ARGONNE NATIONAL LABORATORY
9700 South Cass Avenue
Lemont, Illinois 60439

Minotaur: A Mixed-Integer Nonlinear Optimization Toolkit

Ashutosh Mahajan, Sven Leyffer, Jeff Linderoth, James Luedtke, and Todd Munson

Mathematics and Computer Science Division

Preprint ANL/MCS-P8010-0817

October 16, 2017

This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, under Contract DE-AC02-06CH11357. This work was also supported by the U.S. Department of Energy under grant DE-FG02-05ER25694.

Contents

Noname manuscript No.
(will be inserted by the editor)

Minotaur: A Mixed-Integer Nonlinear Optimization Toolkit

Ashutosh Mahajan · Sven Leyffer · Jeff Linderoth · James Luedtke · Todd Munson

the date of receipt and acceptance should be inserted later

Abstract We present a flexible framework for general mixed-integer nonlinear programming (MINLP), called Minotaur, that enables both algorithm exploration and structure exploitation without compromising computational efficiency. This paper documents the concepts and classes in our framework and shows that our implementations of standard MINLP techniques are efficient compared with other state-of-the-art solvers. We then describe structure-exploiting extensions that we implement in our framework and demonstrate their impact on solution times. Without a flexible framework that enables structure exploitation, finding global solutions to difficult nonconvex MINLP problems will remain out of reach for many applications.

Keywords Mixed-Integer Nonlinear Programming, Global Optimization, Software

Mathematics Subject Classification (2000) 65K05, 90C11, 90C30, 90C26

1 Introduction, Background, and Motivation

Over the past two decades, mixed-integer nonlinear programming (MINLP) has emerged as a powerful modeling paradigm that arises in a broad range of scientific, engineering, and financial applications; see, e.g., [7,27,38, 54,63,68]. MINLP combines the combinatorial complexity of discrete decision variables with the challenges of nonlinear expressions, resulting in a class of difficult nonconvex optimization problems. The nonconvexities can arise from both the integrality restrictions and nonlinear expressions. MINLP problems can be generically expressed as

$$\begin{cases} \underset{x}{\text{minimize}} & f(x), \\ \text{subject to} & c(x) \leq 0, \\ & x \in \mathcal{X},\ x_i \in \mathbb{Z},\ \forall i \in \mathcal{I}, \end{cases} \tag{1.1}$$

A. Mahajan
Industrial Engineering and Operations Research, IIT Bombay, Mumbai, 400076, India, E-mail: amahajan@iitb.ac.in

S. Leyffer
Mathematics and Computer Science Division, Argonne National Laboratory, Lemont, IL 60439, E-mail: leyffer@mcs.anl.gov

J. Linderoth
Department of Industrial and Systems Engineering, University of Wisconsin-Madison, Madison, WI 53706, E-mail: linderoth@wisc.edu

J. Luedtke
Department of Industrial and Systems Engineering, University of Wisconsin-Madison, Madison, WI 53706, E-mail: jim.luedtke@wisc.edu

T. Munson
Mathematics and Computer Science Division, Argonne National Laboratory, Lemont, IL 60439, E-mail: tmunson@mcs.anl.gov

where $f : \mathbb{R}^n \to \mathbb{R}$ and $c : \mathbb{R}^n \to \mathbb{R}^m$ are twice continuously differentiable functions, $\mathcal{X} \subset \mathbb{R}^n$ is a bounded polyhedral set, and $\mathcal{I} \subseteq \{1, \ldots, n\}$ is the index set of the integer variables. Equality and range constraints can be readily included in (1.1).

MINLP problems are at least NP-hard combinatorial optimization problems because they include mixed-integer linear programming (MILP) as a special case [47]. In addition, general nonconvex MINLP problems can be undecidable [46]. In the remainder of this paper, we consider only MINLP problems (1.1) that are decidable by assuming that either $\mathcal{X}$ is compact or the problem functions, f and c, are convex. In reality, the distinction between hard and easy problems in MINLP is far more subtle, and instances of NP-hard problems are routinely solved by state-of-the-art solvers.

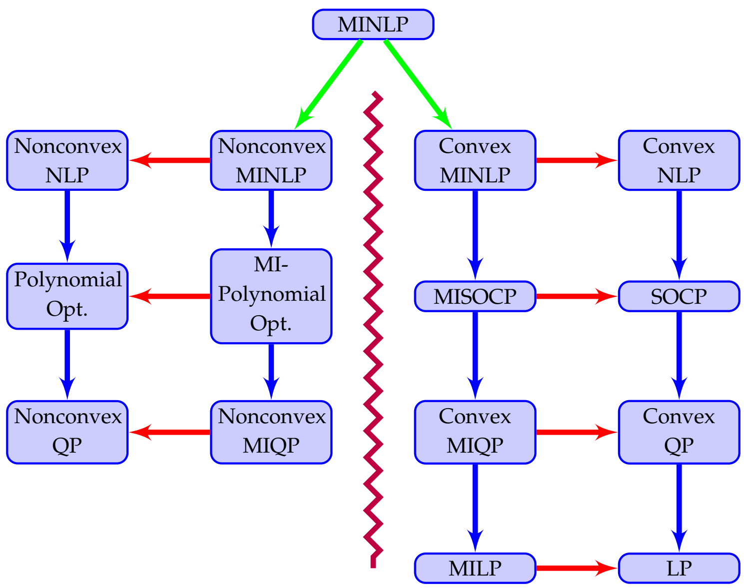

Fig. 1 MINLP problem class tree.

Figure 1 provides an overview of the problem classes within the generic MINLP formulation in (1.1). At the top level, we divide MINLP problems into convex and nonconvex problems (green arrows), where convex refers to problems in which the region defined by the nonlinear constraints, but excluding integrality, is convex. Next, we further divide the problem classes depending on whether they contain discrete variables or not (red arrows). Then we subdivide the problems further by the class of functions that are present. Figure 1 illustrates the broad structural diversity of MINLP. In addition to the standard problem classes of nonlinear programming (NLP), quadratic programming (QP), linear programming (LP), and their mixed-integer (MI) versions, our tree includes second-order cone programming (SOCP) and polynomial optimization, which have received much interest [49, 50, 51]. This tree motivates the development of a flexible software framework for specifying and solving MINLP problems that can be extended to tackle different classes of constraints.

Existing solvers for convex MINLP problems include α-ECP [74], BONMIN [9], DICOPT [72], FilMINT [1], GuRoBi [43] (for convex MIQP problems with quadratic constraints), KNITRO [14, 76], MILANO [8], MINLPBB [52], and SBB [12]. These solvers require only first and second derivatives of the objective function and constraints. The user can either provide routines that evaluate the functions and their derivatives at given points or use modeling languages such as AMPL [29], GAMS [11], or Pyomo [44] to provide them automatically. These solvers are not designed to exploit the structure of nonlinear functions. While the solvers can be applied to nonconvex MINLP problems, they are not guaranteed to find an optimal solution.

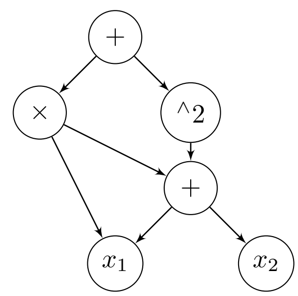

Fig. 2 A directed acyclic graph representing the nonlinear function $x_1(x_1 + x_2) + (x_1 + x_2)^2$.

Existing solvers for nonconvex MINLP include α-BB [4], ANTIGONE [59], BARON [64], COCONUT [60,66], Couenne [6], CPLEX [45] (for MIQP problems with a nonconvex objective function), GloMIQO [58], and SCIP [2, 71]. These solvers require the user to explicitly provide the definition of the objective function and constraints. While linear and quadratic functions can be represented by using data stored in vectors and matrices, other nonlinear functions are usually represented by means of computational graphs. A computational graph is a directed acyclic graph (DAG). A node in the DAG represents either a variable, a constant, or an operation (e.g., $+, -, \times, /, \exp, \log$). An arc connects an operator to its operands. An example of a DAG is shown in Figure 2. Functions that can be represented by using DAGs are referred to as "factorable functions." Modeling languages allow the user to define nonlinear functions in a natural algebraic form, convert these expressions into DAGs, and provide interfaces to read and copy the DAGs.

Algorithmic advances over the past decade have often exploited special problem structure. To exploit these advances, MINLP solvers must be tailored to special problem classes and nonconvex structures. In particular, a solver must be able to evaluate, examine, and possibly modify the nonlinear functions. In addition, a single MINLP solver may require several classes of relaxations or approximations to be solved as subproblems, including LP, QP, NLP, or MILP problems. For example, our QP-diving approach [55] solves QP approximations and NLP relaxations. Different nonconvex structures benefit from tailored branching, bound tightening, cut-generation, and separation routines. In general, nonconvex forms are more challenging and diverse than integer variables, thus motivating a more tailored approach. Moreover, the emergence of new classes of MINLP problems such as MILP problems with second-order cone constraints [20,19] and MINLP problems with partial-differential equation constraints [53] necessitates novel approaches.

These challenges and opportunities motivate the development of our Minotaur software framework for MINLP. Minotaur stands for Mixed-Integer Nonlinear Optimization Toolkit: Algorithms, Underestimators, and Relaxations. Our vision is to enable researchers to implement new algorithms that take advantage of problem structure by providing a general framework that is agnostic of problem type or solvers. Therefore, the goals of Minotaur are to (1) provide reliable, efficient, and usable MINLP solvers; (2) implement a range of algorithms in a common framework; (3) provide flexibility for developing new algorithms that can exploit special problem structure; and (4) reduce the burden of developing new algorithms by providing a common software infrastructure.

The remainder of this paper is organized as follows. In Section 2, we briefly review some fundamental algorithms for MINLP and highlight the main computational and algorithmic components that motivate the design of Minotaur. In Section 3, we describe Minotaur's class structure and introduce the basic building blocks of Minotaur. In Section 4, we show how we can use this class structure to implement the basic algorithms described in Section 2. Section 5 presents some extensions to these basic algorithms that exploit additional problem structure, including a nonlinear presolve and perspective reformulations. This section illustrates how one can take advantage of our software infrastructure to build more complex solvers. Section 6 summarizes our conclusions.

Throughout, we demonstrate the impact of our techniques on sets of benchmark test problems and show that we do not observe decreased performance for increased generality.

2 General Algorithmic Framework

Minotaur is designed to implement a broad range of relaxation-based tree-search algorithms for solving MINLP problems. In this section, we describe the general algorithmic framework and demonstrate how several algorithms for solving MINLP problems fit into this framework. We concentrate on describing single-tree methods for *convex MINLP problems* (i.e., MINLP problems for which the nonlinear functions f and c are convex), such as nonlinear branch-and-bound [48,18,42] and LP/NLP-based branch-and-bound [62]. We also describe how nonconvex MINLP problems fit into our framework. We do not discuss multitree methods such as Benders decomposition [36,69], outer approximation [21,24,9], or the extended cutting plane method [67,75], although these methods can be implemented by using Minotaur.

2.1 Relaxation-Based Tree-Search Framework

The pseudocode in Algorithm 1 describes a basic tree-search algorithm. In the algorithm, P, P', and Q represent subproblems of the form (1.1), which may be obtained by reformulating a problem or by adding restrictions introduced in the branching phase. The set $\mathcal{O}$ maintains a list of open subproblems that need to be processed. When this list is empty, the algorithm terminates. Otherwise, a node is selected from the list and "processed." The results from processing a node are used to (1) determine whether a new feasible solution is found and, if it is better than an existing solution, to update the *incumbent* solution and (2) determine whether the subproblem can be pruned. If the subproblem cannot be pruned, then the subproblem is *branched* to obtain additional subproblems that are added to $\mathcal{O}$.

Algorithm: `Tree Search`
Reformulate P to obtain problem P'.
Place P' in the set $\mathcal{O}$.
repeat
    Node selection: Select and remove a problem Q from $\mathcal{O}$
    (FeasSolutionFound,CanPrune) $\leftarrow$ **Process node** Q
    if FeasSolutionFound **then**
        Update incumbent
    else if not CanPrune **then**
        Branch: Create subproblems and add to $\mathcal{O}$
until $\mathcal{O}$ *is empty*

Algorithm 1: Generic tree-search algorithm for solving a MINLP problem P.

The standard mechanism in Minotaur for node processing is described in Algorithm 2, which constructs relaxations of the subproblem Q. A relaxation of Q is a problem R such that the optimal value of R (when minimizing) is guaranteed to be a lower bound on the optimal value of Q. After the relaxation is constructed, the relaxation problem is solved to obtain its optimal value. If the relaxation problem is infeasible or if its optimal value matches or exceeds the value of a known upper bound, then the subproblem can be pruned. Otherwise, the relaxation may be updated (this step might be skipped in some algorithms); and, if the relaxation solution x^R is no longer feasible, then the updated relaxation is solved.

The relaxation used is a key characteristic of a tree-search algorithm. The basic requirements for a relaxation are that it provides a lower bound on the optimal value of a subproblem and can be solved by an available solver.

Tighter relaxations are preferred because they typically result in smaller branch-and-bound search trees. Creating and updating (refining) relaxations from the description of subproblem Q are critical computational tasks.

