



Konrad-Zuse-Zentrum
für Informationstechnik Berlin

Takustraße 7
D-14195 Berlin-Dahlem
Germany

TIMO BERTHOLD¹
AMBROS M. GLEIXNER

UNDERCOVER – a primal heuristic for MINLP based on sub-MIPs generated by set covering

Zuse Institute Berlin, Department of Optimization, Takustr. 7, 14195 Berlin, Germany, {berthold,gleixner}@zib.de

¹Supported by the DFG Research Center MATHEON *Mathematics for key technologies* in Berlin.

Undercover – a primal heuristic for MINLP based on sub-MIPs generated by set covering

TIMO BERTHOLD
AMBROS M. GLEIXNER

ABSTRACT. We present Undercover, a primal heuristic for mixed-integer nonlinear programming (MINLP). The heuristic constructs a mixed-integer *linear* subproblem (sub-MIP) of a given MINLP by fixing a subset of the variables. We solve a set covering problem to identify a minimal set of variables which need to be fixed in order to linearise each constraint. Subsequently, these variables are fixed to approximate values, e.g. obtained from a linear outer approximation. The resulting sub-MIP is solved by a mixed-integer linear programming solver. Each feasible solution of the sub-MIP corresponds to a feasible solution of the original problem.

Although general in nature, the heuristic seems most promising for mixed-integer quadratically constrained programmes (MIQCPs). We present computational results on a general test set of MIQCPs selected from the MINLPLib [12].

1. Introduction

For mixed-integer programming (MIP) it is well-known that, apart from complete solving methods, general-purpose primal heuristics like the feasibility pump [2, 14, 16] are able to find high-quality solutions for a wide range of problems. Over the years, primal heuristics have become a substantial ingredient of state-of-the-art MIP solvers [5, 8]. For mixed-integer nonlinear programming (MINLP) there have only been a few publications on general-purpose primal heuristics [7, 10, 20, 11].

At the heart of many recently proposed primal MIP heuristics, such as Local Branching [15], RINS [13], DINS [17], and RENS [6], lies large neighbourhood search, hence the paradigm of solving a small sub-MIP which promises to contain good solutions. In this paper, we introduce *Undercover*, a large neighbourhood search start heuristic that constructs and solves a sub-MIP of a given MINLP. We demonstrate its effectiveness on a general test set of mixed-integer quadratically constrained programmes (MIQCPs) taken from the MINLPLib [12].

An MINLP is an optimisation problem of the form

$$\begin{aligned} \min \quad & d^T x \\ \text{s.t.} \quad & g_i(x) \leq 0 \quad \text{for } i = 1, \dots, m, \\ & L_k \leq x_k \leq U_k \quad \text{for } k = 1, \dots, n, \\ & x_k \in \mathbb{Z} \quad \text{for } k \in \mathcal{I}, \end{aligned} \tag{1.1}$$

where $\mathcal{I} \subseteq \{1, \dots, n\}$ is the index set of the integer variables, $d \in \mathbb{R}^n$, $g_i : \mathbb{R}^n \rightarrow \mathbb{R}$ for $i = 1, \dots, m$, and $L \in (\mathbb{R} \cup \{-\infty\})^n$, $U \in (\mathbb{R} \cup \{+\infty\})^n$ are lower and upper bounds on the variables, respectively. Note that a nonlinear objective function can always be reformulated by introducing one additional constraint, hence form (1.1) is general.

2. A generic algorithm

The paradigm of fixing a subset of the variables of a given mixed-integer programme in order to obtain subproblems which are easier to solve has proven successful in many primal MIP heuristics such as RINS [13], DINS [17], and RENS [6]. The core difficulty in MIP solving are the integrality constraints, thus in MIP context “easy to solve” usually takes the meaning of few integer variables. While in MINLP integrality constraints do contribute to the complexity of the problem, a specific difficulty are the nonlinearities. Hence, “easy” in MINLP can be understood as few nonlinear constraints.

This revision is a slightly modified version of the initial ZIB-Report 09-40 and has appeared in the proceedings of the *European Workshop on Mixed Integer Nonlinear Programming*, April 2010, Marseille, France.

