


JAKOB WITZIG¹ TIMO BERTHOLD STEFAN HEINZ

A Status Report on Conflict Analysis in Mixed Integer Nonlinear Programming

¹  [0000-0003-2698-0767](https://orcid.org/0000-0003-2698-0767)

Zuse Institute Berlin
Takustr. 7
14195 Berlin
Germany

Telephone: +49 30-84185-0
Telefax: +49 30-84185-125

E-mail: bibliothek@zib.de
URL: <http://www.zib.de>

ZIB-Report (Print) ISSN 1438-0064
ZIB-Report (Internet) ISSN 2192-7782

A Status Report on Conflict Analysis in Mixed Integer Nonlinear Programming

Jakob Witzig¹, Timo Berthold² and Stefan Heinz²

¹Zuse Institute Berlin, Takustr. 7, 14195 Berlin, Germany
witzig@zib.de

²Fair Isaac Germany GmbH, Takustr. 7, 14195 Berlin, Germany
{timoberthold, stefanheinz}@fico.com

January 18, 2019

Abstract

Mixed integer nonlinear programs (MINLPs) are arguably among the hardest optimization problems, with a wide range of applications. MINLP solvers that are based on linear relaxations and spatial branching work similar as mixed integer programming (MIP) solvers in the sense that they are based on a branch-and-cut algorithm, enhanced by various heuristics, domain propagation, and presolving techniques. However, the analysis of infeasible subproblems, which is an important component of most major MIP solvers, has been hardly studied in the context of MINLPs. There are two main approaches for infeasibility analysis in MIP solvers: *conflict graph analysis*, which originates from artificial intelligence and constraint programming, and *dual ray analysis*.

The main contribution of this short paper is twofold. Firstly, we present the first computational study regarding the impact of dual ray analysis on convex and nonconvex MINLPs. In that context, we introduce a modified generation of infeasibility proofs that incorporates linearization cuts that are only locally valid. Secondly, we describe an extension of conflict analysis that works directly with the nonlinear relaxation of convex MINLPs instead of considering a linear relaxation. This is work-in-progress, and this short paper is meant to present first theoretical considerations without a computational study for that part.

1 Introduction

In this paper, we consider *mixed integer nonlinear programs* (MINLPs) of the form

$$\min\{c^\top x \mid Ax \geq b, g_k(x) \leq 0 \forall k \in \mathcal{K}, \ell \leq x \leq u, x_j \in \mathbb{Z} \forall j \in \mathcal{I}\} \quad (1)$$

with objective coefficient vector $c \in \mathbb{R}^n$, linear constraint matrix $A \in \mathbb{R}^{m \times n}$, nonlinear constraint functions $g_k: \mathbb{R}^n \mapsto \mathbb{R}$, $k \in \mathcal{K} := \{1, \dots, p\}$, continuously differentiable, and possibly nonconvex, and variable bounds $\ell, u \in \mathbb{R}^n$, where $\overline{\mathbb{R}} := \mathbb{R} \cup \{\pm\infty\}$. Furthermore, let $\mathcal{N} = \{1, \dots, n\}$ be the index set of all variables and $\mathcal{I} \subseteq \mathcal{N}$ the set of variables that need to be integral in every feasible solution. Without loss of generality, we assume the objective function to be linear. A nonlinear objective function can be transformed into a constraint bounded by an artificial variable z that needs to be minimized. We call an MINLP *convex* when all of its constraint functions g_k are convex. Otherwise, we call the MINLP *nonconvex*. When omitting the integrality requirements, we obtain the *nonlinear programming* (NLP) relaxation of (1)

$$\min\{c^\top x \mid Ax \geq b, g_k(x) \leq 0 \forall k \in \mathcal{K}, \ell \leq x \leq u, x \in \mathbb{R}^n\}. \quad (2)$$

The *mixed integer programming* (MIP) relaxation of (1) is given by omitting all nonlinear constraints g_k for all $k \in \mathcal{K}$

$$\min\{c^\top x \mid Ax \geq b, \ell \leq x \leq u, x_i \in \mathbb{Z} \forall i \in \mathcal{I}\}. \quad (3)$$

Omitting both, integrality requirements and nonlinear constraints, yields the *linear programming* (LP) relaxation of (1)

$$\min\{c^\top x \mid Ax \geq b, \ell \leq x \leq u, x \in \mathbb{R}^n\}. \quad (4)$$

All three relaxations provide a lower bound on the optimal solution value of the MINLP (1). MINLP combines discrete decisions and nonlinear functions that are potentially nonconvex. In theory, linear and convex smooth nonlinear programs are solvable in polynomial time [27, 48]. In practice, both classes can be solved very efficiently [10, 42]. In contrast to that, nonconvexities as imposed by discrete variables or nonconvex nonlinear functions easily lead to problems that are both \mathcal{NP} -hard in theory and computationally demanding in practice [49].

