The Case for Strategic Oscillation

Fred Glover

OptTek Systems, Inc., 1919 Seventh Street, Boulder, CO 80302 USA

glover@opttek.com

Jin-Kao Hao

LERIA, Université d'Angers
UFR Sciences, 2 Boulevard Lavoisier, 49045 Angers Cedex 01, Angers, France

jin-kao.hao@univ-angers.fr

Abstract

We study a "hard" optimization problem for metaheuristic search, where a natural neighborhood (that consists of moves for flipping the values of zero-one variables) confronts two local optima, separated by a maximum possible number of moves in the feasible space. Once a descent method reaches the first local optimum, all sequences of feasible moves to reach the second, which is the global optimum, must ultimately pass through solutions that are progressively worse until reaching the worst solution of all, which is adjacent to the global optimum.

We show how certain alternative neighborhoods can locate the global more readily, but disclose that each of these approaches encounters serious difficulties by slightly changing the problem formulation. We also identify other possible approaches that seem at first to be promising but turn out to have deficiencies.

Finally, we observe that a strategic oscillation approach for transitioning between feasible and infeasible space overcomes these difficulties, reinforcing recent published observations about the utility of solution trajectories that alternate between feasibility and infeasibility. We also sketch features of such an approach that have implications for future research.

1. Introduction

Consider the following disarmingly simple 0-1 integer linear programming (ILP) problem

(P1) Maximize $$x_0 = nx_1 - \sum_{j \in N_2} x_j \qquad (1.1)$$

subject to

$$(n-1)x_1 - \sum_{j \in N_2} x_j \le 0 \qquad (1.2)$$

$$0 \le x_j \le 1,\ j \in N \qquad (1.3)$$

$$x_j \text{ is integer, } j \in N \qquad (1.4)$$

Here N = {1, …, n } denotes the index set for the x variables and N_2 = {2, 3, …, n } (= N – {1}) is the index set for all variables except x_1. This problem has the interesting property of being "hard" for a search method that starts from the feasible solution where x_j = 0 for all $j \in N$ and uses a search neighborhood that consists of flipping (complementing) the values of the 0-1 variables. As we consider the outcome of applying such a search method, we begin by imposing the requirement of remaining feasible at each step. This requirement reflects a bias found in many search procedures, which favor making feasible moves whenever possible, and in the present case would seem entirely benign, because there is a path through this "feasible flip" neighborhood that leads to the unique globally optimum solution, and in fact it is possible to start from any feasible 0-1 solution and trace a path in this neighborhood to every other feasible solution. (Such a property is not satisfied in many 0-1 ILP problems.)

The feature of this problem that makes it hard for such a search process is that the 0 solution is a local optimum in the feasible space, while the global optimum occurs by setting x_j = 1 for all $j \in N$, and this latter solution is the only locally optimal solution outside of the 0 solution. (The difference in quality between these two solutions can be magnified by giving x_1 a larger coefficient in the objective function.) Still worse, if x' is any other feasible solution except for the "worst" (smallest x_o value) feasible solution that sets x_1' = 0 and x_j' = 1 for all $j \in N_2$, then all moves from x' that lead to an improved feasible solution consist of identifying some x_j' that currently has a value of 1 and setting it equal to 0, hence moving back toward the 0 solution.

In short, a method that starts anywhere in feasible space except at the worst feasible solution, or the global optimum itself, and that tries to make improving moves, will always be driven to revisit the 0 solution. For example, starting from any feasible solution except the global optimum and the worst feasible solution adjacent to it, all improving moves result by flipping some variable with a value of 1 to receive a value of 0, and all improving paths hence likewise consist of such flips until reaching the local optimum with all $x_j = 0$.

A question that quickly comes to mind is whether a straightforward intervention exists to counter this slide back to 0. In fact, the “recency memory” approach often used in simple forms of tabu search, which prevents a certain number of the most recently chosen moves from being reversed, seems usefully suited for this purpose, and on the surface would appear to be capable of removing the difficulty. However, in order to ultimately drive the method to find the global optimum it would be necessary to maintain an “extreme” tabu restriction that forbids reversing each of the last $n – 1$ moves. If the method is started one step away from the worst feasible solution (so that it would first move toward the 0 solution, which it would reach in $n – 2$ steps), then the tabu restriction would need to be retained for a succession of $2n – 3$ moves. This is not an excessive number of moves under most circumstances, but such an extreme tabu restriction creates a very rigid search procedure that severely limits the ability to choose among available moves, and consequently it would provide an ineffective strategy for solving most types of 0-1 ILP problems. We will return to the issue of tabu restrictions later, but for the moment we note that such a recency memory strategy does not appear to be a good resolution of the difficulty faced.

