

MILP Software

Jeffrey T. Linderoth Andrea Lodi

1 Introduction

This article concerns software for solving a general *Mixed Integer Linear Program* (MILP) in the form

$$\min\{c^T x : Ax \geq b, x \geq 0, x_j \in \mathbb{Z} \forall j \in I\}. \quad (1)$$

The algorithmic approach relies on the iterative solution, through general-purpose techniques, of the *Linear Programming* (LP) relaxation

$$\min\{c^T x : Ax \geq b, x \geq 0\}, \quad (2)$$

i.e., the same as problem (1) above but the integrality requirement on the x variables in the set I has been dropped. We denote an optimal solution of problem (2) as x^* . The reason for dropping such constraints is that MILP is NP-hard while LP is polynomially solvable and general-purpose techniques for its solution are efficient in practice.

Entries #1.8.2.1 and #1.8.2.2 cover state-of-the-art LP techniques and solvers, while in the present entry we concentrate on the basic characteristics and components of current, commercial and non-commercial, MILP solvers.

Roughly speaking, using the LP computation as a tool, MILP solvers integrate the *branch-and-bound* and the *cutting plane* algorithms through variations of the general *branch-and-cut* scheme proposed by Padberg & Rinaldi [32] in the context of the *Traveling Salesman Problem* (TSP). Entries #1.3.5.1 and #1.4.1.5 discuss in detail Mathematical-Programming-based approaches for the TSP and the branch-and-cut approach in its full generality, respectively. Nevertheless, with the aim of giving a quick overview and fixing the notation, we briefly discuss in the following both the branch-and-bound and the cutting plane algorithms.

The branch-and-bound algorithm, Land & Doig [26]. In its basic version the branch-and-bound algorithm iteratively partitions the solution space into sub-MILPs (the children nodes) which have the same theoretical complexity of the originating MILP (the father node, or the root node if it is the initial MILP). Usually, for MILP solvers the branching creates two children by using the rounding of the solution of the LP relaxation value of a

fractional variable, say x_j , constrained to be integral

$$x_j \leq \lfloor x_j^* \rfloor \quad \vee \quad x_j \geq \lfloor x_j^* \rfloor + 1. \quad (3)$$

The two children above are often referred to as *left* (or “down”) branch and *right* (or “up”) branch, respectively. On each of the sub-MILPs the integrality requirement on the variables $x_j, \forall j \in I$ is relaxed and the LP relaxation is solved. Despite the theoretical complexity, the sub-MILPs become smaller and smaller due to the partition mechanism (basically some of the decisions are taken) and eventually the LP relaxation is directly integral for all variables in I . In addition, the LP relaxation is solved at every node to decide if the node itself is worthwhile to be further partitioned: if the LP relaxation value is already not smaller than the best feasible solution encountered so far, called *incumbent*, the node can safely be fathomed because none of its children will yield a better solution than the incumbent. Finally, a node is also fathomed if its LP relaxation is infeasible.

The cutting plane algorithm, Gomory [22]. Any MILP can be solved without branching by simply finding its “right” linear programming description, more precisely, the *convex hull* of its (mixed-)integer solutions. In order to do that, one has to iteratively solve the so called *separation problem*

Given a feasible solution x^* of the LP relaxation (2) which is not feasible for the MILP (1), find a linear inequality $\alpha^T x \geq \alpha_0$ which is valid for (1), i.e., satisfied by all feasible solutions \bar{x} of the system (1), while it is violated by x^* , i.e., $\alpha^T x^* < \alpha_0$.

Any inequality solving the separation problem is called a *cutting plane* (or a *cut*, for short) and has the effect of tightening the LP relaxation to better approximate the convex hull.

Gomory [22] has given an algorithm that converges in a finite number of iterations for pure *Integer Linear Programming* (ILP)¹ with integer data. Such an algorithm solves the separation problem above in an efficient and elegant manner in the special case in which x^* is an optimal basis of the LP relaxation. No algorithm of this kind is known for MILPs.

