

Global Optimization of Mixed-Integer Nonlinear Programs with SCIP 8.0

Ksenia Bestuzheva*, Antonia Chmiela†, Benjamin Müller‡,
Felipe Serrano§, Stefan Vigerske¶, Fabian Wegscheider||

December 16, 2022

Abstract

For over ten years, the constraint integer programming framework SCIP has been extended by capabilities for the solution of convex and nonconvex mixed-integer nonlinear programs (MINLPs). With the recently published version 8.0, these capabilities have been largely reworked and extended. This paper discusses the motivations for recent changes and provides an overview of features that are particular to MINLP solving in SCIP. Further, difficulties in benchmarking global MINLP solvers are discussed and a comparison with several state-of-the-art global MINLP solvers is provided.

1 Introduction

Mixed-integer nonlinear programming (MINLP) concerns with the optimization of an objective function such that a finite set of linear or nonlinear constraints and integrality conditions is satisfied. The generality of this problem class means that many real-world applications can be modeled as MINLPs [28, 37, 58, 69], but also that software that can handle this class efficiently becomes extremely complex. Solvers for MINLP [19] are often built on top of or by combining solvers for mixed-integer linear programming (MIP) and solvers that find locally optimal solutions for nonlinear programs (NLP). In fact, one of the first commercial MINLP solvers, SCICONIC [7], extends a MIP solver by piecewise linear approximations of low dimensional nonlinear terms. The first general purpose solver was DICOPT [41], which decomposes the solution of an MINLP into a sequence of MIP and NLP solves [25], thereby building on established software for these two program classes. DICOPT can solve MINLPs where nonlinear constraints are convex to optimality, but works only as a heuristic on nonconvex MINLPs. The first general purpose solvers to solve also nonconvex MINLPs to optimality were α BB, BARON, and GLOP [4, 60, 65], all based on convexification techniques for nonconvex constraints. Also the solver SCIP (Solving Constraint Integer Programs), for which this paper provides an overview, belongs to the latter category.

*Zuse Institute Berlin, Department AIS²T, bestuzheva@zib.de, ORCID: 0000-0002-7018-7099

†Zuse Institute Berlin, Department AIS²T, chmiela@zib.de, ORCID: 0000-0002-4809-2958

‡Zuse Institute Berlin, Department AIS²T, benjamin.mueller@zib.de, ORCID: 0000-0002-4463-2873

§Zuse Institute Berlin, Department AIS²T, serrano@zib.de, ORCID: 0000-0002-7892-3951

¶GAMS Software GmbH, c/o Zuse Institute Berlin, Department AIS²T, svigerske@gams.com

||Zuse Institute Berlin, Department AIS²T

In the following, MINLPs of the form

$$\begin{aligned}
& \min && c^\top x, \\
& \text{such that} && \underline{g} \leq g(x) \leq \bar{g}, \\
& && \underline{b} \leq Ax \leq \bar{b}, \\
& && \underline{x} \leq x \leq \bar{x}, \\
& && x_{\mathcal{I}} \in \mathbb{Z}^{|\mathcal{I}|},
\end{aligned} \tag{MINLP}$$

are considered, where $\underline{x}, \bar{x} \in \bar{\mathbb{R}}^n$, $\bar{\mathbb{R}} := \mathbb{R} \cup \{\pm\infty\}$, $\underline{x} \leq \bar{x}$, $\mathcal{I} \subseteq \{1, \dots, n\}$, $c \in \mathbb{R}^n$, $\underline{g}, \bar{g} \in \bar{\mathbb{R}}^m$, $\underline{g} \leq \bar{g}$, $g: \mathbb{R}^n \rightarrow \bar{\mathbb{R}}^m$ is specified explicitly in algebraic form, $\underline{b}, \bar{b} \in \bar{\mathbb{R}}^{\tilde{m}}$, $\underline{b} \leq \bar{b}$, and $A \in \mathbb{R}^{\tilde{m} \times n}$. The restriction to a linear objective function is a technical detail of SCIP and without loss of generality.

The software SCIP has been designed as a branch-cut-and-price framework to solve different types of optimization problems, most generally *constraint integer programs* (CIPs), and most importantly MIPs and MINLPs. Roughly speaking, CIPs are finite-dimensional optimization problems with arbitrary constraints and a linear objective function that satisfy the following property: if all integer variables are fixed, the remaining subproblem must form a linear or nonlinear program. The problem class of CIP was motivated by the modeling flexibility of constraint programming and the algorithmic requirements of integrating it with efficient solution techniques available for MIP [1].

In order to solve CIPs, SCIP constructs relaxations – typically linear programs (LPs). If the relaxation solution is not feasible for the current subproblem, the plugins that handle the violated constraints need to take measures to eventually render the relaxation solution infeasible for the updated relaxation, for example by branching or separation [1]. A plethora of additional plugin types, e.g., for presolving, finding feasible solutions, or tightening variable bounds, allow accelerating the solution process. After 20 years of development of the framework itself and included plugins, SCIP includes mature solvers for MIP, MINLP, as well as several other problem classes. Since November 2022, SCIP is freely available under an open-source license.

SCIP solves problems like (MINLP) to global optimality via a spatial branch-and-bound algorithm that mixes branch-and-infer and branch-and-cut [9]. Important parts of the solution algorithm are presolving, domain propagation (that is, tightening of variable bounds), linear relaxation, and branching. A distinguishing feature of SCIP is that its capabilities to handle nonlinear constraints are not limited to MINLPs, but can be used for any CIP. For example, problems can be handled where linear and nonlinear constraints are mixed with typical constraints from constraint programming, as long as appropriate constraint handlers have been included in SCIP. Since most constraint handlers in SCIP construct a linear relaxation of their constraints, also the handling of nonlinear constraints focuses on linear relaxations. The emphasis on handling CIPs with nonlinear constraints rather than MINLP only is also a reason that the use of nonlinear relaxations or reformulations of complete MINLPs into other problem types, e.g., mixed-integer conic programs, has not been explored much so far.

The development of SCIP initially focused on solving CIPs where fixing all integer variables resulted in a linear program [1]. However, it was soon realized that this requirement was not actually enforced by the implementation. As long as constraint handlers were able to resolve infeasibilities by separation, branching, or other means, the problem could be handled by SCIP. First experiments to handle nonlinear constraints in continuous variables were conducted for bilinear mixing constraints in mine production planning [18]. The positive results of these experiments motivated the decision to include support for more general nonlinear constraints. With version 1.2 (2009), initial support for quadratic constraints (convex or nonconvex) and solving quadratically constrained

programs (QCPs) to local optimality by Ipopt [75] was added [13]. For version 2.0 (2010), a primal heuristic that solves sub-MIPs was added [10] and other large-neighborhood-search heuristics were extended to create sub-MINLPs [12]. Further, second-order cone constraints in three variables could be handled. More general nonlinear constraints, specified in algebraic form, were first supported by SCIP 2.1 (2011) [73]. Next to the specialized treatment for quadratic constraints, also handlers for signpower constraints ($x|x|^p = z$ for some $p \geq 1$) [32] and 1-convex bivariate constraints ($f(x, y) = z$ for f being convex or concave whenever x or y has been fixed) [6] were added.

With the basic handling of nonlinear constraints in place [74], the next releases were dedicated to adding features that improved performance. SCIP 3.0 (2012) brought optimization-based bound tightening (OBBT) [33] and an NLP diving heuristic. SCIP 3.2 (2015) added a reformulation of general quadratic constraints into second-order cone constraints and separation for edge-concave decompositions of quadratic constraints [30]. With SCIP 4.0 (2017), higher-dimensional second-order cone constraints were disaggregated, KKT conditions for quadratic programs were utilized, multiple starting points were tried for NLP solves, solutions of the LP relaxation were projected onto a convex NLP relaxation, and also OBBT could be performed on the NLP instead of the LP relaxation [47]. Improved convexification of bilinear constraints by use of additional linear constraints [54], a new primal heuristic that solves a sequence of NLP reformulations, and interfaces to the NLP solvers filterSQP and Worhp [27, 20, 56] were added for SCIP 5.0 (2017) [34]. The following two major releases brought a branch-and-price based solver for ring-packing [35] (SCIP 6.0, 2018) and support for convex nonlinear subproblems in Benders Decomposition (SCIP 7.0 [31], 2020).

That versions 6 and 7 added comparatively few features for MINLP was due to an ongoing complete overhaul on the way how nonlinear constraints were handled. The primary motivation for this change, which was released with SCIP 8.0 (2022) [14], was to increase the reliability of the solver and to alleviate numerical issues that arose from problem reformulations and led to SCIP returning solutions that are feasible in the reformulated problem, but infeasible in the original problem. More precisely, previous SCIP versions built an extended formulation of (MINLP) explicitly, with the consequence that the original constraints were no longer included in the presolved problem. Even though the formulations were theoretically equivalent, it was possible that ε -feasible solution for the reformulated problem were not ε -feasible in the original problem. SCIP 8 remedies this by building an implicit extended formulation as an annotation to the original problem. A second motivation for the major changes in SCIP 8 was to reduce the ambiguity of expression and nonlinear structure types by implementing different plugin types for low-level structure types that define expressions, and high-level structure types that add functionality for particular, sometimes overlapping structures. Finally, new features for improving the solver’s performance on MINLPs were introduced with SCIP 8. These include intersection, SDP (semi-definite programming), and RLT (reformulation linearization technique) cuts for quadratic expressions [21, 16], perspective strengthening [15], and symmetry detection [76].

SCIP can read MINLPs from files in the following formats: LP, MPS, NL (AMPL), OSiL, PIP, and ZIMPL. In addition, problems can be passed to SCIP via interfaces to a variety of programming languages and modeling packages, including AMPL, C, GAMS, Java, Julia, Python, and MATLAB.

The following section provides an overview of the MINLP solving capabilities of SCIP. Afterwards, the performance of SCIP is compared with that of other state-of-the-art global solvers for MINLP.

2 MINLP capabilities of SCIP

In the following, an overview of the facilities available in SCIP that are specific to the handling of MINLPs is provided. First, available nonlinear functions are listed and the integration of nonlinear constraints into the branch-and-cut solver of SCIP is discussed. Next, the concept of a *nonlinear handler* is introduced, which is a new plug-in type that has been added with SCIP 8 and facilitates the integration of extensions that handle specific nonlinear structures. The remainder of this section gives an overview of features available in SCIP that increase the efficiency of MINLP solving, e.g., cut generators to tighten the linear relaxation, presolve reductions to simplify the problem, and primal heuristics to find feasible solutions early.

To be concise, the presentation has been limited to high-level descriptions that spare technical details. Unless specified otherwise, more details are often found in [14].

2.1 Framework

2.1.1 Expressions

Algebraic expressions are well-formed combinations of constants, variables, and various algebraic operations such as addition, multiplication, and exponentiation, that are used to describe mathematical functions. They are often represented by a directed acyclic graph with nodes representing variables, constants, and operations and arcs indicating the flow of computation, see Figure 1 for an example.




Figure 1: Expression graph for algebraic expression $\log(x)^2 + 2 \log(x)y + y^2$.

Also in SCIP, expressions are stored as directed acyclic graphs, while all semantics of expression operands are defined by *expression handler* plugins. These handlers implement callbacks that are used by methods in the SCIP core to manage expressions (create, modify, copy, free, parse, print), to evaluate and compute derivatives at a point, to evaluate over intervals, to simplify, to identify common subexpressions, to check curvature and integrality, and to iterate over it. Some additional expression handler callbacks are used by the constraint handler for nonlinear constraints (Section 2.1.2) exclusively.

Expression handlers for the following operators are included in SCIP 8.0:

- **val**: scalar constant;
- **var**: a SCIP variable;
- **sum**: an affine-linear function, $y \mapsto a_0 + \sum_{j=1}^k a_j y_j$ for $y \in \mathbb{R}^k$ with constant coefficients $a \in \mathbb{R}^{k+1}$;
- **prod**: a product, $y \mapsto c \prod_{j=1}^k y_j$ for $y \in \mathbb{R}^k$ with constant factor $c \in \mathbb{R}$;

- **pow**: a power with a constant exponent, $y \mapsto y^p$ for $y \in \mathbb{R}$ and exponent $p \in \mathbb{R}$ (if $p \notin \mathbb{Z}$, then $y \geq 0$ is required);
- **signpower**: a signed power, $y \mapsto \text{sign}(y)|y|^p$ for $y \in \mathbb{R}$ and constant exponent $p \in \mathbb{R}$, $p > 1$;
- **exp**: exponentiation, $y \mapsto \exp(y)$ for $y \in \mathbb{R}$;
- **log**: natural logarithm, $y \mapsto \log(y)$ for $y \in \mathbb{R}_{>0}$;
- **entropy**: entropy, $y \mapsto \begin{cases} -y \log(y), & \text{if } y > 0, \\ 0, & \text{if } y = 0, \end{cases}$ for $y \in \mathbb{R}_{\geq 0}$;
- **sin**: sine, $y \mapsto \sin(y)$ for $y \in \mathbb{R}$;
- **cos**: cosine, $y \mapsto \cos(y)$ for $y \in \mathbb{R}$;
- **abs**: absolute value, $y \mapsto |y|$ for $y \in \mathbb{R}$.

In previous versions of SCIP, also high-level structures such as quadratic functions could be represented as expression types. To avoid ambiguity and reduce complexity, this has been replaced by a recognition of quadratic expressions that is no longer made explicit by a change in the expression type.

2.1.2 Constraint Handler for Nonlinear Constraints

All nonlinear constraints $\underline{g} \leq g(x) \leq \bar{g}$ of (MINLP) are handled by the constraint handler for nonlinear constraints in SCIP, while the linear constraints $\underline{b} \leq Ax \leq \bar{b}$ are handled by the constraint handlers for linear constraints and its specializations (e.g., knapsack, set-covering). A constraint handler is responsible for checking whether solutions satisfy constraints and, if that is not the case, to resolve infeasibility by *enforcing constraints*. This applies in particular to solutions of the LP relaxation. The nonlinear constraint handler currently enforces constraints by the following means:

DOMAINPROP by analyzing the nonlinear constraints with respect to the variable bounds at the current node of the branch-and-bound tree, infeasibility or a bound tightening may be deduced, which allow pruning the node or cutting off the given solution, respectively; this is also known as *domain propagation*;

SEPARATE a cutting plane that is violated by the given solution may be computed;

BRANCH the current node of the branch-and-bound tree is subdivided, that is, a variable x_i and a branching point $\tilde{x}_i \in [\underline{x}_i, \bar{x}_i]$ are selected and two child nodes with x_i restricted to $[\underline{x}_i, \tilde{x}_i]$ and $[\tilde{x}_i, \bar{x}_i]$, respectively, are created.

To decide whether a node can be pruned (DOMAINPROP), an overestimate of the range of $g(x)$ with respect to current variable bounds is computed by means of interval arithmetics [53]. If a constraint k is found such that $g_k([\underline{x}, \bar{x}]) \cap [\underline{g}_k, \bar{g}_k] = \emptyset$, then there exists no point in $[\underline{x}, \bar{x}]$ for which this constraint is feasible. A bound tightening may be computed by applying the same methods in reverse order. That is, interval arithmetic is used to overestimate $g^{-1}([\underline{g}, \bar{g}])$, the preimage of $g(x)$ on $[\underline{g}, \bar{g}]$, and variable bounds are tightened to $[\underline{x}, \bar{x}] \cap g^{-1}([\underline{g}, \bar{g}])$. This is also known as feasibility-based bound tightening (FBBT). In the simplest case, callbacks of expression handlers are used to propagate intervals through expressions. However, in some cases, other methods that take more structure into account or that use additional information to tighten variable bounds and constraint sides are used (see, e.g., Sections 2.3.1 and 2.3.2).

To construct a linear relaxation of the nonlinear constraints (SEPARATE option),

an extended formulation is considered:

$$\begin{aligned}
& \min c^\top x, \\
& \text{such that } h_i(x, w_{i+1}, \dots, w_{\hat{m}}) \lesseqgtr_i w_i, \quad i = 1, \dots, \hat{m}, \\
& \underline{b} \leq Ax \leq \bar{b}, \\
& \underline{x} \leq x \leq \bar{x}, \\
& \underline{w} \leq w \leq \bar{w}, \\
& x_{\mathcal{I}} \in \mathbb{Z}^{|\mathcal{I}|}.
\end{aligned} \tag{MINLP}_{\text{ext}}$$

The functions h_i are obtained from the expressions that define functions g_i by recursively annotating subexpressions with auxiliary variables $w_{i+1}, \dots, w_{\hat{m}}$ for some $\hat{m} \geq m$. Initially, slack variables w_1, \dots, w_m are introduced and assigned to the root of all expressions, i.e., $h_i := g_i$, $\underline{w}_i := \underline{g}_i$, $\bar{w}_i := \bar{g}_i$, for $i = 1, \dots, m$. Next, for each function h_i , subexpressions f may be assigned new auxiliary variables $w_{i'}$, $i' > m$, which results in extending $(\text{MINLP}_{\text{ext}})$ by additional constraints $h_{i'}(x) = w_{i'}$ with $h_{i'} := f$. Bounds $\underline{w}_{i'}$ and $\bar{w}_{i'}$ are initialized to bounds on $h_{i'}$, if available. Since auxiliary variables in a subexpression of h_i always receive an index larger than $\max(m, i)$, the result is referred to by $h_i(x, w_{i+1}, \dots, w_{\hat{m}})$ for any $i = 1, \dots, \hat{m}$. That is, to simplify notation, w_{i+1} is used instead of $w_{\max(i, m)+1}$. If a subexpression appears in several expressions, it is assigned at most one auxiliary variable and reindexing may be necessary to have h_i depend on x and $w_{i+1}, \dots, w_{\hat{m}}$ only.

For the (in)equality sense \lesseqgtr_i , a valid simplification would be to assume equality everywhere. For performance reasons, though, it can be beneficial to relax certain equalities to inequalities if that does not change the feasible space of $(\text{MINLP}_{\text{ext}})$ when projected onto x . Therefore,

$$\lesseqgtr_i := \begin{cases} =, & \text{if } \underline{g}_i > -\infty, \bar{g}_i < \infty, \\ \leq, & \text{if } \underline{g}_i = -\infty, \bar{g}_i < \infty, \\ \geq, & \text{if } \underline{g}_i > -\infty, \bar{g}_i = \infty, \end{cases} \quad \text{for } i = 1, \dots, m.$$

For $i > m$, monotonicity of expressions is taken into account to derive \lesseqgtr_i .

Whether to annotate a subexpression by an auxiliary variable depends on the structures that are recognized. In the simplest case, every subexpression that is not already a variable is annotated with an auxiliary variable. This essentially corresponds to the Smith Normal Form [65]. For every function h_i of $(\text{MINLP}_{\text{ext}})$, the callbacks of the corresponding expression handler can be used to compute linear under- and overestimators, such that a linear relaxation for $(\text{MINLP}_{\text{ext}})$ is constructed. It can, however, be beneficial to not add an auxiliary variable for every subexpression, thus allowing for more complex functions in $(\text{MINLP}_{\text{ext}})$. This will be discussed in Section 2.1.3 below.

Example Recall Figure 1 and the constraint

$$\log(x)^2 + 2\log(x)y + y^2 \leq 4.$$

By annotating the root of the expression graph with a slack variable w_1 and each other non-variable node with an auxiliary variable, the extended formulation

$$\begin{aligned}
w_2 + 2w_3 + w_4 &\leq w_1, \\
w_5^2 &\leq w_2, \\
w_5 y &\leq w_3, \\
y^2 &\leq w_4, \\
\log(x) &= w_5, \\
w_1 &\leq 4.
\end{aligned}$$

is obtained. Bounds on auxiliary variables have been omitted here. The constraints $w_5^2 = w_2$, $w_5 y = w_3$, and $y^2 = w_4$ were relaxed to inequalities because $w_2 + 2w_3 + w_4$ is monotonically increasing in each variable. However, to relax $\log(x) = w_5$ to $\log(x) \leq w_5$, both w_5^2 and $w_5 y$ would need to be monotonically increasing in w_5 . This would be the case if $x \geq 1$ and $y \geq 0$.

If a constraint $h_i(x, w_{i+1}, \dots, w_{\hat{m}}) \leq w_i$ (the \geq -case is analogous) of $(\text{MINLP}_{\text{ext}})$ is violated and h_i is nonconvex, then linear underestimators on h_i can only be as tight as the convex envelope of h_i . Therefore, it may not be possible find a hyperplane that is violated by the solution of the LP relaxation. Since the convex envelope of h_i depends on the bounds of variables appearing in h_i , these variables are candidates for branching (BRANCH). More precisely, when an expression handler computes a linear under- or overestimator for $h_i(x, w_{i+1}, \dots, w_{\hat{m}})$, it also signals for which variables it used current variable bounds. Marked original variables are then added to the list of branching candidates. For an auxiliary variable $w_{i'}$, $i' > i$, the variables in the subexpression that $h_{i'}$ represents are considered for branching instead.

The decision on whether to add a cutting plane that separates the solution of the LP relaxation or to branch is rather complex, but the idea is to branch if either no cutting plane is found or if the violation of available cutting planes in the relaxation solution is rather small when compared to the convexification gap of the under/overestimators that define the cutting planes. In the latter case, it may be beneficial to first reduce the convexification gap by branching. To select one variable from the list of branching candidates, the violation of constraints in $(\text{MINLP}_{\text{ext}})$ and historical information about the effect of branching on a given variable on the optimal value of the LP relaxation (“pseudo costs”) are taken into account. The branching point is a convex combination of the value of the variable in the LP relaxation and the mid-point of the variable’s interval.

2.1.3 Nonlinear Handlers

In the previous example, four auxiliary variables were introduced to construct the extended formulation. This is due to the expression handlers having a rather myopic view, basically, implementing techniques that can handle only their direct children. It is clear that, for this example, an extended formulation that only replaces $\log(x)$ by an auxiliary variable w_2 could be more efficient to solve. However, this requires methods to detect the quadratic (or convex) structure and to either compute linear underestimators for the quadratic (convex) expression $w_2^2 + 2w_2 y + y^2$ or to separate cutting planes for the set defined by $w_2^2 + 2w_2 y + y^2 \leq w_1$.

Such structure detection and handling methods are the task of the new *nonlinear handler* plugins that were introduced with SCIP 8. Nonlinear handlers determine the extended formulation $(\text{MINLP}_{\text{ext}})$ by deciding when to annotate subexpressions with auxiliary variables. That is, given a constraint $h_i(x) \lesseqgtr w_i$, a nonlinear handler analyses the expression that defines h_i and attempts to detect specific structures. At this point, it may also request to introduce additional auxiliary variables, thus changing $h_i(x)$ into $h_i(x, w_{i+1}, \dots, w_{\hat{m}})$. In addition, it informs the constraint handler that it will now provide separation for $h_i(x, w_{i+1}, \dots, w_{\hat{m}}) \leq w_i$, or $\geq w_i$, or both. If none of the nonlinear handlers declare that they will handle $h_i(x) \lesseqgtr w_i$, auxiliary variables are introduced for each argument of the root of the expression h_i and expression handler callbacks are used to construct cutting planes from linear under-/overestimators.

In addition to separation, nonlinear handlers can also contribute to domain propagation. This is implemented analogously to separation by setting up an additional extended formulation similarly to $(\text{MINLP}_{\text{ext}})$, with the main difference that slack and auxiliary variables are not actually created in SCIP and equalities are currently not relaxed to inequalities.

Note that the extended formulations are stored as *annotation* on the original expressions. Thus, for each task, the most suitable formulation can be used. For example, feasibility is checked on the original constraints, domain propagation and separation use the corresponding extended formulations, but branching is performed, by default, with respect to original variables only. With SCIP 7 and earlier, only one extended formulation was constructed explicitly and the connection to the original formulation was no longer available, leading to issues due to not ensuring that solutions are also (ε -)feasible for the original constraints.

In addition to the improved numeric reliability, the nonlinear handlers also allow for a higher flexibility when handling nonlinear structures. For each node in an expression, more than one nonlinear handler can be attached, each one annotating possibly different subexpressions with auxiliary variables. For example, for a nonconvex quadratic constraint $\sum_{i,j} a_{i,j} x_i x_j \leq w$, the nonlinear handler for quadratics can declare that it will provide separation (by intersection cuts, see Section 2.3.5), but that also other means of separation should be tried. However, since no other nonlinear handler declares that it will provide separation, auxiliary variables are introduced for each argument of the sum, that is, an auxiliary variable X_{ij} is assigned to each product $x_i x_j$. For the corresponding constraints $x_i x_j \leq X_{ij}$ (if $a_{i,j} \geq 0$), the well-known McCormick underestimators [49],

$$\begin{aligned} X_{ij} &\geq \underline{x}_i x_j + \underline{x}_j x_i - \underline{x}_i \underline{x}_j, \\ X_{ij} &\geq \bar{x}_i x_j + \bar{x}_j x_i - \bar{x}_i \bar{x}_j, \end{aligned} \tag{1}$$

or other means (see Section 2.3.2) will be used to construct a linear relaxation.

2.1.4 NLP Relaxation

Similar to the central LP relaxation of SCIP, an NLP relaxation is also available. In contrast to constraint handlers, the NLP relaxation uses a common data structure to store its constraints. At the moment, constraint handlers for linear constraints and the constraint handler for nonlinear constraints store a representation of their constraints in the NLP relaxation, so that in case of a MINLP, the NLP relaxation together with the integrality conditions on variables provides a unified view of the problem. For nonlinear constraints, the original (non-extended) form $\underline{g} \leq g(x) \leq \bar{g}$ is added to the NLP. To find local optimal solutions for the NLP relaxation, interfaces to the NLP solvers filterSQP, Ipopt, and Worhp [27, 75, 20] are available. First- and second-order derivatives for these solvers are computed via CppAD [8].

The NLP relaxation is mainly used by some primal heuristics (Section 2.8) and separators (Section 2.4.2) at the moment.

2.2 Presolving

When presolving nonlinear constraints, expressions are simplified and brought into a canonical form. For example, recursive sums and products are flattened and fixed or aggregated variables are replaced by constants or sums of active variables. In addition, it is ensured that if a subexpression appears several times (in the same or different constraints), always the same expression object is used. This ensures that in the extended formulation ($\text{MINLP}_{\text{ext}}$) at most one auxiliary variable is attached to such common subexpressions.

2.2.1 Variable Fixings

Similar to what has been shown by Hansen et al. [38], if a bounded variable x_j does not appear in the objective ($c_j = 0$), but in exactly one constraint $\underline{g}_k \leq g_k(x) \leq \bar{g}_k$ where $g_k(x)$ is convex in x_j for any fixing of other variables and $\bar{g}_k = +\infty$ (or concave in x_j and $\underline{g}_k = -\infty$), then there always exists an optimal solution where $x_j \in \{\underline{x}_j, \bar{x}_j\}$.

For example, if $y \in [0, 1]$ appears only in a constraint $xy + yz - y^2 \leq 5$, then y can be changed to a binary variable.

SCIP recognizes such variables for polynomial constraints (under additional assumptions [14]) and changes the variable type to binary, if $\underline{x}_j = 0$ and $\bar{x}_j = 1$, or adds a bound disjunction constraint $x_j \leq \underline{x}_j \vee x_j \geq \bar{x}_j$. As a consequence, branching on x_j leads to fixing the variable in both children.

2.2.2 Linearization of Products

The introduction emphasized that with SCIP 8, an explicit extended reformulation of nonlinear constraints is avoided. An exception that proves this “rule” is the linearization of products of binary variables in presolving. Doing so has the advantage that more of SCIP’s techniques for MIP solving can be utilized.

In the simplest case, a product $\prod_i x_i$ is replaced by a new variable z and a constraint of type “and” that models $z = \bigwedge_i x_i$ is added. The “and”-constraint handler will then separate a linearization of this product [11]. For a product of only two binary variables, the linearization is added directly.

For a quadratic function in binary variables with many terms, the number of variables introduced may be large. Thus, in this case, a linearization that requires fewer additional variables is used, even though it may lead to a weaker relaxation.