This research was partially funded by the DFG Research Center MATHEON, Project B20. We thank GAMS Development Corp. for providing us with evaluation licenses for BARON. Many thanks to Stefan Vigerske for his valuable comments and to Mark Wallace for his proofreading.

Keywords: MINLP, MIQCP, primal heuristic, large neighbourhood search, set covering.

Our heuristic is based on the simple observation that by fixing certain variables (to some value within their bounds) any given mixed-integer *nonlinear* programme can be reduced to a mixed-integer *linear* subproblem (sub-MIP). Every feasible solution of this sub-MIP is then a feasible solution of the original MINLP.

Whereas in general it holds that many or even all of the variables might need to be fixed in order to arrive at a linear subproblem, our approach is motivated by the experience that for several practically relevant MIQCPs fixing only a comparatively small subset of the variables already suffices to linearise the problem. The computational effort, however, is usually greatly reduced since we can apply the full strength of state-of-the-art MIP solving to the subproblem. Before formulating a first generic algorithm for our heuristic, consider the following definitions.

Definition 1 (cover of a function). *Let a function $g : D \rightarrow \mathbb{R}$, $x \mapsto g(x)$ on a domain $D \subseteq \mathbb{R}^n$ and a point $x^* \in D$ be given. We call a set $\mathcal{C} \subseteq \{1, \dots, n\}$ of variable indices an x^* -cover of g if and only if the set*

$$\{(x, g(x)) \mid x \in D, x_k = x_k^* \text{ for all } k \in \mathcal{C}\} \quad (2.1)$$

is affine. We call \mathcal{C} a (global) cover of g if and only if \mathcal{C} is an x^ -cover of g for all $x^* \in D$.*

Definition 2 (cover of an MINLP). *Let P be an MINLP of form (1.1), let $x^* \in [L, U]$, and $\mathcal{C} \subseteq \{1, \dots, n\}$ be a set of variable indices of P . We call \mathcal{C} an x^* -cover of P if and only if \mathcal{C} is an x^* -cover for g_1, \dots, g_m . We call \mathcal{C} a (global) cover of P if and only if \mathcal{C} is an x^* -cover of P for all $x^* \in [L, U]$.*

A first generic algorithm for our heuristic is given in Figure 1. The hinge of the algorithm is clearly found in Line 5: finding a suitable cover of the given MINLP. Section 3 elaborates on this in detail with special emphasis on the case of MIQCPs.

FIGURE 1. Generic algorithm

```

1 Input: MINLP  $P$  as in (1.1)
2 begin
3   compute a solution  $x^*$  of an approximation or relaxation of  $P$ 
4   round  $x_k^*$  for  $k \in \mathcal{I}$ 
5   determine an  $x^*$ -cover  $\mathcal{C}$  of  $P$ 
6   solve the sub-MIP of  $P$  given by  $x_k = x_k^*$ ,  $k \in \mathcal{C}$ 
7 end

```

To obtain suitable fixing values for the selected variables, an approximation or relaxation of the original MINLP is used. For integer variables the approximate values are rounded. Most exact solvers for MINLP are based on branch-and-bound. If the heuristic is embedded within a branch-and-bound solver, using its (linear or nonlinear) relaxation appears as a natural choice for obtaining approximate variable values.

Large neighbourhood search heuristics which rely on fixing variables typically have to trade off between eliminating many variables in order to make the sub-MIP tractable versus leaving enough degrees of freedom such that the sub-MIP is still feasible and contains good solutions. Often their implementation inside a MIP solver demands a sufficiently large percentage of variables to be fixed to arrive at an easy to solve sub-MIP [5, 6, 13, 17].

For our heuristic, the situation is different since we do not aim at eliminating integrality constraints, but nonlinearities. In order to linearise a given MINLP, in general we may be forced to fix integer and continuous variables. Especially the fixation of continuous variables in an MINLP can introduce a significant error, even rendering the subproblem infeasible. Thus our heuristic will aim at fixing as *few* variables as possible to obtain as large a linear subproblem as possible.

