

Principles of Tabu Search

FRED GLOVER, MANUEL LAGUNA

Leeds School of Business, University of Colorado, Campus Box 419, Boulder, CO
80309 Fred.Glover@colorado.edu, Laguna@colorado.edu

RAFAEL MARTÍ

Dpto. de Estadística e Investigación Operativa, Facultad de Matemáticas, Universidad
de Valencia, Dr. Moliner 50, 46100 Burjassot (Valencia) Spain. Rafael.Marti@uv.es

Abstract

Tabu Search is a meta-heuristic that guides a local heuristic search procedure to explore the solution space beyond local optimality. One of the main components of Tabu Search is its use of adaptive memory, which creates a more flexible search behavior. Memory-based strategies are therefore the hallmark of tabu search approaches, founded on a quest for “integrating principles,” by which alternative forms of memory are appropriately combined with effective strategies for exploiting them. A novel finding is that such principles are sometimes sufficiently potent to yield effective problem solving behavior in their own right, with negligible reliance on memory. Over a wide range of problem settings, however, strategic use of memory can make dramatic differences in the ability to solve problems. Pure and hybrid Tabu Search approaches have set new records in finding better solutions to problems in production planning and scheduling, resource allocation, network design, routing, financial analysis, telecommunications, portfolio planning, supply chain management, agent-based modeling, business process design, forecasting, machine learning, data mining, biocomputation, molecular design, forest management and resource planning, among many other areas.

Key Words: Intelligent Problem Solving, Memory Structures, Adaptive Memory Programming.

1. Introduction

The term *tabu search* was coined in the same paper that introduced the term *meta-heuristic* (Glover, 1986). Tabu search (TS) is based on the premise that problem solving, in order to qualify as intelligent, must incorporate *adaptive memory* and *responsive exploration*. The adaptive memory feature of TS allows the implementation of procedures that are capable of searching the solution space economically and effectively. Since local choices are guided by information collected during the search, TS contrasts with memoryless designs that heavily rely on semi-random processes that implement a form of sampling. The emphasis on responsive exploration (and hence purpose) in tabu search, whether in a deterministic or probabilistic implementation, derives from the supposition that a bad strategic choice can often yield more information than a good random choice.

TS can be directly applied to virtually any kind of optimization problem. We can state most of these problems in the following form, where “optimize” means to minimize or maximize::

$$\begin{array}{l} \text{Optimize } f(x) \\ \text{subject to} \\ x \in X \end{array}$$

The function $f(x)$ may be linear, nonlinear or even stochastic, and the set \mathbf{X} summarizes constraints on the vector of decision variables x . The constraints may similarly include linear, nonlinear or stochastic inequalities, and may compel all or some components of x to receive discrete values.

While this representation is useful for discussing a number of problem solving considerations, we emphasize that in many applications of combinatorial optimization, the problem of interest may not be easily formulated as an objective function subject to a set of constraints. The requirement $x \in \mathbf{X}$, for example, may specify logical conditions or interconnections that would be cumbersome to formulate mathematically, but may be better left as verbal stipulations that can be then coded as rules.

The TS technique is rapidly becoming the method of choice for designing solution procedures for hard combinatorial optimization problems. A comprehensive examination of this methodology can be found in the book by Glover and Laguna (1997). Widespread successes in practical applications of optimization have spurred a rapid growth of the method as a means of identifying extremely high quality solutions efficiently. TS methods have also been used to create hybrid procedures with other heuristic and algorithmic methods, to provide improved solutions to problems in scheduling, sequencing, resource allocation, investment planning, telecommunications and many other areas. Some of the diversity of tabu search applications is shown in Table 1.

