

SCATTER SEARCH AND PATH RELINKING

Fred Glover

MediaOne Professor
Graduate School of Business, CB 419
University of Colorado
Boulder, CO 80309

fred.glover@colorado.edu

URL: <http://spot.colorado.edu/~glover>

February 1999

Abstract:—The evolutionary approach called scatter search, and its generalized form called path relinking, originated from strategies for creating composite decision rules and surrogate constraints. Recent studies demonstrate the practical advantages of these approaches for solving a diverse array of optimization problems from both classical and real world settings.

Scatter search and path relinking contrast with other evolutionary procedures, such as genetic algorithms, by providing unifying principles for joining solutions based on generalized path constructions (in both Euclidean and neighborhood spaces) and by utilizing strategic designs where other approaches resort to randomization. Additional advantages are provided by intensification and diversification mechanisms that exploit adaptive memory, drawing on foundations that link scatter search and path relinking to tabu search.

The goal of this paper is to clarify the connection between these developments in evolutionary methods, and to highlight key ideas and research issues that offer promise of yielding future advances.

1. Introduction

Scatter search and path relinking have recently been investigated in a number of studies, disclosing the promise of these methods for solving difficult problems in discrete and nonlinear optimization. Recent applications of these methods (and of selected component strategies within these methods) are shown in Table 1.

Illustrative Applications of Scatter Search and Path Relinking Strategies	
Vehicle Routing	Rochat and Taillard (1995); Taillard (1996), Rego (1999)
Quadratic Assignment	Cung <i>et al.</i> (1996, 1977)
Financial Product Design	Consiglio and Zenios (1996)
Neural Network Training	Kelly, Rangaswamy and Xu (1996)
Job Shop Scheduling	Yamada and Nakano (1996), Jain and Meeran (1998a)
Flow Shop Scheduling	Yamada and Reeves (1998, 1999), Jain and Meeran (1998b)
Crew Scheduling	Laurenço, Paixao and Portugal (1998)
Graph Drawing	Laguna and Marti (1997)
Linear Ordering	Laguna, Marti and Campos (1997, 1999)
Unconstrained Optimization	Fleurent <i>et al.</i> (1996)
Bit Representation	Rana and Whitley (1997)
Multi-objective Assignment	Laguna, Laurenço and Marti (1999)
Optimizing Simulation	Glover, Kelly and Laguna (1996)
Complex Control Systems	Laguna (1997)
Mixed Integer Programming	Glover, Løkketangen and Woodruff (1999)

Table 1

Improved benchmarks for solving a variety of classical problems have resulted from these applications, along with new advances for solving a significant range of commercial problems, particularly those attended by uncertainty and complex nonlinearities.

Scatter search and path relinking derive their foundations from earlier strategies for combining decision rules and constraints, with the goal of enabling a solution procedure based on the combined elements to yield better solutions than one based only on the original elements. An examination of these origins sheds light on the character of these methods.

1.1 Combining Decision Rules

Historically, the antecedent strategies for combining decision rules were introduced in the context of scheduling methods to obtain improved local decision rules for job shop scheduling problems (Glover, 1963). New rules were generated by creating numerically weighted combinations of existing rules, suitably restructured so that their evaluations embodied a common metric.

The approach was motivated by the supposition that information about the relative desirability of alternative choices is captured in different forms by different rules, and that this information can be exploited more effectively when integrated by means of a combination mechanism than when treated by the standard strategy of selecting different rules one at a time, in isolation from each other. In addition, the method departed from the customary approach of stopping upon reaching a local optimum, and instead continued to vary the parameters that determined the combined rules, as a basis for producing additional trial solutions. (This latter strategy also became a fundamental component of tabu search. See, e.g., Glover and Laguna, 1997.)

The decision rules created from such combination strategies produced better empirical outcomes than standard applications of local decision rules, and also proved superior to a “probabilistic learning approach” that selected different rules probabilistically at different junctures, but without the integration effect provided by generating combined rules (Crowston, *et al.*, 1963).

1.2 Combining Constraints

The associated procedures for combining constraints likewise employed a mechanism of generating weighted combinations, in this case applied in the setting of integer and nonlinear programming, by introducing nonnegative weights to create new constraint inequalities, called *surrogate constraints* (Glover, 1965). The approach isolated subsets of constraints that were gauged to be most critical, relative to trial solutions based on the surrogate constraints, and produced new weights that reflected the degree to which the component constraints were satisfied or violated.

A principal function of surrogate constraints, in common with the approaches for combining decision rules, was to provide ways to evaluate choices that could be used to generate and modify trial solutions. From this foundation, a variety of heuristic processes evolved that made use of surrogate constraints and their evaluations. Accordingly, these processes led to the complementary strategy of combining solutions, as a *primal* counterpart to the *dual* strategy of combining constraints, which became manifest in scatter search and its path relinking generalization. (The *primal/dual* distinction stems from the fact that surrogate constraint methods give rise to a mathematical duality theory associated with their role as relaxation methods for optimization. E.g., see Greenberg and Pierskalla, 1970, 1973; Glover, 1965, 1975; Karwan and Rardin, 1976, 1979; Freville and Plateau, 1986, 1993.)

2. Elements of Scatter Search and Path Relinking

2.1 Scatter Search

The scatter search process, building on the principles that underlie the surrogate constraint design, is organized to capture information not contained separately in the original vectors, and to take advantage of auxiliary heuristic methods both for selecting the elements to be combined and for generating new vectors.¹

The original form of scatter search (Glover, 1977) may be sketched as follows.

Scatter Search Procedure

1. Generate a starting set of solution vectors by heuristic processes designed for the problem considered, and designate a subset of the best vectors to be *reference solutions*. (Subsequent iterations of this step, transferring from Step 3 below, incorporate advanced starting solutions and best solutions from previous history as candidates for the reference solutions.)
2. Create new points consisting of linear combinations of subsets of the current reference solutions. The linear combinations are chosen to produce points both inside and outside the convex regions spanned by the reference solutions, modified by generalized rounding processes to yield integer values for integer-constrained vector components.
3. Extract a collection of the best solutions generated in Step 2 to be used as starting points for a new application of the heuristic processes of Step 1. Repeat these steps until reaching a specified iteration limit.

Three particular features of scatter search deserve mention. First, the linear combinations are structured according to the goal of generating weighted centers of selected subregions, allowing for nonconvex combinations that project these centers into regions external to the original reference solutions. (The dispersion

¹ One group of researchers has argued that the coupling of heuristic improvement with solution combination strategies should be given an entirely new name, and accordingly has inaugurated the term *memetic algorithms* to designate such a coupling (see, e.g., Moscato (1989), Radcliffe and Surrey (1994), Burke, Newall and Weare (1996)).

patterns created by such centers and their external projections is particularly useful for mixed integer optimization.) Second, the strategies for selecting particular subsets of solutions to combine in Step 2 are designed to make use of clustering, which allows different types of strategic variation by generating new solutions “within clusters” and “across clusters”. Third, the method is organized to use supporting heuristics that are able to start from infeasible solutions, and hence removes the restriction that solutions selected as starting points for re-applying the heuristic processes must be feasible.

