

Integrating SQP and branch-and-bound for Mixed Integer Nonlinear Programming*

Sven Leyffer^{†‡}

August 6, 1998

Abstract

This paper considers the solution of Mixed Integer Nonlinear Programming (MINLP) problems. Classical methods for the solution of MINLP problems decompose the problem by separating the nonlinear part from the integer part. This approach is largely due to the existence of packaged software for solving Nonlinear Programming (NLP) and Mixed Integer Linear Programming problems.

In contrast, an integrated approach to solving MINLP problems is considered here. This new algorithm is based on branch-and-bound, but does not require the NLP problem at each node to be solved to optimality. Instead, branching is allowed after each iteration of the NLP solver. In this way, the nonlinear part of the MINLP problem is solved whilst searching the tree. The nonlinear solver that is considered in this paper is a Sequential Quadratic Programming solver.

A numerical comparison of the new method with nonlinear branch-and-bound is presented and a factor of about 3 improvement over branch-and-bound is observed.

Keywords: Mixed Integer Nonlinear Programming, Branch-and-Bound, Sequential Quadratic Programming

1 Introduction

This paper considers the solution of Mixed Integer Nonlinear Programming (MINLP) problems. Problems of this type arise when some of the variables are required to take integer values. Applications of MINLP problems include the design of batch plants (e.g. [14] and [18]), the synthesis of processes (e.g. [6]), the design of distillation sequences (e.g. [24]), the optimal positioning of products in a multi-attribute space (e.g. [6]), the minimization of waste in paper cutting [26] and the optimization of core reload patterns for nuclear reactors [22]. For a comprehensive survey of applications of MINLP applications see Grossmann and Kravanja [13].

MINLP problems can be modelled in the following form

$$(P) \begin{cases} \underset{x,y}{\text{minimize}} & f(x,y) \\ \text{subject to} & g(x,y) \leq 0 \\ & x \in X, y \in Y \text{ integer,} \end{cases}$$

*University of Dundee Numerical Analysis Report NA/182

[†]University of Dundee, Department of Mathematics, Dundee, U.K., sleyffer@mcs.dundee.ac.uk.

[‡]Support by the EPSRC under grant number GR/K51204 is gratefully acknowledged.

where y are the integer variables modelling for instance decisions ($y \in \{0, 1\}$), numbers of equipment ($y \in \{0, 1, \dots, N\}$) or discrete sizes and x are the continuous variables.

Classical methods for the solution of (P) “decompose” the problem by separating the nonlinear part from the integer part. This approach is largely due to the existence of packaged software for solving Nonlinear Programming (NLP) and Mixed Integer Linear Programming (MILP) problems. For instance, branch-and-bound [5] performs a tree-search (integer part) solving an NLP problem at each node (nonlinear part). The decomposition is even more apparent in Outer Approximation ([6] and [8]) and Benders Decomposition ([12] and [11]) where an alternating sequence of MILP master problems and NLP subproblems obtained by fixing the integer variables is solved. The recent branch-and-cut method for 0-1 convex MINLP of Stubbs and Mehrotra [23] also separates the nonlinear (lower bounding and cut generation) and integer part of the problem.

In contrast to these approaches, the present paper considers an integrated approach to MINLP problems. The algorithm developed here is based on branch-and-bound, but instead of solving an NLP problem at each node of the tree, the tree-search and the iterative solution of the NLP are interlaced. Thus the nonlinear part of (P) is solved whilst searching the tree. The nonlinear solver that is considered in this paper is a Sequential Quadratic Programming (SQP) solver. SQP solves an NLP problem via a sequence of quadratic programming (QP) problem approximations.

The basic idea underlying the new approach is to branch early – possibly after a single QP iteration of the SQP solver. This idea was first presented by Borchers and Mitchell [1]. The present paper improves on their idea in a number of ways:

1. In order to derive lower bounds needed in branch-and-bound Borchers and Mitchell propose to evaluate the Lagrangian dual. This is effectively an unconstrained nonlinear optimization problem. In Section 3 it is shown that there is no need to compute a lower bound if the linearizations in the SQP solver are interpreted as cutting planes.
2. In [1] *two* QP problems are solved before branching. This restriction is due to the fact that a packaged SQP solver is used. By using our own solver we are able to remove this unnecessary restriction, widening the scope for early branching.
3. A new heuristic for deciding when to branch early is introduced which does not rely on a *fixed* parameter, but takes the second order rate of convergence of the SQP solver into account when deciding whether to branch or continue with the SQP iteration.