TABLE 1. ILLUSTRATIVE TABU SEARCH APPLICATIONS

<p>Scheduling Flow-Time Cell Manufacturing Heterogeneous Processor Scheduling Workforce Planning Rostering Machine Scheduling Flow Shop Scheduling Job Shop Scheduling Sequencing and Batching</p> <p>Design Computer-Aided Design Fault Tolerant Networks Transport Network Design Architectural Space Planning Diagram Coherency Fixed Charge Network Design Irregular Cutting Problems Lay-Out Planning</p> <p>Logic and Artificial Intelligence Maximum Satisfiability Probabilistic Logic Pattern Recognition/Classification Data Mining Clustering Statistical Discrimination Neural Network Training Neural Network Design</p> <p>Location and Allocation Multicommodity Location/Allocation Quadratic Assignment Quadratic Semi-Assignment Multilevel Generalized Assignment Large-Scale GAP Problems</p> <p>Technology Seismic Inversion Electrical Power Distribution Engineering Structural Design Minimum Volume Ellipsoids Space Station Construction Circuit Cell Placement Off-Shore Oil Exploration</p>	<p>Telecommunications Call Routing Bandwidth Packing Hub Facility Location Path Assignment Network Design for Services Customer Discount Planning Failure Immune Architecture Synchronous Optical Networks</p> <p>Production, Inventory and Investment Supply Chain Management Flexible Manufacturing Just-in-Time Production Capacitated MRP Part Selection Multi-item Inventory Planning Volume Discount Acquisition Project Portfolio Optimization</p> <p>Routing Vehicle Routing Capacitated Routing Time Window Routing Multi-Mode Routing Mixed Fleet Routing Traveling Salesman Traveling Purchaser Convoy Scheduling</p> <p>Graph Optimization Graph Partitioning Graph Coloring Clique Partitioning Maximum Clique Problems Maximum Planner Graphs</p> <p>General Combinational Optimization Zero-One Programming Fixed Charge Optimization Nonconvex Nonlinear Programming All-or-None Networks Bilevel Programming Multi-Objective Discrete Optimization General Mixed Integer Optimization</p>
--	---

The tabu search emphasis on adaptive memory makes it possible to exploit the types of strategies that underlie the best of human problem-solving, instead of being confined to mimicking the processes found in lower orders of natural phenomena and behavior. The basic elements of tabu search have several important features, summarized in Table 2. Tabu search is concerned with finding new and more effective ways of taking advantage of the concepts embodied in Table 2, and with identifying associated principles that can expand the foundations of intelligent search

TABLE 2. PRINCIPAL TABU SEARCH FEATURES
<p>Adaptive Memory</p> <p>Selectivity (including strategic forgetting)</p> <p>Abstraction and decomposition (through explicit and attributive memory)</p> <p>Timing: recency of events frequency of events differentiation between short term and long term</p> <p>Quality and impact: relative attractiveness of alternative choices magnitude of changes in structure or constraining relationships</p> <p>Context: regional interdependence structural interdependence sequential interdependence</p> <p>Responsive Exploration</p> <p>Strategically imposed restraints and inducements (<i>tabu conditions and aspiration levels</i>)</p> <p>Concentrated focus on good regions and good solution features (<i>intensification processes</i>)</p> <p>Characterizing and exploring promising new regions (<i>diversification processes</i>)</p> <p>Non-monotonic search patterns (<i>strategic oscillation</i>)</p> <p>Integrating and extending solutions (<i>path relinking</i>)</p>

In this chapter we will describe some key aspects of this methodology, as the use of memory structures and search strategies, and illustrate them in an implementation to solve the linear ordering problem.

2. Memory Structures

Tabu search begins in the same way as ordinary local or neighborhood search, proceeding iteratively from one point (solution) to another until a chosen termination criterion is satisfied. Each solution x has an associated neighborhood $N(x) \subset X$, and each solution $x' \in N(x)$ is reached from x by an operation called a *move*.

We may contrast TS with a simple descent method where the goal is to minimize $f(x)$. Such a method only permits moves to neighbor solutions that improve the current objective function value and ends when no improving solutions can be found. The final x obtained by a descent method is called a local

optimum, since it is at least as good as or better than all solutions in its neighborhood. The evident shortcoming of a descent method is that such a local optimum in most cases will not be a global optimum, i.e., it usually will not minimize $f(x)$ over all $x \in \mathbf{X}$.

Tabu search permits moves that deteriorate the current objective function value but the moves are chosen from a modified neighborhood $N^*(x)$. Short and long term memory structures are responsible for the specific composition of $N^*(x)$. In other words, the modified neighborhood is the result of maintaining a selective history of the states encountered during the search. In the TS strategies based on short term considerations, $N^*(x)$ characteristically is a subset of $N(x)$, and the tabu classification serves to identify elements of $N(x)$ excluded from $N^*(x)$. In TS strategies that include longer term considerations, $N^*(x)$ may also be expanded to include solutions not ordinarily found in $N(x)$, such as solutions found and evaluated in past search, or identified as high quality neighbors of these past solutions. Characterized in this way, TS may be viewed as a dynamic neighborhood method. This means that the neighborhood of x is not a static set, but rather a set that can change according to the history of the search.

