

Scatter Search and Local NLP Solvers: A Multistart Framework for Global Optimization

Zsolt Ugray

Management Information Systems Department, Utah State University, Logan, Utah 84322,
 zsolt.ugray@usu.edu

Leon Lasdon

Department of Information, Risk, and Operations Management, McCombs School of Business,
 The University of Texas at Austin, Austin, Texas 78712, lasdon@mail.utexas.edu

John Plummer

Department of Computer Information Systems and Quantitative Methods, Texas State University,
 San Marcos, Texas 78666, jp05@txstate.edu

Fred Glover

University of Colorado, Boulder, Colorado 80309, fred.glover@colorado.edu

James Kelly

OptTek Systems, Inc., Boulder, Colorado 80302, kelly@opttek.com

Rafael Martí

Departamento de Estadística e Investigación Operativa, University of Valencia,
 46100 Burjassot, Valencia, Spain, rafael.marti@uv.es

The algorithm described here, called OptQuest/NLP or OQNLP, is a heuristic designed to find global optima for pure and mixed integer nonlinear problems with many constraints and variables, where all problem functions are differentiable with respect to the continuous variables. It uses OptQuest, a commercial implementation of scatter search developed by OptTek Systems, Inc., to provide starting points for any gradient-based local solver for nonlinear programming (NLP) problems. This solver seeks a local solution from a subset of these points, holding discrete variables fixed. The procedure is motivated by our desire to combine the superior accuracy and feasibility-seeking behavior of gradient-based local NLP solvers with the global optimization abilities of OptQuest. Computational results include 155 smooth NLP and mixed integer nonlinear program (MINLP) problems due to Floudas et al. (1999), most with both linear and nonlinear constraints, coded in the GAMS modeling language. Some are quite large for global optimization, with over 100 variables and 100 constraints. Global solutions to almost all problems are found in a small number of local solver calls, often one or two.

Key words: global optimization; multistart heuristic; mixed integer nonlinear programming; scatter search; gradient methods

History: Accepted by Michel Gendreau, Area Editor for Heuristic Search and Learning; received August 2002; revised January 2005, September 2005; accepted January 2006. Published online in *Articles in Advance* July 20, 2007.

1. Introduction

This paper describes OQNLP, a multistart heuristic algorithm designed to find global optima of smooth constrained nonlinear programs (NLPs) and mixed integer nonlinear programs (MINLPs). It uses the OptQuest Callable Library (OCL) implementation of scatter search (Laguna and Martí 2002) to generate trial points, which are candidate starting points for a local NLP solver. These are filtered to provide a smaller subset from which the solver attempts to find a local optimum. Our GAMS implementation can use any GAMS NLP solver, and the stand-alone version

uses the generalized reduced gradient NLP solver LSGRG2 (Smith and Lasdon 1992).

The most general problem this algorithm can solve has the form

$$\text{minimize } f(x, y), \quad (1)$$

subject to the nonlinear constraints

$$gl \leq G(x, y) \leq gu, \quad (2)$$

and to the linear constraints

$$l \leq A_1x + A_2y \leq u \quad (3)$$

$$x \in S, \quad y \in Y, \quad (4)$$

where x is an n -dimensional vector of continuous decision variables, y is a p -dimensional vector of discrete decision variables, and the vectors gl , gu , l , and u contain upper and lower bounds for the nonlinear and linear constraints, respectively. The matrices A_1 and A_2 are m_2 by n and m_2 by p , respectively, and contain the coefficients of any linear constraints. The set S is defined by simple bounds on x , and we assume that it is closed and bounded, i.e., that each component of x has a finite upper and lower bound. This is required by the OptQuest scatter-search procedure. The set Y is assumed to be finite, and is often the set of all p -dimensional binary or integer vectors y that satisfy finite bounds. The objective function f and the m_1 -dimensional vector of constraint functions G are assumed to have continuous first partial derivatives at all points in $S \times Y$. This is necessary so that a gradient-based local NLP solver can be applied to the relaxed NLP sub-problems formed from (1)–(4) by allowing the y variables to be continuous.

2. Multistart Algorithms for Global Optimization

In this section, which reviews past work on multistart algorithms, we focus on unconstrained problems where there are no discrete variables, since to the best of our knowledge multistart algorithms have been investigated theoretically only in this context. These problems have the form of (1)–(4) with no y variables and no constraints except the bounds $x \in S$ in (4). All global minima of f are assumed to occur in the interior of S . By multistart we mean any algorithm that attempts to find a global solution by starting a local NLP solver, denoted by L , from multiple starting points in S . The most basic multistart method generates uniformly distributed points in S , and starts L from each of these. This converges to a global solution with probability one as the number of points approaches infinity—in fact, the best of the starting points converges as well. However, this procedure is very inefficient because the same local solution is located many times. A convergent procedure that largely overcomes this difficulty is called multilevel single linkage (MLSL) (Rinnooy Kan and Timmer 1987a, b). MLSL uses a simple rule to exclude some potential starting points. A uniformly distributed sample of N_s points in S is generated, and the objective f is evaluated at each point. The points are sorted according to their f values, and the qN_s best points are retained, where q is an algorithm parameter between 0 and 1. L is started from each point of this reduced sample, except if there is another sample point within a certain critical distance that has a lower f value. L is also not started from sample points that are too near the boundary of S , or

too close to a previously discovered local minimum. Then, N_s additional uniformly distributed points are generated, and the procedure is applied to the union of these points and those retained from previous iterations. The critical distance referred to above decreases each time a new set of sample points is added. The authors show that, if the sampling continues indefinitely, each local minimum of f will be located, but the total number of local searches is finite with probability one. They also develop Bayesian stopping rules, which incorporate assumptions about the costs and potential benefits of further function evaluations, to determine when to stop the procedure.

When the critical distance decreases, a point from which L was previously not started may become a starting point in the next cycle. Hence all sample points generated must be saved. This also makes the choice of the sample size N_s important, since too small a sample leads to many revised decisions, while too large a sample will cause L to be started many times. Random-linkage (RL) multistart algorithms introduced by Locatelli and Schoen (1999) retain the good convergence properties of MLSL, and do not require that past starting decisions be revised. Uniformly distributed points are generated one at a time, and L is started from each point with a probability given by a nondecreasing function $\phi(d)$, where d is the distance from the current sample point to the closest of the previous sample points with a better function value. Assumptions on this function that give RL methods the same theoretical properties as MLSL are derived in the above reference.