2. What Can Be Done?

Admittedly, this is a “cooked” example that is purposely designed to be perverse for the simple 0-1 flip neighborhood used. However, if we cannot find an approach that overcomes the difficulty posed by a problem of such simplicity, then we may also anticipate that we will be highly vulnerable to failing when confronted with more complex problems. Granting that an example always exists that can make any given method perform badly, we may nevertheless ask whether there is at least a relatively simple algorithmic design that will resist being confounded so readily, and that therefore gives a chance of performing more effectively than in the current illustration.

Some of the strategies we explore in our following discussions include:

1) Specially structured neighborhoods
2) Multiple flip neighborhoods
3) Linear programming relaxations
4) Feasible/infeasible strategies (drawing on strategic oscillation)

As a last ditch effort to avoid having to come up with a more substantive response, we might be tempted to “reverse engineer the formulation” – that is, by drawing on the knowledge of a preferable solution to start the search from, we might complement the 0-1 variables x_j, $j \in N_2$ (replacing them by variables $y_j = 1 - x_j$). Then the initial "0 solution" would be the worst feasible solution that is one move away from the globally optimal solution and the difficulty of this problem would go away. But by allowing this kind of trickery we could just as well complement all variables and start at the global optimum itself. We seek a more useful type of response that does not depend on knowing in advance the characteristics of the problem or the location of the best solution.

3. Reasonable Possibilities and Loopholes

A conspicuous first possibility to consider is to change the neighborhood employed in attempting to solve the example problem. To be sure, we do not want to propose a neighborhood that is too arcane, because then no one will use it and its value would be limited. Instead, we are motivated to find a neighborhood that employs some natural search principle, and hence that could "make sense" as the basis for a more general approach.

There is a neighborhood that seems to overcome the difficulties encountered called the CX neighborhood (as examined, for example, in Reeves, 2006), where the neighbors of a given solution consist of the n solutions that simultaneously flip all the variables indexed k through n, as k takes the values $k = 1$ to n. Then the "all 1" solution is adjacent to the 0 solution (being the neighbor that arises when $k = 1$), and hence the CX neighborhood will uncover the global optimum in short order.

However, the CX neighborhood runs into two limitations that eliminate it as useful in this context. First, it can require much more effort to evaluate than the single flip neighborhood (for example requiring $O(n^2)$ operations versus $O(n)$ operations in the application of Glover et al. 1999). Second, and more to the point, we can simply change our formulation slightly so that the best solution is not the complement of the 0 solution, and then the CX neighborhood loses its ability to uncover the global optimum. For instance, it suffices to add a variable x_{n+1} which has a coefficient of –1 in the objective function (1.1) (in common with the variables x_2 through x_n) and that has a coefficient of 1 (instead of –1) in the constraint (1.2). Then x_{n+1} takes a value of 0 rather than 1 in the global optimum, and the CX neighborhood will not help to find this solution. A variety of other kinds of formulation variations can also thwart this change of neighborhoods.

Instead of looking for a neighborhood that will stumble on the global optimum by blind luck, we are motivated to look for one that can help to find this solution by a more systematic means. One of the most commonly used "alternative neighborhood strategies" in 0-1 optimization is to flip the value not just of a single variable, but of 2 or more variables simultaneously. In its naïve form such an approach is rather costly to execute, since the number of ways to flip k variables out of n can become large even for fairly small values of k. To combat this effect, in practice multi-flip neighborhoods are employed by restricting attention to promising subsets of possibilities through the use of candidate list strategies (see, e.g., Chapter 3 of Glover and Laguna, 1997). Regardless of whether such refinements are employed, however, the resulting compound neighborhoods turn out to be of little use in the present example. Even if all possible "k flips" could be examined in a reasonable time (say for $k \leq 4$), such a procedure would not be immune to the defect of repeatedly falling back into the 0 solution when applied to solving problem (1).

An approach that contains a still larger element of "strategy," and appears to offer more hope in the present setting, is to employ a neighborhood based on linear programming (LP). In this approach, we would temporarily relax the integer requirement (1.4) and solve the problem (1) as an LP problem in the quest for a solution that is in some sense closer to an optimal integer solution. The method would then continue to solve various amended LP problems, or carry out a search by pivoting to adjacent extreme point solutions, in an effort to close the gap between the present solution and one that is integer feasible.