The idea behind integrating the two algorithms above is that LP relaxations (2) do not naturally well approximate, in general, the convex hull of (mixed-)integer solutions of MILPs (1). Thus, some extra work to devise a better approximation by tightening any relaxation with additional linear inequalities (cutting planes) increases the chances that fewer nodes in the search tree are needed. On the other hand, pure cutting plane algorithms

¹ILPs are the special case of MILPs where all variables belong to I , i.e., are constrained to be integer.

show, in general, a slow convergence and the addition of too many cuts can lead to very large LPs which in turn present numerical difficulties for the solvers. The branch-and-cut algorithm has been proven to be very effective initially for combinatorial optimization problems (like TSP) with special-purpose cuts based on a polyhedral analysis and later on in the general MILP context.

The paper is split into two parts. In Section 2 we discuss the basic components of nowadays MILP solvers. In Section 3 we list the available, commercial and non-commercial, solvers with special emphasis to their flexibility in terms of being usable within different modeling and development environments.

Some of the solvers that will be discussed in Section 3 have their own modeling environment and modeling language. Moreover, a recent trend for MILP software producers has been adding additional capabilities to the solvers, like that of solving some special classes of (Mixed Integer) Non Linear Programs (MI)NLPs. Although modeling languages and the possibility of solving MINLPs within a MILP enumerative framework are very interesting topics, they are not discussed in detail in the present paper. The reader is referred to entry #1.8.1.2 for modeling environments and languages and to entries #1.8.2.3, #1.8.2.5 and #1.8.2.6 for NLP and MINLP software.

2 Basic Components of MILP Solvers

Nowadays MILP solvers incorporates key ideas developed during the first 50 years of Integer Programming. In the next sections we will discuss the main ingredients in a concise way. For more details, the reader is referred to Achterberg [1] and Lodi [29] and to many entries in the present encyclopedia.

2.1 Presolving

In the presolving (often called preprocessing) phase the solver tries to detect certain changes in the input that will probably lead to a better performance of the solution process. This is generally done without “changing” the set of optimal solutions of the problem at hand² and it affects two main situations.

On the one side, it is often the case that MILP models can be improved with respect to their given formulation³ by either removing redundant information (variables or constraints) or strengthening the variable bounds generally by exploiting the integrality of variables in I . Modern MILP solvers have the capability of “cleaning up” the models so as to create a presolved

²In fact, the set of optimal solutions might change due to presolve in case, for example, of symmetry breaking reductions, see Margot [30] and entry #1.4.5.1.

³This is especially true for models originated from real-world applications and created by using modeling languages.

instance associated with the original one on which the MILP technology is then applied. The advantage is twofold: first the LP relaxation becomes smaller (then, generally, quicker to solve), and second such relaxations become stronger, i.e., better approximating the convex hull of (mixed-)integer solutions. Finally, more sophisticated presolve mechanisms are also able to discover important implications and sub-structures that might be of fundamental importance later on in the computation for both branching purposes and cutting plane generation.

For a detailed overview of presolving techniques the reader is referred to Savelsbergh [36] and Martin [31].

2.2 Cutting plane generation

The importance of cutting planes has been already pointed out in the introduction: without iteratively strengthening the approximation of the convex hull of (mixed-)integer solutions provided by the LP relaxation, current enumerative schemes would fail in solving difficult MILP instances.

The arsenal of separation algorithms has been continuously enlarged over the years. A large group of cuts, with strong relationship among each other, includes Chvátal-Gomory cuts, Gomory mixed-integer cuts, mixed-integer rounding cuts, $\{0, \frac{1}{2}\}$ cuts, lift-and-project cuts, and split cuts. Essentially, all these inequalities are obtained by applying a disjunctive argument, exploiting the integrality, on an integer or mixed-integer set given by a single constraint generally derived by aggregating many others. This group of cuts is presented in a brilliant and unified way by Cornuéjols [12].

Among the more ‘combinatorial’ cutting planes, i.e., those whose derivation is not directly associated with disjunctions on the integer variables, certainly the most used and useful ones are clique and cover inequalities. Cover inequalities played a fundamental role in pioneering MILP solvers [13, 38].