Commonly used methods to solve convex MINLPs (1) include the extended cutting plane algorithm (ECP) [52], the extended supporting hyperplane algorithm [31], outer approximation (OA) [17, 19], NLP-based branch-and-bound [23], and LP/NLP-based branch-and-bound [45]. The most commonly used method to solve nonconvex MINLPs is a combination of OA [29, 50] and spatial branch-and-bound [34, 35, 24]. Different MINLP solvers either use LP or MIP relaxations or both during the tree search. For example, *Couenne* [14] and *SCIP* [49] derive valid lower bounds by solving LP relaxations only, whereas *BARON* [28, 5] and *BONMIN* [11, 12] solve both LP and MIP relaxations. In contrast to that, only a handful of MINLP solvers provide the possibility to exclusively use NLP relaxations, e.g., *BONMIN* and *FICO Xpress Optimizer* [18]. For a detailed overview of MINLP solvers that can handle convex and/or nonconvex MINLPs and the implemented algorithm, we refer to [30].

In the following, we will focus on MINLP solvers that use a combination of OA and spatial branch-and-bound. Spatial branch-and-bound is – analogous

to LP-based branch-and-bound [15, 33] – a divide-and-conquer method which splits the search space sequentially into smaller subproblems that are intended to be easier to solve. Additionally, convex relaxations are used to compute lower bounds on the individual subproblems. Based on the computed lower bound, a subproblem can be pruned earlier if the lower bound already exceeds the currently best-known solution. To divide the search space into smaller pieces, spatial branch-and-bound branches on discrete variables with a fractional solution value in the relaxation solution. In addition to that, spatial branch-and-bound uses continuous variables for branching if they appear in nonconvex terms of nonlinear constraints that are violated by the current relaxation solution. During this procedure, infeasible subproblems may be encountered. Infeasibility can either be detected by contradicting variable bounds, derived by domain propagation, or by an infeasible convex relaxation. In contrast to modern MIP solvers that can refer to a variety of well-studied techniques, e.g., [16, 46, 2], to ‘learn’ from infeasible subproblems, similar techniques for MINLPs exist for certain special cases only.

2 Conflict Analysis in MINLP

In this section, we will briefly describe conflict analysis for MIPs of type (3) and the drawbacks when applying these techniques to general MINLP.

2.1 Technical Background: Conflict Analysis in MIP

Conflict analysis for MIP has a long history and has its origin in artificial intelligence [47] and solving satisfiability problems (SAT) [36]. Similar ideas are used in constraint programming (CP), see, e.g., [21, 25]. Integrations of these techniques into MIP were independently suggested by [16], [46], and [2].

If infeasibility is encountered by domain propagation, modern SAT and MIP solvers construct a directed acyclic graph which represents the logic of how the set of branching decisions led to the detection of infeasibility. This graph is called the *conflict graph*. Valid *conflict constraints* can be derived from cuts in the graph that separate the branching decisions from an artificial vertex representing the infeasibility. Based on such a cut, a conflict constraint consists of a set of variables with associated bounds, requiring that in each feasible solution at least one of the variables has to take a value outside these bounds.

If the LP relaxation of a subproblem with local bounds ℓ' and u' turns out to be infeasible, it is necessary to identify a set of variables and bound changes that are sufficient to render the infeasibility. Such a set, the so-called *Farkas proof* [44, 53], can be constructed by using LP duality theory that states that exactly one of the systems

$$Ax \geq b, \ell' \leq x \leq u' \tag{5}$$

$$y^T A + r^T \{\ell', u'\} = 0, y^T b + r^T \{\ell', u'\} > 0, y \geq 0 \tag{6}$$

where $r^\top\{\ell', u'\} := \sum_{j \in \mathcal{N}: r_j > 0} r_j \ell'_j + \sum_{j \in \mathcal{N}: r_j < 0} r_j u'_j$, can be satisfied. System (6) implies a proof of infeasibility w.r.t. to the local bounds

$$0 < y^\top b + r^\top\{\ell', u'\} = y^\top b - (y^\top A)\{\ell', u'\} \iff (y^\top A)\{\ell', u'\} < y^\top b. \quad (7)$$