Algorithm: `Process Node` Q

Input: Incumbent value v^I and subproblem Q

Output: FeasSolutionFound, CanPrune, and possibly updated incumbent value v^I

FeasSolutionFound $\leftarrow$ FALSE, CanPrune $\leftarrow$FALSE

Construct a relaxation R of Q

repeat

- **Solve relaxation:** Solve problem R to obtain value v^R and solution x^R
- **if** *R is infeasible* **or** $v^R \geq v^I$ **then**
  - CanPrune $\leftarrow$ TRUE
- **else if** *x^R is feasible for Q* **then**
  - FeasSolutionFound $\leftarrow$ TRUE
- **else**
  - **Update Relaxation** R

until CanPrune **or** FeasSolutionFound **or** *x^R feasible for updated R*

Algorithm 2: Generic relaxation-based algorithm for processing a node Q.

Minotaur provides a basic infrastructure for managing and storing the open nodes in the tree-search algorithm (the *tree*), for interfacing to modeling languages and subproblems solvers, and for performing basic housekeeping tasks, such as timing and statistics. Section 3 shows how these computational components are implemented in Minotaur's class structure. The remaining subsections illustrate how these components are used to build standard solvers and to develop more advanced MINLP solvers.

2.2 Nonlinear Branch and Bound for Convex MINLPs

In nonlinear branch and bound (NLPBB) for convex MINLP problems, a relaxation R of a subproblem Q is obtained by relaxing the constraints $x_j \in \mathbb{Z}$ to $x_j \in \mathbb{R}$ for all $j \in I$. The resulting problem is a continuous NLP problem; and when all functions defining the (reformulated) MINLP P' are smooth and convex, Q can be solved to global optimality with standard NLP solvers.

If the solution of the relaxation x^R is integer feasible (i.e., $x_j^R \in \mathbb{Z}$ for all $j \in I$), then the relaxation solution is feasible and the node processor sets FeasSolutionFound to TRUE. If the relaxation is infeasible or its optimal value is at least as large as the incumbent optimal value, then the subproblem can be pruned. Otherwise, branching must be performed. In this case, branching is performed by choosing a variable x_j with $j \in I$ such that $x_j^R \notin \mathbb{Z}$. Then, two new subproblems are created by adding new bounds $x_j \leq \lfloor x_j^R \rfloor$ and $x_j \geq \lceil x_j^R \rceil$, respectively, and these subproblems are added to $\mathcal{O}$.

2.3 LP/NLP-Based Branch-and-Bound Algorithms for Convex MINLPs

Modern implementations of the LP/NLP-based branch-and-bound method [62] are among the most powerful solvers [9,1] for convex MINLP problems. The basic idea is to replace the NLP relaxation used in NLPBB with an LP relaxation. This LP relaxation is constructed by relaxing the constraints $x_j \in \mathbb{Z}$ to $x_j \in \mathbb{R}$ for all $j \in I$ and by replacing the nonlinear functions f and c with piecewise-linear lower bounds obtained from first-order Taylor-series approximations about a set of points $x^{(l)}$ for $l \in \mathcal{L}$. The convexity of the problem functions ensures that this linearization provides an outer approximation. As usual, if this relaxation is infeasible or its objective value is at least as large as the incumbent objective value, then the subproblem can be pruned.

Feasibility of the relaxation solution x^R is checked with respect to both the integrality constraints and the relaxed nonlinear constraints.

1. If x^R is feasible for both, then the incumbent is updated, and the node is pruned.
2. If x^R is integer feasible, but violates a nonlinear constraint, then the relaxation is updated by fixing the integer variables $x_j = x_j^R$ for all $j \in I$ and solving the resulting continuous NLP subproblem. If the NLP subproblem is feasible and improves upon the best-known solution, then the incumbent is updated. Whether the NLP subproblem is feasible or not, the set of linearization points $x^{(l)}$ for $l \in \mathcal{L}$ is updated so that the LP relaxation is refined.
3. If x^R is not integer feasible, then either the LP relaxation can be refined (e.g., by updating the set of linearization points so that the relaxation solution x^R is no longer feasible), or we can choose to exit the node processing.

If the node processing terminates with a relaxed solution that is not integer feasible, then, as in NLPBB, the subproblem is subdivided by choosing an integer variable j with $x_j^R \notin \mathbb{Z}$ and updating the bounds in the two subproblems.

2.4 Branch and Bound for Nonconvex MINLPs

If the problem functions f or c are nonconvex, then standard NLP solvers are not guaranteed to solve the continuous relaxation of (1.1) to global optimality. In order to ensure that the relaxations remain solvable, convex relaxations of the nonconvex feasible set must be created. In such relaxations, the quality of the outer approximation depends on the tightness of the variable bounds. The details of such a relaxation scheme for nonconvex quadratically constrained quadratic programs are described in Section 4. A key difference is that in addition to branching on integer variables, this algorithm requires branching on continuous variables that appear in nonconvex expressions in (1.1). Thus, in the branching step, subproblems may be created by subdividing the domain of a continuous variable. The updated lower and upper bounds are then used when these subproblems are processed to obtain tighter convex outer approximations of the nonconvex feasible region.

3 Software Classes in Minotaur

The Minotaur framework is written in C++ by using a class structure that allows developers to easily implement new functionality and exploit structure. By following a modular approach, the components remain interoperable and compatible if the functions they implement are compatible. Thus, developers can customize only a few selected components and use the other remaining components to produce new solvers. In particular, a developer can override the default implementation of only a few specific functions of a class by creating a “derived C++ class” that implements these functions using methods or algorithms different from the base class. This approach also facilitates easy development of extensions to solvers, such as MINLP solvers with nonconvex nonlinear expressions.

Our framework has three main parts: (1) core, (2) engine, and (3) interface. The core includes all methods and data structures used while solving a problem, for example, those to store, modify, analyze, and presolve problems, create relaxations and add cuts, and implement the tree search and heuristic searches. The engine includes routines that call various external solvers for LP, QP, or NLP relaxations or approximations. The interface contains routines that read input files in different formats and construct an instance. We first describe the most commonly used classes in these three parts and then demonstrate how some can be overridden.

Table 1 Base classes used in defining a MINLP problem.

Name of Class	Description
`Variable`	Store a representation of a variable including its index in the instance, its type (integer, binary or continuous), its lower and upper bounds, and other useful information, such as the list of constraints where it appears and whether it appears in any nonlinear functions.
`Function`	Store a mathematical function of the variables and define C++ functions to evaluate the function, gradient, and Hessian at a given point. It also has routines to infer bounds on function values, check the type of function, add another mathematical function, replace a variable by another one, remove variables, and implement other tasks.
`LinearFunction`	Store a linear function of variables and implements methods to evaluate it at a given point and to query and modify the coefficients.
`QuadraticFunction`	Store a quadratic function of variables and implements methods to evaluate the function, gradient, and Hessian at a given point and query and modify the coefficients.
`NonlinearFunction`	Store a nonlinear function of variables and implements methods to evaluate the function, gradient, and Hessian at a given point.
`Constraint`	Store and modify properties of a constraint including its index in the problem, its lower and upper bounds, and the functions used by the constraint.
`Objective`	Store and modify properties of the objective including the functions used by the objective, its sense (maximize or minimize), and the constant offset.
`Problem`	Store a representation of a MINLP problem. The object stores `Variable`, `Constraint`, and `Objective` objects. Additionally it has routines to store and query the Jacobian and Hessian of the Lagrangian and other problem-specific objects, such as `SOS` variables.

3.1 `Core`

The C++ classes in the core can be classified into four types based on their use.

3.1.1 Classes Used to Represent and Modify a Problem

A MINLP problem is represented by means of the `Problem` class. This class stores pointers to all the variables and constraints and the objective and provides methods to query and modify them. Each variable is an object of the `Variable` class. Similarly, constraints and the objective function are objects of the `Constraint` and the `Objective` classes, respectively. A separate `SOS` class is provided for storing special ordered sets (e.g., SOS-1 and SOS-2). This scheme provides a natural and intuitive representation of the problem that is easy to modify. Table 1 lists the main classes used in defining a MINLP problem in Minotaur and their brief description.

Since the objective and constraints of the MINLP problem may have general nonlinear functions, we require specialized data structures for these functions. Among other operations, the `Constraint` and `Objective` classes each store a pointer to an object of the `Function` class. The `Function` class in turn contains pointers to objects of the `LinearFunction`, `QuadraticFunction`, and `NonlinearFunction` classes and provides other operations. Thus, we store the mathematical function of a constraint or objective as a sum of a linear component, a quadratic component, and a general nonlinear component. A linear constraint in the MINLP problem, for example, is represented by a `Constraint` object whose `Function` class has a pointer to only a `LinearFunction`; the pointers to `QuadraticFunction` and `NonlinearFunction` are null. The `NonlinearFunction` class has several derived classes that we describe next.

The `CGraph` Class is a derived class of the `NonlinearFunction` class used to store any factorable function. As described in Section 1, it stores a nonlinear function in the form of a directed acyclic graph. Thus, this class has DAG-specific methods, such as adding or deleting a node or changing a variable or a constant. These methods can be used to create and modify any factorable function by using a given set of operators. For instance, Figure 3 shows an excerpt of code that can be used to create an object of `CGraph` class corresponding to the example

```
ProblemPtr p = (ProblemPtr) new Problem();
VariablePtr x1 = p->newVariable(0,1,Binary);
VariablePtr x2 = p->newVariable(0,1,Binary);

CGraphPtr cg = (CGraphPtr) new CGraph();
CNode *nx1 = cg->newNode(x1);
CNode *nx2 = cg->newNode(x2);
CNode *np  = cg->newNode(OpPlus, nx1, nx2);
CNode *nm  = cg->newNode(OpMult, nx1, np);
CNode *ns  = cg->newNode(OpSqr, np);

cg->setOut(cg->newNode(OpPlus, nm, ns));
cg->finalize(); cg->write(std::cout);
```

Fig. 3 Excerpt of code used to create and display the nonlinear function in two variables $x_1(x_1 + x_2) + (x_1 + x_2)^2$.

DAG from Figure 2. A more complicated example is shown in Figure 4 that constructs the function needed for an approximation of the perspective formulation of a given nonlinear expression; see Section 5.

Being a derived class of `NonlinearFunction`, the `CGraph` class also contains routines for evaluating the gradient and Hessian of the function it stores. We have implemented automatic differentiation techniques [15,35, 37] for these purposes. In particular, the gradient is evaluated by using reverse mode. The Hessian evaluation of a function $f : \mathbb{R}^n \to \mathbb{R}$ uses at most n evaluations, one for each column of the Hessian matrix. In each iteration, say i, we first evaluate $\nabla f(x)^T e_i$ in a forward-mode traversal and then perform a reverse-mode traversal to compute the ith column of the Hessian (see, e.g., [61, Ch. 7]).

Besides computing the derivatives, the `CGraph` class is used for finding bounds on the values that the nonlinear function can assume over the given ranges of variables. Conversely, it can deduce bounds on the values that a variable can assume from given lower or upper bounds on the nonlinear function and other variables. These techniques are called feasibility-based bound tightening [6].

The `MonomialFunction` Class is a derived class of the `NonlinearFunction` class used for representing monomial functions of the form $a \prod_{i \in J} x_i^{p_i}$, where $a \in \mathbb{R}, p_i \in \mathbb{Z}_+, i \in J$, and the set J are given. This class stores the pointer to the variables and the powers in a C++ map data structure.

The `PolynomialFunction` Class is a derived class of the `NonlinearFunction` class used for representing polynomial functions. It keeps a vector of objects of the `MonomialFunction` class to represent the polynomial.

3.1.2 Classes Used in Branch-and-Bound Algorithms

To keep the design of branch-and-bound algorithms modular and flexible, a base class is defined for every step of the algorithm described in Section 2. Table 2 lists some of the main classes and their functionality. The user can derive a new class for any of these steps without modifying the others.

We illustrate the design principle by means of the `NodeProcessor` class. The class implements the methods `processRootNode()` and `process()`, and other helper functions. Figure 5 depicts two ways of processing a node. In a simple NLPBB solver for convex MINLP problems, we may only need to solve an NLP relaxation of the problem at the current node. This procedure is implemented in the `BndProcessor` class derived from the `NodeProcessor` class. Based on the solution of this NLP, we may either prune the node or branch. For other algorithms, we may need a more sophisticated node processor that can call a cut-generator to add cuts or invoke presolve. The `PCBProcessor` derived class implements this scheme.

