



Sparsity-Driven Aggregation of Mixed Integer Programs

Liding Xu¹  

Zuse Institute Berlin, Germany

Gioni Mexi  

Zuse Institute Berlin, Germany

Ksenia Bestuzheva  

Zuse Institute Berlin, Germany

Abstract

Cutting planes are crucial for the performance of branch-and-cut algorithms for solving mixed-integer programming (MIP) problems, and linear row aggregation has been successfully applied to better leverage the potential of several major families of MIP cutting planes. This paper formulates the problem of finding good quality aggregations as an ℓ_0 -norm minimization problem and employs a combination of the lasso method and iterative reweighting to efficiently find sparse solutions corresponding to good aggregations. A comparative analysis of the proposed algorithm and the state-of-the-art greedy heuristic approach is presented, showing that the greedy heuristic implements a stepwise selection algorithm for the ℓ_0 -norm minimization problem. Further, we present an example where our approach succeeds, whereas the standard heuristic fails to find an aggregation with desired properties. The algorithm is implemented within the constraint integer programming solver SCIP, and computational experiments on the MIPLIB 2017 benchmark show that although the algorithm leads to slowdowns on relatively “easier” instances, our aggregation approach decreases the mean running time on a subset of challenging instances and leads to smaller branch-and-bound trees.

2012 ACM Subject Classification Mathematics of computing → Solvers; Mathematics of computing → Mixed discrete-continuous optimization; Mathematics of computing → Linear programming

Keywords and phrases mixed integer linear programming, cutting plane, valid inequality, separation, aggregation, projection, sparse optimization

Digital Object Identifier 10.4230/LIPIcs.SEA.2025.27

Supplementary Material *Software (Source code)*: https://github.com/lidingxu/aggregation_lp [45]; archived at `swb:1:dir:a7313b858eba5b70d7d70e5041dad8041700ae60`

Funding The work for this article has been conducted in the Research Campus MODAL funded by the German Federal Ministry of Education and Research (BMBF) (fund numbers 05M14ZAM, 05M20ZBM).

Acknowledgements We want to thank Felipe Serrano for his insightful comments on this paper.

1 Introduction

Mixed integer programs (MIPs) are an important class of mathematical programs combining continuous and discrete decision variables with linear objective functions and constraints. The branch-and-cut algorithm is the backbone technique used in many MIP solvers, e.g., SCIP [10]. A critical component of the branch-and-cut algorithm’s success is the use of cutting planes. There are rich classes of valid inequalities for the mixed-integer linear program (MILP):

¹ corresponding author



$$\min\{c^\top x \mid x \in \mathbb{Q}^p \times \mathbb{Z}^q, Ax - b \leq 0\}, \quad (1)$$

where \mathbb{Q} denotes the set of rationals, \mathbb{Z} denotes the set of integers, $A \in \mathbb{Q}^{m \times (p+q)}$, and $b \in \mathbb{Q}^m$.

In this paper, we investigate the common aggregation procedure [37] used in the generation of valid inequalities for MILPs. We formulate the aggregation problem as a cardinality optimization problem, and we exploit the sparse optimization technique called least absolute shrinkage and selection operator (lasso) [42]. Our primary contribution is the development of a sparsity-driven and optimization-based aggregation algorithm tailored for MILPs. This allows us to derive more effective complemented Mixed-Integer-Rounding (c-MIR) cuts [36, 40] than the previous heuristic [37].

We next review the aggregation procedure for MILPs. Aggregation constructs a single-row surrogate relaxation, defined by the so-called base inequality, in the form: $\min\{c^\top x \mid x \in \mathbb{Q}^p \times \mathbb{Z}^q, \lambda^\top (Ax - b) \leq 0\}$, where λ is a vector of non-negative factors. This surrogate relaxation simplifies the structure of the original problem, facilitating the derivation of valid inequalities. An illustrative example is the Chvátal-Gomory (CG) cut [20, 25]. For pure integer linear programs (ILPs), CG cuts take the form of $\lfloor \lambda^\top A \rfloor x - \lfloor \lambda^\top b \rfloor \leq 0$ with $\lambda \in \mathbb{Q}_+^m$, and these cuts are a specific type of split inequality [21], inferred from the aggregated single-row surrogate relaxation. Subclasses of CG cuts include zero-half cuts [15], where factors of λ are restricted to zero or one-half, and mod- k cuts [16], with factors of λ in $\{0, 1/k, \dots, k - 1/k\}$. For general MILPs, projected CG cuts can be used. Similar to the principles of Benders decomposition [7], Bonami et al. [11] proposed projecting out continuous variables from MILPs to reduce them to ILPs, i.e., the first p elements of $\lambda^\top A$ corresponding integer variables are zeros, then CG cuts can be applied effectively.

In addition to split cuts, Mixed-Integer-Rounding (MIR) cuts [40] form another fundamental class of valid inequalities, particularly relevant for single-row MIP models involving continuous variables. While MIR cuts are mathematically equivalent to Gomory's mixed-integer cuts [28], split cuts, and a special case of disjunctive cuts [4], their separation procedures differ significantly, particularly when combined with aggregation, resulting in distinct implementations in optimization solvers. Aggregation is also used for generating complemented MIR (c-MIR) cuts, as introduced in the heuristic by Marchand and Wolsey (MW) [36, 37]. The MW heuristic has been shown to be effective in the computational studies realized in SCIP [1, 44], MOPS [19], and CPLEX [29]. The MW heuristic uses a stepwise selection process to create a sparse base inequality that has as few continuous variables as possible, i.e., some continuous variables are greedily projected out. For MILP instances with network structures, such as network design problems, a tailored heuristic [2] has been proposed to find aggregations corresponding to sparse cut sets in networks, and the resulting c-MIR cuts are the well-known flow-cutset inequalities [3, 9]. Note that the technique of projecting out continuous variables via row aggregations before applying c-MIR cuts is also crucial for enabling other components, such as cut-based conflict analysis [38], to generate conflict constraints in the presence of continuous variables.