2.2.3 KKT Strengthening for QPs

A presolving method that aims to tighten the relaxation of a quadratic program by adding redundant constraints derived from Karush-Kuhn-Tucker (KKT) conditions is available. Consider a quadratic program of the form

$$\begin{aligned} \min \quad & \frac{1}{2} x^\top Q x + c^\top x, \\ \text{such that} \quad & Ax \leq b, \end{aligned} \tag{QP}$$

where $Q \in \mathbb{R}^{n \times n}$ is symmetric, $c \in \mathbb{R}^n$, $A \in \mathbb{R}^{m \times n}$, and $b \in \mathbb{R}^m$. If (QP) is bounded, then all optima of (QP) satisfy the following KKT conditions:

$$\begin{aligned} Qx + c + A^\top \mu &= 0, \\ Ax &\leq b, \\ \mu_i (Ax - b)_i &= 0, \quad i \in \{1, \dots, m\}, \\ \mu &\geq 0, \end{aligned} \tag{KKT}$$

where μ is the vector of Lagrangian multipliers of the constraints $Ax \leq b$.

In a specialized presolver, SCIP recognizes whether (MINLP) is equivalent to (QP) by checking whether a quadratic objective function has been reformulated into a constraint. If a (QP) has been found and all variables are bounded, then the equations (KKT) are added as redundant constraints to the problem, whereby the complementarity constraints are formulated via special ordered sets of type 1. The redundant constraints can help to strengthen the linear relaxation and prioritize branching decisions to satisfy the complementarity constraints, which focuses the search more on the local optima of (QP).

In addition to (QP), the implementation can also handle mixed-binary quadratic programs. For all details, see [47, 26]. When this presolver was added to SCIP 4.0, it has shown to be very beneficial for box-constrained quadratic programs. Due to the many changes and extensions in SCIP 8, in particular for the handling of quadratic constraints (Section 2.3), it needs to be reevaluated under which conditions this presolver should be enabled. Currently, it is disabled by default.

2.2.4 Symmetry Detection

Symmetries in a MINLP are automorphisms on \mathbb{R}^n that map optimal solutions to optimal solutions. Such symmetries have an adverse effect on the performance of branch-and-bound solvers, because symmetric subproblems may be treated repeatedly. Therefore, SCIP can enforce lexicographically maximal solutions from an orbit of symmetric solutions via bound tightening and separation of linear inequalities [39, 34, 31, 14].

Since optimal solutions are naturally not known in advance, the symmetry detection resorts to find permutations of variables that map the feasible set onto itself and map each point to one with the same objective function value [48]. These permutations are given by isomorphisms in an auxiliary symmetry detection graph, which is constructed from the problem data (c , A , \mathcal{I} , and the expressions that define $g(x)$) [43, 76].

2.3 Quadratics

Since quadratic functions frequently appear in MINLPs (every second instance of MINLPLib [50]) has only linear and quadratic constraints), a number of techniques have been added to SCIP to handle this structure. Next to the presolving methods that were discussed in the previous section, three nonlinear handlers and four separators deal with quadratic structures. When none of the nonlinear handlers are active, then for each square and bilinear term in a quadratic function, an auxiliary variable is added in the extended formulation and gradient, secant, and McCormick under- and overestimators (see (1)) are generated.

2.3.1 Domain Propagation

If variables appear more than once in a quadratic function, then a term-wise domain propagation does not necessarily yield the best possible results, due to suffering from the so-called *dependency problem* of interval arithmetics. For example, it is easy to compute the range for $x^2 + x$ for given bounds on x , or bounds on x for a given interval on $x^2 + x$, but standard interval arithmetics would treat the terms x^2 and x separately, which can lead to overestimating the result.

Therefore, a specialized nonlinear handler in SCIP provides a domain propagation procedure for quadratics that aims to reduce overestimation. For this, the detection routine of the nonlinear handler writes a quadratic expression as

$$q(y) = \sum_{i=1}^k q_i(y) \quad \text{with} \quad q_i(y) = a_i y_i^2 + c_i y_i + \sum_{j \in P_i} b_{i,j} y_i y_j, \quad (2)$$

where y_i is either an original variable (x) or another expression, $a_i, c_i \in \mathbb{R}$, $b_{i,j} \in \mathbb{R} \setminus \{0\}$, $j \in P_i \Rightarrow i \notin P_j$ for all $j \in P_i$, $P_i \subset \{1, \dots, k\}$, $i = 1, \dots, k$. For functions q_i with at least two terms (at least two of a_i , $b_{i,j}$, $j \in P_i$, and c_i are nonzero), a relaxation is obtained by replacing each y_j by $[\underline{y}_j, \bar{y}_j]$, $j \in P_i$. For this univariate quadratic interval-term in y_i , tight bounds can be computed [24].

In addition, bounds on variables y_j , $j \in P_i$, are computed by considering

$$\sum_{j \in P_i} b_{i,j} y_j \in ([\underline{q}, \bar{q}] - \sum_{i' \neq i} q_{i'}(y)) / y_i - a_i y_i - c_i, \quad y_i \in [\underline{y}_i, \bar{y}_i], \quad (3)$$

where $[\underline{q}, \bar{q}]$ are given bounds on $q(y)$. After relaxing each $q_{i'}$ to an interval, bounds on the right-hand side of (3) are computed, which are then used to calculate bounds on each y_j , $j \in P_i$.

2.3.2 Bilinear Terms

For a product $y_1 y_2$, where y_1 and y_2 are either non-binary variables or other expressions, the expression handler for products already provides linear under- and overestimators and domain propagation that is best possible when considering the bounds $[y_1, \bar{y}_1] \times [y_2, \bar{y}_2]$ only. However, if linear inequalities in y_1 and y_2 are available, then possibly tighter linear estimates and variable bounds can be computed. In SCIP, this is done by a specialized nonlinear handler that implements the algorithm by Locatelli [45]. The inequalities are found by projection of the LP relaxation onto variables (y_1, y_2) . For more details, see [54]. An alternative method that uses linear constraints to tighten the relaxation of quadratic constraints are the RLT cuts described in the following.

2.3.3 RLT Cuts

The Reformulation-Linearization Technique (RLT) [2, 3] has proven very useful to tighten relaxations of polynomial programming problems. In SCIP, a separator of cuts that are computed via RLT for bilinear product relations in $(\text{MINLP}_{\text{ext}})$ is available.

For simplicity, let us denote by X_{ij} the auxiliary variable that is associated with a constraint $x_i x_j \leq X_{ij}$ of $(\text{MINLP}_{\text{ext}})$ (X_{ji} denotes the same variable as X_{ij}). Recall that it is valid to replace \leq by $=$. Given $X_{ij} = x_i x_j$, where $x_i \in [\underline{x}_i, \bar{x}_i]$, $x_j \in [\underline{x}_j, \bar{x}_j]$, and a linear constraint $a^\top x \leq b$, RLT cuts are derived by first multiplying the constraint by a nonnegative bound factors $(x_i - \underline{x}_i)$, $(\bar{x}_i - x_i)$, $(x_j - \underline{x}_j)$, or $(\bar{x}_j - x_j)$. For instance, consider multiplication by the factor $(x_i - \underline{x}_i)$, which yields a valid nonlinear inequality:

$$a^\top x (x_i - \underline{x}_i) \leq b (x_i - \underline{x}_i). \quad (4)$$

This is referred to as the reformulation step.

The linearization step is then performed for all terms $x_k x_i$ in (4). If a product relation $X_{ki} = x_k x_i$ exists, then the product is replaced with X_{ki} . If x_k and x_i are contained in the same clique, the product is replaced with an equivalent linear expression. Otherwise, it is replaced by a linear under- or overestimator such as (1).

In addition, the RLT separator can reveal linearized products between binary and continuous variables. To do so, it checks whether pairs of linear inequalities that are defined in the same triple of variables (one of them binary, the other two continuous) imply a product relation. These implicit products can then be used in the linearization step of RLT cut generation [16].

2.3.4 SDP Cuts

As in the previous section, denote by X_{ij} the auxiliary variable that is associated with a constraint $x_i x_j \leq X_{ij}$ of $(\text{MINLP}_{\text{ext}})$. A popular convex relaxation of the condition $X = xx^\top$ is given by requiring $X - xx^\top$ to be positive semidefinite. Separation for the set $\{(x, X) : X - xx^\top \succeq 0\}$ itself is possible, but cuts are typically dense and may include variables X_{ij} for products that do not exist in the problem. Therefore, only principal 2×2 minors of $X - xx^\top$, which also need to be positive semidefinite, are considered. By Schur's complement, this means that the condition

$$A_{ij}(x, X) := \begin{bmatrix} 1 & x_i & x_j \\ x_i & X_{ii} & X_{ij} \\ x_j & X_{ij} & X_{jj} \end{bmatrix} \succeq 0 \quad (5)$$

needs to hold for any i, j , $i \neq j$. A separator in SCIP detects minors for which X_{ii} , X_{jj} , X_{ij} exist in $(\text{MINLP}_{\text{ext}})$ and enforces $A_{ij}(x, X) \succeq 0$. To do so for a solution (\hat{x}, \hat{X}) that violates (5), an eigenvector $v \in \mathbb{R}^3$ of $A_{ij}(\hat{x}, \hat{X})$ with $v^\top A_{ij}(\hat{x}, \hat{X}) v < 0$ is computed and the globally valid linear inequality $v^\top A_{ij}(x, X) v \geq 0$ is added.

2.3.5 Intersection Cuts

Intersection cuts [70, 5] have shown to be an efficient tool to strengthen relaxations of MIPs. Recently, Muñoz and Serrano showed how to compute the tightest possible intersection cuts for quadratic programs [55]. This method has been implemented in SCIP [21].

Assume a nonconvex quadratic constraint of $(\text{MINLP}_{\text{ext}})$ is $q(y) \leq w$ with $q(y)$ as in (2) and w an auxiliary variable. The separation of intersection cuts is implemented for the set $S := \{(y, w) \in \mathbb{R}^k : q(y) \leq w\}$ that is defined by this constraint.

Let (\hat{y}, \hat{w}) be a basic feasible LP solution violating $q(y) \leq w$. First, a convex inequality $g(y, w) < 0$ is build that is satisfied by (\hat{y}, \hat{w}) , but by no point of S . This defines a so-called *S-free set* $C = \{(y, w) \in \mathbb{R}^{k+1} : g(y, w) \leq 0\}$, that is, a convex set with $(\hat{y}, \hat{w}) \in \text{int}(C)$ containing no point of S in its interior. The quality of the resulting cut highly depends on which *S-free set* is used, but using *maximal S-free sets* yield the tightest possible intersection cuts [55].

By using the conic relaxation K of the LP-feasible region defined by the nonbasic variables at (\hat{y}, \hat{w}) , the intersection points between the extreme rays of K and the boundary of C are computed. The intersection cut is then defined by the hyperplane going through these points and successfully separates (\hat{y}, \hat{w}) and S . See Figure 2 for an illustration. To obtain even better cuts, there is also a strengthening procedure implemented that uses the idea of negative edge extension of the cone K [36].




Figure 2: An intersection cut (red) separating the basic feasible LP solution (\hat{y}, \hat{w}) from S (blue). The cut is computed using the intersection points of an *S-free set* C (orange) and the rays of a simplicial cone $K \supseteq S$ (boundary in green) with apex $(\hat{y}, \hat{w}) \notin S$.

In addition to the separation of intersection cuts for a set S given by a constraint $q(y) \leq w$, SCIP can also generate intersection cuts for implied quadratic equations. Recall the matrix of auxiliary variables X as introduced in Section 2.3.3. The condition $X = xx^\top$ implies that X needs to have rank 1. Therefore, any 2×2 minor $\begin{pmatrix} X_{i_1 j_1} & X_{i_1 j_2} \\ X_{i_2 j_1} & X_{i_2 j_2} \end{pmatrix}$ of X needs to have determinant zero. That is, for any set of variable indices i_1, i_2, j_1, j_2 with $i_1 \neq i_2$ and $j_1 \neq j_2$, the condition $X_{i_1 j_1} X_{i_2 j_2} = X_{i_1 j_2} X_{i_2 j_1}$ needs to hold. If all variables in this condition exist in $(\text{MINLP}_{\text{ext}})$ and a solution violates this condition, then the previously described procedure to generate intersection cuts is applied to the set defined by this condition.

Since intersection cuts can be rather dense, it is not clear yet how to decide when it will be beneficial to generate such cuts. Their separation is therefore currently disabled by default. For more details, see [21].

2.3.6 Edge-Concave Cuts

Another method to obtain a linear outer-approximation for a quadratic constraint is by utilizing an edge-concave decomposition of the quadratic function. This has shown to be particularly useful for randomly generated quadratic instances [51, 52]. A function is edge-concave over the variables' domain (e.g., $[\underline{x}, \bar{x}]$) if it is componentwise concave.

Given a quadratic function, the separator for edge-concave cuts solves an auxiliary MIP to partition the square and bilinear terms into a sum of edge-concave functions and a remaining function. Since the convex envelope of edge-concave functions is *vertex-polyhedral* [67], that is, it is a polyhedral function with vertices corresponding to the vertices of the box of variable bounds, facets on the convex envelope of each edge-concave function can be computed by solving an auxiliary linear program (see also Section 2.4.1). For the function of remaining terms, term-wise linear underestimators such as (1) are summed up.

Since the current implementation of edge-concave cuts in SCIP has not shown to be particularly useful for general MINLP, this separator is disabled for now.

2.3.7 Second-Order Cones

An important connection between MINLP and conic programming is the detection of constraints that can be represented as a second-order cone (SOC) constraint, since the latter defines a convex set, while the original constraint may use a nonconvex constraint function.

A specialized nonlinear handler aims to detect SOC representable structures. In the detection phase, a constraint $h_i(x) \leq w_i$ (the case \geq is handled similarly) of the extended formulation ($\text{MINLP}_{\text{ext}}$) is passed to the nonlinear handler. For this constraint, it is checked whether it defines a bound on an Euclidian norm ($\sqrt{\sum_{j=1}^k (a_j y_j^2 + b_j y_j)} + c \leq w_i$ for some coefficients $a_j, b_j, c \in \mathbb{R}$, $a_j > 0$, where y_j is either an original variable or some subexpression of $h_i(\cdot)$), or is a quadratic constraint that is SOC-representable [46]. Since the introduction of slack variables w_i , $i \leq m$, may prevent such a detection, the equivalent constraint $h_i(x) \leq \bar{w}_i$ is considered instead.

A detected SOC constraint is stored in the form

$$\sqrt{\sum_{j=1}^k (v_j^\top y + \beta_j)^2} \leq v_{k+1}^\top y + \beta_{k+1} \quad (6)$$

with $v_j \in \mathbb{R}^\ell$, $j = 1, \dots, k+1$, where y_1, \dots, y_ℓ are variables of ($\text{MINLP}_{\text{ext}}$). Since the left-hand side of (6) is convex, a solution \hat{y} that violates (6) can be separated by linearization of the left-hand side of (6).

However, if there are many terms on the left-hand side of (6) (k being large), then it can require many cuts to provide a tight linear relaxation of (6). Thus, a disaggregation of the cone [72] is used if $k \geq 3$:

$$(v_j^\top y + \beta_j)^2 \leq z_j (v_{k+1}^\top y + \beta_{k+1}), \quad j = 1, \dots, k, \quad (7)$$

$$\sum_{j=1}^k z_j \leq v_{k+1}^\top y + \beta_{k+1}, \quad (8)$$

where variables z_1, \dots, z_k are new variables. A solution (\hat{y}, \hat{z}) that violates (6) needs to violate also (7) for some $j \in \{1, \dots, k\}$ or (8). The latter is already linear and can be added as a cut. If a rotated second-order cone constraint (7) is violated for some j , then it is transformed into the standard form

$$\sqrt{4(v_j^\top y + \beta_j)^2 + (v_{k+1}^\top y + \beta_{k+1} - z_j)^2} \leq v_{k+1}^\top y + \beta_{k+1} + z_j$$

and a gradient cut is constructed by linearization of the left-hand side.

2.4 Convexity

2.4.1 Convex and Concave Constraints

For the linear underestimation of functions like $x \exp(x)$ or $x^2 + 2xy + y^2$, the construction of an extended formulation ($xw, \exp(x) = w; w_1 + 2w_2 + w_3, w_1 = x^2, w_2 = xy, w_3 = y^2$) is not advisable. Instead, hyperplanes that support the epigraph of a convex function can be used if convexity is recognized. In SCIP, specialized nonlinear handlers are available to detect for a function $h_i(x)$ of $(\text{MINLP}_{\text{ext}})$ the subexpressions that need to be replaced by auxiliary variables $w_{i+1}, \dots, w_{\hat{m}}$ such that the remaining expression $h_i(x, w_{i+1}, \dots, w_{\hat{m}})$ is convex or concave. The detection utilizes the often applied rules for convexity/concavity of function compositions (e.g., f convex and monotone decreasing, g concave $\Rightarrow f \circ g$ convex), but applies them in reverse order. That is, instead of deciding whether a function is convex/concave based on information on the convexity/concavity and monotonicity of its arguments, the algorithm formulates conditions on the convexity/concavity of the function arguments given a convexity/concavity requirement on the function itself. When a condition on an argument cannot be fulfilled, it is replaced by an auxiliary variable.

Next to “myopic” rules for convexity/concavity that are implemented by the expression handlers, also rules for product compositions ($af(bg(x) + c)g(x)$ with constants a, b, c and repeating subexpression $g(x)$), signomials ($c \prod_{j=1}^k f_j^{p_j}(x)$ with $c, p_j \in \mathbb{R}$ and subexpressions $f_j(x)$, $j = 1, \dots, k$), and quadratic forms are available. The latter may check for definiteness of its Hessian by calculating its eigenvalues. Further, it has been shown that for a composition of convex functions $f \circ g$, it can be beneficial for the linear relaxation to consider the extended formulation $f(w)$, $w \geq g(x)$, instead of the composition $f(g(x))$ [68]. This is enforced by a small variation of the detection algorithm.

When a convex constraint $h_i(x, w_{i+1}, \dots, w_{\hat{m}}) \leq w_i$ of $(\text{MINLP}_{\text{ext}})$ is violated at a point (\hat{x}, \hat{w}) , a tangent on the graph of h_i at (\hat{x}, \hat{w}) is used to compute a separating hyperplane. The slope of the tangent is given by the gradient of h_i at (\hat{x}, \hat{w}) , which is calculated via automatic differentiation on the expression graph. If, however, h_i is univariate, that is, $h_i(x, w_{i+1}, \dots, w_{\hat{m}}) = f(y)$ for some variable y , and y is integral, then taking the hyperplane through the points $(\lfloor \hat{y} \rfloor, f(\lfloor \hat{y} \rfloor))$ and $(\lfloor \hat{y} \rfloor + 1, f(\lfloor \hat{y} \rfloor + 1))$ can give a tighter underestimator.

For a concave function $h_i(x, w_{i+1}, \dots, w_{\hat{m}})$, any hyperplane $\alpha x + \beta w + \gamma$ that underestimates $h_i(x, w_{i+1}, \dots, w_{\hat{m}})$ in all vertices of the box $[\underline{x}, \bar{x}] \times [\underline{w}_{i+1}, \bar{w}_{i+1}] \times \dots \times [\underline{w}_{\hat{m}}, \bar{w}_{\hat{m}}]$ is a valid linear underestimator, since h_i is vertex-polyhedral with respect to the box. Maximizing $\alpha \hat{x} + \beta \hat{w} + \gamma$ such that $\alpha x + \beta w + \gamma$ does not exceed $h_i(x, w_{i+1}, \dots, w_{\hat{m}})$ for all vertices gives an underestimator that is as tight as possible at a given reference point (\hat{x}, \hat{w}) . For the frequent cases $k = 1$ and $k = 2$, routines that directly compute such an underestimator are available. For $k > 2$, a linear program is solved. Since the size of this LP is exponential in k , underestimators for concave functions in more than 14 variables are currently not computed.

2.4.2 Tighter Gradient Cuts

The separating hyperplanes generated for convex functions of $(\text{MINLP}_{\text{ext}})$ as discussed in the previous section are, in general, not supporting for the feasible region of $(\text{MINLP}_{\text{ext}})$, because the point where the functions are linearized is not at the boundary of the feasible region (which is the reason why it needs to be separated). Therefore, often several rounds of cut generation and LP solving are required until the relaxation solution satisfies the convex constraints. Solvers for convex MINLP have handled this problem in various ways [25, 42], but the basic idea is to build gradient cuts at a suitable boundary point of the feasible region.

In SCIP, three procedures for building tighter and/or deeper gradient cuts for convex relaxations are included. The first two methods compute a point on the boundary of the set defined by all convex constraints of (MINLP) that is close to the point to be separated. The first method solves an additional nonlinear program to project the point to be separated onto the convex set. Since solving an NLP for every point to be separated can be quite expensive, the second method, going back to an idea by Veinott [71], does a binary search between an interior point of the convex set and the point to be separated. The interior point is computed once in the beginning of the search by solving an auxiliary NLP. For more details, see [47].

The third method does not aim to separate a given point, but utilizes the feasible points that are found by primal heuristics of SCIP. When a new solution is found, gradient cuts are generated at this solution for convex constraints of (MINLP_{ext}) and added to the cutpool. If such a cut is later found to separate the relaxation solution, it is added to the LP.

All methods are currently disabled as they require more tuning to be efficient in general.

2.5 Quotients

Note that the available expression handlers (see Section 2.1.1) do not include a handler for quotients, since they can equivalently be written using a product and a power expression. Therefore, the default extended formulation for an expression $y_1 y_2^{-1}$ is given by replacing y_2^{-1} by a new auxiliary variable w . The linear outer-approximation is then obtained by estimating $y_1 w$ and y_2^{-1} separately. However, tighter linear estimates are often possible. Therefore, a specialized nonlinear handler checks whether a given function $h_i(x)$ can be cast as

$$f(y) = \frac{ay_1 + b}{cy_2 + d} + e \quad (9)$$

with $a, b, c, d, e \in \mathbb{R}$, $a, c \neq 0$, and y_1 and y_2 being either original variables or subexpressions of $h_i(x)$.

Tight linear estimators for (9) are computed by distinguishing a number of cases. For example, for $ay_1 + b \geq 0$ and $cy_2 + d > 0$ (if $c > 0$), a linear underestimator is obtained by computing a tangent on the graph of the convex underestimator of f that is given by [78]. A linear overestimator is obtained by computing a facet on the concave envelope of f , which is easy since $-f$ is vertex-polyhedral. Furthermore, in the univariate case ($y_1 = y_2$), f is either convex or concave on $[y_1, \bar{y}_1]$ if $-d/c \notin [y_2, \bar{y}_2]$.

Since in the univariate case the same variable appears twice, also a specialized domain propagation method that avoids the dependency problem of interval arithmetic is available.

2.6 Perspective Strengthening

Perspective reformulations have shown to efficiently tighten relaxations of convex mixed-integer nonlinear programs with on/off-structures, which are often modeled via big-M constraints or semi-continuous variables [29]. A variable x_j is semi-continuous with respect to the binary indicator variable $x_{j'}$, $j' \in \mathcal{I}$, if it is restricted to the domain $[x_j^1, \bar{x}_j^1]$ when $x_{j'} = 1$ and has a fixed value x_j^0 when $x_{j'} = 0$.

In SCIP, a strengthening of under- and overestimators for functions that depend on semi-continuous variables is available. Consider a constraint $h_i(x, w_{i+1}, \dots, w_{\hat{m}}) \lesseqgtr w_i$ of (MINLP_{ext}) and write h_i as a sum of its nonlinear and linear parts:

$$h_i(x, w_{i+1}, \dots, w_{\hat{m}}) = h_i^{\text{nl}}(x_{\text{nl}}, w_{\text{nl}}) + h_i^1(x_1, w_1),$$

where h_i^{nl} is a nonlinear function, h_i^{l} is a linear function, x_{nl} and w_{nl} are the vectors of variables x and w , respectively, that appear only in the nonlinear part of h_i , and x_{l} and w_{l} are the vectors of variables x and w , respectively, that appear only in the linear part of h_i . A strengthening of under- or overestimators for $h_i(x, w_{i+1}, \dots, w_{\hat{m}})$ is attempted if x_{nl} and w_{nl} are semi-continuous with respect to the same indicator variable $x_{j'}$.

To determine whether a variable x_j is semi-continuous, bounds on x_j that are implied by fixing a binary variable are analyzed. The implied bounds can be obtained either from linear constraints directly or by probing, and are stored by SCIP in a globally available data structure. If a pair of implied bounds on x_j with the same binary variable $x_{j'}$ is found, i.e.,

$$\begin{aligned} x_j &\leq \alpha^{(u)} x_{j'} + \beta^{(u)}, \\ x_j &\geq \alpha^{(\ell)} x_{j'} + \beta^{(\ell)}, \end{aligned}$$

and $\beta^{(u)} = \beta^{(\ell)}$, then x_j is a semi-continuous variable with $x_j^0 = \beta^{(u)}$, $\underline{x}_j^1 = \alpha^{(\ell)} + \beta^{(\ell)}$, and $\bar{x}_j^1 = \alpha^{(u)} + \beta^{(u)}$. In addition, an auxiliary variable w_i is found to be semi-continuous if function $h_i(x, w_{i+1}, \dots, w_{\hat{m}})$ depends only on semi-continuous variables with the same indicator variable.

Assume that a linear underestimator $\ell(x, w_{i+1}, \dots, w_{\hat{m}})$ of $h_i(x, w_{i+1}, \dots, w_{\hat{m}})$ has been computed and split it into parts corresponding to the nonlinear and linear variables of h_i , respectively:

$$\ell(x, w_{i+1}, \dots, w_{\hat{m}}) = \ell^{\text{nl}}(x_{\text{nl}}, w_{\text{nl}}) + \ell^{\text{l}}(x_{\text{l}}, w_{\text{l}}).$$

The perspective strengthening consists of extending the part of the underestimator that corresponds to the nonlinear part such that it is tight for $x_{j'} = 0$:

$$\ell^{\text{nl}}(x_{\text{nl}}, w_{\text{nl}}) + (h_i^{\text{nl}}(x_{\text{nl}}^0, w_{\text{nl}}^0) - \ell^{\text{nl}}(x_{\text{nl}}^0, w_{\text{nl}}^0)) (1 - x_{j'}) + \ell^{\text{l}}(x_{\text{l}}, w_{\text{l}}).$$

The linear part remains unchanged, since it shares none of the variables with the nonlinear part. This extension ensures that the estimator is equal to $h_i(x, w_{i+1}, \dots, w_{\hat{m}})$ for $x_{j'} = 0$, $x_{\text{nl}} = x_{\text{nl}}^0$, and $w_{\text{nl}} = w_{\text{nl}}^0$, and equal to $\ell(x, w_{i+1}, \dots, w_{\hat{m}})$ for $x_{j'} = 1$. If h_i is convex, cuts obtained this way are equivalent to the classic perspective cuts [29]. For more details on the implementation in SCIP, see [15]. An example is given in Figure 3.




Figure 3: The original set $\{(x, y, w) : x^2 \leq w, y \in \{0, 1\}, y = 0 \rightarrow x = 0\}$ (left) and a continuous relaxation given by $\{(x, y, w) : x^2 \leq wy, y \in [0, 1], w \geq 0\}$ (right). From the original set, cuts of the form $\hat{x}^2 + 2\hat{x}(x - \hat{x}) \leq w$ for some reference point \hat{x} would be generated. With perspective strengthening, a linearization on the right set is obtained instead, i.e., $\hat{x}^2 + 2\hat{x}(x - \hat{x}) + \hat{x}^2(1 - y) \leq w$. The latter is typically better as it is tight for $y = 0$ as well.

2.7 Optimization-Based Bound Tightening

Optimization-Based Bound Tightening (OBBT) is a domain propagation technique which minimizes and maximizes each variable over the feasible set of the problem or a relaxation thereof [59]. Whereas FBBT (see Section 2.1.2) propagates the nonlinearities individually, OBBT considers (a relaxation of) all constraints together, and may hence compute tighter bounds. However, it is rather expensive compared to FBBT.