The ideas presented in this paper have a similar motivation as a recent algorithm by Quesada and Grossmann [21]. Their algorithm is related to outer approximation but avoids the resolution of related MILP master problems by interrupting the MILP branch-and-bound solver each time an integer node is encountered. At this node an NLP problem is solved (obtained by fixing all integer variables) and new outer approximations are added to all problems on the MILP branch-and-bound tree. Thus the MILP is updated and the tree-search resumes. The difference to the approach considered here is that Quesada and Grossmann still solve NLP problems at some nodes, whereas the new solver presented here usually solves one quadratic programming (QP) problem at each node.

The paper is organized as follows. Section 2 briefly reviews the SQP method and branch-and-bound. In Section 3 the new algorithm is developed for the special case of convex MINLP problems. Section 3.4 introduces a heuristic for handling nonconvex MINLP problems. Finally,

Section 4 presents a comparison of the new method with nonlinear branch-and-bound. These results are very encouraging, often improving on branch-and-bound by a factor of up to 3.

2 Background

2.1 Sequential Quadratic Programming

A popular iterative method for solving NLP problems is Sequential Quadratic Programming (SQP). The basic form of SQP methods date back to Wilson [27] and were popularized by Han [16] and Powell [20], see Fletcher [7, Chapter 12.4] and Conn, Gould and Toint [4] for a selection of other references. In its basic form, SQP solves an NLP problem via a sequence of quadratic programming (QP) approximations obtained by replacing the nonlinear constraints by a linear first order Taylor series approximation and the nonlinear objective by a second order Taylor series approximation augmented by second order information from the constraints. It can be shown under certain conditions that the SQP method converges quadratically near a solution.

It is well known that the SQP method may fail to converge if it is started far from a local solution. In order to induce convergence, many popular methods use a *penalty function* which is a linear combination of the objective function f and some measure of the constraint violation. A related idea is an *augmented Lagrangian function* in which a weighted penalty term is added to a Lagrangian function. A step in an SQP method is then accepted when it produces a sufficient decrease in the penalty function.

Two frameworks exist which enforce sufficient decrease, namely line-search in the direction of the QP solution or a trust-region that limits the QP step that is computed (see e.g. [7, Chapter 12.4] and references therein). In our implementation global convergence is promoted through the use of a trust-region and the new concept of a “filter” [9] which accepts a trial point whenever the objective or the constraint violation is improved compared to all previous iterates. The size of the trust-region is reduced if the step is rejected and increased if it is accepted.

2.2 Branch-and-bound

Branch-and-bound dates back to Land and Doig [19]. The first reference to nonlinear branch-and-bound can be found in Dakin [5]. It is most conveniently explained in terms of a tree-search.

Initially, all integer restrictions are relaxed and the resulting NLP relaxation is solved. If all integer variables take an integer value at the solution then this solution also solves the MINLP. Usually, some integer variables take a non-integer value. The algorithm then selects one of those integer variables which take a non-integer value, say y_i with value \hat{y}_i , and branches on it. Branching generates two new NLP problems by adding simple bounds $y_i \leq [\hat{y}_i]$ and $y_i \geq [\hat{y}_i] + 1$ respectively to the NLP relaxation (where $[a]$ is the largest integer not greater than a).

One of the two new NLP problems is selected and solved next. If the integer variables take non-integer values then branching is repeated, thus generating a branch-and-bound tree whose nodes correspond to NLP problems and where an edge indicates the addition of a branching bound. If one of the following fathoming rules is satisfied, then no branching is required, the corresponding node has been fully explored (fathomed) and can be abandoned. The fathoming rules are

- FR1** An infeasible node is detected. In this case the whole subtree starting at this node is infeasible and the node has been fathomed.
- FR2** An integer feasible node is detected. This provides an upper bound on the optimum of the MINLP; no branching is possible and the node has been fathomed.
- FR3** A lower bound on the NLP solution of a node is greater or equal than the current upper bound. In this case the node is fathomed, since this NLP solution provides a lower bound for all problems in the corresponding sub-tree.

Once a node has been fathomed the algorithm backtracks to another node which has not been fathomed until all nodes are fathomed.

Many heuristics exist for selecting a branching variable and for choosing the next problem to be solved after a node has been fathomed (see surveys by Gupta and Ravindran [15] and Volkowitz et.al. [25]).

Branch-and-bound can be inefficient in practice since it requires the solution of one NLP problem per node which is usually solved iteratively through a sequence of quadratic programming problems. Moreover, a large number of NLP problems are solved which have often no physical meaning if the integer variables do not take integer values.