Our innovation is centered on the use of sparse optimization to find aggregations. The corresponding sparse optimization problem can be formulated as a cardinality optimization problem [43] involving the nonconvex and non-differentiable ℓ_0 -norm, which is a hard combinatorial optimization problem. As a practical alternative to ℓ_0 -norm minimization, we employ approximations using sparsity-inducing convex norms. In the case of the ℓ_1 -norm, the approximation problem is lasso [42], which can be solved efficiently by linear

programming (LP). Under specific conditions, the lasso's solutions can exactly recover the original ℓ_0 -norm solution [12, 18]. We refer to [5, 6, 23, 26, 31] for other (approximations of) optimization-based separators, some of which have particular emphasis on sparsity.

The remaining part of this paper is organized as follows. In Sect. 2, we introduce the notation and the definition of c-MIR cuts. In Sect. 3, we revise the MW heuristic and its implementation, and we show its limitations in an example. In Sect. 4, we provide the cardinality optimization formulation of the aggregation, we show that the MW aggregation heuristic is a stepwise selection heuristic, which is a relatively ancient approach in statistical community [24, 32], and we introduce the lasso approximation of the cardinality optimization problem. In Sect. 5, using the lasso approximation, we propose an aggregation algorithm that solves several linear programs (LPs) to pursue a sparse aggregated base inequality and project out integer variables. In Sect. 6, we present the experiment results. In Sect. 7, we conclude this paper with the future research prospect. This perspective not only sheds light on the structure of the aggregation process but also provides a foundation for our further refinements.

2 Notation and c-MIR cuts

In this section, we provide the necessary notation and the definition of c-MIR cuts [37]. Let $[t_1 : t_2]$ denote the sequence (t_1, \dots, t_2) , and let $[t]$ be an abbreviation for $[1 : t]$. Denote $x^+ = \max(x, 0)$. Let $\|\cdot\|_0$ and $\|\cdot\|_1$ denote the ℓ_0 -norm and ℓ_1 -norm, respectively. Given a matrix $A \in \mathbb{Q}^{m \times n}$, $A_{i,j}$ is the entry at row index $i \in [m]$ and column index $j \in [n]$. To index its submatrices, i and/or j can be replaced with index sets $I \subseteq [m]$ and/or $J \subseteq [n]$. Additionally, when using a matrix index, the symbol $:$ denotes the full index set $[m]$ or $[n]$, depending on the context.

Next, we define the c-MIR cut for the following mixed knapsack set:

$$X = \left\{ (s, z) \in \mathbb{Q}_+ \times \mathbb{Z}_+^q : \sum_{j \in [q]} a_j z_j \leq b + s, z_j \leq u_j \text{ for } j \in [q] \right\}. \quad (2)$$

Let (T, U) be a partition of $[q]$ and let the scaling parameter $\delta > 0$. We can complement variables in U and scale the constraint $\sum_{j \in [q]} a_j z_j \leq b + s$ by dividing it by δ . This results in the single-row relaxation:

$$X' = \left\{ (s, z) \in \mathbb{Q}_+ \times \mathbb{Z}_+^q : \sum_{j \in T} \frac{a_j}{\delta} z_j + \sum_{j \in U} \frac{-a_j}{\delta} (u_j - z_j) \leq \frac{b - \sum_{j \in U} a_j u_j}{\delta} + \frac{s}{\delta} \right\}. \quad (3)$$

Then, the c-MIR inequality derived from X' is as follows:

$$\sum_{j \in T} G\left(\frac{a_j}{\delta}\right) z_j + \sum_{j \in U} G\left(\frac{-a_j}{\delta}\right) (u_j - z_j) \leq \lfloor \beta \rfloor + \frac{s}{\delta(1-f)}, \quad (4)$$

where $\beta = (b - \sum_{j \in U} a_j u_j)/\delta$, $f = \beta - \lfloor \beta \rfloor$, and $G(d) = (\lfloor d \rfloor + (f_d - f)^+ / (1 - f))$ with $f_d = d - \lfloor d \rfloor$. The c-MIR inequality is valid for X' (hence for X), and it becomes the so-called MIR inequality when $U = \emptyset$ and $\delta = 1$ [37, 40]. The c-MIR inequalities can represent mixed cover inequalities, residual capacity inequalities, and weight inequalities.

3 The Marchand-Wolsey Heuristic

In this section, we discuss the MW heuristic [37], a procedure designed to generate c-MIR inequalities, and several enhancements of its implementations [19, 44].

As we have mentioned, continuous variables are generally not preferred in the base inequality, since a base inequality consisting solely of continuous variables does not yield additional useful inequalities. Although continuous variables can be included when deriving c-MIR cuts (4), it is often beneficial to project out some continuous variables to enhance the effectiveness of the cuts. If we enforce the use of a row in the aggregated base inequality, since MILP instances are usually sparse, then there may be only a few nonzero continuous variables in the enforced row to eliminate. Moreover, the complete elimination of continuous variables is not necessary. This characteristic marks a key distinction from Benders decomposition.