In SCIP, OBBT solves two auxiliary LPs for each variable x_k that could be subject to spatial branching:

$$\min / \max \{x_k : D^x x + D^w w \leq d, c^\top x \leq U, x \in [\underline{x}, \bar{x}], w \in [\underline{w}, \bar{w}]\} \quad (10)$$

where $D^x x + D^w w \leq d$, $D^x \in \mathbb{R}^{\ell \times n}$, $D^w \in \mathbb{R}^{\ell \times \hat{m}}$, $d \in \mathbb{R}^\ell$ is the linear relaxation of the feasible region of $(\text{MINLP}_{\text{ext}})$, and $c^\top x \leq U$ is an objective cutoff constraint that excludes solutions with objective value worse than the current incumbent. The optimal value of (10) may then be used to tighten the lower / upper bound of variable x_k . A variable is subject to spatial branching if cut separation routines use the bounds of the variable at a node of the branch-and-bound tree.

SCIP, by default, applies OBBT at the root node to tighten bounds globally. It restricts the computational effort by limiting the amount of LP iterations spent for solving the auxiliary LPs and interrupting for cheaper domain propagation techniques to be called between LP solves.

Further, SCIP does not only use the optimal objective values of (10) to tighten the bounds on x_k , but it also applies a computationally cheap approximation of OBBT during the branch-and-bound search by exploiting the dual solutions from solves of (10) at the root node. Suppose the maximization LP is solved and feasible dual multipliers $\lambda_1, \dots, \lambda_\ell, \mu \geq 0$ for $D^x x + D^w w \leq d$, $c^\top x \leq U$, respectively, and the corresponding reduced cost vectors r^x and r^w are obtained. Then

$$x_k \leq \sum_j r_j^x x_j + \sum_j r_j^w w_j + \lambda^\top d + \mu U \quad (11)$$

is a valid inequality, which is called *Lagrangian variable bound* (LVB), and

$$\sum_{j:r_j^x < 0} r_j^x \underline{x}_j + \sum_{j:r_j^x > 0} r_j^x \bar{x}_j + \sum_{j:r_j^w < 0} r_j^w \underline{w}_j + \sum_{j:r_j^w > 0} r_j^w \bar{w}_j + \lambda^\top d + \mu U \quad (12)$$

is a valid upper bound for x_k that equals the OBBT bound if the dual multipliers are optimal. SCIP learns LVBs at the root node and propagates them during the tree search whenever the bounds of variables on the right-hand side of (11) become tighter or an improved primal solution is found. For further details, see [33].

In addition to OBBT with respect to the LP relaxation, also a variant is available that optimizes single variables over the potentially tighter convex NLP relaxation that is given by all linear and convex nonlinear constraints of (MINLP) . Also for this variant, linear Lagrangian variable bounds similar to (11) can be constructed by taking constraint convexity and KKT conditions into account. Because of the potentially high computational cost of solving many NLPs, this variant of OBBT is deactivated by default. For more details, see [47].

2.8 Primal Heuristics

The purpose of primal heuristics is to find high quality feasible solutions early in the search. When given an MINLP, up to 40 primal heuristics are active in SCIP by default. Many of them aim to find an integer-feasible solution to the LP relaxation. In the following, primal heuristics that are only active in the presence of nonlinear constraints are discussed.

2.8.1 subNLP

A primal heuristic like **subNLP** is implemented in virtually any global MINLP solver. Given a point \tilde{x} that satisfies the integrality requirements ($\tilde{x}_{\mathcal{I}} \in \mathbb{Z}^{|\mathcal{I}|}$), the heuristic starts by fixing all integer variables in (MINLP) to the values given by \tilde{x} . It then calls the SCIP presolver on this subproblem for possible simplifications. Finally, it triggers a solution of the remaining NLP, using \tilde{x} as the starting point. If the NLP solver, such as Ipopt, finds a solution that is feasible (and often also locally optimal) for the NLP relaxation, then a feasible point for (MINLP) has been found.

The starting point \tilde{x} can be the current solution of the LP relaxation if integer-feasible, a point found by a primal heuristic that searches for integer-feasible solutions of the LP relaxation, or a point that is passed on by other primal heuristics for MINLP, such as those mentioned in the next sections.

How frequently the heuristic should run and how much effort to spend on an NLP solve is a nontrivial decision. In the current implementation, the heuristic uses a fixed number for the iteration limit of the NLP solver for its first run. For the following calls, the limit is set to twice the average number of iterations required in previous runs. If, however, many of the previous runs hit the iteration limit, then an increased iteration limit is used. Whether to run the heuristic at a node of the branch-and-bound tree depends on the number of nodes processed since it ran the last time, the iteration limit that would be used, and how successful the heuristic has been in finding feasible points in previous calls.

2.8.2 Multistart

If (MINLP) is nonconvex after fixing all integer variables, then several local optima may exist for the NLPs solved by heuristic **subNLP**. The success of the NLP solver then strongly depends on the starting point. Therefore, the multistart heuristic aims to compute several starting points that are passed to the **subNLP** heuristic.

The algorithm, originally developed in [66], tries to approximate the boundary of the feasible set of the NLP relaxation by sampling points from $[\underline{x}, \bar{x}]$ and pushing them towards the feasible set by the use of an inexpensive gradient descent method. Afterwards, points that are relatively close to each other are grouped into clusters. Ideally, each cluster approximates the boundary of some connected component of the feasible set. For each cluster, a linear combination of the points is passed as a starting point to **subNLP**. For integer variables x_i , $i \in \mathcal{I}$, the value in the starting point is rounded to an integral value.

To reduce infeasibility of a point \hat{x} , the *constraint consensus* method [66] is used. The algorithm computes a descent direction for each violated constraint of (MINLP). For example, if $g_i(\hat{x}) > \bar{g}_i$ for some $i \in \{1, \dots, m\}$, then the descent direction is given by $-\frac{g_i(\hat{x})}{\|\nabla g_i(\hat{x})\|^2} \nabla g_i(\hat{x})$. Point \hat{x} is then updated by adding the average of the descent directions for all violated linear and nonlinear constraints. This step is iterated until \hat{x} becomes feasible, or a stopping criterion has been fulfilled.

The multistart heuristic currently runs for continuous problems ($\mathcal{I} = \emptyset$) only by default, since rounding and fixing integer variables most likely lead to infeasible NLP subproblems. For more details, see [47].

2.8.3 NLP Diving

As an alternative to finding a good fixing for all integer variables of (MINLP), the NLP diving heuristic starts by solving the NLP relaxation at the current branch-and-bound node with an NLP solver. It then iteratively fixes integer variables with fractional value and resolves both the LP and NLP relaxations, thereby simulating a depth-first-search in a branch-and-bound tree. By default, variables for which the sum of the distances

from the solutions of the LP and NLP relaxations to a common integer value is minimal are rounded to the nearest integer value. Further, binary variables and nonlinear variables are preferred. If the resulting NLP is found to be (locally) infeasible, one-level backtracking is applied, that is, the last fixing is undone, and the opposite fixing is tried. If this is infeasible, too, the heuristic aborts.

2.8.4 MPEC

While the NLP diving heuristic either completely omits or enforces integrality restrictions in the NLP relaxation, the MPEC heuristic adds a relaxation of the integrality restriction to the NLP and tightens this relaxation iteratively. The heuristic is only applicable to mixed-binary nonlinear programs at the moment.

The basic idea of the heuristic, originally developed in [61], is to reformulate (MINLP) as a mathematical program with equilibrium constraints (MPEC) and to solve this MPEC to local optimality. The MPEC is obtained from (MINLP) by rewriting the condition $x_i \in \{0, 1\}$, $i \in \mathcal{I}$, as complementarity constraint $x_i \perp 1 - x_i$. This reformulation is again reformulated to an NLP by writing it as $x_i(1 - x_i) = 0$. However, since these reformulated complementarity constraints will not, in general, satisfy constraint qualifications, solving this NLP reformulation with a generic NLP solver will often fail.

Therefore, in order to increase the chances of solving the NLP reformulation, the heuristic solves regularized versions of the NLP by relaxing $x_i(1 - x_i) = 0$ to $x_i(1 - x_i) \leq \theta$, for different, ever smaller $\theta > 0$. The solution of one NLP is thereby used as the starting point for the next solve. If the NLP solution is close to satisfying $x_{\mathcal{I}} \in \{0, 1\}^{|\mathcal{I}|}$, it is passed as starting point to the `subNLP` heuristic. If an NLP is (locally) infeasible, the heuristic does two more attempts where the values for binary variables that are already close to 0 or 1 are flipped to 1 or 0, respectively. For more details, see [34].

2.8.5 Undercover

While the previous heuristics focused mainly on enforcing the integrality condition on an NLP, heuristic **undercover** [10] starts from a completely different angle. The heuristic is based on the observation that it sometimes suffices to fix only a comparatively small number of variables of (MINLP) to yield a subproblem with all constraints being linear. For example, for a bilinear term, only one of the variables needs to be fixed. The variables to fix are chosen by solving a set covering problem, which aims at minimizing the number of variables to fix. The values for the fixed variables are taken from the solution of the LP or NLP relaxation or a known feasible solution of the MINLP.

The resulting sub-MIP is less complex to solve, and does not need to be solved to proven optimality. The solutions of the sub-MIP are immediately feasible for (MINLP). However, the best one is also passed as starting point to heuristic `subnlp` to try for further improvement. For more details, see [10].

3 Benchmark

The following aims to present a fair comparison of SCIP with several other state-of-the-art solvers for general MINLP. Doing so is not trivial at all. First, a set of instances needs to be selected that is suitable as a benchmark set. Second, solver parameters have to be set such that all solvers solve the same instances with the same working limits and the same requirements on feasibility and optimality – this goal could not be reached completely. Third, the solver’s results have to be checked for correctness, or, when this is not possible, plausibility.

GAMS was used for the experiments, as it provides various facilities to help on solver comparisons and comes with current versions of SCIP and the commercial solvers BARON [40], Lindo API [44], and Octeract included. ANTIGONE has not been included in the comparison, as its development seems to have stopped years ago.

All computations were run on a Linux cluster with Intel Xeon E5-2670 v2 CPUs (20 cores). The GAMS version is 41.2.0, which includes SCIP 8.0.2, BARON 22.9.30, Lindo API 14.0.5099.162, and Octeract 4.5.1. A GAMS license with all solvers enabled was used, so that SCIP uses CPLEX 22.1.0.0 as LP solver and Ipopt with HSL MA27 as NLP solver, BARON can choose between all LP/MIP/NLP solvers that it interfaces with, and Octeract uses CPLEX 22.1.0.0 as LP/MIP/QP/QCP solver.

3.1 Test Set

To construct a test set suitable for benchmarking, the MINLPLib [50] collection of 1595 MINLP instances was used as source. First, all instances that could not be handled by some of the considered solvers were excluded, e.g., instances with trigonometric functions, as they are not supported by BARON. All solvers were then run in serial mode (that is, with parallelization features disabled) on the remaining 1505 instances and using the parameter settings described below. The results of these runs were then used to select a set of 200 instances that could be solved by at least one solver, that were not all trivial, had a varying degree of integrality and nonlinearity, and such that having many instances with a similar name is avoided. The latter was done to avoid overrepresentation of optimization problems for which many instances were added to MINLPLib.

Since small changes to an instance can lead to large variations in the solver’s performance, the benchmark’s reliability is improved by considering for each instance four additional variants where the order of variables and equations has been permuted. The permuted instances were generated with GAMS/Convert. Thus, a test set of 1000 instances is obtained.

The following approach was used to select the 200 instances before permutation: Let I be the set of 1505 instances, d_i be the fraction of integer variables in instance $i \in I$, and e_i be the fraction of nonzeros in the Jacobian and objective function gradient that correspond to nonlinear terms. Next, assign to each instance an identifier $f_i \in F$ such that instances that seem to come from the same model are assigned the same identifier. This goal is approximated by mapping i to the name of the instance until the first digit, underscore, or dash, except for the block layout design instances `fo*`, `m*`, `no*`, `o*`, which were all assigned to the same identifier. $|F| = 230$ different identifiers were found this way.

Further, let \bar{t}_i be the largest time in seconds that any solver who did not produce wrong results on instance i spend on instance i . Finally, let S be the number of instances that could be solved by at least one solver.

To ensure that instances with a varying amount of integer variables and nonlinearity are included, the interval $[0, 1]$ was split once at breakpoints 0.05, 0.25, 0.5, 0.9 and once at 0.1, 0.25, 0.5. Let D and E be the resulting partitions of $[0, 1]$. For every interval from D and E , the aim is to have roughly the same number of instances with d_i and e_i in the respective intervals. For the choice of breakpoints that define D and E , the distribution of d_i and e_i , $i \in I$, have been taken into account. For example, MINLPLib contains many purely continuous and purely discrete instances, but not many instances that are mostly linear or completely nonlinear.

To avoid including too many instances originating from the same model, including more than two instances for each identifier in F is discouraged. Further, instances that seem trivial, i.e., which are solved by all solvers in no more than five seconds, or could not be solved by any solver are excluded. Introducing penalty terms, the following

optimization problem for instance selection is obtained:

$$\begin{aligned}
& \min \sum_{d \in D} \lambda_d^2 + \sum_{e \in E} \lambda_e^2 + 10 \sum_{f \in F} \lambda_f^2 \\
& \text{such that} \quad \sum_{i \in I: d_i \in d} z_i = \left\lfloor \frac{N}{|D|} \right\rfloor + \lambda_d \quad \forall d \in D, \\
& \quad \sum_{i \in I: e_i \in e} z_i = \left\lfloor \frac{N}{|E|} \right\rfloor + \lambda_e \quad \forall e \in E, \\
& \quad \sum_{i \in I: f_i = f} z_i \leq 2 + \lambda_f \quad \forall f \in F, \\
& \quad z_i = 0 \quad \forall i \in I : \bar{t}_i \leq 5, \\
& \quad z_i = 0 \quad \forall i \in I : i \notin S, \\
& \quad z \in \{0, 1\}^{|I|}, \lambda \in \mathbb{Z}^{|D|+|E|+|F|}
\end{aligned}$$

This problem was solved for N varying between 180 and 220. For $N = 208$, this yield a selection of 200 instances with an acceptable penalty value of 106. See Section A for a list of all selected instances. Table 1 shows the number of instances for each element of $D \times E$. For five identifiers from F , three instead of two instances were selected, i.e., $\lambda_f = 1$ for five $f \in F$.

$E \downarrow \mid D \rightarrow \mid$	$[0, 0.05)$	$[0.05, 0.25)$	$[0.25, 0.5)$	$[0.5, 0.9)$	$[0.9, 1]$	$[0, 1]$
$[0, 0.1)$	3	7	19	15	6	50
$[0.1, 0.25)$	8	22	9	7	4	50
$[0.25, 0.5)$	8	8	6	10	18	50
$[0.5, 1]$	25	2	5	7	11	50
$[0, 1]$	44	39	39	39	39	200

Table 1: Number of instances selected with “discreteness” d_i and “nonlinearity” e_i in intervals from D and E .

3.2 Parameter Settings

3.2.1 Missing Variable Bounds

To compute a lower bound on the optimal value of a minimization problem, all solvers considered here construct a convex relaxation of the given problem. For nonconvex constraints, this often relies on the computation of valid convex underestimators or concave overestimators. As these typically depend on variables’ bounds (recall the McCormick underestimators (1)), missing or very large bounds on variables in nonconvex terms can mean that an instance will be very hard or impossible to solve.

Even when the user forgot to specify some variable bounds, the solver may still be able to derive bounds via domain propagation. Further, once a feasible solution \hat{x} has been found, additional bounds may be derived from the inequality $c^\top x \leq c^\top \hat{x}$. However, as there are always cases where bounds are still missing after presolve, solvers invented different ways to deal with this obstacle.

If SCIP cannot construct an under- or overestimator because of missing variable bounds, it continues by branching on an unbounded variable. This way, there will eventually be a node in the branch-and-bound tree where all variables are bounded. Nodes that still contain unbounded variable domains may be pruned due to a derived lower bound on the objective function exceeding the incumbents objective function

value. But it may also be the case that pruning will not be possible and SCIP does not terminate. However, variable bounds after branching cannot grow indefinitely in SCIP, but are limited by $\pm 10^{20}$ by default. That is, SCIP does not search for solutions with variable values beyond this value.

The other solvers considered here add variable bounds based on a heuristic decision. If BARON is still missing bounds on variables in nonconvex terms after presolve, it sets the bound to a value that depends on the type of nonlinearity involved. Typically, this value is around $\pm 10^{10}$. BARON also prints a warning to the log and no longer claim to have solved a problem to global optimality, i.e., it does not return a lower bound. Lindo API adjusts the bounds for all variables that are involved in convexification to be within $[-10^{10}, 10^{10}]$. At termination, it returns the lower bound for the restricted problem. Octeract proceeds similarly and introduces a bound of $\pm 10^7$ for every missing bound and returns the lower bound for the restricted problem at termination.

Evidently, just passing an instance with unbounded variables to a solver with default settings may mean that each solver solves a different subproblem of the actual problem and often also reports a lower bound that corresponds to the solved subproblem only. Fortunately, for every solver considered here, parameters are available to adjust the treatment of unbounded variables. A first impulse could be to tell all solvers to set missing bounds to infinity, but this is not possible as each solver treats values beyond a certain finite value as “infinity” (BARON: 10^{50} , Octeract: 10^{308} , SCIP: 10^{20}). Changing this value is either not possible or not advisable.

We therefore decided to aim for $\pm 10^{12}$ as replacement for a missing variable bound. For BARON and SCIP, the GAMS interface can replace any missing bound by $\pm 10^{12}$ before the instance is passed to the solver. BARON will hence also return a lower bound for this restricted problem. For Lindo API, a solver parameter can be changed so that bounds for all variables subject to convexification are bounded by $\pm 10^{12}$ (instead of $\pm 10^{10}$). Finally, also for Octeract, all missing bounds are set to $\pm 10^{12}$ (instead of $\pm 10^7$) by changing of a solver parameter. Note, that this still does not ensure that all solvers solve the same instance, since Lindo API would still change initial finite bounds beyond 10^{12} and may also not set any bounds for variables that are not involved in convexification.

Next to missing bounds on problem variables, also singularities in functions (e.g., $1/x$, $\log(x)$) can prevent finite under- or overestimators from being available. Unfortunately, there are no parameters available to ensure a uniform treatment of this case in all solvers. SCIP ensures that a variable x in x^p , $p < 0$, or $\log(x)$ is bounded away from zero by 10^{-9} , and terminates with a lower bound for this modified problem. BARON applies the same method as the one for missing bounds on problem variables to choose a suitable bound on x . No lower bound is returned at termination then. The methods in Lindo API and Octeract are not known to us.

3.2.2 Solution Quality

To ensure that all solvers return solutions of the same quality, constraints of (MINLP) are required to be satisfied with an absolute tolerance of 10^{-6} . This applies to linear and nonlinear equations, variable bounds, and integrality.

In addition, a tolerance on the proof of optimality is set. For this purpose, typically, solvers are allowed to stop when the absolute or relative gap between lower and upper bounds on the optimal value are sufficiently small. Since the test set is diverse and has optimal values of varying magnitude, setting only a relative gap limit and no absolute gap limit would be preferable. Unfortunately, Octeract does not permit different values for these limits. As a compromise, BARON, Lindo API, and SCIP are run with 10^{-4} as relative gap limit and 10^{-6} as absolute gap limit, while for Octeract, 10^{-6} is used for both the absolute and relative gap limit. Below, the impact of using a tighter optimality tolerance for Octeract is analyzed in a separate comparison.

3.2.3 Working Limits

As working limits, a time limit of two hours is used and the jobs on the cluster are restricted to 50 GB of RAM. Further, the amount parallelization (multiple threads or processes) that a solver is allowed to use is limited in varying degrees. To simplify the presentation, the term “threads” is used also for Octeract, even though it uses multiple processes instead of threads to parallelize its solving process.

3.2.4 Summary

To summarize, the following parameters are used:

GAMS (applied to all solvers): `optcr=1e-4`, `optca=1e-6`, `reslim=7200`, `workspace=50000`, `threads` $\in \{1, 4, 8, 16\}$

BARON: `InfBnd=1e12`, `AbsConFeasTol=1e-6`, `AbsIntFeasTol=1e-6`

Lindo API: `GOP_BNDLIM=1e12`, `SOLVER_FEASTOL=1e-6`

Octeract: `INFINITY=1e12`, `INTEGRALITY_VIOLATION_TOLERANCE=1e-6`

SCIP: `gams/infbound=1e12`, `constraints/nonlinear/linearizeheursol=o` (this undoes a change in the algorithmic settings of SCIP that is part of the GAMS/SCIP interface)

3.3 Correctness Checks

The GAMS/Examiner 2.0 tool is used to evaluate the violation of constraints, bounds, and integrality in the solutions reported by the solver. Examiner generates for each solver a file that contains for each instance the solving time, returned lower and upper bound, and solution infeasibility.

A run of a solver on an instance is marked as *failed* if the solver terminated abnormally, the solution is not feasible with respect to the feasibility tolerance, or the lower or upper bound contradicts with the bounds on the optimal value that are specified on the MINLPLib page. Note, that the primal and dual bounds on the MINLPLib page were calculated without enforcing the $\pm 10^{12}$ limit on unbounded variables. However, in order for an instance to be accepted into the test set, one of the solvers considered here must have solved the instance and found an optimal value that fits within the lower and upper bounds given at MINLPLib. It is therefore acceptable to use these bounds for checking.

A run that has not failed is marked as *solved* if the relative or absolute limits on the gap between lower and upper bound are satisfied. If a solver stopped without closing the gap before the time limit, then the solver time is changed to the time limit. The only exception here is BARON, which stops on two instances before the time limit without reporting a lower bound due to singularities in functions (see Section 3.2.1). To be consistent with the treatment of other solvers, these two instances were accounted as solved by BARON with the original solver time.

3.4 Results

3.4.1 Serial Mode

For the main comparison, all parallelization features in the solvers were disabled, that is, GAMS was run with option `threads` set to 1. In addition to the solver itself, results for the *virtual best* and *virtual worst* solver are reported, which are obtained by picking for each instance the fastest or the slowest solver, respectively.

Table 2 shows for each solver the number of instances that could be solved, the number of times the time limit was reached, and the number of runs that were marked as failed. Further, the shifted geometric mean of the running time of the solver is

provided. The shift has been set to 1 second. Here, instances that failed are accounted with the time limit. The performance profile [23] in Figure 4 shows the number of instances a solver solved within a time that is at most a factor of the fastest solvers time. Section B.1 provides detailed results.

	solved	timeout	fail	time
BARON	790	183	27	75.4
Lindo API	538	323	139	489.1
Octeract	671	279	50	184.1
SCIP	776	183	41	85.2
virt. worst	368	405	227	1505.2
virt. best	967	33	0	19.7

Table 2: Aggregated performance data for all solvers on test set of 1000 instances with parallelization disabled.




Figure 4: Performance profile comparing all solvers with parallelization disabled.

The results show a small lead of BARON before SCIP with respect to both number of instances solved and average time. Since the number of timeouts is almost equal, one could argue that it is the higher stability of BARON that moves it onto the first place here. In fact, the 41 fails of SCIP are due to returning a wrong optimal value 16 times, returning an infeasible solution 23 times, and aborts due to numerical troubles for two instances. For BARON, fails are due to returning a wrong optimal value 26 times and an infeasible solution only once. While SCIP 8.0 has made a large step forward in ensuring that nonlinear constraints are satisfied in the non-presolved problem, violations in linear constraints or variable bounds still occur for a few instances. These are typically due to variables being aggregated during presolve.

Even though Octeract and Lindo API solved considerably fewer instances than BARON and SCIP, which also results in an increased mean time, it is noteworthy that each of the two is also the fastest solver on 270 and 66 instances, respectively. Octeract also produced correct results for 95% of the test set, while for Lindo API a relatively high number of wrong optimal values, infeasible solutions, or aborts is observed.

The large differences between the real and virtual solvers show that none of the solvers dominates all others or is dominated.

Next, the effect of changing the gap limit for Octeract has been investigated. Recall from Section 3.2.2 that relative and absolute gap limits of 10^{-6} and 10^{-4} , respectively, were used for all solvers except for Octeract. Since Octeract does not allow choosing

these limits separately, it had been run with the tighter relative gap limit of 10^{-6} . To check whether this lead to a considerable disadvantage for this solver, the solver was rerun on the 200 non-permuted instances with both relative and absolute gap limit set to 10^{-4} . The table in Section B.2 shows that the change in the convergence tolerance had essentially no effect on the solver’s performance. In both cases, the same 134 instances could be solved. The mean time changed from 178.6 for a limit of 10^{-6} to 179.0 for a limit of 10^{-4} .

3.4.2 Parallel Mode

In the next comparison, each solver is allowed to use multiple threads or processes. Since SCIP’s use of multiple threads is limited to presolving MIPs, checking quadratic functions for convexity, and the linear algebra in Ipopt, FiberSCIP [64] is used to run SCIP in parallel mode. FiberSCIP is a shared-memory instantiation of the UG framework [62] for the parallelization of branch-and-bound based solvers. The framework parallelizes the search of the branch-and-bound tree by collecting and distributing open problems between independent instances of SCIP. In addition, the first seconds of the solving process are used for a “racing ramp-up” phase. Here, multiple SCIP instances with differing parameter sets are run concurrently, and the one with the best lower bound is used for the remaining solve. The UG version was 1.0.0 beta3.

For the runs in serial mode, reaching the memory limit of 50 GB was not observed for any solver. But since parallelization often increases memory requirements, a memory limit of 100 GB has been used for the runs in parallel mode. Since this meant a reduction in available computing resources, only the 200 non-permuted instances are used for comparisons.

Table 3 shows, for an increasing number of threads, the number of instances that could be solved by each solver and the mean time spent. In addition, Figure 5 provides a performance profile that compares SCIP and FiberSCIP only. Section B.3 gives detailed results.

	1 thread		4 threads		8 threads		16 threads	
	solved	time	solved	time	solved	time	solved	time
BARON	161	64.3	160	58.2	160	57.1	158	58.6
Lindo API	114	423.6	114	379.2	106	459.5	107	456.4
Octeract	134	178.6	133	146.9	138	118.1	135	123.2
(Fiber)SCIP	161	76.9	145	94.3	147	77.8	152	74.8

Table 3: Aggregated performance data for all solvers on test set of 200 instances when run with parallelization allowed.

Apparently, enabling parallelization seldom has a considerable advantage on this test set. For Octeract, where parallelization was part of its original design, a small increase in the number of instances that could be solved and a reduction in time by 34% when using up to 8 parallel processes is observed. As far as we know, BARON’s use of multiple threads is currently limited to enabling this feature in the solver for a MIP relaxation. As a consequence, only moderate improvement of running time by up to 11% are seen. For Lindo API, an improvement due to parallelization seems to be impeded by a further increase in fails when using multiple threads (1 thread: 24, 4: 28, 8: 35, 16: 43). Finally, for SCIP/FiberSCIP the additional overhead due to the parallelization being build on top of the solver instead of being tightly integrated is not compensated by the use of multiple threads. However, in contrast to other solvers, a monotonous improvement in both number of instances solved and mean solving time when increasing from 4 to 16 threads is observed. Further, the virtual solvers in the




Figure 5: Performance profile comparing SCIP and FiberSCIP.

performance profile show that FiberSCIP can solve instances that SCIP on one thread couldn't solve.

Finally, note that a benefit due to parallelization can usually only be expected for rather challenging instances because of the additional overhead in duplicating and synchronizing data and processes. However, the test set deliberately included only instances that could already be solved by some solver in serial mode, and only instances that were trivial for all solvers, though they may be solved quickly by some, were excluded. As a small experiment, for each solver only those instances that required at least 10 or 100 seconds to solve in serial mode were considered. Unfortunately, this essentially repeated the trends shown in Table 3, so details are omitted here. A more thorough analysis of the parallelization capabilities of MINLP solvers using a set of challenging instances only would be necessary, but exceeds the scope of this paper.

4 Conclusion

The development of the MINLP solver in SCIP has come a long way. In a recent version-to-version comparison [57, slides 49-51], a steady improvement in the performance of SCIP on MINLP over the last ten years has been measured, resulting in SCIP 8 solving twice as many instances as SCIP 3 and a speed-up of factor three. Partially, this improvement has been achieved by improving and adding features particular for MINLP. However, due to the generality of SCIP as a CIP solver, also many developments that targeted MIP solving were immediately available for MINLP solving.

With version 8.0, the MINLP solving capabilities of SCIP have been largely reworked and extended, which resulted in a considerable improvement in both robustness and performance [14, 57]. As a result, SCIP's performance is currently on par with the state-of-the-art commercial solver BARON.