Remark 1. Note that in general a minimum cover does not necessarily yield a dimension-wise largest sub-MIP that can be obtained by fixing variables in a given MINLP. First, this is because in our definition we do not look at the feasible region given by a constraint, but at the graph of

its left hand side. Second, we do not take into account the interrelation of constraints with each other and with variable bounds and integrality constraints. Through these interrelations, fixing one variable may lead to further domain reductions and variable fixations which can not immediately be foreseen when looking at each nonlinearity separately.

Hence, searching for a minimum cover may be understood as an approximate method for determining dimension-wise maximal sub-MIPs. Propagation routines using these interrelations can, however, be very effectively integrated within the heuristic, see “fix-and-propagate” in Section 4.

Remark 2. We point out the links of our general-purpose approach for MINLP to the works on *bilinearly constrained bilinear programmes* in global optimization. These are defined as quadratically constrained quadratic programmes which allow for a partition of the variable set into two parts such that each quadratic term is bilinear with one variable from each part. Holding the variables in either set fixed, per definition one obtains a linear programme, a simple property which has been used extensively in various solution approaches.

3. Finding minimum covers

This section describes our method for determining a minimum cover of an MINLP, i.e. a minimal subset of variables to fix in order to linearise each constraint. The idea for Undercover originated from our work on solving MIQCPs. Since its application also appears most promising for this class of problems, we start by presenting conditions for covers of quadratic constraints.

Covering quadratic functions. Suppose we are given a quadratic function $g : \mathbb{R}^n \rightarrow \mathbb{R}, x \mapsto x^T Q x$, with $Q \in \mathbb{R}^{n \times n}$ symmetric. Let $x^* \in \mathbb{R}^n$ and $\mathcal{C} \subseteq \{1, \dots, n\}$. Fixing $x_k = x_k^*$ for all $k \in \mathcal{C}$ transforms $x^T Q x$ into $y^T \tilde{Q} y + \tilde{q}^T y + \tilde{c}$ with variable vector $y = (x_k)_{k \notin \mathcal{C}} \in \mathbb{R}^{n-|\mathcal{C}|}$, the restricted matrix $\tilde{Q} = (Q_{uv})_{u,v \notin \mathcal{C}}$ of dimension $(n - |\mathcal{C}|) \times (n - |\mathcal{C}|)$, the vector $\tilde{q} = (2 \sum_{u \in \mathcal{C}} Q_{uk} x_u^*)_{k \notin \mathcal{C}} \in \mathbb{R}^{n-|\mathcal{C}|}$, and offset $\tilde{c} = \sum_{u,v \in \mathcal{C}} Q_{uv} x_u^* x_v^*$. Thus, the set (2.1) is affine if and only if \tilde{Q} vanishes, i.e. if $q_{uv} = 0$ for all $u, v \notin \mathcal{C}$. In reverse, this means that for \mathcal{C} to be a cover of g , it is necessary and sufficient to contain at least one out of u or v for all nonzero matrix entries Q_{uv} .

This can be interpreted as a *set covering problem*, where items are given by those $(u, v) \in \{1, \dots, n\} \times \{1, \dots, n\}$ with nonzero Q_{uv} , and sets are given by $S(k) := \{(u, v) \mid u = k \text{ or } v = k\}$ for each variable index $k = 1, \dots, n$.

Remark 3. Note that in the quadratic case, any x^* -cover is already a global cover, thus the distinction made in Definitions 1 and 2 is void.

Covering MIQCPs. An MIQCP is an MINLP as in (1.1) where each constraint $i = 1, \dots, m$ takes the form $g_i(x) = x^T A_i x + b_i^T x + c_i \leq 0$ with $A_i \in \mathbb{R}^{n \times n}$ symmetric, $b_i \in \mathbb{R}^n$, and $c_i \in \mathbb{R}$. Matrices A_i are not required to be positive semidefinite, i.e. we allow for nonconvex constraints.