The bound distance is a measure used in [37, 44] to determine “bad” continuous variables that will be preferred for elimination. To illustrate this, we divide the linear constraints $Ax - b \leq 0$ of the MILP (1) into two parts: the bound constraint part and the normal constraint part. For brevity, we denote $p + q$ by n . The bound constraints, which include variable bounds and implied variable bounds, are expressed in the set

$$X_1 = \{x \in \mathbb{Q}^n : \forall j \in [n], x_j \leq u_j, \text{ and } \forall j \in [p], j' \in J_j, x_j \leq u_{jj'} + d_{jj'}x_{j'}\}, \quad (5)$$

where J_j is a subset of $[p+1 : p+q]$. For the sake of simplicity, we only consider upper bounds on variables, and our results are applicable for variables with lower bounds as well. The remaining normal constraints define the set $X_2 = \{x \in \mathbb{Q}^n, A'x - b' \leq 0\}$, where $A' = A_{\bar{I}}$, and $b' = b_{\bar{I}}$ with $\bar{I} \subseteq I$ being the index set for the remaining constraints. Thus, the MILP (1) reads as $\min_{x \in \mathbb{Q}^p \times \mathbb{Z}^q, x \in X_1 \cap X_2} c^\top x$. Given a point \tilde{x} to separate, for $j \in [p]$, the bound distance of \tilde{x}_j is defined as

$$\text{bdd}(\tilde{x}_j) = \min(\{u_j\} \cup \{u_{jj'} + d_{jj'}\tilde{x}_{j'}\}_{j' \in J_j}) - \tilde{x}_j. \quad (6)$$


In the heuristic, eliminating continuous variables with large bound distances is preferred, which we will explain later. A continuous variable with a nonzero bound distance is called a “bad variable”. We next revise the three basic steps of the MW heuristic, which was presented at a high level in [37]:

1. **Normal constraint aggregation:** The MW heuristic identifies bad continuous variables, then it employs a stepwise selection algorithm to determine a few nonzero factors $\lambda_{\bar{I}}^\top$. The aggregated inequality $\lambda_{\bar{I}}^\top (A'x - b') \leq 0$ should be sparse, and some bad continuous variables are projected out.
2. **Bound substitution:** Each remaining continuous variable x_j in the aggregated inequality is substituted by an expression of its bound and a nonnegative continuous slack variable y_j : $x_j = u_j - y_j$ or $x_j = u_{jj'} + d_{jj'}z_{j'} - y_j$. This yields the inequality defining the mixed knapsack set X , where s in (2) is taken as the weighted sum of all continuous variables y , and z in (2) contains all integer variables present in the base inequality.
3. **Separation:** The parameters (T, U) and δ are determined, and the associated c-MIR inequality (4) is generated.

The MW heuristic’s preference for eliminating continuous variables with large bound distances is explained as follows. We can reduce the set X' in (3) to a simple two-variable mixed-integer set $\{(\sigma, \zeta) \in \mathbb{Q}_+ \times \mathbb{Z} : \zeta \leq \beta + \sigma\}$, and the c-MIR inequality reduces to the split inequality $\zeta \leq \lfloor \beta \rfloor + \sigma / (1 - (\beta - \lfloor \beta \rfloor))$ valid for this set [22, 37]. See Fig. 1 for an example. The gray area indicates the set of nonnegative points satisfying $\zeta \leq \beta + \sigma$, and the vertical lines represent sets of integer feasible solutions. The line passing through $(0, 1)$ and $(2, 0.6)$ shows the boundary of the set of solutions satisfying the split inequality.

Observe that feasible solutions of the LP relaxation that are cut off by the split inequality have relatively small values of σ . Further, the smaller the value of σ , the larger the corresponding interval that is cut off. In the extreme case, consider a point $(\tilde{\sigma}, \tilde{\zeta})$. If $\tilde{\sigma}$

is zero (i.e., at its bound) and the base inequality $\zeta \leq \beta + \sigma$ has zero slack (i.e., $\tilde{\zeta} = \tilde{\beta}$), then the point is cut off by the split inequality. An aggregation that eliminates continuous variables with large bound distances, followed by the substitution of remaining continuous variables using their bounds, represents a shift into a space where the remaining nonnegative continuous variable (σ in the example above) tends to be close to 0. This results in a higher likelihood that a strongly violated MIR inequality will be found.



■ **Figure 1** Example of a split cut.

The MW heuristic differentiates between normal constraints and variable bound constraints. However, the second bound substitution step can be interpreted as a specialized constraint aggregation procedure applied to variable bound constraints, where the absolute values of the corresponding factors are restricted to either zero or one by the MW heuristic. We will show that, in Sect. 5, our aggregation algorithm can unify the aggregation of two types of constraints, so the algorithm has no restrictions on factors for bounds.

SCIP's implementation of the aggregation algorithm includes several enhancements, such as preprocessing and the use of row scores. Normal constraints' rows containing bad variables are identified as "useful rows" (that is, useful for eliminating bad variables). The preprocessing step reduces the LP's rows to these useful rows, effectively pruning the search space of λ . Row scores are used to select rows to combine. The score of a row is a weighted sum of its dual value, density, slack, fractionalities of its integer variables, and "fractionalities" of its continuous variables defined through slacks of implied variable bounds.