In contrast to the commercial solvers considered here, SCIP offers a variety of possibilities for a user, developer, or researcher to interact with the solving process. In particular, the newly added “nonlinear handler” plugin type sets SCIP apart from most other MINLP solvers, as it allows focusing on experimenting with new algorithms to handle certain structures in nonlinear functions without modifications to the solvers code.

The rather large number of features that are disabled by default shows that tuning and improving the existing code base has become increasingly necessary. Future work will of course also include the addition of new features, e.g., an improved separation for signomial functions [77], the use of alternative relaxations for polynomial functions [17],

or monoidal strengthening of intersection cuts for quadratic constraints [22].

The increasing number of cores in present-day CPUs means that to fully utilize an ordinary desktop computer, a solver needs to be parallelized. While the UG framework provides such a possibility for SCIP in both shared and distributed memory environments, the experiments with FiberSCIP on up to 16 threads show that more tuning is necessary to ensure that the additional overhead can be compensated by the use of additional computing resources. Since the development of UG was initially motivated and has focused primarily on the use of large-scale parallel computing environments [63], an investigation on using UG with SCIP to solve challenging MINLPs in distributed memory environments with many CPU cores could be interesting as well.

Acknowledgments

We are very much in all SCIP developers’ debt – the extensions to support nonlinear constraints and solve MINLPs would not have been possible without the framework’s existence and the powerful MIP solver that we could build upon. While the authors of this paper are the main developers of the new MINLP features in SCIP 8, many have contributed to the MINLP capabilities in previous releases of SCIP, namely Martin Ballerstein, Timo Berthold, Tobias Fischer, Thorsten Gellermann, Ambros Gleixner, Renke Kuhlmann, Dennis Michaels, Marc Pfetsch, and Stefan Weltge. Further, we thank Yuji Shinano for the development of FiberSCIP and swiftly reacting to our request for the possibility to set gap limits. Last but not least, we are very grateful to Franziska Schlösser for the setup and maintenance of benchmarking and testing facilities for the infamous “consexpr” development branch of SCIP.

The work for this article has been conducted within the Research Campus Modal funded by the German Federal Ministry of Education and Research (BMBF grant numbers 05M14ZAM, 05M20ZBM). Additional funding has been received from the German Federal Ministry for Economic Affairs and Energy within the project EnBA-M (ID: 03ET1549D).

References

- [1] T. Achterberg. *Constraint Integer Programming*. PhD thesis, Technische Universität Berlin, 2007.
- [2] W. P. Adams and H. D. Sherali. A tight linearization and an algorithm for zero-one quadratic programming problems. *Management Science*, 32(10):1274–1290, 1986. doi:10.1287/mnsc.32.10.1274.
- [3] W. P. Adams and H. D. Sherali. Linearization strategies for a class of zero-one mixed integer programming problems. *Operations Research*, 38(2):217–226, 1990. doi:10.1287/opre.38.2.217.
- [4] C. Adjiman and C. Floudas. Rigorous convex underestimators for general twice-differentiable problems. *Journal of Global Optimization*, 9(1):23–40, 1996. doi:10.1007/BF00121749.
- [5] E. Balas. Intersection cuts – a new type of cutting planes for integer programming. *Operations Research*, 19(1):19–39, 1971. doi:10.1287/opre.19.1.19.
- [6] M. Ballerstein, D. Michaels, and S. Vigerske. Linear underestimators for bivariate functions with a fixed convexity behavior. ZIB-Report 13-02, Zuse Institute Berlin, 2015. urn:nbn:de:0297-zib-17641.

- [7] E. Beale. Branch and bound methods for numerical optimization of non-convex functions. In M. Barritt and D. Wishart, editors, *COMPSTAT 80 Proceedings in Computational Statistics*, pages 11–20, Vienna, 1980. Physica-Verlag.
- [8] B. Bell. CppAD: A package for differentiation of C++ algorithms. <https://github.com/coin-or/CppAD/>.
- [9] P. Belotti, C. Kirches, S. Leyffer, J. Linderoth, J. Luedtke, and A. Mahajan. Mixed-integer nonlinear optimization. *Acta Numerica*, 22:1–131, 2013. doi:10.1017/S0962492913000032.
- [10] T. Berthold and A. Gleixner. Undercover: a primal MINLP heuristic exploring a largest sub-MIP. *Mathematical Programming*, 144(1-2):315–346, 2014. doi:10.1007/s10107-013-0635-2.
- [11] T. Berthold, S. Heinz, and M. E. Pfetsch. Nonlinear pseudo-boolean optimization: relaxation or propagation? In O. Kullmann, editor, *Theory and Applications of Satisfiability Testing – SAT 2009*, number 5584 in LNCS, pages 441–446. Springer, 2009. doi:10.1007/978-3-642-02777-2_40.
- [12] T. Berthold, S. Heinz, M. E. Pfetsch, and S. Vigerske. Large neighborhood search beyond MIP. In L. D. Gaspero, A. Schaerf, and T. Stützle, editors, *Proceedings of the 9th Metaheuristics International Conference (MIC 2011)*, pages 51–60, 2011.
- [13] T. Berthold, S. Heinz, and S. Vigerske. Extending a CIP framework to solve MIQCPs. In J. Lee and S. Leyffer, editors, *Mixed Integer Nonlinear Programming*, volume 154 of *The IMA Volumes in Mathematics and its Applications*, pages 427–444. Springer, 2012. doi:10.1007/978-1-4614-1927-3_15.
- [14] K. Bestuzheva, M. Besançon, W.-K. Chen, A. Chmiela, T. Donkiewicz, J. van Doornmalen, L. Eifler, O. Gaul, G. Gamrath, A. Gleixner, L. Gottwald, C. Graczyk, K. Halbig, A. Hoen, C. Hojny, R. van der Hulst, T. Koch, M. Lübbecke, S. J. Maher, F. Matter, E. Mühmer, B. Müller, M. E. Pfetsch, D. Rehfeldt, S. Schlein, F. Schlösser, F. Serrano, Y. Shinano, B. Sofranac, M. Turner, S. Vigerske, F. Wegscheider, P. Wellner, D. Weninger, and J. Witzig. The SCIP Optimization Suite 8.0. ZIB Report 21-41, Zuse Institute Berlin, 2021. urn:nbn:de:0297-zib-85309.
- [15] K. Bestuzheva, A. Gleixner, and S. Vigerske. A computational study of perspective cuts. ZIB Report 21-07, Zuse Institute Berlin, 2021.
- [16] K. Bestuzheva, A. Gleixner, and T. Achterberg. Efficient separation of RLT cuts for implicit and explicit bilinear products. Technical Report 2211.13545, arXiv, 2022. doi:10.48550/ARXIV.2211.13545.
- [17] K. Bestuzheva, A. Gleixner, and H. Völker. Strengthening SONC relaxations with constraints derived from variable bounds. In preparation, 2022.
- [18] A. Bley, T. Koch, and L. Niu. Experiments with nonlinear extensions to SCIP. ZIB-Report 08-28, Zuse Institute Berlin, 2008. urn:nbn:de:0297-zib-8300.
- [19] M. R. Bussieck and S. Vigerske. MINLP solver software. In J. J. Cochran, L. A. Cox, Jr., P. Keskinocak, J. P. Kharoufeh, and J. C. Smith, editors, *Wiley Encyclopedia of Operations Research and Management Science*. Wiley & Sons, Inc., 2010. doi:10.1002/9780470400531.eorms0527.
- [20] C. Büskens and D. Wassel. The ESA NLP solver WORHP. In G. Fasano and J. D. Pintér, editors, *Modeling and Optimization in Space Engineering*, volume 73, pages 85–110. Springer New York, 2013. doi:10.1007/978-1-4614-4469-5_4.

- [21] A. Chmiela, G. Muñoz, and F. Serrano. On the implementation and strengthening of intersection cuts for QCQPs. In M. Singh and D. P. Williamson, editors, *Integer Programming and Combinatorial Optimization*, pages 134–147. Springer, 2021. doi:10.1007/978-3-030-73879-2_10.
- [22] A. Chmiela, G. Muñoz, and F. Serrano. Monoidal strengthening and unique lifting in MIQCPs. Submitted, preprint available at <https://www.gonzalomunoz.org/publications/>, 2022.
- [23] E. D. Dolan and J. J. Moré. Benchmarking optimization software with performance profiles. *Mathematical Programming*, 91(2):201–213, 2002. doi:10.1007/s101070100263.
- [24] F. Domes and A. Neumaier. Constraint propagation on quadratic constraints. *Constraints*, 15(3):404–429, 2010. doi:10.1007/s10601-009-9076-1.
- [25] 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. doi:10.1007/BF02592064.
- [26] T. Fischer. *Branch-and-cut for complementarity and cardinality constrained linear programs*. PhD thesis, Technical University of Darmstadt, 2017.
- [27] R. Fletcher and S. Leyffer. User manual for filterSQP. Numerical Analysis Report NA/181, Department of Mathematics, University of Dundee, Scotland, 1998.
- [28] C. A. Floudas. *Nonlinear and Mixed Integer Optimization: Fundamentals and Applications*. Oxford University Press, New York, 1995.
- [29] A. Frangioni and C. Gentile. Perspective cuts for a class of convex 0–1 mixed integer programs. *Mathematical Programming*, 106(2):225–236, 2006. doi:10.1007/s10107-005-0594-3.
- [30] G. Gamrath, T. Fischer, T. Gally, A. Gleixner, G. Hendel, T. Koch, S. J. Maher, M. Miltenberger, B. Müller, M. E. Pfetsch, C. Puchert, D. Rehfeldt, S. Schenker, R. Schwarz, F. Serrano, Y. Shinano, S. Vigerske, D. Weninger, M. Winkler, J. T. Witt, and J. Witzig. The SCIP Optimization Suite 3.2. ZIB Report 15-60, Zuse Institute Berlin, 2016. urn:nbn:de:0297-zib-57675.
- [31] G. Gamrath, D. Anderson, K. Bestuzheva, W.-K. Chen, L. Eifler, M. Gasse, P. Gemander, A. Gleixner, L. Gottwald, K. Halbig, G. Hendel, C. Hojny, T. Koch, P. L. Bodic, S. J. Maher, F. Matter, M. Miltenberger, E. Mühmer, B. Müller, M. E. Pfetsch, F. Schlösser, F. Serrano, Y. Shinano, C. Tawfik, S. Vigerske, F. Wegscheider, D. Weninger, and J. Witzig. The SCIP Optimization Suite 7.0. ZIB Report 20-10, Zuse Institute Berlin, 2020. urn:nbn:de:0297-zib-78023.
- [32] A. Gleixner, H. Held, W. Huang, and S. Vigerske. Towards globally optimal operation of water supply networks. *Numerical Algebra, Control and Optimization*, 2(4):695–711, 2012. doi:10.3934/naco.2012.2.695.
- [33] A. Gleixner, T. Berthold, B. Müller, and S. Weltge. Three enhancements for optimization-based bound tightening. *Journal of Global Optimization*, 67(4):731–757, 2017. doi:10.1007/s10898-016-0450-4.
- [34] A. Gleixner, L. Eifler, T. Gally, G. Gamrath, P. Gemander, R. L. Gottwald, G. Hendel, C. Hojny, T. Koch, M. Miltenberger, B. Müller, M. E. Pfetsch, C. Puchert, D. Rehfeldt, F. Schlösser, F. Serrano, Y. Shinano, J. M. Viernickel, S. Vigerske, D. Weninger, J. T. Witt, and J. Witzig. The SCIP Optimization Suite 5.0. ZIB Report 17-61, Zuse Institute Berlin, 2017. urn:nbn:de:0297-zib-66297.

- [35] A. Gleixner, S. J. Maher, B. Müller, and J. P. Pedroso. Price-and-verify: a new algorithm for recursive circle packing using Dantzig–Wolfe decomposition. *Annals of Operations Research*, 284(2):527–555, 2020. doi:10.1007/s10479-018-3115-5.
- [36] F. Glover. Polyhedral convexity cuts and negative edge extensions. *Zeitschrift für Operations Research*, 18:181–186, 1974. doi:10.1007/BF02026599.
- [37] I. E. Grossmann and Z. Kravanja. Mixed-integer nonlinear programming: A survey of algorithms and applications. In A. R. Conn, L. T. Biegler, T. F. Coleman, and F. N. Santosa, editors, *Large-Scale Optimization with Applications, Part II: Optimal Design and Control*, volume 93 of *The IMA Volumes in Mathematics and its Applications*, pages 73–100. Springer, New York, 1997. doi:10.1007/978-1-4612-1960-6_5.
- [38] P. Hansen, B. Jaumard, M. Ruiz, and J. Xiong. Global minimization of indefinite quadratic functions subject to box constraints. *Naval Research Logistics (NRL)*, 40(3):373–392, 1993. doi:10.1002/1520-6750(199304)40:3<373::AID-NAV3220400307>3.0.CO;2-A.
- [39] C. Hojny and M. E. Pfetsch. Polytopes associated with symmetry handling. *Mathematical Programming*, 175(1):197–240, 2019. doi:10.1007/s10107-018-1239-7.
- [40] A. Khajavirad and N. V. Sahinidis. A hybrid LP/NLP paradigm for global optimization relaxations. *Mathematical Programming Computation*, 10(3):383–421, 2018. doi:10.1007/s12532-018-0138-5.
- [41] G. Kocis and I. Grossmann. Computational experience with DICOPT: Solving MINLP problems in process systems engineering. *Computers & Chemical Engineering*, 13(3):307–315, 1989. doi:10.1016/0098-1354(89)85008-2.
- [42] 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. doi:10.1007/s10898-015-0322-3.
- [43] L. Liberti. Reformulations in mathematical programming: automatic symmetry detection and exploitation. *Mathematical Programming*, 131(1):273–304, 2012. doi:10.1007/s10107-010-0351-0.
- [44] Y. Lin and L. Schrage. The global solver in the LINDO API. *Optimization Methods & Software*, 24(4–5):657–668, 2009. doi:10.1080/10556780902753221.
- [45] M. Locatelli. Convex envelopes of bivariate functions through the solution of KKT systems. *Journal of Global Optimization*, 72(2):277–303, 2018. doi:10.1007/s10898-018-0626-1.
- [46] A. Mahajan and T. Munson. Exploiting second-order cone structure for global optimization. Technical Report ANL/MCS-P1801-1010, Argonne National Laboratory, 2010.
- [47] S. J. Maher, T. Fischer, T. Gally, G. Gamrath, A. Gleixner, R. L. Gottwald, G. Hendel, T. Koch, M. E. Lübbecke, M. Miltenberger, B. Müller, M. E. Pfetsch, C. Puchert, D. Rehfeldt, S. Schenker, R. Schwarz, F. Serrano, Y. Shinano, D. Weninger, J. T. Witt, and J. Witzig. The SCIP Optimization Suite 4.0. ZIB Report 17-12, Zuse Institute Berlin, 2017. urn:nbn:de:0297-zib-62170.
- [48] 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*, pages 647–686. Springer, 2010. doi:10.1007/978-3-540-68279-0_17.

- [49] G. P. McCormick. Computability of global solutions to factorable nonconvex programs: Part I – convex underestimating problems. *Mathematical Programming*, 10(1):147–175, 1976. doi:10.1007/bf01580665.
- [50] MINLPLib. A library of mixed-integer and continuous nonlinear programming instances. <https://www.minplib.org>, 2022-10-14.
- [51] R. Misener and C. A. Floudas. Global optimization of mixed-integer quadratically-constrained quadratic programs (MIQCQP) through piecewise-linear and edge-concave relaxations. *Mathematical Programming*, 136(1):155–182, 2012. doi:10.1007/s10107-012-0555-6.
- [52] R. Misener, J. B. Smadbeck, and C. A. Floudas. Dynamically generated cutting planes for mixed-integer quadratically constrained quadratic programs and their incorporation into GloMIQO 2. *Optimization Methods and Software*, 30(1):215–249, 2015. doi:10.1080/10556788.2014.916287.
- [53] R. E. Moore. *Interval Analysis*. Englewood Cliffs, NJ: Prentice Hall, 1966.
- [54] B. Müller, F. Serrano, and A. Gleixner. Using two-dimensional projections for stronger separation and propagation of bilinear terms. *SIAM Journal on Optimization*, 30(2):1339–1365, 2020. doi:10.1137/19m1249825.
- [55] G. Muñoz and F. Serrano. Maximal quadratic-free sets. In D. Bienstock and G. Zambelli, editors, *Integer Programming and Combinatorial Optimization*, pages 307–321. Springer, 2020. doi:10.1007/978-3-030-45771-6_24.
- [56] B. Müller, R. Kuhlmann, and S. Vigerske. On the performance of NLP solvers within global MINLP solvers. In *Operations Research Proceedings 2017*, pages 633–639, Cham, 2018. Springer International Publishing. doi:10.1007/978-3-319-89920-6_84.
- [57] M. Pfetsch. SCIP: Past, Present, Future. Presentation at workshop *Let’s SCIP it!*, November 2022. URL <https://scipopt.org/20years/slides/pfetsch.pdf>.
- [58] J. D. Pintér, editor. *Global Optimization: Scientific and Engineering Case Studies*, volume 85 of *Nonconvex Optimization and Its Applications*. Springer, 2006. doi:10.1007/0-387-30927-6.
- [59] I. Quesada and I. E. Grossmann. Global optimization algorithm for heat exchanger networks. *Industrial & Engineering Chemistry Research*, 32(3):487–499, 1993. doi:10.1021/ie00015a012.
- [60] N. Sahinidis. BARON: A general purpose global optimization software package. *Journal of Global Optimization*, 8(2):201–205, 1996. doi:10.1007/BF00138693.
- [61] L. Schewe and M. Schmidt. Computing feasible points for binary MINLPs with MPECs. *Mathematical Programming Computation*, 11(1):95–118, 2019. doi:10.1007/s12532-018-0141-x.
- [62] Y. Shinano. UG – Ubiquity Generator framework v1.0.0beta. <https://ug.zib.de>, 2021. doi:10.12752/8521.
- [63] Y. Shinano, T. Achterberg, T. Berthold, S. Heinz, T. Koch, and M. Winkler. Solving open MIP instances with ParaSCIP on supercomputers using up to 80,000 cores. In *2016 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, pages 770–779, 2016. doi:10.1109/IPDPS.2016.56.

- [64] Y. Shinano, S. Heinz, S. Vigerske, and M. Winkler. FiberSCIP – a shared memory parallelization of SCIP. *INFORMS Journal on Computing*, 30(1):11–30, 2018. doi:10.1287/ijoc.2017.0762.
- [65] E. Smith and C. Pantelides. A symbolic reformulation/spatial branch-and-bound algorithm for the global optimisation of nonconvex MINLPs. *Computers & Chemical Engineering*, 23(4-5):457–478, 1999. doi:10.1016/s0098-1354(98)00286-5.
- [66] L. Smith, J. Chinneck, and V. Aitken. Improved constraint consensus methods for seeking feasibility in nonlinear programs. *Computational Optimization and Applications*, 54(3):555–578, 2013. doi:10.1007/s10589-012-9473-z.
- [67] F. Tardella. On the existence of polyhedral convex envelopes. In C. A. Floudas and P. Pardalos, editors, *Frontiers in Global Optimization*, pages 563–573, Boston, MA, 2004. Springer US. doi:10.1007/978-1-4613-0251-3_30.
- [68] M. Tawarmalani and N. V. Sahinidis. A polyhedral branch-and-cut approach to global optimization. *Mathematical Programming*, 103(2):225–249, 2005. doi:10.1007/s10107-005-0581-8.
- [69] F. Trespalcios and I. Grossmann. Review of mixed-integer nonlinear and generalized disjunctive programming methods. *Chemie Ingenieur Technik*, 86(7):991–1012, 2014. doi:10.1002/cite.201400037.
- [70] H. Tuy. Concave programming with linear constraints. *Doklady Akademii Nauk*, 159(1):32–35, 1964.
- [71] A. F. Veinott. The supporting hyperplane method for unimodal programming. *Operations Research*, 15(1):147–152, 1967. doi:10.1287/opre.15.1.147.
- [72] J. P. Vielma, I. Dunning, J. Huchette, and M. Lubin. Extended formulations in mixed integer conic quadratic programming. *Mathematical Programming Computation*, 9(3):369–418, 2016. doi:10.1007/s12532-016-0113-y.
- [73] S. Vigerske. *Decomposition of Multistage Stochastic Programs and a Constraint Integer Programming Approach to Mixed-Integer Nonlinear Programming*. PhD thesis, Humboldt-Universität zu Berlin, 2013. urn:nbn:de:kobv:11-100208240.
- [74] 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, 2017. doi:10.1080/10556788.2017.1335312.
- [75] A. Wächter and L. Biegler. On the implementation of a primal-dual interior point filter line search algorithm for large-scale nonlinear programming. *Mathematical Programming*, 106(1):25–57, 2006. doi:10.1007/s10107-004-0559-y.
- [76] F. Wegscheider. Exploiting symmetry in mixed-integer nonlinear programming. Master’s thesis, Zuse Institute Berlin, 2019. urn:nbn:de:0297-zib-77055.
- [77] L. Xu, C. D’Ambrosio, L. Liberti, and S. H. Vanier. On cutting planes for extended formulation of signomial programming. Technical report, arXiv, 2022. doi:10.48550/ARXIV.2212.02857.
- [78] J. M. Zamora and I. E. Grossmann. Continuous global optimization of structured process systems models. *Computers and Chemical Engineering*, 22(12):1749–1770, 1998. doi:10.1016/S0098-1354(98)00244-0.

A Test Set

The following table provides details on the test set of 200 instances that was constructed by the selection process described in Section 3.1. For each instance, the number of variables (n), the number of discrete variables ($|\mathcal{I}|$), the number of constraints ($m + \tilde{m}$), the number of nonzeros in the Jacobian and objective function gradient (nz), and the number of nonzeros that correspond to nonlinear terms (nlz) is given.

instance	n	$ \mathcal{I} $	$m + \tilde{m}$	nz	nlz
alan	8	4	7	23	3
autocorr_bern20-05	20	20	0	20	20
autocorr_bern35-04	35	35	0	35	35
ball_mk2_10	10	10	1	20	10
ball_mk2_30	30	30	1	60	30
ball_mk3_10	10	10	1	20	10
batch0812_nc	76	36	205	472	232
batches101006m	278	129	1019	2865	49
batches121208m	406	203	1511	4255	59
bayes2_20	86	0	77	615	440
bayes2_30	86	0	77	618	440
blend029	102	36	213	542	64
blend146	222	87	624	1721	256
camshape100	199	0	200	696	299
cardqp_inlp	50	50	1	100	50
cardqp_iqp	50	50	1	100	50
carton7	328	256	687	3979	678
carton9	360	288	893	4917	758
casctanks	500	40	517	1605	514
cecil_13	840	180	898	2811	360
celar6-sub0	640	640	16	1280	640
chakra	62	0	41	142	41
chem	11	0	4	36	11
chenery	43	0	38	132	56
chimera_k64maxcut-01	1101	1101	0	1101	1101
chimera_mis-01	2032	2032	0	2032	2032
chp_shorttermplan1a	1008	144	2068	6118	576
chp_shorttermplan2a	1584	240	3896	10160	1152
chp_shorttermplan2b	1392	192	2552	7672	1440
clay0204m	52	32	90	284	64
clay0205m	80	50	135	430	80
color_lab3_3x0	316	316	80	632	237
crossdock_15x7	210	210	44	630	210
crossdock_15x8	240	240	46	720	240
crudeoil_lee1_07	749	56	1776	8124	896
crudeoil_pooling_ct2	403	108	732	2523	140
csched1	76	63	22	173	8
csched1a	28	15	22	77	7
cvxnonsep_psig20	20	10	0	20	20
cvxnonsep_psig30	30	15	0	30	30
du-opt	20	13	9	46	20
du-opt5	20	13	9	46	20
edgexcross10-040	90	90	480	1530	90
edgexcross10-080	90	74	480	1528	88
eg_all_s	7	7	27	219	196
eigena2	2500	0	1275	127500	127500
elec50	150	0	50	300	300
elf	54	24	38	177	30
eniplac	141	24	189	510	48
enpro56pb	127	73	191	650	24
ex1244	95	23	129	468	52
ex1252a	24	9	34	93	36
faclay20h	190	190	2280	7030	190
faclay80	3160	3160	164320	496120	3160
feedtray	97	7	91	450	282
fin2bb	588	175	618	9413	42
flay04m	42	24	42	154	4
flay05m	62	40	65	242	5
flay06m	86	60	93	350	6
fo7_ar25_1	112	42	269	1054	14
fo7_ar3_1	112	42	269	1054	14
forest	236	73	309	1013	178
gabriel01	215	72	467	1789	512
gabriel02	261	71	597	2608	1024
gasnet	90	10	69	266	130
gasprod_sarawak16	1526	38	2252	6453	1088
gastrans582_cold13_95	2186	250	3732	8538	2139

instance	n	$ \mathcal{I} $	$m+\tilde{m}$	nz	nlz
gastrans582_mild11	2186	250	3732	8538	2139
gear	4	4	0	4	4
gear2	28	24	4	32	4
gear4	6	4	1	8	4
genpooling_lee1	49	9	82	369	128
genpooling_lee2	53	9	92	453	192
ghg_1veh	29	12	37	130	91
gilbert	1000	0	1	2000	2000
graphpart_2g-0066-0066	108	108	36	216	108
graphpart_clique-60	180	180	60	360	180
gsg_0001	77	0	111	368	44
hadamard_5	25	25	0	25	25
heatexch_spec1	56	12	64	224	32
heatexch_spec2	76	16	90	300	42
hhfair	29	0	25	80	21
himmel16	18	0	21	96	84
house	8	0	8	25	9
hs62	3	0	1	6	6
hvb11	9817	9537	10251	36005	64
hybridynamic_var	81	10	100	286	61
hybridynamic_varcc	151	0	110	388	101
hydroenergy1	288	96	428	1212	120
ibs2	3010	1500	1821	13510	3000
johnall	194	190	192	957	573
kall_circles_c6b	17	0	53	148	86
kall_congruentcircles_c72	17	0	59	160	86
kissing2	772	0	10000	154400	154400
kport20	101	40	27	189	116
kriging-peaks-red020	2	0	0	2	2
kriging-peaks-red100	2	0	0	2	2
lop97icx	986	899	87	1890	704
mathopt5_7	1	0	0	1	1
mathopt5_8	1	0	0	1	1
maxcsp-geo50-20-d4-75-36	1000	1000	50	2000	1000
meanvar-orl400_05_e_7	2000	400	2003	7200	1600
meanvar-orl400_05_e_8	1600	400	1603	6400	800
mhw4d	5	0	3	13	10
milinfract	1000	500	501	502000	1000
minlphi	64	0	79	206	36
multiplants_mtg1a	193	93	256	1972	95
multiplants_mtg2	229	112	306	2689	126
nd_netgen-3000-1-1-b-b-ns_7	15000	3000	12155	48000	9000
netmod_kar1	456	136	666	1848	4
netmod_kar2	456	136	666	1848	4
nous1	50	2	43	196	122
nous2	50	2	43	196	122
nvs02	8	5	3	19	16
nvs06	2	2	0	2	2
oil2	936	2	926	2214	440
optmass	30010	0	25005	80020	10006
ortez	87	18	74	268	54
p_ball_10b_5p_3d_m	95	50	129	518	150
p_ball_15b_5p_2d_m	105	75	139	523	150
parabol5_2_3	40400	0	40200	240004	601
parallel	205	25	115	751	155
pedigree_ex485	485	426	296	1925	485
pedigree_ex485_2	485	426	296	1925	485
pointpack06	12	0	20	86	60
pointpack08	16	0	35	155	112
pooling_epa1	214	30	340	1154	257
pooling_epa2	331	45	524	1913	554
portfol_buyin	17	8	19	58	16
portfol_card	17	8	20	66	16
powerflow0014r	118	0	197	652	461
powerflow0057r	440	0	725	2462	1795
prob07	14	0	35	109	63
process	10	0	7	27	11
procurement1mot	784	60	749	2444	12
procurement2mot	796	60	761	2480	12
product	1553	107	1925	5555	264
product2	2842	128	3125	8249	1056
prolog	20	0	22	128	14
qp3	100	0	52	2747	100
qspp_0_10_0_1_10_1	180	180	100	540	180
qspp_0_11_0_1_10_1	220	220	121	660	220
radar-2000-10-a-6_lat_7	10000	2000	8001	28000	6000
radar-3000-10-a-8_lat_7	15000	3000	12001	42000	9000
ravempb	112	54	186	610	28

instance	n	$ \mathcal{I} $	$m+\tilde{m}$	nz	nlz
risk2bpb	463	14	580	2288	3
routingdelay_bigm	1123	396	2977	7739	1827
rsyn0815m	205	79	347	909	11
rsyn0815m03m	705	282	1647	4120	33
sfacloc2_2.95	186	39	239	595	76
sfacloc2_3.90	291	75	496	1282	135
sjup2	1696	8	17085	151716	88800
slay06m	102	60	135	462	12
slay07m	140	84	189	644	14
smallinvDAXr1b010-011	30	30	3	120	30
smallinvDAXr1b020-022	30	30	3	120	30
sonet17v4	136	136	2057	6527	272
sonet18v6	153	153	2466	7802	306
sonetgr17	152	152	152	694	302
spectra2	69	30	72	408	240
sporttournament24	276	276	0	276	276
sporttournament30	435	435	0	435	435
sssd12-05persp	95	75	52	305	45
sssd18-06persp	150	126	66	474	54
st_testgr1	10	10	5	51	10
st_testgr3	20	20	20	181	20
steenbrf	468	0	108	972	108
stockcycle	480	432	97	1008	48
supplychainp1_022020	2940	460	5300	15040	40
supplychainp1_030510	445	70	835	2330	15
supplychainr1_022020	1440	460	1840	7000	40
supplychainr1_030510	230	70	280	1005	15
syn15m04m	340	120	806	1986	44
syn30m02m	320	120	604	1502	40
synheat	56	12	64	224	28
tanksize	46	9	73	290	63
telecomsp_pacbell	3570	3528	2940	121302	74088
tln5	35	35	30	155	50
tln7	63	63	42	287	98
tls2	37	33	24	209	8
tls4	105	89	64	613	32
topopt-mbb_60x40_50	33600	2400	14363	259956	33600
toroidal2g20_5555	400	400	0	400	400
toroidal3g7_6666	343	343	0	343	343
transswitch0009r	69	9	103	346	255
tricip	169	0	190	1493	1140
tspn08	44	28	18	136	60
tspn15	135	105	34	502	165
unitcommit1	960	720	5329	12404	240
unitcommit2	960	720	5329	12404	480
wager	156	84	142	532	240
waste	2484	400	1991	9242	2736
wastepaper3	52	27	30	177	108
wastepaper4	76	44	38	274	176
wastepaper6	136	90	54	528	360
water4	195	126	137	756	46
waternd1	74	20	83	301	114
waterno2_02	332	18	410	1088	202
waterno2_03	498	27	616	1635	303
waterund01	40	0	38	152	78

B Detailed Computational Results

The following tables show the outcome from running each solver on instances from the test set. If an instance has been solved to optimality, the time spend is reported. Note that due to differences in formulas for the relative gap in the various solvers, an instance may be accounted as solved even though the solver stopped at the time limit. If a run has been flagged as failed, the reason for this decision is given: “abort” if the solver did not return with a result, “nonopt” if the reported upper or lower bound were not consistent with those given by MINLPLib, and “infeas” if the reported solution is not feasible with respect to the feasibility tolerance. Otherwise, the relative gap at termination is reported, which is ∞ if no feasible solution or lower bound has been computed. An exception here is BARON, where an instance is considered as solved if the solver only decided to not return a lower bound due to singularities in functions (see Section 3.2.1). This is the case for instances `mhw4d` and `multiplants_mtg2` and

their permutations. For each instance, a time or gap that is at most 10% worse than the one from the best solver on this instance is printed in bold font.