Modularity enables us to select at runtime the implementation of each component of the branch-and-bound algorithm. For instance, depending on the type of problem we are solving, we can select at runtime one of

```
CGraph* PerspRefUT::getPerspRec(CGraph* f, VariablePtr y, double eps, int *err)
{
  CNode *ynode = 0, *anode = 0;
  CGraph* p = new CGraph();

  ynode = p->newNode(y);
  anode = p->newNode(1.0-eps);
  ynode = p->newNode(OpMult, anode, ynode); // y*(1-eps)
  anode = p->newNode(eps);
  ynode = p->newNode(OpPlus, anode, ynode); // eps + y*(1-eps)

  anode = getPerspRec_(p, f->getOut(), ynode, f); // start recursion

  anode = p->newNode(OpMult, anode, ynode); // final multiplication
  p->setOut(anode); p->finalize();
  return p;
}

CNode* PerspRefUT::getPerspRec_(CGraph* p, const CNode *node, CNode *znode, CGraph* f)
{
  CNode *newl  = 0, *newr = 0;
  if (OpVar == node->getOp()) {
    newl = p->newNode(f->getVar(node));
    return (p->newNode(OpDiv, newl, znode));          // x/(y(1-eps)+eps)
  } else if (OpNum == node->getOp()) {
    return (p->newNode(node->getVal()));
  } else if (1 == node->numChild()) {
    newl = getPerspRec_(p, node->getL(), znode, f);   // recurse
    return (p->newNode(node->getOp(), newl, newl));
  } else if (2 == node->numChild()) {
    newl = getPerspRec_(p, node->getL(), znode, f);   // recurse
    newr = getPerspRec_(p, node->getR(), znode, f);   // recurse
    return (p->newNode(node->getOp(), newl, newr));
  } else if (2 < node->numChild()) {
    CNode **childrn = new CNode*[node->numChild()];   // array of children
    CNode** c1 = node->getListL();
    CNode** c2 = node->getListR();
    int i = 0;
    while(c1 < c2) {
      childrn[i] = getPerspRec_(p, *c1, znode, f);    // recurse
    }
    return (p->newNode(node->getOp(), childrn, node->numChild()));
  }
  return 0;
}
```

Fig. 4 Excerpt of code used to obtain a CGraph of $p(x, y) = (y(1-\epsilon)+\epsilon)f(\frac{x}{y(1-\epsilon)+\epsilon})$ from a given CGraph of $f(x)$ by recursively traversing it.

five branchers: `ReliabilityBrancher`, for implementing reliability branching [3]; `MaxVioBrancher`, for selecting the maximum violation of the nonconvex constraint; `LexicoBrancher`, for selecting the variable with the smallest index; `MaxFreqBrancher`, for selecting a variable that appears as a candidate most often; and `RandomBrancher`, for selecting a random variable. Minotaur also has branchers for global optimization that select continuous variables in order to refine the convex relaxations.

Table 2 Base classes used in implementing branch-and-bound algorithms for solving MINLP problems.

Name of Class	Description
`Brancher`	Determine how to branch a node of the branch-and-bound tree to create new subproblems.
`Engine`	Solve LP, QP, NLP, or some other problem using an external solver such as FilterSQP, IPOPT, or CLP.
`Node`	Store information about a node of the branch-and-bound tree such as the lower bound on the optimal value, the branching variable (or object), and pointers to child and parent nodes.
`NodeRelaxer`	Create or modify a relaxation of the problem at a given node in the branch-and-bound tree.
`TreeManager`	Store, add, and delete nodes of the branch-and-bound tree and select the node for processing.
`ActiveNodeStore`	Store active or open nodes of the branch-and-bound tree in an appropriate data structure and compute the global lower bound from the bounds of the individual nodes.
`CutManager`	Store, add, and remove cuts from the problem.
`NodeProcessor`	Process a node of the branch-and-bound tree by solving the relaxation and/or adding cuts, presolving, branching, or applying heuristics.

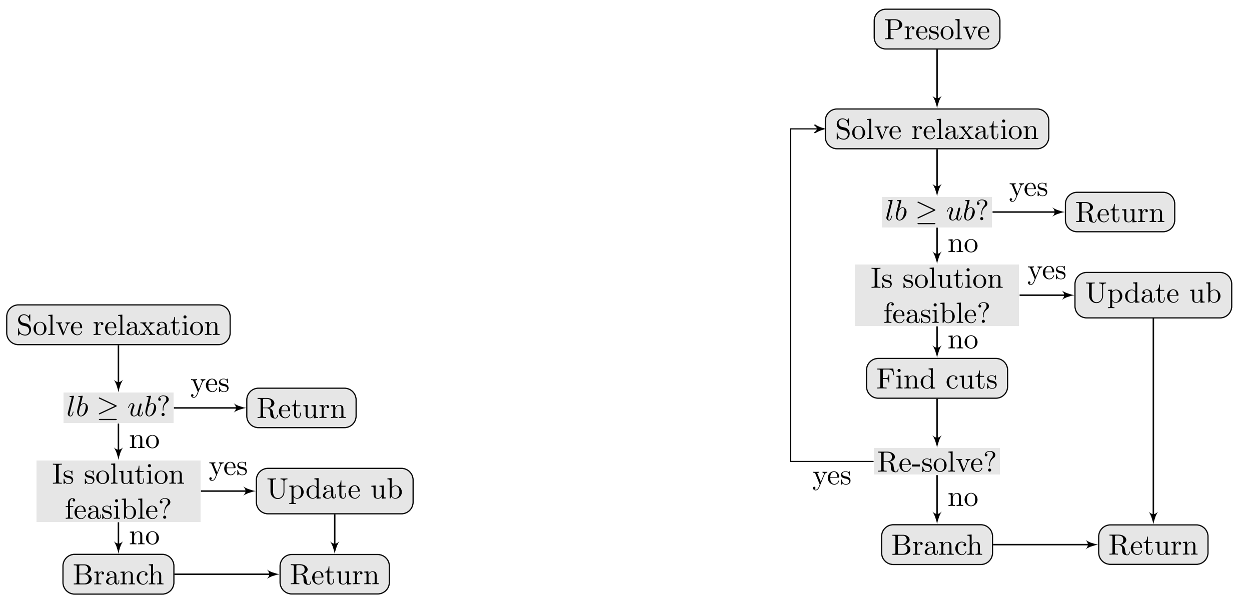

Fig. 5 Two ways of implementing a `NodeProcessor` class in the branch-and-bound method. Left: standard NLPBB; right: LP/NLP-based branch and bound.

3.1.3 Classes Used to Exploit Problem Structure

The classes mentioned in the preceeding sections implement general methods of a branch-and-bound algorithm and do not explicitly exploit the "structure" or "specific difficulties" of a problem. By structure, we mean those special properties of the constraints and variables that can be exploited (and, sometimes, must be exploited) to solve the problem. For instance, if we have integer decision variables and linear constraints, then we can generate valid inequalities to make the relaxation tighter by using the techniques developed by the MILP community. Similarly, if some constraints of the problem are nonlinear but convex, then we can create their linear relaxation by deriving linear underestimators from a first-order Taylor approximation.

We use a `Handler` class to enable implementation of routines that exploit specific problem structures. This idea is inspired from the "Constraint Handlers" used in SCIP [2] that ensure that a solution satisfies all the constraints of a problem. Since special properties or structures can be exploited in a branch-and-bound algorithm at many different steps, `Handler` instances are invoked at all these steps: creating a relaxation, presolving, finding valid inequalities, checking whether a given point is feasible, shortlisting candidates for branching, and creating branches. Table 3 lists some of the important functions of a `Handler`.

Table 3 Main functions of a `Handler`.

Name of the Function	Description
`relaxNodeFull`	Create a new relaxation of the constraints being handled.
`relaxNodeInc`	Create a relaxation by adding relevant modifications to the relaxation obtained at the parent node.
`presolve`	Presolve the specific structure.
`presolveNode`	Presolve the specific structure at a given node.
`isFeasible`	Check whether a point is feasible to the constraints being handled.
`separate`	Store, add, and remove valid inequalities for the specific structure.
`getBranchingCandidates`	Shortlist candidates for branching in a node of the branch-and-bound tree.
`branch`	Create branches of the current node using the branching candidate shortlisted by this handler and selected by the brancher.

The `Handler` base class in Minotaur is abstract; it declares only the functions that every `Handler` instance must have and leaves the implementation of the structure-specific methods to the developer. We describe a few commonly used `Handler` types implemented for our solvers.

`IntVarHandler` is one of the simplest handlers. It ensures that a candidate accepted as a solution satisfies all integer constraints at a given point. It implements only three of the functions listed in Table 3: `isFeasible`, `getBranchingCandidates`, and `branch`. The first function checks whether the value of each integer-constrained variable is integer within a specified tolerance. The second function returns a list of all the integer variables that do not satisfy their integer constraints. The last function creates two branches of a node if the `Brancher` selects an integer variable for branching.

`SOS2Handler` is used when the MINLP problem contains SOS-2 variables [5]. Its `isFeasible` routine checks whether at most two consecutive variables in the set are nonzero. The `getBranchingCandidates` routine returns two subsets, one for each branch. In the first branch all variables of the first subset are fixed to zero, and in the second all variables of the second subset are fixed to zero.

`QGHandler` is a more complicated handler used when solving convex MINLP problems with the LP/NLP-based algorithm of Quesada and Grossmann [62] (QG stands Quesada-Grossmann.) This handler does not implement any `branch` or `getBranchingCandidates` routines. The `isFeasible` routine checks whether the nonlinear constraints are satisfied by the candidate solution. If the problem is not feasible, then the `separate` routine solves a continuous NLP relaxation and obtains linear inequalities that cut off the given candidate. These valid inequalities are added to the LP relaxation, which is then solved by the `NodeProcessor`.

`NlPresHandler` is for applying presolving techniques to nonlinear functions in the constraints and objective. Therefore, it implements only the `presolve` and `presolveNode` functions. This handler tries to compute tighter bounds on the variables in nonlinear constraints by performing feasibility-based bound tightening on one constraint at a time. It also checks whether a nonlinear constraint is redundant or infeasible.

`QuadHandler` is for constraints of the form

$$x_i^2 = x_k, \quad \text{or} \tag{3.2}$$

$$x_i x_j = x_k \tag{3.3}$$

for some $i, j, k \in \{1, \ldots, n\}$, with bounds on all variables given by $l_i \leq x_i \leq u_i$ for all $i \in \{1, \ldots, n\}$. This handler ensures that a solution satisfies all constraints of type (3.2) and (3.3) of a MINLP problem. It implements all the functions in Table 3. The `RelaxNodeFull` and `RelaxNodeInc` functions create linear relaxations of these constraints by means of the McCormick inequalities [57],

$$(l_i + u_i)x_i - x_k \geq l_i u_i \tag{3.4}$$

for constraint (3.2) and

$$\begin{aligned}
l_j x_i + l_i x_j - x_k &\leq l_i l_j \\
u_j x_i + u_i x_j - x_k &\leq u_i u_j \\
l_j x_i + u_i x_j - x_k &\geq u_i l_j \\
u_j x_i + l_i x_j - x_k &\geq l_i u_j
\end{aligned} \tag{3.5}$$

for the bilinear constraint (3.3). The `presolve` and `nodePresolve` routines find tighter bounds, $l_i \leq u_i$, on the variables based on the bounds of the other two variables. A lower bound on x_k, for example, is $\min\{l_i l_j, u_i l_j, l_i u_j, u_i u_j\}$. The `isFeasible` routine checks whether a given candidate solution satisfies all the constraints of the form (3.2) and (3.3). If the given point $\hat{x}$ has $\hat{x}_i^2 > \hat{x}_k$, then the `separate` routine generates a valid inequality

$$2\hat{x}_i x_i - x_k \leq \hat{x}_i^2.$$

If $\hat{x}_i^2 < \hat{x}_k$ for constraint (3.2) or $\hat{x}_i \hat{x}_j \neq \hat{x}_k$ for constraint (3.3), then the `getBranchingCandidates` routine returns x_i and x_j as candidates for branching. If one of these is selected for branching by the `Brancher`, then the `branch` routine creates two branches by modifying the bounds on the branching variable.

3.1.4 Transformer

It is frequently the case that the specific structure a handler supports, such as the quadratic constraints (3.2) and (3.3), may occur in more complicated functions of a given problem formulation. To enable application of handlers for those structures, we first need to transform a problem into a form that is equivalent to the original problem, but exposes these simpler structures. The `Transformer` class performs this task. The default implementation of the `Transformer` traverses the DAG of a nonlinear function in a depth-first search and adds new variables to create simple constraints of the desired form. The process can be better explained with an example. Consider the constraint

$$x_1(x_1 + x_2) + (x_1 + x_2)^2 \leq 1.$$

Its computational graph was earlier shown in Figure 2. The `Transformer` reformulates it as

$$\begin{aligned}
x_3 &= x_1 + x_2, \\
x_4 &= x_1 x_3, \\
x_5 &= x_3^2, \\
x_4 + x_6 &\leq 1,
\end{aligned}$$

where we have used unary and binary operations for simplicity. The `Transformer` also maintains a list of new variables it introduces and the corresponding expression they represent. These variables are then reused if the same expression is observed in other constraints or the objective function. In this example, if the expression $x_1 + x_2$ is observed in some other computational graph, then x_3 will be reused. Similarly, x_4 and x_5 will be reused. The code uses a simple hashing function to detect these common subexpressions. In addition to applying the transformations, the `Transformer` assigns every nonlinear constraint to a specific `Handler`. Since there are often many alternative reformulations of a problem, a different implementation of the `Transformer` class may lead to a different reformulation and can have significant impact on algorithm performance.