The structure of a neighborhood in tabu search differs from that used in local search in an additional manner, by embracing the types of moves used in constructive and destructive processes (where the foundations for such moves are accordingly called *constructive neighborhoods* and *destructive neighborhoods*). Such expanded uses of the neighborhood concept reinforce a fundamental perspective of TS, which is to define neighborhoods in dynamic ways that can include serial or simultaneous consideration of multiple types of moves.

TS uses attributive memory for guiding purposes (i.e., to compute $N^*(x)$). Instead of recording full solutions, attributive memory structures are based on recording attributes. This type of memory records information about solution properties (attributes) that change in moving from one solution to another. The most common attributive memory approaches are recency-based memory and frequency-based memory. Recency, as its name suggests, keeps track of solutions attributes that have changed during the recent past. Frequency typically consists of ratios about the number of iterations a certain attribute has changed or not (depending whether it is a transition or a residence frequency). Some examples of recency and frequency based memory are shown in Tables 3 and 4 respectively.

Context	Attributes	To record the last time ...
Binary problems	Variable index (i)	variable i changed its value from 0 to 1 or 1 to 0 (depending on its current value).
Job sequencing	Job index (j)	job j changed positions.
	Job index (j) and position (p)	job j occupied position p .
Graphs	Pair of job indexes (i, j)	job i exchange positions with job j .
	Arc index (i)	arc i was added to the current solution. arc i was dropped from the current solution.

Table 3. Examples of recency-based memory

Context	Residence measure	Transition measure
Binary problems	Number of times variable i has been assigned the value of 1.	Number of times variable i has changed values.
Job sequencing	Number of times job j has occupied position p .	Number of times job i has exchanged positions with job j .
	Average objective function value when job j occupies position p .	Number of times job j has been moved to an earlier position in the sequence.
Graphs	Number of times arc i has been part of the current solution.	Number of times arc i has been deleted from the current solution when arc j has been added.
	Average objective function value when arc i is part of the solution.	Number of times arc i has been added during improving moves.

Table 4. Examples of frequency-based memory

Characteristically, a TS process based strictly on short term strategies may allow a solution x to be visited more than once, but it is likely that the corresponding reduced neighborhood $N^*(x)$ will be different each time. With the inclusion of longer term considerations, the likelihood of duplicating a previous

neighborhood upon revisiting a solution, and more generally of making choices that repeatedly visit only a limited subset of X , is all but nonexistent.

Recency-based memory is the most common memory structure used in TS implementations. As its name suggests, this memory structure keeps track of solutions attributes that have changed during the recent past. To exploit this memory, selected attributes that occur in solutions recently visited are labeled *tabu-active*, and solutions that contain tabu-active elements, or particular combinations of these attributes, are those that become tabu. This prevents certain solutions from the recent past from belonging to $N^*(x)$ and hence from being revisited. Other solutions that share such tabu-active attributes are also similarly prevented from being visited. Note that while the tabu classification strictly refers to solutions that are forbidden to be visited, by virtue of containing tabu-active attributes (or more generally by violating certain restriction based on these attributes), moves that lead to such solutions are also often referred to as being tabu.

Frequency-based memory provides a type of information that complements the information provided by recency-based memory, broadening the foundation for selecting preferred moves. Like recency, frequency often is weighted or decomposed into subclasses. Also, frequency can be integrated with recency to provide a composite structure for creating penalties and inducements that modify move evaluations.

Frequencies typically consist of ratios, whose numerators represent counts expressed in two different measures: a *transition measure* — the number of iterations where an attribute changes (enters or leaves) the solutions visited on a particular trajectory, and a *residence measure* — the number of iterations where an attribute belongs to solutions visited on a particular trajectory, or the number of instances where an attribute belongs to solutions from a particular subset. The denominators generally represent one of three types of quantities: (1) the total number of occurrences of all events represented by the numerators (such as the total number of associated iterations), (2) the sum (or average) of the numerators, and (3) the maximum numerator value. In cases where the numerators represent weighted counts, some of which may be negative, denominator (3) is expressed as an absolute value and denominator (2) is expressed as a sum of absolute values (possibly shifted by a small constant to avoid a zero denominator). The ratios produce *transition frequencies* that keep track of how often attributes change, and *residence frequencies* that keep track of how often attributes are members of solutions generated. In addition to referring to such frequencies, thresholds based on the numerators alone can be useful for indicating when phases of greater diversification are appropriate.