Recently, Fylstra et al. have implemented a version of MLSL that can solve constrained problems (<http://www.solver.com/xlspresolv3.htm>). Limited to problems with no discrete variables y , it uses the L_1 exact penalty function, defined as

$$P_1(x, w) = f(x) + \sum_{i=1}^m w_i \text{viol}(g_i(x)), \quad (5)$$

where the w_i are nonnegative penalty weights, $m = m_1 + m_2$, and the vector g has been extended to include the linear constraints (4). The function $\text{viol}(g_i(x))$ is the absolute amount by which the i th constraint is violated at the point x . It is well known (Nash and Sofer 1996) that if x^* is a local optimum of (1)–(4), u^* is a corresponding optimal multiplier vector, the second order sufficiency conditions are satisfied at (x^*, u^*) and

$$w_i > \text{abs}(u_i^*), \quad (6)$$

then x^* is a local unconstrained minimum of P_1 . If (1)–(4) has several local minima, and each w_i is larger than the maximum of all absolute multipliers for constraint i over all these optima, then P_1 has a local minimum at each of these local constrained minima. Even

though P_1 is not a differentiable function of x , MLSL can be applied to it, and when a randomly generated trial point satisfies the MLSL criterion to be a starting point, any local solver for the smooth NLP problem can be started from that point. The local solver need not make any reference to the exact penalty function P_1 , whose only role is to provide function values to MLSL. We will use P_1 in the same way in our OQNLP algorithm. We are not aware of any theoretical investigations of this extended MLSL procedure, so it must currently be regarded as a heuristic.

3. The OQNLP Algorithm

3.1. The Global Phase—Scatter Search

Scatter search (ScS) is a population-based meta-heuristic algorithm devised to intelligently perform a search on the problem domain (Glover 1998). It operates on a set of solutions called the *reference set* or *population*. Elements of the population are maintained and updated from iteration to iteration. ScS differs from other population-based evolutionary heuristics like genetic algorithms (GAs) mainly in its emphasis on generating new elements of the population mostly by deterministic combinations of previous members of the population as opposed to the more extensive use of randomization. ScS was founded on strategies that were proposed as augmentations to GAs more than a decade after their debut in ScS. It embodies principles and strategies that are still not emulated by other evolutionary methods and prove to be advantageous for solving a variety of complex optimization problems. For the most recent and complete description of ScS, see Laguna and Marti (2003). Also see the Online Supplement to this paper on the journal's website.

3.2. The Local Phase—Gradient-Based NLP Solvers

Many papers and texts discuss gradient-based NLP solvers, e.g., Nash and Sofer (1996), Nocedal and Wright (1999), and Edgar et al. (2001). These solve problems of the form (1)–(4), but with no discrete (y) variables. They require a starting point as input, and use values and gradients of the problem functions to generate a sequence of points that, under fairly general smoothness and regularity conditions, converge to a local optimum. The main classes of algorithms in widespread use are successive quadratic programming (SQP) and generalized reduced gradient (GRG); see Edgar et al. (2001, chapter 8). The algorithm implemented in the widely used MINOS solver (Murtagh and Saunders 1982) is similar to SQP. If there are nonlinear constraints, SQP and MINOS generate a sequence of points that usually violate the nonlinear constraints, with the violations decreasing to

within a specified feasibility tolerance as the sequence converges to a local optimum. GRG algorithms have a simplex-like phase 1–phase 2 structure. Phase 1 begins with the given starting point and, if it is not feasible, attempts to find a feasible point by minimizing the sum of constraint violations. If this effort terminates with some constraints violated, the problem is assumed to be infeasible. However, this local optimum of the phase 1 objective may not be global, so a feasible point may exist. If a feasible point is found, phase 2 uses it as its starting point, and proceeds to minimize the true objective. Both phases consist of a sequence of line searches, each of which produces a feasible point with an objective value not worse (and usually better) than its predecessor.

Several good commercially available implementations of GRG and SQP solvers exist; see Nash (1998) for a review. As with any numerical-analysis software, a local NLP solver can fail to find a local solution from a specified starting point. The problem may be too badly conditioned, badly scaled, or too large for the solver, causing it to terminate at a point (feasible or infeasible) that is not locally optimal. While the reliability of the best current NLP solvers is quite high, these difficulties occurred in our computational testing, and we discuss this in more detail later.

Let L be a local NLP solver capable of solving (1)–(4), and assume that L converges to a local optimum for any starting point $x_0 \in S$. Let $L(x_0)$ be the locally optimal solution found by L starting from x_0 , and let x_i^* , $i = 1, 2, \dots, n_{loc}$ be all the local optima of the problem. The *basin of attraction* of the i th local optimum relative to L , denoted by $B(x_i^*)$, is the set of all starting points in S from which the sequence of points generated by L converges to x_i^* : $B(x_i^*) = \{x_0 \mid x_0 \in S, L(x_0) = x_i^*\}$.

One measure of difficulty of a global optimization problem with unique global solution x_1^* is the volume of $B(x_1^*)$ divided by the volume of the rectangle, S , the *relative volume* of $B(x_1^*)$. The problem is trivial if this relative volume is 1, as it is for convex programs, and problem difficulty increases as this relative volume approaches zero.

3.3. Comparing Heuristic Search Methods and Gradient-Based NLP Solvers

For smooth problems, the relative advantages of a heuristic search method like ScS over a gradient-based NLP solver are its ability to locate an approximation to a good local solution (often the global optimum), and the fact that it can handle discrete variables. Gradient-based NLP solvers converge to the “nearest” local solution, and have no facilities for discrete variables, unless they are embedded in a rounding heuristic or branch-and-bound method. Relative disadvantages of heuristic search methods are

their limited accuracy, and their weak abilities to deal with equality constraints (more generally, narrow feasible regions). They find it difficult to satisfy many nonlinear constraints to high accuracy, but this is a strength of gradient-based NLP solvers. Search methods also require an excessive number of iterations to find approximations to local or global optima accurate to more than two or three significant figures, while gradient-based solvers usually achieve four- to eight-digit accuracy rapidly.

The motivation for combining search and gradient-based solvers in a multistart procedure is to achieve the advantages of both while avoiding the disadvantages of either. Surprisingly, we have been unable to locate any published efforts in this direction, besides the Frontline extended MLSL method discussed in Section 2.