Consequently, every feasible solution has to satisfy

$$(y^\top A)x \geq y^\top b, \quad (8)$$

which is called Farkas proof; it is a globally valid constraint because it is a nonnegative combination of all globally valid constraints. Thereby, Farkas proofs are a special case of Benders cuts [6]. The Farkas proof is used as a starting point for conflict graph analysis or dual ray analysis. Note, in MIP conflict graph analysis yields at least one conflict that does not need to be linear, whereas dual ray analysis yields exactly one linear constraint.

2.2 Conflict Analysis in MINLP

Only a few publications are dealing with infeasibility in MINLP. Most of the literature is restricted to a certain class of MINLPs, e.g., conic certificates for convex MINLPs [13] which has been proven to be very successful on *mixed-integer second-order cone* (MISOCP) problems. Purely theoretical results for *mixed integer semidefinite programs* (MISDP) were recently published in [26]. Both publications deal with MINLPs that are infeasible as a whole, and not with the analysis of infeasible subproblems to learn information.

For MINLP algorithms that are based on solving LP relaxations, in particular, for OA- and ECP-based solvers, conflict analysis methods for MIP can be applied under certain conditions. To this end, let us first recap the idea of constructing an LP relaxation for an MINLP.

During the tree search, nonlinear functions are approximated by linear functions if they are violated by a relaxation solution. Let \tilde{x} be a relaxation solution with $g_k(\tilde{x}) > 0$. If g_k is convex, a so-called *gradient cut*

$$g_k(\tilde{x}) + \nabla g_k(\tilde{x})(x - \tilde{x}) \leq 0$$

is added. If g_k is nonconvex, convex underestimators are added, see, e.g., [49]. For quadratic functions, e.g., these are the so-called McCormick underestimators [37]. More general nonlinear functions are typically decomposed into functions of a single variable, for which explicit underestimators are known. Note that gradient cuts are globally valid, while underestimators for non-convex functions typically involve the local bounds and are hence not globally valid.

For a subproblem s during the tree search, let $\mathcal{G}^s := \{l_1^s, \dots, l_q^s\}$ be the index set of all linear approximations of all g_k with $k \in \mathcal{K}$ that have been added at the node corresponding to s or any of its ancestors. Hence, it is the current set of (local) *linear relaxation cuts*; all are valid at s . Let G^s be the matrix containing all of these linearizations and d^s be the corresponding right-hand sides. Thus, the LP relaxation solved for subproblem s reads as

$$\min\{c^\top x \mid Ax \geq b, G^s x \geq d^s, \ell \leq x \leq u\}. \quad (9)$$

We denote the set of linearizations added at the root node by \mathcal{G}^0 . During the (spatial) branch-and-bound the set of linearizations expands along each path of the tree: It holds that $\mathcal{G}^0 \subseteq \mathcal{G}^{s_1} \subseteq \dots \subseteq \mathcal{G}^{s_p} \subseteq \mathcal{G}^s$ for each path $(0, s_1, \dots, s_p, s)$. In analogy to solving MIPs, if (9) is infeasible each ray (y, w, r) in its dual can be used to construct a proof of local infeasibility. Here, y_i are the dual variables corresponding to A_i , w_l are the dual variables corresponding to G_l^s for all $l \in \mathcal{G}^s$, and r_j denotes the reduced costs (the duals of the bound constraints) of every variable x_j . Note that $r_j = c_j - y^\top A_{.j} - w^\top G_{.j}^s$.

Hence, a local infeasibility proof w.r.t. the local bounds ℓ' and u' is given by

$$y^\top b + w^\top d^s + r^\top \{\ell', u'\} > 0, \quad (10)$$

In contrast to (8) the constraint $y^\top Ax + w^\top G^s x \geq y^\top b + w^\top d^s$ is not globally valid in general because linearizations of nonlinear constraints might rely on intermediate local bounds. Conflict analysis as introduced in [1, 53] only considers globally valid reasons of infeasibility. Therefore, every local certificate of infeasibility (10) needs to be relaxed to consider \mathcal{G}^0 only

$$y^\top b + \bar{w}^\top d^s + \bar{r}^\top \{\ell', u'\} > 0, \quad (11)$$

where $\bar{w}_l := w_l$, if $l \in \mathcal{G}^0$, and $\bar{w}_l := 0$, otherwise, and $\bar{r}_j := c_j - y^\top A_{.j} - \bar{w}^\top G_{.j}^s$. As a consequence, the relaxed certificate (11) might not provide an infeasibility proof anymore and cannot be used to generate a conflict constraint. If, however, (11) is a valid proof of local infeasibility, all conflict analysis techniques known from MIP can be applied.