3 Integrating SQP and branch-and-bound

An alternative to nonlinear branch-and-bound for *convex* MINLP problems is due to Borchers and Mitchell [1]. They observe that it is not necessary to solve the NLP at each node to optimality before branching and propose an *early branching rule*, which branches on an integer variable before the NLP has converged. Borchers and Mitchell implement this approach and report encouraging numerical results compared to Outer Approximation [2]. In their implementation the SQP solver is interrupted after two QP solves¹ and branching is done on an integer variable which is more than a given tolerance away from integrality.

The drawback of the *early branching rule* is that since the NLP problems are *not* solved to optimality, there is no guarantee that the current value of $f(x, y)$ is a lower bound. As a consequence, bounding (fathoming rule **FR3**) is no longer applicable. Borchers and Mitchell seek to overcome this difficulty by evaluating the Lagrangian dual for a given set of multiplier estimates, λ . The evaluation of the Lagrangian dual amounts to solving

$$\begin{cases} \underset{x, y}{\text{minimize}} & f(x, y) + \lambda^T g(x, y) \\ \text{subject to} & x \in X, y \in \hat{Y} \end{cases}$$

where the $\hat{Y} \subset Y$ now also includes bounds that have been added during the branching.

Note that this evaluation requires the solution of a bound constrained nonlinear optimization problem. In order to diminish the costs of these additional calculations, Borchers and Mitchell only evaluate the dual every sixth QP solve.

Clearly, it would be desirable to avoid the solution of this problem and at the same time obtain a lower bounding property at each QP solve. In the remainder of this section it is shown how this can be achieved by integrating the iterative NLP solver with the tree search.

¹Ideally one would prefer to interrupt the SQP method after each QP solve. Borchers and Mitchell use an SQP method from the NAG library which does not allow this.

Throughout this section it is assumed that both f and g are smooth, *convex* functions. The nonconvex case is discussed in Section 3.4 where a simple heuristic is proposed. The ideas are first presented in the context of a basic SQP method *without* globalization strategy. Next this basic algorithm is modified to allow for the use of a line-search or a trust-region. It is then shown how the early branching rule can take account of the quadratic rate of convergence of the SQP method. Finally, a simple heuristic for nonconvex MINLP problems is suggested.

3.1 The basic SQP algorithm

This section shows how the solution of the Lagrangian dual can be avoided by interpreting the constraints of the QP approximation as supporting hyperplanes. If f and g are convex, then the linear approximations of the QP are outer approximations and this property is exploited to replace the lower bounding.

Consider an NLP problem at a given node of the branch-and-bound tree,

$$(\hat{P}) \begin{cases} \underset{x,y}{\text{minimize}} & f(x,y) \\ \text{subject to} & g(x,y) \leq 0 \\ & x \in X, y \in \hat{Y}, \end{cases}$$

where the integer restrictions have been relaxed and $\hat{Y} \subset Y$ contains bounds that have been added by the branching. Let \hat{f} denote the solution of (\hat{P}) .

Applying the SQP method to (\hat{P}) results in solving a sequence of QP problems of the form

$$(QP^k) \begin{cases} \underset{d}{\text{minimize}} & f^{(k)} + \nabla f^{(k)T} d + \frac{1}{2} d^T W^{(k)} d \\ \text{subject to} & g^{(k)} + \nabla g^{(k)T} d \leq 0 \\ & x^k + d_x \in X, y^k + d_y \in \hat{Y}. \end{cases}$$

where $f^{(k)} = f(x^{(k)}, y^{(k)})$ etc. and

$$W^{(k)} \simeq \nabla^2 \mathcal{L}^{(k)} = \nabla^2 f^{(k)} + \sum \lambda_i \nabla^2 g_i^{(k)}$$

approximates the Hessian of the Lagrangian.

Unfortunately, the solution of (QP^k) does not underestimate \hat{f} , even if all functions are assumed to be convex. This implies that fathoming rule **FR3** cannot be used if early branching is employed. However, it turns out that it is not necessary to compute an underestimator explicitly. Instead, a mechanism is needed of terminating the sequence of (QP^k) problems once such an underestimator would become greater than the current upper bound, U , of the branch-and-bound process. This can be achieved by adding the surrogate constraint

$$f^{(k)} + \nabla f^{(k)T} d \leq U - \epsilon$$

to (QP^k) , where $\epsilon > 0$ is the optimality tolerance of branch-and-bound. Denote the new QP with this surrogate constraint by (QP_s^k) .

$$(QP_s^k) \begin{cases} \underset{d}{\text{minimize}} & f^{(k)} + \nabla f^{(k)T} d + \frac{1}{2} d^T W^{(k)} d \\ \text{subject to} & g^{(k)} + \nabla g^{(k)T} d \leq 0 \\ & f^{(k)} + \nabla f^{(k)T} d \leq U - \epsilon \\ & x^k + d_x \in X, y^k + d_y \in \hat{Y}. \end{cases}$$