■ **Algorithm 1** The preprocessing algorithm.

Require: the LP data (A, b) .

- 1: compute bound distances of continuous variables;
- 2: identify the index set J of bad variables and the index set I of useful rows;
- 3: compute scores of useful rows;
- 4: sort the indices of bad variables in decreasing order of bound distance;
- 5: sort the indices of useful rows in decreasing order of score;

The whole preprocessing algorithm is presented in Algorithm 1. Finally, we outline SCIP's implementation of the MW aggregation heuristic in Algorithm 2.

The MW heuristic cannot adjust the factor λ_i after the row r_i is added into the aggregated row. Consequently, it cannot explore all possible combinations of rows. The following example illustrates a limitation of the MW heuristic, where it fails to eliminate all bad variables, even if there exists such an aggregation.

■ **Algorithm 2** The MW aggregation heuristic.

Require: relaxation solution \tilde{x} , the LP data (A, b) , the starting row index i_0 , and the maximum round MAXAGGR of aggregations.

- 1: call the preprocessing Algorithm 1;
- 2: in the aggregated constraint $\alpha x - \beta \leq 0$, set $\alpha \leftarrow A_{i_0,:}$ and $\beta \leftarrow b_{i_0}$
- 3: set the aggregation counter $c \leftarrow 0$;
- 4: set the set I' of used rows to $\{i_0\}$;
- 5: set the set J' of eliminated bad variables empty;
- 6: **for** each sorted index $j \in J \setminus J'$ of bad variables **do**
- 7: **for** each sorted index $i \in I \setminus I'$ of useful rows for which $A_{ij} \neq 0$ **do**
- 8: find the factor λ_i such that the coefficient vector $\alpha^{\text{new}} = \alpha + \lambda_i A_{i,:}$ has $\alpha_j^{\text{new}} = 0$;
- 9: **if** there exists $j \in J'$ such that $\alpha_j^{\text{new}} \neq 0$ **then**
- 10: **continue** ; \triangleright do not add rows that reintroduce eliminated bad variables
- 11: **end if**
- 12: update: $\alpha \leftarrow \alpha^{\text{new}}$, $\beta \leftarrow \beta + \lambda_i b_i$, $I' \leftarrow I' \cup \{i\}$, $J' \leftarrow J' \cup \{j\}$, and $c \leftarrow c + 1$;
- 13: try bound substitution and separation on the aggregated constraint $\alpha^\top x - \beta \leq 0$;
- 14: **if** $c < \text{MAXAGGR}$ **then**
- 15: **break** ; \triangleright proceed to next bad variable
- 16: **else**
- 17: **return** ;
- 18: **end if**
- 19: **end for**
- 20: **end for**

► **Example 1.** Consider the MILP's constraint $Ax \leq b, x \geq 0$ as follows:

$$1/3x_1 + x_2 - 2/3x_3 \leq 1, \tag{7a}$$

$$2/3x_1 - 1/3x_2 - 4/3x_3 + x_4 \leq 1, \tag{7b}$$

$$-1/3x_2 + x_3 \leq 1, \tag{7c}$$

$$x \geq 0, \tag{7d}$$

where all rows are useful, x_2 and x_3 are bad continuous variables, x_1 and x_4 are integer variables. Assume that the objective of the MILP is $\min \sum_{j \in [4]} -x_j$. Thus, the solution to the LP relaxation is $(0, 2.143, 1.714, 4)$, and the slacks of constraints are all zeros. A vector $\lambda = (1, 1, 2)$ exists such that the corresponding aggregated inequality $x_1 + x_4 \leq 4$ eliminates all bad variables. However, the MW heuristic cannot find a projected inequality starting from any row. For example, starting from the first row (7a), assume that the MW heuristic chooses the second row (7b) to eliminate x_2 . In this case, the heuristic greedily selects the factor 3 for (7b) and does not select any factor for (7c), then we obtain the aggregated inequality $7/3x_1 - 14/3x_3 + 3x_4 \leq 4$ using $\lambda = (1, 3, 0)$.

Hence, a systematic method is required to explore the space of λ , rather than relying on the stepwise activation and fixing of nonzero factors. In the following sections of this paper, we propose an optimization-based approach to address this need.

4 Cardinality Formulation of the Aggregation Problem

In this section, we formulate the aggregation problem as a cardinality optimization problem, more precisely a cosparsity problem [39]. We cast the problem as the standard weighted regularized cardinality optimization problem. We show that the MW aggregation heuristic is a stepwise selection heuristic, which was used to solve the ℓ_0 minimization problem [32]. Then, we present the lasso approximation of the cardinality optimization problem.