3. Search Strategies

The use of recency and frequency memory in tabu search generally fulfills the function of preventing searching processes from *cycling*, i.e., from endlessly executing the same sequence of moves (or more generally, from endlessly and exclusively revisiting the same set of solutions). More broadly, however, the various manifestations of these types of memory are designed to impart additional robustness or vigor to the search.

A key element of the adaptive memory framework of tabu search is to create a balance between search intensification and diversification. Intensification strategies are based on modifying choice rules to encourage move combinations and solution features historically found good. They may also initiate a return to attractive regions to search them more thoroughly. Diversification strategies, on the other hand, seek to incorporate new attributes and attribute combinations that were not included within solutions previously generated. In one form, these strategies undertake to drive the search into regions dissimilar to those already examined. It is important to keep in mind that intensification and diversification are not mutually opposing, but are rather mutually reinforcing.

Most types of intensification strategies require a means for identifying a set of elite solutions as basis for incorporating good attributes into newly created solutions. Membership in the elite set is often determined by setting a threshold that is connected to the objective function value of the best solution found during the search. A simple instance of the intensification strategy is shown in Figure 1. Two simple variants for elite solution selection have proved quite successful. One introduces a diversification measure to assure the solutions recorded differ from each other by a desired degree, and then erases all short term memory before resuming from the best of the recorded solutions. The other keeps a bounded length sequential list that adds a new solution at the end only if it is better than any previously seen, and the short term memory that accompanied this solution is also saved.

```

Apply TS short term memory
Apply an elite selection strategy.
do {
    Choose one of the elite solutions.
    Resume short term memory TS from chosen solution.
    Add new solutions to elite list when applicable.
} while (iterations < limit and list not empty)

```

Figure 1. Simple TS intensification approach.

Diversification is automatically created in TS (to some extent) by short term memory functions, but is particularly reinforced by certain forms of longer term memory. TS diversification strategies are often based on modifying choice rules to bring attributes into the solution that are infrequently used. Alternatively, they may introduce such attributes by periodically applying methods that assemble subsets of these attributes into candidate solutions for continuing the search, or by partially or fully restarting the solution process. Diversification strategies are particularly helpful when better solutions can be reached only by crossing barriers or “humps” in the solution space topology.

The incorporation of modified choice rules can be moderated by using the following penalty function:

$$MoveValue' = MoveValue + d * Penalty.$$

This type of penalty approach is commonly used in TS, where the *Penalty* value is often a function of frequency measures such as those indicated in Table 2, and *d* is an adjustable diversification parameter. Larger *d* values correspond to a desire for more diversification.

4. Advanced Designs: Strategic Oscillation and Path Relinking

There are many forms in which a simple tabu search implementation can be improved by adding long term elements. In this paper we restrict our attention to two of the most used methods, namely strategic oscillation and path relinking, which constitute the core of many adaptive memory programming algorithms.

Strategic oscillation operates by orienting moves in relation to a critical level, as identified by a stage of construction or a chosen interval of functional values. Such a critical level or *oscillation boundary* often represents a point where the method would normally stop. Instead of stopping when this boundary is reached, however, the rules for selecting moves are modified, to permit the region defined by the critical level to be crossed. The approach then proceeds for a specified depth beyond the oscillation boundary, and turns around. The oscillation boundary again is approached and crossed, this time from the opposite direction, and the method proceeds to a new turning point (see Figure 2).




Figure 2. Strategic oscillation

The process of repeatedly approaching and crossing the critical level from different directions creates an oscillatory behavior, which gives the method its name. Control over this behavior is established by generating modified evaluations and rules of movement, depending on the region navigated and the direction of search. The possibility of retracing a prior trajectory is avoided by standard tabu search mechanisms, like those established by the recency-based and frequency-based memory functions.

When the level or functional values in Figure 2 refer to degrees of feasibility and infeasibility, a vector-valued function associated with a set of problem constraints can be used to control the oscillation. In this case, controlling the search by bounding this function can be viewed as manipulating a parameterization of the selected constraint set. A preferred alternative is often to make the function a Lagrangean or surrogate constraint penalty function, avoiding vector-valued functions and allowing tradeoffs between degrees of violation of different component constraints.