A simple two-dimensional illustration of a type of distribution of points that might be generated by scatter search is shown in Figure 1, by a criterion of selecting weighted centers of subregions. Here, for example, the parent solutions might consist of the three heavily shaded points or the three unshaded points, depending on whether an emphasis is given to generating points inside or outside the region spanned by the reference solutions. Generally speaking, the perspectives of these two types of emphasis are combined, although only a subset of the points produced by patterns such as depicted in Figure 1 are generated, and the offspring are skewed or displaced by the influence of the heuristics applied.² Larger number of reference points can be selected, and can be used to influence the character of the displacement.




Fig. 1 Illustration of a Scatter Search Pattern

In sum, scatter search is founded on the following premises.

- (P1) Useful information about the form (or location) of optimal solutions is typically contained in a suitably diverse collection of elite solutions.
- (P2) When solutions are combined as a strategy for exploiting such information, increased diversity and quality become possible by including combinations that extrapolate beyond the regions spanned by the reference solutions, and further by incorporating heuristic processes to map combined solutions into new points.
- (P3) The opportunity to exploit information contained in the union of elite solutions is enhanced by taking account of multiple reference solutions simultaneously.

The fact that the heuristic processes of scatter search are not restricted to a single uniform design, but represent a varied collection of procedures, affords additional strategic possibilities.

² In some contexts, high quality solutions are found more often near the boundaries of a feasible region than deep in the interior of the region. For problems that technically have no interior due to the presence of equality constraints, the indicated phenomenon nevertheless can occur relative to a subset of constraints that are inequalities, or relative to a space created by identifying and removing variables that may take the role of slack variables for certain equalities, thus transforming them into inequalities. Then linear combinations may reasonably be biased to generate points that lie within a chosen proximity to the feasible boundary.

2.2. Scatter Search and Parametric Search

The approach described in Section 1 for creating combined decision rules, which restructures component rules for compatibility and then joins them by weighted combinations, can be viewed not only as a precursor to scatter search, but also as a model for wider applications of the method. In particular, scatter search strategies can be applied to a variety of problems by the use of approaches that parameterize rules or problem data, and which then generate solutions based on strategic manipulation of the parameters.

A strategy of this type has proved effective in a workforce planning application by Glover and McMillan (1986), where useful gains in solution quality are obtained by replacing the original objective function by one containing a *smoothing criterion*, whose role is gradually diminished until the original objective dominates. The notion of parametrically modifying an objective function to create a smoothing effect has more recently been examined in the traveling salesman setting by Gu and Huang (1994), Schneider et al. (1997) and Coy et al. (1998). These types of approaches can be extended by the use of strategic oscillation, as described later, which periodically phases in and phases out the modified objective.

The application of scatter search with smoothed objectives, and likewise with smoothed constraints, can be formalized by means of the class of approaches called *ghost image (GI) processes* (Glover, 1994c). In the form applicable to our present discussion, the GI approach modifies constraint coefficients, objective function coefficients and/or bounds, to generate a model that initially has an idealized structure which is easy to solve. For example, in covering or matching problems, constraint coefficients can be modified so that they all initially equal 1. In graph partitioning problems, weighted clique problems and traveling salesman problems, the weights (or costs) of edges can be modified so that all initially take the same value, accompanied by adding edges where necessary (penalized in the “true” formulation) so that the graph has an easily identified optimal or near optimal solution. As illustrated in Glover (1994c), these types of starting points can be progressively transformed in a variety of ways to gradually recover the true problem structure, while generating trial solutions for the original problem based on associated solutions to the transformed problems.




Fig. 2 All Costs (or Capacities, Resource Levels, etc.) Above – or Alternately Below – the Moving Threshold are Treated as if Equal the Threshold Value

Figure 2 illustrates how selected types of problem data can be progressively transformed from an initial idealized state where subsets of coefficients are made uniform (hence “smoothed” on chosen dimensions) and gradually made to recover their original unmodified form.

The *moving threshold* in Figure 2 consists of a single value such that all associated data values (costs, resource availabilities, etc.) that lie on a specified side of the threshold are treated "as if" they are the same as the threshold. For example, the moving threshold can begin at the maximum of the data values, with the stipulation that all values smaller than the threshold are treated as equal to the threshold, while all values larger than the threshold are not affected. Therefore, to begin, all values equal the maximum data value. Subsequently, on successive steps the threshold is dropped and an increasing number of data values receive their true values, until finally all do so. The same approach can be applied in reverse, by starting the threshold at the minimum data value and treating all values larger than the threshold as equal to the threshold.

Such a procedure may be applied in conjunction with rules that initially either "spread out" or "contract" the data values, again adopting the ghost image approach to gradually recover the true representation. For example, a special set covering problem described in Glover (1977) proves to be more readily solvable by spreading out the original cost data, in this case by squaring its values. Raising data (and weighted sums of data) to powers and fractional powers has also proved effective in creating data normalizations for surrogate constraint methods in optimization (Løkketangen and Glover, 1997). Similarly, Coy, Golden and Wasil (1998) have shown the ideas of spreading out and contracting data to be useful in smoothing methods for the traveling salesman problem.

The thresholding approach can be applied more generally by reference to paired upper and lower thresholds, as also depicted in Figure 2. In this case, data values above the upper threshold are treated as equal to the upper threshold and those below the lower threshold are treated as equal to the lower threshold, while all other values are unaffected. The paired threshold approach allows the two thresholds to begin equal (in the "mid-range," or in the lower or upper quartile of values, for example) and then to be gradually separated. This approach includes the single moving threshold approach as a special case. To refine the process, modifications of the data based on recency and frequency memory provide a set of strategies that are able to generate changes with additional selectivity.

Within this framework, scatter search can then be applied to the different sets of changed data exactly as it would be applied to the solutions created on the basis of this data. Thus, by such a design, the data sets that have led to the best solutions become the reference points, and selected collections of these points are joined by linear combinations to generate new data sets, which in turn are processed by the associated heuristics to generate new trial solutions.

3. Path Relinking

Features that have been added to scatter search, by extension of its basic philosophy, are captured in the path relinking framework. From a spatial orientation, the process of generating linear combinations of a set of reference solutions may be characterized as generating paths between and beyond these solutions, where solutions on such paths also serve as sources for generating additional paths. This leads to a broader conception of the meaning of creating *combinations* of solutions. By natural extension, such combinations may be conceived to arise by generating paths between and beyond selected solutions in neighborhood space, rather than in Euclidean space (Glover 1989, 1994a; Glover and Laguna, 1993).

This conception is reinforced by the fact that a path between solutions in a neighborhood space will generally yield new solutions that share a significant subset of attributes contained in the parent solutions, in varying "mixes" according to the path selected and the location on the path that determines the solution currently considered. The character of such paths is easily specified by reference to solution attributes that are added, dropped or otherwise modified by the moves executed in neighborhood space. 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.