We consider the following cardinality optimization problem:

$$\min_{\lambda_I \in \mathbb{Q}_+^I, \lambda_{i_0} \geq 1} \|W \lambda_I^\top A_{I,J}\|_0, \quad (8)$$

where $A_{I,J}$ is the submatrix of A with useful rows indexed by I and bad variables indexed by J , $W = \text{Diag}(w)$ is the diagonal weight matrix with diagonal entries $w = (w_j)_{j \in J} = (\text{bdd}(\tilde{x}_j))_{j \in J}$ (see (6)). Note that we enforce the starting row r_{i_0} by setting $\lambda_{i_0} \geq 1$ (otherwise, the minimum value is zero). To separate a sufficient number of cuts, we will generate multiple aggregations in a single separation round of cuts, each of which is obtained by approximately solving a cardinality optimization problem enforcing a different starting row r_{i_0} . The objective in (8) is a weighted sum of nonzeros in the aggregated row:

$$\|W \lambda_I^\top A_{I,J}\|_0 = \sum_{j \in J} w_j \mathbf{1}\left\{\left(\sum_{i \in I} A_{i,j} \lambda_i\right) \neq 0\right\}. \quad (9)$$

The problem (8) is a cosparsity problem, which presumes that the linear transformation $\lambda_I^\top A_{I,J}$ of λ_I is sparse. The problem is nonconvex and is known to be \mathcal{NP} -hard [43].

Moreover, we want the aggregated base inequality to be sparse, so this translates to a requirement of sparse factors. Thus, we add ℓ_0 regularization term on λ_I to the objective, and we obtain the following regularized cosparsity formulation of the aggregation problem:

$$\min_{\lambda_I \in \mathbb{Q}_+^I, \lambda_{i_0} \geq 1} \|W \lambda_I^\top A_{I,J}\|_0 + \|W' \lambda_I\|_0, \quad (10)$$

where $W' = \text{Diag}(w')$ is a diagonal weight matrix, and the diagonal entries w'_i ($i \in I$) is a user-defined penalty. For example, w'_i could be the score of row r_i .

We next cast the problem as the standard weighted regularized cardinality optimization problem. Let us introduce variable μ representing $\lambda_I^\top A_{I,J}$, and this leads to the weighted version of the ℓ_0 minimization problem:

$$\min_{\lambda_I \in \mathbb{Q}_+^I, \lambda_{i_0} \geq 1, \lambda_I^\top A_{I,J} = \mu} \|W \mu\|_0 + \|W' \lambda_I\|_0. \quad (11)$$

Forward stepwise selection [24, 32] and the lasso [42, 18] (aka ℓ_1 -norm approximation) are two popular methods for approximately solving the ℓ_0 minimization problem (11). Forward stepwise selection can be seen as a “greedy” algorithm, updating the active set by one variable at each step, instead of re-optimizing over all possible subsets of a given size; in turn, the lasso can be seen as a more “democratic” version of forward stepwise, updating the coefficients to maintain equal absolute correlation of all active variables with the objective.

The MW heuristic can be viewed as a forward stepwise selection algorithm for the ℓ_0 formulation (10) of the aggregation problem. In the statistical community, there are several known drawbacks of the forward stepwise selection heuristic. The number of the selected candidate variables may be smaller than the total number of final model variables [33].

Models that are created may be over-simplifications of the real models of the data [41]. These drawbacks imply that the MW heuristic cannot find complicated combinations of rows, which was already demonstrated in Example 1.

We adopt the lasso approach to tackle the cardinality formulation (10), and this yields our new heuristic to solve the aggregation problem. As the ℓ_1 -norm promotes sparsity, the lasso approximation uses the ℓ_1 -norm to approximate the ℓ_0 -norm in (11):

$$\min_{\lambda_I \in \mathbb{Q}_+^I, \lambda_{i_0} \geq 1, \lambda_I^\top A_{I,J} = \mu} \|W\mu\|_1 + \|W'\lambda_I\|_1. \quad (12)$$

In our setting, we let the penalty w'_i be the nonnegative slack $b_i - A_{i,:}^\top \tilde{x}$ of the row. This penalty and ℓ_1 approximation promote zero slacks, i.e., tight rows. This casts $\|W'\lambda_I\|_1$ as $\sum_{i \in I} \lambda_i (b_i - A_{i,:}^\top \tilde{x})$, and the lasso approximation (12) becomes:

$$\min_{\lambda_I \in \mathbb{Q}_+^I, \lambda_{i_0} \geq 1, \lambda_I^\top A_{I,J} = \mu} \|W\mu\|_1 + \sum_{i \in I} \lambda_i (b_i - A_{i,:}^\top \tilde{x}). \quad (13)$$

Note that the term $\sum_{i \in I} \lambda_i (b_i - A_{i,:}^\top \tilde{x})$ is precisely the slack of the aggregated base inequality.

Thus, viewing instead $\|W\lambda_I^\top A_{I,J}\|_1$ as the regularization term, we reach another interpretation of the lasso approximation: the ℓ_1 approximation aims at finding a sparse aggregated base inequality with a small slack, where the sparsity means that the (weighted) sum $\|W\mu\|_0$ of the number of non-zero bad variables is small. The lasso is a convex optimization problem, and it can be converted to an LP by introducing some extra variables modeling the ℓ_1 norm, so we can solve it to optimality using the off-the-shelf LP solvers. Note that we can convert an equality row $\alpha^\top x = \beta$ into two inequality rows $\alpha^\top x - \beta \leq 0$ and $-\alpha^\top x + \beta \leq 0$ both of zero slacks. An optimal basic solution to (3) usually has at most one non-zero factor for these two rows.

5 An Optimization-Driven Aggregation Algorithm

In this section, we present an LP-based algorithm to solve the aggregation problem. The algorithm solves several LPs to promote the sparsity of the solution.

Our aggregation algorithm can unify the normal constraint aggregation and bound substitution steps of the MW heuristic. To this end, we modify the preprocessing algorithm by treating variable bound constraints as normal constraints, i.e., useful rows can include variable bounds. Thus, the index set I of useful rows is a subset of $[m]$ (not \bar{I}). This exploits more possibilities to project out bad variables. We next present the main aggregation algorithm.