These classes of inequalities are discussed in entries #1.4.3 and #1.4.4.

2.3 Sophisticated branching strategies

The branching mechanism introduced in Section 1 requires to take two independent and important decisions at any step: node and variable selections.

Node Selection. This is very classical: one extreme is the so called *best-bound first* strategy in which one always considers the most promising node, i.e., the one with the smallest LP value, while the other extreme is *depth first* where one goes deeper and deeper in the tree and starts backtracking only once a node is fathomed, i.e., it is either (mixed-)integer feasible, or LP infeasible or it has a lower bound not better (smaller) than the incumbent. The pros and cons of each strategy are well known: the former explores less nodes but generally maintains a larger tree in terms of memory while

the latter might need to explore a very large number of nodes. (See entry #1.4.1.4 for more details.) Modern software packages typically by default employ a hybrid of these two search techniques.

Variable Selection. The variable selection problem is the one of deciding how to partition the current node, i.e., on which variable to branch on in order to create the two children. For a long time, a classical choice has been branching on the most fractional variable, i.e., in the 0-1 case the closest to 0.5. That rule has been computationally shown [2] to be worse than a random selection. However, it is of course cheap to evaluate. In order to devise stronger criteria one has to do much more work. The extreme is the so called *strong branching* technique (see, e.g., Linderoth & Savelsbergh [28]) in which at any node one has to tentatively branch on each fractional variable and select the one on which the increase in the bound on the left branch times the one on the right branch is maximum. In such a version strong branching is clearly unpractical but its computational effort can be limited by (i) defining a small candidate set of variables to branch on, and (ii) heuristically solving each LP by limiting *a priori* the number of simplex pivots to be performed. Other sophisticated techniques are *pseudocost branching* (see, e.g., Benichou *et al.* [10]) and, the recently introduced *reliability branching* (see, Achterberg, Koch & Martin [2]).

2.4 Primal heuristics

Although primal heuristics have been part of the arsenal of the MILP solvers almost since the beginning, only recently have they become a crucial component. This is likely in response to user demands. Solving an MILP to optimality might be less important in application than providing “good” solutions within short computing times.

Entry #1.7.6.3 is devoted to heuristics for Mixed Integer Programming, and two main classes of primal heuristics are considered.

On the one side, at the root node and often at the internal nodes of the search tree, the solvers apply heuristics with the aim of converting a solution of the LP relaxation into an integer feasible solution. This is obtained through *rounding* an LP solution either component by component (without any backtracking or additional LP solve) or by *diving*, i.e., rounding one or more variables and solving the LP relaxation again, iteratively. The most successful of the algorithms falling in this category is the *feasibility pump* heuristic proposed (for 0-1 MILPs) by Fischetti, Glover & Lodi [17].

On the other hand, once one or more feasible solutions are available, *improvement* heuristics are executed to improve the quality of the current incumbent solution. These heuristics are usually local search methods and we have recently seen a variety of algorithms which solve sub-MILPs for exploring the neighborhood of the incumbent or of a set of solutions. When these sub-MILPs are solved in a general-purpose fashion, i.e., through a

nested call to an MILP solver, this is referred to as *MIPping* [19]. The most successful of the algorithms falling in this category are *local branching* by Fischetti & Lodi [18] and *relaxation induced neighborhood search* by Danna, Rothberg & Le Pape [14].

2.5 Parallel Implementation

The branch-and-bound algorithm is a natural one to parallelize, as nodes of the search tree may be processed independently. There is a long body of research and implementations of parallel branch-and-bound, outlined in the survey [21]. The two types of parallel integer programming research can be loosely categorized based on the type of parallel computing architecture used. Distributed-memory architectures rely on message passing to communicate results of the algorithm. Shared-memory computers communicate information among CPU's by reading from and writing to a common memory pool.