In fact, this approach appears to be a splendid strategy in the present case, because the primal simplex LP method will quickly find the global ILP solution, without bothering to seek to impose the integer requirement (1.4). Specifically, starting from the 0 extreme point solution, the first pivot with the primal method is degenerate (leading to an LP basis that still gives the 0 solution), but afterward the method marches along a path of improving fractional extreme points until reaching the optimum that sets all variables to 1. [1]

Regrettably, however, this strategy is likewise easily thwarted by a minor change in the problem formulation. To confound this LP approach, all that is required is to replace the constraint (1.2) by the constraint

$$(n+h-1)x_1 - \sum_{j \in N_2} x_j \leq h$$

where h is selected to be some small positive integer > 1. (Constraint (1.2) itself

[1] It is possible to use a different formulation that replaces the constraint (1.2) by the collection of constraints $x_1 \leq x_j$, $j \in N_2$, then the simplex method would make $n - 1$ degenerate pivots followed by a single pivot going from the 0 solution to the "all 1" solution. This alternative formulation responds identically to the method introduced in the next section for resolving the complications noted here.

results when $h = 0$.) Then the optimum LP solution sets x_1 to the fractional value $h/(n + h - 1)$ together with setting all $x_j = 0$ for $j \in N_2$, giving an objective value of $x_o = nh/(n + h - 1)$, which exceeds 1 when $h > 1$. For example, when $h = 5$ and $n = 100$, the optimum LP solution yields $x_1 = 5/104$ and $x_o = 500/104$. Rounding this solution to its nearest integer neighbor gives the 0 solution again. Moreover, using the LP solution as a starting local optimum, this approach has all the deviously bad features for an LP pivoting neighborhood that the “all 0” local optimum has for the flip neighborhood.

We can overcome the immediate difficulty if we augment the LP solution approach by incorporating special cutting planes (see, e.g., Eckstein and Nediak, 2007; Glover, 2006, 2008), but it is correspondingly easy to create a formulation that disrupts this strategy as well.

Taking stock

At this point, in view of the multiple difficulties that emerge when we try to manufacture different types of neighborhoods or to manipulate the problem into an exploitable form, we are strongly motivated to step back and take a broader perspective. It would be valuable if we could find an approach that retains the neighborhood that flips 0-1variables, but rescues it from its deficiencies. The rationale underlying such a perspective is that the 0-1 flip neighborhood is analogous to a variety of neighborhoods for other kinds of combinatorial problems (especially neighborhoods that build on a basic design of “adding” and “dropping” solution components), and these neighborhoods additionally arise in applications that do not conveniently lend themselves to an ILP formulation. (More precisely, such a formulation exists for many of these combinatorial problems, but creates a model that is exceedingly difficult to solve using customary ILP methods.) Consequently, we confront the question of whether there is any way to salvage the 0-1 flip neighborhood by modifying the way it is used.

4. Feasibility/Infeasibility Asymmetry and Strategic Oscillation

A useful clue concerning a way to effectively restructure our use of the 0-1 flip moves is provided by a recent observation concerning bounded ILP problems which has implications for the broad range of combinatorial problems having equivalent integer programming formulations. In particular, bounded ILP problems exhibit a novel asymmetry, embodied in contrasting properties of their feasible and infeasible regions. As demonstrated in Glover (2007), the infeasible space for such problems is always connected when using trajectories that change the values of the integer variables by unit amounts (such as the 0-1 flip neighborhood in the case of 0-1 problems) – a property that lies in stark contrast to the situation for the feasible space, which offers no such guarantee. Even where trajectories exist between all integer solutions in feasible space, as they do in our illustrated problem (1), the structure of infeasible space can be conducive to creating simpler and more direct types of solution paths. In addition, paths that have the latitude to cross back and forth through feasible and infeasible space have features that make them particularly attractive.