In the first factor search stage, our algorithm solves the LP (13) to find a λ_I . Denote by $I' = \{i \in I : \lambda_i > 0\}$ the index set of active factors. Then, the aggregated base inequality is $\lambda_{I'} A_{I',J} x - \lambda_{I'} b_{I'} \leq 0$. As lasso is sparse inducing, the nonzero factors in the solution λ_I are usually quite sparse. This implies that the aggregated base inequality is sparse. Moreover, the sparsity of $\lambda_{I'}^\top A_{I',J}$ is also promoted by the LP, and some bad variables are eliminated in the aggregated base inequality.

If $\lambda_{I'}^\top A_{I',J}$ is not sparse enough, i.e., the density of bad variables exceeds a threshold, then our algorithm moves to the second stage, wherein it applies iterative reweighting. The iterative reweighted algorithm [14] performs exceptionally well in locating sparse solutions to ℓ_0 minimization, and it even outperforms ℓ_1 approximation. The key feature of the reweighted algorithm is the solution of a series of reweighted ℓ_1 approximation problems. In this stage, we use the iterative reweighted algorithm to construct and solve several modified LPs to find a sparser λ supported on I' . As the active set I' is already known, we are only interested in projecting out the bad variables, so our new LP has only the second term in its objective:

$$\min_{\lambda_{I'} \in \mathbb{Q}_+^I, \lambda_{i_0} \geq 1} \|W \lambda_{I'}^\top A_{I',J}\|_1. \quad (14)$$

In each round of iterative reweighting, the algorithm dynamically adjusts the weight W and resolves the modified LP (14). The weights are updated with the use of the previous solution λ to (14):

$$w_i = \begin{cases} w_i/(\epsilon + \lambda_i), & \lambda_i \neq 0, \\ 0, & \lambda_i = 0, \end{cases} \quad (15)$$

where ϵ is a small positive constant introduced to avoid numeric overflow. We solve the iteratively reweighted LPs (14) until $\lambda_{I'}^\top A_{I',J}$ is sparse enough.

We next present our sparsity-driven LP-based aggregation algorithm in Algorithm 3.

■ **Algorithm 3** Sparsity-driven LP-based aggregation algorithm.

Require: relaxation solution \tilde{x} , the LP data (A, b) , the starting row index i_0 , the maximum round MAXAGGR of aggregations, and the density threshold η .

```

1: call the preprocessing Algorithm 1;
2: solve the lasso approximation (13);
3: set the set  $I'$  of used rows as support of nonzeros in  $\lambda_{I'}$ ;
4: set the aggregation counter  $c \leftarrow 0$ ;
5: while  $c$  is less than MAXAGGR do
6:   let  $a = \lambda_{I'}^\top A_{I',:}$ ;
7:   try separation on the aggregated base inequality  $a^\top x - \lambda_{I'}^\top b_{I'} \leq 0$ ;
8:   if the density  $|\{j \in J : a_j \neq 0\}|/|J|$  of bad variables is greater than the threshold  $\eta$ 
       then
9:     let  $c \leftarrow c + 1$ ;
10:    update the weight  $w$  as in (15);
11:    solve the reweighted LP (14);
12:   else
13:     return ;
14:   end if
15: end while

```

► **Example 2.** We revise Example 1 and apply the LP-based aggregation algorithm instead of the MW heuristic. Recall that the slacks of all constraints in (7) are zeros, and x_2 and x_3 are bad variables. We can express the corresponding lasso approximation as:

$$\min_{\lambda \in \mathbb{Q}_+^3, \lambda_{i_0} \geq 1} w_2 |\lambda_1 - 1/3\lambda_2 - 1/3\lambda_3| + w_3 |-2/3\lambda_1 - 4/3\lambda_2 + \lambda_3|, \quad (16)$$

where w_2 and w_3 are non-negative weights (set to be equal to bound distances). Note that the term $\sum_{i \in I} \lambda_i^\top (b_i - A_{i,J} \tilde{x}_J)$ is null due to the slacks being zero. The solution $(1, 1, 2)$ is optimal to the lasso with objective value zero regardless of values of w_2, w_3 , and this solution eliminates all bad variables. Thus, the LP approach finds a suitable aggregation, and it is stable under the perturbation of weights.

6 Experimental Comparison

In this section, we evaluate the performance of SCIP on MIP problems and compare it with the performance achieved with the addition of our new aggregation-based cut separator.

6.1 Implementation Details

The Algorithm 2 and Algorithm 3 have been implemented in the open-source MIP solver SCIP [10]. At each node of the branch-and-bound tree, several rounds of LPs are solved, and separators are called to strengthen the LP relaxations. SCIP has the default aggregation cut separator based on the MW heuristic, and we implement the new aggregation cut separator based on Algorithm 3. In a single separation round, the cut separators generate multiple aggregations using different starting rows and try separations on the aggregated base inequalities. The generation of flow cover inequalities [30], knapsack cover inequalities [34], and c-MIR inequalities on an aggregated base inequality is tried, with the first two types of inequalities being concrete models of c-MIR inequalities.

