number of elite solutions (usually a small fraction with respect to the problem size) and any better elite solution found will replace the poorest recorded solution. The effect of Radar A is to approximate the “goodness” of the region where the search is heading towards. Generally, if Radar A shows a curve gradually moving upward, it indicates that the search is diversifying from all the elite solutions. On the other hand, if the curve is moving downwards, it means that the search is intensifying onto the elite solutions.

Radar B displays the sorted elite solutions by their recency in descending order. The quality of the elite solutions is represented by the intensity of their colour (the darkest showing the best objective value). Typically the number of recorded solutions is set to be the same as the tabu tenure. Radar B can be seen as a long-term memory mechanism complementing the tabu list (short term memory). As cycling usually occurs around these solutions (especially local optimal), Radar B detects cycling in them quickly.

In addition, Radar A and B also complement each other to detect possible poor regions. The collaboration of both radars has benefited many important search strategies as they provide an alternative way of enhancing the transparency of the search. Figure 1 shows some incidents observed via one or both radars.

3.2 Selecting the choice of remedial actions

Generally, a good search trajectory has the following characteristics: intensifying the search on good region to yield better solutions and diversifying when the region is depleted of its potential. Hence, it is important to select a *correct* remedial action whenever the search experienced a negative incident (such as solution cycling). V-MDF provides an insight into the

Vienna, Austria, August 22–26, 2005

search space by allowing human to diagnose the encountered incidents. Figure 2 provides an example on how incidents can assist the selection of a remedial action. In this example, three elite solutions (local optimal) are found and recorded as Local Optimal 1, 2 and 3. Suppose from the 3rd local optimal to current solution, the search has experienced a series of non-improving solutions and triggered a remedial action **{cause}**. At this point, the logistician may attempt to improve the search by applying one of his strategies. Under this scenario, the recorded Local Optimal 1, 2 and 3 behave as “signposts” or *anchor points* to our current solution. Let Solution X and Y be the solutions reached after applying his remedial action **{outcome}**.

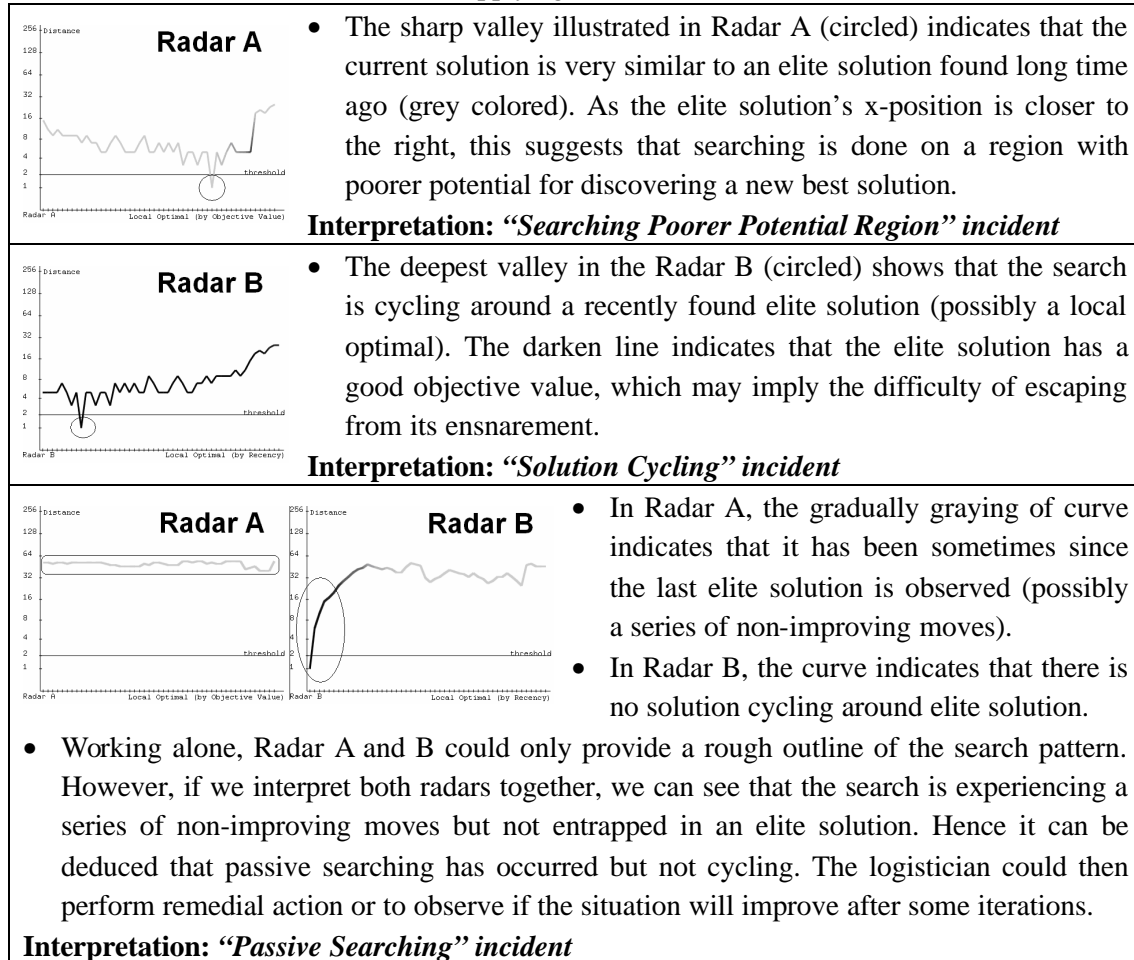


Figure 1: Interpretations from observing the Radars.