In order to find a cover of a given MIQCP P , we solve the set covering problem outlined above, taking into account all constraints. We introduce auxiliary binary variables $\alpha_k, k = 1, \dots, n$, equal to 1 if and only if x_k is fixed in P . As explained above, $\mathcal{C}(\alpha) := \{k \mid \alpha_k = 1\}$ is a cover of P if and only if

$$\alpha_k = 1 \quad \text{for all } i \in \{1, \dots, m\}, k \in \{1, \dots, n\}, A_{kk}^i \neq 0, L_k \neq U_k, \quad (3.1)$$

$$\alpha_k + \alpha_j \geq 1 \quad \text{for all } i \in \{1, \dots, m\}, k \neq j \in \{1, \dots, n\}, A_{kj}^i \neq 0, L_k \neq U_k, L_j \neq U_j, \quad (3.2)$$

i.e. we require all square terms and one variable in each bilinear term to be fixed. Our heuristic tries to identify as large a linear subproblem as possible. Therefore, we solve the binary programme

$$\min \left\{ \sum_{k=1}^n \alpha_k : (3.1), (3.2), \alpha \in \{0, 1\}^n \right\}. \quad (3.3)$$

The following lemma summarises our discussion from above:

Lemma 1. *Let an MIQCP P be given. Then $\alpha \mapsto \mathcal{C}(\alpha) = \{k \mid \alpha_k = 1\}$ gives a one-to-one correspondence between the feasible solutions of (3.3) and covers of P . A solution α^* of (3.3) is optimal if and only if $\mathcal{C}(\alpha^*)$ is a minimum cardinality cover of P .*

Remark 4. Note that the above covering problem is an optimisation version of 2-SAT, hence there is a polynomial-time algorithm for its solution. In our computational experiments the binary programme (3.3) could always be solved within a fraction of a second by a general MIP solver, hence we did not employ a specialised polynomial-time algorithm. Moreover, in the case of general MINLP, a cover generating problem contains more involved constraints.

Covering general MINLPs. Our approach for computing covers of quadratic constraints can be applied to more general nonlinearities. The immediate generalisations are multilinear and polynomial constraints. Sufficient conditions for a global cover of a monomial $x_1^{p_1} \cdots x_n^{p_n}$, $p_1, \dots, p_n \in \mathbb{N}_0$, are $\alpha_k = 1$ for all $k \in \{1, \dots, n\}$, $p_k \geq 2$, $L_k \neq U_k$, and $\sum_{k:p_k=1, L_k \neq U_k} (1 - \alpha_k) \leq 1$, similar to (3.1) and (3.2).

As can be seen from this example, with more and more general nonlinearities present, more and more variables need to be fixed to arrive at a linear subproblem. However, note that now the notion of an x^* -cover may be much weaker than that of a global cover.¹

4. Variants and extensions

The generic algorithm in Figure 1 can be extended and modified in several ways in order to make Undercover more efficient in practise. This section outlines a few of them, with main focus on avoiding infeasibility of the sub-MIPs.

Fix-and-propagate. Fixing a variable can have great impact on the original problem and the approximation we use. Therefore, we do not fix the variables simultaneously, but sequentially one by one, propagating the bound changes after each fixing. If by that, the approximation solution falls outside the feasible domain of a variable, we instead fix it to the closest bound.² This fix-and-propagate method resembles a method described in [16]. Additionally, we apply it for continuous variables and apply backtracking in case of infeasibility.

Backtracking. If the fix-and-propagate procedure deduces some variable domain to be empty, hence the subproblem to be infeasible, we apply a one-level backtracking, i.e. we undo the last bound change and try other fixing values instead.³

Recovering. During the fix-and-propagate routine, variables outside the precomputed cover may also be fixed. In this case, the fixing of some of the yet unfixed variables in the cover might become redundant and recomputing the cover may yield a smaller number of variable fixings still necessary.

Different covers. Our initial motivation for fixing as few variables as possible was to minimise the impact on the original MINLP. Other measures for the impact of fixing a variable could be used in the objective function of (3.3), such as domain size, appearance in nonlinear terms or nonlinear constraints violated by the approximation solution, variable type, or hybrid measures. In particular, if a minimum cardinality cover yields an infeasible sub-MIP, we may try a cover minimising a different impact measure.

NLP postprocessing. In the spirit of the QCP local search heuristic described in [7], we try to further improve the best sub-MIP solution \tilde{x} by solving the NLP which results from fixing all integer variables to their values in \tilde{x} .