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.

The incorporation of attributes from elite parents in partially or fully constructed solutions was

foreshadowed by another aspect of scatter search, embodied in an accompanying proposal to assign preferred values to subsets of *consistent* and *strongly determined* variables. The theme is to isolate assignments that frequently or influentially occur in high quality solutions, and then to introduce compatible subsets of these assignments into other solutions that are generated or amended by heuristic procedures. (Such a process implicitly relies on a form of frequency based memory to identify and exploit variables that qualify as consistent.)

Multiparent path generation possibilities emerge in path relinking by considering the combined attributes provided by a set of guiding solutions, where these attributes are weighted to determine which moves are given higher priority. The generation of such paths in neighborhood space characteristically "relinks" previous points in ways not achieved in the previous search history, hence giving the approach its name.

3.1 Initial Steps

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 leaves a reduced number of moves remaining to reach x'' (or more aggressively, a "fewest" number of moves). This policy, even if applied without exception, can permit a significant number of alternative choices for generating the next solution at each step. Consequently, additional criteria are relevant to creating the path, as indicated shortly.

It is possible, as in applying scatter search, that x' and x'' were previously joined by a search trajectory produced by a heuristic method (or by a metaheuristic such as tabu search). In this event, the new trajectory created by path relinking is likely to be somewhat different than the one initially established, representing a 'more direct route' between the solutions. An illustration of this is given in Fig. 3.




Fig. 3 Path relinking. Original path shown by heavy line; relinked path (one possibility) shown by dotted line

It may also be that x' and x'' were not previously joined by a search path at all, but were generated on different search paths, which may have been produced either by a heuristic or by a previous relinking process. Such a situation is depicted in Figure 4.




Fig. 4 Path Relinking: Previously generated paths shown by heavy lines; relinked path shown by dotted line. (Multiple additional points in the space are not shown.)

In this case, the path between x and x' performs a relinking function by changing the connections that generated x and x' originally. The relinking path of this diagram is shown as extending beyond the points x and x' . We discuss this type of construction subsequently under the heading of *extrapolated relinking*.

To choose among the different paths that may be possible in going from x to x' , let $c(x)$ denote an objective function which is to be minimized. Selecting unattractive moves relative to $c(x)$, from the moves that are candidates to generate the path at each step, will tend to produce a final series of strongly improving moves to complete the path. Correspondingly, selecting attractive moves at each step will tend to produce lower quality moves at the end. (The last move, however, will be improving, or leave $c(x)$ unchanged, if x' is selected to be a local optimum.) Thus, choosing best, worst or average moves, provides options that produce contrasting effects in generating the indicated sequence. An aspiration criterion may be used as in tabu search to override choices in the last two cases if a sufficiently attractive solution is available. (In general, it appears reasonable to select best moves at each step, and then to allow the option of reinitiating the process in the opposite direction by interchanging x and x' .)

The choice of one or more solutions $x(i)$ to become reference points for launching a new search phase will preferably be made to depend not only on $c(x(i))$ but also on the values $c(x)$ of those solutions x that can be reached by a move from $x(i)$. The process can additionally be varied to allow solutions to be evaluated other than those that yield $x(i+1)$ closer to x' . Aspiration criteria again are relevant for deciding whether such solutions qualify as candidates for selection.

To elaborate the process, let $x^*(i)$ denote a neighbor of $x(i)$ that yields a minimum $c(x)$ value during an evaluation step, excluding $x^*(i) = x(i+1)$. If the choice rules do not automatically eliminate the possibility $x^*(i) = x(h)$ for $h < i$, then a simple tabu restriction can be used to do this (e.g., see Glover and Laguna, 1997). Then the method selects a solution $x^*(i)$ that yields a minimum value for $c(x^*(i))$ as a new point to launch the search. If only a limited set of neighbors of $x(i)$ are examined to identify $x^*(i)$, then a superior least cost solution $x(i)$, excluding x and x' , may be selected instead. Early termination becomes possible (though is not compulsory) upon encountering an $x^*(i)$ that yields $c(x^*(i)) < \min(c(x'), c(x), c(x(p)))$, where $x(p)$ is the minimum cost $x(h)$ for all $h < i$. The procedure will continue if $x(i)$, in contrast to $x^*(i)$, yields a smaller $c(x)$ value than x and x' , since $x(i)$ effectively adopts the role of x .

3.2 Variation and Tunnelling

A variant of the path relinking approach starts with both endpoints x and x' simultaneously producing two sequences $x = x(1), \dots, x(r)$ and $x' = x(1), \dots, x(s)$. The choices in this case are designed to yield $x(r) = x(s)$, for final values of r and s . To progress toward this outcome when $x(r) \neq x(s)$, either $x(r)$ is selected to create $x(r+1)$, as by the criterion of minimizing reducing the number of moves remaining to reach

$x(s)$, or $x(r)$ is chosen to create $x(s+1)$, as by the criterion of minimizing (reducing) the number of moves remaining to reach $x(r)$. From these options, the move is selected that produces the smallest $c(x)$ value, thus also determining which of r or s is incremented on the next step. Useful variation is also produced by basing the relinking process on more than one neighborhood.

The path relinking approach benefits by a tunnelling strategy that often encourages a different neighborhood structure to be used than in the standard search phase. For example, moves for path relinking may be periodically allowed that normally would be excluded due to creating infeasibility. Such a practice is protected against the possibility of becoming ‘lost’ in an infeasible region, since feasibility evidently must be recovered by the time x is reached. The tunnelling effect therefore offers a chance to reach solutions that might otherwise be bypassed. In the variant that starts from both x and x , at least one of $x(r)$ and $x(s)$ may be kept feasible. An example of tunnelling is shown in Figure 5.




Fig. 5 Feasible region consists of disconnected components. The path “tunnels through” the infeasible region to regain feasibility.

As in tabu search strategies for achieving intensification and diversification, it is appropriate to select the points x and x by reference to clusters of solutions that are created according to criteria of similarity or affinity. Choosing x and x from the same cluster then stimulates intensification, while choosing them from two ‘widely separated’ clusters stimulates diversification. Alternately, parents can be chosen by “anti-clustering,” where each parent is selected to be as far as possible from those previously chosen within the same “family” of parents. A separation criterion such as maximizing the minimum distance to previous points can be used, for example.

3.3 Extrapolated Relinking

The path relinking approach goes beyond consideration of points ‘between’ x and x in the same way that linear combinations extend beyond points that are expressed as convex combinations of two endpoints. In seeking a path that continues beyond x (starting from the point x) we invoke a tabu search concept, referring to sets of attributes associated with the solutions generated, as a basis for choosing a move that ‘approximately’ leaves the fewest moves remaining to reach x . Let $A(x)$ denote the set of solution attributes associated with (‘contained in’) x , and let A_{drop} denote the set of solution attributes that are dropped by moves performed to reach the current solution $x(i)$, starting from x . (Such attributes may be components of the x vectors themselves, or may be related to these components by appropriately defined mappings.)