3.1.5 *Utility Classes*

Utility classes provide commonly required functions such as measuring the CPU time (`Timer` class), writing logs (`Logger`), parsing and storing user-supplied options (`Options`), and commonly used operations on vectors, matrices, and intervals (`Operations`).

3.2 `Engine`

The Minotaur framework calls external libraries for solving relaxations or simpler problems. A user can choose to link the Minotaur libraries with several external libraries. Minotaur libraries and executables can also be compiled without any of these external libraries.

The Open-Solver Interface (OSI) library provided by COIN-OR [17] is used to link to the CLP [28], GuRoBi [43], and CPLEX [45] solvers for LP problems. The BQPD [23] and qpOASES [22] solvers can be used to solve QP problems. FilterSQP [26,25] and IPOPT [73] can be used to solve NLP problems using active-set and interior-point methods, respectively.

The interface to each solver is implemented in a separate class derived from the `Engine` abstract base class. For instance, the `BQPDEngine` class implements an interface with the BQPD solver. The two main functions of an `Engine` class are to (1) convert a problem represented by the `Problem` class to a form required by the particular solver and (2) convert and convey the solution and solution status to the Minotaur routines. The conversion of LP and QP problems is straightforward. The engine sets up the matrices and vectors in the format required by the solver before passing them. More general NLP solvers such as FilterSQP and IPOPT require routines to evaluate derivatives. These engine classes implement callback functions. For instance, the `FilterSQPEngine` class implements the functions `objfun`, `confun`, `gradient`, and `hessian` that the FilterSQP solver requires, and the `IpoptEngine` class implements the `eval_f`, `eval_grad_f`, and `eval_h` functions required by the `IpoptFunInterface` class.

In a branch-and-bound algorithm, these engines are called repeatedly with only minor changes to the problem. Therefore, one can use the solution information from the previous invocation to *warm-start* the next. The `BQPDEngine`, `FilterSQPEngine`, and `IpoptEngine` classes implement methods to save and use the warm-starting information from previous calls. The LP solvers already have sophisticated warm-starting techniques, and hence their engines do not need such routines.

3.3 `Interface`

The `Interface` consists of classes that convert MINLP problems written in a modeling language or a software environment to Minotaur's `Problem` class and other associated classes. In the current version of the Minotaur framework, we have an interface only to AMPL. This interface can be used in two modes.

In the first mode, the `AMPLInterface` class reconstructs the full computational graph of each nonlinear function in the problem. The class uses AMPL library functions to parse each graph and then converts it into the form required by the `CGraph` class. Derivatives are provided by the `CGraph` class. Once the problem instance is created, the `AMPLInterface` is no longer required to solve the instance.

In the second mode, we do not store the computational graph of the nonlinear functions. Rather, the `AMPLProblem` class is derived from the `Problem` class and stores linear constraints and objective using the default implementation. Nonlinear constraints are not stored by using the `CGraph` class. Instead, pointers to the nonlinear functions stored by the AMPL library are placed in the `AMPLNonlinearFunction` class derived from the `NonlinearFunction` class. This class calls methods provided by the AMPL solver library to evaluate the function or its gradient at a given point. In this mode, the `AMPLInterface` class provides an object of class

`AMPLJacobian` to evaluate the Jacobian of the constraint set and `AMPLHessian` to evaluate the Hessian of the Lagrangian of the continuous relaxation. This mode is useful when implementing an algorithm that only requires evaluating the values and derivatives of nonlinear functions in the problem.

The interface also implements routines to write solution files so that the user can query the solution and solver status within AMPL.

4 Implementing Basic Solvers in Minotaur

We now describe how the three algorithms presented in Section 2 can be implemented by combining different classes of the Minotaur framework. Our goal here is not to develop the fastest possible solver, but rather to demonstrate that our flexible implementation does not introduce additional computational overhead. These demonstrations are implemented as examples in our source code and are only simplistic versions of the more sophisticated solvers available.

The general approach for implementing a solver is to first read the problem instance by using an `Interface`. The next step is to initialize the appropriate objects of the base or derived classes of `NodeRelaxer`, `NodeProcessor`, `Brancher`, `Engine`, and `Handler`. Then an object of the `BranchAndBound` class is set using these components, and the `solve` method of the `BranchAndBound` object is called to solve the instance. The code for these steps can be written in a C++ `main()` function that can be compiled and linked with the Minotaur libraries to obtain an executable.

4.1 Nonlinear Branch-and-Bound

The NLPBB algorithms for convex MINLPs (Section 2.2) is the simplest to implement. It needs only one handler: `IntVarHandler` to check whether the solution of a relaxation satisfies integer constraints and to return a branching candidate if it does not satisfy the constraints. The `BndProcessor` described in Section 3.1.2 can be used as the node processor, since we need only compare the lower bound of the relaxation with the upper bound at each node. It needs a pointer to an `Engine` to solve the relaxation. We use `FilterSQP` in this example. Since only the bounds on variables of the relaxation change during the tree search, `NodeIncRelaxer` is used to update the relaxation at every node by changing the bounds on appropriate variables. We need not create any relaxation since the variables, constraints, and objective are the same as in the original problem. We use `ReliabilityBrancher` as the brancher, which implements a nonlinear version of reliability branching [3,10]. The pointer to the `Engine` used for `BndProcessor` can be used for this brancher as well. Figure 6 contains the `main()` function used to implement this solver. Additional lines containing directives to include header files are omitted for brevity. We refer to our NLPBB solver as `simple-bnb`.

4.2 LP/NLP-Based Branch-and-Bound

Implementing the LP/NLP-based branch-and-bound method [62] requires us to solve a LP relaxation that is different from that of the original problem. The `QGHandler` described in Section 3.1.3 solves the required NLP to find the point around which linearization constraints are added to the relaxation. `LinearHandler` is used to copy the linear constraints from the original problem to the relaxation. The `NodeProcessor` is required to solve several relaxations at each node if cuts are added. Thus we use the `PCBProcessor` described in Section 3.1.2. The remaining setup is similar to that of the NLPBB algorithm. Figure 7 shows the main portion of the code for implementing this solver. We refer to our LP/NLP-based solver as `simple-qg`.

```
int main(int argc, char** argv) {
  EnvPtr env = (EnvPtr) new Environment();
  HandlerVector handlers;
  int err = 0;

  env->startTimer(err); assert(err==0); // Start timer
  env->getOptions()->findBool("use_native_cgraph")->setValue(true);
  AMPLInterface* iface = new AMPLInterface(env, "bnb");
  ProblemPtr p = iface->readInstance(argv[1]); // read the problem
  p->setNativeDer();

  // create branch-and-bound objects
  BranchAndBound *bab = new BranchAndBound(env, p);
  IntVarHandlerPtr v_h = (IntVarHandlerPtr) new IntVarHandler(env, p);
  handlers.push_back(v_h); // only one handler required

  // setup engine
  EnginePtr e = (FilterSQPEnginePtr) new FilterSQPEngine(env);
  ReliabilityBrancherPtr rel_br = (ReliabilityBrancherPtr) new
                                  ReliabilityBrancher(env, handlers);
  rel_br->setEngine(e);

  // node processor and relaxer
  NodeProcessorPtr npr = (BndProcessorPtr)
                         new BndProcessor(env, e, handlers);
  npr->setBrancher(rel_br); bab->setNodeProcessor(npr);
  NodeIncRelaxerPtr nr = (NodeIncRelaxerPtr)
                         new NodeIncRelaxer(env, handlers);
  RelaxationPtr rel = (RelaxationPtr) new Relaxation(p);
  nr->setRelaxation(rel); nr->setEngine(e); nr->setModFlag(false);
  bab->setNodeRelaxer(nr);
  bab->shouldCreateRoot(false);

  bab->solve();
  bab->writeStats(std::cout);
  bab->getSolution()->writePrimal(std::cout);
}
```

Fig. 6 Excerpt of code for implementing our NLPBB algorithm, `simple-bnb`, using the Minotaur framework and the AMPL function, gradient, and Hessian evaluation routines.

4.3 Global Optimization of Quadratically Constrained Problems

A simple implementation of the branch-and-bound method for solving nonconvex quadratically-constrained quadratic programming (QCQP) problems demonstrates how the Minotaur framework can be used for global optimization. A `Transformer` first converts a given QCQP problem to a form where each quadratic term is assigned to a new variable by means of constraints of the form (3.2) and (3.3). The `QuadHandler` creates the linear relaxations of each of these constraints. Other components of the solver are similar to those described earlier: the `LinearHandler` copies all the linear constraints in the problem, the `IntVarHandler` is used to check integrality constraints, the `NodeIncRelaxer` is used to update the relaxation at each node, and the `PCBProcessor` is used for processing each node. Any of the branchers available in the toolkit can be used for branching. Figure 8 shows the implementation of this solver. We refer to our simple global optimization solver as `simple-glob`.

```
// create branch-and-bound object
BranchAndBound *bab = new BranchAndBound(env, p);
EnginePtr nlp_e = (FilterSQPEnginePtr) new FilterSQPEngine(env);
EnginePtr e = (OsiLPEnginePtr) new OsiLPEngine(env);

// setup handlers
IntVarHandlerPtr v_h = (IntVarHandlerPtr) new IntVarHandler(env, p);
LinearHandlerPtr l_h = (LinearHandlerPtr) new LinearHandler(env, p);
QGHandlerPtr q_h = (QGHandlerPtr) new QGHandler(env, p, nlp_e);
l_h->setModFlags(false, true);
q_h->setModFlags(false, true);
handlers.push_back(v_h);
handlers.push_back(l_h);
handlers.push_back(q_h);

// setup engine for solving relaxations and branching
ReliabilityBrancherPtr rel_br = (ReliabilityBrancherPtr) new
                                ReliabilityBrancher(env, handlers);
rel_br->setEngine(e);

// node processor
NodeProcessorPtr np = (PCBProcessorPtr) new PCBProcessor(env, e,
                                                         handlers);
np->setBrancher(rel_br);
bab->setNodeProcessor(np);

// node relaxer
NodeIncRelaxerPtr nr = (NodeIncRelaxerPtr)
                       new NodeIncRelaxer(env, handlers);
nr->setEngine(e); nr->setModFlag(false);
bab->setNodeRelaxer(nr);
bab->shouldCreateRoot(true);

bab->solve();
```

Fig. 7 Excerpt of code for implementing our LP/NLP-based algorithm, `simple-qg`, using the Minotaur framework.

The `QuadHandler` requires bounds on variables in the quadratic constraints to create a linear relaxation. A `Presolver` is used to obtain bounds on all variables before the relaxation is created. The `Presolver` class, in turn, calls the `presolve` function of each `Handler`.

4.4 Performance of Minotaur

We now present experimental results of the performance of the algorithms presented in this section. We have divided the results into three parts: (1) using automatic differentiation, (2) solving convex MINLPs and (3) solving nonconvex MINLPs. Before describing the experiments we first explain the common setup used for all the numerical results.

Experimental Setup: All experiments were done on a computer with a 2.9MHz Intel Xeon CPU E5-2670 processor and 128GB RAM. Hyperthreading was switched off. A single compute core was used to solve each problem instance within a specified time limit of one hour. Debian-8 Linux was the operating system. The Minotaur, IPOPT, AMPL Solver Library, and OSI-CLP code was compiled with the Gnu "g++" version 4.9.2, while "gfor-

```
ProblemPtr p, newp;

// create branch-and-bound object
BranchAndBound *bab = new BranchAndBound(env, p);
EnginePtr nlp_e = (FilterSQPEnginePtr) new FilterSQPEngine(env);
EnginePtr e = (OsiLPEnginePtr) new OsiLPEngine(env);

// Call transformer. It creates the required handlers.
SimpTranPtr trans = (SimpTranPtr) new SimpleTransformer(env, p);
trans->reformulate(newp, handlers, err); assert(0==err);

// presolve
PresolverPtr pres = (PresolverPtr) new Presolver(newp, env, handlers);
pres->solve();

// brancher
ReliabilityBrancherPtr rel_br = (ReliabilityBrancherPtr) new
                                ReliabilityBrancher(env, handlers);
rel_br->setEngine(e);

// node processor and relaxer
NodeProcessorPtr nproc = (PCBProcessorPtr) new PCBProcessor(env, e,
                         handlers);
nproc->setBrancher(rel_br);
bab->setNodeProcessor(nproc);
NodeIncRelaxerPtr nr = (NodeIncRelaxerPtr) new NodeIncRelaxer(env,
                       handlers);
nr->setEngine(e);
nr->setModFlag(false);
bab->setNodeRelaxer(nr);
bab->shouldCreateRoot(true);

// start solving
bab->solve();
```

Fig. 8 Excerpt of code for implementing our simple global optimization algorithm, `simple-glob`, for QCQP problems using the Minotaur framework.

tran" version 4.9.2 was used for the fortran code including FilterSQP, BQPD, and MUMPS (used by IPOPT). The optimization flag was set to '-O3' for both compilers.

Automatic Differentiation: We first compare the effect of using our own implementation of automatic differentiation on the performance of NLP engines. We set up an experiment in which an NLP engine solves a sequence of NLPs that differ only in bounds on variables. This setup closely mimics the use of NLP engines in a MINLP solver while ensuring that other MINLP routines such as heuristics, presolving, and valid inequalities do not affect the observations.