Eckstein [15, 16] was among the first to write an industrial strength parallel branch-and bound-code for general mixed integer programming. Over the past few years, there has been an emergence of *multi-core* computer architectures, so that nearly all modern CPU's contain more than one unit on which computation may be performed. MILP software has followed this trend, so that many packages surveyed below have the ability to run multiple threads of computation, thus making use of multiple cores. For many applications, an important characteristic of the branch-and-bound algorithm is to be able to produce solutions that are *reproducible*. In parallel branch-and-bound, it is well known that the *order* in which node computations are completed can have a significant impact on performance, and often lead to anomalous behavior [25]. An (often) undesirable consequence of this is that users can run the same instance, with the same parameter settings, and achieve very different results in terms of nodes evaluated and CPU time. To combat this undesirable behavior, modern (shared-memory-based) MIP software has introduced appropriate synchronization points in the algorithm to ensure reproducible behavior in a parallel environment. Some overhead is introduced by these synchronization mechanisms.

3 MILP Software

In this section we concentrate on MILP software by first presenting an historical perspective and briefly discussing the issue of the user interface. Then, we review the main characteristics of both commercial and noncommercial MILP software packages.

3.1 Historical Perspective

The earliest commercial grade implementations of branch-and-bound algorithms for solving MILP were done by Beale & Small [9]. For a historical perspective of the development and features of early MILP software packages, such as UMPIRE, SCICONIC, and MPSX/370, the reader is referred to the article of Forrest & Tomlin [20]. These early MILP software packages developed many effective branching and node selection techniques, some of which are still in use today. A major step forward in MILP solution technology occurred in the works of Crowder, Johnson & Padberg [13] and Van Roy & Wolsey [38]. These papers introduced many of the preprocessing and structure-specific cutting plane techniques in use today. Another important work from a computational perspective was by Balas *et al.* [7], demonstrating that the inequalities of Gomory [23] could be effectively used as cutting planes in a branch-and-cut procedure.

Modern MILP codes synthesized many of the ideas present in the literature, taking the best points of these works and engineering them into commercial-grade software packages. The article of Bixby & Rothberg [11] offers a nice historical exposition of how MILP software improved by orders of magnitude over a period of a decade, starting in late 1990's.

3.2 User Interfaces

An important aspect of the design of software for solving MILPs is the user interface. The range of purposes for MILP software is quite large, so it stands to reason that the number of user interface types is also large. In general, users may wish to solve MILPs using the solver as a “black box”, they may wish to call the software from a third-party package, or they may wish to embed the solver into custom applications, which would require software to have a callable library. Often, the user may wish to adapt certain aspects of the algorithm. For instance, the user may wish to provide a custom branching rule or problem-specific valid inequalities. The customization is generally accomplished either through the use of language callback functions, or through an abstract interface wherein the user must derive certain base classes and override default implementations for the desired functions. Nearly all of the software packages surveyed below have standalone executables, callable libraries, and allow the user to customize the branch-and-bound algorithm to varying degrees.

3.3 MILP Software Packages

Any review of software features is inherently limited by the temporal nature of software itself. In fact, algorithmic advances in the field of computational integer programming, many previously described in this article, can quickly render obsolete MILP software that was once state-of-the-art.

Here we give a brief overview of what are the most widely-used MILP software packages at the time of publication. Hans Mittelmann has for many years run independent benchmarks of MILP software, and been publishing the results. In general, the commercial software significantly outperforms non-commercial software on the instances in the Mittelmann suite, but it is difficult to draw strong conclusions about the relative performance of different commercial systems. Results of these benchmarks can be found at <http://plato.asu.edu/ftp/milpf.html> and <http://plato.asu.edu/ftp/milpc.html>. Many of these solvers are available for use free of charge on the NEOS server website <http://www-neos.mcs.anl.gov>.

3.3.1 Commercial Software Packages

Naturally, exact implementation details of specific commercial MILP software packages are not known, as revealing such details would be a competitive disadvantage. In this article, information from the various software user manuals was used to compile features specific to each package. Further, it is impossible to detail all features of a software package, so we focus on those features that are, in the opinion of the authors, particularly noteworthy. Many of the commercial software packages listed here have free or limited cost licensing options available for academic users.