3.4. The OQNLP Algorithm

A pseudo-code description of the simplest OQNLP algorithm follows:

INITIALIZATION

```
Read_Problem_Parameters (n, p, m1, m2, bounds,
starting point);
Setup_OQNLP and OptQuest_parameters and Op-
tions (problem size, stage 1 and 2 iteration limits,
population size, accuracy, names and types of vari-
ables and constraints, bounds on variables and con-
straints);
Initialize_OptQuest_Population;
Stage 1 iterations = Stage 2 iterations = 0;
```

STAGE 1: INITIAL OPTQUEST ITERATIONS AND FIRST *L* CALL

WHILE (Stage 1 iterations < Stage 1 iteration limit)

```
DO {
    Get (trial solution from OptQuest);
    Evaluate (objective and nonlinear constraint
    values at trial solution,);
    Put (trial solution, objective and constraint
    values to OptQuest database);
    Stage 1 iterations = Stage 1 iterations + 1;
} ENDDO
```

```
Get_Best_Point_from_OptQuest_database (starting
point);
```

Call *L* (starting point, local solution);

Threshold = default value;

```
IF (local solution feasible) THEN {
    Insert local solution in linked list;
    Penalty weights = max(positive lower limit,
    absolute multiplier values from L call);
    threshold = P1 value of local solution; }
```

```
Penalty weights = max(positive lower limit,
absolute multiplier values from L call)
```

STAGE 2: MAIN ITERATIVE LOOP

WHILE (Stage 2 iterations < Stage 2 iteration limit)

```
DO {
    Get (trial solution from OptQuest);
    Evaluate (objective and nonlinear constraint
    values at trial solution,);
    Put (trial solution, objective and constraint
    values to OptQuest database);
    Calculate_Penalty_Function (trial solution,
    Penalty weights, P1);
    IF (distance and merit filter criteria are
    satisfied) THEN {
        Replace threshold with current P1 value;
        Call L (trial solution, local solution);
        IF (local solution feasible) THEN {
            Insert local solution in linked list;
            Penalty weights = max(positive lower
            limit, absolute multiplier values
            from L call); }
        }
    ELSE IF (P1 > threshold for waitcycle
    consecutive iterations) {increase threshold}
    Stage 2 iterations = Stage 2 iterations + 1;
} ENDDO
```

After initialization, there are two main stages. In the “initial OptQuest iterations” stage, the objective and constraint values at all trial points generated by the initial OptQuest population (including the population points themselves) are evaluated, and these values are returned to OptQuest, which computes its penalty function, P_{OQ} , at each point. The point with the best P_{OQ} value is selected, and L is started from this point. If there are any discrete variables y , they are fixed at their current values during the L solution process. Figure D in the Online Supplement shows a graph of these trial points for a two variable unconstrained problem. In general, they are scattered within the rectangle defined by the bounds on the variables, so choosing the best corresponds to performing a coarse search over this rectangle. If the best point falls inside the basin of attraction of the global optimum relative to L (as it often does), then if the subsequent L call is successful, it will find a global optimum. This call also determines optimal Lagrange multiplier values u^* for the constraints. These are used to determine initial values for the penalty weights w_i satisfying (6), which are used in the exact penalty function, P_1 , defined in (5). All local optima found are stored in a linked list, along with the associated Lagrange multipliers and objective values. Whenever a new local optimum is found, the penalty weights are updated so that (6) is satisfied over all known local optima.

The main iterative loop of stage 2 obtains trial points from OptQuest, and starts L from the subset of these points determined by two filters. The distance filter helps insure that these starting points are

diverse, in the sense that they are not too close to any previously found local solution. Its goal is to prevent L from starting more than once within the basin of attraction of any local optimum, so it plays the same role as the rule in the MLSL algorithm of Section 2, which does not start at a point if it is within a critical distance of a better point. When the final point found by L is feasible, it is stored in a linked list, ordered by its objective value, as is the Euclidean distance between it and the starting point that led to it. If a local solution is located more than once, the maximum of these distances, $maxdist$, is updated and stored. For each trial point t , if the distance between t and any local solution already found is less than $distfactor * maxdist$, L is not started from the point, and we obtain the next trial solution from OptQuest.

This distance filter implicitly assumes that the attraction basins are spherical, with radii at least $maxdist$. The default value of $distfactor$ is 0.75, and it can be set to any positive value. As $distfactor$ approaches zero, the filtering effect vanishes, as would be appropriate if there were many closely spaced local solutions. As it increases, the filtering effect increases until eventually L is never started in stage 2. Its default value is chosen empirically to achieve a reasonable compromise between these extremes.

The merit filter helps insure that the L starting points have high quality, by not starting from candidate points whose exact penalty function value P_1 in (5) is greater than a threshold. This threshold is set initially to the P_1 value of the best candidate point found in the first stage of the algorithm. If trial points are rejected by this test for more than $waitcycle$ consecutive iterations, the $threshold$ is increased by the updating rule:

$$threshold \leftarrow threshold + threshfactor * (1.0 + abs(threshold)),$$

where the default value of $threshfactor$ is 0.2 and that for $waitcycle$ is 20. The $threshfactor$ value is selected to provide a significant but not-too-large increase in $threshold$, while the $waitcycle$ value insures that these increases happen often enough but not too often. Using the default values of 200 stage 1 and 800 stage 2 candidate points, $threshold$ can be increased in stage 2 at most 40 times. The additive 1.0 term is included so that $threshold$ increases by at least $threshfactor$ when its current value is near zero. When a trial point is accepted by the merit filter, $threshold$ is decreased by setting it to the P_1 value of that point.

The combined effect of these two filters is that L is started at only a few percent of the OptQuest trial points, yet global optimal solutions are found for a very high percentage of the test problems. Some insight is gained by examining Figure 1, which shows




Figure 1 Local Optima and Ten L Starting Points for Six-Hump Camelback Function

the stationary point at the origin and the six local minima of the two-variable six-hump camelback function (Dixon and Szegö 1975) as dark squares, labeled with their objective value. The ten points from which OQNLP starts the local solver are shown as nine white diamonds, plus the origin. The local minima occur in pairs with equal objective value, located symmetrically about the origin. There were 144 trial points generated in the “initial OptQuest iterations” stage, and these, plus the ten points in the initial population, are shown in Figure D in the Online Supplement. The best of these 154 points is the population point $(0, 0)$, so this becomes the first starting point for the local solver. This happens to be a stationary point of F , so it satisfies the optimality test (that the norm of the gradient of the objective be less than the optimality tolerance), and the local solver terminates there. The next local solver start is at iteration 201, and this locates the global optimum at $(0.0898, -0.7127)$, which is located twice. The other global optimum at $(-0.0898, 0.7127)$ is found first at iteration 268, and is located six times.