Define a *to-attribute* of a move to be an attribute of the solution produced by the move, but not an attribute of the solution that initiates the move. Similarly, define a *from-attribute* to be an attribute of the initiating solution but not of the new solution produced. Then we seek a move at each step to maximize the number of *to-attributes* that belong to $A(x'') - A(x(i))$, and subject to this to minimize the number that belong to $A_{drop} - A(x'')$. Such a rule generally can be implemented very efficiently by appropriate data structures.

Once $x(r) = x$ is reached, the process continues by modifying the choice rule as follows. The criterion

now selects a move to maximize the number of its *to-attributes* not in A_{drop} minus the number of its *to-attributes* that are in A_{drop} , and subject to this to minimize the number of its *from-attributes* that belong to $A(x)$. The combination of these criteria establishes an effect analogous to that achieved by the standard algebraic formula for extending a line segment beyond an endpoint. (The secondary minimization criterion is probably less important in this determination.) The path then stops whenever no choice remains that permits the maximization criterion to be positive. The maximization goals of these two criteria are of course approximate, and can be relaxed.


For neighborhoods that allow relatively unrestricted choices of moves, this approach yields a path extending beyond x that introduces new attributes, without reincorporating any old attributes, until no move remains that satisfies this condition. The ability to go beyond the limiting points x and x creates a form of diversification analogous to that provided by the original scatter search approach. At the same time the exterior points are influenced by the trajectory that links x and x .

3.4 Multiple Parents

New points can be generated from multiple parents as follows. Instead of moving from a point x to (or through) a second point x , we replace x by a collection of solutions X . Upon generating a point $x(i)$, the options for determining a next point $x(i+1)$ are given by the union of the solutions in X , or more precisely, by the union A of the attribute sets $A(x)$, for $x \in X$. A takes the role of $A(x)$ in the attribute-based approach previously described, with the added stipulation that each attribute is counted (weighted) in accordance with the number of times it appears in elements $A(x)$ of the collection. Still more generally, we may assign a weight to $A(x)$, which thus translates into a sum of weights over A applicable to each attribute, creating an effect analogous to that of creating a weighted linear combination in Euclidean space. Parallel processing can be applied to operate on an entire collection of points $x \in X$ relative to a second collection $x \in \hat{X}$ by this approach. Further considerations that build on these ideas, and that go beyond the scope of our present development, are detailed in Glover (1994b).

This multiparent path relinking approach generates new elements by a process that emulates the strategies of the original scatter search approach at a higher level of generalization. The reference to neighborhood spaces makes it possible to preserve desirable solution properties (such as complex feasibility conditions in scheduling and routing), without requiring artificial mechanisms to recover these properties in situations where they may otherwise become lost.

Promising regions may be searched more thoroughly in path relinking by modifying the weights attached to attributes of guiding solutions, and by altering the bias associated with solution quality and selected solution features. Figure 6 depicts the type of variation that can result, where the point X represents an initiating solution and the points A , B , and C represent guiding solutions. For appropriate choices of the reference points (and neighborhoods for generating paths from them), principles such as those discussed in Glover and Laguna (1997) suggest that additional elite points are likely to be found in the regions traversed by the paths, upon launching new searches from high quality points on these paths.



*X = Solution selected to generate a relinked path.
(A, B and C may also interchange roles with X,
alternately or simultaneously.)*

Fig. 6 Neighborhood Space Paths with Different Attribute Trade-offs.

3.5 Constructive Neighborhoods and Vocabulary Building

A natural variation of path relinking occurs by using constructive neighborhoods for creating offspring from a collection of parent solutions. In this case the guiding solutions consist of subsets of elite solutions, as before, but the initiating solution begins as a partial (incomplete) solution or even as a null solution, where some of the components of the solutions, such as values for variables, are not yet assigned. The use of a constructive neighborhood permits such an initiating solution to "move toward" the guiding solutions, by a neighborhood path that progressively introduces elements contained in the guiding solutions, or that are evaluated as attractive based on the composition of the guiding solutions.

The evaluations can be conceived as produced by a process where the guiding solutions *vote* for attributes to be included in the initiating solution. It is possible, for example, that a certain partial configuration may be reached where none of the attributes of the guiding solutions can be incorporated within the existing solution, relative to a given constructive neighborhood. Then it is important to still be able to select a next constructive step, by relying upon the voting process for evaluating moves. This same consideration can arise in transition neighborhoods, though it is encountered less frequently there.

Combinations created in this way are called *structured combinations*, and their generation rests upon three properties.

Property 1 (representation property). Each guiding solution represents a vector of votes for particular decisions (e.g., the decision of assigning a specific value to a particular variable).

Property 2 (trial solution property). The votes prescribed by a guiding solution translate into a trial solution to the problem of interest by a well-defined process (determined by the neighborhood structure).

Property 3 (update property). If a decision is made according to the votes of a given vector, a clearly defined rule exists to update all voting vectors for the residual problem so that Properties 1 and 2 continue to hold.

Features of these properties in particular contexts may be clarified as follows.

Elaboration of Property 1: Standard solution vectors for many problems can directly operate as voting vectors, or can be expanded in a natural way to create such vectors. For instance, a solution vector for a job shop scheduling problem can be interpreted as a set of 0-1 votes for predecessor decisions in scheduling specific jobs on particular machines.

Elaboration of Property 2: A set of "yes-no" votes for items to include in a knapsack, for instance, can be translated into a trial solution according to a designated sequence for processing the votes (such as determined by benefit-to-weight ratios), until either the knapsack is full or all votes are considered. More general numerical votes for the same problem may additionally prescribe the sequence to be employed, as where knapsack items are rearranged so the votes occur in descending order. (The voting vectors are not required to represent feasible solutions to the problems considered, or even represent solutions in a customary sense at all. Thus, for example, the scheme can also operate to combine decision rules as in the approach for doing this described in section 1.1.)

Elaboration of Property 3. Upon assigning a specific value to a particular variable, all votes for assigning different values to this variable effectively become cancelled. Property 3 then implies that the remaining updated votes of each vector retain the ability to be translated into a trial solution for the residual problem in which the assignment has been made.

Concrete illustrations of processes for generating structured combinations by reference to these properties are provided in Glover (1994b). These same kinds of processes can be implemented by reference to destructive neighborhoods – that is, neighborhoods that allow the removal of less attractive elements. Typically, destructive processes are applied to solutions that begin with an “excessive assignment” (such as too many elements to satisfy cardinality or capacity restrictions).

3.6 Vocabulary Building

Vocabulary building creates structured combinations not only by using the primitive elements of customary neighborhoods, but also building and joining more complex assemblies of such elements. The process receives its name by analogy with the process of building words progressively into useful phrases, sentences and paragraphs, where valuable constructions at each level can be visualized as represented by “higher order words,” just as natural languages generate new words to take the place of collections of words that embody useful concepts.