B.1 Serial Mode

The following table shows the outcome from running each solver on the test set of 200 instances and their permutations in serial mode.

instance	perm	BARON	Lindo API	Ocateract	SCIP
alan	-	0.33	0.89%	0.05	0.19
	1	0.30	0.89%	0.05	0.14
	2	0.19	0.89%	0.13	0.11
	3	0.19	0.89%	0.05	0.11
	4	0.18	0.89%	0.05	0.13
autocorr_bern20-05	-	7.29	100.43	4.36	16.34
	1	5.18	90.50	3.73	13.26
	2	6.19	86.68	4.58	23.15
	3	6.01	70.76	3.30	15.95
	4	5.35	63.71	4.33	19.75
autocorr_bern35-04	-	29.91	13.5%	12.7%	107.46
	1	45.39	infeas	3524.97	93.55
	2	55.21	infeas	2425.34	104.74
	3	79.60	infeas	6572.33	88.83
	4	39.39	infeas	4446.69	77.29
ball_mk2_10	-	0.07	3.98	0.01	0.05
	1	0.07	3.96	0.01	0.10
	2	0.07	3.98	0.01	0.01
	3	0.07	3.66	0.01	0.01
	4	0.07	3.66	0.01	0.01
ball_mk2_30	-	0.25	100.0%	0.02	0.06
	1	0.08	100.0%	0.02	0.08
	2	0.09	100.0%	0.02	0.03
	3	0.08	100.0%	0.02	0.03
	4	0.13	100.0%	0.04	0.03
ball_mk3_10	-	10.58	3.58	0.04	0.00
	1	10.30	3.85	0.04	0.00
	2	10.47	3.88	0.04	0.00
	3	10.15	3.68	0.04	0.00
	4	10.20	3.56	0.04	0.00
batch0812_nc	-	6.57	25.46	8.37	1.45
	1	8.95	22.13	9.06	1.18
	2	5.85	20.34	8.84	1.01
	3	6.41	20.35	8.65	1.45
	4	70.42	23.51	9.21	1.05
batches101006m	-	6.91	134.04	4.84	7.63
	1	18.23	59.63	5.28	6.55
	2	10.39	51.85	5.20	6.44
	3	10.03	81.83	5.13	6.31
	4	8.77	57.22	4.80	6.19
batches121208m	-	47.33	198.76	79.1%	6.99
	1	27.10	221.28	abort	14.64
	2	30.58	180.28	abort	15.32
	3	33.04	184.44	abort	14.49
	4	26.62	152.58	60.6%	12.04
bayes2_20	-	67.77	0.033%	0.033%	0.033%
	1	80.24	0.033%	0.033%	0.033%
	2	69.41	0.033%	0.033%	0.033%
	3	58.54	0.033%	6051.51	0.033%
	4	372.69	0.033%	0.033%	0.033%
bayes2_30	-	21.02	7200.00	7200.00	7200.00
	1	57.89	7200.00	7200.00	7200.00
	2	30.04	7200.00	7200.00	7200.00
	3	20.17	7200.00	7200.00	25.1%
	4	33.53	7200.00	7200.00	7200.00
blend029	-	1.03	abort	29.33	3.67
	1	1.07	abort	30.37	2.79
	2	1.28	abort	36.35	3.29
	3	0.72	2512.99	27.62	3.86
	4	1.15	abort	23.74	2.66
blend146	-	3648.64	abort	8.8%	471.12
	1	4300.65	abort	3.6%	860.35
	2	2597.99	4.9%	8.4%	770.63
	3	2642.15	5.5%	8.8%	380.89
	4	3212.94	abort	7.9%	534.07
camshape100	-	9.2%	9.2%	19.40	5.4%
	1	9.9%	9.2%	110.89	5.2%

instance	perm	BARON	Lindo API	Octeract	SCIP
cardqp_inlp	2	9.5%	9.3%	97.31	5.3%
	3	11.1%	9.5%	111.83	5.2%
	4	9.9%	9.6%	113.55	5.6%
	-	12.40	134.08	792.06	2751.70
	1	12.39	133.88	878.19	3093.49
cardqp_iqp	2	12.35	137.71	747.37	3131.70
	3	12.34	134.87	765.74	3163.05
	4	12.21	136.55	797.50	3091.96
	-	12.92	133.22	785.02	2769.87
	1	12.23	133.62	878.44	3105.68
carton7	2	12.38	138.12	744.48	3145.64
	3	12.26	135.01	765.38	3174.45
	4	12.19	136.78	807.90	3101.12
	-	987.64	110.53	6.03	13.54
	1	1174.14	66.95	5.57	121.71
carton9	2	3846.89	84.70	9.25	13.62
	3	516.36	173.14	12.03	14.55
	4	2021.71	75.91	8.99	12.96
	-	45.0%	287.45	344.94	36.30
	1	37.9%	197.12	9.94	31.5%
casctanks	2	38.6%	375.81	9.28	78.81
	3	43.8%	684.50	19.96	67.78
	4	31.2%	296.68	27.07	187.44
	-	256.46	3.4%	11.8%	199%
	1	333.16	3.5%	10.2%	196%
cecil_13	2	338.11	3.1%	10.5%	199%
	3	357.11	3.5%	10.1%	198%
	4	447.59	3.6%	11.3%	196%
	-	280.62	0.32%	188.40	605.03
	1	227.21	0.6%	157.90	626.61
celar6-sub0	2	207.11	0.32%	242.70	571.07
	3	230.85	0.32%	94.69	648.09
	4	181.85	0.32%	113.29	614.13
	-	100.0%	100%	∞	1503.66
	1	100.0%	100%	3042.60	92.3%
chakra	2	100.0%	100%	4763.43	6435.62
	3	100.0%	100%	6198.12	93.7%
	4	100.0%	100%	2291.77	79.2%
	-	0.14	1.28	7200.00	0.04
	1	0.31	1.27	7200.00	0.05
chem	2	0.15	1.28	7200.00	0.04
	3	0.15	1.01	7200.00	0.11
	4	0.30	1.27	7200.00	0.05
	-	0.04	1172.22	0.18	932.15
	1	0.04	1159.31	0.19	914.43
chenery	2	0.04	1168.01	0.18	844.62
	3	0.04	1169.18	0.18	887.05
	4	0.04	1168.36	0.18	847.04
	-	1.40	0.43%	200.71	2.10
	1	1.11	0.43%	239.04	4.92
chimera_k64maxcut-01	2	1.39	0.78%	259.88	1.26
	3	1.15	0.9%	262.94	7.66
	4	1.77	0.43%	261.43	1.50
	-	560.62	16.4%	168.60	450.42
	1	499.78	17.8%	142.22	435.54
chimera_mis-01	2	1969.20	19.0%	112.38	552.15
	3	605.37	18.7%	106.61	595.36
	4	738.03	14.9%	71.59	431.57
	-	13.69	nonopt	1.14	6.76
	1	16.00	nonopt	1.58	7.06
chp_shorttermplan1a	2	14.94	nonopt	1.46	7.71
	3	14.89	nonopt	1.57	7.48
	4	15.41	nonopt	1.30	7.54
	-	2296.84	∞	64.74	17.55
	1	2708.40	∞	27.60	20.99
chp_shorttermplan2a	2	4998.33	∞	61.12	15.40
	3	0.092%	∞	27.43	19.76
	4	2693.91	∞	62.06	15.63
	-	937.38	∞	17.59	27.16
	1	2.6%	∞	19.04	16.19
chp_shorttermplan2b	2	1181.77	∞	19.44	27.58
	3	5678.85	∞	19.96	26.63
	4	5971.28	∞	19.61	16.91
	-	0.34%	16.6%	0.67%	7200.00
	1	nonopt	0.27%	0.7%	0.08%
	2	0.3%	0.32%	0.67%	0.11%
	3	0.31%	48.9%	0.67%	0.08%
	4	0.3%	0.18%	0.67%	0.067%

instance	perm	BARON	Lindo API	Octeract	SCIP
clay0204m	-	0.47	4.75	1.47	1.67
	1	0.67	6.32	1.24	1.63
	2	0.31	5.77	1.33	1.59
	3	0.45	5.32	1.55	1.96
	4	0.47	4.80	1.32	1.67
clay0205m	-	4.92	51.80	15.10	8.40
	1	4.22	61.13	9.78	8.07
	2	4.75	63.96	8.70	9.07
	3	6.57	77.14	7.83	9.11
	4	5.30	62.76	11.03	6.53
color_lab3_3x0	-	2654.00	65.6%	∞	15.3%
	1	5176.64	65.6%	∞	13.7%
	2	5761.65	65.7%	∞	15.1%
	3	5225.72	65.5%	∞	13.0%
	4	0.51%	65.6%	∞	11.5%
crossdock_15x7	-	385.56	124%	∞	32.0%
	1	418.01	123%	∞	28.1%
	2	233.85	123%	∞	39.1%
	3	180.42	124%	∞	31.1%
	4	183.64	123%	∞	33.5%
crossdock_15x8	-	788.81	124%	6306.65	42.8%
	1	2830.45	121%	∞	59.2%
	2	681.63	124%	∞	52.7%
	3	1084.33	121%	∞	57.8%
	4	1414.43	121%	∞	58.7%
crudeoil_lee1_07	-	7.50	63.26	8.56	4.72
	1	12.30	272.58	10.36	6.96
	2	12.42	nonopt	10.48	7.01
	3	13.79	111.62	4.11	7.31
	4	7.69	103.43	4.54	3.09
crudeoil_pooling_ct2	-	4.95	1.8%	nonopt	18.94
	1	22.28	3.5%	nonopt	26.23
	2	3.24	2.9%	5.85	17.42
	3	5.98	3%	nonopt	17.02
	4	28.56	2.4%	3.84	30.75
csched1	-	13.03	66.81	8.1%	1.59
	1	10.52	91.91	8.1%	2.44
	2	12.10	57.13	8.7%	3.25
	3	11.92	60.21	8.2%	2.29
	4	14.74	77.16	8.1%	3.06
csched1a	-	0.70	4.85	17.77	6.53
	1	0.98	5.76	17.34	5.08
	2	0.50	4.48	2.49	6.85
	3	0.89	4.33	18.52	5.70
	4	0.90	4.69	19.54	5.65
cvxnonsep-psig20	-	0.99	0.48	35.0%	12.04
	1	1.05	0.24	35.1%	14.75
	2	0.96	0.25	34.9%	15.16
	3	0.92	0.50	35.1%	16.56
	4	1.04	0.41	34.9%	10.64
cvxnonsep-psig30	-	4.25	3.26	45.6%	65.40
	1	4.29	3.06	45.7%	367.26
	2	4.25	2.98	45.5%	239.09
	3	4.17	3.22	45.8%	127.14
	4	4.31	3.28	45.5%	88.59
du-opt	-	2.37	5.35	122.38	2.38
	1	2.37	4.83	60.86	2.16
	2	1.86	5.11	52.34	2.68
	3	2.05	4.97	44.64	2.40
	4	2.25	4.79	44.04	4.24
du-opt5	-	3.61	4.70	101.32	1.69
	1	4.48	4.48	130.24	1.50
	2	4.86	4.64	94.97	1.56
	3	3.97	4.88	103.22	1.68
	4	5.33	4.77	132.03	1.31
edgecross10-040	-	7.25	0.29	4.16	8.36
	1	5.42	nonopt	4.57	6.42
	2	5.71	nonopt	4.84	3.65
	3	6.30	nonopt	5.29	5.03
	4	6.34	nonopt	4.82	7.68
edgecross10-080	-	51.58	nonopt	6%	74.51
	1	42.04	nonopt	7.2%	58.81
	2	41.55	nonopt	6.2%	77.79
	3	50.34	nonopt	6.8%	43.25
	4	51.39	nonopt	7.1%	43.84
eg_all_s	-	infeas	abort	102%	3095.77
	1	90.7%	abort	109%	4758.26
	2	90.9%	abort	89.9%	6829.05

instance	perm	BARON	Lindo API	Octeract	SCIP
eigena2	3	85.5%	abort	98.1%	5101.84
	4	89.7%	abort	119%	1520.48
	-	1146.69	∞	416.89	∞
	1	1131.74	∞	101.71	∞
	2	1147.70	∞	117.89	∞
elec50	3	1028.72	∞	288.60	∞
	4	1134.27	∞	269.29	∞
	-	66.5%	798.72	66.4%	45.1%
	1	66.5%	919.20	66.5%	44.9%
	2	66.5%	1073.13	66.4%	44.6%
elf	3	66.5%	1191.25	66.5%	44.8%
	4	66.5%	930.48	66.5%	44.7%
	-	5.91	nonopt	2.54	1.01
	1	6.24	2.09	infeas	1.67
	2	4.51	3.47	infeas	2.42
eniplac	3	4.39	infeas	1.76	1.28
	4	4.05	infeas	infeas	1.74
	-	4.54	3%	1.54	2.19
	1	4.01	2.7%	2.01	2.64
	2	5.11	3.6%	1.67	2.61
enpro56pb	3	6.47	3.9%	1.25	2.58
	4	3.63	3.7%	1.64	1.76
	-	1.84	1454.20	1.98	3.40
	1	2.00	2040.82	2.11	2.27
	2	2.21	infeas	2.08	3.78
ex1244	3	4.95	2838.41	2.08	4.34
	4	5.63	1628.60	2.07	2.44
	-	4.86	9.69	81.99	7.67
	1	1.34	11.66	2.74	64.25
	2	2.62	12.32	nonopt	26.26
ex1252a	3	2.06	11.75	3.98	2.07
	4	2.34	13.07	nonopt	15.20
	-	9.12	infeas	245.41	7200.00
	1	16.38	51.58	334.58	0.031%
	2	2.94	36.26	260.68	0.06%
fac1ay20h	3	1812.77	infeas	313.87	358.50
	4	5.17	55.21	292.16	0.2%
	-	349.36	nonopt	785.28	741.15
	1	251.18	nonopt	721.19	720.40
	2	524.94	nonopt	713.21	934.75
fac1ay80	3	238.74	nonopt	693.80	780.05
	4	478.08	nonopt	1071.27	834.45
	-	120%	5096.67	∞	160%
	1	120%	abort	∞	160%
	2	120%	abort	∞	159%
feedtray	3	120%	abort	∞	93.9%
	4	∞	abort	∞	160%
	-	80.5%	13.79	82.1%	1.70
	1	80.5%	22.18	80.5%	80.5%
	2	80.5%	74.65	83.2%	80.5%
fin2bb	3	80.5%	19.32	80.5%	nonopt
	4	80.5%	19.37	80.5%	80.5%
	-	114.16	1294.03	100.0%	15.92
	1	69.25	1827.21	100.0%	7.78
	2	55.05	1645.40	100.0%	9.54
flay04m	3	108.22	512.39	100.0%	7.80
	4	64.77	2036.16	∞	138.92
	-	10.56	647.41	0.87	3.93
	1	10.74	503.62	0.85	4.29
	2	10.35	583.37	0.91	4.17
flay05m	3	10.15	494.25	0.95	4.35
	4	11.30	552.23	0.78	3.83
	-	394.85	0.54%	infeas	120.82
	1	404.17	0.65%	infeas	123.07
	2	396.72	0.58%	infeas	116.15
flay06m	3	411.54	0.58%	infeas	117.29
	4	439.41	0.56%	infeas	123.48
	-	7%	14.3%	infeas	4931.28
	1	6.6%	14.0%	infeas	4727.15
	2	6.1%	13.9%	infeas	5259.68
fo7_ar25_1	3	7.2%	14.5%	infeas	4811.50
	4	6%	13.7%	3154.62	5047.64
	-	17.00	1282.08	8.95	39.74
	1	9.09	1719.51	7.78	41.70
	2	13.13	1424.43	9.90	53.04
fo7_ar3_1	3	13.99	1979.11	6.24	36.41
	4	10.36	1738.14	9.99	45.67
	-	25.69	1847.41	18.33	88.79

instance	perm	BARON	Lindo API	Octeract	SCIP
forest	1	51.20	2351.00	7.57	63.32
	2	55.38	940.00	10.98	59.84
	3	46.00	1796.83	18.54	65.72
	4	41.03	2416.13	12.23	59.26
	-	747.74	1.18	nonopt	538.83
gabriel01	1	1244.39	∞	nonopt	nonopt
	2	27.0%	infeas	nonopt	nonopt
	3	0.36%	nonopt	nonopt	nonopt
	4	654.45	nonopt	nonopt	nonopt
	-	1760.52	5.4%	2%	351.17
gabriel02	1	1596.81	abort	3.5%	331.53
	2	1676.43	8.6%	4.9%	292.51
	3	1584.57	8.7%	3.7%	245.91
	4	2622.95	10.5%	2%	286.68
	-	11.1%	44.0%	7078.04	1435.31
gasnet	1	16.6%	∞	5618.78	3606.13
	2	17.4%	∞	6015.73	2175.98
	3	19.2%	57.3%	5439.17	973.13
	4	16.8%	∞	5941.14	2639.64
	-	22.38	64.9%	96.7%	42.5%
gasprod_sarawak16	1	nonopt	64.9%	96.5%	41.0%
	2	nonopt	64.9%	96.3%	41.7%
	3	56.9%	65.1%	96.8%	41.1%
	4	nonopt	65.0%	96.5%	41.9%
	-	4578.18	infeas	0.74%	0.39%
gastrans582_cold13_95	1	0.18%	1.5%	1.1%	0.92%
	2	0.18%	0.39%	0.71%	0.94%
	3	0.29%	0.43%	0.88%	0.8%
	4	1258.01	infeas	0.75%	0.92%
	-	1442.44	∞	∞	55.01
gastrans582_mild11	1	828.50	∞	∞	60.20
	2	2066.06	∞	∞	36.26
	3	1801.06	∞	∞	60.01
	4	652.52	∞	∞	∞
	-	687.72	∞	∞	7.41
gear	1	1845.78	∞	∞	13.16
	2	1635.11	∞	∞	20.77
	3	2784.89	∞	∞	10.81
	4	2150.14	∞	∞	10.53
	-	0.12	0.05	0.11	12.55
gear2	1	0.12	0.04	0.07	13.78
	2	0.12	0.03	0.07	13.68
	3	0.08	0.03	0.07	16.84
	4	0.11	0.05	0.07	12.58
	-	0.64	0.23	0.14	18.25
gear4	1	0.28	0.23	0.21	16.14
	2	0.50	0.20	0.09	17.22
	3	0.43	0.62	0.11	17.01
	4	0.31	0.28	0.12	15.55
	-	1.63	infeas	17.56	4.05
genpooling_lee1	1	1.20	infeas	18.07	0.82
	2	1.61	infeas	18.56	0.75
	3	0.88	infeas	17.84	0.80
	4	1.46	infeas	18.12	0.56
	-	9.37	0.41%	117.04	2.09
genpooling_lee2	1	10.02	0.94%	192.20	2.44
	2	4.20	1.1%	39.59	2.40
	3	5.08	1.1%	73.26	3.12
	4	8.35	0.29%	52.15	2.43
	-	85.93	3724.02	186.63	6.88
ghg_1veh	1	25.74	4288.32	254.82	6.48
	2	76.31	1402.41	236.81	5.23
	3	54.53	3502.43	278.49	5.03
	4	54.56	2305.10	281.04	6.53
	-	2.77	106.61	12.42	32.25
gilbert	1	4.52	105.03	12.61	31.41
	2	1.84	104.86	13.26	34.17
	3	6.61	100.45	14.83	32.06
	4	4.62	98.82	12.80	33.13
	-	2.35	14.25	0.9%	1.13
graphpart_2g-0066-0066	1	2.98	15.68	0.99%	1.16
	2	2.85	15.75	1%	1.03
	3	3.49	15.54	0.99%	1.29
	4	3.00	16.13	0.95%	1.24
	-	0.65	8.7%	0.75	1.39
	1	0.93	8.7%	0.69	0.78
	2	1.08	8%	0.42	1.73
	3	0.62	8.2%	0.44	2.06

instance	perm	BARON	Lindo API	Octeract	SCIP
graphpart_clique-60	4	0.83	8.7%	0.45	1.12
	-	34.6%	83.3%	2890.51	3962.88
	1	32.0%	83.1%	7034.14	58.2%
	2	39.7%	80.3%	3824.11	48.7%
	3	36.1%	83.3%	5323.55	55.6%
gsg_0001	4	36.8%	80.3%	1751.71	54.5%
	-	13.68	351.30	8.58	32.26
	1	11.72	412.60	75.14	30.21
	2	10.34	372.91	95.72	28.77
	3	12.57	567.84	62.47	27.70
hadamard_5	4	13.55	379.44	115.60	28.15
	-	54.92	129.27	20.45	21.56
	1	65.13	117.17	44.80	27.78
	2	15.25	121.23	34.82	66.81
	3	17.17	121.44	30.28	23.27
heatexch_spec1	4	10.01	141.44	33.75	22.94
	-	1.39	0.34%	17.7%	7.1%
	1	46.44	0.29%	15.0%	214.24
	2	29.74	0.26%	15.8%	infeas
	3	2001.04	0.46%	14.1%	7.3%
heatexch_spec2	4	0.53%	0.3%	15.4%	infeas
	-	6.60	0.041%	5%	infeas
	1	8.17	0.055%	5.2%	7.08
	2	5.50	0.048%	5%	15.69
	3	7.98	0.043%	5.1%	9.42
hhfair	4	9.97	0.053%	5.2%	8.75
	-	0.43	100.0%	∞	100.0%
	1	0.20	100.0%	∞	100.0%
	2	0.35	100.0%	∞	∞
	3	0.27	100.0%	∞	100.0%
himmel16	4	0.92	100.0%	∞	100.0%
	-	41.46	20.42	12.95	6.13
	1	19.50	22.20	13.13	7.20
	2	29.07	21.96	16.00	5.29
	3	23.68	23.39	13.04	6.44
house	4	22.73	22.24	14.29	6.12
	-	0.37	0.89	108.41	0.33
	1	0.38	1.24	73.21	0.29
	2	0.37	infeas	80.09	0.21
	3	0.46	1.18	115.14	0.58
hs62	4	0.44	0.88	81.46	0.30
	-	1.07	4.04	0.023%	2.65
	1	0.95	4.02	0.021%	3.23
	2	1.07	4.61	0.021%	3.22
	3	1.02	3.79	0.023%	2.29
hvb11	4	0.87	3.91	0.021%	3.00
	-	175.10	40.7%	7.3%	103.26
	1	431.91	45.5%	7.2%	76.87
	2	176.62	29.0%	5.9%	868.06
	3	616.86	47.9%	5.7%	37.39
hybriddynamic_var	4	158.69	36.5%	6.9%	77.85
	-	0.75	4.11	0.32%	2.21
	1	0.83	4.04	3775.94	2.08
	2	1.05	4.04	2838.97	1.87
	3	0.89	4.14	349.49	2.13
hybriddynamic_varcc	4	0.95	4.08	433.50	1.66
	-	0.91	12.00	158.14	2.40
	1	2.08	13.02	730.83	100%
	2	0.97	13.51	7200.00	100%
	3	2.38	13.18	7200.00	100%
hydroenergy1	4	1.20	12.90	7200.00	100%
	-	1026.99	0.93%	0.65%	6295.41
	1	1305.68	0.93%	0.66%	5883.90
	2	3528.61	0.92%	0.65%	5487.34
	3	2844.90	0.92%	0.66%	4605.84
ibs2	4	3324.70	0.93%	0.67%	6541.83
	-	nonopt	abort	5.1%	20.64
	1	nonopt	∞	40.8%	12.25
	2	nonopt	abort	17.4%	8.54
	3	nonopt	abort	14.6%	8.39
johnall	4	nonopt	abort	13.3%	8.14
	-	2.85	22.36	44.23	30.44
	1	3.54	37.05	43.07	31.43
	2	2.30	40.92	44.03	29.52
	3	2.86	40.69	43.02	31.01
kall_circles_c6b	4	2.59	30.65	43.35	30.52
	-	310.13	62.28	463.62	183.83
	1	274.68	55.04	309.04	179.76