There are two main differences between our aggregation-based cut separator and the default one, which result in our separator generating fewer cuts:

- The first difference is by design to prevent the generation of repetitive cuts. Different starting rows in Algorithms 2 and 3 may lead to the same aggregated base inequality. To avoid redundancy and conserve computational resources, our separator only utilizes a starting row if it has not been used in previous aggregations.
- The second difference does not directly involve aggregation. Unlike our method, the default separator attempts to separate cuts on every row of a MILP instance, even if the instance is purely an ILP. For example, the default separator might separate cuts for a knapsack constraint.

The two separators share the same parameter settings, defined as follows along with their respective values:

- SEPA_FREQ (10): the frequency defines the depth levels at which the separation is called. A frequency of 10 means that the separation is executed in depth 0, 10, 20, ... of the branching tree.
- MAXAGGR (6 at root node, 3 at non-root nodes): the maximum round of aggregations, see Algorithms 2 and 3.

For the cut separator based on our new Algorithm 3, we limit the numbers of bad variables and useful rows to 50 and 5000, respectively. During a separation round, we will solve multiple aggregation LPs and separate cuts on the aggregated base inequality. The aggregation LPs are solved using the primal simplex algorithm of CPLEX, so that our algorithm can benefit from warm-start after changes of objectives and lower bounds of factor variables of aggregation LPs.

6.2 Results

All experiments were conducted on a cluster equipped with Intel Xeon Gold 5122 CPUs, operating at 3.60 GHz with 96 GB of RAM. Each test run was executed on a single thread with a time limit of two hours. To ensure accurate comparisons despite performance variability [35], we used a large and diverse test set, specifically the MIPLIB 2017 benchmark set [27]. We used random seeds ranging from 0 to 4, resulting in a total of 1200 instance-seed pairs.

Table 1 compares the performance of default SCIP (“SCIP-default”), with SCIP using the new aggregation-based cut separator (“SCIP+aggr.”). It shows the numbers of solved instances, the shifted geometric means of running time, and the number of branch-and-bound

tree nodes. In the computation of geometric means, we use a shift of 1 second and 100 nodes, respectively. Relative differences in time and nodes for each category are shown in the last two columns of the table, with relative differences that exceed 1% highlighted.

We present these values for various subsets of the test instances. The category “all” encompasses all 1200 test instances. “solved-by-both” refers to instances that were solved to optimality by both configurations. The “affected” category highlights instances where performance varied between the two configurations, indicated by a difference in the total number of LP iterations. Further, we compare the performance on different categories of affected instances: instances that both configurations solved in under 100 seconds (“solved-under-100s”), instances that took more than 100 seconds to solve by either configuration (“solved-over-100s”), and instances that only one of the configurations managed to solve (“solved-by-one”).

■ **Table 1** Performance Comparison of SCIP-default with SCIP+aggr.

| Subset | instances | SCIP-default | | | SCIP+aggr | | | relative | |
|-------------------|-----------|--------------|--------|-------|------------|--------|-------|-------------|-------------|
| | | solved | time | nodes | solved | time | nodes | time | nodes |
| all | 1200 | 714 | 936.9 | 6436 | 716 | 939.9 | 6151 | 1.00 | – |
| solved-by-both | 706 | 706 | 226.4 | 2711 | 706 | 228.6 | 2584 | 1.01 | 0.95 |
| affected | 330 | 320 | 287.9 | 5135 | 322 | 289.9 | 4629 | 1.01 | – |
| solved-under-100s | 98 | 98 | 17.7 | 375 | 98 | 20.9 | 320 | 1.18 | 0.85 |
| solved-over-100s | 214 | 214 | 789.4 | 13133 | 214 | 751.6 | 11931 | 0.95 | 0.91 |
| solved-by-one | 18 | 8 | 5476.5 | 39786 | 10 | 4719.1 | 37720 | 0.86 | – |

Across the entire data set, there is minimal difference in time or number of solved instances, with only two additional instances solved by SCIP+aggr. In the category “solved-by-both”, we observe a 5% decrease in the number of nodes. In the category “solved-under-100s”, there is a notable 17% increase in time, despite a 15% reduction in the number of nodes. This increase is expected, as the separation algorithm is expensive for easy-to-solve problems. In contrast, in the category “solved-over-100s,” we observed a 5% decrease in time and a 9% reduction in nodes, which indicates that our cut separator helps solve harder problems more efficiently. To assess the statistical significance in the reduction in the number of nodes from “SCIP-default” to “SCIP+aggr”, we performed a version of the Wilcoxon signed-rank test [8]. The test, applied to instances solved by both models, yielded a p-value below 0.01, indicating a significant difference in the number of nodes between the two settings. For the limited subset of affected instances “solved-by-one” SCIP+aggr was 14% faster.

Finally, to complement our performance analysis, we examine the sparsity of the aggregated rows produced by Algorithm 2 and Algorithm 3. In Table 2, we report the average number of bad columns in an aggregated row, the average of the total number of distinct bad columns across all rows used for aggregation, their ratio, and the average number of rows used for aggregation.