The limit on total OQNLP iterations in this run was 1,000. L was started at only nine of the 846 OptQuest trial points generated in the main iterative loop of stage 2. All but two of the starting points are in the basin of attraction of one of the two global optima. This is mainly due to the merit filter. In particular, the threshold values are always less than 1.6071, so no starts are ever made in the basin of attraction of the two local optima with this objective value. The merit filter alone rejected 498 points, the distance filter alone 57, and both rejected 281.

Figure 2 illustrates the dynamics of the merit-filtering process for iterations 155 to 407 of this problem, displaying the objective values for the trial points as white diamonds, and the threshold values as dark lines. All objective values greater than 2.0 are set to 2.0.

The initial threshold value is zero, and it is raised twice to a level of 0.44 at iteration 201, where the




Figure 2 Objective and Threshold Values for Six-Hump Camelback Function for Iterations 155 to 407

trial-point objective value of -0.29 falls below it. L is then started and locates the global optimum at $(0.0898, -0.7127)$, and the threshold is reset to -0.29 . This cycle then repeats. Nine of the ten L starts are made in the 252 iterations shown in the graph. In this span, there are 12 points where the merit filter allows a start and the threshold is decreased, but L is not started at three of these because the distance filter rejects them.

Figure 3 shows the same information for iterations 408 to 1,000. There is only one L start in this span. This is not due to a lack of high-quality trial points: there are more good points than previously, many with values near or equal to -1.0310 (the global minimum is -1.0316), and the merit threshold is usually -1.0310 as well. Every time this threshold is raised, the merit

filter accepts one of the next trial points, but 51 of the 52 accepted points are too near one of the two global optima, and they are rejected by the distance filter.

This simple example illustrates a number of important points:

1. Setting the bounds on the continuous or discrete variables to be too large in magnitude is likely to slow the $OQNLP$ algorithm (or any search algorithm) and may lead to a poorer final solution. In the above example, if the variable bounds had been $[-2, 2]$ rather than $[10, 10]$, the trial points generated by the initial population would have had much lower objective values. OptQuest can overcome this when the initial population is updated.

2. L found a highly accurate approximation to the global solution of this unconstrained problem at its




Figure 3 Objective and Threshold Values for Six-Hump Camelback Function for Iterations 408 to 1,000

second call. OptQuest alone would have taken many more iterations to achieve this accuracy.

3. The best trial point generated by the initial population may not have as good an objective value as those generated from the second or succeeding ones, especially if the variable bounds are too large. Using the best “first-generation” point as the initial L starting point may not lead to as good a local solution as if some “second-generation” points had been considered. For this reason our base-case computational results use a first stage of 200 OptQuest trial points, which in this example would include all 144 first-generation points and 56 from the second generation.

3.5. Filtering Logic for Problems with Discrete Variables

The filtering logic described above must be extended when there are discrete variables (the y variables in the problem statement (1)–(4)). There are two distinct modes: (1) Optquest is aware of both the x and y variables and all problem constraints, and (2) Optquest is aware only of the y variables and the constraints involving y only. Tests thus far do not show conclusively that one of these is preferred, so both modes are described and comparisons thus far are presented later.

In mode (1), when a trial point (x_i, y_i) provided by OptQuest passes the two filtering tests and is passed to the local solver L , x_i acts as a starting point and is changed by L , but the y_i values are fixed and are not changed. Each new set of y_i values defines a different NLP for L to solve, say $NLP(y_i)$, with its own set of local minima in x space, so both filters must be made specific to $NLP(y_i)$. For the distance filter, it is irrelevant if x_i is close to any local minima (in x space) previously found that correspond to problems $NLP(y)$ with y different from y_i . Hence the distance filter is based on the distance from x_i to local minima of $NLP(y_i)$ only. Similarly, the tests and threshold values in the merit filter must be specific to the problem $NLP(y_i)$ currently being solved. However, the weights w in the exact penalty function $P_1(x, y, w)$ used in the merit filter are based on the maximum absolute multipliers over *all* local optima for *all* vectors y_i , because these weights are large enough to ensure that this function is exact for all problems $NLP(y)$.

Therefore, in stage 2 of the algorithm, the exact penalty function, $P_1(x_i, y_i, w)$, is calculated at each trial point (x_i, y_i) , and L is started at (x_i, y_i) if P_1 is smaller than the current threshold for $NLP(y_i)$. This threshold is initialized to $+\infty$, so if the values y_i have not occurred in a previous stage 2 trial point, L will be called at this point. This leads to many more local solver calls in problems with discrete variables, as we show later in the computational-results sections.

In mode (2), Optquest presents candidate y vectors only to L , which are fixed while L finds corresponding (locally) optimal x values. The starting values for x can be chosen to minimize computational effort. We are experimenting with an option that obtains all possible trial points for the current population, sorts them in terms of their distance from each other, and calls L in that sorted order, starting each call of L from the previous optimum. It is expected that y 's that are close to one another will have x 's with that property, so the previous optimum will be a good starting point.

In mode (2), there is no stage 1, and L must be called at each y vector that has not been produced previously. As a result, the local solver call where the best value is found typically comes later than with mode (1). On the other hand, OptQuest's effort is reduced since it processes a much smaller problem, and the information returned to it by the local solver (the optimal objective value over the continuous variables) is of much higher quality than in the base case (the penalized objective value at OptQuest's trial point).

An important option involves the return of information from the local solver to OptQuest, which is absent in the above procedure, i.e., local solutions found by the local solver are not returned to OptQuest. Such solutions are generally of very high quality, and might aid the search process if they were incorporated into the OptQuest population, because at least a subset would likely be retained there. However, this should be done so as to preserve the diversity of the population.

4. Computational Results

The algorithm described in the previous section has been implemented as a callable C-language function. In this form, the user supplies a C function that evaluates the objective and constraint functions, an optional routine that evaluates their first partial derivatives (finite-difference approximations are used otherwise), and a calling program that supplies problem size, bounds, an initial point, and invokes the algorithm. Algorithm parameters and options all are set initially to default values, and any changes are specified in an options file. The local NLP solver is the LSGRG2 implementation of Smith and Lasdon (1992). We have also developed an interface between our C implementation and the GAMS algebraic modeling language (<http://www.gams.com/>), using C library routines provided by GAMS Development Corporation. The user function routine is replaced by one that calls the GAMS interpreter, and a special derivative routine accesses and evaluates expressions developed by GAMS for first derivatives of all nonlinear problem