instance	perm	BARON	Lindo API	Octeract	SCIP
kall_congruentcircles_c72	2	335.78	72.49	417.47	170.66
	3	323.33	52.58	374.48	284.95
	4	309.15	58.86	459.34	195.62
	-	26.45	4.74	49.42	22.40
	1	20.00	10.50	41.78	29.79
kissing2	2	24.93	4.45	38.33	30.96
	3	20.36	10.36	44.57	28.19
	4	14.04	5.11	40.23	26.06
	-	184.78	867.71	100.0%	100.0%
	1	286.21	911.33	100.0%	100.0%
kport20	2	249.09	927.99	100.0%	∞
	3	278.40	916.33	100.0%	1623.09
	4	212.31	909.74	100.0%	397.82
	-	6.1%	13.1%	3596.06	1001.03
	1	1387.42	10.4%	4236.68	4384.07
kriging_peaks-red020	2	3.5%	12.0%	2599.71	2587.04
	3	3294.62	11.9%	2950.96	1982.41
	4	6.3%	11.4%	1998.33	1456.99
	-	10.33	21.79	85.61	117.65
	1	10.52	21.32	90.28	117.03
kriging_peaks-red100	2	8.61	21.31	90.38	116.84
	3	10.46	21.11	90.31	117.11
	4	10.84	21.33	89.54	117.06
	-	196.08	140.14	629.86	7200.00
	1	131.90	142.45	616.46	7200.00
lop97icx	2	133.05	140.39	616.91	7200.00
	3	132.12	138.68	612.72	7200.00
	4	133.57	142.08	613.09	7200.00
	-	1709.08	39.83	5.90	25.38
	1	1963.46	92.39	8.45	29.05
mathopt5_7	2	2294.75	74.88	6.74	30.77
	3	1893.25	68.66	8.80	45.83
	4	1623.45	45.81	8.93	36.10
	-	0.29	12.35	0.08	0.17
	1	0.12	12.49	0.21	0.17
mathopt5_8	2	0.12	12.35	0.08	0.19
	3	0.12	12.22	0.08	0.33
	4	0.22	12.27	0.08	0.18
	-	0.26	9.72	0.06	0.14
	1	0.26	10.66	0.07	0.22
maxcsp-geo50-20-d4-75-36	2	0.26	10.83	0.07	0.14
	3	0.26	10.71	0.07	0.26
	4	0.25	10.79	0.07	0.13
	-	16.94	101%	7.71	52.80
	1	20.06	101%	5.53	335.96
meanvar-orl400_05_e_7	2	27.23	101%	7.01	49.99
	3	8.13	103%	25.70	296.74
	4	22.97	102%	25.64	186.01
	-	95.4%	23.95	∞	4080.76
	1	94.3%	31.77	6630.62	5981.88
meanvar-orl400_05_e_8	2	94.2%	25.01	6570.19	0.59%
	3	95.0%	nonopt	∞	5120.84
	4	94.5%	nonopt	6663.61	5806.16
	-	600.24	19.32	6.00	2404.35
	1	571.98	∞	6.89	2126.66
mhw4d	2	549.46	nonopt	6.24	2366.11
	3	645.73	∞	6.61	2115.85
	4	407.85	∞	6.32	2353.26
	-	0.83	0.78	0.43	0.44
	1	0.82	0.62	1.64	0.44
milinfract	2	0.69	0.51	1.26	0.40
	3	0.82	0.78	1.61	0.74
	4	0.74	0.65	1.31	0.37
	-	55.06	infeas	75.5%	68.8%
	1	55.42	infeas	75.4%	75.2%
minlphi	2	55.04	infeas	75.5%	72.0%
	3	54.76	infeas	76.0%	75.3%
	4	54.45	infeas	75.7%	70.6%
	-	∞	1.86	100%	100.0%
	1	∞	1.65	100%	100.0%
multiplants_mtg1a	2	∞	1.74	100%	0.20
	3	∞	1.65	100%	0.07
	4	∞	1.24	100%	0.16
	-	2245.95	15.6%	5.1%	4320.86
	1	1669.44	17.9%	149.26	18.8%
	2	212.57	15.1%	288.46	3067.67
	3	2199.49	17.9%	98.00	4467.25
	4	1918.40	17.3%	2301.73	3773.36

instance	perm	BARON	Lindo API	Octeract	SCIP
multiplants_mtg2	-	2115.70	0.12%	2%	21.0%
	1	717.45	0.15%	1.8%	0.4%
	2	276.56	0.13%	469.69	0.27%
	3	132.47	0.041%	1%	6.6%
	4	212.58	2.2%	1.9%	1%
nd_netgen-3000-1-1-b-b-ns_7	-	∞	680.80	3.67	31.10
	1	99.7%	infeas	infeas	32.20
	2	∞	infeas	infeas	47.17
	3	99.7%	∞	infeas	40.17
	4	99.7%	infeas	infeas	47.92
netmod_kar1	-	4.53	5948.45	41.89	5.51
	1	11.3%	110.53	51.27	4.90
	2	11.1%	311.30	39.29	4.88
	3	4420.66	576.52	42.64	5.64
	4	15.4%	1050.14	40.84	5.11
netmod_kar2	-	5.26	5917.59	42.01	5.77
	1	11.3%	110.81	51.55	5.48
	2	11.1%	311.71	39.14	4.90
	3	4429.25	576.74	43.00	5.64
	4	15.4%	1051.59	40.95	5.12
nous1	-	51.20	2.9%	52.93	7.70
	1	39.45	1.7%	48.64	5.19
	2	38.92	0.88%	62.63	2.62
	3	36.38	1.5%	37.63	5.75
	4	15.51	1.5%	42.83	4.75
nous2	-	0.40	2.16	8.29	0.80
	1	0.43	infeas	10.50	0.91
	2	0.77	2.21	8.61	0.65
	3	0.51	1.76	6.92	0.87
	4	0.59	2.25	5.81	0.69
nvs02	-	0.04	69.68	0.22	0.06
	1	0.05	63.08	0.11	0.04
	2	0.12	54.34	0.28	0.03
	3	0.09	39.48	0.11	0.03
	4	0.05	46.77	0.15	0.04
nvs06	-	0.08	6.64	0.06	0.02
	1	0.08	3.90	0.13	0.03
	2	0.09	3.79	0.06	0.03
	3	0.08	3.83	0.06	0.03
	4	0.08	3.84	0.06	0.03
oil2	-	3.30	43.56	nonopt	4.64
	1	nonopt	infeas	nonopt	4.55
	2	3.55	infeas	nonopt	8.76
	3	5.05	22.24	nonopt	7.57
	4	3.64	39.41	nonopt	6.34
optmass	-	∞	11.3%	463.14	1.8%
	1	∞	∞	480.52	9.9%
	2	∞	∞	493.97	7.6%
	3	∞	12.8%	502.07	9.9%
	4	∞	12.8%	507.91	9.9%
ortez	-	nonopt	2.42	9.48	0.14
	1	nonopt	3.56	18.07	0.12
	2	nonopt	1.50	0.98	0.14
	3	nonopt	3.47	1.26	0.26
	4	0.25	2.95	7.15	0.15
p_ball_10b_5p_3d_m	-	32.22	26.41	49.66	4.26
	1	23.79	30.34	52.39	3.94
	2	26.62	35.67	infeas	4.35
	3	30.37	29.09	70.38	4.35
	4	31.69	33.60	49.71	4.43
p_ball_15b_5p_2d_m	-	90.52	101.62	infeas	5.44
	1	93.01	110.13	infeas	5.30
	2	59.07	83.88	infeas	5.01
	3	70.12	103.04	infeas	5.71
	4	73.50	153.02	infeas	4.21
parabol5_2_3	-	0.051%	15.7%	4095.70	0.051%
	1	7200.00	infeas	6.6%	7200.00
	2	0.051%	16.1%	6.6%	7200.00
	3	6.7%	16.4%	406.60	7200.00
	4	7200.00	infeas	6.6%	7200.00
parallel	-	19.19	38.93	102.24	10.26
	1	17.35	37.95	295.19	8.15
	2	19.55	36.08	291.45	8.78
	3	22.27	40.15	299.21	8.40
	4	20.80	37.29	295.03	7.73
pedigree_ex485	-	68.37	5.8%	276.45	105.90
	1	85.93	6.5%	248.15	104.00
	2	535.45	9.1%	161.40	56.53

instance	perm	BARON	Lindo API	Octeract	SCIP
pedigree_ex485_2	3	75.96	5.7%	243.80	77.55
	4	54.46	4.8%	240.39	79.44
	-	17.30	192.97	3.66	35.24
	1	14.91	221.77	1.86	53.86
	2	17.00	245.18	1.86	53.18
pointpack06	3	15.79	238.26	1.77	40.26
	4	16.50	231.25	1.83	39.39
	-	5.59	27.80	4.76	3.05
	1	5.68	30.47	5.09	2.99
	2	6.00	31.82	4.68	3.54
pointpack08	3	5.71	47.47	5.55	3.41
	4	6.45	30.25	5.18	3.12
	-	159.31	3611.87	208.18	61.22
	1	172.19	3492.98	205.47	67.88
	2	194.38	4171.95	201.51	64.73
pooling_epa1	3	210.56	2597.64	234.46	60.70
	4	170.46	3025.00	178.05	63.89
	-	12.84	infeas	21.07	12.87
	1	10.13	infeas	17.86	1054.75
	2	14.46	infeas	20.75	19.10
pooling_epa2	3	7.62	infeas	18.78	21.67
	4	10.55	2.4%	20.88	35.76
	-	2.4%	0.041%	1694.37	1.9%
	1	1030.44	1.4%	2186.70	1.8%
	2	1.8%	0.088%	1797.05	1.8%
portfol_buyin	3	3.2%	6096.04	1992.47	1.7%
	4	4.3%	0.39%	1615.89	2%
	-	0.34	2.34	3.56	0.50
	1	0.56	2.27	0.85	0.35
	2	0.30	2.25	1.72	0.29
portfol_card	3	0.22	2.28	3.42	0.41
	4	0.31	2.26	3.49	0.57
	-	0.31	2.15	5.29	0.58
	1	0.22	2.34	1.93	0.95
	2	0.43	2.33	4.81	0.45
powerflow0014r	3	0.36	2.09	7.60	0.33
	4	0.31	2.33	7.36	0.75
	-	50.50	100.0%	100.0%	7079.42
	1	31.00	100.0%	66.6%	2519.01
	2	47.00	100.0%	100.0%	0.025%
powerflow0057r	3	59.10	100.0%	100.0%	6935.50
	4	50.99	100.0%	88.1%	7200.00
	-	∞	∞	∞	80.19
	1	∞	∞	∞	102.21
	2	∞	∞	∞	78.25
prob07	3	∞	∞	∞	78.13
	4	∞	∞	∞	68.48
	-	75.62	68.41	16.19	76.05
	1	74.75	67.34	15.70	10.94
	2	73.53	69.74	15.49	11.11
process	3	80.90	69.79	16.44	11.47
	4	78.75	71.26	16.69	10.54
	-	0.51	1.38	0.93	0.4%
	1	0.39	1.53	0.64	0.46
	2	0.66	1.44	0.95	0.39%
procurement1mot	3	0.41	1.39	0.89	0.39%
	4	0.42	1.36	0.58	0.4%
	-	59.57	86.4%	79.5%	85.0%
	1	114.45	87.8%	79.7%	84.2%
	2	177.11	86.6%	79.7%	80.6%
procurement2mot	3	144.06	87.4%	79.7%	83.9%
	4	148.48	87.6%	79.7%	84.5%
	-	4.13	13.59	2.71	2.03
	1	8.12	12.54	2.54	2.39
	2	5.86	12.99	2.80	1.76
product	3	7.10	15.11	3.40	2.37
	4	7.46	19.46	2.67	3.02
	-	48.09	465.36	nonopt	16.80
	1	nonopt	471.35	nonopt	11.71
	2	nonopt	2467.73	nonopt	13.45
product2	3	nonopt	278.34	nonopt	18.04
	4	107.32	446.81	nonopt	21.42
	-	2.90	145.38	2.53	3.85
	1	79.81	185.85	nonopt	infeas
	2	3.41	184.60	2.92	infeas
prolog	3	26.03	467.30	2.62	infeas
	4	242.41	133.79	nonopt	infeas
	-	0.31	100.0%	100.0%	0.06

instance	perm	BARON	Lindo API	Octeract	SCIP
qp3	1	1.16	100.0%	100.0%	abort
	2	0.30	100.0%	100.0%	0.08
	3	1.22	100.0%	100.0%	0.05
	4	0.36	100%	100.0%	0.06
	-	13.7%	0.01	4.6%	14.1%
	1	12.2%	0.01	4.6%	6%
	2	11.9%	0.01	4.6%	6%
	3	12.7%	0.01	4.6%	13.8%
qspp_0_10_0_1_10_1	4	13.2%	0.01	4.6%	13.4%
	-	24.2%	632.32	26.39	180.85
	1	22.5%	653.67	25.59	190.34
	2	21.3%	643.24	23.37	174.14
	3	22.5%	661.55	23.97	207.82
qspp_0_11_0_1_10_1	4	23.7%	650.22	22.88	158.80
	-	30.5%	3676.78	124.07	623.76
	1	31.5%	3955.95	138.43	710.85
	2	29.0%	3821.06	143.49	770.51
	3	31.4%	3870.07	131.00	679.01
radar-2000-10-a-6_lat_7	4	31.1%	3857.65	116.44	756.40
	-	75.1%	33.14	33.34	129.82
	1	96.8%	64.95	39.42	140.45
	2	95.9%	65.73	39.22	142.02
	3	6060.60	65.56	39.32	141.78
radar-3000-10-a-8_lat_7	4	96.8%	64.08	10.04	139.60
	-	100.0%	558.69	209.09	422.76
	1	100.0%	732.73	14.72	291.33
	2	100.0%	724.04	14.73	769.42
	3	99.9%	754.45	13.79	287.51
ravempb	4	100.0%	730.86	14.73	307.35
	-	0.82	3438.61	1.26	3.09
	1	0.98	560.99	1.09	2.22
	2	0.73	infeas	0.83	1.54
	3	0.91	182.23	0.80	2.26
risk2bpb	4	0.85	990.98	1.06	2.70
	-	0.70	1.46	10.70	infeas
	1	0.93	1.72	2.49	infeas
	2	0.92	1.49	47.77	infeas
	3	0.49	1.51	2.72	infeas
routingdelay_bigm	4	0.46	0.99	2.39	infeas
	-	26.08	nonopt	18.4%	nonopt
	1	14.94	2.5%	54.3%	nonopt
	2	16.96	nonopt	20.3%	3223.45
	3	25.96	2.9%	18.9%	0.97%
rsyn0815m	4	24.02	nonopt	23.3%	nonopt
	-	0.29	92.08	1.02	1.44
	1	0.29	80.27	1.02	1.45
	2	0.35	62.63	1.00	1.88
	3	0.34	72.91	0.82	1.57
rsyn0815m03m	4	0.34	112.79	1.01	1.61
	-	31.23	nonopt	44.8%	14.06
	1	47.45	nonopt	abort	20.78
	2	40.66	abort	33.6%	14.84
	3	44.05	nonopt	abort	14.36
sfacloc2_2_95	4	38.36	nonopt	abort	17.09
	-	0.31	5.25	1.11	2.78
	1	0.74	77.82	0.58	2.18
	2	0.30	69.08	1.02	2.19
	3	0.38	129.88	0.64	2.05
sfacloc2_3_90	4	0.39	22.24	0.57	2.36
	-	29.12	588.76	67.86	20.32
	1	28.72	3934.46	39.19	17.48
	2	27.15	11.7%	47.61	27.05
	3	25.26	1353.20	28.54	25.05
sjup2	4	16.92	39.7%	94.54	27.50
	-	∞	3456.91	∞	469.64
	1	3049.75	2931.24	∞	1.6%
	2	∞	3562.09	∞	254.22
	3	∞	3131.78	100.0%	308.94
slay06m	4	∞	2986.81	∞	345.25
	-	0.46	5.66	0.69	1.85
	1	0.81	5.20	0.94	2.42
	2	0.65	5.53	0.91	0.94
	3	0.64	6.30	0.93	1.25
slay07m	4	0.78	5.21	0.63	1.57
	-	0.62	6.96	0.61	2.05
	1	0.81	10.48	1.12	2.01
	2	0.80	8.22	1.14	2.01
	3	0.73	11.43	0.81	3.39

instance	perm	BARON	Lindo API	Octeract	SCIP
smallinvDAXr1b010-011	4	0.85	11.45	1.26	3.28
	-	1.62	3.51	8.69	2.17
	1	1.72	3.08	7.52	2.31
	2	1.75	3.57	7.69	2.29
	3	2.79	3.05	6.83	2.74
smallinvDAXr1b020-022	4	1.65	3.60	9.48	2.19
	-	1.95	3.46	205.80	3.37
	1	2.43	3.95	209.56	2.80
	2	1.94	3.56	185.73	2.92
	3	2.26	4.04	192.10	0.75
sonet17v4	4	2.22	3.57	158.16	0.75
	-	339.62	nonopt	160.55	949.25
	1	414.29	nonopt	188.66	1453.04
	2	324.27	nonopt	143.09	926.01
	3	238.47	nonopt	138.66	916.91
sonet18v6	4	358.04	nonopt	186.01	927.66
	-	429.30	nonopt	184.16	1445.90
	1	439.64	nonopt	154.31	1353.80
	2	421.41	nonopt	171.20	1426.08
	3	557.19	nonopt	180.39	1211.53
sonetgr17	4	624.38	nonopt	206.98	1208.09
	-	346.77	nonopt	65.24	1255.92
	1	327.68	29.5%	129.02	1229.96
	2	370.11	nonopt	81.73	1228.20
	3	424.67	5378.39	36.48	1043.24
spectra2	4	324.24	4886.49	30.05	1290.93
	-	1.37	infeas	105%	9.20
	1	1.35	infeas	104%	8.46
	2	1.59	infeas	105%	12.31
	3	1.33	5.75	105%	11.51
sporttournament24	4	1.33	infeas	104%	11.05
	-	4982.31	6.7%	21.48	45.75
	1	1.3%	8%	27.89	59.26
	2	2.5%	7.9%	15.68	81.50
	3	1.3%	4.1%	20.86	81.00
sporttournament30	4	4451.83	9.8%	21.45	121.43
	-	11.8%	12.5%	1946.51	0.72%
	1	11.7%	14.3%	2934.20	0.84%
	2	12.5%	13.3%	2795.89	0.55%
	3	12.0%	13.6%	2178.10	6530.54
sssd12-05persp	4	11.8%	15.3%	4553.97	7010.55
	-	27.8%	378.13	780.87	4.5%
	1	28.1%	abort	572.29	4.6%
	2	27.7%	∞	691.17	4.8%
	3	28.4%	nonopt	695.56	4.7%
sssd18-06persp	4	28.5%	abort	100.0%	5.1%
	-	41.4%	7200.00	19.9%	15.1%
	1	41.4%	nonopt	21.4%	15.9%
	2	41.4%	abort	25.6%	16.6%
	3	41.3%	abort	23.2%	15.9%
st_testgr1	4	41.6%	abort	100.0%	16.5%
	-	0.24	5.51	0.07	0.12
	1	0.34	4.73	0.06	0.24
	2	0.37	4.93	0.07	0.13
	3	0.13	4.50	0.06	0.12
st_testgr3	4	0.12	4.45	0.07	0.13
	-	0.24	3.67	0.07	0.24
	1	0.23	3.66	0.07	0.17
	2	0.23	3.52	0.06	0.17
	3	0.24	3.28	0.07	0.17
steenbrf	4	0.24	3.33	0.07	0.17
	-	14.39	0.75	98.0%	12.42
	1	7.59	1.87	97.9%	0.038%
	2	9.81	1.21	97.9%	0.33%
	3	8.40	0.99	98.0%	13.85
stockcycle	4	4.69	9.33	97.8%	0.34%
	-	7.82	3.18	39.57	0.79
	1	6.15	1.36	35.36	0.56
	2	5.74	1.85	35.21	0.96
	3	6.14	1.57	36.67	0.79
supplychainp1_022020	4	7.59	1.63	36.22	0.71
	-	nonopt	20.8%	8.5%	1056.69
	1	nonopt	33.7%	11.5%	abort
	2	nonopt	25.5%	11.6%	0.098%
	3	nonopt	36.8%	11.4%	0.077%
supplychainp1_030510	4	nonopt	18.6%	11.1%	4696.71
	-	nonopt	63.91	3.08	3.54
	1	nonopt	92.12	3.56	5.41

instance	perm	BARON	Lindo API	Octeract	SCIP
supplychainr1_022020	2	nonopt	114.12	3.52	2.84
	3	nonopt	85.05	3.39	4.46
	4	1.48	63.43	3.35	3.68
	-	6.13	12.5%	2966.92	21.97
	1	11.57	15.8%	0.88%	27.18
supplychainr1_030510	2	19.28	15.9%	1.1%	1.64
	3	6.66	15.5%	1.2%	2.39
	4	10.05	15.2%	1.9%	26.06
	-	0.28	7.19	0.74	0.12
	1	0.40	7.87	0.73	0.28
syn15m04m	2	0.41	9.32	0.81	0.20
	3	0.38	9.23	0.78	0.16
	4	0.42	6.65	0.79	0.16
	-	0.54	42.37	1.04	2.75
	1	0.67	43.23	1.44	2.05
syn30m02m	2	0.58	41.87	1.48	1.13
	3	0.49	44.00	1.13	0.94
	4	0.50	43.84	1.40	1.86
	-	2.51	239.18	1.83	1.58
	1	2.14	239.12	2.28	1.39
synheat	2	1.92	249.03	1.88	2.18
	3	2.27	187.53	2.53	2.27
	4	2.71	226.18	2.04	1.62
	-	35.92	0.36%	17.1%	infeas
	1	734.68	0.33%	16.6%	3.9%
tanksize	2	5112.71	0.34%	15.3%	infeas
	3	33.77	0.39%	14.7%	7.2%
	4	26.63	0.34%	15.9%	infeas
	-	4.81	13.85	54.61	4.12
	1	5.53	14.77	59.53	3.22
telecomsp_pacbell	2	5.07	14.41	57.79	3.13
	3	4.37	16.30	57.19	3.40
	4	5.57	14.46	58.71	3.34
	-	3393.96	2.8%	1125.10	0.63%
	1	2926.69	8.2%	2376.86	0.81%
tln5	2	3854.14	5.4%	2301.23	0.96%
	3	0.52%	5%	1194.43	0.78%
	4	2094.65	6.2%	1483.59	0.58%
	-	5.26	30.1%	2.08	0.50
	1	8.30	30.1%	1.20	0.78
tln7	2	5.75	31.1%	1.91	0.69
	3	3.67	30.1%	1.74	0.75
	4	2.76	29.1%	1.63	0.44
	-	2935.35	61.4%	274.33	188.05
	1	1666.87	70.9%	357.03	122.53
tls2	2	1.3%	65.1%	301.36	302.03
	3	1622.52	69.1%	328.32	970.70
	4	4263.19	66.5%	319.16	338.06
	-	0.13	5.43	0.18	0.93
	1	0.14	8.16	0.14	1.08
tls4	2	0.14	7.82	0.14	0.59
	3	0.13	4.81	0.15	0.84
	4	0.14	7.82	0.14	0.79
	-	35.15	833.24	36.82	41.19
	1	30.56	870.73	34.63	44.52
topopt-mbb_60x40_50	2	42.93	abort	31.00	41.94
	3	26.01	abort	29.09	40.57
	4	38.94	865.53	32.69	47.23
	-	∞	195.38	∞	∞
	1	∞	nonopt	∞	∞
toroidal2g20_5555	2	∞	nonopt	∞	∞
	3	∞	nonopt	∞	∞
	4	∞	481.77	∞	∞
	-	3747.17	nonopt	3.23	5.70
	1	833.26	nonopt	3.58	6.35
toroidal3g7_6666	2	2465.17	nonopt	2.57	5.91
	3	914.43	nonopt	3.40	4.66
	4	1190.57	nonopt	3.07	5.89
	-	6.2%	nonopt	67.16	141.64
	1	4.6%	nonopt	74.53	174.00
transswitch0009r	2	5.7%	nonopt	86.25	148.01
	3	6%	nonopt	122.43	122.20
	4	5.4%	nonopt	88.10	89.17
	-	8%	16.0%	5.6%	6975.38
	1	6.5%	18.1%	5%	5869.65
	2	7.7%	23.3%	4%	6188.29
	3	11.9%	19.6%	5.6%	0.43%
	4	7.4%	21.9%	3.8%	6364.97

instance	perm	BARON	Lindo API	Octeract	SCIP
trcp	-	0.82	infeas	100.0%	100.0%
	1	7.10	infeas	0.02%	100.0%
	2	182.36	infeas	100.0%	100.0%
	3	167.41	infeas	2.02	100.0%
	4	405.43	infeas	100.0%	100.0%
tspn08	-	10.81	10.9%	2.3%	17.8%
	1	13.75	10.9%	2.3%	17.4%
	2	10.93	10.6%	2.3%	19.1%
	3	11.68	9.5%	2.4%	19.8%
	4	8.78	10.8%	2.3%	19.5%
tspn15	-	884.44	12.8%	81.4%	50.0%
	1	967.84	12.8%	79.9%	52.0%
	2	1235.57	12.6%	82.7%	52.7%
	3	1129.22	12.8%	79.7%	49.7%
	4	1088.07	12.8%	78.4%	52.7%
unitcommit1	-	0.27%	81.17	1082.69	infeas
	1	1.4%	65.40	1838.25	infeas
	2	1.5%	85.31	1538.88	infeas
	3	1.3%	60.10	1517.16	infeas
	4	1.5%	69.59	1419.76	infeas
unitcommit2	-	7200.00	18.0%	11.20	439.70
	1	7200.00	19.6%	8.42	23.15
	2	7200.00	19.3%	9.70	10.83
	3	7200.00	18.2%	8.15	infeas
	4	7200.00	21.1%	7.83	infeas
wager	-	15.30	∞	252.73	3.46
	1	17.16	28.9%	482.07	53.79
	2	8.88	28.6%	1694.38	1.53
	3	15.85	∞	324.73	9.28
	4	14.15	37.2%	644.93	11.81
waste	-	22.55	46.3%	99.9%	38.85
	1	34.81	46.2%	100.0%	34.62
	2	26.85	36.9%	100.0%	infeas
	3	38.54	47.0%	100.0%	44.26
	4	24.72	46.6%	100.0%	41.89
wastepaper3	-	9.19	33.90	22.76	4.22
	1	11.99	33.19	27.77	4.01
	2	7.51	32.25	24.13	3.55
	3	8.27	39.31	26.56	5.57
	4	11.26	31.73	23.90	2.41
wastepaper4	-	871.50	2320.60	2099.99	167.16
	1	328.30	1712.61	1708.66	144.72
	2	365.29	2919.38	1800.71	113.53
	3	229.96	2723.08	1593.40	117.95
	4	349.52	2223.22	1838.80	213.66
wastepaper6	-	0.022%	abort	0.16%	0.023%
	1	7200.00	abort	0.22%	0.044%
	2	7200.00	7200.00	0.13%	0.044%
	3	7200.00	0.022%	0.2%	0.051%
	4	7200.00	7200.00	0.13%	0.058%
water4	-	44.4%	46.1%	1648.53	nonopt
	1	42.9%	51.9%	1946.91	nonopt
	2	47.4%	51.1%	2214.42	nonopt
	3	45.8%	51.2%	2181.62	nonopt
	4	44.1%	53.2%	2084.97	nonopt
waternd1	-	5.50	1130.72	1.34	7.21
	1	5.06	1101.35	1.30	4.58
	2	7.94	1150.81	1.36	5.74
	3	22.42	1046.61	1.40	10.78
	4	10.55	1181.39	1.61	7.26
waterno2_02	-	5.20	26.45	4.18	2.89
	1	3.89	60.16	6.69	4.15
	2	5.13	28.66	6.66	3.72
	3	2.36	28.36	6.69	3.91
	4	4.53	24.33	4.40	3.38
waterno2_03	-	260.33	1886.92	440.65	nonopt
	1	254.85	1847.35	479.46	nonopt
	2	270.87	1970.97	494.92	43.44
	3	313.74	2146.47	681.78	nonopt
	4	356.97	2060.53	550.59	38.73
waterund01	-	0.35%	infeas	0.18%	1525.91
	1	0.34%	0.84%	0.18%	1988.24
	2	0.34%	infeas	0.18%	3003.09
	3	0.35%	0.98%	0.18%	0.025%
	4	0.35%	infeas	0.18%	16.83

B.2 Octeract Gap Limit

The following table shows the outcome from running Octeract in serial mode with both gap limits set to either 10^{-6} or 10^{-4} on the test set of 200 non-permuted instances.