Note that (QP_s^k) is the QP that is generated by the SQP method when applied to the following NLP problem

$$\begin{cases} \text{minimize}_{x,y} & f(x,y) \\ \text{subject to} & g(x,y) \leq 0 \\ & f(x,y) \leq U - \epsilon \\ & x \in X, y \in \hat{Y}, \end{cases}$$

where a nonlinear surrogate constraint has been added to (\hat{P}) . This NLP problem is infeasible, if fathoming rule **FR3** holds at the solution to (\hat{P}) .

It is now possible to replace fathoming rule **FR3** applied to (\hat{P}) by a test for feasibility of (QP_s^k) . This results in the following lemma.

Lemma 1 *Let f and g be smooth, convex functions. A sufficient condition for fathoming rule **FR3** applied to (\hat{P}) to be satisfied is that any QP problem (QP_s^k) generated by the SQP method in solving (\hat{P}) is infeasible.*

Proof: If (QP_s^k) is infeasible, then it follows that there exists no step d such that

$$f^{(k)} + \nabla f^{(k)T} d \leq U - \epsilon \quad (1)$$

$$g^{(k)} + \nabla g^{(k)T} d \leq 0. \quad (2)$$

Since (2) is an outer approximation of the nonlinear constraints of (\hat{P}) and (1) underestimates f , it follows that there exists no point $(\hat{x}, \hat{y}) \in X \times \hat{Y}$ such that

$$f(\hat{x}, \hat{y}) \leq U - \epsilon.$$

Thus any lower bound on f at the present node has to be larger than $U - \epsilon$. Thus to within the tolerance ϵ , $f \geq U$ and fathoming rule **FR3** holds. \square

In practice, an SQP method would use a primal active set QP solver which establishes feasibility first and then optimizes. Thus the test of Lemma 1 comes at no extra cost. The basic integrated SQP and branch-and-bound algorithm can now be stated.

Algorithm 1: Basic SQP and branch-and-bound

Initialization: Place the continuous relaxation of (P) on the tree and set $U = \infty$.

while (there are pending nodes in the tree) **do**

Select an unexplored node.

repeat (SQP iteration)

1. Solve (QP_s^k) for a step $d^{(k)}$ of the SQP method.

2. **if** $((QP_s^k)$ infeasible) **then** fathom node, **exit**

3. Set $(x^{(k+1)}, y^{(k+1)}) = (x^{(k)}, y^{(k)}) + (d_x^{(k)}, d_y^{(k)})$.

4. **if** $((x^{(k+1)}, y^{(k+1)})$ NLP optimal) **then**

if $(y^{(k+1)})$ integral) **then**

Update current best point by setting

$(x^*, y^*) = (x^{(k+1)}, y^{(k+1)})$, $f^* = f^{(k+1)}$ and $U = f^*$.

else

Choose a non-integral $y_i^{(k+1)}$ and branch.

endif

endif

```

    exit
  endif
5. Compute the integrality gap  $\theta = \max_i |y_i^{(k+1)} - \text{anint}(y_i^{(k+1)})|$ .
6. if ( $\theta > \tau$ ) then
    Choose a non-integral  $y_i^{(k+1)}$  and branch, exit
  endif
end while

```

Here $\text{anint}(y)$ is the nearest integer to y . Step 2. implements the fathoming rule of Lemma 1 and step 6. is the early branching rule. In [1] a value of $\tau = 0.1$ is suggested for the early branching rule and this value has also been chosen here.

A convergence proof for this algorithm is given in the next section where it is shown how the algorithm has to be modified if a line-search or a trust-region is used to enforce global convergence for SQP. The key idea in the convergence proof is that (as in nonlinear branch-and-bound), the union of the child problems that are generated by branching is equivalent to the parent problem. The fact, that only a single QP has been solved in the parent problem does not alter this equivalence. Finally, the new fathoming rule implies **FR3** and hence, any node that has been fathomed by Algorithm 1 can also be considered as having been fathomed by nonlinear branch-and-bound. Thus convergence follows from the convergence of branch-and-bound.

Algorithm 1 has two important advantages over the work in [1]. Firstly, the lower bounding can be implemented at no additional cost (compared to the need to solve the Lagrangian dual in [1]). Secondly, the lower bounding is available at *all* nodes of the branch-and-bound tree (rather than at every sixth node).