To benchmark the derivatives produced with our native `CGraph` class, we compare its computation time with that of the AMPL interface derivatives. In particular, we modified `simple-bnb` to replace the `Reliability Brancher` with `LexicoBrancher`. `LexicoBrancher` simply selects the candidate for branching with the smallest index from the list of branching candidates. This solver thus spends almost the entire time solving the NLP relaxations. Table 4 reports the time taken to process 100 nodes of branch and bound (i.e., 100 NLPs) with derivatives from both `CGraph` and `AMPLInterface` on selected instances. We observe no significant difference between the two for nearly 80% of the instances. `CGraph` is slower by a factor of 8 or less for QCQP problems that have a dense Hessian of the Lagrangian as in the case of the smallinvDAXrx and smallinvSNPrx instances. The

Table 4 Time taken to solve 100 NLP relaxations. Derivatives are obtained from Minotaur's AMPL interface (column "AMPL") and Minotaur's `CGraph` class (column "CGraph").

	Time (s)	
Instance	AMPL	CGraph
batchs151208m	28.52	29.96
clay0304m	1.51	1.68
clay0305h	42.36	41.00
flay05h	5.41	5.24
flay06m	1.13	1.15
fo7_2	3.76	3.83
fo9_ar2_1	8.38	8.26
m6	2.88	2.87
m7_ar4_1	5.78	5.93
no7_ar3_1	4.97	5.06
o7_ar5_1	5.34	5.27
rsyn0820m04m	44.47	44.88
rsyn0830m04h	119.22	119.75
slay07m	1.75	1.80
slay09h	10.55	10.82
smallinvDAXr5b150-165	0.71	1.43
smallinvSNPr3b020-022	2.85	17.35
sssd22-08	2.24	2.20
syn30m03m	18.53	18.74
syn40m04h	71.18	75.16

times for these problems could be significantly reduced by extracting and storing the vectors and matrices for the quadratic forms. While speed is important, the main goal of `CGraph` is to enable a user to easily and reliably manipulate nonlinear functions. `CGraph` never failed in evaluating derivatives in all runs, thus demonstrating its reliability in solving MINLPs.

Convex MINLP: We now benchmark the performance of our `simple-bnb` solver described in Section 4.1. We compare its performance in two settings: one using `CGraph` and the other using `AMPLInterface`. We also compare the results with Bonmin's implementation of the branch and bound method. In all three settings, we use the IPOPT solver (with MUMPS and ATLAS-BLAS). We set a limit of one hour for each solver. Of the 356 convex MINLPs available in the MINLPLib-2 collection [13,70], we consider only the 333 instances that have integer variables.

The time taken to solve the instances is compared by using the extended performance profiles [56] shown in Figure 9. The figure shows that our `simple-bnb` method is able to solve almost the same number of instances as is Bonmin (nearly 73% of all instances) within the time limit. On 10% of instances Bonmin is faster than the Minotaur solvers. The simple implementation with the `AMPLInterface` is slower by a factor of 2 or more as compared with the fastest solver for about 12% of the instances. The same is true for 24% of the instances for `CGraph` and 8% for Bonmin.

Figure 10 compares the performance of our `simple-qg` method described in Section 4.2. In this experiment, Bonmin's QG algorithm is selected for solving the problems. The Minotaur solvers use IPOPT for solving NLP problems and OSI-CLP for solving LP problems. The impact of using `CGraph` in this algorithm is much smaller when compared with that of our NLPBB algorithm since fewer NLP relaxations are solved. The performance of the simple implementation is comparable to that of Bonmin even though a gap is discernible in the left half of the graph. This gap is expected because the two examples of Minotaur are nearly identical in behavior and hence only in a few instances is one of them is faster than all other solvers by a large factor.

Global Optimization: We compare the performance of our `simple-glob` method described in Section 4.3 with that of three other solvers: BARON-15.2 (with AMPL interface), Couenne-0.5.6 (with AMPL interface, CLP as the LP solver, IPOPT as the NLP solver, and Cbc as the MILP solver), and SCIP-3.2.0 (with pip interface, SoPlex as the LP solver, and IPOPT as the NLP solver). All solvers were run using default settings. A time limit of one hour was set on each instance. Tests were performed on 266 continuous QP and QCQP problems taken from the MINLPLib-2 collection [13,70]. Figure 11 plots the extended performance profiles. Our `simple-glob` method was able to solve over 50% of the problems in the time limit without using any heuristics, convexity-detection routines, cutting planes, or advanced presolving. Other solvers that use these techniques were able to solve nearly 70% of the instances.

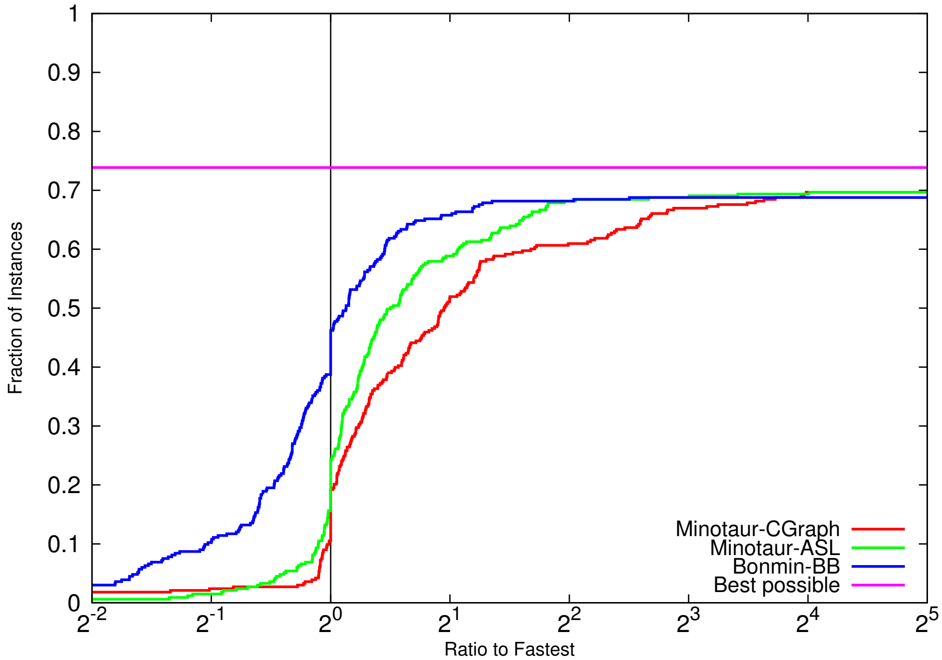

Fig. 9 Extended performance profile based on CPU time for simple NLPBB solvers on 333 convex MINLP problems.

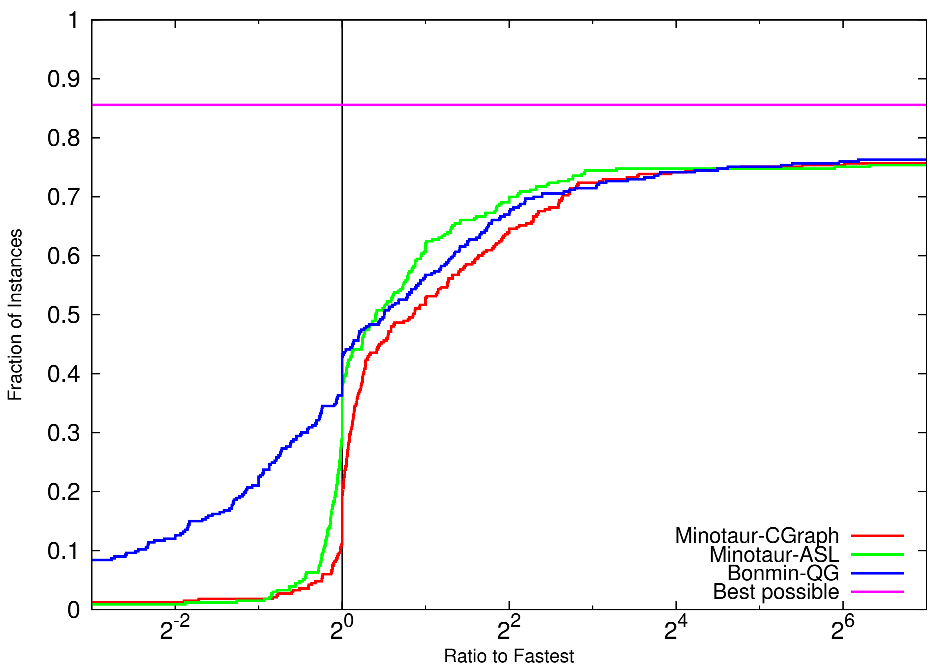

Fig. 10 Extended performance profile based on CPU time of LP/NLP-based solvers on 333 convex MINLP problems.

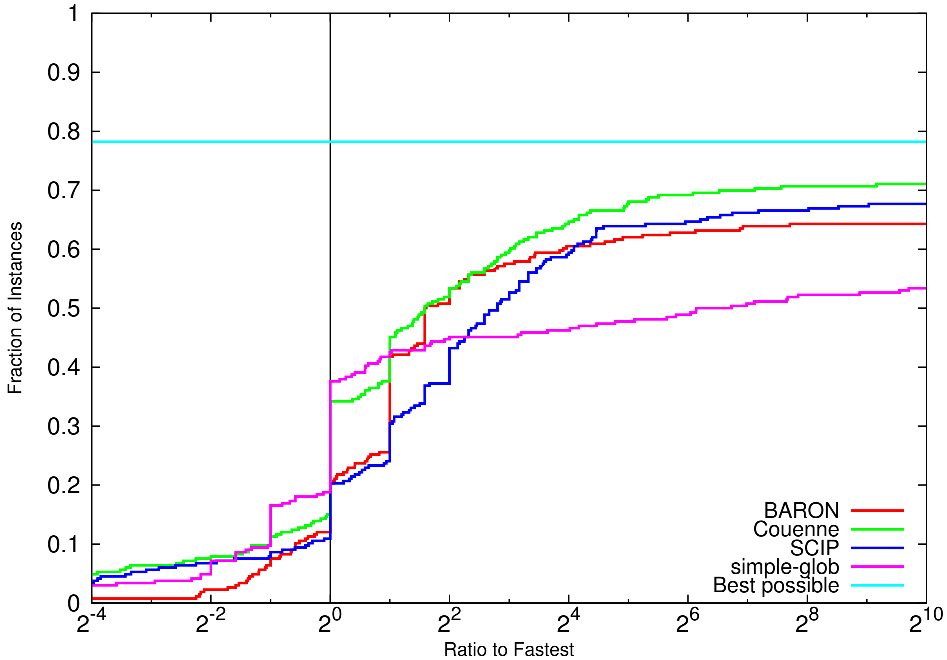

Fig. 11 Extended performance profile based on CPU time of global optimization solvers on 266 QCQP problems.

5 Extensions of Minotaur

In this section, we provide a small set of examples that illustrate how Minotaur can be used and extended to exploit problem specific structures. In each case, we present the main algorithmic idea, show how it is implemented in Minotaur, and give numerical results from our experiments. The experimental setup used for this section is same as that described in Section 4.4.

5.1 Nonlinear Presolve

Our first example illustrates how the availability of a native computational graph (see Section 3) allows us to discover nonlinear structure and to perform certain nonlinear reformulations that are simple but have a dramatic effect on the solution of an important class of MINLP problems.

We consider a chemical engineering problem from the IBM/CMU collection [16], namely, `syn20M04M`. This problem has 160 integer variables, 260 continuous variables, 56 convex nonlinear constraints, and 996 linear constraints. The NLP relaxation can be solved in a fraction of a second. However, standard NLPBB and LP/NLP-based solvers fail to solve this MINLP problem in a two-hour time limit.

Analysis of these models reveals that these problems contain big-M constraints of the form

$$\begin{aligned} c(x_0, x_1, x_2, \ldots, x_k) &\leq M_0(1 - x_0), \\ 0 \leq x_i &\leq M_i x_0, \;\; i = 1, \ldots, k, \\ x_0 \in \{0, 1\}, \;\; & x_i \in \mathbb{R}, \;\; i = 1, \ldots, k. \end{aligned} \tag{5.6}$$

This structure is common in MINLP. The binary variable x_0 acts as a "switch", such that, when $x_0 = 1$ the continuous variables $x_1, \ldots, x_k$ can take nonzero values and the nonlinear constraint is enforced, and when $x_0 = 0$ the continuous variables are forced to zero, and M_0 is large enough so that the nonlinear constraint is redundant. The difficulty for MINLP solvers arises because the upper bound M_0 is not always "tight" when the nonlinear constraint is switched off. In particular, if M_0 is chosen to be too large, the continuous relaxation allows more noninteger solutions, thus making the relaxation weak.