The motive underlying vocabulary building is to take advantage of those contexts where certain partial configurations of solutions often occur as components of good complete solutions. A strategy of seeking “good partial configurations” – good vocabulary elements – can help to circumvent the combinatorial explosion that potentially results by manipulating only the most primitive elements by themselves. The process also avoids the need to reinvent (or rediscover) the structure of a partial configuration as a basis for building a good complete solution. (The same principle operates in mathematical analysis generally, where basic premises are organized to produce useful lemmas, which in turn facilitate the generation of more complex theorems.)

Vocabulary building has an additional useful feature in some problem settings by providing compound elements linked by special neighborhoods that are more exploitable than the neighborhoods which operate on the primitive elements. For example, a vocabulary building proposal of Glover (1992) discloses that certain subassemblies (partial "tours") for traveling salesman problems can be linked by exact algorithms to produce optimal unions of these components. Variants of this strategy have more recently been introduced by Aggarwal, Orlin and Tai (1997) as a proposal for modifying traditional genetic algorithms, and have also been applied to weighted clique problems by Balas and Niehaus (1998). A particularly interesting application occurs in the work of Lourenço, Paixao and Portugal (1998), who use such concepts to create “perfect children” for crew scheduling problems.

In general, vocabulary building relies on destructive as well as constructive processes to generate desirable partial solutions, as in the early proposals for exploiting strongly determined and consistent variables – which essentially “break apart” good solutions to extract good component assignments, and then subject these assignments to heuristics to rebuild them into complete solutions. Construction and destruction therefore operate hand in hand in these approaches. An illustration of vocabulary building is depicted in Figure 7.

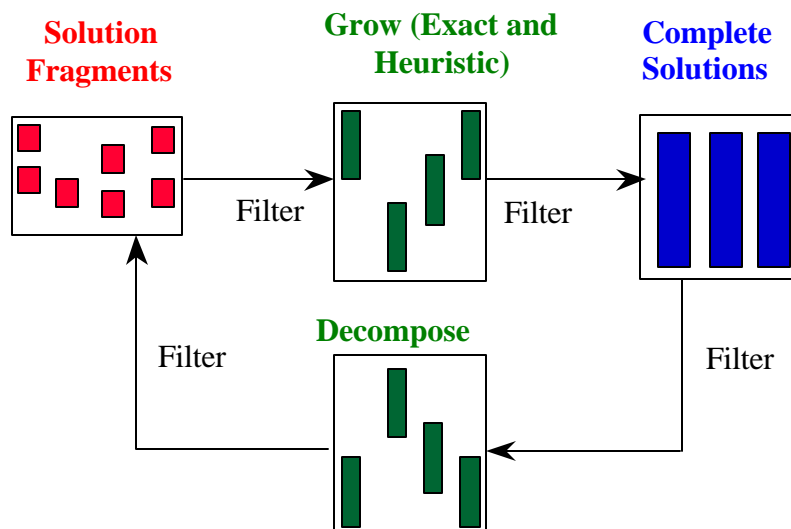


Fig. 7 Vocabulary Building Process

4. Implications for Future Developments

The focus and emphasis of the scatter search and path relinking approaches have a number of specific implications for the goal of designing improved optimization procedures. To understand these implications, it is useful to consider certain contrasts between the highly exploitable meaning of “solution combination” provided by path relinking and the rather amorphous concept of “crossover” used in genetic algorithms. Originally, GAs were founded on precise notions of crossover, using definitions based on binary strings and motivated by analogies with genetics. Although there are still many GA researchers who favor the types of crossover models originally proposed with genetic algorithms – since these give rise to the theorems that have helped to popularize GAs – there are also many who have largely abandoned these ideas and who have sought, on a case-by-case basis, to replace them with something different. The well-defined earlier notions of crossover have not been abandoned without a price. The literature is rife with examples where a new problem (or a new variant of an old one) has compelled the search for an appropriate “crossover” to begin anew.³

As a result of this lack of an organizing principle, many less-than-suitable modes of combination have been produced, some eventually replacing others, without a clear basis for taking advantage of context – in contrast to the strong context-exploiting emphasis embodied in the concept of search neighborhoods. The difficulty of devising a unifying basis for understanding or exploiting context in GAs was inherited from its original theme, which had the goal of making GAs *context free*.

A few of the more conspicuous features of “genetic crossover” and path relinking that embody such contrasts appear in the following table.

³ The disadvantage of lacking a clear and unified model for combining solutions has had its compensations for academic researchers, since each new application creates an opportunity to publish another form of crossover! The resulting abundance of papers has done nothing to tarnish the image of a dynamic and prospering field.

“Genetic Crossover” Features	Contrasting Path Relinking Features
<ul style="list-style-type: none"> • Contains No Integrated Framework • Each New “Crossover” is Separate, with No Guidance for the Next • No Basis Exists to Systematically Exploit Context • Advances are Piecemeal, Without Clear Sources of Potential for Transfer • There is No Design Plan that Is Subject to Analysis or Improvement 	<ul style="list-style-type: none"> • Embodies a Unifying “Path Combination” Principle • Each Implementation of Path Relinking Derives From a Common Foundation • Context Inheres in Neighborhood Structures and is Directly Exploitable by them • Advances in Neighborhood Search Foster Advances in Path Relinking (and Reciprocally) • A Cohesive Framework exists for Developing Progressively Improved Methods

Table 2

The differences identified in Table 2 have important consequences for research to yield improved methods. Specific areas of research for developing improved solution strategies that emerge directly from the path relinking orientation are catalogued in Table 3.

Research Areas Providing Opportunities for Improved Methods
<ul style="list-style-type: none"> • Connections and Complementarities Between Neighborhoods for Search Methods and Neighborhoods for Path Relinking • Rules for Generating Paths to Different Depths and Thresholds of Quality • Strategies for Generating Multiple Paths Between and Beyond Reference Solutions (with Parallel Processing Applications) • Path Interpolations and Extrapolations that are Effective for Intensification and Diversification Goals • Strategies for Clustering and Anti-Clustering, to Generate Candidate Sets of Solutions to be Combined • Rules for Multi-Parent Compositions • Isolating and Assembling Solution Components by Means of Constructive Linking and Vocabulary Building

Table 3

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 strategic orientation underlying path relinking is motivated by connections with the tabu search setting where the path relinking ideas were first proposed, and invites the use of adaptive memory structures in determining the strategies produced. The learning approach called target analysis (Glover and Laguna, 1997) gives a particularly useful basis for pursuing such research.

5. Intensification and Diversification

A significant feature that distinguishes scatter search and path relinking from other evolutionary approaches is the fact that intensification and diversification processes are not conceived to be embedded solely within the mechanisms for combining solutions, or within supplementary "mutation" strategies based on randomly varying offspring to produce new solutions.⁴

Evidently, except where hybrids are being created with tabu search, alternative evolutionary computation approaches do not undertake to control search paths by adaptive memory strategies such as those based on measures of recency, frequency and influence.