Table 1 Floudas et al. (1999) Test Problem Set Characteristics

Series	Problems	Max. variables	Max. discrete variables	Max. linear constraints	Max. nonlinear constraints	Problem type
EX2_1_x	14	24	0	10	0	Concave QP (min)
EX3_1_x	4	8	0	4	6	Quadratic obj and constraints
EX4_1_x	9	2	0	0	2	Obj or constraints polynomial
EX5_2_x	2	32	0	8	11	Bilinear-pooling
EX5_3_x	2	62	0	19	34	Distillation column sequencing
EX5_4_x	3	27	0	13	6	Heat exchanger network
EX6_1_x	4	12	0	3	6	Gibbs free energy min
EX6_2_x	10	9	0	3	0	Gibbs free energy min
EX7_2_x	4	8	0	3	12	Generalized geometric prog
EX7_3_x	6	17	0	10	11	Robust stability analysis
EX8_1_x	8	6	0	0	5	Small unconstrained, constrained
EX8_2_x	5	55	0	6	75	Batch plant design-uncertainty
EX8_3_x	14	141	0	43	65	Reactor network synthesis
EX8_4_x	8	62	0	0	40	Constrained least squares
EX8_5_x	6	6	0	2	2	Min tangent plane distance
EX8_6_1	N	$3N - 6$	0	0	0	Lennard-Jones energy min
EX8_6_2	N	$3N - 6$	0	0	0	Morse energy min
EX9_1_x	10	29	6	27	5	Bilevel LP
EX9_2_x	9	16	3	11	6	Bilevel QP
EX12_2_x	6	11	8	9	4	MINLP
EX14_1_x	9	10	0	4	17	Infinity norm solution of equations
EX14_2_x	9	7	0	1	10	Infinity norm solution of equations
Total:	$142 + 2N$					

functions. GAMS identifies all linear terms in each function, and supplies their coefficients separately, thus identifying all linear constraints. This enables us to invoke the OptQuest option that maps each trial point into a point that satisfies the linear constraints. The derivative information supplied by GAMS significantly enhances the performance of the local solver since only nonconstant derivatives are re-evaluated, and these are always available to full machine precision. As mentioned earlier, this GAMS version can call any GAMS NLP solver.

For our computational experiments we used the large set of global optimization test problems coded in GAMS from Floudas et al. (1999). Table 1 shows the characteristics of 142 individual and two groups of problems.

Most problems arise from chemical engineering, but some are from other sources. Most are small, but a few have over 100 variables and a comparable numbers of constraints, and 13 have both continuous and discrete variables. Almost all of the problems without discrete variables have local solutions distinct from the global solution, and the majority of problems have constraints. Sometimes all constraints are linear, as with the concave quadratic programs of series EX2_1_x, but many problems have nonlinear constraints, and these are often the source of the nonconvexities. The best known objective value and (in most cases) the corresponding variable values are provided in Floudas et al. (1999). The symbol N in the

rows for the series EX8_6_1 and EX8_6_2 is the number of particles in a cluster whose equilibrium configuration is sought via potential-energy minimization. (For more details on these problems see Section 4.2.)

4.1. Continuous Variables—The Base Case

This section describes results obtained when OQNLP is applied to 128 of the problems in the Floudas et al. (1999) test set with no discrete variables. A few problems for which no GAMS NLP solver can find a feasible solution in 800 solver calls are omitted. Computations were performed on a DELL OptiPlex PC with a 1.2 GHz Pentium IV processor and 261 Mbytes of RAM, running under Windows 2000.

The options and main algorithm parameters used are shown in Table 2 (see Section 3.4 for definitions). The filter parameter values (*waitcycle*, *threshfactor*, *distfactor*) correspond to fairly tight filters, and these must be loosened to solve some problems. The OptQuest “use linear constraints” option, which projects trial points onto the linear constraints, is not used because it is very time-consuming for larger problems. SNOPT, an SQP implementation, was used for the largest problems because many calls to the GRG solvers CONOPT (Drud 1994) and LSGRG2 terminate infeasible on these problems. The 8_3_x problems include many “pooling” constraints, which have bilinear terms. In almost all these terminations, the GRG solvers find a local minimum of the phase 1 objective. SNOPT has no phase 1, and never terminates infeasible.

Table 2 Base Case OQNLP and OptQuestGRG Parameters and Options Used

Total iterations = 1,000
Stage 1 iterations = 200
Waitcycle = 20
Threshfactor = 0.2
Distfactor = 0.75
Use linear constraints = no
OptQuest search type = <i>boundary</i>
Boundary search parameter = 0.5
NLP solver = LSGRG2, except SNOPT for 110 to 141 range

Table 3 shows outcomes and average-effort statistics for 128 of the Floudas et al. (1999) collection of test problems with continuous variables only, sorted into six groups by number of variables. Geometric rather than arithmetic means are used to reduce the effects of outliers: function calls, iterations, and times for the larger problem sets typically include a few problems with values much larger than all others. Computational effort is measured by OptQuest iterations, solver calls, function calls (each function call evaluates the objective and all constraint functions), and computation time. The three “to best” columns show the effort required to find the best OQNLP objective value. Function calls are not available for the largest problems because the SNOPT interface does not yet make them available.

Since all problems have known solutions, we define “failures” as problems with a solution gap of more than 1%. This gap is the percentage difference between the best feasible OQNLP objective value, f_{OQNLP} , and the best known feasible objective value f_{best} , defined for minimization problems as $gap = 100(f_{\text{OQNLP}} - f_{\text{best}}) / (1 + |f_{\text{best}}|)$, and the negative of the above for maximization, so positive gaps indicate that the best known solution was not reached. Nine of the 128 problems failed to achieve gaps smaller than 1%, with gaps ranging from 2.2% to 80%. All these are solved with more iterations or by loosening the filters. Percentage gaps for almost all 119 “solved” problems are less than 10^{-4} , and the largest gap among solved problems is 0.37%.

Computational effort needed to achieve these results is quite low, and increases slowly with problem size, except for the geometric mean solution time for the largest problems. The best OQNLP value is also found very early: in the first solver call in 80 of the 118 solved problems, and the second call in 19 more. This shows that, for these test problems, stage 1 of OQNLP is very effective in finding a point in the basin of attraction of the global optimum. The ratio of the “to best” effort to total effort is also small. For iterations, since there are always 200 stage 1 iterations, we subtract 200 before computing the ratio, giving $51.3/800 = 0.06$. The solver call ratio is 0.17 and the time ratio 0.23. This implies that, for these problems, a criterion that stops OQNLP when the fractional change in the best feasible objective value found thus far is below a small tolerance for some (reasonably large) number of successive iterations would rarely terminate the algorithm before the best solution was found. The ratio of total solver calls to locals found, a measure of filter efficiency, varies from 3 to 5, and is nearly 1 for the largest problems.