To compute a tighter upper bound M_0, we exploit the implication constraints $0 \le x_i \le M_i x_0, i = 1, \dots, k$. If $x_0 = 0$, then each $x_i = 0$ and we can replace the coefficient M_0 by

$$c^u = c(0, \dots, 0)$$

if $c^u \le M_0$. If $c^u > M_0$, we can fix $x_0 = 1$ since (5.6) is infeasible if $x_0 = 0$. The bound c^u obtained by exploiting the above structure is the tightest possible. The effect of improving the coefficient is dramatic. As shown in Table 5, the solution time for `syn20M04M` is reduced to 2 minutes. To compute the tightest upper bound in the general case, we need to solve the following nonconvex optimization problem:

$$\underset{x}{\text{maximize}} \;\; c(x_0, x_1, x_2, \dots, x_k) \quad \text{subject to} \;\; x_0 = 0, x \in X,$$

where X is a feasible region of the MINLP problem obtained by removing constraint (5.6). This maximization problem can be efficiently solved in some cases (e.g., by exploiting the structure of the constraints).

Table 5 Effects of presolve on the size of instance `syn20M04M` and its solution time.

	No Presolve	Basic Presolve	Full Presolve
Variables	420	328	292
Binary Vars.	160	144	144
Constraints	1052	718	610
Nonlin. Constr.	56	56	56
Bonmin-Bnb (sec)	>3600	NA	NA
Bonmin-QG (sec)	>3600	NA	NA
simple-bnb (sec)	>3600	>3600	66.76
simple-qg (sec)	>3600	>3600	0.49

We also compared the effect of this simple presolve technique on all 96 instances of `RSyn-*` and `Syn-*`. The results are shown in the performance profile in Figure 12. This presolve clearly has a dramatic effect on the solution times for almost all of these instances, increasing the robustness by nearly 20% compared with that of Bonmin and our `simple-bnb` without presolve.

We have implemented these and other more standard presolve techniques, well known from MILP [65], into the core Minotaur library. The `Presolver` class implements the main routine of calling various handlers for exploiting specific structures. In its simplest form, it calls the `Handler::presolve()` function of each of the handlers one by one repeatedly. It stops when no changes are made in the last k calls to handlers, where k is the number of handlers, or when a certain overall iteration limit is reached (see Figure 13 for pseudocode). The handlers used in presolving may be specifically designed for presolving alone and need not implement every function in the `Handler` class (e.g., checking feasibility of a given point or separating a given point). For example, `NlPresHandler` class does not implement any other `Handler` functions besides presolving techniques for general nonlinear constraints and the objective function.

`NlPresHandler` applies four presolve methods to the problem: bound-improvement, identifying redundant constraints, coefficient improvement, and linearization of bilinear expressions on binary variables. In bound improvement, we consider all nonlinear constraints of the type $l_i \le c_i(x) \le u_i$. Using available bounds on the variables x, we find lower and upper bounds (L_i, U_i) of the nonlinear function by propagating the bounds in the forward-mode transversal of the computational graph. We then traverse in reverse mode to update bounds on each node of the graph based on the constraint bounds l_i and u_i. The bounds on variables are thus tightened. Additionally, if $l_i \le L_i$ and $u_i \ge U_i$, then the constraint is identified as redundant and may be dropped. Similarly, if $L_i > u_i$ or $U_i < l_i$ for some i, then no point can satisfy this constraint, and the problem is infeasible.

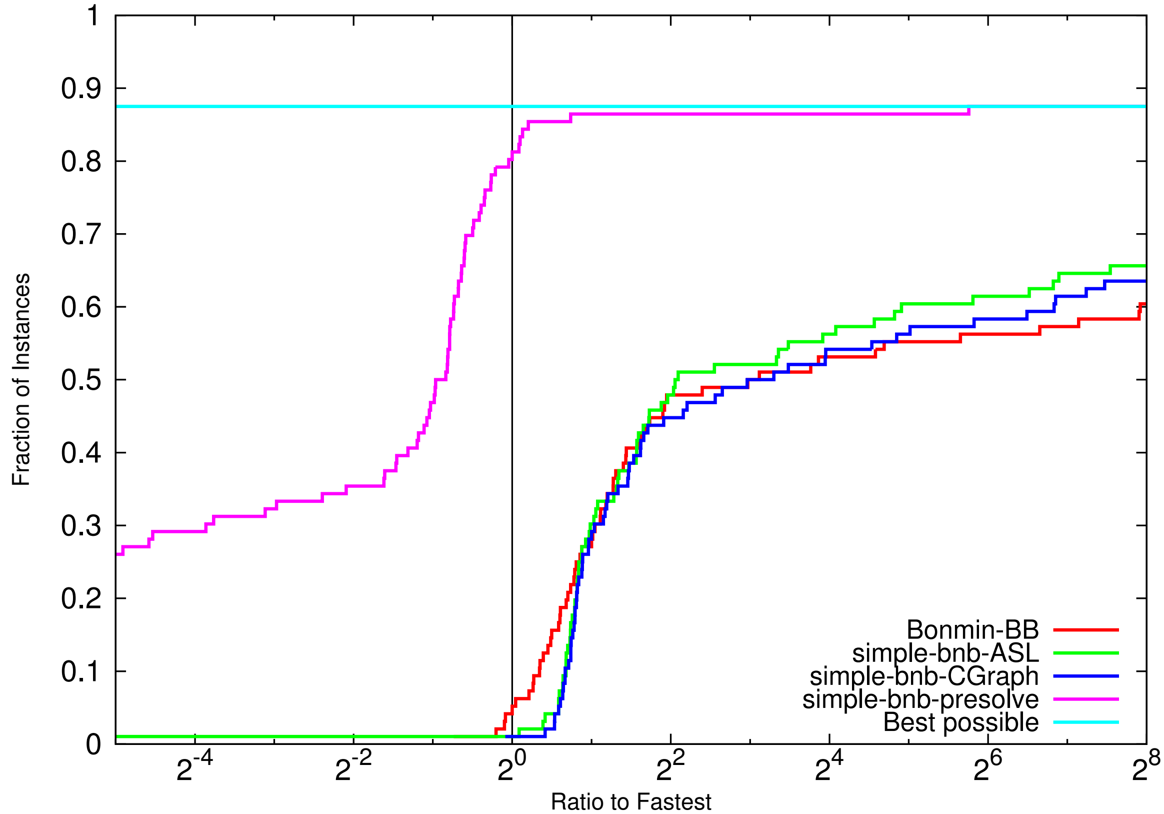

Fig. 12 Performance profile comparing presolve on `RSyn-X` and `Syn-X` instances.

When we are unable to find implications to fix all variables of the function c in (5.6), we find upper bounds on c^u by traversing the computational graph of c as in the bound-improvement technique. If we have a quadratic function of the form $\sum_{i,j} q_{ij}x_ix_j$ in which all variables have finite bounds and every term of the sum contains at least one binary variable, then `NlPresHandler` replaces this expression by an equivalent set of linear constraints. For example, the constraint $y_{1,2} = x_1x_2$, where $x_1 \in \{0,1\}$ and $x_2 \in [l_2, u_2]$, is reformulated as

$$\begin{aligned} y_{1,2} &\leq u_2x_1, \\ y_{1,2} &\geq l_2x_1, \\ x_2 - y_{1,2} &\leq u_2(1-x_1), \\ x_2 - y_{1,2} &\geq l_2(1-x_1). \end{aligned}$$

Coefficient improvement, bound tightening, and identification of redundancy are also applied to linear constraints by the `LinearHandler`. In addition to these functions, the `LinearHandler` class implements dual fixing and identification of duplicate columns (variables) and constraints.

Figure 14 shows the effect of presolving on `simple-bnb` when applied to the 333 convex MINLP instances from MINLPLib-2. In this example, we call the presolver routines only once before the root node is solved in the branch-and-bound portion. Our `simple-bnb` solver with initial presolving was able to solve approximately 10% more instances than did the solver without presolve in the one-hour time limit. Similarly our `simple-qg` solver was able to solve 5% more instances with initial presolving (see Figure 15).

```
SolveStatus Presolver::solve()
{
  iters=0; subiters=0; last_ch_subiter = 0;
  while (iters<iterLimit_){
    ++iters;
    for (HandlerIterator h=handlers_.begin();h!=handlers_.end();++h){
      ++subiters; changed = false;
      h_status = (*h)->presolve(&mods_, &changed);
      if (h_status==SolvedOptimal || h_status==SolvedInfeasible ||
          h_status==SolvedUnbounded) {
        return h_status;
      }
      if (changed == true) last_ch_subiter = subiters;
      else if (subiters>=last_ch_subiter + handlers_.size()){
        return Finished
      }
    }
  }
  return Finished;
}
```

Fig. 13 Code snippet illustrating `Presolver:solve` function from base classes.

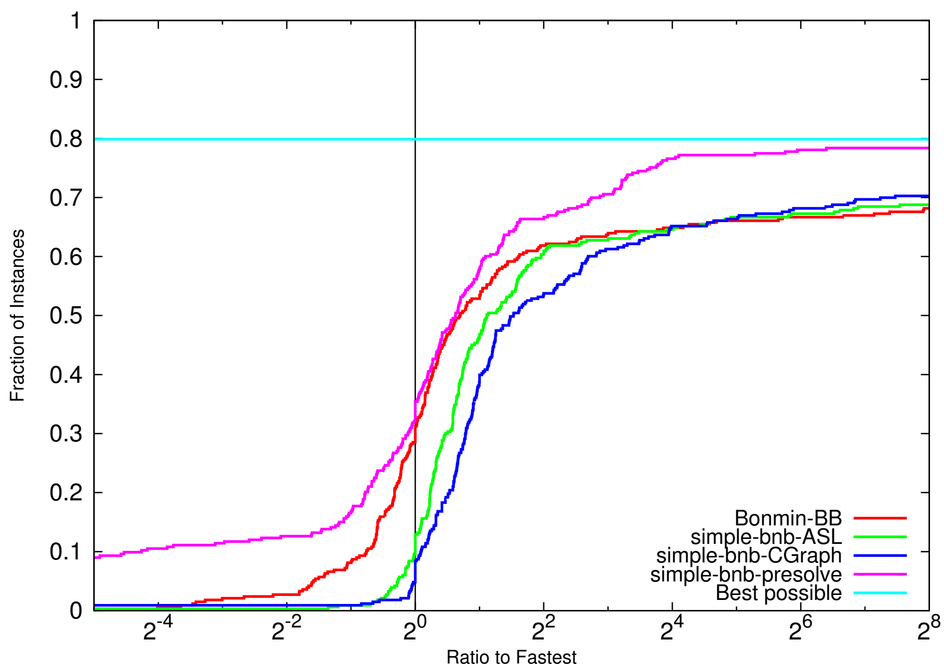

Fig. 14 Performance profile showing the impact of presolving on NLPBB solvers.

5.2 Nonlinear Perspective Formulations

Consider the mixed binary set $\mathcal{S} = \{(x, z) \in \mathbb{R}^n \times \{0, 1\} : c(x) \leq 0, lz \leq x \leq uz\}$, where c is a convex function. Setting the single binary variable z to zero forces all the other continuous variables x to zero. Moreover, the continous relaxation of $\mathcal{S}$ is a convex set. The convex hull of $\mathcal{S}$ can be obtained by taking a perspective reformulation of the nonlinear constraint [31]. More precisely, $conv(\mathcal{S}) = \{(x, z) \in \mathbb{R}^n \times [0, 1] : zc(\frac{x}{z}) \leq 0, lz \leq x \leq uz\}$. Perspective cuts using this reformulation can be derived and were shown to be effective in structured problem

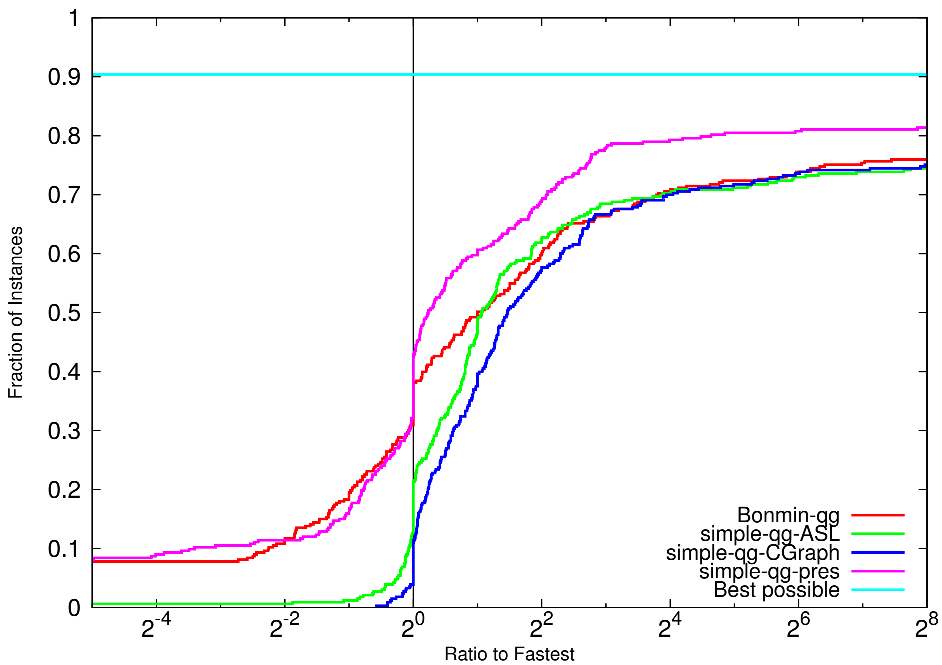

Fig. 15 Performance profile showing the impact of presolving on LP/NLP-based solvers.

instances [31]. Perspective reformulations, cuts derived from these reformulations, and their applications have been studied extensively (see, e.g., [30,32,33,39,40,41]).