Convexification. The main idea of Undercover is to reduce the computational effort by finding easier to solve subproblems. While here we have focused on sub-MIPs, for nonconvex MINLPs, already a convex sub-MINLP may be significantly easier to solve and contain more and better solutions than a sub-MIP. This modification of Undercover only requires to weaken the constraints in the cover generating problem (3.3) suitably.

¹As a simple example consider the multilinear term $x_1 \cdots x_n$ with no variable bounds. The minimum cardinality of an x^* -cover is 1 as soon as $x_k^* = 0$ for some k . In contrast, the smallest global cover has size $n - 1$.

²Alternatively, we could recompute our approximation to obtain values within the current bounds.

³In our implementation, we try the bounds of the variable, if finite, and the midpoint between the approximation value and each finite bound. For unbounded variables, we try zero and twice the approximation value.

Domain reduction. Instead of fixing the variables in a cover, we could also merely reduce their domains. Especially on continuous variables this leaves significantly more freedom to the subproblem. Since domain propagation is an essential ingredient in MINLP solvers, this might still reduce the computational effort significantly on some problems.

5. Computational experiments

Only few solvers exist that handle general nonconvex MINLPs, such as BARON [21], COUENNE [4], and LindoGlobal [19]. Others, such as BONMIN [9] and SBB [3], guarantee global optimality only for convex problems, but can be used as heuristic solvers for nonconvex problems. Recently, the solver SCIP [1] was extended to solve nonconvex MIQCPs to global optimality [7].

The target of our computational experiments is to demonstrate the potential of Undercover as a start heuristic for MINLPs applied at the root node. We implemented Undercover within the branch-cut-and-price framework SCIP [1] and used SCIP’s linear outer approximation solution for the fixing values. We incorporated the fix-and-propagate, backtracking, and NLP postprocessing features described in Section 4. To perform the fix-and-propagate procedure, we used the standard propagation engine of SCIP. Secondary SCIP instances were used to solve both the cover generating problem and the Undercover sub-MIP.

In our experiments, we ran SCIP with all heuristics other than Undercover switched off, set a node limit of 1, and deactivated cut generation. We set a node limit of 500 both for the covering problem and the sub-MIP. For solving the sub-MIP, we used “emphasis feasibility” and “fast presolving” settings. We used SCIP 1.2.0.4 with CPLEX 12.1 [18] as LP solver and IPOPT 3.7 [22] as NLP solver for the postprocessing. This configuration we refer to as UC.

For comparison, we ran SCIP 1.2.0.4 with CPLEX 12.1 and IPOPT 3.7 in default mode, which applies ten primal heuristics at the root node. We further compared with the state-of-the-art solvers BARON [21] (commercial) and COUENNE [4] (open source). For all solvers, we used node limit 1. Our goal is thus not to compare the Undercover heuristic with SCIP, BARON, and COUENNE as complete solvers – a comparison rather insignificant –, but specifically with the performance of their root heuristics.

As test set we used a selection of 33 MIQCP instances from MINLPLib [12]. We excluded `lop97ic`, `lop97icx`, `pb302035`, `pb351535`, `qap`, and `qapw`, which are linear after the default presolving of SCIP. On the `nuclear` instances, the root LP relaxation of SCIP is often unbounded due to unbounded variables in nonconvex terms of the constraints. In this case, we cannot apply Undercover since no fixing values are available. Due to this, we only included two of those instances, `nuclear14a` and `nuclear14b`, for which the root LP of SCIP is bounded.

Results. The results are shown in Table 1. In column “nnz/var” we state the average number of nonlinear nonzeros, i.e. the number of quadratic terms, per variable as an indication of the nonlinearity of the problem.⁴ In column “% cov” we report the relative size of the cover used by UC as percentage of the total number of variables after preprocessing, and the objective value of the best solution found by Undercover. A star indicates that the sub-MIP was solved to optimality. For all other solvers, we provide the objective value of the best solution found during root node processing. The best solution among the four solvers is marked bold.

The computational results seem to confirm our expectation that often a low fixing rate suffices to obtain a linear subproblem: 12 of the instances in our test set allow a cover of at most 5% of the variables, further 10 instances of at most 15%. On the remaining third of the test set, a minimum cover contains 19–96% of the variables.