3.2 The globalized SQP algorithm

It is well known that the basic SQP method may fail to converge if started far from a solution. In order to obtain a globally convergent SQP method, descent in some merit function has to be enforced. There are two concepts which enforce descent: a line-search and a trust-region. This section describes how the basic algorithm of the previous section has to be modified to accommodate these global features.

The use of a line-search requires only to replace $d^{(k)}$ by $\alpha^{(k)}d^{(k)}$, where $\alpha^{(k)}$ is the step-length chosen in the line-search. Similarly a Levenberg-Marquardt approach would require minimal change to the algorithm.

An alternative to using a line-search is to use a trust region. Trust-region SQP methods usually possess stronger convergence properties and this leads us to explore this approach. In this type of approach the step that is computed in (QP_s^k) is restricted to lie within a trust-region. Thus the QP that is being solved at each iteration becomes

$$(QP_{s,\infty}^k) \left\{ \begin{array}{l} \underset{d}{\text{minimize}} \quad f^{(k)} + \nabla f^{(k)T} d + \frac{1}{2} d^T W^{(k)} d \\ \text{subject to} \quad g^{(k)} + \nabla g^{(k)T} d \leq 0 \\ \quad \quad \quad f^{(k)} + \nabla f^{(k)T} d \leq U - \epsilon \\ \quad \quad \quad \|d\|_\infty \leq \rho^k \\ \quad \quad \quad x^k + d_x \in X, \quad y^k + d_y \in \hat{Y} \end{array} \right.$$

where ρ^k is the trust-region radius which is adjusted according to how well the quadratic model of $(QP_{s,\infty}^k)$ agrees with the true function (see [7, Chapter 14.5] for details).

The drawback of using a trust-region in the present context is that the trust-region may cause $(QP_{s,\infty}^k)$ to be infeasible even though the problem without a trust-region (QP_s^k) has a non-empty feasible region. Thus Lemma 1 is no longer valid. However, if $(QP_{s,\infty}^k)$ is infeasible and the trust-region is inactive, then the argument of Lemma 1 still holds and the node can be fathomed.

Otherwise, the *feasibility restoration phase* has to be entered and this part of the SQP method is briefly described next. The aim of the restoration phase is to get closer to the feasible region by minimizing the violation of the constraints $\|g^+(x, y)\|$, where $g_i^+(x, y) = \max(0, g_i(x, y))$ for any norm, subject to the linear constraints $x \in X$, $y \in \hat{Y}$. In this phase, an SQP algorithm is applied to minimize the constraint violation. One possible outcome of the restoration phase is that a nonzero minimum of $\|g^+(x, y)\|$ is found (case (d) in Algorithm 2). This is taken as an indication that (\hat{P}) has no feasible point with an objective that is lower than the current upper bound and the corresponding node can be fathomed. The other possible outcome of the restoration phase is that a point is found at which the QP subproblem is feasible (case (c) below). In this case the normal SQP method can be resumed. If $\|d\| < \rho$ during the restoration phase (case (b) below) then the node can be fathomed by an argument similar to Lemma 1.

The modifications required to make Algorithm 1 work for trust-region SQP methods are as follows.

Algorithm 2: Trust-region SQP and branch-and-bound

Only steps 1. and 2. need to be changed to:

- 1'. Compute an acceptable step $d^{(k)}$ of the trust-region SQP method.
- 2'. **if** $((QP_{s,\infty}^k)$ infeasible and $\|d^{(k)}\|_\infty < \rho^k$) **then**
 fathom node **exit**
 elseif $((QP_{s,\infty}^k)$ infeasible and $\|d^{(k)}\|_\infty = \rho^k$) **then**
 Enter feasibility restoration phase:
 repeat (SQP iteration for feasibility restoration)
 (a) Compute a step of the feasibility restoration.
 (b) **if** $(\|d^{(k)}\|_\infty < \rho^k)$ **then** fathom node, **exit**
 (c) **if** $((QP_{s,\infty}^k)$ feasible) **then** return to normal SQP, **exit**
 (d) **if** $(\min \|g^+(x, y)\| > 0)$ **then** fathom node, **exit**
 endif

The following theorem shows that Algorithm 2 converges under standard assumptions (e.g. [6] or [8]).

Theorem 1 *If f and g are smooth and convex functions, if X and Y are convex sets, if Y integer is finite and if the underlying trust-region SQP method is globally convergent, then Algorithm 2 terminates after visiting a finite number of nodes at the solution (x^*, y^*) of (P) .*

Proof: The finiteness of Y integer implies that the branch-and-bound tree is finite. Thus the algorithm must terminate after visiting a finite number of nodes. Nodes are only fathomed if they are infeasible, integer feasible or if the lower bounding of Lemma 1 holds. Thus an optimal node must eventually be solved and convergence follows from the convergence of the trust-region SQP method. \square

3.3 Quadratic convergence and early branching

When implementing Algorithm 2 the following adverse behaviour has been observed on some problems.