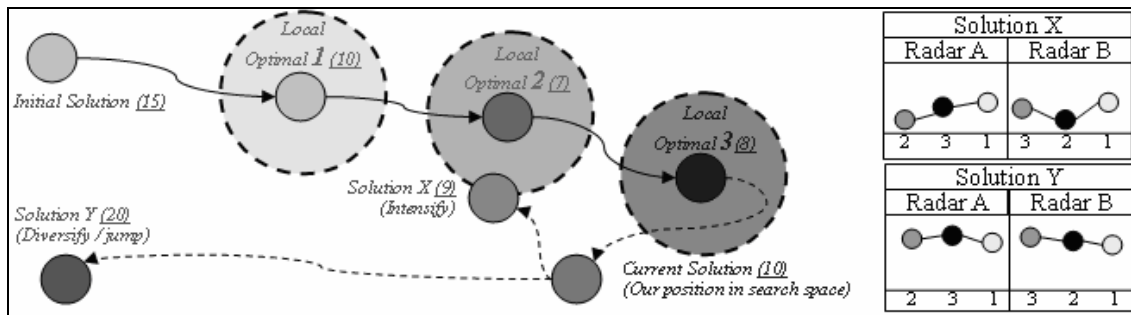


Figure 2: A visualization of the search trajectory of a minimizing problem.

For Solution X, Radar A shows that the search is heading to current best local optimal solution and Radar B shows that the nearest local optimal solution is the one that is 2nd most recently found. Together, Radar A and B has shown the logistician that after his strategy, the search is heading towards good recently found local optimal. Hence if the logistician is performing intensification, he is considered to be “on the right track”; otherwise moving towards Solution X may not be his desired trajectory.

For Solution Y, Radar A and B shows an upward moving horizontal line. This indicates that the current solution is moving away from all known local optimal solutions, which is the “correct” outcome if the logistician is conducting diversification. In short, Solution X and Y have shown the possible outcomes of the logistician actions. If the displayed outcome matches his intended trajectory, his decision **{action}** is deemed to be a correct remedial strategy. This **{cause–action–outcome}** tuple is then stored in the knowledge base. The procedure is repeated for all incidents encountered in the training test cases. When the training is completed, rules are extracted from knowledge base on tuples with highest frequency.

4 Experimental Results

The Military Transport Planning (MTP) problem is an NP-hard optimization problem and can be formally defined as follows: “Given service level q and a set of n requests from military units with the attributes $\{number_of_vehicle_required, start_time, end_time\}$, satisfy at least q out of n requests while minimizing the number of vehicles used.”. (see [7]).

In our experiments, we apply V-MDF on a set of MTP training instances to yield adaptive rules for our tabu search. Due to the page constraint, in this paper, we can only show one training example (see Figure 3) with two possible remedial actions: (1) Adjustment of tabu tenure, and (2) Conducting a diversifying restart by randomly changing some of the allocation. From the two remedial actions we draw the adaptive rules as recorded in Table 1. We then applied adaptive rules on four mock test cases and recorded their effectiveness in Table 2.

5 Conclusions and Future Work

In this paper, we presented a visualizer for tuning adaptive rules and parameters for tabu search. V-MDF is not restricted to a single metaheuristic, but rather to metaheuristics in general.

Vienna, Austria, August 22–26, 2005

By allowing the algorithm designer to visualize search, we allow him to evolve (improve) his metaheuristic. At the moment we are only able to confine to rectification of predefined incidents. We believe that it is possible to extend the human intelligence to develop a more intelligent system that can adaptively react to unknown scenarios.