A procedure called strategic oscillation, which was initially proposed with the purpose of crossing back and forth between feasible and infeasible space, is well suited to taking advantage of this asymmetry. Occupying a key position among tabu search strategies, strategic oscillation has notably also been used in settings for transitioning between multiple neighborhoods, decision rules and search regions (Glover, 1977; Freville and Plateau, 1986; Kelly, Golden and Assad, 1993; Osman, 1993; Vasquez and Hao, 2001a, 2001b; Hvattum, Lokketangen and Glover, 2005). One of the explanations suggested for the success of the approach lies in its ability to integrate diversification with intensification, without resorting to "randomized" forms of diversification. (Randomized diversification, in spite of its popularity in some metaheuristic approaches, will clearly have little value in the context of the challenge previously illustrated.)

5. A Rudimentary Strategic Oscillation Method

To show the utility of the strategic oscillation approach in the present context, we sketch a rudimentary version and examine how it performs in addressing the problem (1).

As a basis for such a method, we will evaluate a prospective move in relation both to its "quality" and to the degree of its potential "infeasibility". For simplicity, we will define the quality of a move by reference to the change it induces in the objective function. Hence specifically, relative to a specific set of move options, we define a "best move" to be a member of the set that improves the objective function the most or (in case no improvement is possible) causes it to deteriorate the least. The infeasibility evaluation will be expressed as a simple function of the constraint violations, such as a weighted sum of such violations. In the setting of problem (1), such an evaluation can be taken to be the amount by which the solution produced by the move violates the constraint (1.2).

With these conventions, our simple approach for solving problem (1), starting from any feasible or infeasible 0-1 solution, may be described as follows.

I. While the current solution is feasible:

A. If a move exits to a feasible solution better than the current one, select a best move from the set of these options.

B. If no move exists to a better feasible solution, choose a best move from all moves available. (The resulting move may enter infeasible space.)

II. While the current solution is infeasible:

A. If a move exists that improves the infeasibility measure, select a best move from the set of these options. (The resulting move may enter feasible space.)

B. If no move exists that improves the infeasibility measure, select a best move from the set that causes this measure to deteriorate the least.

It may be noted that this method does not treat feasibility and infeasibility symmetrically, and in this respect accords with the principle underscored in our earlier discussion concerning the desirability of a procedure that operates differently in feasible and infeasible space.

Applied to problem (1), if we start from an initial feasible solution, the preceding method first executes a series of improving moves employing step I.A. that takes it to a locally optimal solution. In this case, assuming the starting solution was not itself the global optimum or the solution adjacent to it, the solution attained will be the 0 solution. Upon reaching this juncture the method applies step I.B. to immediately cross into infeasible space, since setting $x_1 = 1$ is the best of all available moves (by the simplified definition of "best" used here). The next series of moves applies step II.A. to systematically march back to feasibility, setting the variables x_2 through x_n successively equal to 1 (since these are the best moves from those that improve the infeasibility measure). At this point the globally optimal solution is obtained.

If instead the method starts from an infeasible solution, it will likewise employ a series of moves to systematically march back to feasibility, and will obtain the global optimum when feasibility is reached. It is also readily seen that this method will perform essentially in the same way when applied to any of the modified formulations previously discussed that thwart various attempts to create alternative neighborhoods and alternative solution strategies.

In fact, if we change the definition of a best move to use a common "ratio definition" (identified below) the method will make the same moves as before when starting in infeasible space – or from a local optimum or the worst feasible solution – but will require even fewer moves than before in all other cases. It suffices to define a best move to be one that yields a maximum ratio of the objective function improvement to the change in resources consumed by the move when the current solution is feasible, and to be the one that yields the minimum ratio of the objective function deterioration to the reduction of the infeasibility measure when the current solution is infeasible. (This ratio definition should appropriately become more subtle for general applications, as discussed in the Appendix.) Then starting from any feasible solution the best move will be the one that sets $x_1 = 1$, and if this does not immediately give the global optimum the solution will then be in infeasible space and will proceed directly to the global optimum by the same path previously indicated.

The message provided by this illustrated solution process is not that we are able to obtain the global optimum. We have already seen that this solution can be obtained by an intervention using a tabu restriction (although a restriction that will not work well for solving most problems). Rather, the moral is that we have found the global optimum by a strategy that does not make use of advanced knowledge about the nature of the optimum (as choosing an extreme tabu restriction does). In addition, this method can readily be extended to handle other more general problems. We examine some of the details for structuring a convenient and effective form of such an extension in the Appendix.

6. Conclusion

Evidently no single strategy is going to prevail when attempting to solve hard combinatorial problems. Judicious use of all the weapons in our arsenal for battling with recalcitrant problems is essential. At the same time, it can be helpful to analyze how to overcome obstacles presented by simple problems that are resistant to solution, by means of strategies having broader applicability.