2.3 Locally Valid Certificates of Infeasibility

In MIP both conflict graph analysis and dual ray analysis rely on globally valid proofs. In most MIP solvers, local cuts are applied rarely, if at all. This is very different for non-convex MINLP solvers which rely on local linearization cuts. A computational study within the constraint integer programming and MINLP solver SCIP showed that the impact of conflict graph analysis for general MINLPs is almost negligible [49]. A computational study regarding the impact of dual ray analysis on an MINLP solver has – to the best of our knowledge – never been conducted before. We present such a computational study in Section 3.

The observation that conflict graph analysis on MINLP instances has a much smaller impact than on MIP instances [8, 49] led to the assumption that a substantial amount of infeasibility proofs of form (11) were not globally valid. Hence, they are not suitable for conflict graph analysis as known from the literature and implemented in SCIP. These results indicate that locally added linearization cuts are, non-surprisingly, important to render infeasibility w.r.t. local bounds.

Table 1: Aggregated results on MINLPLIB

	#	solved	time	nodes	time _Q	nodes _Q	confs _{glb}	confs _{loc}
all								
noconflict	1170	689	79.11	3014.25	1.000	1.000	–	–
confgraph	1170	694	77.94	2952.07	0.985	0.979	9679.01	–
dualray	1170	695	76.78	2871.86	0.970	0.953	1359.92	–
dualray-loc	1170	698	76.35	2841.90	0.965	0.943	1338.65	3192.50
[100,tilim]								
noconflict	99	83	638.34	86860.54	1.000	1.000	–	–
confgraph	99	88	563.06	74251.69	0.882	0.855	23653.88	–
dualray	99	89	458.28	62890.08	0.718	0.724	2019.46	–
dualray-loc	99	92	429.31	59629.05	0.673	0.686	2086.62	3177.67

To incorporate local linearizations of nonlinear constraints we propose to generalize dual infeasibility proofs of subproblem s with local bounds ℓ' and u' as described in Section 2.1 to locally valid certificates of form

$$y^\top b + \hat{w}^\top d^s + \hat{r}^\top \{\ell', u'\} > 0, \quad (12)$$

incorporating linearizations $\hat{\mathcal{G}}$ with $\mathcal{G}^0 \subseteq \hat{\mathcal{G}} \subseteq \mathcal{G}^{s_p}$, $\hat{w}_l := w_l$, if $l \in \hat{\mathcal{G}}$, and $\hat{w}_l := 0$, otherwise, and $\hat{r}_j := c_j - y^\top A_{.j} - \hat{w}^\top G_{.j}^s$. The certificate (12) is valid for the search tree induced by subproblem q , where q is chosen to satisfy

$$q = \min_{q \in \{1, \dots, s_p\}} \{\mathcal{G}^{q-1} \subseteq \hat{\mathcal{G}}, \hat{\mathcal{G}} \cap (\mathcal{G}^{q+1} \setminus \mathcal{G}^q) = \emptyset\}. \quad (13)$$

Hence, the infeasibility proof might be lifted to an ancestor q of the subproblem s it was created for, if all local information used for the proof were already available at q . Note that it would be possible to apply conflict graph analysis to (12), too. However, this would introduce a computational overhead because the order of locally apply bound changes and separated local linearizations needs to be tracked and maintained. Since conflict graph analysis already comes with an overhead due to maintaining the so-called delta-tree, i.e., complete information about bound deductions and its reasons within the tree, we omit applying conflict graph analysis on locally valid infeasibility certificates.