Table 4 shows results obtained in solving the nine “failed” problems with looser filters and an OptQuest boundary search parameter of 1. Seven of these nine, the 2_1_x series, are quadratic programs (QPs) with concave objectives (to be minimized), so each has an optimal extreme-point solution, and every extreme point is a local solution. The base case and new parameter values are in Table 5.

The looser merit filter increases its threshold every ten iterations, replacing the old value by $old\ value + 1.0 \times (1 + |old\ value|)$. The looser distance filter rejects a trial solution if its distance from any previously found local solution is less than $0.1 \times maxdist$, where $maxdist$ is the largest distance traveled to reach that solution. A search parameter of 1 causes more OptQuest trial points to have values on the boundary of the rectangle defined by the variable bounds, which helps solve the seven concave QPs.

Eight of the nine “unsolved” problems are solved with these new parameters, and the other, EX14_1_8, achieves a gap of 1.15%. It is solved by using 1,000 stage 1 iterations and 5,000 total, with all other parameters as in the base case.

Table 3 Results and Effort Statistics for 128 Continuous-Variable Floudas et al. (1999) Problems

Variable range	Problems	Variables	Constraints	Iterations to best	Solver calls to	Solver calls	Locals found	Function calls to best	Function calls	Time to best	Total time	Failed	First L call	Second L call
1 to 4	32	2.5	1.7	206.3	1.1	7.5	1.9	263.9	2,158.0	0.1	0.5	1	27	3
4 to 7	31	5.5	5.7	214.4	1.3	5.8	1.9	381.5	4,766.7	0.2	0.6	0	22	6
8 to 12	21	9.4	8.1	238.2	1.5	13.2	3.0	575.6	19,698.0	0.1	0.8	3	10	4
13 to 20	18	15.9	11.7	303.5	2.4	7.4	2.6	968.2	5,211.9	0.3	0.7	4	8	1
22 to 78	13	37.9	27.6	259.6	2.1	14.1	3.1	1,562.4	23,077.9	0.6	2.5	1	5	3
110 to 141	13	116.4	80.1	305.0	2.7	23.7	22.7	NA	NA	6.6	64.1	0	7	2
Total/avg.	128			251.3	1.8	10.5	3.5	537.9	7,190.9	0.4	1.7	9	80	19

Table 4 Solving Nine “Failed” Problems with Looser Filters and Boundary Parameter = 1

Problem name	Variables	Constraints	Iterations to best	Solver calls to best	Total solver calls	Base-case solver calls	Locals found	Function calls to best	Total function calls	Time to best	Total time	Gap
EX14_1_8	3	4	338	21	122	17	2	38,141	1,218,680	0.49	2.25	1.15
EX9_2_5	8	7	344	2	6	45	3	690	2,832	0.21	0.6	0
EX2_1_6	10	5	203	3	107	7	17	389	138,749	0.2	1.38	0
EX2_1_9	10	1	201	1	111	66	39	252	342,345	0.16	1.45	0
EX2_1_7_1	20	10	279	8	27	2	19	5,051	60,043	0.39	1.25	0
EX2_1_7_3	20	10	269	4	65	22	36	1,321	286,213	0.3	1.92	0
EX2_1_7_4	20	10	253	6	29	4	9	3,676	83,989	0.39	1.41	0
EX2_1_7_5	20	10	254	5	29	3	15	2,620	71,004	0.33	1.34	0
EX2_1_8	24	10	226	3	120	8	25	981	1,071,730	0.29	2.87	0
Means(geom)			258.6	4.1	48.7	10.5	12.8	1,760.7	137,876.8	0.3	1.5	

Table 5 Base Case and Loosened Filter Parameter Values

Parameter	Base-case value	Looser value
<i>Waitcycle</i>	20	10
<i>Threshold_factor</i>	0.2	1.0
<i>Distance_factor</i>	0.75	0.1
<i>Boundary search parameter</i>	0.5	1.0

Comparing the “total solver calls” and “base-case solver calls” columns shows that the new parameters represent a substantial loosening of both filters. The looser filters result in many more solver calls in all but problem 9_2_5, and the geometric mean solver calls is 48.7 with the loose filters vs. 10.5 with the tighter ones. The behavior of 9_2_5 is surprising (six solver calls with loose filters versus 45 with tighter ones), but the run with looser filters finds the global minimum at iteration 344, and after that its merit thresholds and set of local solutions differ from those of the base-case run.

Table 6 shows the geometric performance means and totals obtained from solving the 14 concave QP problems with base-case parameters, with and without the OptQuest “use linear constraints” option, which maps each trial point into a nearest point feasible for the linear constraints. Since these are linearly constrained problems, invoking this option guarantees that all trial points are feasible.

Clearly, this option helps: there are roughly twice as many solver calls on average when using it, and only two failures, vs. seven in the base case. The gaps for the two unsolved problems (2_1_7_5 and 2_1_9) are between 1% and 3.5% in both cases. However, this option increases run times here by about a factor of 30, so it is currently off by default.

4.2. The Lennard-Jones and Morse Energy-Minimization Problems

The Floudas et al. (1999) set of test problems includes two GAMS models that choose the locations of a cluster of N particles to minimize the potential energy of the cluster, using two different potential energy functions, called Lennard-Jones and Morse. The decision variables are the (x, y, z) coordinates of each particle. Particle 1 is located at the origin, and three position components of particles 2 and 3 are fixed, so each family of problems has $3N - 6$ variables. These problems have many local minima, and their number increases rapidly with problem size, so they constitute a good test for global optimization algorithms.

Results of applying OQNLP to 14 of these problems, using 200 stage 1 and 1,000 total iterations, are shown in Tables 7 and 8. Each problem set was solved with LSGRG2 and CONOPT. These results use CONOPT for the Lennard-Jones problems and LSGRG2 for the Morse, because they provide slightly better results, illustrating the value of being able to call several solvers. Because of the many distinct local minima, the number of local minima found is equal to the number of solver calls for the three largest Lennard-Jones problems and for all the Morse problems.