Path Relinking, as a strategy of creating trajectories of moves passing through high quality solutions was first proposed in connection with tabu search in Glover (1989). The approach was then elaborated in greater detail as a means of integrating TS intensification and diversification strategies, and given the name *path relinking* (PR), in Glover and Laguna (1993). PR generally operates by starting from an *initiating solution*, selected from a subset of high quality solutions, and generating a path in the neighbourhood space that leads toward the other solutions in the subset, which are called *guiding solutions*. This is accomplished by selecting moves that introduce attributes contained in the guiding solutions.

Path relinking can be considered an extension of the Combination Method of Scatter Search (Glover and Laguna, 1993; Laguna and Martí, 2003). Instead of directly producing a new solution when combining two or more original solutions, PR generates paths between and beyond the selected solutions in the neighborhood space. The character of such paths is easily specified by reference to solution attributes that are added, dropped or otherwise modified by the moves executed. Examples of such attributes include edges and nodes of a graph, sequence positions in a schedule, vectors contained in linear programming basic solutions, and values of variables and functions of variables.

The approach may be viewed as an extreme (highly focused) instance of a strategy that seeks to incorporate attributes of high quality solutions, by creating inducements to favor these attributes in the moves selected. However, instead of using an inducement that merely encourages the inclusion of such attributes, the path relinking approach subordinates other considerations to the goal of choosing moves that introduce the attributes of the guiding solutions, in order to create a “good attribute composition” in the current solution. The composition at each step is determined by choosing the best move, using customary choice criteria, from a restricted set — the set of those moves currently available that incorporate a maximum number (or a maximum weighted value) of the attributes of the guiding solutions. (Exceptions are provided by aspiration criteria, as subsequently noted.) The approach is called path relinking either by virtue of generating a new path between solutions previously linked by a series of moves executed during a search, or by generating a path between solutions previously linked to other solutions but not to each other.

To generate the desired paths, it is only necessary to select moves that perform the following role: upon starting from an *initiating solution*, the moves must progressively introduce attributes contributed by a *guiding solution* (or reduce the distance between attributes of the initiating and guiding solutions). The roles of the initiating and guiding solutions are interchangeable; each solution can also be induced to move simultaneously toward the other as a way of generating combinations. First consider the creation of paths that join two selected solutions x' and x'' , restricting attention to the part of the path that lies ‘between’ the solutions, producing a solution sequence $x' = x(1), x(2), \dots, x(r) = x''$. To reduce the number of options to be considered, the solution $x(i + 1)$ may be created from $x(i)$ at each step by choosing a move that minimizes the number of moves remaining to reach x'' . The relinked path may encounter solutions that may not be better than the initiating or guiding solution, but that provide fertile “points of access” for reaching other, somewhat better, solutions. For this reason it is valuable to examine neighboring solutions along a relinked path, and keep track of those of high quality which may provide a starting point for launching additional searches.

As described in Martí et al. (2004), we can apply different PR elements to perform more elaborated designs. Some examples are: simultaneous relinking, tunneling strategy, extrapolated relinking, multiple guiding solutions, constructive neighborhoods or vocabulary building.

5. The Linear Ordering Problem

Given a matrix of weights $E = \{e_{ij}\}_{m \times m}$, the linear ordering problem (LOP) consists of finding a permutation p of the columns (and rows) in order to maximize the sum of the weights in the upper triangle. In mathematical terms, we seek to maximize:

$$C_E(p) = \sum_{i=1}^{m-1} \sum_{j=i+1}^m e_{p_i p_j} .$$

where p_i is the index of the column (and row) in position i in the permutation. Note that in the LOP, the permutation p provides the ordering of both the columns and the rows. Solution methods for this NP-hard problem have been proposed since 1958, when Chenery and Watanabe outlined some ideas on how to obtain solutions for this problem (Reinelt, 1985). In this section we describe a tabu search implementation (Laguna et al. 1999) for the LOP.

The LOP has a wide range of applications in several fields. Perhaps, the best know application occurs in the filed of economics. In this application, the economy (regional or national) is first subdivided into sectors. Then, an input/output matrix is created, in which the entry (i,j) represents the flow of money from sector i to sector j . Economists are often interested in ordering the sectors so that suppliers tend to come first followed by consumers. This is achieved by permuting the rows and columns of the matrix so that the sum of entries above the diagonal is maximized, which is the objective of the LOP.