CPLEX

Version	12.1
Website	http://www-01.ibm.com/software/integration/optimization/cplex-optimizer/
Interfaces	C, C++, Java, .NET, Matlab, Python, Microsoft Excel

CPLEX is a commercial MILP software package owned and distributed by IBM. The branch-and-cut algorithm contains a full suite of presolving techniques, cutting planes, search strategies, and heuristic techniques for finding feasible solutions. A special search algorithm in CPLEX, called *dynamic search* can be used instead of traditional branch-and-cut. CPLEX also includes a *mipemphasis* metaparameter that adjusts many algorithmic parameters simultaneously to help users achieve a high-level goal. CPLEX contains facilities for automatically tuning MILP algorithm parameters for a particular class of instances. CPLEX may parallelize the branch-and-bound search using multiple threads in either a deterministic or non-deterministic fashion. A final notable feature of CPLEX is its ability to enumerate multiple solutions that are close to optimality and store them in a *solution pool*.

Gurobi

Version	3.0
Website	www.gurobi.com
Interfaces	C, C++, Java, Python, .NET, Matlab

Gurobi Optimizer contains a relatively new MILP solver that was built from scratch to exploit modern multi-core processing technology. Gurobi contains all advanced algorithmic features, including a variety of cutting planes, heuristics, and search techniques. Notable features of Gurobi include a **MIPFocus** metaparameter that simultaneously controls many algorithmic parameters to achieve a desired goal. Gurobi may execute the branch-and-bound algorithm using multiple computational threads. The parallel search is done in a deterministic fashion. The Gurobi interactive shell is implemented as a set of extensions to the Python shell. Gurobi is also available “on demand” using the Amazon Elastic Compute Cloud.

LINDO

Version	6.1
Website	www.lindo.com
Interfaces	C, Visual Basic, Matlab, Ox

LINDO Systems offers a MILP solver as part of its LINDO API. The MILP solver offers fifteen different types of cutting planes, six different node selection rules, and a variety of preprocessing techniques and branching rules.

Mosek

Version	6.0
Website	www.mosek.com
Interfaces	C, C++, Java, .NET, Python

MOSEK ApS is a company specializing in generic mathematical optimization software, and their suite of software includes a solver for MILP. The solver has six node selection rules, eleven types of cutting planes, the feasibility pump and other heuristics, as well as advanced branching methodologies such as strong branching. Mosek is available for use, though a GAMS interface, on the NEOS Server.

XPRESS-MP

Version	7.0
Website	http://www.fico.com/en/Products/DMTools/xpress-overview/Pages/Xpress-Optimizer.aspx
Interfaces	C, C++, Java, .NET, VBA

FICO Xpress Optimization Suite 7 offers a branch-and-cut-based software for solving MILP. XPRESS-MP offers many advanced cutting planes, including an effective implementation of lift-and project cuts [8]. Modern preprocessing, heuristic, and branching techniques are also available. A unique feature of XPRESS-MP is that it offers an option to branch into general (split) disjunctions, or to search for special structures on which to branch. XPRESS-MP has features to enumerate multiple feasible solutions, and also to tune algorithmic parameters. XPRESS-MP is available for use, with three different input formats, on the NEOS Server.

3.3.2 Noncommercial MILP Packages

Here, we give a cursory description of the most popular non-commercial MILP software packages. The paper of Linderoth & Ralphs [27] contains a more in-depth treatment of non-commercial software packages.

BLIS

License	Common Public License
Version	0.91
Website	https://projects.coin-or.org/CHiPPS
Language	C++

BLIS is an open-source MILP solver available as part of the COIN-OR project. It is built on top of the COIN-OR High-Performance Parallel Search Framework (CHiPPS), so it has been designed to run on distributed memory platforms (one of the few available MILP software packages with this feature). Linear programs are solved using the COIN-OR linear programming Solver (Clp), and cutting planes from the Coin Cut Generation Library (CGL) are used.