UC found a feasible solution for 21 test instances: on 16 out of the 22 instances with a cover of at most 15% of the variables, and on 5 instances in the remaining third of the test set. In comparison, BARON found a feasible solution in 15 cases, COUENNE in 9, SCIP in 14. We note that on 7 instances UC found a solution, although none of BARON, COUENNE, and SCIP did. UC could solve `ex1266` to optimality and `util` to 0.1% primal-dual gap.

⁴Note that since the cover generating problem contains one constraint for each nonlinear nonzero, this corresponds to the ratio of items to sets of the set covering problem solved.

TABLE 1. Computational results on MIQCP instances.

instance	nnz/var	% cov	UC	SCIP	BARON	COUENNE
du-opt	0.95	95.24	4233.8709*	4233.8709	108.331477	41.3038865
du-opt5	0.95	94.74	3407.05415*	14.1684489	–	1226.02232
ex1263	0.34	3.88	30.1*	–	–	–
ex1264	0.36	4.26	15.1*	–	–	–
ex1265	0.38	3.52	15.1*	–	–	15.1
ex1266	0.40	3.03	16.3*	–	–	–
fac3	0.81	78.26	130653857*	–	38328601.6	–
feedtray2	10.70	3.26	–	0	0	–
meanvarx	0.19	23.33	16.9968975*	14.3692129	14.3692321	18.701873
netmod_dol1	0.00	0.30	0*	-0.317294979	0	–
netmod_dol2	0.00	0.38	-0.0780227488*	-0.50468289	0	–
netmod_kar1	0.01	0.88	0*	-0.132809562	0	–
netmod_kar2	0.01	0.88	0*	-0.132809562	0	–
nous1	2.39	19.44	–	–	–	1.567072
nous2	2.39	19.44	–	1.38431729	0.625967412	1.38431741
nuclear14a	4.98	6.43	–	–	–	–
nuclear14b	2.42	6.43	–	–	–	-1.11054393
nvs19	8.00	88.89	–	0	-1098	–
nvs23	9.00	90.00	–	0	-1124.8	–
product	0.17	30.87	–	–	–	–
product2	0.37	26.15	–	–	–	–
sep1	0.40	10.53	-510.080984*	–	-510.080984	-510.080984
space25	0.12	1.04	–	–	–	–
space25a	0.29	5.84	–	–	–	–
spectra2	3.43	35.71	26.6076018*	23.2840887	119.8743	–
tln5	1.39	9.09	15.1*	–	–	14.5
tln6	1.47	7.69	32.3*	–	–	–
tln7	1.53	6.67	30.3*	–	–	–
tln12	1.70	3.99	–	–	–	–
tloss	1.47	7.89	27.3*	–	–	–
tltr	1.10	12.50	61.1333333*	–	–	–
util	0.07	3.13	999.578743*	1000.48517	1006.50609	–
waste	1.10	5.65	693.392795	693.290675	712.301232	–

To evaluate the solution quality of Undercover, for each other solver consider the instances on which both UC and this solver found a solution: On those instances the solution found by Undercover is better than the one found by SCIP in 1 case (equal in 1, worse in 8), better than BARON 4 times (equal 4, worse 3), and better than COUENNE in 1 case (equal in 2, worse in 3 cases).

The overall time for SCIP preprocessing, solving the root LP and applying Undercover was always less than two seconds. Thereof, the time for applying Undercover was always less than 0.2 seconds, except for the instance **waste**, where Undercover ran for 1.1 seconds. The major amount of time was usually spent in solving the sub-MIP. Although the polytope described by (3.3) is not integral, the covering instance could always be solved to optimality in the root node by SCIP’s default heuristics and cutting plane algorithms. We note that in 10 out of the 14 cases where the resulting sub-MIP was infeasible, the infeasibility was already detected during the fix-and-propagate stage. Thus in most cases, no time was invested in vain to try and find a solution to an infeasible subproblem. Also, except for instance **waste**, all feasible sub-MIPs could be solved to optimality within the imposed node limit of 500, which indicates that – with a state-of-the-art MIP solver at hand – the generated subproblems are indeed significantly easier than the full MINLP.