Insertions are used as the primary mechanism to move from one solution to another in Laguna's et al method for the LOP. $INSERT_MOVE(p_j, i)$ consist of deleting p_j from its current position j to be inserted in position i (i.e., between the current sectors p_{i-1} and p_i). This operation results in the ordering p' , as follows:

$$p' = \begin{cases} (p_1, \dots, p_{i-1}, p_j, p_i, \dots, p_{j-1}, p_{j+1}, \dots, p_m) & \text{for } i < j \\ (p_1, \dots, p_{j-1}, p_{j+1}, \dots, p_i, p_j, p_{i+1}, \dots, p_m) & \text{for } i > j \end{cases}$$

The neighborhood N consists of all permutations resulting from executing general insertion moves as:

$$N = \{p' : INSERT_MOVE(p_j, i), \text{ for } j = 1, \dots, m \text{ and } i = 1, 2, \dots, j-1, j+1, \dots, m\} ,$$

and N is partitioned into m N^j neighborhoods associated with each sector p_j , for $j = 1, \dots, m$.

$$N^j = \{p' : INSERT_MOVE(p_j, i), i = 1, 2, \dots, j-1, j+1, \dots, m\}$$

Starting from a randomly generated permutation p , the basic TS procedure alternates between an intensification and a diversification phase. An iteration of the *Intensification Phase* begins by randomly selecting a sector. The probability of selecting sector j is proportional to its weight w_j according to:

$$w_j = \sum_{i \neq j} (e_{ij} + e_{ji})$$

The move $INSERT_MOVE(p_j, i) \in N^j$ with the largest move value is selected. (Note that this rule may result in the selection of a non-improving move.) The move is executed even when the move value is not positive, resulting in a deterioration of the current objective function value. The moved sector becomes tabu-active for *TabuTenure* iterations, and therefore it cannot be selected for insertions during this time.

The number of times that sector j has been chosen to be moved is accumulated in the value $freq(j)$. This frequency information is used for diversification purposes. The intensification phase terminates after *MaxInt* consecutive iterations without improvement. Before abandoning this phase, a local search procedure based in the same neighborhood is applied to the best solution found (during the current intensification). We denote this solution as $p^\#$, in contrast to p^* (the best solution found over the entire search). By applying this greedy procedure (without tabu restrictions), a local optimum is guaranteed as the output of the intensification phase.


The *Diversification Phase* is performed for $MaxDiv$ iterations. In each iteration, a sector is randomly selected, where the probability of selecting sector j is inversely proportional to the frequency count $freq(j)$. The chosen sector is placed in the best position, as determined by the move values associated with the insert moves in N^j . The procedure stops when $MaxGlo$ global iterations are performed without improving $C_E(p^*)$. A global iteration is an application of the intensification phase followed by the application of the diversification phase.

An additional intensification is introduced by implementing a long term path relinking phase. Specifically, the best solution found at the end of an intensification phase $p^\#$ (which not necessarily represents p^* , the best solution overall) is subjected to a relinking process. The process consists of making moves starting from $p^\#$ (the initiating solution) in the direction of a set of elite solutions (also referred to as guiding solutions). The set of elite solutions consists of the *EltSol* best solutions found during the entire search. The insertions used to move the initiating solution closer to the guiding solutions can be described as follows. For each sector p_j in the current solution:

- 1) Find the position i for which the absolute value of $(j-i)$ is minimized, where i is the position that p_j occupies in at least one of the guiding solutions.
- 2) Perform $INSERT_MOVE(p_j, i)$.

A long term diversification phase is also implemented to complement the diversification phase in the basic procedure. The long-term diversification is applied after $MaxLong$ global iterations have elapsed without improving $C_E(p^*)$. For each sector p_j , a rounded average position $\alpha(p_j)$ is calculated using the positions occupied by this sector in the set of elite solutions and the solutions visited during the last intensification phase. Then, m diversification steps are performed which insert each sector p_j in its complementary position $m-\alpha(p_j)$, i.e., $INSERT_MOVE(p_j, m-\alpha(p_j))$ is executed for $j = 1, \dots, m$.