The Lennard-Jones problems are the more difficult of the two. The three largest problems have gaps of roughly 1% to 2%, using the looser filter parameters in Table 5. The default filter parameters led to positive gaps for the last three problems totaling 7.8%, while this sum in Table 7 is 3.8%. The objective approaches infinity as the distance between any two particles approaches zero, so its unconstrained minimization for $N = 20$ leads to about 50,000 domain violations

Table 6 Solving Concave QP Problems With and Without “Use Linear Constraints”

Case	Iterations to best	Solver calls to best	Total solver calls	Locals found	Fcn calls to best	Total fcn calls	Time to best	Total time	Failed
No use	284.8	2.3	6.6	3.7	643.8	3,875.1	0.3	0.6	7
Use	247.1	2.1	12.1	3.1	437.7	3,827.6	6.9	19.0	2

Table 7 Solving Six Lennard-Jones Problems Using CONOPT and Loose Filters

Problem name	Variables	Constraints	Solver calls to best	Total solver calls	Locals found	Time to best	Total	Gap, %
EX8_6_1_5	9	10	1	152	39	1.09	21.83	0.00
EX8_6_1_10	24	45	21	130	114	18.56	68.34	0.00
EX8_6_1_15	39	105	6	104	100	13.63	165.09	0.00
EX8_6_1_20	54	190	67	118	118	257.62	396.21	1.12
EX8_6_1_25	69	300	42	94	94	325.82	730.68	1.84
EX8_6_1_30	84	435	16	59	59	134.35	434.56	0.88

Table 8 Solving Eight Morse Problems Using LSGRG2 and Default Parameters

Problem name	Variables	Solver calls to best	Total solver calls	Locals found	Time to best	Total time	Gap
EX8_6_2_5	9	1	5	5	0.23	0.61	0.0000
EX8_6_2_10	24	1	15	15	0.57	4.44	0.0000
EX8_6_2_15	39	1	6	6	1.41	6.43	0.0000
EX8_6_2_20	54	2	43	43	4.20	51.20	0.0000
EX8_6_2_25	69	4	20	20	13.44	58.38	0.0000
EX8_6_2_30	84	17	43	43	68.56	160.19	0.0000
EX8_6_2_40	114	7	33	33	66.29	273.91	0.0000
EX8_6_2_50	144	20	25	25	337.20	403.96	0.1251

(either divide by zero or integer power overflow), and this number grows rapidly with N . Hence we added constraints lower bounding this distance by 0.1 for all distinct pairs of points, and the number of these constraints is shown in the table. None are active at the best solution found.

Table 8 shows the Morse potential results using the LSGRG2 solver and the default OQNLP parameters shown in Table 2. The objective here has no singularities, so there are no difficulties with domain violations, and the only constraints are variable bounds. All problems are solved to very small gaps except the largest (144 variables), which has a gap of 0.125%. The number of solver calls is much smaller than for the Lennard-Jones problems, because the filters are much tighter. Each call leads to a different local optimum. The largest problem is solved to a gap less than $10^{-4}\%$ with 5,000 total and 1,000 stage 1 iterations and the same filter parameters. This run terminated because the 3,000-second time limit was exceeded, took 4,083 iterations, and found 210 distinct local optima in 210 solver calls, compared to only 25 in the base case.

4.3. Problems with Discrete Variables

There are 11 MINLP problems in the Floudas et al. (1999) test set, with the total number of variables ranging from three to 29 and the number of binary variables ranging from one to eight. Two of these, EX12_2_3 and EX12_2_4, had been reformulated so that all binaries appeared linearly, and we restored them to their original state where the binaries appear nonlinearly. OQNLP allows such representations, while the other GAMS MINLP solvers do not. The final test set contains 13 problems. These are far too small to yield meaningful inferences about the power of OQNLP on problems of practical size, but allow preliminary testing of the two MINLP modes described in Section 3.5. The geometric means of some measures of computational outcomes and effort for both modes are shown in Table 9, using the LSGRG2 NLP solver.

The first table row is for mode (2), where OptQuest manipulates only the discrete variables. Each NLP problem was “cold started” from the same initial point in these runs, so the number of function calls could be reduced substantially by warm starts. All runs are terminated by OptQuest after the small number of possible binary-variable combinations have been completely enumerated. The optimal solution is found on average about midway through the solution process, but we expect that this will occur earlier as the number of discrete variables increases. The OptQuest logic requires that at least one population of binary solutions be evaluated before any learning can occur, and the average number of solver calls to find the best solution here is about equal to the population size of ten.

The last three rows of Table 9 show results for mode (1), where OptQuest manipulates both binary and continuous variables. In rows 2 and 3, we do

Table 9 Solution Statistics for 13 Problems with Discrete Variables

Discrete var mode	Options and parameters	Iterations to best	L calls to best	Total L calls	Fcn calls to best	Total fcn calls	Time to best	Total time	Failures
Discrete only	Default	10.2	10.2	20.1	2,065.8	5,794.1	0.8	1.5	0
All	Default	261.7	4.0	32.1	1,065.7	25,022.2	0.3	0.8	7
All	(1,000, 5,000)	1,551.9	12	89.1	10,196.8	224,025.7	0.8	3.7	1
All	Default, use	272.4	3.6	115.8	1,178.0	88,983.0	9.9	26.1	0

Table 10 Average Solver Calls vs. Number of Binary Variables

Binaries	1	3	4	5	6	8
Problems	1	3	2	1	4	2
Discretes only	2	7	12.5	27	52.5	81
All	52	35.7	17.5	109	195.5	551.5

not require trial points to satisfy linear constraints, while in row 4 we do. Without using linear constraints, the default number of stage 1 and total iterations of (200, 1,000), are not enough to find the best solution for about seven of the 13 problems. This is because many OptQuest trial points have the same values for the binary variables but different values for the continuous variables, so complete enumeration takes far longer than in mode (2). However, 1,000 stage 1 and 5,000 total iterations solve all but one problem (its gap is 9.4%), and the ratio (solver calls to best)/(total solver calls) of about 1/9 is favorably small. Row 4 shows that, if trial points are required to satisfy linear constraints, all problems are solved in 1,000 total iterations. This is because these problems have mostly linear constraints (geometric mean of linear constraints is 9.1 and of total constraints is 9.9), so the projected trial points tend to contain an optimal set of binary variable values earlier, after only 3.6 solver calls on average. However, solving the MILPs, which map trial points into nearest points that satisfy the linear constraints, increases total solution time by a factor of about 30 (compare the times in rows 2 and 4).

Table 10 shows that total solver calls increase quickly with the number of binary variables for the two discrete-variable modes, especially the “all” mode. When more than one problem has the same number of binaries, averages over those problems are given. The values for the “Discretes only” mode are the number of feasible binary vectors, averaged over the number of problems shown.