The initial connections between scatter search and strategies involving these types of measures have already been noted in reference to exploiting *consistent* and *strongly determined* variables. Such strategies naturally fall in the category of intensification strategies, in the sense that they undertake to take advantage of features associated with good solutions. They are predicated on highly explicit analysis of the frequencies that attributes belong to high quality solutions, supplemented by considerations such as clustering the solutions to give increased meaning to the frequencies. This stands in notable contrast to the philosophy of other mainstream evolutionary procedures, where the relevance of attribute membership in solutions is left to be "discovered" chiefly by the device of shuffling and combining solutions.

An approach called *strategic oscillation* introduced with the original scatter search proposal is important for linking intensification and diversification. The basic idea of this approach is to identify critical regions of search, and to induce the search pattern to visit these regions to various depths within their boundaries, by a variable pattern that approaches and retreats from the boundaries in oscillating waves. Examples of such regions and their associated boundaries are indicated in the following table.

Regions	Boundary
Feasible and Infeasible	Determined by Constraints
Partial Constructions and (sometimes) "Excessive" Constructions	A complete Construction (tree, clique, tour, etc.)
Underfilled or Overfilled Schedules	Satisfied Schedules (all or an appropriate set of jobs assigned)
Local Optima and Suboptima	Solutions with no Immediate Improvement
Elite Solution Clusters (or Partitioned Spaces)	Zones between Clusters (or between Partitions)
Alternative Neighborhoods	Transitions among Move Types

Table 4: Applications of Strategic Oscillation

A number of variations of the regions and associated boundaries are evident, such as replacing the

⁴ Within the last few years, some researchers in the evolutionary computation field have begun to adopt aspects of scatter search and path relinking by incorporating systematic strategies for achieving intensification and diversification, instead of relying on randomization to achieve less purposeful forms of variation. However, some of the latest literature still disallows this type of approach as a legitimate feature of evolutionary computation. For example, Fogel (1998) says that the main disciplines of evolutionary computation all involve a process whereby "New solutions are created by randomly varying the existing solutions."

feasible/infeasible dichotomy by a focus on selected critical constraints to define varying domains of "partial infeasibility." In each case, the strategic oscillation approach operates by moving through one region to approach the boundary, and then either crosses the boundary or reverses direction to move back into the region just traversed (in the case of a "one-sided" oscillation).

Often there are advantages to crossing boundaries to descend for varying depths within different regions and then doubling back to return again to the boundary. For example, in a number of discrete optimization problems the solutions that are most readily accessible from the feasible and infeasible regions differ – and, quite significantly the types of moves and choice criteria for traversing feasible and infeasible regions also differ. The ability to exploit these differences by rules that are specific to the regions traversed and the direction of movement within these regions (e.g., toward or away from their boundaries) provides an enriched set of options for carrying out the search. Similar characteristics are found in processes that build solutions by constructive processes and then dismantle them by destructive processes. In many settings, classical heuristics have been restricted to constructive processes for generating solutions and in these cases strategic oscillation entails the creation of additional destructive moves to complement the constructive moves.

From the perspective of intensification and diversification, greater diversification is normally achieved by penetrating to greater depths beyond regional boundaries, while greater intensification is normally achieved by spending more time in the vicinity of such boundaries. However, the spatial image of remaining close to a boundary is misleading, because oscillations of small depths can create significant changes. For example, even when only a few destructive moves are made to reverse a constructive process, the portions of the construction dismantled can have a significant influence on the solution structure and composition, and this influence can become magnified after a few oscillation cycles. The guidance of memory as used in tabu search allows the oscillations to avoid becoming mired in local optima or in a process that unproductively examines similar points in a common locale.

An extreme application of strategic oscillation in the context of a constructive approach is to employ destructive steps to completely dismantle the solution built by the constructive phase, which thus simply reduces the approach to a "restart" procedure. However, the use of memory to guide the successively restarted constructions produces significantly different outcomes than those produced by customary restarting procedures based on randomization. These outcomes also contrast sharply with those produced by "randomized greedy" construction schemes. The comparative advantages of memory-based strategies for rebuilding solutions are documented, for example, in Fleurent and Glover (1999).

Intensification is often associated with shallow oscillations because in many settings the best solutions are found on or near the boundary. This is clearly true for multidimensional knapsack and covering problems, for instance, and it is also true by definition where the boundary is established to segregate local optima from suboptimal solutions. In such cases the oscillation process is augmented by spending additional time in the proximity of the boundary, as by shifting from a simple neighborhood to a more complex neighborhood. For instance, simple "add/drop" or "increment/decrement" moves may be augmented by a series of "exchange" or "paired increment/decrement" moves upon reaching (or drawing close to) the boundary. Candidate list strategies are important when complex neighborhoods are used, in order to achieve proper tradeoffs between time spent looking for moves and the quality of the moves found. (Principal approaches of this type are described in Glover and Laguna, 1997.)

5.1 Randomization and the Intensification/Diversification Dichotomy

The emphasis on systematic strategies in achieving intensification and diversification does not preclude the use of randomized selection schemes, which are often motivated by the fact that they require little thought or sophistication to apply. By the same token, deterministic rules that are constructed with no more reflection than devoted to creating a simple randomized rule can be quite risky, because they can easily embody oversights that will cause them to perform poorly. A randomized rule can then offer a safety net, by preventing a bad decision from being applied persistently and without exception.

Yet a somewhat different perspective suggests that deterministic rules can offer important advantages in the longer run. A "foolish mistake" incorporated into a deterministic rule becomes highly visible by its consequences, whereas such a mistake in a randomized rule may be buried from view – obscured by the patternless fluctuations that surround it. Deterministic rules afford the opportunity to profit by mistakes and

learn to do better. The character of randomized rules, that provides the chance to escape from repetitive folly, also inhibits the chance to identify more effective decisions.

The concepts of intensification and diversification are predicated on the view that intelligent variation and randomized variation are rarely the same.⁵ This clearly contrasts with the prevailing perspective in the literature of evolutionary methods although, perhaps surprisingly, the intensification and diversification terminology has been appearing with steadily increasing frequency in this literature. Nevertheless, a number of the fundamental strategies for achieving the goals of intensification and diversification in scatter search and path relinking applications have still escaped the purview of other evolutionary methods.

Perhaps one of the factors that is slowing a more complete assimilation of these ideas is a confusion between the terminology of intensification and diversification and the terminology of “exploitation versus exploration” popularized in association with genetic algorithms. The exploitation/exploration distinction comes from control theory, where exploitation refers to following a particular recipe (traditionally memoryless) until it fails to be effective, and exploration then refers to instituting a series of random changes – typically via multi-armed bandit schemes – before reverting to the tactical recipe. The issue of exploitation versus exploration concerns how often and under what circumstances the randomized departures are launched.

By contrast, intensification and diversification are mutually reinforcing (rather than being mutually opposed), and can be implemented in conjunction as well as in alternation. In longer term strategies, intensification and diversification are both activated when simpler tactics lose their effectiveness. Characteristically, they are designed to profit from memory, rather than to rely solely on indirect “inheritance effects.”