6. Conclusion and future work

Altogether, Undercover seems to be a fast start heuristic, that often produces feasible solutions of reasonable quality. On the chosen test set, the experiments confirmed our expectation that a low fixing rate often suffices to obtain a feasible linear subproblem which is easy to solve. The computational results indicate, that it complements nicely with existing root node heuristics in different solvers.

Future research will focus on fully implementing and testing the described features and variants in Section 4 and experimenting with fixing values from other approximations, especially solutions of standard NLP relaxations.

Bibliography

- [1] T. Achterberg. *Constraint Integer Programming*. PhD thesis, Technische Universität Berlin, 2007.
- [2] T. Achterberg and T. Berthold. Improving the Feasibility Pump. *Disc. Opt.*, Special Issue 4(1):77–86, 2007.
- [3] ARKI Consulting & Development A/S and GAMS Inc. SBB. <http://www.gams.com/solvers/solvers.htm#SBB>.
- [4] P. Belotti, J. Lee, L. Liberti, F. Margot, and A. Waechter. Branching and bounds tightening techniques for non-convex MINLP. Research Report RC24620, IBM, 2008.
- [5] T. Berthold. Primal heuristics for mixed integer programs. Diploma thesis, Technische Universität Berlin, 2006.
- [6] T. Berthold. RENS – Relaxation Enforced Neighborhood Search. ZIB-Report 07-28, Zuse Institute Berlin, 2007.
- [7] T. Berthold, S. Heinz, and S. Vigerske. Extending a CIP framework to solve MIQCPs. ZIB-Report 09-23, Zuse Institute Berlin, 2009.
- [8] R.E. Bixby, M. Fenelon, Z. Gu, E. Rothberg, and R. Wunderling. MIP: Theory and practice – closing the gap. In M.J.D. Powell and S. Scholtes, editors, *Systems Modelling and Optimization: Methods, Theory, and Applications*, pages 19–49. Kluwer Academic Publisher, 2000.
- [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. *Disc. Opt.*, 5:186–204, 2008.
- [10] P. Bonami, G. Cornuéjols, A. Lodi, and F. Margot. A feasibility pump for mixed integer nonlinear programs. *Math. Prog.*, 119(2):331–352, 2009.
- [11] P. Bonami and J.P.M. Gonçalves. Primal heuristics for mixed integer nonlinear programs. Research Report RC24639, IBM, 2008.
- [12] M.R. Bussieck, A.S. Drud, and A. Meeraus. MINLPLib – A Collection of Test Models for Mixed-Integer Nonlinear Programming. *INFORMS Journal on Computing*, 15(1):114–119, 2003.
- [13] E. Danna, E. Rothberg, and C. Le Pape. Exploring relaxation induced neighborhoods to improve MIP solutions. *Math. Prog. A*, 102(1):71–90, 2004.
- [14] M. Fischetti, F. Glover, and A. Lodi. The feasibility pump. *Math. Prog. A*, 104(1):91–104, 2005.
- [15] M. Fischetti and A. Lodi. Local branching. *Math. Prog. B*, 98(1-3):23–47, 2003.
- [16] M. Fischetti and D. Salvagnin. Feasibility pump 2.0. *Math. Prog. C*, 1:201–222, 2009.
- [17] S. Ghosh. DINS, a MIP improvement heuristic. In *Proc. of 12th IPCO*, pages 310–323, 2007.
- [18] ILOG, Inc. CPLEX. <http://www.ilog.com/products/cplex>.
- [19] Lindo Systems, Inc. LindoGlobal. <http://www.lindo.com>.
- [20] G. Nannicini, P. Belotti, and L. Liberti. A local branching heuristic for MINLPs. *ArXiv e-prints*, 2008.
- [21] M. Tawarmalani and N.V. Sahinidis. Global optimization of mixed-integer nonlinear programs: A theoretical and computational study. *Math. Prog.*, 99:563–591, 2004.
- [22] A. Wächter and L.T. Biegler. On the implementation of a primal-dual interior point filter line search algorithm for large-scale nonlinear programming. *Math. Prog.*, 106(1):25–57, 2006.