Replacing the original constraint $c(x) \leq 0$ with $zc(\frac{x}{z}) \leq 0$ creates difficulties for convex NLP solvers since the new function and its gradient need to be specially defined at $z = 0$ and constraint qualifications fail. The following approximation [34] overcomes these difficulties:

$$(z(1-\epsilon)+\epsilon)\, c\left(\frac{x}{z(1-\epsilon)+\epsilon}\right) \leq 0. \tag{5.7}$$

The function in (5.7) has the same value as that of c when $z = 0$ and $z = 1$. Moreover, this function is convex, and its gradient is well defined at $z = 0$. For small values of ϵ, this reformulation provides a good approximation to the convex hull while being amenable to solution by standard convex NLP solvers.

We implemented a function in the `NlPresHandler` class of Minotaur to identify this structure automatically. For each nonlinear constraint in the problem, Algorithm 3 is used to identify a binary variable for applying the reformulation. A set $\mathcal{C}$ of candidates is initially populated with all binary variables in the problems. For each variable x_i of this nonlinear constraint, we try to find those binary variables that turn x_i off by visiting all linear constraints that reference x_i. All other binary variables are removed from $\mathcal{C}$, and the procedure is repeated for the remaining variables in the nonlinear constraint. If a required binary variable is found, the original function is replaced by its approximate perspective reformulation (5.7) by updating the `CGraph` using a procedure similar to that in Figure 4. This routine is applied only during the initial presolve phase.

Figures 16 and 17 show the effect of using the perspective reformulation with `simple-bnb` and `simple-qg`, respectively, on the convex problems in the MINLPLib-2 collection. Six instances that were not solved in the one-hour time limit with presolve using the `simple-bnb` code were solved after perspective reformulation. As one might expect from a tighter formulation, there was generally a reduction in the number of nodes. However, the reduction in time to solve the problems was less than the reduction in the number of nodes, as the average time spent for solving the NLP relaxations generally increased. Table 6 reports the time spent per NLP solve, the total number of nodes, and the total time to solve such instances. The reported time spent per NLP problem is not

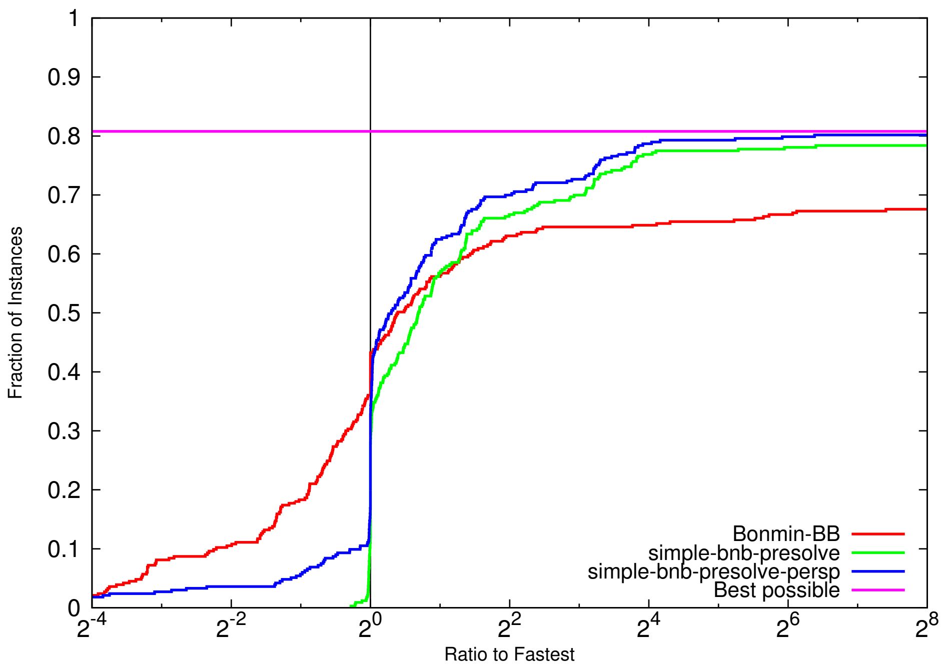

Fig. 16 Performance profile showing the impact of using perspective reformulations with NLPBB algorithms.

equal to the total time divided by the number of nodes because more than one NLP problem is solved at some nodes to find branching candidates.

Algorithm: `Detecting Perspective Structure`

input : A MINLP (1.1) and an index r of a nonlinear constraint in (1.1).
output: A binary variable that can be used for perspective reformulation of $c_r(x) \leq 0$
begin
- $\mathcal{C} \longleftarrow \{t : x_t$ is a binary variable in (1.1)$\}$
- **for** *each* i *such that* x_i *appears in constraint* $c_r \leq 0$ **do**
  - $\mathcal{F}^i \longleftarrow \phi$ (Set of binary variables that turn off x_i)
  - **for** *each* j *such that* x_i *appears in constraint* $c_j(x) \leq 0$ *and* c_j *is a linear function* **do**
    - Let $\mathcal{K}$ be the set of all binary variables in $\mathcal{C}$ that also appear in c_j
    - **for** *each* k *in* $\mathcal{K}$ **do**
      - **if** *Fixing* x_k *to zero in* c_j *forces* x_i *to zero* **then**
        - $\mathcal{F}^i \longleftarrow \mathcal{F}^i \cup \{k\}$
  - $\mathcal{C} \longleftarrow \mathcal{C} \cap \mathcal{F}^i$
  - **if** $\mathcal{C} = \phi$ **then**
    - **return** *no variable found*
- **return** *first element of* $\mathcal{C}$

Algorithm 3: Algorithm for finding a binary variable that can be used to perform a perspective reformulation of a given nonlinear constraint.

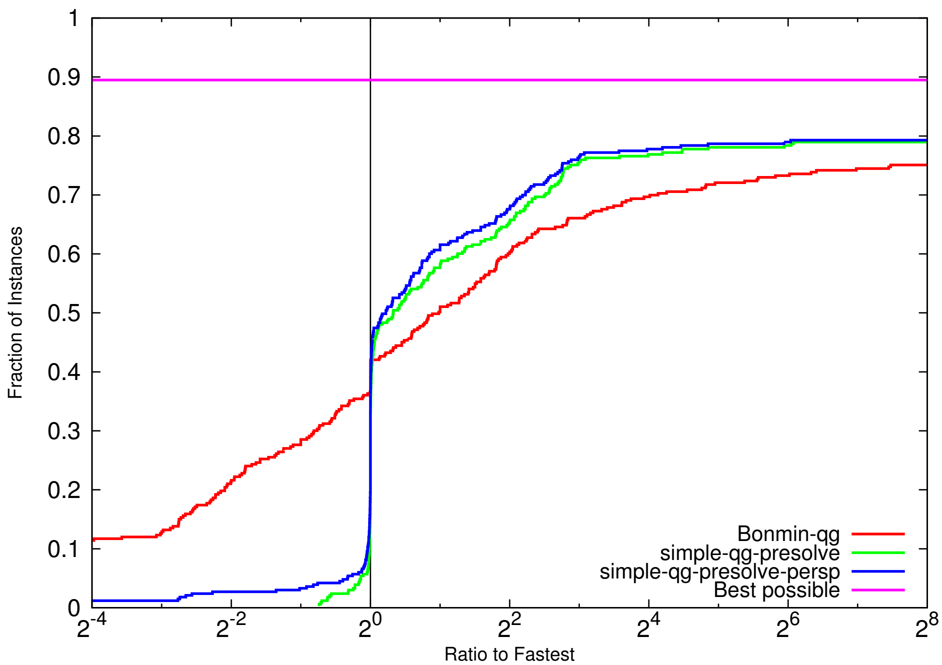

Fig. 17 Performance profile showing the impact of using perspective reformulations with LP/NLP-based methods.

6 Conclusions

Identifying and exploiting structure in MINLP problems are essential when solving difficult problems. A flexible and extensible framework for developing MINLP solvers is necessary in order to rapidly adopt new ideas and techniques and to specialize the methods. The modular class structure of our Minotaur framework provides such capabilities to developers and is the vehicle by which we deliver the resulting numerical methods to users. This flexibility does not come at the cost of speed and efficiency. As demonstrated by our nonlinear presolve and perspective formulation extensions, exploiting the problem structure in Minotaur can result in more reliable and efficient solvers. The source code for Minotaur is available from `https://github.com/minotaur-solver/minotaur`. Because of its availability and extensibility, Minotaur opens the door for future research to further advance the state of the art in algorithms and software for MINLP.

Acknowledgements This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, under contract DE-AC02-06CH11357. This work was also supported by the U.S. Department of Energy under grant DE-FG02-05ER25694.

References

1. K. Abhishek, S. Leyffer, and J. T. Linderoth. FilMINT: An outer-approximation-based solver for nonlinear mixed integer programs. *INFORMS Journal on Computing*, 22:555–567, 2010. DOI:10.1287/ijoc.1090.0373.
2. T. Achterberg. SCIP: Solving constraint integer programs. *Mathematical Programming Computation*, 1(1):1–41, 2009.
3. T. Achterberg, T. Koch, and A. Martin. Branching rules revisited. *Operations Research Letters*, 33:42–54, 2004.
4. C. S. Adjiman, I. P. Androulakis, and C. A. Floudas. A global optimization method, αBB, for general twice-differentiable constrained NLPs - II. implementation and computational results. *Computers & Chemical Engineering*, 22:1159–1179, 1998.
5. E. W. L. Beale and J. A. Tomlin. Special facilities in a general mathematical programming system for non-convex problems using ordered sets of variables. In J. Lawrence, editor, *Proceedings of the 5th International Conference on Operations Research*, pages 447–454, 1969.
6. P. Belotti. COUENNE: A user's manual. Technical report, Lehigh University, 2009.

Table 6 Effect of using perspective reformulation on three measures: the total time taken (seconds) to solve an instance, the number of nodes processed in branch and bound, and the average time taken per NLP problem (milliseconds).

	Without Perspective Ref.			With Perspective Ref.		
Instance	Total Time (s)	Nodes	Time/NLP (ms)	Total Time (s)	Nodes	Time/NLP (ms)
rsyn0805m	33.38	1103	24.35	30.59	802	29.96
rsyn0805m02m	873.88	7444	103.20	774.96	5387	121.67
rsyn0805m03m	2259.38	12303	163.44	1943.15	8485	195.62
rsyn0805m04m	>3600.00	13832	227.78	>3600.00	8814	332.13
rsyn0810m	26.96	786	26.53	31.44	680	36.07
rsyn0810m02m	1030.63	8404	108.99	1014.40	5874	146.51
rsyn0810m03m	>3600.00	16642	194.67	1879.59	6257	239.25
rsyn0810m04m	>3600.00	10875	273.70	2049.40	3741	368.59
rsyn0815m	23.70	505	30.09	29.38	495	41.32
rsyn0815m02m	1259.49	8617	128.12	1342.10	6294	185.64
rsyn0815m03m	>3600.00	15640	206.27	2783.14	8113	289.00
rsyn0815m04m	>3600.00	9568	296.64	>3600.00	6407	428.55
rsyn0820m	35.34	822	30.35	25.02	406	42.52
rsyn0820m02m	1742.62	10084	152.14	1190.27	5027	194.82
rsyn0820m03m	>3600.00	13476	230.98	>3600.00	9099	335.58
rsyn0820m04m	>3600.00	7378	355.32	>3600.00	5380	489.58
rsyn0830m	30.86	576	35.67	27.46	301	55.55
rsyn0830m02m	3214.76	17294	171.81	1007.66	3358	225.94
rsyn0830m03m	>3600.00	10329	289.15	2993.20	5492	423.63
rsyn0830m04m	>3600.00	5420	439.46	>3600.00	4063	609.82
rsyn0840m	71.12	1292	44.93	43.04	452	67.40
rsyn0840m02m	>3600.00	16059	202.69	1234.59	3670	266.33
rsyn0840m03m	>3600.00	8205	336.70	1404.32	1924	415.15
rsyn0840m04m	>3600.00	4148	485.52	>3600.00	3438	680.76
syn05m	0.07	5	7.78	0.06	3	12.00
syn05m02m	0.32	7	15.24	0.20	3	22.22
syn05m03m	0.75	11	22.73	0.41	5	31.54
syn05m04m	1.25	15	30.49	0.68	7	40.00
syn10m	0.06	3	12.00	0.09	3	18.00
syn10m02m	1.32	13	24.91	0.58	5	34.12
syn10m03m	3.48	21	40.94	1.33	9	53.20
syn10m04m	6.11	27	57.01	2.39	13	72.42
syn15m	0.12	5	13.33	0.12	3	24.00
syn15m02m	0.60	5	35.29	0.47	3	52.22
syn15m03m	1.64	9	56.55	0.99	5	76.15
syn15m04m	2.97	13	80.27	1.87	7	110.00
syn20m	0.50	15	15.15	0.26	5	23.64
syn20m02m	2.60	15	43.90	1.87	7	60.32
syn20m03m	20.48	79	72.85	4.06	13	94.42
syn20m04m	66.76	150	109.38	7.33	19	133.27
syn30m	1.11	27	19.47	0.54	7	31.76
syn30m02m	6.88	23	64.30	3.39	9	82.68
syn30m03m	57.29	148	100.49	7.93	11	130.00
syn30m04m	210.63	389	146.65	16.06	17	188.94
syn40m	2.00	43	25.97	1.04	11	41.60
syn40m02m	23.04	80	81.13	9.00	11	103.45
syn40m03m	272.23	859	135.93	18.70	34	179.81
syn40m04m	318.25	460	194.50	49.74	63	260.42

7. P. Belotti, C. Kirches, S. Leyffer, J. T. Linderoth, J. Luedtke, and A. Mahajan. Mixed-integer nonlinear optimization. *Acta Numerica*, 22:1–131, 5 2013.