6. Conclusion

It is not possible within the limited scope of this paper to detail completely the aspects of scatter search and its path relinking generalization that warrant further investigation. Additional implementation considerations, including associated intensification and diversification processes, and the design of accompanying methods to improve solutions produced by combination strategies, may be found in the *template* for scatter search and path relinking in Glover (1997).

However, a key observation deserves to be stressed. The literature often contrasts evolutionary methods – especially those based on combining solutions – with local search methods, as though these two types of approaches are fundamentally different. In addition, evolutionary procedures are conceived to be independent of any reliance on memory, except in the very limited sense where solutions forged from combinations of others carry the imprint of their parents. Yet as previously noted, the foundations of scatter search strongly overlap with those of tabu search and, in addition, path relinking was initiated as a strategy to be applied with the guidance of adaptive memory processes. By means of these connections, a wide range of strategic possibilities exist for implementing scatter search and path relinking.

Very little computational investigation of these methods has been done by comparison to other evolutionary methods, and a great deal remains to be learned about the most effective implementations for various classes of problems. The highly promising outcomes of studies such as those cited in Section 1 suggest that these approaches may offer a useful potential for applications in areas beyond those investigated up to now.

⁵ Intelligence can sometimes mean quickly doing something mildly clever, rather than slowly doing something profound. This can occur where the quality of a single move obtained by extended analysis is not enough to match the quality of multiple moves obtained by more superficial analysis. Randomized moves, which are quick, sometimes gain a reputation for effectiveness because of this phenomenon. In such setting, a different perspective may result by investigating comparably fast mechanisms that replace randomization with intelligent variation.

REFERENCES

1. Aggarwal, C.C., J.B. Orlin and R.P. Tai (1997). "Optimized Crossover for the Independent Set Problem," *Operations Research* 45, 226-234.
2. Balas, E., and W. Niehaus (1998). "Optimized Crossover-Based Genetic Algorithms for the Maximum Cardinality and Maximum Weight Clique Problems," *Journal of Heuristics*, Vol 4:2. 107-124.
3. Burke, E.K., J.P. Newall and R.F. Weare (1996). "A Memetic Algorithm for University Timetabling," in *The Practice and Theory of Automated Timetabling*, (E.K. Burke and P. Ross eds.). Springer-Verlag Lecture Notes in Computer Science. 1153.
4. Consiglio, A. and S.A. Zenios (1996). "Designing Portfolios of Financial Products via Integrated Simulation and Optimization Models," Report 96-05, Department of Public and Business Administration, University of Cyprus, Nicosia, CYPRUS, to appear in *Operations Research*.
5. Consiglio, A. and S.A. Zenios (1997). "A Model for Designing Callable Bonds and its Solution Using Tabu Search," *Journal of Economic Dynamics and Control* 21, 1445-1470.
6. Coy, S., B. Golden, G. Runger, and E. Wasil (1998). "See the Forest Before the Trees: Finetuned Learning and its Application to the Traveling Salesman Problem," *IEEE Transactions on Systems, Man, and Cybernetics*, 28(4), 454-464.
7. Coy, S. P., B. L. Golden and E. A. Wasil (1998). "A computational Study of Smoothing Heuristics for the Traveling Salesman Problem," Research report, University of Maryland, to appear in *European Journal of Operational Research*.
8. Crowston, W.B., F. Glover, G.L.Thompson and J.D. Trawick (1963). "Probabilistic and Parametric Learning Combinations of Local Job Shop Scheduling Rules," ONR Research Memorandum No. 117, GSIA, Carnegie Mellon University, Pittsburgh, PA.
9. Cung, V-D., T. Mautor, P. Michelon, A. Tavares (1996). "Scatter Search for the Quadratic Assignment Problem," Laboratoire PRISM-CNRS URA 1525.
10. Fleurent, C., F. Glover, P. Michelon and Z. Valli (1996). "A Scatter Search Approach for Unconstrained Continuous Optimization," *Proceedings of the 1996 IEEE International Conference on Evolutionary Computation*, 643-648.
11. Fleurent, C. and F. Glover (1999). "Improved Constructive Multistart Strategies for the Quadratic Assignment Problem Using Adaptive Memory," *INFORMS Journal on Computing*, Vol. 11, No. 2, 198-204.
12. Fogel, David B. (1998). "An Introduction to Evolutionary Computation," Chapter 1 of *Evolutionary Computation: The Fossil Record*, D.B. Fogel, ed., IEEE Press, 1 - 2.
13. Freville, A. and G. Plateau (1986). "Heuristics and Reduction Methods for Multiple Constraint 0-1 Linear Programming Problems," *European Journal of Operational Research*, 24, 206-215.
14. Freville, A. and G. Plateau (1993). "An Exact Search for the Solution of the Surrogate Dual of the 0-1 Bidimensional Knapsack Problem," *European Journal of Operational Research*, 68, 413-421.
15. Glover, F. (1963). "Parametric Combinations of Local Job Shop Rules," Chapter IV, ONR Research Memorandum no. 117, GSIA, Carnegie Mellon University, Pittsburgh, PA.
16. Glover, F. (1965). "A Multiphase Dual Algorithm for the Zero-One Integer Programming Problem," *Operations Research*, Vol 13, No 6, 879.
17. Glover, F. (1975). "Surrogate Constraint Duality in Mathematical Programming," *Operations Research*, 23, 434-451.
18. Glover, F. (1977). "Heuristics for Integer Programming Using Surrogate Constraints," *Decision Sciences*, Vol 8, No 1, 156-166.
19. Glover, F. (1992). "Ejection Chains, Reference Structures and Alternating Path Methods for Traveling Salesman Problems," University of Colorado. Shortened version published in *Discrete Applied Mathematics*, 1996, 65, 223-253.
20. Glover, F. (1994a). "Genetic Algorithms and Scatter Search: Unsuspected Potentials," *Statistics and Computing*, 4, 131-140.
21. Glover, F. (1994b). "Tabu Search for Nonlinear and Parametric Optimization (with links to genetic algorithms)," *Discrete Applied Mathematics*, 49, 231-255.
22. Glover, F. (1994c). "Optimization by Ghost Image Processes in Neural Networks," *Computers and Operations Research*, Vol 21, No. 8, 801-822.
23. Glover, F. (1997). "A Template for Scatter Search and Path Relinking," in *Lecture Notes in Computer Science*,