1. Some integer variables converge to a non-integral solution, namely $y_i^{(k)} \rightarrow \hat{y}_i$ but $\theta \simeq 0.02 < \tau$, i.e. the integrality gap remains bounded away from zero.
2. The SQP solver converges at second order rate during this time and $\theta \gg \|d^{(k)}\|_\infty \rightarrow 0$.

In other words the early branching heuristic is not activated during the second order convergence of the SQP method so that no early branching takes place in the final stages of the SQP solver (as $\tau = 0.1$ is too large).

This observation led to a new early branching heuristic which takes the quadratic rate of convergence of the SQP method into account. Steps 5. and 6. of Algorithm 2 are replaced by

- 5'. Compute the integrality gap $\theta = \max_i |y_i^{(k+1)} - \text{anint}(y_i^{(k+1)})|$
and the experimental order of convergence $p := \ln(\|d^{(k)}\|_\infty) / \ln(\|d^{(k-1)}\|_\infty)$.
- 6'. **if** ($\theta > \tau$ or ($p > 1.5$ and $\|d^{(k)}\|_\infty < \theta$)) **then**
 Choose a non-integral $y_i^{(k+1)}$ and branch, **exit**
endif

For a quadratically convergent method one expects that $p \simeq 2$ and once quadratic convergence sets in, that $\|d^{(k+1)}\|_\infty \simeq \|d^{(k)}\|_\infty^p$. Thus if $\|d^{(k)}\|_\infty < \theta$, one cannot expect to reduce the integrality gap to zero in the remaining SQP iterations.

The new early branching rule resulted in a small but consistent reduction in the number of QP problems solved (roughly a 5% reduction) compared to the original heuristic.

3.4 Nonconvex MINLP problems

In practice many MINLP problems are nonconvex. In this case, the underestimating property of Lemma 1 no longer holds, as linearizations of nonconvex functions are not necessarily underestimators. Thus it may be possible that a node is wrongly fathomed on account of Lemma 1 and this could prevent the algorithm from finding the global minimum.

A rigorous way of extending Algorithm 2 to classes of nonconvex MINLP problems would be to replace the linearizations by linear underestimators (e.g. Horst and Tuy [17, Chapter VI]). However, it is not clear what effect this would have on the convergence properties of the SQP method. Instead the following heuristic is proposed, which consists of replacing 2'. by 2'':

- 2''. **if** ($(QP_{s,\infty}^k)$ infeasible) **then**
 Enter feasibility restoration phase:
 repeat (SQP iteration)
 (a) Compute a step of the feasibility restoration.
 (c) **if** ($(QP_{s,\infty}^k)$ feasible) **then** return to normal SQP, **exit**
 (d) **if** ($\min \|g^+(x, y)\| > 0$) **then** fathom node, **exit**
 endif

The main difference between 2' and 2'' is that a node is no longer fathomed if the QP problem is infeasible and the step lies strictly inside the trust-region ($\|d\| < \rho$). Instead a feasibility problem has to be solved to optimality (step (d) above) or until a feasible QP problem is encountered (step (b) above). This heuristic is motivated by the success of nonlinear branch-and-bound in solving some nonconvex MINLP problems. It ensures that a node is only fathomed if it would have been fathomed by nonlinear branch-and-bound.

4 Numerical Experience

Nonlinear branch-and-bound and Algorithm 2 have been implemented in Fortran 77 using filterSQP [10] as the underlying NLP solver. The implementation of Algorithm 2 uses a callback routine that is invoked after each step of the SQP method. This routine manages the branch-and-bound tree stored on a stack, decides on when to branch and which problem to solve next. The use of this routine has kept the necessary changes to filterSQP to a minimum.

Header	Description
problem	The SIF name of the problem being solved.
n	Total number of variables of the problem.
n_i	Number of integer variables of the problem.
m	The total number of constraints (excl. simple bounds).
m_n	Number of nonlinear constraints.
k	The dimension of the null-space at the solution.
# NLPs	Number of NLP problems solved (branch-and-bound).
# nodes	Number of nodes visited by Algorithm 2.
# QPs	Number of QP (and LP) problems solved.
CPU	Seconds of CPU time needed for the solve.