CBC

License	Common Public License
Version	2.5
Website	https://projects.coin-or.org/Cbc
Language	C++

CBC is an open-source MILP solver distributed under the COIN-OR project. It is built from many COIN components, including the COIN-OR linear programming Solver (Clp) and the Coin Cut Generation Library (CGL). It is available both as a library and a standalone solver. CBC may be parallelized using shared-memory parallelism, and there are a large number of examples demonstrating its use. CBC is available for use on the NEOS Server.

GLPK

License	GNU General Public License (GPL)
Version	4.44
Website	http://www.gnu.org/software/glpk/
Language	C

GLPK is the GNU Linear Programming Kit, a set of subroutines comprising a callable library and black box solver for MILP. The software distinguishes itself by being available under the GNU GPL and through the large number of community-built interfaces available for its use. GLPK is available for use, though a GAMS interface, on the NEOS Server.

lp_solve

License	GNU lesser general public license (LGPL)
Version	5.5
Website	http://lpsolve.sourceforge.net/5.5/
Language	C

The software `lp_solve` is an open source linear and integer programming solver. There are a large number of interfaces available through which `lp_solve` can be used, including Java, AMPL, MATLAB, O-Matrix, Sysquake, Scilab, Octave, Freemat, Euler, Python, Sage, PHP, and R.

MINTO

License	Given as library only
Version	3.1
Website	http://coral.ie.lehigh.edu/minto/
Language	C

MINTO (Mixed INTeger Optimizer) is a black box solver and solver framework for MILP. MINTO is available only in library form. MINTO relies on external software to solve the linear programming relaxations that arise during the algorithm. Primary development of the software was done in the 1990's. The code was noteworthy at that time for a dynamic preprocessing engine [6] and effective lifted cover inequalities [24]. MINTO is available for use, though an AMPL interface, on the NEOS Server.

SCIP

License	ZIB Academic License
Version	1.2
Website	http://scip.zib.de/
Language	C

SCIP is a MILP software package developed and distributed by a team of researchers at Konrad-Zuse-Zentrum für Informationstechnik Berlin (ZIB). SCIP is also a framework for Constraint Integer Programming and branch-cut-and-price, allowing the user significant control of the algorithm. Development of SCIP was awarded the Beale Orchard Hayes Prize in 2009 [4]. Current benchmarks indicate that SCIP is likely the fastest noncommercial MILP solver. Other references for SCIP include [3, 5]. SCIP is available for use, though a variety of interfaces, on the NEOS Server.

SYMPHONY

License	Common Public License
Version	5.2
Website	http://www.coin-or.org/SYMPHONY/index.htm
Language	C

SYMPHONY is a black box solver and callable library for MILPs. The core solution methodology of SYMPHONY is a customizable branch, cut, and price algorithm that can be executed sequentially or in parallel [34]. SYMPHONY calls on several other open source libraries for specific functionality, including COIN-OR's Cut Generation Library (CGL), and COIN-OR's LP Solver (Clp). There are several unique features of SYMPHONY that are worthy of mention. First, SYMPHONY contains an implementation of an algorithm for solving bicriteria MILPs and methods for approximating the set of Pareto outcomes [33]. SYMPHONY also contains functions for local sensitivity analysis [37]. Second, SYMPHONY has the capability to warm start the branch-and-bound process from a previously calculated

branch-and-bound tree, even after modifying the problem data. The work [35] gives an overview of the SYMPHONY callable library. SYMPHONY is available for use via MPS files on the NEOS Server.

4 Conclusions

As decision processes become increasingly complex, the need for formal tools to aid in both the planning and operation of the underlying system becomes more and more important. Many disciplines have embraced the use of optimization as a central tool to enhance the quality of decisions made. A large fraction of the models used in commercial settings are MILPs. Powerful, reliable, and flexible software systems (both commercial and non-commercial) were a key component of this broad-based acceptance. The field of MILP and computational aspects therein remain an active area of research. Given the high commercial penetration of MILP planning models, undoubtedly future research advances will find themselves implemented in software.