8. H. Y. Benson. Mixed integer nonlinear programming using interior point methods. *Optimization Methods and Software*, 26(6):911–931, 2011.

9. P. Bonami, L. T. Biegler, A. R. Conn, G. Cornuéjols, I. E. Grossmann, C. D. Laird, J. Lee, A. Lodi, F. Margot, N. Sawaya, and A. Wächter. An algorithmic framework for convex mixed integer nonlinear programs. *Discrete Optimization*, 5(2):186–204, 2008.

10. P. Bonami, J. Lee, S. Leyffer, and A. Wächter. On branching rules for convex mixed-integer nonlinear optimization. *J. Exp. Algorithmics*, 18:2.6:2.1–2.6:2.31, November 2013.

11. A. Brooke, D. Kendrick, A. Meeraus, and R. Raman. *GAMS, A User's Guide*. GAMS Development Corporation, 1992.

12. M. R. Bussieck and A. Drud. SBB: A new solver for mixed integer nonlinear programming. Talk, OR 2001, Section *Continuous Optimization*, 2001.
13. M. R. Bussieck, A. Drud, and A. Meeraus. MINLPLib – A collection of test models for mixed-integer nonlinear programming. *INFORMS Journal on Computing*, 15(1), 2003.
14. R. H. Byrd, J. Nocedal, and W. A. Richard. KNITRO: An integrated package for nonlinear optimization. In G. Pillo and M. Roma, editors, *Large-Scale Nonlinear Optimization*, volume 83 of *Nonconvex Optimization and Its Applications*, pages 35–59. Springer, 2006.
15. B. Christianson. Automatic Hessians by reverse accumulations. *IMA Journal of Numerical Analysis*, 12(2):135–150, 1992.
16. CMU-IBM cyber-infrastructure for MINLP, 2009. `http://www.minlp.org/`.
17. COIN-OR: Computational Infrastructure for Operations Research, 2014. `http://www.coin-or.org`.
18. R. J. Dakin. A tree search algorithm for mixed programming problems. *Computer Journal*, 8:250–255, 1965.
19. S. Drewes. *Mixed Integer Second Order Cone Programming*. PhD thesis, Technische Universität Darmstadt, 2009.
20. S. Drewes and S. Ulbrich. Subgradient based outer approximation for mixed integer second order cone programming. In *Mixed Integer Nonlinear Programming*, volume 154 of *The IMA Volumes in Mathematics and Its Applications*, pages 41–59. Springer, New York, 2012.
21. M. A. Duran and I. Grossmann. An outer-approximation algorithm for a class of mixed-integer nonlinear programs. *Mathematical Programming*, 36:307–339, 1986.
22. H. J. Ferreau, C. Kirches, A. Potschka, H. G. Bock, and M. Diehl. qpOASES: A parametric active-set algorithm for quadratic programming. *Mathematical Programming Computation*, 6(4):327–363, 2014.
23. R. Fletcher. User manual for BQPD, 1995. University of Dundee.
24. R. Fletcher and S. Leyffer. Solving mixed integer nonlinear programs by outer approximation. *Mathematical Programming*, 66:327–349, 1994.
25. R. Fletcher and S. Leyffer. User manual for filterSQP, 1998. University of Dundee Numerical Analysis Report NA-181.
26. R. Fletcher and S. Leyffer. Nonlinear programming without a penalty function. *Mathematical Programming*, 91:239–270, 2002.
27. C. A. Floudas. *Deterministic Global Optimization: Theory, Algorithms and Applications*. Kluwer Academic Publishers, 2000.
28. J. Forrest. CLP, 2014. `http://www.coin-or.org/`.
29. R. Fourer, D. M. Gay, and B. W. Kernighan. *AMPL: A Modeling Language for Mathematical Programming*. The Scientific Press, 1993.
30. A. Frangioni, F. Furini, and C. Gentile. Projected perspective relaxations: A project and lift approach. *Computational Optimization and Applications*, 2015. Published Online.
31. A. Frangioni and C. Gentile. Perspective cuts for a class of convex 0-1 mixed integer programs. *Mathematical Programming*, 106:225–236, 2006.
32. A. Frangioni and C. Gentile. SDP diagonalizations and perspective cuts for a class of nonseparable MIQP. *Operations Research Letters*, 35:181–185, 2007.
33. A. Frangioni and C. Gentile. A computational comparison of reformulations of the perspective relaxation: SOCP vs cutting planes. *Operations Research Letters*, 37(3):206–210, 2009.
34. K. Furman, I. Grossmann, and N. Sawaya. An exact MINLP formulation for nonlinear disjunctive programs based on the convex hull. Presented at the 20th International Symposium on Mathematical Programming, Chicago, IL, 2009.
35. D. M. Gay. More AD of nonlinear AMPL models: Computing Hessian information and exploiting partial separability. In M. Berz, C. Bischof, G. Corliss, and A. Griewank, editors, *Computational Differentiation Techniques Applications and Tools*. SIAM, 1996.
36. A. M. Geoffrion. Generalized Benders decomposition. *Journal of Optimization Theory and Applications*, 10(4):237–260, 1972.
37. A. Griewank and A. Walther. *Evaluating Derivatives Principles and Techniques of Algorithmic Differentiation*. SIAM, Philadelphia, second edition, 2008.
38. I. E. Grossmann and Z. Kravanja. 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*, Springer, New York, 1997.
39. O. Günlük and J. Linderoth. Perspective relaxation of mixed integer nonlinear programs with indicator variables. In A. Lodi, A. Panconesi, and G. Rinaldi, editors, *IPCO 2008: The Thirteenth Conference on Integer Programming and Combinatorial Optimization*, volume 5035, pages 1–16, 2008.
40. O. Günlük and J. T. Linderoth. Perspective relaxation of mixed integer nonlinear programs with indicator variables. *Mathematical Programming Series B*, 104:186–203, 2010.
41. O. Günlük and J. T. Linderoth. Perspective reformulation and applications. In *IMA Volumes*, volume 154, pages 61–92, 2012.
42. O. K. Gupta and A. Ravindran. Branch and bound experiments in convex nonlinear integer programming. *Management Science*, 31:1533–1546, 1985.
43. Gurobi Optimization, Inc. *Gurobi Optimizer Reference Manual, Version 5.6*, 2014.
44. W. E. Hart, J-P. Watson, and D. L. Woodruff. Pyomo: modeling and solving mathematical programs in Python. *Mathematical Programming Computations*, 3:219–260, 2011.
45. IBM Corp. *IBM ILOG CPLEX V12.6: User's Manual for CPLEX*, 2014.
46. R. G. Jeroslow. There cannot be any algorithm for integer programming with quadratic constraints. *Operations Research*, 21(1):221–224, 1973.

47. R. Kannan and C. L. Monma. On the computational complexity of integer programming problems. In R. Henn, B. Korte, and W. Oettli, editors, *Optimization and Operations Research*, volume 157 of *Lecture Notes in Economics and Mathematical Systems*, pages 161–172. Springer, 1978.
48. A. H. Land and A. G. Doig. An automatic method for solving discrete programming problems. *Econometrica*, 28:497–520, 1960.
49. J. Lasserre. An explicit exact SDP relaxation for nonlinear 0-1 programs. In K. Aardal and A. M. H. Gerards, editors, *Integer Programming and Combinatorial Optimization 2001*, Lecture Notes in Computer Science, Vol. 2081, pages 293–303, Berlin, Germany, 2001. Springer.
50. J. Lasserre. Global optimization with polynomials and the problem of moments. *SIAM Journal on Optimization*, 11(3):796–817, 2001.
51. M. Laurent. A comparison of the Sherali-Adams, Lovász-Schrijver, and Lasserre relaxations for 0–1 programming. *Mathematics of Operations Research*, 28(3):470–496, 2003.
52. S. Leyffer. User manual for MINLP-BB, 1998. University of Dundee.
53. S. Leyffer. Mixed-integer PDE-constrained optimization. Technical report, Argonne, 2015.
54. S. Leyffer, J. T. Linderoth, J. Luedtke, A. Miller, and T. Munson. Applications and algorithms for mixed integer nonlinear programming. In *Journal of Physics: Conference Series, SciDAC 2009*, volume 180, page 012014, 2009.
55. A. Mahajan, S. Leyffer, and C. Kirches. Solving convex mixed-integer nonlinear programs by QP-diving. Preprint ANL/MCS-P1801-101, Argonne National Laboratory, October 2010.
56. A. Mahajan, S. Leyffer, and C. Kirches. Solving mixed-integer nonlinear programs by QP-diving. Preprint ANL/MCS-2071-0312, Argonne National Laboratory, Mathematics and Computer Science Division, 2012.
57. G. P. McCormick. Computability of global solutions to factorable nonconvex programs: Part I — Convex underestimating problems. *Mathematical Programming*, 10:147–175, 1976.
58. R. Misener and C. A. Floudas. GloMIQO: Global mixed-integer quadratic optimizer. *Journal of Global Optimization*, pages 1–48, 2012.
59. R. Misener and C. A. Floudas. ANTIGONE: Algorithms for coNTinuous / Integer Global Optimization of Nonlinear Equations. *Journal of Global Optimization*, 59:503–526, 2014.
60. A. Neumaier. Complete search in continuous global optimization and constraint satisfaction. *Acta Numerica*, 13:271–369, 5 2004.
61. J. Nocedal and S. J. Wright. *Numerical Optimization*. Springer-Verlag, New York, 1999.
62. I. Quesada and I. E. Grossmann. An LP/NLP based branch-and-bound algorithm for convex MINLP optimization problems. *Computers & Chemical Engineering*, 16:937–947, 1992.
63. H. S. Ryoo and N. V. Sahinidis. Global optimization of nonconvex NLPs and MINLPs with applications in process design. *Computers & Chemical Engineering*, 19:552–566, 1995.
64. N. V. Sahinidis. BARON: A general purpose global optimization software package. *Journal of Global Optimization*, 8:201–205, 1996.
65. M. W. P. Savelsbergh. Preprocessing and probing techniques for mixed integer programming problems. *ORSA Journal on Computing*, 6:445–454, 1994.
66. H. Schichl. Global optimization in the COCONUT Project. In R. Alt, A. Frommer, R. Baker Kearfott, and W. Luther, editors, *Numerical Software with Result Verification*, volume 2991 of *Lecture Notes in Computer Science*, pages 243–249. Springer Berlin Heidelberg, 2004.
67. C. Still and T. Westerlund. Solving convex MINLP optimization problems using a sequential cutting plane algorithm. *Computational Optimization and Applications*, 34(1):63–83, 2006.
68. M. Tawarmalani and N. V. Sahinidis. *Convexification and Global Optimization in Continuous and Mixed-Integer Nonlinear Programming: Theory, Algorithms, Software, and Applications*. Kluwer Academic Publishers, Boston MA, 2002.
69. T. J. Van Roy. Cross decomposition for mixed integer programming. *Mathematical Programming*, 25:145–163, 1983.
70. S. Vigerske. MINLPLib 2. In *Proceedings of the XII global optimization workshop: Mathematical and Applied Global Optimization*, pages 137–140, 2014.
71. S. Vigerske and A. Gleixner. SCIP: global optimization of mixed-integer nonlinear programs in a branch-and-cut framework. *Optimization Methods and Software*, 2017. Online article.
72. J. Viswanathan and I. E. Grossmann. A combined penalty function and outer-approximation method for MINLP optimization. *Computers & Chemical Engineering*, 14(7):769 – 782, 1990.
73. A. Wächter and L. T. Biegler. On the implementation of a primal-dual interior point filter line search algorithm for large-scale nonlinear programming. *Mathematical Programming*, 106(1):25–57, 2006.
74. T. Westerlund and K. Lundqvist. Alpha-ECP, Version 5.01: An interactive MINLP-solver based on the Extended Cutting Plane Method. Technical Report 01-178-A, Process Design Laboratory at Åbo University, 2001.
75. T. Westerlund and F. Pettersson. A cutting plane method for solving convex MINLP problems. *Computers & Chemical Engineering*, 19:131–136, 1995.
76. Ziena Optimization. *KNITRO Documentation*, Dec. 2012.