Table 1: Description of headers of the result tables

The test problems are from a wide range of backgrounds. Their characteristics are displayed in Table 2. The SYNTHES* problems are small process synthesis problems [6]. OPTPRLOC is a problem arising from the optimal positioning of a product in a multi-attribute space [6]. BATCH is a batch plant design problem that has been transformed to render it convex [18]. FEEDLOC is a problem arising from the optimal sizing and feed tray location of a distillation column [24], TRIMLOSS is a trim loss optimization problem that arises in the paper industry [26] and TRIMLON4 is an equivalent nonconvex formulation of the same problem. All problems are input as SIF files [3] plus an additional file to identify the integer variables. The test-problems are available at http://www.mcs.dundee.ac.uk:8080/~sleyffer/MINLP_TP/.

Table 1 lists a description of the headers of the result tables. Table 2 list the problem characteristics. The performance of nonlinear branch-and-bound and Algorithm 2 is compared in Table 3. All runs were performed on a SUN SPARCstation 4 with 128 Mb memory, using the compiler options `-r8 -o` and SUN's `f77` Fortran compiler.

Table 3 shows that branch-and-bound only requires about 2-6 QP solves per node that is solved making it a very efficient solver. By comparison, Algorithm 2 often visits more nodes in the branch-and-bound tree. That is not surprising, as the early termination rule is based on a

problem	n	n_i	m	m_n	Description
SYNTHES1	6	3	6	2	Process synthesis
SYNTHES2	11	5	14	3	Process synthesis
SYNTHES3	17	8	23	4	Process synthesis
OPTPRLOC	30	25	30	25	Optimal product location
BATCH	49	24	73	1	Batch plant design
FEEDLOC	90	37	259	89	Distillation column design (nonconvex)
TRIMLOSS	142	122	75	4	Trimloss optimization
TRIMLON4	24	24	27	4	Trimloss optimization (nonconvex)

Table 2: Problem characteristics of test problems

problem	Branch-and-bound			Algorithm 2		
	# NLPs	# QPs	CPU	# nodes	# QPs	CPU
SYNTHES1	5	18	0.1	5	13	0.1
SYNTHES2	17	53	0.4	17	40	0.3
SYNTHES3	25	80	0.8	25	64	0.8
OPTPRLOC	109	491	8.5	112	232	5.3
BATCH	143	371	8.9	181	391	14.8
FEEDLOC	47	340	70.8	55	103	29.0
TRIMLOSS	1336	3820	254.3	944	1624	90.3
TRIMLON4	887	1775	19.0	566	900	13.3

Table 3: Results for branch-and-bound and Algorithm 2

weaker lower bound than **FR2**. However, the overall number of QP problems that are being solved is reduced by 20% – 80%. This results in a dramatic reduction in terms of CPU time compared to nonlinear branch-and-bound. For the largest problem, Algorithm 2 is 3 times faster than nonlinear branch-and-bound.

Even though problem TRIMLON4 is nonconvex, Algorithm 2 terminated with the global minimum. For the nonconvex problem FEEDLOC, Algorithm 2 did not find an integer feasible point. This behaviour is due to the fact that the linearization are *not* outer approximations in the nonconvex case and shows that the early termination rule does not hold for nonconvex MINLP problems. Running Algorithm 2 in nonconvex mode (see Section 3.4) produced the same minimum as branch-and-bound. As one would expect, the performance of Algorithm 2 in nonconvex mode is not as competitive as in convex mode (65 nodes visited, 159 QPs solved and 55.9 s CPU time) but still significantly faster than nonlinear branch-and-bound. However, better heuristics for nonconvex MINLP problems remain an open problem.

5 Conclusions

We have presented an algorithm for MINLP problems which interlaces the iterative solution of each NLP node with the branch-and-bound tree search. An implementation of this idea gave a

factor of about 3 improvement in terms of CPU time compared to nonlinear branch-and-bound.