In this spirit, drawing on lessons learned from the challenge posed by problem (1), we have identified the limitations of certain popular strategies (such as intervention by a rigid tabu restriction, the use of various alternative neighborhoods, incorporation of a strategy based on linear programming, and making recourse to randomized diversification). To cap these observations, we have disclosed that a simple strategy for crossing the feasibility/infeasibility boundary can be useful in overcoming these limitations, and can be readily generalized to broader settings.

Acknowledgment: The work is partially supported by a “Chaire d’excellence” from “Pays de la Loire” Region (France)” and regional MILES (2007-2009) and RaDaPop projects (2008-2011). We are grateful to the reviewers of this paper and to Cesar Rego for his assistance in improving our presentation.

References

J. Beasley (1993) “Lagrangian Relaxation,” *Modern Heuristic Techniques for Combinatorial Problems*, John Wiley and Sons, pp. 243-303.

J. Eckstein and M. Nediak (2007) “Pivot, Cut and Dive: a Heuristic for Mixed 0-1 Integer Programming,” *Journal of Heuristics*, 13(5), pages 471-503.

A. Freville and G. Plateau (1986) “Heuristics and Reduction Methods for Multiple Constraint 0-1 Linear Programming Problems,” *European Journal of Operational Research*, 24, pp. 206-215.

P. Galinier and J.-K. Hao (2004) “A general approach for constraint solving by local search,” *Journal of Mathematical Modelling and Algorithms* 3(1), pp. 73-88.

F. Glover (1977) “Heuristics for Integer Programming Using Surrogate Constraints," *Decision Sciences*, 8(1), pp. 156-166.

F. Glover (1995) "Tabu Thresholding: Improved Search by Nonmonotonic Trajectories," *ORSA Journal on Computing*, 7(3), pp. 426-442.

F. Glover (1996). "Tabu Search and Adaptive Memory Programming: Advances, Applications and Challenges," *Interfaces in Computer Science and Operations Research*, R. Barr, R. Helgason and J. Kennington (eds.) Kluwer Academic Publishers, pp. 1-75.

F. Glover (2000) “Multi-Start and Strategic Oscillation Methods – Principles to Exploit Adaptive Memory,” *Computing Tools for Modeling, Optimization and Simulation: Interfaces in Computer Science and Operations Research*, M. Laguna and J.L. Gonzales Velarde, eds., Kluwer Academic Publishers, pp. 1-24.

F. Glover (2003) “Tutorial on Surrogate Constraint Approaches for Optimization in Graphs,” *Journal of Heuristics*, 9(3), pp 175-227.

F. Glover (2006) “Parametric Tabu Search for Mixed Integer Programs,” *Computers and Operations Research*, 33(9), pp. 2449-2494

F. Glover (2007) “Infeasible/Feasible Search Trajectories and Directional Rounding in Integer Programming,” *Journal of Heuristics*, 13(6), pp. 505-542.

F. Glover (2008) “Inequalities and Target Objectives for Metaheuristic Search – Part I: Mixed Binary Optimization,” in *Advances in Metaheuristics for Hard Optimization*, P. Siarry and Z. Michalewicz, eds. Springer, New York, pp. 439-474.

F. Glover and M. Laguna (1997) *Tabu Search*, Kluwer Academic Publishers, 1997.

F. Glover, G. Kochenberger, B. Alidaee, and M.M. Amini (1999) “Tabu with Search Critical Event Memory: An Enhanced Application for Binary Quadratic Programs,” In: *MetaHeuristics: Advances and Trends in Local Search Paradigms for Optimization,* (Eds.) S. Voss, S. Martello, I. Osman, and C. Roucairol. Kluwer Academic Publisher, Boston, pp. 93-110.

L.M. Hvattum, A. Lokketangen and F. Glover (2005) “New Heuristics and Adaptive Memory Procedures for Boolean Optimization Problems,” *Integer Programming Theory and Practice*, J. K. Karlof, ed.. pp. 1-18.

J. Kelly, B. Golden and A. Assad (1993) “Large Scale Controlled Rounding Using Tabu Search with Strategic Oscillation,” *Annals of Operations Research* 41, pp. 69-84.

Z. Lü and J.-K. Hao (2008) “Adaptive Tabu Search for Course Timetabling,” *European Journal of Operational Research* (In press).