3 Computational Study

For our computational study, we implemented the generation of locally valid infeasibility certificates in the academic constraint integer programming solver SCIP [22]. In the following, we refer to SCIP with (global) conflict graph analysis as `confgraph` and SCIP with (global) dual ray analysis as `dualray`. Moreover, we refer to `dualray` extended by locally valid infeasibility proofs as `dualray-loc`. As a baseline we use SCIP with deactivated conflict analysis (`noconflict`). As a test set we use the MINLPLIB [40] without instances for which at least one setting finished with numerical violations. This yields a test set of 1170 instances. The experiments were run on a cluster of Intel Xeon

E5-2690 2.6 GHz machines with 128 GB of RAM; a time limit of 3600 seconds was set.

Aggregated results of all four settings are shown in Table 1. Here, [100,tilim] denotes the set of instances for which all settings need at least 100 seconds and are solved by at least one setting [4].

All settings with activated conflict analysis improve both the running time of SCIP, the number of branch-and-bound nodes, and the number of solved instances. Moreover, there seems to be a clear ordering: `dualray-loc` is superior to `dualray` which in turn is superior to `confgraph`. Further, the harder the instances are, the more performance is gained by `dualray` and `dualray-loc` compared to `confgraph`. The number of locally added conflict constraints (`confsloc`) by `dualray-loc` is on average larger than the amount of globally added conflict constraints (`confsglb`) but in the same order of magnitude. On the set of nonconvex MINLPs, however, `dualray-loc` constructs 11.08 times more locally than globally valid conflict constraints. These results indicate that locally added linearizations of nonlinear constraints are important to render local infeasibility.

When looking into the generation of local proofs into detail, we could observe that in 5% of all analyzed infeasible LPs no local cut was needed to construct a valid infeasibility certificate, i.e., we could lift the local conflict to a global one. For 14% of all local proofs we found a set of local cuts such that $q = \lfloor s/2 \rfloor$, the conflict could be lifted up at least half of the depth. 78% of the local proofs could not be lifted. Since a lot of infeasibility information is lost, we propose to use a nonlinear relaxation instead. The theoretical base for nonlinear conflict analysis will be discussed in the following section, whereas the implementation and a computational study is future work.

4 Outlook and Theoretical Thoughts

In this final section, we will discuss theoretical considerations how conflict analysis can be directly applied to a nonlinear relaxation of convex MINLPs. The content described in the following is work-in-progress. At the beginning of this paper, we argued that after LP/MIP-based branch-and-bound, another common method to solve MINLPs is NLP-based branch-and-bound. We will briefly sketch how a generalization of LP infeasibility analysis can be derived from the KKT-conditions of convex NLPs. Given a convex MINLP of form

$$\min_{x \in X} \{f(x) \mid g_k(x) \leq 0 \ \forall k \in \mathcal{K}, \ h_e(x) = 0 \ \forall e \in \mathcal{E}\}, \quad (14)$$

where f, g_k are convex, continuously differentiable functions over \mathbb{R}^n and h_e are affine functions. For every optimal solution x^* of (14) of the (convex) NLP relaxation of (14) there exist $\lambda \geq 0$ such that it holds that

$$\nabla f(x^*) + \sum_{k \in \mathcal{K}} \lambda_k \nabla g_k(x^*) + \sum_{e \in \mathcal{E}} \mu_e \nabla h_e(x^*) = 0, \quad \lambda_k g_k(x^*) = 0. \quad (15)$$

These conditions raise from the so-called *Karush-Kuhn-Tucker-Conditions* [32]. Equality (15) is the gradient of the *Lagrangian dual* that reads as

$$\mathcal{L}(x, \lambda, \mu) := f(x) + \sum_{k \in \mathcal{K}} \lambda_k g_k(x) + \sum_{e \in \mathcal{E}} \mu_e h_e(x), \quad (16)$$

with $\lambda \geq 0$ and $\mu \in \mathbb{R}^{|\mathcal{E}|}$. By duality theory, the *Lagrangian dual function* which reads as $q(\lambda, \mu) := \sup_{\lambda, \mu} \mathcal{L}(x, \lambda, \mu)$ yields a lower bound on the optimal value of (14). Maximizing $q(\lambda, \mu)$ would give the tightest lower bound of (14), and strict duality of convex optimization tells us that this is equivalent to the optimal value of (14). Consequently, if there exists (λ^*, μ^*) such that $\sum_{k \in \mathcal{K}} \lambda_k^* g_k(x) + \sum_{e \in \mathcal{E}} \mu_e^* h_e(x) > 0$, then the dual is unbounded and thus (λ^*, μ^*) proofs infeasibility of (14). Even though Slater regularity does not hold for infeasible points ¹,

$$\sum_{k \in \mathcal{K}} \lambda_k^* g_k(x) + \sum_{e \in \mathcal{E}} \mu_e^* h_e(x) \leq 0 \quad (17)$$

is a valid inequality for (14); it is a convex combination (defined by the dual multipliers) of the constraints of (14). Inequality (17) is the convex optimization equivalent of the Farkas proof (8).