After preliminary experimentation, the search parameters are set to $MaxGlo = 100$, $MaxLong = 50$, $EltSol = 4$, $TabuTenure = 2\sqrt{m}$, $MaxInt = m$, and $MaxDiv = 0.5m$ and $EltSol = 4$. In the 49 instances of the public domain LOLIB library, the method obtains the optimal solution within 1 second of computer time run on a Pentium IV at 3Ghz. The method is also compared with a previous procedure due to Chanas and Kobylanski (1996) and a greedy procedure based on the N local search. The methods were run in a way that the best solution found was reported every 0.5 seconds. These data points were used to generate the performance graph in Figure 3. The superior performance of TS_LOP is made evident by Figure 3.



6. The Tabu Cycle and Conditional Probability Methods

In this section, we describe the implementation and testing of the tabu cycle method and two variants of the conditional probability method (Laguna, 2005). These methods were originally described in Glover (1990) and again in the book by Glover and Laguna (1997) but have been largely ignored in the tabu search literature. The *tabu cycle method* is a short-term memory mechanism that is based on partitioning the elements (i.e., move attributes) of a tabu list. The methodology is general and capable of accommodating multi-attribute tabu search memory, as described in Glover and Laguna (1997). In its most basic form, the tabu cycle method divides the short-term memory list into *TabuGroups* groups, where group k consists of elements that were added to the list between a specified range of iterations ago. While in some variants of tabu search (e.g., probabilistic tabu search) it is common to progressively relax the tabu status of elements as they become older, the tabu cycle method, by contrast, allows the elements of some groups to fully escape their tabu status according to certain frequencies that increase with the age of the groups. The method is based on the use of iteration intervals called *tabu cycles*, which are made smaller for older groups than for younger groups (with the exception of a small buffer group). Specifically, if group k has a tabu cycle of $TC(k)$ iterations, then at each occurrence of this many iterations, on average, the elements of group k escape their tabu status and are free to be chosen. In other words, on average, group k is designated as FREE every $TC(k)$ iterations. Mechanisms and data structures that are useful for achieving this are described in Laguna (2005).

The *conditional probability method* is a variant of the tabu cycle method that chooses elements by establishing the probability that a group will be FREE on a given iteration. The probability assigned to group k may be viewed conceptually as the inverse of the tabu cycle value. That is, $P(k) = 1/TC(k)$. Analogous to the tabu cycle method, group k is FREE if all older groups likewise are FREE. The method employs a *conditional probability*, $CP(k)$, as a means of determining whether a particular group k can be designated as FREE. The conditional probability values are fixed and that at each iteration the status of a group is determined by a probabilistic process that is not affected by previous choices. Consequently, the approach ignores the possibility that actual tabu cycle values may be far from their targets for some groups. This may happen, for example, when for a number of iterations no elements are chosen from a particular set of FREE groups. The conditional probability method also makes use of a buffer group, for which no element is allowed to escape its tabu status.

A variant of the conditional probability method uses substitute probability values to keep the expected number of elements per iteration chosen from groups no older than any given group k close to $P(k)$. The substitute probabilities replace the original $P(k)$ values in the determination of the conditional probabilities. These substitute probabilities make use of cycle counts, which are also used in connection with the tabu cycle method.

Laguna (2005) uses a single machine scheduling problem to test the merit of implementations of the tabu cycle method and both variants of the conditional probability method. The problem consists of minimizing the sum of the setup costs and linear delay penalties when n jobs, arriving at time zero, are to be scheduled for sequential processing on a continuously available machine. Several variants of tabu search for this problem have been reported in the literature (Laguna, Barnes and Glover, 1991 and 1993; Glover and Laguna, 1991; Laguna and Glover, 1993). Experiments with more than 300 problem instances with up to 200 jobs were performed to compare a simple static and dynamic short-term memory schemes with a tabu cycle implementation (Cycle), a conditional probability implementation (C-Prob) and an implementation of the conditional probability method with substitute probabilities (S-Prob). The static short-term memory assigns a constant tabu tenure to all attributes during the search. The dynamic short term memory randomly selects a tabu tenure from a specified range. Therefore, the tabu tenure assigned to an attribute in a give iteration may not be the same as the tabu tenure assigned to another attribute in a different iteration. Table 5 shows the number of best solutions found by each method in each set of 100 problems.