References

- [1] Borchers, B. and Mitchell, J.E. An improved branch and bound algorithm for Mixed Integer Nonlinear Programming. *Computers and Operations Research*, 21(4):359–367, 1994.
- [2] Borchers, B. and Mitchell, J.E. A comparison of branch and bound and outer approximation methods for 0–1 MINLPs. Talk given at IFORS 4th specialized Conference on OR and Engineering Design, October 24–27, 1995, St. Louis, Missouri., October 1995.
- [3] Conn, A.R. Gould, N.I.M. and Toint, Ph.L. *LANCELOT: a Fortran package for large-scale nonlinear optimization (Release A)*. Springer Verlag, Heidelberg, New York, 1992.
- [4] Conn, A.R. Gould, N.I.M. and Toint, Ph.L. Methods for nonlinear constraints in optimization calculations. Report RAL-TR-96-042, Rutherford Apleton Laboratory, 1996. (To appear in "The State of the Art in Numerical Analysis", I. Duff and G.A. Watson (eds.)).
- [5] Dakin, R.J. A tree search algorithm for mixed integer programming problems. *Computer Journal*, 8:250–255, 1965.
- [6] Duran, M. and Grossmann, I.E. An outer-approximation algorithm for a class of mixed–integer nonlinear programs. *Mathematical Programming*, 36:307–339, 1986.
- [7] Fletcher, R. *Practical Methods of Optimization, 2nd edition*. John Wiley, Chichester, 1987.
- [8] Fletcher, R. and Leyffer, S. Solving mixed integer nonlinear programs by outer approximation. *Mathematical Programming*, 66:327–349, 1994.
- [9] Fletcher, R. and Leyffer, S. Nonlinear programming without a penalty function. Numerical Analysis Report NA/171, Department of Mathematics, University of Dundee, September 1997.
- [10] Fletcher, R. and Leyffer, S. User manual for filterSQP. Numerical Analysis Report NA/181, Dundee University, April 1998.
- [11] Flippo, O.E. and Rinnoy Kan, A.H.G. Decomposition in general mathematical programming. *Mathematical Programming*, 60:361–382, 1993.
- [12] Geoffrion, A.M. Generalized benders decomposition. *Journal of Optimization Theory and Applications*, 10:237–260, 1972.
- [13] Grossmann, I.E. and Kravanja, Z. Mixed–integer nonlinear programming: A survey of algorithms and applications. In A.R. Conn L.T. Biegler, T.F. Coleman and F.N. Santosa, editors, *Large–Scale Optimization with Applications, Part II: Optimal Design and Control*, New York, Berlin, 1997. Springer.
- [14] Grossmann, I.E. and Sargent, R.W.H. Optimal design of multipurpose batch plants. *Ind. Engng. Chem. Process Des. Dev.*, 18:343–348, 1979.

- [15] Gupta, O.K. and Ravindran, A. Branch and bound experiments in convex nonlinear integer programming. *Management Science*, 31:1533–1546, 1985.
- [16] Han, S.P. A globally convergent method for nonlinear programming. *Journal of Optimization Theory and Applications*, 22(3):297–309, 1977.
- [17] Horst, R. and Tuy, H. *Global Optimization: Deterministic Approaches*, 2nd edition. Springer–Verlag, Berlin, 1993.
- [18] Kocis, G.R. and Grossmann, I.E. Global optimization of nonconvex mixed–integer nonlinear programming (MINLP) problems in process synthesis. *Industrial Engineering Chemistry Research*, 27:1407–1421, 1988.
- [19] Land, A.H. and Doig, A.G. An automatic method for solving discrete programming problems. *Econometrica*, 28:497–520, 1960.
- [20] Powell, M.J.D. A fast algorithm for nonlinearly constrained optimization calculations. In G.A. Watson, editor, *Numerical Analysis, 1977*, pages 144–157, Berlin, 1978. Springer–Verlag.
- [21] Quesada, I. and Grossmann, I.E. An LP/NLP based branch–and–bound algorithm for convex MINLP optimization problems. *Computers and Chemical Engineering*, 16:937–947, 1992.
- [22] Quist, A.J. van Geemert, R. Hoogenboom, J.E. Illés, T. de Klerk, E. Roos, C. and Terlaky, T. Optimization of a nuclear reactor core reload pattern using nonlinear optimization and search heuristics. Draft paper, Delft University of Technology, Faculty of Applied Mathematics, Department of Operation Research, Mekelweg 4, 2628 CD Delft, The Netherlands, September 1997.
- [23] Stubbs, R.A. and Mehrotra, S. A branch–and–cut method for 0–1 mixed convex programming. Technical report 96-01, Department of IE/MS, Northwestern University, Evanston, IL 60208, USA, January 1996.
- [24] Viswanathan, J. and Grossmann, I.E. Optimal feed location and number of trays for distillation columns with multiple feeds. *I&EC Research*, 32:2942–2949, 1993.
- [25] Volkovich, O.V. Roshchin, V.A. and Sergienko, I.V. Models and methods of solution of quadratic integer programming problems. *Cybernetics*, 23:289–305, 1987.
- [26] Westerlund, T. Isaksson, J. and Harjunkski, I. Solving a production optimization problem in the paper industry. Report 95–146–A, Department of Chemical Engineering, Abo Akademi, Abo, Finland, 1995.
- [27] Wilson, R.B. *A simplicial algorithm for concave programming*. PhD thesis, Harvard University Graduate School of Business Administration, 1963.