Assume that constraint (17) is given as proof of infeasibility for a subproblem within an NLP-based branch-and-bound. If no local cuts are involved in the infeasibility proof, inequality (17) is a globally valid convex nonlinear constraint. Note in this context that gradient cuts are globally valid.

Clearly, inequality (17) holds for all non-negative λ^* . The following observation makes the concrete (λ^*, μ^*) from the infeasibility proof interesting to use as global information inside a branch-and-bound tree search for convex MINLP. Consider the linearization at an infeasible point x^*

$$\nabla g_k(x^*)^\top (x - x^*) \leq 0 \quad \Leftrightarrow \quad \nabla g_k(x^*)^\top x \leq \nabla g_k(x^*)^\top x^* \quad \forall k \in \mathcal{K}. \quad (18)$$

Then, the corresponding dual multipliers λ^* give the (linear) Farkas proof

$$\sum_{k \in \mathcal{K}} \lambda_k^* \nabla g_k(x^*) + \sum_{e \in \mathcal{E}} \mu_e^* \nabla h_e(x^*) = 0 \quad (19)$$

$$\sum_{k \in \mathcal{K}} \lambda_k^* \nabla g_k(x^*)^\top x^* + \sum_{e \in \mathcal{E}} \mu_e^* \nabla h_e(x^*)^\top x^* < 0. \quad (20)$$

Hence, as in the case of dual ray analysis for MIP, inequality (17) is a single inequality that would have provided the infeasibility proof from its derivative. The hope (which is true for MIP) is that it is a good candidate to detect infeasibility by propagation (under the use of integrality information) in other parts of the search tree, and might be a meaningful aggregation of problem constraints to create cuts from.

¹If one wanted to assume regularity on the constraint functions of (14), linear independence constraint classification would be applicable.

For many NLP solvers, in particular *dual active set methods* [41, 43, 20] and barrier algorithms [38, 39, 51], dual multipliers will be readily available. The added advantage of active set methods is that they typically yield a sparse dual weight-vector (λ, μ) . This might come in handy when the local bounds involved in the infeasibility proof should be used to seed a conflict graph analysis. Like in the linear case, the problem is that the initial reason will typically be too large to be meaningful.

All of this is subject to further investigation. We plan to implement NLP-based conflict analysis into the academic constraint integer programming solver SCIP and to study its impact on solver behavior. As in the MIP case, infeasibility information might be used in several other contexts, consider hybrid branching [3], conflict-driven diving heuristics [54], and also rapid learning [7, 9].

Acknowledgments

We thank Zsolt Csizmadia for his valuable comments on Section 4. The work for this article has been conducted within the Research Campus Modal funded by the German Federal Ministry of Education and Research (fund number 05M14ZAM). We thank three anonymous reviewers for their valuable suggestions and helpful comments.

References

- [1] T. Achterberg. Conflict analysis in mixed integer programming. *Discrete Optimization*, 4(1):4–20, 2007.
- [2] T. Achterberg. Constraint integer programming, 2007.
- [3] T. Achterberg and T. Berthold. Hybrid branching. In W.-J. van Hoes and J. N. Hooker, editors, *Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems, 6th International Conference, CPAIOR 2009*, volume 5547 of *Lecture Notes in Computer Science*, pages 309–311. Springer Berlin Heidelberg, May 2009.
- [4] T. Achterberg and R. Wunderling. Mixed integer programming: Analyzing 12 years of progress. In M. Jünger and G. Reinelt, editors, *Facets of Combinatorial Optimization: Festschrift for Martin Grötschel*, pages 449–481. Springer Berlin Heidelberg, 2013.
- [5] BARON. <https://minlp.com/baron>.
- [6] J. F. Benders. Partitioning procedures for solving mixed-variables programming problems. *Numerische Mathematik*, 4(1):238–252, 1962.
- [7] T. Berthold, T. Feydy, and P. J. Stuckey. Rapid learning for binary programs. In A. Lodi, M. Milano, and P. Toth, editors, *Integration of AI and*