instance	10^{-6}	10^{-4}
alan	0.05	0.05
autocorr_bern20-05	4.36	4.39
autocorr_bern35-04	12.7%	12.7%
ball_mk2_10	0.01	0.01
ball_mk2_30	0.02	0.02
ball_mk3_10	0.04	0.04
batch0812.nc	8.37	8.63
batches101006m	4.84	4.92
batches121208m	79.1%	79.1%
bayes2_20	0.033%	0.033%
bayes2_30	7200.00	7200.00
blend029	29.33	30.14
blend146	8.8%	8.8%
camshape100	19.40	19.42
cardqp_inlp	792.06	788.89
cardqp_iqp	785.02	793.77
carton7	6.03	6.34
carton9	344.94	346.42
casctanks	11.8%	12.4%
cecil_13	188.40	205.63
celar6-sub0	∞	∞
chakra	7200.00	7200.00
chem	0.18	0.18
chenery	200.71	200.09
chimera_k64maxcut-01	168.60	169.03
chimera_mis-01	1.14	1.45
chp_shorttermplan1a	64.74	65.18
chp_shorttermplan2a	17.59	17.69
chp_shorttermplan2b	0.67%	0.69%
clay0204m	1.47	1.45
clay0205m	15.10	14.97
color_lab3_3x0	∞	∞
crossdock_15x7	∞	∞
crossdock_15x8	6306.65	6336.71
crudeoil_lee1_07	8.56	7.58
crudeoil_pooling_ct2	nonopt	nonopt
csched1	8.1%	8.1%
csched1a	17.77	17.43
cvxnonsep_psig20	35.0%	35.0%
cvxnonsep_psig30	45.6%	45.6%
du-opt	122.38	122.77
du-opt5	101.32	102.03
edgexcross10-040	4.16	4.11
edgexcross10-080	6%	5.6%
eg_all_s	102%	102%
eigena2	416.89	523.38
elec50	66.4%	66.4%
elf	2.54	2.54
eniplac	1.54	1.30
enpro56pb	1.98	2.31
ex1244	81.99	82.60
ex1252a	245.41	248.28
faclay20h	785.28	781.38
faclay80	∞	∞
feedtray	82.1%	82.1%
fin2bb	100.0%	100.0%
flay04m	0.87	0.90
flay05m	infeas	infeas
flay06m	infeas	infeas
fo7_ar25_1	8.95	8.99
fo7_ar3_1	18.33	18.57
forest	nonopt	nonopt
gabriel01	2%	4.3%
gabriel02	7078.04	6811.48
gasnet	96.7%	96.7%
gasprod_sarawak16	0.74%	0.68%
gastrans582_cold13_95	∞	∞
gastrans582_mild11	∞	∞
gear	0.11	0.07
gear2	0.14	0.13
gear4	17.56	17.33
genpooling_lee1	117.04	117.72

instance	10^{-6}	10^{-4}
genpooling_lee2	186.63	186.60
ghg_1veh	12.42	12.36
gilbert	0.9%	0.9%
graphpart_2g-0066-0066	0.75	0.43
graphpart_clique-60	2890.51	2884.75
gsg_0001	8.58	8.56
hadamard_5	20.45	20.23
heatexch_spec1	17.7%	17.7%
heatexch_spec2	5%	5%
hhfair	∞	∞
himmel16	12.95	13.27
house	108.41	108.71
hs62	0.023%	0.023%
hvb11	7.3%	7.3%
hybridynamic_var	0.32%	0.32%
hybridynamic_varcc	158.14	158.12
hydroenergy1	0.65%	0.65%
ibs2	5.1%	5.1%
johnall	44.23	43.95
kall_circles_c6b	463.62	466.28
kall_congruentcircles_c72	49.42	49.78
kissing2	100.0%	100.0%
kport20	3596.06	3593.09
kriging_peaks-red020	85.61	85.98
kriging_peaks-red100	629.86	630.04
lop97icx	5.90	5.90
mathopt5_7	0.08	0.08
mathopt5_8	0.06	0.07
maxcsp-geo50-20-d4-75-36	7.71	7.78
meanvar-orl400_05_e_7	∞	∞
meanvar-orl400_05_e_8	6.00	5.98
mhw4d	0.43	0.62
milinfract	75.5%	75.5%
minlphi	100%	100%
multiplants_mtg1a	5.1%	5.1%
multiplants_mtg2	2%	2%
nd_netgen-3000-1-1-b-b-ns_7	3.67	3.59
netmod_kar1	41.89	41.79
netmod_kar2	42.01	42.04
nous1	52.93	52.90
nous2	8.29	8.04
nvs02	0.22	0.15
nvs06	0.06	0.06
oil2	nonopt	nonopt
optmass	463.14	463.57
ortez	9.48	9.48
p_ball_10b_5p_3d_m	49.66	49.32
p_ball_15b_5p_2d_m	infeas	infeas
parabol5_2_3	4095.70	4110.54
parallel	102.24	102.42
pedigree_ex485	276.45	277.27
pedigree_ex485_2	3.66	3.94
pointpack06	4.76	4.77
pointpack08	208.18	209.02
pooling_epa1	21.07	21.10
pooling_epa2	1694.37	2039.75
portfol_buyin	3.56	3.37
portfol_card	5.29	5.26
powerflow0014r	100.0%	100.0%
powerflow0057r	∞	∞
prob07	16.19	15.83
process	0.93	0.92
procurement1mot	79.5%	79.5%
procurement2mot	2.71	2.69
product	nonopt	nonopt
product2	2.53	2.84
prolog	100.0%	100.0%
qp3	4.6%	4.6%
qspp_0_10_0_1_10_1	26.39	26.51
qspp_0_11_0_1_10_1	124.07	124.05
radar-2000-10-a-6_lat_7	33.34	32.97
radar-3000-10-a-8_lat_7	209.09	208.61
ravempb	1.26	0.96
risk2bpb	10.70	10.68
routingdelay_bigm	18.4%	18.4%
rsyn0815m	1.02	1.05
rsyn0815m03m	44.8%	44.8%
sfacloc2_2_95	1.11	0.79

instance	10^{-6}	10^{-4}
sfacloc2_3_90	67.86	67.78
sjup2	∞	∞
slay06m	0.69	0.48
slay07m	0.61	0.91
smallinvDAXr1b010-011	8.69	8.99
smallinvDAXr1b020-022	205.80	208.43
sonet17v4	160.55	160.19
sonet18v6	184.16	184.08
sonetgr17	65.24	65.03
spectra2	105%	105%
sporttournament24	21.48	21.47
sporttournament30	1946.51	1959.43
sssd12-05persp	780.87	780.44
sssd18-06persp	19.9%	20.2%
st_testgr1	0.07	0.07
st_testgr3	0.07	0.07
steenbrf	98.0%	98.0%
stockcycle	39.57	39.61
supplychainp1_022020	8.5%	8.8%
supplychainp1_030510	3.08	3.10
supplychainr1_022020	2966.92	2966.61
supplychainr1_030510	0.74	0.74
syn15m04m	1.04	1.36
syn30m02m	1.83	1.56
synheat	17.1%	17.1%
tanksize	54.61	54.45
telecomsp_pacbell	1125.10	1111.55
tln5	2.08	2.10
tln7	274.33	274.65
tls2	0.18	0.22
tls4	36.82	37.19
topopt-mbb_60x40_50	∞	∞
toroidal2g20_5555	3.23	2.98
toroidal3g7_6666	67.16	67.08
transswitch0009r	5.6%	5.5%
tricip	100.0%	100.0%
tspn08	2.3%	2.3%
tspn15	81.4%	81.4%
unitcommit1	1082.69	1092.08
unitcommit2	11.20	11.23
wager	252.73	252.96
waste	99.9%	99.9%
wastepaper3	22.76	22.60
wastepaper4	2099.99	2110.22
wastepaper6	0.16%	0.21%
water4	1648.53	2062.46
waternd1	1.34	1.08
waterno2_02	4.18	4.11
waterno2_03	440.65	440.06
waterund01	0.18%	0.18%

B.3 Parallel Mode

The following table shows the outcome from running BARON in serial and parallel mode.

instance	1 thread	4 threads	8 threads	16 threads
alan	0.33	0.38	0.37	0.36
autocorr_bern20-05	7.29	1.91	1.33	1.31
autocorr_bern35-04	29.91	8.44	5.71	4.66
ball_mk2_10	0.07	0.07	0.07	0.20
ball_mk2_30	0.25	0.08	0.08	0.08
ball_mk3_10	10.58	11.19	11.35	11.71
batch0812.nc	6.57	5.33	6.47	5.32
batches101006m	6.91	13.82	9.41	11.02
batches121208m	47.33	26.40	24.95	32.16
bayes2_20	67.77	68.42	70.96	68.64
bayes2_30	21.02	22.46	21.94	23.08
blend029	1.03	1.08	1.25	0.95
blend146	3648.64	2175.95	2854.41	3654.96
camshape100	9.2%	9.2%	9.2%	9.2%
cardqp_inlp	12.40	10.82	12.25	18.77
cardqp_iqp	12.92	10.80	12.49	18.70
carton7	987.64	1964.21	1441.91	832.10

instance	1 thread	4 threads	8 threads	16 threads
carton9	45.0%	36.0%	46.9%	37.9%
casctanks	256.46	251.20	248.18	251.93
cecil13	280.62	238.03	201.60	195.64
celar6-sub0	100.0%	100.0%	100.0%	100.0%
chakra	0.14	0.15	0.14	0.15
chem	0.04	0.04	0.04	0.04
chenery	1.40	1.42	1.23	1.16
chimera_k64maxcut-01	560.62	183.06	119.55	171.40
chimera_mis-01	13.69	12.39	11.84	12.54
chp_shorttermplan1a	2296.84	2329.95	0.061%	1276.12
chp_shorttermplan2a	937.38	1002.49	999.57	956.66
chp_shorttermplan2b	0.34%	0.31%	0.3%	0.31%
clay0204m	0.47	0.37	0.24	1.28
clay0205m	4.92	3.78	2.63	4.94
color_lab3_3x0	2654.00	606.50	383.87	346.70
crossdock_15x7	385.56	87.79	49.36	38.69
crossdock_15x8	788.81	322.11	128.24	102.94
crudeoil_lee1.07	7.50	7.47	27.59	6.98
crudeoil_pooling_ct2	4.95	41.75	21.47	4.28
csched1	13.03	12.42	29.59	31.28
csched1a	0.70	0.76	0.70	0.78
cvxnonsep-psig20	0.99	0.92	0.91	0.95
cvxnonsep-psig30	4.25	4.06	4.13	4.21
du-opt	2.37	2.18	1.95	2.26
du-opt5	3.61	3.34	3.32	3.54
edgexcross10-040	7.25	2.73	2.13	2.54
edgexcross10-080	51.58	15.21	10.25	9.23
eg_all.s	infeas	infeas	infeas	90.9%
eigena2	1146.69	1150.94	1136.30	1132.85
elec50	66.5%	66.5%	66.5%	66.5%
elf	5.91	4.52	4.24	5.65
eniplac	4.54	4.04	5.48	9.66
enpro56pb	1.84	3.38	2.01	2.92
ex1244	4.86	4.56	4.49	4.23
ex1252a	9.12	9.19	8.97	10.85
fac1ay20h	349.36	91.42	98.32	47.33
fac1ay80	120%	120%	120%	120%
feedtray	80.5%	80.5%	80.5%	80.5%
fin2bb	114.16	83.45	97.44	91.84
flay04m	10.56	9.18	8.80	8.94
flay05m	394.85	370.14	353.95	411.75
flay06m	7%	7.2%	6.4%	6.6%
fo7_ar25_1	17.00	4.55	5.17	3.37
fo7_ar3_1	25.69	52.69	30.43	21.73
forest	747.74	1311.21	877.73	1470.56
gabriel01	1760.52	1619.79	1272.26	1131.82
gabriel02	11.1%	18.1%	9%	15.6%
gasnet	22.38	57.1%	50.7%	50.7%
gasprod_sarawak16	4578.18	5171.96	2524.08	2214.82
gastrans582_cold13_95	1442.44	732.29	954.39	1435.95
gastrans582_mild11	687.72	1206.29	1283.65	2130.36
gear	0.12	0.11	0.11	0.28
gear2	0.64	0.62	0.69	0.67
gear4	1.63	1.53	1.90	2.39
genpooling_lee1	9.37	5.21	5.37	5.46
genpooling_lee2	85.93	84.33	83.67	85.79
ghg_1veh	2.77	2.96	3.11	2.98
gilbert	2.35	2.60	2.56	2.38
graphpart_2g-0066-0066	0.65	0.76	0.46	0.37
graphpart_clique-60	34.6%	100.0%	100.0%	100.0%
gsg_0001	13.68	16.00	16.69	11.94
hadamard_5	54.92	17.35	8.54	9.77
heatexch_spec1	1.39	1.64	1.76	1.32
heatexch_spec2	6.60	7.91	8.50	7.88
hhfair	0.43	0.39	0.42	0.35
himmel16	41.46	36.26	36.42	36.46
house	0.37	0.38	0.38	0.57
hs62	1.07	1.08	0.94	0.88
hvb11	175.10	174.88	179.60	177.88
hybriddynamic_var	0.75	0.75	0.98	1.09
hybriddynamic_varcc	0.91	0.75	0.77	0.78
hydroenergy1	1026.99	1057.36	1054.54	1057.15
ibs2	nonopt	nonopt	nonopt	nonopt
johnall	2.85	2.51	2.50	2.23
kall_circles_c6b	310.13	313.79	307.19	306.85
kall_congruentcircles_c72	26.45	27.78	24.66	22.91
kissing2	184.78	207.89	184.05	207.42
kport20	6.1%	3.3%	6.8%	5.7%

instance	1 thread	4 threads	8 threads	16 threads
kriging_peaks-red020	10.33	9.90	9.00	9.43
kriging_peaks-red100	196.08	198.30	195.92	193.18
lop97icx	1709.08	450.05	315.56	317.81
mathopt5_7	0.29	0.13	0.12	0.12
mathopt5_8	0.26	0.43	0.40	0.26
maxcsp-geo50-20-d4-75-36	16.94	18.01	15.92	17.40
meanvar-ori400_05_e_7	95.4%	94.7%	94.3%	94.8%
meanvar-ori400_05_e_8	600.24	541.81	508.87	455.48
mhw4d	0.83	0.70	0.67	0.80
milinfract	55.06	54.78	54.71	55.20
minlphi	∞	∞	∞	∞
multiplants_mtgl1a	2245.95	2338.78	4026.31	1138.55
multiplants_mtgl2	2115.70	2189.56	2175.33	2125.52
nd_netgen-3000-1-1-b-b-ns_7	∞	∞	∞	∞
netmod_kar1	4.53	5.70	5.71	6.81
netmod_kar2	5.26	5.46	5.75	6.88
nous1	51.20	47.60	44.39	53.14
nous2	0.40	0.35	0.34	0.36
nvs02	0.04	0.05	0.05	0.05
nvs06	0.08	0.08	0.10	0.10
oil2	3.30	3.02	3.00	2.67
optmass	∞	∞	∞	∞
ortez	nonopt	nonopt	nonopt	nonopt
p_ball_10b_5p_3d_m	32.22	23.87	20.81	25.42
p_ball_15b_5p_2d_m	90.52	63.59	63.02	77.58
parabol5_2_3	0.051%	0.051%	0.051%	0.051%
parallel	19.19	21.16	20.85	18.15
pedigree_ex485	68.37	39.15	79.06	414.20
pedigree_ex485_2	17.30	14.20	16.23	14.65
pointpack06	5.59	6.60	4.94	5.96
pointpack08	159.31	153.67	154.66	151.67
pooling_epa1	12.84	11.97	14.83	14.15
pooling_epa2	2.4%	0.44%	5.5%	3%
portfol_buyin	0.34	0.33	0.34	0.32
portfol_card	0.31	0.30	0.29	0.30
powerflow0014r	50.50	51.54	50.22	53.19
powerflow0057r	∞	∞	∞	∞
prob07	75.62	81.05	76.59	71.80
process	0.51	0.52	0.42	0.42
procurement1mot	59.57	59.24	63.12	61.09
procurement2mot	4.13	4.24	3.78	3.65
product	48.09	24.31	49.29	nonopt
product2	2.90	225.61	76.64	289.08
prolog	0.31	0.32	0.33	0.36
qp3	13.7%	14.0%	14.2%	14.4%
qspp_0_10_0_1_10_1	24.2%	23.8%	23.7%	23.8%
qspp_0_11_0_1_10_1	30.5%	30.1%	30.1%	30.1%
radar-2000-10-a-6_lat_7	75.1%	91.0%	91.0%	91.0%
radar-3000-10-a-8_lat_7	100.0%	100.0%	100.0%	100.0%
ravempb	0.82	0.80	1.05	0.94
risk2bpb	0.70	0.43	0.46	0.49
routingdelay_bigm	26.08	23.23	23.91	24.85
rsyn0815m	0.29	0.26	0.26	0.26
rsyn0815m03m	31.23	38.94	34.49	39.53
sfacloc2_2_95	0.31	0.28	0.41	0.27
sfacloc2_3_90	29.12	24.28	14.83	16.30
sjup2	∞	∞	∞	∞
slay06m	0.46	0.41	0.48	0.59
slay07m	0.62	0.64	0.58	0.58
smallinvDAXr1b010-011	1.62	1.52	1.44	1.52
smallinvDAXr1b020-022	1.95	2.26	1.63	1.75
sonet17v4	339.62	62.99	33.05	33.46
sonet18v6	429.30	97.11	62.50	60.56
sonetgr17	346.77	74.19	44.37	29.83
spectra2	1.37	1.29	1.47	1.38
sporttournament24	4982.31	1521.61	722.06	763.20
sporttournament30	11.8%	67.7%	67.7%	67.7%
sssd12-05persp	27.8%	28.0%	27.8%	27.6%
sssd18-06persp	41.4%	40.5%	41.4%	41.3%
st_testgr1	0.24	0.14	0.16	0.22
st_testgr3	0.24	0.28	0.35	0.50
steenbrf	14.39	15.32	15.07	17.55
stockcycle	7.82	6.48	6.94	7.08
supplychainp1_022020	nonopt	nonopt	nonopt	nonopt
supplychainp1_030510	nonopt	nonopt	nonopt	nonopt
supplychainr1_022020	6.13	5.67	5.47	5.25
supplychainr1_030510	0.28	0.26	0.25	0.37
syn15m04m	0.54	0.45	0.47	0.49

instance	1 thread	4 threads	8 threads	16 threads
syn30m02m	2.51	2.20	2.38	2.21
synheat	35.92	28.74	26.21	25.37
tanksize	4.81	4.17	4.75	4.07
telecomsp_pacbell	3393.96	1396.50	1189.08	1401.72
tlm5	5.26	1.33	5.03	5.53
tlm7	2935.35	1739.45	3085.42	2.7%
tls2	0.13	0.12	0.22	0.12
tls4	35.15	16.87	14.67	20.79
topopt-mbb_60x40_50	∞	∞	∞	∞
toroidal2g20_5555	3747.17	114.31	81.54	190.59
toroidal3g7_6666	6.2%	74.7%	74.7%	74.7%
transswitch0009r	8%	8%	8%	8%
tripc	0.82	1.25	1.11	1.14
tspn08	10.81	9.28	11.25	9.49
tspn15	884.44	1003.34	412.58	410.34
unitcommit1	0.27%	0.27%	0.27%	0.27%
unitcommit2	7200.00	725.90	6169.57	732.94
wager	15.30	27.54	25.87	27.99
waste	22.55	24.58	23.21	29.11
wastepaper3	9.19	14.02	14.63	12.67
wastepaper4	871.50	417.33	763.71	556.85
wastepaper6	0.022%	0.022%	7200.00	0.022%
water4	44.4%	43.7%	47.6%	49.2%
waternd1	5.50	6.83	6.98	6.07
waterno2_02	5.20	4.18	4.33	5.03
waterno2_03	260.33	262.06	272.21	272.70
waterund01	0.35%	0.35%	0.35%	0.35%

The following table shows the outcome from running Lindo API in serial and parallel mode.

instance	1 thread	4 threads	8 threads	16 threads
alan	0.89%	0.89%	0.89%	0.89%
autocorr_bern20-05	100.43	33.83	23.88	21.91
autocorr_bern35-04	13.5%	2655.83	1788.33	2461.81
ball_mmk2_10	3.98	2.17	1.28	1.43
ball_mmk2_30	100.0%	100.0%	100.0%	100.0%
ball_mmk3_10	3.58	1.79	2.05	2.17
batch0812_nc	25.46	18.73	17.51	17.68
batches101006m	134.04	abort	abort	abort
batches121208m	198.76	abort	abort	abort
bayes2_20	0.033%	0.033%	0.033%	infeas
bayes2_30	7200.00	abort	7200.00	infeas
blend029	abort	205.94	166.19	96.01
blend146	abort	5.8%	4.2%	6.8%
camshape100	9.2%	8.9%	8.7%	8.4%
cardqp_inlp	134.08	nonopt	nonopt	nonopt
cardqp_iqp	133.22	nonopt	nonopt	nonopt
carton7	110.53	110.35	110.48	110.50
carton9	287.45	287.68	287.13	287.83
casctanks	3.4%	3.4%	3.4%	3.4%
cecil13	0.32%	0.16%	0.11%	5951.36
celar6-sub0	100%	100%	100%	100%
chakra	1.28	0.34	infeas	infeas
chem	1172.22	471.44	314.57	329.14
chenery	0.43%	0.43%	0.43%	abort
chimera_k64maxcut-01	16.4%	16.4%	16.1%	16.7%
chimera_mis-01	nonopt	nonopt	nonopt	nonopt
chp_shorttermplan1a	∞	∞	∞	∞
chp_shorttermplan2a	∞	∞	∞	∞
chp_shorttermplan2b	16.6%	6045.30	5089.35	3782.12
clay0204m	4.75	2.51	2.56	2.44
clay0205m	51.80	18.76	11.84	10.23
color_lab3_3x0	65.6%	65.6%	65.6%	65.6%
crossdock_15x7	124%	124%	124%	124%
crossdock_15x8	124%	124%	124%	124%
crudeoil_lee1_07	63.26	63.10	63.51	63.54
crudeoil_pooling_ct2	1.8%	3.1%	1.6%	1.5%
csched1	66.81	65.05	46.49	36.59
csched1a	4.85	2.82	2.44	2.36
cvxnonsep_psig20	0.48	0.25	0.25	0.25
cvxnonsep_psig30	3.26	3.00	3.05	2.95
du-opt	5.35	5.94	5.92	6.74
du-opt5	4.70	4.77	4.63	5.17
edgexcross10-040	0.29	0.12	0.12	0.12

instance	1 thread	4 threads	8 threads	16 threads
edgexcross10-080	nonopt	nonopt	nonopt	nonopt
eg_all.s	abort	abort	abort	abort
eigena2	∞	∞	∞	∞
elec50	798.72	648.65	757.22	656.49
elf	nonopt	2.56	infeas	3.02
eniplac	3%	3952.97	3085.99	2683.08
enpro56pb	1454.20	564.02	92.37	42.35
ex1244	9.69	6.99	6.71	6.71
ex1252a	infeas	infeas	infeas	infeas
faclay20h	nonopt	nonopt	nonopt	nonopt
faclay80	5096.67	5092.40	5092.51	5120.74
feedtray	13.79	13.86	13.58	13.71
fin2bb	1294.03	3136.42	0.59%	3467.31
flay04m	647.41	231.73	166.42	123.11
flay05m	0.54%	0.038%	5375.73	3837.92
flay06m	14.3%	9.5%	8.2%	6.8%
fo7_ar25_1	1282.08	662.73	661.65	661.84
fo7_ar3_1	1847.41	796.80	752.45	954.82
forest	1.18	1.18	1.23	1.53
gabriel01	5.4%	4.5%	4.2%	3.8%
gabriel02	44.0%	43.0%	42.2%	42.1%
gasnet	64.9%	64.1%	63.9%	63.8%
gasprod_sarawak16	infeas	0.4%	0.4%	1.2%
gastrans582_cold13_95	∞	∞	∞	∞
gastrans582_mild11	∞	∞	∞	∞
gear	0.05	0.06	0.06	0.06
gear2	0.23	0.23	0.24	0.24
gear4	infeas	infeas	infeas	infeas
genpooling_lee1	0.41%	680.62	136.93	infeas
genpooling_lee2	3724.02	169.90	111.71	86.41
ghg_1veh	106.61	60.51	46.93	45.15
gilbert	14.25	28.89	25.67	34.60
graphpart_2g-0066-0066	8.7%	8.7%	8.7%	8.7%
graphpart_clique-60	83.3%	83.3%	83.3%	83.3%
gsg_0001	351.30	578.61	338.79	308.84
hadamard_5	129.27	73.47	59.47	35.89
heatexch_spec1	0.34%	0.14%	0.16%	0.25%
heatexch_spec2	0.041%	0.052%	0.039%	0.038%
hhfair	100.0%	100.0%	100.0%	100.0%
himmel16	20.42	8.10	5.33	3.63
house	0.89	infeas	infeas	infeas
hs62	4.04	3.33	2.13	1.81
hvb11	40.7%	27.4%	20.5%	nonopt
hybriddynamic_var	4.11	1.54	1.26	0.78
hybriddynamic_varcc	12.00	4.55	3.17	2.71
hydroenergy1	0.93%	0.74%	0.69%	0.59%
ibs2	abort	abort	abort	abort
johnall	22.36	22.56	22.26	22.48
kall_circles_c6b	62.28	120.62	90.79	80.68
kall_congruentcircles_c72	4.74	7.27	5.22	5.33
kissing2	867.71	866.79	868.19	869.92
kport20	13.1%	13.1%	12.4%	11.5%
kriging_peaks-red020	21.79	7.24	4.12	3.65
kriging_peaks-red100	140.14	59.63	41.23	32.35
lop97icx	39.83	13.81	nonopt	nonopt
mathopt5_7	12.35	10.93	10.78	11.87
mathopt5_8	9.72	8.27	8.20	9.06
maxcsp-geo50-20-d4-75-36	101%	101%	101%	101%
meanvar-orl400_05_e_7	23.95	16.40	14.93	nonopt
meanvar-orl400_05_e_8	19.32	19.40	abort	19.05
mhw4d	0.78	0.63	0.41	0.65
milinfract	infeas	68.0%	67.8%	68.1%
minlphi	1.86	1.14	infeas	infeas
multiplants_mtg1a	15.6%	11.9%	5.6%	7%
multiplants_mtg2	0.12%	1303.99	1336.22	1444.62
nd_netgen-3000-1-1-b-b-ns_7	680.80	infeas	760.69	infeas
netmod_kar1	5948.45	5.7%	6.2%	4040.34
netmod_kar2	5917.59	6.5%	7.6%	infeas
nous1	2.9%	0.79%	6836.04	7023.44
nous2	2.16	1.27	1.08	1.36
nvs02	69.68	70.03	69.93	70.27
nvs06	6.64	6.36	6.68	6.68
oil2	43.56	43.27	43.63	43.29
optmass	11.3%	12.8%	12.8%	12.8%
ortez	2.42	2.38	2.28	2.38
p_ball_10b_5p_3d_m	26.41	11.49	7.87	9.79
p_ball_15b_5p_2d_m	101.62	32.66	13.84	19.15
parabol5_2.3	15.7%	15.6%	15.5%	15.5%

instance	1 thread	4 threads	8 threads	16 threads
parallel	38.93	153.77	89.39	66.77
pedigree_ex485	5.8%	5.7%	5.1%	2.4%
pedigree_ex485_2	192.97	128.78	83.00	81.80
pointpack06	27.80	11.29	6.94	10.78
pointpack08	3611.87	1638.66	2129.26	1356.41
pooling_epa1	infeas	infeas	infeas	abort
pooling_epa2	0.041%	4885.54	0.081%	6389.62
portfol_buyin	2.34	2.27	2.28	2.29
portfol_card	2.15	2.19	1.96	2.20
powerflow0014r	100.0%	82.1%	71.3%	60.6%
powerflow0057r	∞	abort	abort	abort
prob07	68.41	155.23	105.41	100.26
process	1.38	1.88	1.10	0.96
procurement1mot	86.4%	88.5%	89.2%	89.3%
procurement2mot	13.59	abort	abort	abort
product	465.36	448.67	267.76	infeas
product2	145.38	145.40	145.40	145.42
prolog	100.0%	0.40	infeas	1.50
qp3	0.01	0.01	0.01	0.01
qspp_0_10_0_1_10_1	632.32	624.71	631.86	623.84
qspp_0_11_0_1_10_1	3676.78	3669.93	3677.76	3673.84
radar-2000-10-a-6_lat_7	33.14	33.30	33.04	33.00
radar-3000-10-a-8_lat_7	558.69	495.51	498.02	538.74
ravempb	3438.61	305.70	26.47	infeas
risk2bpb	1.46	1.20	1.20	1.19
routingdelay_bigm	nonopt	nonopt	nonopt	nonopt
rsyn0815m	92.08	18.96	15.88	14.40
rsyn0815m03m	nonopt	nonopt	infeas	nonopt
sfacloc2_2_95	5.25	5.49	5.55	5.50
sfacloc2_3_90	588.76	496.17	400.13	0.24%
sjup2	3456.91	2496.08	2490.04	2495.30
slay06m	5.66	4.64	4.32	5.06
slay07m	6.96	6.39	6.64	6.65
smallinvDAXr1b010-011	3.51	4.58	4.67	5.10
smallinvDAXr1b020-022	3.46	4.79	nonopt	nonopt
sonet17v4	nonopt	nonopt	nonopt	nonopt
sonet18v6	nonopt	nonopt	nonopt	nonopt
sonetgr17	nonopt	1464.38	nonopt	nonopt
spectra2	infeas	infeas	infeas	infeas
sporttournament24	6.7%	5.4%	4.7%	4.7%
sporttournament30	12.5%	10.4%	11.2%	10.8%
sssd12-05persp	378.13	abort	abort	abort
sssd18-06persp	7200.00	7200.00	abort	abort
st_testgr1	5.51	2.49	2.78	nonopt
st_testgr3	3.67	3.20	3.79	4.19
steenbrf	0.75	0.73	0.77	1.52
stockcycle	3.18	3.97	4.50	3.75
supplychainp1_022020	20.8%	13.9%	11.9%	11.8%
supplychainp1_030510	63.91	66.23	63.40	48.63
supplychainr1_022020	12.5%	8.8%	7%	2.9%
supplychainr1_030510	7.19	8.73	8.30	8.18
syn15m04m	42.37	19.93	13.59	14.97
syn30m02m	239.18	123.93	79.51	85.36
synheat	0.36%	0.19%	0.28%	0.16%
tanksize	13.85	7.22	5.12	4.44
telecomsp_pacbell	2.8%	3.5%	2.4%	5.1%
tln5	30.1%	26.2%	7.8%	1682.28
tln7	61.4%	60.0%	9.9%	28.3%
tls2	5.43	3.96	4.13	3.86
tls4	833.24	319.16	276.69	302.31
topopt-mbb_60x40_50	195.38	194.93	194.58	194.68
toroidal2g20_5555	nonopt	nonopt	nonopt	nonopt
toroidal3g7_6666	nonopt	nonopt	nonopt	nonopt
transswitch0009r	16.0%	abort	abort	abort
tricip	infeas	infeas	infeas	infeas
tspn08	10.9%	10.8%	10.8%	10.7%
tspn15	12.8%	12.6%	12.4%	12.4%
unitcommit1	81.17	68.02	67.44	78.21
unitcommit2	18.0%	18.0%	17.8%	17.7%
wager	∞	32.4%	∞	∞
waste	46.3%	46.1%	46.1%	46.1%
wastepaper3	33.90	26.89	0.83%	23.07
wastepaper4	2320.60	1388.31	1206.13	1041.52
wastepaper6	abort	7200.00	7200.00	7200.00
water4	46.1%	39.1%	42.4%	34.1%
waternd1	1130.72	853.76	525.28	560.05
waterno2_02	26.45	12.97	9.72	10.04
waterno2_03	1886.92	737.09	532.41	465.62

instance	1 thread	4 threads	8 threads	16 threads
waterund01	infeas	infeas	infeas	infeas