References

- [1] T. Achterberg. *Constraint Integer Programming*. PhD thesis, ZIB, Berlin, 2007.
- [2] T. Achterberg, T. Koch, and A. Martin. Branching roles revisited. *Operations Research Letters*, 33:42–54, 2005.
- [3] T. Achterberg. *Constraint Integer Programming*. PhD thesis, Technischen Universität Berlin, 2007.
- [4] T. Achterberg. SCIP: solving constraint integer programs. *Mathematical Programming Computation*, 1(1):1–41, 2009.
- [5] T. Achterberg, T. Berthold, T. Koch, and K. Wolter. Constraint integer programming: a new approach to integrate CP and MIP. In *Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems*, volume 5015 of *Lecture Notes in Computer Science*, pages 6–20, 2008.
- [6] A. Atamtürk, G. Nemhauser, and M. W. P. Savelsbergh. Conflict graphs in solving integer programming problems. *European Journal on Operations Research*, 121:40–55, 2000.
- [7] E. Balas, S. Ceria, G. Cornuéjols, and N. Natraj. Gomory cuts revisited. *Operations Research Letters*, 19:1–9, 1996.

- [8] E. Balas and M. Perregaard. A precise correspondence between lift-and-project cuts, simple disjunctive cuts, and mixed integer Gomory cuts for 0-1 programming. *Mathematical Programming*, 94:221–245, 2003.
- [9] E. M. L. Beale and R. E. Small. Mixed integer programming by a branch and bound method. In W. H. Kalenich, editor, *Proceedings IFIP Congress 65*, volume 2, pages 450–451, 1965.
- [10] M. Benichou, J. M. Gauthier, P. Girodet, and G. Hentges. Experiments in mixed-integer programming. *Mathematical Programming*, 1:76–94, 1971.
- [11] R. Bixby and E. Rothberg. Progress in computational mixed integer programming – a look back from the other side of the tipping point. *Annals of Operations Research*, 149:37–41, 1007.
- [12] G. Cornuéjols. Valid inequalities for mixed integer linear programs. *Mathematical Programming*, 112:3–44, 2008.
- [13] H. Crowder, E. Johnson, and M. W. Padberg. Solving large scale zero-one linear programming problem. *Operations Research*, 31:803–834, 1983.
- [14] E. Danna, E. Rothberg, and C. Le Pape. Exploiting relaxation induced neighborhoods to improve MIP solutions. *Mathematical Programming*, 102:71–90, 2005.
- [15] J. Eckstein. Parallel branch-and-bound algorithms for general mixed integer programming on the CM-5. *SIAM Journal on Optimization*, 4:794–814, 1994.
- [16] J. Eckstein. Parallel branch-and-bound methods for mixed integer programming. *SIAM News*, 27:12–15, 1994.
- [17] M. Fischetti, F. Glover, and A. Lodi. The feasibility pump. *Mathematical Programming*, 104:91–104, 2005.
- [18] M. Fischetti and A. Lodi. Local branching. *Mathematical Programming*, 98:23–47, 2002.
- [19] M. Fischetti, A. Lodi, and D. Salvagnin. Just MIP it! In V. Maniezzo, T. Stützle, and S. Voss, editors, *MATHEURISTICS: Hybridizing meta-heuristics and mathematical programming*, pages 39–70. Operations Research/Computer Science Interfaces Series. Springer, 2009.
- [20] J. J. H. Forrest and J. A. Tomlin. Branch and bound, integer and non-integer programming. *Annals of Operations Research*, 149:81–87, 2007.