K. Nonobe and T. Ibaraki (1998) “A tabu search approach for the constraint satisfaction problem as a general problem solver,” *European J. of Operational Research*, Vol. 106, pp.599-623.

K. Nonobe and T. Ibaraki, (2001) “An improved tabu search method for the weighted constraint satisfaction problem,'' *INFOR*, 39, pp.131-151.

I.H. Osman (1993) “Metastrategy Simulated Annealing and Tabu Search Algorithms for the Vehicle Routing Problem,” Annals of Operations Research 41, pp.421-451.

C. R. Reeves (2006) “Fitness Landscapes,” Chapter 19 of *Search Methodologies*, Springer.

C. Rego (2005) “RAMP: A New Metaheuristic Framework for Combinatorial Optimization,” in

Metaheuristic Optimization via Memory and Evolution: Tabu Search and Scatter Search, C. Rego and B. Alidaee (Eds.), Kluwer Academic Publishers, pp. 441-460.

C. Rego and B. Alidaee (2005) *Metaheuristic Optimization via Memory and Evolution: Tabu Search and Scatter Search*, Kluwer Academic Publishers.

M. Vasquez and J.-K. Hao (2001a) “A heuristic approach for antenna positioning in cellular networks,” *Journal of Heuristics*. 7(5), pp. 443-472.

M. Vasquez and J.-K. Hao (2001b) “A ‘logic-constrained’ knapsack formulation and a Tabu algorithm for the daily photograph scheduling of an earth observation satellite,” *Computational Optimization and Applications,* 20(2), pp. 137-157.

Appendix: A More General Strategic Oscillation Procedure

We briefly sketch a few of the features of a more general strategic oscillation procedure for crossing the feasibility boundary, by looking at some of the possible functions of such a procedure. More detailed information on strategic oscillation can be found in Glover and Laguna (1997) and Glover (1995, 2000).

Exploiting Asymmetry and the Direction of Oscillation

Our comments will focus on observations that are straightforward in nature, but that are nevertheless often overlooked and that can make a significant difference in the effectiveness of a strategic oscillation procedure.

Strategic oscillation is often structured to penetrate varying distances beyond the boundary between feasible and infeasible space. As a result, there can be a difference in the nature of an evaluation that depends not only on the region in which a solution lies, but also on whether the search is currently undertaking to approach or move away from the boundary. Thus, it is generally advisable to employ different evaluations according to direction of search in relation to the feasibility boundary as well as according to the region in which the current solution lies.

To set the stage for considering the nature of useful evaluations, and how they change under different circumstances, we observe that the notion of moving "toward" or "away from" a feasibility boundary must be chiefly restricted to considering the role of inequality constraints when the search is in feasible space, but will also make reference to equality constraints when the search is in infeasible space. If the neighborhood employed does not assure that the equality constraints will remain satisfied once feasibility is attained, the oscillation will be primarily one-sided, spending most of its time in the infeasible region.

Multi-move combinations are often useful for procedures that seek to satisfy equality constraints, and it is particularly worth noting that such move combinations become increasingly relevant for finding improved solutions (even in the absence of equality constraints) as the search trajectory gets close to the feasibility boundary. The importance of "near boundary" conditions for triggering more intensive search is highlighted in Glover (1995, 2000) and in Hvattum, Lokketangen and Glover (2005).

Features of Evaluations

Evaluations within feasible space typically involve creating measures of objective function gain (or loss) by comparison to the amount of resources of various constraints that are consumed or made available by a move. Evaluations within infeasible space involve similar measures but more particularly emphasize the changes in an infeasibility measure associated with the constraints. In forming such evaluations, it is relevant to differentiate constraints by their relative importance, which can be determined by duality analysis involving the use of problem relaxations. Both Lagrangean relaxation and surrogate constraint relaxation can be of value in such applications, as demonstrated in Beasley (1993), Glover (2003) and Rego (2005).

Differentiating constraints by their relative importance can be further aided by a form of learning that keeps track of which constraints are most limiting on high quality solutions found throughout the search. Useful proposals for doing this are given in Rego and Alidaee (2005). Upon approaching feasibility boundaries, however, such measures of relative importance should be given progressively less weight in comparison with measures that reflect by the current restrictiveness of each constraint.