The following table shows the outcome from running Octeract in serial and parallel mode.

instance	1 thread	4 threads	8 threads	16 threads
alan	0.05	0.06	0.08	0.09
autocorr_bern20-05	4.36	4.91	4.88	7.71
autocorr_bern35-04	12.7%	4064.07	1891.23	2241.67
ball_mk2.10	0.01	0.01	0.02	0.02
ball_mk2.30	0.02	0.02	0.02	0.03
ball_mk3.10	0.04	0.04	0.04	0.04
batch0812.nc	8.37	4.85	3.34	4.12
batches101006m	4.84	5.55	5.40	7.16
batches121208m	79.1%	79.1%	79.1%	abort
bayes2.20	0.033%	0.033%	3777.02	6089.69
bayes2.30	7200.00	7200.00	7200.00	7200.00
blend029	29.33	13.45	10.27	11.02
blend146	8.8%	3.4%	6.8%	2.6%
camshape100	19.40	9.17	7.12	9.91
cardqp_inlp	792.06	258.89	147.93	124.77
cardqp_iqp	785.02	259.33	148.69	124.80
carton7	6.03	1.00	0.89	0.76
carton9	344.94	7.49	3.62	3.55
casctanks	11.8%	12.1%	11.6%	12.5%
cecil_13	188.40	99.55	134.75	166.38
celar6-sub0	∞	1637.26	1641.73	1562.09
chakra	7200.00	0.021%	0.022%	7200.00
chem	0.18	0.19	0.19	0.28
chenery	200.71	77.26	35.42	25.05
chimera_k64maxcut-01	168.60	55.75	40.84	39.93
chimera_mis-01	1.14	1.50	1.38	1.94
chp_shorttermplan1a	64.74	0.028%	0.023%	0.042%
chp_shorttermplan2a	17.59	16.76	16.30	21.88
chp_shorttermplan2b	0.67%	0.67%	0.71%	0.7%
clay0204m	1.47	2.10	2.28	3.10
clay0205m	15.10	13.45	13.77	29.70
color_lab3.3x0	∞	∞	5460.67	∞
crossdock_15x7	∞	6047.29	2087.00	1851.98
crossdock_15x8	6306.65	4889.74	2772.27	2837.68
crudeoil_lee1.07	8.56	8.03	4.93	5.83
crudeoil_pooling_ct2	nonopt	nonopt	nonopt	nonopt
csched1	8.1%	4306.97	2116.58	1592.40
csched1a	17.77	11.89	5.44	5.07
cvxnonsep-psig20	35.0%	28.1%	24.8%	25.2%
cvxnonsep-psig30	45.6%	36.6%	35.3%	34.4%
du-opt	122.38	18.42	15.18	8.15
du-opt5	101.32	48.22	17.33	14.29
edgexross10-040	4.16	0.79	0.44	0.63
edgexross10-080	6%	3.8%	5205.96	3.8%
eg_all.s	102%	76.0%	41.6%	198%
eigena2	416.89	461.71	602.17	∞
elec50	66.4%	66.4%	66.4%	66.3%
elf	2.54	infeas	infeas	infeas
eniplac	1.54	1.06	1.01	1.61
enpro56pb	1.98	2.48	2.25	3.52
ex1244	81.99	35.00	0.056%	0.082%
ex1252a	245.41	94.34	42.83	31.72
faclay20h	785.28	413.15	463.78	418.22
faclay80	∞	∞	∞	∞
feedtray	82.1%	82.1%	82.2%	82.1%
fin2bb	100.0%	100.0%	100.0%	100.0%
flay04m	0.87	0.34	0.27	0.28
flay05m	infeas	infeas	infeas	infeas
flay06m	infeas	infeas	566.98	413.02
fo7_ar25_1	8.95	3.12	2.85	1.49
fo7_ar3_1	18.33	3.23	2.13	2.18
forest	nonopt	nonopt	nonopt	nonopt
gabriel01	2%	0.28%	3548.92	2168.66
gabriel02	7078.04	2457.11	1088.23	588.74
gasnet	96.7%	96.4%	96.2%	96.1%
gasprod_sarawak16	0.74%	0.68%	0.5%	0.4%
gastrans582_cold13-95	∞	∞	∞	∞
gastrans582_mild11	∞	∞	∞	∞
gear	0.11	0.09	0.09	0.14

instance	1 thread	4 threads	8 threads	16 threads
gear2	0.14	0.13	0.13	0.26
gear4	17.56	7.25	3.29	2.45
genpooling_lee1	117.04	40.77	20.27	15.25
genpooling_lee2	186.63	77.47	36.61	23.73
ghg_1veh	12.42	infeas	3.58	4.05
gilbert	0.9%	9.7%	9.7%	22.5%
graphpart_2g-0066-0066	0.75	0.19	0.21	0.22
graphpart_clique-60	2890.51	1502.07	287.82	802.57
gsg_0001	8.58	4.67	3.04	2.54
hadamard_5	20.45	23.43	23.33	38.20
heatexch_spec1	17.7%	17.0%	14.2%	12.7%
heatexch_spec2	5%	5.3%	5.2%	4.9%
hhfair	∞	∞	∞	∞
himmel16	12.95	5.89	3.28	2.22
house	108.41	35.02	20.60	12.89
hs62	0.023%	0.02%	7200.00	7200.00
hvb11	7.3%	4.1%	2.1%	1.5%
hybriddynamic_var	0.32%	166.64	73.29	53.94
hybriddynamic_varcc	158.14	61.30	27.37	19.22
hydroenergy1	0.65%	0.48%	0.36%	0.3%
ibs2	5.1%	6.4%	5%	5%
johnall	44.23	50.09	51.26	88.02
kall_circles_c6b	463.62	158.77	77.24	55.71
kall_congruentcircles_c72	49.42	19.60	8.03	6.64
kissing2	100.0%	100.0%	100.0%	100.0%
kport20	3596.06	1590.23	472.21	381.94
kriging_peaks-red020	85.61	35.87	18.22	16.08
kriging_peaks-red100	629.86	277.25	140.33	113.86
lop97icx	5.90	3.76	3.67	4.60
mathopt5_7	0.08	0.07	0.07	0.11
mathopt5_8	0.06	0.08	0.06	0.09
maxcsp-geo50-20-d4-75-36	7.71	9.59	21.01	23.19
meanvar-orl400_05_e_7	∞	2015.07	1112.27	1205.72
meanvar-orl400_05_e_8	6.00	5.57	5.87	8.45
mhw4d	0.43	0.26	0.39	0.40
millinfract	75.5%	75.2%	75.0%	75.7%
minlphi	100%	100%	100%	100%
multiplants_mtg1a	5.1%	303.62	153.51	64.70
multiplants_mtg2	2%	1.5%	1.4%	1.6%
nd_netgen-3000-1-1-b-b-ns_7	3.67	infeas	infeas	6.50
netmod_kar1	41.89	28.58	25.95	33.54
netmod_kar2	42.01	28.58	26.39	33.30
nous1	52.93	32.90	19.28	18.50
nous2	8.29	6.12	3.68	4.58
nvs02	0.22	0.25	0.25	0.22
nvs06	0.06	0.09	0.07	0.10
oil2	nonopt	nonopt	nonopt	nonopt
optmass	463.14	237.53	155.15	239.61
ortez	9.48	7.09	1.10	1.18
p_ball_10b_5p_3d_m	49.66	infeas	infeas	infeas
p_ball_15b_5p_2d_m	infeas	infeas	infeas	infeas
parabol5_2_3	4095.70	1851.01	1883.43	0.11%
parallel	102.24	39.89	18.25	15.46
pedigree_ex485	276.45	65.18	45.86	46.16
pedigree_ex485_2	3.66	2.61	2.57	3.08
pointpack06	4.76	1.93	0.90	1.28
pointpack08	208.18	70.21	34.30	23.21
pooling_epa1	21.07	14.15	13.18	15.01
pooling_epa2	1694.37	463.42	233.31	242.45
portfol_buyin	3.56	1.60	1.12	1.47
portfol_card	5.29	2.64	1.68	5.86
powerflow0014r	100.0%	100.0%	100.0%	28.4%
powerflow0057r	∞	∞	∞	∞
prob07	16.19	6.21	2.91	2.55
process	0.93	0.30	0.21	0.33
procurement1mot	79.5%	77.1%	74.6%	73.4%
procurement2mot	2.71	2.88	2.92	4.47
product	nonopt	nonopt	nonopt	nonopt
product2	2.53	2.51	2.40	3.29
prolog	100.0%	100.0%	100.0%	100.0%
qp3	4.6%	3.8%	3.2%	3%
qspp_0_10_0_1_10_1	26.39	9.61	7.43	7.73
qspp_0_11_0_1_10_1	124.07	39.87	27.50	21.59
radar-2000-10-a-6_lat_7	33.34	8.06	6.80	11.05
radar-3000-10-a-8_lat_7	209.09	17.99	17.84	138.18
ravempb	1.26	1.06	1.04	1.65
risk2bpb	10.70	nonopt	nonopt	nonopt
routingdelay_bigm	18.4%	18.6%	18.6%	18.6%

instance	1 thread	4 threads	8 threads	16 threads
rsyn0815m	1.02	0.81	0.81	1.26
rsyn0815m03m	44.8%	abort	abort	abort
sfacloc2_2_95	1.11	1.44	1.34	1.72
sfacloc2_3_90	67.86	42.11	42.06	46.51
sjup2	∞	∞	∞	∞
slay06m	0.69	0.84	0.62	1.00
slay07m	0.61	0.82	0.92	1.24
smallinvDAXr1b010-011	8.69	4.27	4.64	4.82
smallinvDAXr1b020-022	205.80	62.49	23.58	17.20
sonet17v4	160.55	32.25	17.94	17.26
sonet18v6	184.16	37.12	22.59	20.81
sonetgr17	65.24	18.42	13.25	10.75
spectra2	105%	103%	104%	104%
sporttournament24	21.48	5.01	7.20	5.26
sporttournament30	1946.51	476.70	242.25	285.89
sssd12-05persp	780.87	100.0%	100.0%	100.0%
sssd18-06persp	19.9%	100.0%	100.0%	infeas
st_testgr1	0.07	0.08	0.09	0.12
st_testgr3	0.07	0.10	0.11	0.10
steenbrf	98.0%	97.6%	97.6%	97.0%
stockcycle	39.57	34.72	19.66	infeas
supplychainp1_022020	8.5%	8.7%	8.7%	8.7%
supplychainp1_030510	3.08	2.71	2.68	3.86
supplychainr1_022020	2966.92	752.44	333.16	225.20
supplychainr1_030510	0.74	0.66	0.64	0.94
syn15m04m	1.04	1.08	1.08	1.77
syn30m02m	1.83	1.72	2.24	2.89
synheat	17.1%	16.2%	14.0%	13.1%
tanksize	54.61	20.94	9.66	7.03
telecomsp_pacbell	1125.10	966.13	632.56	1152.71
tlm5	2.08	1.50	1.36	1.60
tlm7	274.33	144.17	112.57	78.00
tls2	0.18	0.18	0.18	0.27
tls4	36.82	18.27	10.79	17.84
topopt-mbb_60x40_50	∞	∞	∞	∞
toroidal2g20_5555	3.23	2.43	2.23	2.38
toroidal3g7_6666	67.16	25.61	21.88	23.43
transswitch0009r	5.6%	3%	1.5%	0.97%
trisp	100.0%	1236.23	100.0%	100.0%
tspn08	2.3%	2%	1.6%	1.4%
tspn15	81.4%	82.3%	75.3%	80.3%
unitcommit1	1082.69	422.70	208.33	152.08
unitcommit2	11.20	10.85	10.38	12.71
wager	252.73	100.0%	100.0%	100.0%
waste	99.9%	99.9%	99.9%	99.9%
wastepaper3	22.76	14.47	11.86	12.59
wastepaper4	2099.99	853.02	460.96	443.80
wastepaper6	0.16%	0.2%	0.16%	0.12%
water4	1648.53	704.24	302.89	201.96
waternd1	1.34	1.08	1.08	1.81
waterno2_02	4.18	5.63	6.82	8.84
waterno2_03	440.65	214.98	111.60	71.05
waterund01	0.18%	0.13%	0.11%	0.097%

The following table shows the outcome from running SCIP (1 thread) and FiberSCIP (4, 8, 16 threads).

instance	1 thread	4 threads	8 threads	16 threads
alan	0.19	1.05	1.06	1.13
autocorr_bern20-05	16.34	33.08	17.10	18.19
autocorr_bern35-04	107.46	158.09	101.10	64.17
ball_mk2_10	0.05	1.05	1.06	1.12
ball_mk2_30	0.06	1.06	1.06	1.14
ball_mk3_10	0.00	0.00	0.00	0.00
batch0812_nc	1.45	1.06	1.08	1.13
batches101006m	7.63	abort	abort	abort
batches121208m	6.99	abort	7.15	abort
bayes2_20	0.033%	0.97%	0.033%	0.033%
bayes2_30	7200.00	7200.00	7200.00	7200.00
blend029	3.67	3.06	3.09	2.12
blend146	471.12	280.13	119.13	176.20
camshape100	5.4%	10.5%	10.4%	9.7%
cardqp_inlp	2751.70	2408.46	1083.38	1285.54
cardqp_iqp	2769.87	2406.76	1190.43	758.44
carton7	13.54	10.89	6.94	8.69

instance	1 thread	4 threads	8 threads	16 threads
carton9	36.30	34.26	23.15	20.50
casctanks	199%	114%	114%	114%
cecil13	605.03	0.36%	72.23	71.31
celar6-sub0	1503.66	639.66	599.98	685.13
chakra	0.04	1.05	1.06	1.13
chem	932.15	376.09	545.11	6288.63
chenery	2.10	16.06	4.06	3.13
chimera_k64maxcut-01	450.42	435.40	352.61	387.45
chimera_mis-01	6.76	8.20	7.31	10.29
chp_shorttermplan1a	17.55	7.21	7.24	7.36
chp_shorttermplan2a	27.16	17.68	12.85	11.82
chp_shorttermplan2b	7200.00	23.3%	17.9%	11.5%
clay0204m	1.67	2.04	1.09	2.12
clay0205m	8.40	11.07	9.07	6.15
color_lab3_3x0	15.3%	66.8%	63.3%	59.4%
crossdock_15x7	32.0%	88.8%	87.6%	83.5%
crossdock_15x8	42.8%	100.0%	89.6%	92.3%
crudeoil_lee1.07	4.72	2.20	2.16	2.23
crudeoil_pooling_ct2	18.94	4.5%	0.22%	0.26%
csched1	1.59	2.08	2.09	3.10
csched1a	6.53	2.07	2.08	4.11
cvxnonsep-psig20	12.04	0.086%	163.10	119.16
cvxnonsep-psig30	65.40	3.4%	8.2%	52.11
du-opt	2.38	3.30	3.20	3.30
du-opt5	1.69	1.15	1.16	1.23
edgexcross10-040	8.36	4.11	5.12	6.13
edgexcross10-080	74.51	53.15	62.17	47.26
eg_all.s	3095.77	3082.62	1220.55	1492.45
eigena2	∞	7200.00	∞	∞
elec50	45.1%	44.9%	44.9%	44.9%
elf	1.01	1.04	1.07	1.13
eniplac	2.19	3.07	3.08	3.13
enpro56pb	3.40	3.07	3.07	3.13
ex1244	7.67	2.06	1.09	2.05
ex1252a	7200.00	13.05	4.06	4.13
faclay20h	741.15	589.88	513.92	832.87
faclay80	160%	∞	∞	∞
feedtray	1.70	87.7%	80.5%	80.5%
fin2bb	15.92	100.0%	100.0%	100.0%
flay04m	3.93	5.04	4.05	3.11
flay05m	120.82	44.06	15.6%	16.12
flay06m	4931.28	27.6%	25.1%	10.4%
fo7_ar25_1	39.74	60.06	24.04	15.12
fo7_ar3_1	88.79	65.08	37.08	24.13
forest	538.83	0.21%	abort	0.19%
gabriel01	351.17	353.14	116.14	146.23
gabriel02	1435.31	1638.34	782.25	267.30
gasnet	42.5%	65.0%	64.4%	63.6%
gasprod_sarawak16	0.39%	0.4%	0.39%	0.39%
gastrans582_cold13_95	55.01	6.96	4.97	6.99
gastrans582_mild11	7.41	257.72	13.73	4.04
gear	12.55	24.05	5.06	39.16
gear2	18.25	51.07	7.07	597.17
gear4	4.05	1.05	1.06	1.11
genpooling_lee1	2.09	3.06	2.07	2.13
genpooling_lee2	6.88	4.06	4.08	4.15
ghg_1veh	32.25	26.07	22.10	21.14
gilbert	1.13	5.12	11.14	19.72
graphpart_2g-0066-0066	1.39	1.10	1.11	1.19
graphpart_clique-60	3962.88	36.3%	5234.58	2596.77
gsg_0001	32.26	799.12	1209.24	7.12
hadamard_5	21.56	16.07	11.07	8.15
heatexch_spec1	7.1%	3.8%	13.09	88.12
heatexch_spec2	infeas	12.08	110.11	5.15
hhfair	100.0%	∞	∞	100.0%
himmel16	6.13	6.08	4.07	4.04
house	0.33	1.06	1.07	1.12
hs62	2.65	2.05	3.06	2.11
hvb11	103.26	27.11	27.13	2285.34
hybriddynamic_var	2.21	1.05	1.07	1.10
hybriddynamic_varcc	2.40	3.05	2.08	2.14
hydroenergy1	6295.41	5582.56	1646.23	3446.41
ibs2	20.64	16.24	17.02	22.00
johnall	30.44	55.01	46.41	44.71
kall_circles_c6b	183.83	95.07	48.08	30.12
kall_congruentcircles_c72	22.40	21.05	12.07	7.13
kissing2	100.0%	100.0%	100.0%	100.0%
kport20	1001.03	nonopt	334.10	28.12

instance	1 thread	4 threads	8 threads	16 threads
kriging_peaks-red020	117.65	36.08	39.08	69.16
kriging_peaks-red100	7200.00	384.19	325.21	1028.67
lop97icx	25.38	118.18	84.13	23.21
mathopt5_7	0.17	1.04	1.05	1.10
mathopt5_8	0.14	1.05	1.06	1.11
maxcsp-geo50-20-d4-75-36	52.80	63.05	74.36	74.19
meanvar-orl400_05_e_7	4080.76	abort	2.3%	2.4%
meanvar-orl400_05_e_8	2404.35	0.54%	abort	4849.82
mhw4d	0.44	1.05	1.07	1.09
milinfract	68.8%	76.0%	76.0%	75.9%
minlphi	100.0%	34.04	100%	100%
multiplants_mtgl1a	4320.86	35.6%	6.4%	5.5%
multiplants_mtgl2	21.0%	98.1%	97.2%	97.0%
nd_netgen-3000-1-1-b-b-ns_7	31.10	38.65	42.30	52.25
netmod_kar1	5.51	4.09	3.11	5.16
netmod_kar2	5.77	4.08	5.13	6.19
nous1	7.70	26.07	16.07	16.13
nous2	0.80	1.05	2.07	1.12
nvs02	0.06	1.04	1.06	1.09
nvs06	0.02	1.04	1.06	1.11
oil2	4.64	5.23	5.25	7.34
optmass	1.8%	9.6%	7.1%	7.1%
ortez	0.14	1.06	1.07	1.14
p_ball_10b_5p_3d_m	4.26	5.08	5.09	4.17
p_ball_15b_5p_2d_m	5.44	6.11	5.10	6.17
parabol5_2_3	0.051%	∞	∞	∞
parallel	10.26	7.07	8.09	8.16
pedigree_ex485	105.90	95.32	95.39	153.99
pedigree_ex485_2	35.24	76.58	20.67	19.54
pointpack06	3.05	4.06	3.07	3.11
pointpack08	61.22	25.05	12.07	9.14
pooling_epa1	12.87	29.09	49.12	30.19
pooling_epa2	1.9%	1.8%	1.8%	1.6%
portfol_buyin	0.50	1.06	1.07	1.12
portfol_card	0.58	1.07	1.07	1.11
powerflow0014r	7079.42	0.36%	0.47%	0.59%
powerflow0057r	80.19	57.15	59.19	61.29
prob07	76.05	11.05	6.07	706.26
process	0.4%	3.05	3.07	3.11
procurement1mot	85.0%	90.8%	90.8%	90.7%
procurement2mot	2.03	2.19	2.14	2.22
product	16.80	31.11	11.13	13.14
product2	3.85	4.45	5.48	6.88
prolog	0.06	abort	abort	abort
qp3	14.1%	100%	100%	91.0%
qspp_0_10_0_1_10_1	180.85	133.44	80.79	75.94
qspp_0_11_0_1_10_1	623.76	208.02	154.94	115.04
radar-2000-10-a-6_lat_7	129.82	148.80	162.20	185.15
radar-3000-10-a-8_lat_7	422.76	0.37%	0.026%	68.0%
ravempb	3.09	2.06	2.08	2.11
risk2bpb	infeas	infeas	infeas	infeas
routingdelay_bigm	nonopt	6439.89	1.1%	2.8%
rsyn0815m	1.44	2.06	3.04	3.14
rsyn0815m03m	14.06	81.16	14.17	12.19
sfacloc2_2_95	2.78	2.06	2.09	2.14
sfacloc2_3_90	20.32	62.08	28.09	22.16
sjup2	469.64	107.74	114.44	146.11
slay06m	1.85	1.06	1.08	1.13
slay07m	2.05	1.05	1.08	1.14
smallinvDAXr1b010-011	2.17	1.08	1.10	1.14
smallinvDAXr1b020-022	3.37	1.07	1.09	1.15
sonet17v4	949.25	829.28	807.35	777.52
sonet18v6	1445.90	906.34	815.39	816.65
sonetgr17	1255.92	651.90	320.18	368.98
spectra2	9.20	3.10	3.14	3.23
sporttournament24	45.75	71.17	58.16	69.26
sporttournament30	0.72%	4%	3.4%	1.5%
sssd12-05persp	4.5%	19.6%	2774.49	256.15
sssd18-06persp	15.1%	47.7%	46.1%	47.6%
st_testgr1	0.12	1.05	1.07	1.10
st_testgr3	0.24	1.05	1.08	1.10
steenbrf	12.42	6.08	5.09	7.14
stockcycle	0.79	1.25	2.15	1.23
supplychainp1_022020	1056.69	32.8%	31.2%	26.7%
supplychainp1_030510	3.54	3.11	3.13	4.19
supplychainr1_022020	21.97	80.80	14.68	2.81
supplychainr1_030510	0.12	1.09	1.11	1.17
syn15m04m	2.75	1.10	1.12	1.18

instance	1 thread	4 threads	8 threads	16 threads
syn30m02m	1.58	1.09	1.10	2.16
synheat	infeas	0.83%	1.8%	148.14
tanksize	4.12	3.07	3.07	3.12
telecomsp_pacbell	0.63%	6523.90	6267.26	0.45%
tln5	0.50	7.05	5.08	10.13
tln7	188.05	69.5%	30.4%	66.4%
tls2	0.93	1.07	1.07	1.10
tls4	41.19	18.06	24.09	13.14
topopt-mbb_60x40_50	∞	∞	∞	∞
toroidal2g20_5555	5.70	6.27	6.40	6.33
toroidal3g7_6666	141.64	142.26	157.46	169.61
transswitch0009r	6975.38	2602.29	1289.17	959.22
tricp	100.0%	100.0%	100.0%	100.0%
tspn08	17.8%	31.2%	18.1%	18.6%
tspn15	50.0%	62.5%	27.4%	29.4%
unitcommit1	infeas	2.38	4.58	6.25
unitcommit2	439.70	3.81	5.85	6.92
wager	3.46	2.14	1.16	2.22
waste	38.85	57.1%	55.4%	55.0%
wastepaper3	4.22	6.07	5.09	5.14
wastepaper4	167.16	abort	abort	40.14
wastepaper6	0.023%	0.04%	0.034%	7200.00
water4	nonopt	nonopt	nonopt	3473.45
waternd1	7.21	10.06	4.08	5.13
waterno2_02	2.89	2.10	2.11	1.19
waterno2_03	nonopt	33.13	nonopt	39.3%
waterund01	1525.91	1.5%	1.5%	abort