5. Summary and Future Research

The results of Section 4 show that OQNLP is a promising approach for smooth nonconvex NLPs with continuous variables. It solves all 142 of the test problems with no discrete variables with very reasonable solution effort. While there is no guarantee of optimality and no “gap” is available, it can be combined with other algorithms that provide this information, e.g., LGO (Pinter 2004) or BARON (Tawarmalani and Sahinidis 2002). Information on these solvers is also available at <http://www.gamsworld.org/> (link to *global world* and then to *global solvers*). The lower bounds provided by these procedures can be used to estimate the gap for the OQNLP solution, and the solution produced by any algorithm can be used as a warm start for any other.

Future research includes enhancing the filter logic. As described above, the filters needed to be loosened to solve nine of the Floudas et al. (1999) test problems, and this loosening could be done automatically. The merit filter parameter *threshfactor* (new threshold = *threshfactor* × (1 + old threshold)) could be calculated dynamically. Each time a penalty value is above the threshold, calculate the value of *threshfactor* that would cause the new threshold to equal the penalty value. If this happens for *waitcycle* consecutive iterations, set *threshfactor* to the smallest of these values, so the new threshold would have just accepted the lowest of the penalty values. Similar logic can be developed for the distance filter, reducing a basin radius *maxdist* if that basin’s distance filter rejects trial points for *waitcycle* consecutive iterations.

Also, the current distance-filter logic allows overlap of the spherical approximations to the attraction basins. The true basins can have no points in common, so we can impose this condition on the spheres. If the spherical model basins for any 2 local solutions x_i and x_j have radii r_i and r_j , these must satisfy

$$r_i + r_j \leq d(i, j),$$

where $d(i, j)$ is the Euclidean distance between x_i and x_j . If this inequality is violated, the radii r_i and r_j can be reduced by the same scale factor so that it holds as equality. We plan to test these options soon.

Another important aspect is the communication of locally optimal solutions back to OptQuest, to improve its search process. These solutions usually have substantially lower penalty values than do typical OptQuest trial points, so they are likely ultimately to be included in OptQuest’s population. However, their penalty values often become the merit-filter thresholds, causing most other trial points to be rejected. Also, the local optima and their nearby “children” will be rejected by the distance filter. We have seen these effects in preliminary tests.

NLP algorithms can fail by failing to find a feasible point in cases where the problem instance is feasible. With GRG algorithms, this usually happens when phase 1 terminates at a local optimum of the phase 1 objective. OQNLP can be applied to such problems, if they are reformulated by dropping the true objective, adding deviation variables into all constraints, and minimizing the sum of these deviation variables. This approach could greatly improve the ability of existing NLP solvers to diagnose infeasibility. More generally, OQNLP can improve NLP solver reliability by starting the solver from as many points as desired, while insuring that these points balance diversity and quality.

The performance of OQNLP in solving MINLPs is less clear, because the 13 MINLP test problems used

here are so small. More extensive testing is needed, which should clarify the relative merits of the two MINLP “modes” discussed in Section 4.3. If OptQuest manipulates only the discrete variables, then all trial points generated by the current population may be generated at once, and the solver calls at these points may be done in any order. The points can be sorted by increasing distance from their nearest neighbor, and each NLP call can be started from the previous optimum. The NLPs can also be solved in parallel.

Finally, comparative studies of OQNLP and other global and MINLP solvers are needed. This testing is facilitated by the existing GAMS interfaces for BARON, LGO, DICOPT, and SBB. The “MINLP World” and “Global World” websites developed by GAMS Development Corporation (<http://www.gamsworld.org/>) provide solver information and test problems with known solutions.

Acknowledgments

This research was partially supported by the Ministerio de Educación y Ciencia of Spain under Reference Code TIN2006-02696.

References

- Dixon, L., G. P. Szegő. 1975. Towards global optimization. *Proc. Workshop at the University of Cagliari, Italy*, North Holland, Amsterdam, The Netherlands.
- Drud, A. 1994. CONOPT—A large-scale GRG-code. *ORSA J. Comput.* **6** 207–218.
- Edgar, T. F., D. M. Himmelblau, L. S. Lasdon. 2001. *Optimization of Chemical Processes*. McGraw-Hill, New York.
- Floudas, C. A., P. M. Pardalos, C. S. Adjiman, W. R. Esposito, Z. Gumus, S. T. Harding, J. L. Klepeis, C. A. Meyer, C. A. Schweiger. 1999. *Handbook of Test Problems for Local and Global Optimization*. Kluwer Academic Publishers, Boston, MA.
- Glover, F. 1998. A template for scatter search and path relinking. J.-K. Hao, E. Lutton, E. Ronald, M. Schoenauer, D. Snyers, eds. *Artificial Evolution, Lecture Notes in Computer Science*, Vol. 1363. Springer Verlag, New York, 13–54.
- Laguna, M., R. Marti. 2002. The OptQuest callable library. S. Voss, D. Woodruff, eds. *Optimization Software Class Libraries*. Kluwer Academic Publishers, Boston, MA, 193–218.
- Laguna, M., R. Marti. 2003. *Scatter Search: Methodology and Implementations in C*. Kluwer Academic Publishers, Boston, MA.
- Locatelli, M., F. Schoen. 1999. Random linkage: A family of acceptance/rejection algorithms for global optimization. *Math. Programming* **85** 379–396.
- Murtagh, B. A., M. A. Saunders. 1982. A projected Lagrangian algorithm and its implementation for sparse nonlinear constraints. *Math. Program. Study* **16** 84–117.
- Nash, S. G. 1998. Nonlinear programming. *OR/MS Today* **25** 36–45.
- Nash, S. G., A. Sofer. 1996. *Linear and Nonlinear Programming*. McGraw-Hill, New York.
- Nocedal, J., S. J. Wright. 1999. *Numerical Optimization*. Springer Series in Operations Research, New York.
- Pinter, J. D. 2004. *Computational Global Optimization in Nonlinear Systems. An Interactive Tutorial*. Lionheart Publishing, Atlanta, GA.
- Rinnooy Kan, A. H. G., G. T. Timmer. 1987a. Stochastic global optimization methods; part I: Clustering methods. *Math. Programming* **37** 27–56.
- Rinnooy Kan, A. H. G., G. T. Timmer. 1987b. Stochastic global optimization methods; part II: Multi level methods. *Math. Programming* **37** 57–78.
- Smith, S., L. Lasdon. 1992. Solving large sparse nonlinear programs using GRG. *ORSA J. Comput.* **4** 3–15.
- Tawarmalani, M., N. Sahinidis. 2002. *Convexification and Global Optimization in Continuous and Mixed-Integer Nonlinear Programming*. Kluwer Academic Publishers, Boston, MA.