- OR Techniques in Constraint Programming for Combinatorial Optimization Problems, 7th International Conference, CPAIOR 2010*, volume 6140 of *LNCS*, pages 51–55. Springer Berlin Heidelberg, June 2010.
- [8] T. Berthold, A. M. Gleixner, S. Heinz, and S. Vigerske. Analyzing the computational impact of MIQCP solver components. *Numerical Algebra, Control and Optimization*, 2(4):739 – 748, 2012.
 - [9] T. Berthold, P. J. Stuckey, and J. Witzig. Local rapid learning for integer programs. Technical Report 18-56, ZIB, Takustr. 7, 14195 Berlin, 2018.
 - [10] R. E. Bixby. Solving real-world linear programs: A decade and more of progress. *Operations research*, 50(1):3–15, 2002.
 - [11] P. Bonami, L. T. Biegler, A. R. Conn, G. Cornuéjols, I. E. Grossmann, C. D. Laird, J. Lee, A. Lodi, F. Margot, N. Sawaya, et al. An algorithmic framework for convex mixed integer nonlinear programs. *Discrete Optimization*, 5(2):186–204, 2008.
 - [12] Bonmin. <https://projects.coin-or.org/Bonmin>.
 - [13] C. Coey, M. Lubin, and J. P. Vielma. Outer approximation with conic certificates for mixed-integer convex problems. *arXiv preprint arXiv:1808.05290*, 2018.
 - [14] Couenne. <https://www.coin-or.org/Couenne/>.
 - [15] R. J. Dakin. A tree-search algorithm for mixed integer programming problems. *The Computer Journal*, 8(3):250–255, 1965.
 - [16] B. Davey, N. Boland, and P. J. Stuckey. Efficient intelligent backtracking using linear programming. *INFORMS Journal of Computing*, 14(4):373–386, 2002.
 - [17] M. A. Duran and I. E. Grossmann. An outer-approximation algorithm for a class of mixed-integer nonlinear programs. *Mathematical programming*, 36(3):307–339, 1986.
 - [18] FICO Xpress Optimizer. <https://www.fico.com/de/products/fico-xpress-optimization>.
 - [19] R. Fletcher and S. Leyffer. Solving mixed integer nonlinear programs by outer approximation. *Mathematical programming*, 66(1-3):327–349, 1994.
 - [20] A. Forsgren, P. E. Gill, and E. Wong. Primal and dual active-set methods for convex quadratic programming. *Mathematical programming*, 159(1-2):469–508, 2016.
 - [21] M. L. Ginsberg. Dynamic backtracking. *Journal of Artificial Intelligence Research*, 1:25–46, 1993.

- [22] A. Gleixner, M. Bastubbe, L. Eifler, T. Gally, G. Gamrath, R. L. Gottwald, G. Hendel, C. Hojny, T. Koch, M. E. Lübbecke, S. J. Maher, M. Miltenberger, B. Müller, M. E. Pfetsch, C. Puchert, D. Rehfeldt, F. Schlösser, C. Schubert, F. Serrano, Y. Shinano, J. M. Viernickel, M. Walter, F. Wegscheider, J. T. Witt, and J. Witzig. The SCIP Optimization Suite 6.0. Technical Report 18-26, ZIB, Takustr. 7, 14195 Berlin, 2018.
- [23] O. K. Gupta and A. Ravindran. Branch and bound experiments in convex nonlinear integer programming. *Management science*, 31(12):1533–1546, 1985.
- [24] R. Horst and H. Tuy. *Global optimization: Deterministic approaches*. Springer Science & Business Media, 2013.
- [25] Y. Jiang, T. Richards, and B. Richards. No-good backmarking with min-conflict repair in constraint satisfaction and optimization. In *PPCP*, volume 94, pages 2–4. Citeseer, 1994.
- [26] K. Kellner, M. E. Pfetsch, and T. Theobald. Irreducible infeasible subsystems of semidefinite systems. *arXiv preprint arXiv:1804.01327*, 2018.
- [27] L. G. Khachiyan. A polynomial algorithm in linear programming. In *Doklady Akademii Nauk SSSR*, volume 244, pages 1093–1096, 1979.
- [28] M. R. Kılınç and N. V. Sahinidis. Exploiting integrality in the global optimization of mixed-integer nonlinear programming problems with BARON. *Optimization Methods and Software*, 33(3):540–562, 2018.
- [29] G. R. Kocis and I. E. Grossmann. Global optimization of nonconvex mixed-integer nonlinear programming (MINLP) problems in process synthesis. *Industrial & engineering chemistry research*, 27(8):1407–1421, 1988.
- [30] J. Kronqvist, D. Bernal, A. Lundell, and I. Grossmann. A review and comparison of solvers for convex MINLP. *Preprint, Optimization Online*, 2018.
- [31] J. Kronqvist, A. Lundell, and T. Westerlund. The extended supporting hyperplane algorithm for convex mixed-integer nonlinear programming. *Journal of Global Optimization*, 64(2):249–272, 2016.
- [32] H. W. Kuhn and A. W. Tucker. Nonlinear programming. In *Traces and emergence of nonlinear programming*, pages 247–258. Springer, 2014.
- [33] A. H. Land and A. G. Doig. An automatic method of solving discrete programming problems. *Econometrica*, 28(3):497–520, 1960.
- [34] A. H. Land and A. G. Doig. An automatic method for solving discrete programming problems. In *50 Years of Integer Programming 1958-2008*, pages 105–132. Springer, 2010.