- [21] B. Gendron and T. G. Crainic. Parallel branch and bound algorithms: Survey and synthesis. *Operations Research*, 42:1042–1066, 1994.
- [22] R.E. Gomory. Outline of an algorithm for integer solutions to linear programs. *Bulletin of the American Mathematical Society*, 64:275–278, 1958.
- [23] R.E. Gomory. An algorithm for the mixed integer problem. Technical Report RM-2597, The Rand Corporation, 1960.
- [24] Z. Gu, G. L. Nemhauser, and M. W. P. Savelsbergh. Cover inequalities for 0-1 linear programs: Computation. *INFORMS Journal on Computing*, 10:427–437, 1998.
- [25] T. H. Lai and S. Sahni. Anomalies in parallel branch and bound algorithms. In *Proceedings of the 1983 International Conference on Parallel Processing*, pages 183–190, 1983.
- [26] A. H. Land and A. G. Doig. An automatic method of solving discrete programming problems. *Econometrica*, 28:497–520, 1960.
- [27] J. T. Linderoth and T. K. Ralphs. Noncommercial software for mixed-integer linear programming. In *Integer Programming: Theory and Practice*, pages 253–303. CRC Press Operations Research Series, 2005.
- [28] J. T. Linderoth and M. W. P. Savelsbergh. A computational study of search strategies for mixed integer programming. *INFORMS Journal on Computing*, 11:173–187, 1999.
- [29] A. Lodi. MIP computation. In M. Jünger, T.M. Liebling, D. Naddef, G.L. Nemhauser, W.R. Pulleyblank, G. Reinelt, G. Rinaldi, and L.A. Wolsey, editors, *50 Years of Integer Programming 1958-2008*, pages 619–645. Springer-Verlag, 2009.
- [30] F. Margot. Symmetry in Integer Linear Programming. In M. Jünger, T.M. Liebling, D. Naddef, G.L. Nemhauser, W.R. Pulleyblank, G. Reinelt, G. Rinaldi, and L.A. Wolsey, editors, *50 Years of Integer Programming 1958-2008*, pages 647–686. Springer-Verlag, 2009.
- [31] A. Martin. General mixed integer programming: Computational issues for branch-and-cut algorithms. In M. Jünger and D. Naddef, editors, *Computational Combinatorial Optimization*, volume 2241 of *Lecture Notes in Computer Science*, pages 1–25. Springer-Verlag, Berlin Heidelberg, 2001.
- [32] M. W. Padberg and G. Rinaldi. A branch and cut algorithm for the resolution of large-scale symmetric traveling salesmen problems. *SIAM Review*, 33:60–100, 1991.

- [33] T. Ralphs, M. Saltzman, and M. Wiecek. An improved algorithm for biobjective integer programming. *Annals of Operations Research*, 147:43–70, 2006.
- [34] T. K. Ralphs, L. Ladányi, and M. J. Saltzman. Parallel branch, cut, and price for large-scale discrete optimization. *Mathematical Programming*, 98:253–280, 2003.
- [35] T. K. Ralphs and M. Guzelsoy. The SYMPHONY callable library for mixed integer programming. In *The Proceedings of the Ninth Conference of the INFORMS Computing Society*, pages 61–71, 2005.
- [36] M. P. W. Savelsbergh. Preprocessing and probing techniques for mixed integer programming problems. *ORSA Journal on Computing*, 6:445–454, 1994.
- [37] L. Schrage and L. A. Wolsey. Sensitivity analysis for branch and bound linear programming. *Operations Research*, 33:1008–1023, 1985.
- [38] T. J. Van Roy and L. A. Wolsey. Solving mixed integer programming problems using automatic reformulation. *Operations Research*, 35:45–57, 1987.

Author Addresses:

Jeffrey T. Linderoth
Department of Industrial and Systems Engineering
University of Wisconsin-Madison
1513 University Avenue
3226 Mechanical Engineering Bldg.
Madison, WI 53706-1572, USA.
E-mail: linderoth@wisc.edu

Andrea Lodi
DEIS, University of Bologna
Viale Risorgimento, 2
40136 Bologna, Italy.
E-mail: andrea.lodi@unibo.it

Abstract: In this article, we give a brief overview of state-of-the-art software for the solution of Mixed Integer Linear Programs (MILP). We begin with a brief description of important algorithmic features and conclude with concise individual descriptions of software packages.

Key Words: Mixed Integer Linear Programming — software — branch-and-cut.