References

- [1] Adenso-Díaz, Belarmino and Laguna, Manuel (2005): “Fine-tuning of Algorithms using Fractional Experimental Designs and Local Search”. To appear in *Operations Research*.
- [2] Battiti, Roberto and Tecchiolli, Giampietro (1994): “The Reactive Tabu Search”. *ORSA Journal on Computing*, 6(2):126-140.
- [3] Birattari, Mauro (2004): “The Problem of Tuning Metaheuristics: as seen from a machine learning perspective”. PhD thesis. Universite Libre de Bruxelles, Belgium.
- [4] Fonlupt, Cyril, Robilliard, Denis, Preux, Philippe, and Talbi, El-Ghazali (1999): “Fitness Landscapes and Performance of Meta-heuristics”. In *Meta-Heuristics: Advances and Trends in Local Search Paradigms for Optimization*, 18:255-266. Reading. Kluwer Academic Press.
- [5] Jones, Christopher V (1996). “Visualization and Optimization”. Reading. Kluwer Academic Publishers.
- [6] Klau, Gunnar W, Lesh, Neal, Marks, Joe, and Mitzenmacher, Michael (2002): “Human-Guided Tabu Search”. In *Proceedings of AAAI 2002*: 41-47.
- [7] Lau, Hoong Chuin, Ng, Kien Ming, and Wu, Xiaotao (2004): “Transport Logistics Planning with Service-Level Constraints”. In *Proceedings of AAAI 2004*: 519-524.
- [8] Lau, Hoong Chuin, Wan, Wee Chong, Lim, Min Kwang, and Halim, Steven (2004): “A Development Framework for Rapid Meta-Heuristics Hybridization”. In *Proceedings of COMPSAC 2004*: 362-367.
- [9] Merz, Peter and Freisleben, Bernhard (1999): “Fitness Landscapes and Memetic Algorithm Design”. In *New Ideas in Optimization*: 245-260. Reading. McGraw-Hill.

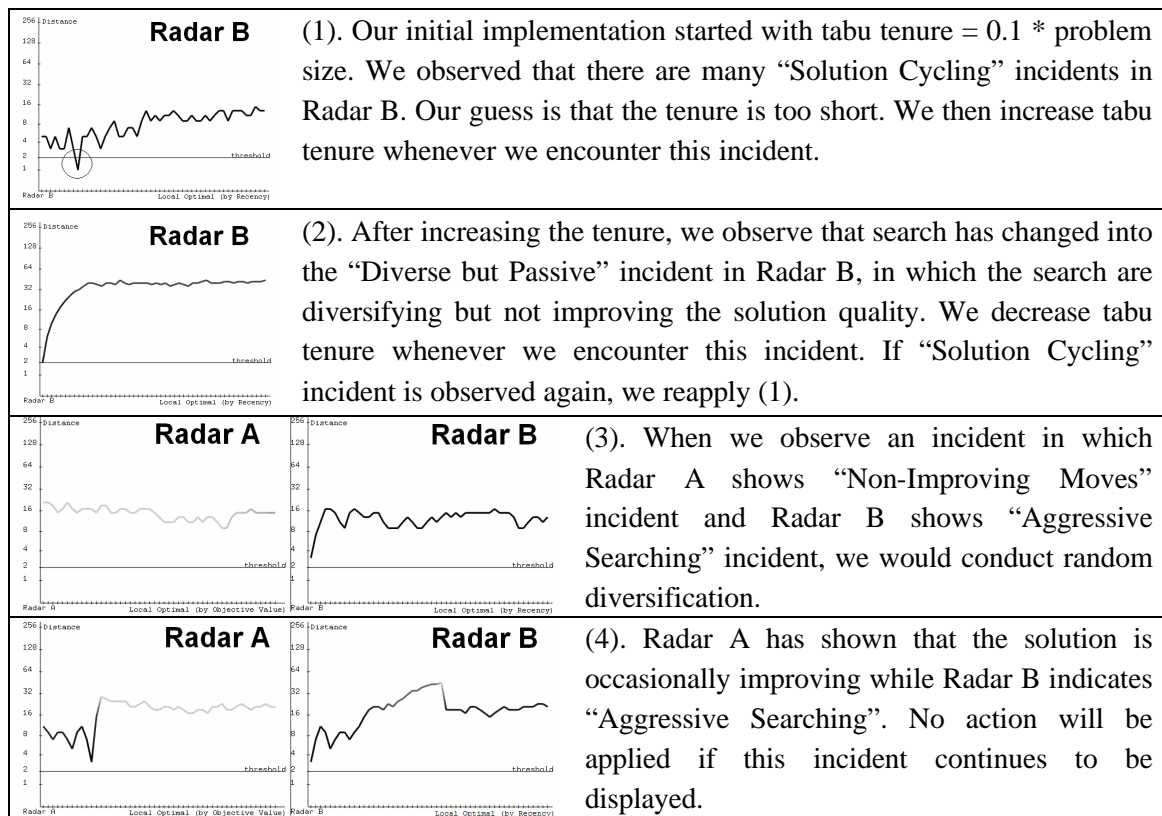


Figure 3: Illustration on training test case.

Table 1: Knowledge base of derived rules.

Cause	Action	Desired Outcome
Solution Cycling	Incr Tabu tenure	No Solution Cycling
Diverse but Passive	Decr Tabu tenure	Aggressive Searching
Non-Improving Moves	Diversify	Good solutions found

Table 2: Table of required vehicles.

MTP test case	TS	TS + Rules
1. n:302, q:250	149	125
2. n:39, q:31	6	2
3. n:283, q:226	60	48
4. n:287, q:258	247	238