- [35] L. Liberti and C. C. Pantelides. Convex envelopes of monomials of odd degree. *Journal of Global Optimization*, 25(2):157–168, 2003.
- [36] J. P. Marques-Silva and K. Sakallah. Grasp: A search algorithm for propositional satisfiability. *Computers, IEEE Transactions on*, 48(5):506–521, 1999.
- [37] G. P. McCormick. Computability of global solutions to factorable nonconvex programs: Part I—Convex underestimating problems. *Mathematical programming*, 10(1):147–175, 1976.
- [38] S. Mehrotra. On the implementation of a primal-dual interior point method. *SIAM Journal on optimization*, 2(4):575–601, 1992.
- [39] C. Mészáros. The BPMPD interior point solver for convex quadratic problems. *Optimization Methods and Software*, 11(1-4):431–449, 1999.
- [40] MINLPLib. Githash 033934c0. <http://www.minlplib.org/>.
- [41] K. G. Murty and F.-T. Yu. *Linear complementarity, linear and nonlinear programming*, volume 3. Citeseer, 1988.
- [42] J. Nocedal and S. Wright. Springer series in operations research and financial engineering. *Numerical Optimization, 2nd edn. Springer, Berlin*, 2006.
- [43] J. Nocedal and S. J. Wright. *Nonlinear Equations*. Springer, 2006.
- [44] I. Pólik. Some more ways to use dual information in MILP. In *International Symposium on Mathematical Programming*, Pittsburgh, PA, 2015.
- [45] I. Quesada and I. E. Grossmann. An LP/NLP based branch and bound algorithm for convex MINLP optimization problems. *Computers & chemical engineering*, 16(10-11):937–947, 1992.
- [46] T. Sandholm and R. Shields. Nogood learning for mixed integer programming. In *Workshop on Hybrid Methods and Branching Rules in Combinatorial Optimization, Montréal*, 2006.
- [47] R. M. Stallman and G. J. Sussman. Forward reasoning and dependency-directed backtracking in a system for computer-aided circuit analysis. *Artificial intelligence*, 9(2):135–196, 1977.
- [48] S. A. Vavasis. Complexity issues in global optimization: a survey. In *Handbook of global optimization*, pages 27–41. Springer, 1995.
- [49] S. Vigerske and A. Gleixner. SCIP: Global optimization of mixed-integer nonlinear programs in a branch-and-cut framework. *Optimization Methods and Software*, 33(3):563–593, 2018.

- [50] 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.
- [51] A. Wächter. Short tutorial: getting started with Ipopt in 90 minutes. In *Dagstuhl Seminar Proceedings*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2009.
- [52] T. Westerlund and F. Pettersson. An extended cutting plane method for solving convex MINLP problems. *Computers & Chemical Engineering*, 19:131–136, 1995.
- [53] J. Witzig, T. Berthold, and S. Heinz. Experiments with conflict analysis in mixed integer programming. In *Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems, 14th International Conference, CPAIOR 2017*, volume 10335 of *LNCS*, pages 211–220. Springer Berlin Heidelberg, May 2017.
- [54] J. Witzig and A. Gleixner. Conflict-Driven Heuristics for Mixed Integer Programming. In preparation.