Problem Set	Static	Dynamic	Cycle	C-Prob	S-Prob
$n = 50$	2	50	9	31	65
$n = 100$	0	10	28	17	47
$n = 200$	0	8	37	26	29

Table 5. Number of best solutions (out of 100) found by each method

The results in Table 5 show the merit of the tabu cycle and the conditional probability variants as the problem size increases. In addition to these results, the S-Prob is able to find 17 new best solutions to 20 problems used for experimentation in Glover and Laguna (1991). For problems with up to 60 jobs, for which a lower bound can be computed, S-Prob produces a maximum gap of 3.56% in relation to this optimistic bound..

These results confirm that a tabu search procedure based solely on a static tabu list is not a robust method, because it is incapable of maintaining an acceptable level of diversity during the search. The dynamic short-term memory continues to be an appealing alternative, because it is easy to implement and provides a good balance between diversification and intensification. The results also show that improved outcomes are possible with the additional effort required to implement the tabu cycle or conditional probability methods.

Additional strategies identified in Glover and Laguna (1997) can be valuable for exploiting other aspects of intensification and diversification, but this example demonstrates the importance of handling short-term memory in a strategic way, especially when faced with larger and more difficult problems.

Conclusions

The focus and emphasis of tabu search have a number of implications for the goal of designing improved optimization procedures. These research opportunities carry with them an emphasis on producing systematic and strategically designed rules, rather than following the policy of relegating decisions to random choices, as often is fashionable in evolutionary methods. The highly attractive results provided by the adaptive memory structures underlying tabu search are producing an evident impact on the design of metaheuristic methods in general, and are motivating the emergence of new hybrids of TS with other procedures.

Acknowledgments

This research is partially supported by the Spanish Government under codes TIC2002-10886E and TIC2003-C05-01.

References

- Chanas, S. and P. Kobylanski (1996) "A New Heuristic Algorithm Solving the Linear Ordering Problem," *Computational Optimization and Applications*, Vol. 6, pp. 191-205.
- Glover, F. (1986) "Future Paths for Integer Programming and Links to Artificial Intelligence," *Computers and Operations Research*, Vol. 13, pp. 533-549.
- Glover, F., 1989. Tabu Search, Part I, *ORSA Journal on Computing*, vol. 1, no. 3, 190-206.
- Glover, F., 1990. Tabu Search, Part II, *ORSA Journal on Computing*, vol. 2, no. 1, 4-32.
- Glover, F. and M. Laguna (1991) "Target Analysis to Improve a Tabu Search Method for Machine Scheduling," *The Arabian Journal for Science and Engineering*, vol. 16, no. 2B, pp. 239-253.
- Glover, F. and M. Laguna (1993) "Tabu Search," *Modern Heuristic Techniques for Combinatorial Problems*, C. Reeves (ed.), Blackwell Scientific Publishing, Oxford, pp. 70-150.
- Glover, F. and Laguna, M., (1997). *Tabu Search*, Kluwer Academic Publishers, Boston.
- Laguna, M. (2005) "Implementing and Testing the Tabu Cycle and Conditional Probability Methods," <http://leeds-faculty.colorado.edu/laguna/articles/tabucycle.html>
- Laguna, M., J. W. Barnes and F. Glover (1993) "Intelligent Scheduling with Tabu Search: An Application to Jobs with Linear Delay Penalties and Sequence Dependent Setup Costs and Times," *Journal of Applied Intelligence*, vol. 3, pp. 159-172.
- Laguna, M. and F. Glover (1993) "Integrating Target Analysis and Tabu Search for Improved Scheduling Systems," *Expert Systems with Applications*, vol. 6, pp. 287-297.

Laguna, M., Martí, R., 2003. Scatter Search – Methodology and Implementations in C, Kluwer Academic Publishers, Boston.

Laguna, M., Martí, R. and Campos, V. (1999), Intensification and Diversification with Elite Tabu Search Solutions for the Linear Ordering Problem, *Computers and Operations Research*, 26, 1217-1230

LOLIB (1997) <http://www.iwr.uni-heidelberg.de/iwr/comopt/soft/LOLIB/LOLIB.html>.

Martí, R., Laguna, M. and Glover, F. (2004), Principles of Scatter Search, *European Journal of Operational Research*, forthcoming.

Reinelt, G. (1985) *The Linear Ordering Problem: Algorithms and Applications*, Research and Exposition in Mathematics, Vol. 8, H. H. Hofmann and R. Wille (Eds.), Heldermann Verlag Berlin.