- 1363, J.K. Hao, E. Lutton, E. Ronald, M. Schoenauer, D. Snyers (Eds.), pp. 13-54. (Expanded version available on request.)
24. Glover, F., J. P. Kelly and M. Laguna (1996). "New Advances and Applications of Combining Simulation and Optimization," *Proceedings of the 1996 Winter Simulation Conference*, J. M. Charnes, D. J. Morrice, D. T. Brunner, and J. J. Swain (Eds.), 144-152.
 25. Glover, F. and M. Laguna (1997). *Tabu Search*, Kluwer Academic Publishers.
 26. Glover, F., A. Løkketangen and D. Woodruff (1999). "Scatter Search to Generate Diverse MIP Solutions," Research Report, University of Colorado, Boulder.
 27. Glover, F. and C. McMillan (1986). "The General Employee Scheduling Problem: An Integration of MS and AI," *Computers and Operations Research*, Vol 13, No. 5, 563-573.
 28. Greenberg, H. J. and W. P. Pierskalla (1970). "Surrogate Mathematical Programs," *Operations Research*, 18, 924-939.
 29. Greenberg, H. J. and W. P. Pierskalla (1973). "Quasi-conjugate Functions and Surrogate Duality," *Cahiers du Centre d'Etudes de Recherche Operationelle*, 15, 437-448.
 30. Gu, J. and X. Huang (1994). "Efficient Local Search With Search Space Smoothing: A Case Study of the Traveling Salesman Problem," *IEEE Transactions on Systems, Man, and Cybernetics*, 24(5) 728-735.
 31. Holland, J.H. (1975). *Adaptation in Natural and Artificial Systems*. University of Michigan Press, Ann Arbor, MI.
 32. Jain A.S. and S. Meeran (1998a). "An Improved Search Template for Job-Shop Scheduling," INFORMS Spring Meeting, Montreal, Quebec, Canada, April 26-29, Department of Applied Physics, Electronic and Mechanical Engineering, University of Dundee, Dundee, Scotland Department of Applied Physics, Electronic and Mechanical Engineering, University of Dundee, Dundee, Scotland.
 33. Jain A.S. and S. Meeran (1998b). "A Multi-Level Hybrid Framework for the General Flow-Shop Scheduling Problem," Department of Applied Physics, Electronic and Mechanical Engineering, University of Dundee, Dundee, Scotland.
 34. Karwan, M.H. and R.L. Rardin (1976). "Surrogate Dual Multiplier Search Procedures in Integer Programming," School of Industrial Systems Engineering, Report Series No. J-77-13, Georgia Institute of Technology.
 35. Karwan, M.H. and R.L. Rardin (1979). "Some Relationships Between Lagrangean and Surrogate Duality in Integer Programming," *Mathematical Programming*, 17, 230-334.
 36. Kelly, J., B. Rangaswamy and J. Xu (1996). "A Scatter Search-Based Learning Algorithm for Neural Network Training," *Journal of Heuristics*, Vol. 2, pp. 129-146.
 37. Laguna, M. (1997). "Optimizing Complex Systems with OptQuest," Research Report, University of Colorado.
 38. Laguna, M. and R. Martí (1997). "GRASP and Path Relinking for 2-Layer Straight Line Crossing Minimization," Research Report, University of Colorado.
 39. Laguna, M., R. Martí and V. Campos (1997). "Tabu Search with Path Relinking for the Linear Ordering Problem," Research Report, University of Colorado.
 40. Laguna, M., H. Lourenço and R. Martí (1999). "Multi-objective scatter search for efficient assignments of proctors to examinations," Working Paper, University of Colorado.
 41. Løkketangen, A. and F. Glover (1997). "Surrogate Constraint Analysis – New Heuristics and Learning Schemes for Satisfiability Problems," DIMACS Series in *Discrete Mathematics and Theoretical Computer Science*, Vol. 35, 537-572.
 42. Lourenço, H., J. Paixao and R. Portugal, (1998). "Metaheuristics for the bus-driver scheduling problem," Economic Working Papers Series, no. 304, Universitat Pompeu Fabra, Spain.
 43. Moscato, P. (1989), "On Evolution, Search, Optimization Algorithms and Martial Arts: Towards Memetic Algorithms," Report 826, Caltech Concurrent Computation Program, California Institute of Technology, Pasadena.
 44. Radcliffe, N., and P.D. Surrey (1994), "Formal Memetic Algorithms", *AISB94*.
 45. Rana, S. and D. Whitley (1997). "Bit Representations with a Twist," Proc. 7th International Conference on Genetic Algorithms, T. Baeck ed., pp. 188-196, Morgan Kaufman.

46. Reeves, C.R. (1997). "Genetic Algorithms for the Operations Researcher," *Journal on Computing*, Vol. 9, No. 3, 231-250 (with commentaries and rejoinder).
47. Rego, C. (1999). "Integrating Advanced Principles of Tabu Search for the Vehicle Routing Problem," Working Paper, Faculty of Sciences, University of Lisbon.
48. Rochat, Y. and É. D. Taillard (1995). "Probabilistic diversification and intensification in local search for vehicle routing," *Journal of Heuristics* 1, 147-167.
49. Schneider, J., M. Dankersreiter, W. Fettes, I. Morgenstern, M. Schmid, and J. Singer (1997). "Search-space Smoothing for Combinatorial Optimization Problems," *Physica A Statistical and Theoretical Physics*, 243(1-2) 77-112.
50. Taillard, É. D. (1996). "A heuristic column generation method for the heterogeneous VRP," Publication CRT-96-03, Centre de recherche sur les transports, Université de Montréal. To appear in *RAIRO-OR*.
51. Ulder, N.L.J., E. Pesch, P.J.M. van Laarhoven, H.J. Bandelt and E.H.L. Aarts (1991). "Genetic local search algorithm for the traveling salesman problem," *Parallel Problem solving from Nature* (Edited by R. Maenner and H.P. Schwefel), Lectures in Computer Science, Vol. 496, pp. 109-116. Springer, Berlin.
52. Whitley, D. (1989). "The GENITOR Algorithm and Selective Pressure: Why Rank Based Allocation of Reproductive Trials is Best," Morgan Kaufmann, J. D. Schaffer, ed., pp. 116-121.
53. Whitley, D. and J. Kauth, (1988). "GENITOR: A Different Genetic Algorithm," Proceedings of the 1988 Rocky Mountain Conference on Artificial Intelligence.
54. Yamada, T. and R. Nakano (1996). "Scheduling by Genetic Local Search with Multi-Step Crossover," 4th International Conference on Parallel Problem Solving from Nature, 960-969.
55. Yamada, T. and C. Reeves (1997). "Permutation Flowshop Scheduling by Genetic Local Search," 2^d IEE/IEEE Int. Conf. on Genetic Algorithms in Engineering Systems (GALESIA '97), pp. 232-238, Glasgow, UK.

Filename: Scatter Search.doc
Directory: G:\Bartley
Template: C:\Program Files\Microsoft Office\Templates\Normal.dot
Title: SCATTER SEARCH AND PATH RELINKING
Subject:
Author: Schaefer
Keywords:
Comments:
Creation Date: 01/22/99 8:22 AM
Change Number: 166
Last Saved On: 03/15/00 10:47 AM
Last Saved By: Bartley
Total Editing Time: 1,710 Minutes
Last Printed On: 03/15/00 10:55 AM
As of Last Complete Printing
Number of Pages: 20
Number of Words: 9,304 (approx.)
Number of Characters: 53,034 (approx.)