An interesting approach for handling constraints in the setting of constraint satisfaction problems is given in Galinier and Hao (2004), by introducing a way of defining a penalty function for each predefined constraint. In essence, each penalty measures the minimum number of variables that need to be modified to reach a consistent assignment. The penalties are then exploited by a tabu search approach that creates an evaluation function consisting of a weighted aggregation of the penalty terms. This technique helps the search to approach feasibility more effectively than other “hand-made” penalties.

Oscillation Over Parameter Settings

The fact that different parameters can be attached to the multiple components of evaluations made during strategic oscillation invites the use of an associated type of strategic oscillation that varies the values of these parameters.

An extremely simple form of such an approach can be illustrated by considering the use of just two basic measures, one identifying the quality of an objective function change and the other identifying the “net” satisfaction or violation of the constraints. The oscillation then varies the weight attached to the second measure, so that a decreased weight on the degree of satisfying constraints will allow the search within a feasible region to move toward and ultimately cross the feasibility boundary, while an increased weight on satisfying constraints (i.e., on reducing their violations) will induce the search to turn around inside the infeasible region and again head toward the feasibility boundary. Depending on the nature of the feasible region, the search may then penetrate "further into" this region until the weight is shifted to the point where progress back toward the boundary is again initiated.

Appreciably more elaborate strategies than this "single parameter oscillation" approach are possible, and deserve to be considered when challenging problems are confronted. An interesting variant of a multi-parameter procedure has been created by Nonobe and Ibaraki (2001) to give a general purpose method that they demonstrate in the context of weighted constraint satisfaction problems. In their approach, weights are given to all constraints individually, and controlled independently of others. The weight given to a constraint is increased if the solutions currently being searched stay infeasible for the constraint, while it is decreased in the other case. In this way, weights make changes up and down during computation, establishing an automatic control that creates a strategic oscillation between feasibility and infeasibility without setting a preplanned scenario. When applied to the example problem illustrated in this paper their method likewise succeeds in uncovering the global optimum solution.

Transitioning Between Ratio Evaluations and Threshold Evaluations

Many oscillation searches make use of evaluations involving ratios that divide changes in the objective function by changes in the degree of satisfying or violating constraints. As the search draws closer to the feasibility boundary this type of evaluation should be shifted into the background, to be replaced by evaluation that instead makes use of thresholds. Thus, for example, it can become preferable in the vicinity of the boundary to select best moves based strictly on the objective function change, subject to satisfying a threshold that limits the amount of deterioration in the level of meeting or violating the constraints. Alternatively, the criterion for selecting a best move may be based on the degree to which constraints are satisfied or violated, subject to meeting some threshold of improvement (or of limited deterioration) in the objective function.

Implications for Tabu Restrictions and Associated Memories

The fact that problem (1) can be solved by using an extreme tabu restriction, even though such a restriction would be a poor strategy to apply generally, suggests the merit of a form of strategic oscillation that varies this restriction – or more accurately, that varies the *tabu tenure* which specifies the number of moves that a tabu restriction is maintained in force.

Such an oscillating tenure strategy would consist of periodically electing to increase the tenure by chosen increments until reaching an extreme value, and then to return the tenure to its customary range (either immediately or by a more gradual reduction schedule). This type of approach would appropriately be executed by superimposing it on ordinary dynamic tabu tenure strategies (Glover and Laguna, 1997) and it can likewise be used in conjunction with self-adjusting tabu tenure approaches, such as those proposed by Nonobe and Ibaraki (1998) and Lü and Hao (2008).

A useful supplement to such an approach occurs by maintaining a *reference set* that consists of high quality solutions (or local optima) previously found. The choice of new moves during diversification phases is then made to favor those that endow new solutions with attributes (e.g., values of variables) that infrequently occur among solutions in the reference set, until a preferred level of difference is attained between the current solution and the reference set solutions.

This approach can be further supplemented by keeping a record of moves that were often attractive in the past, but that were infrequently or never chosen because they did not manage to become the "top pick." Such a record may include moves that were excluded based on some screening rule but that nevertheless had a feature deemed attractive, either in terms of improving the objective or of enhancing some other measure of interest. (For example, in the strategy first illustrated for solving (1) the move of setting $x_1 = 1$ was excluded on the grounds of creating infeasibility, but apart from this basis for exclusion, the move would be considered attractive.) Attributes of these moves that did not therefore become part of solutions previously visited are valuable to consider as components of new moves.

The use of critical event memory and conditional critical event memory, along with associated processes for giving improved control by strategic oscillation, are discussed in Glover (2000).