The results indicate that rows aggregated by Algorithm 3 contain significantly fewer bad columns on average – only 0.37 compared to 2.45 by Algorithm 2. This corresponds to 20% of bad columns remaining after aggregation compared to 52% with the MW heuristic. Additionally, we consider the number of rows used in the aggregation, which typically implies the sparsity of the resulting aggregated row, as its number of nonzeros is upper-bounded by the total number of nonzeros of the used rows. The statistics show that our method utilizes fewer rows in aggregation. This reveals that our aggregation process naturally produces sparse cuts, which is a desirable property. Both aggregation algorithms utilize only a small subset of rows for aggregation, with the number being significantly smaller than the total

■ **Table 2** Sparsity of rows aggregated by Algorithm 2 (SCIP-default) and Algorithm 3 (SCIP+aggr).

| | Algorithm 2 | Algorithm 3 |
|----------------|-------------|-------------|
| bad-cols | 2.45 | 0.37 |
| total-bad-cols | 4.72 | 1.87 |
| ratio | 0.52 | 0.20 |
| used-rows | 3.10 | 2.24 |

number of useful rows (note that our separator imposes a limit of 5000 useful rows). This behavior may be linked to the inherent sparsity of instances in the MIPLIB 2017 benchmark. This suggests that one may obtain a potential performance improvement using LP solvers that can effectively exploit sparsity.

All the numbers in Table 2 are small, but the size of lasso approximation (13) can be much larger than these numbers. This indicates that important (non-zero) factors are confined to a low-dimensional subspace of \mathbb{Q}^I . Similar scenarios arise in fields such as signal processing, machine learning, and statistics. Notably, the lasso method has been demonstrated [13] to effectively recover such low-dimensional subspaces, given the sparsity of the model. This further justifies our approach.

7 Conclusion

We developed a new approach for aggregating MILP constraints to generate mixed-integer knapsack sets for which strong c-MIR cutting planes can be constructed. Similarly to the standard greedy heuristic, our algorithm aims to project out as many bad continuous variables as possible within given working limits, the “badness” being defined as slacks of variable bound or implied variable bound constraints. The algorithm approximately solves the optimization problem of finding such aggregation weights that minimize the number of non-zeroes corresponding to bad continuous variables. This problem is formulated as an ℓ_0 norm minimization problem and solved with the use of the lasso technique. If the number of bad variables in the aggregated row still exceeds a threshold, iterative reweighting is performed to project out further bad variables.

We give examples to 1) illustrate the connection between the continuous variable bound distance and the potential for generating a strong rounding cut for a given mixed-integer set, and 2) show that our approach has a greater capacity to project out bad variables than the standard heuristic. A unifying perspective on the standard and new aggregation algorithms is presented, demonstrating that both represent different approaches to approximately solving the ℓ_0 norm minimization problem: in particular, the standard algorithm is a stepwise selection heuristic for this problem.

The algorithm is implemented within the constraint integer programming solver SCIP and tested by running SCIP with the standard and the new algorithms on the MIPLIB 2017 test set. When considering averages over all instances, we observe only a very small impact on mean performance, but dividing the instances into subsets of easier instances where SCIP solves the problem in under 100s with both configurations compared, and challenging instances which take at least 100s with at least one of the configurations, reveals a connection between instance difficulty and performance impact. While the new algorithm results in an 18% increase in time on easier instances, likely due to the increased computational cost of aggregation, it yields a 5% decrease in mean solver running time on challenging instances.

This indicates that a flexible approach to select the right aggregation algorithm based on estimated instance difficulty could leverage the strength of both methods, and could be a subject of future work.

Further work will, on the one hand, delve deeper into the theory of the connection between the properties of mixed-knapsack sets and the quality of c-MIR cutting planes constructed for these sets, and perform a theoretical analysis of the behavior of aggregation algorithms and the quality of the formulations they produce. On the other hand, it will extend these ideas to the mixed-integer nonlinear programming case, starting with conic MIPs which, on the one hand, have a well-defined structure, and, on the other hand, encompass a rich variety of problems beyond MILPs. Further, minimization of the sum of absolute values in (13) is a classical problem in nonsmooth convex optimization. Due to this, we want to find a more efficient algorithm than the primal simplex for faster aggregation, and a possible candidate is the primal-dual hybrid gradient descent algorithm [17]. A more open problem is how to avoid generating repetitive aggregations.

References

- 1 Tobias Achterberg. *Constraint Integer Programming*. PhD thesis, Technische Universität Berlin, 2007. URL: opus4.kobv.de/opus4-zib/frontdoor/index/index/docId/1018.
- 2 Tobias Achterberg and Christian Raack. The MCF-separator: detecting and exploiting multi-commodity flow structures in MIPs. *Mathematical Programming Computation*, 2(2):125–165, 2010. doi:10.1007/s12532-010-0015-3.
- 3 Alper Atamtürk. On capacitated network design cut-set polyhedra. *Mathematical Programming*, 92:425–437, 2002. doi:10.1007/s101070100284.
- 4 Egon Balas. Disjunctive programming. *Annals of Discrete Mathematics*, 5:3–51, 1979. doi:10.1016/S0167-5060(08)70342-X.
- 5 Egon Balas and Anureet Saxena. Optimizing over the split closure. *Mathematical Programming*, 113(2):219–240, 2008. doi:10.1007/s10107-006-0049-5.
- 6 Berkay Becu, Santanu S Dey, Feng Qiu, and Alinson S Xavier. Approximating the gomory mixed-integer cut closure using historical data. *arXiv preprint arXiv:2411.15090*, 2024.
- 7 Jacques F Benders. Partitioning procedures for solving mixed-variables programming problems. *Numerische Mathematik*, 4(1):238–252, 1962. doi:10.1007/bf01386316.
- 8 Timo Berthold. *Heuristic algorithms in global MINLP solvers*. Verlag Dr. Hut, 2015.
- 9 Daniel Bienstock and Oktay Günlük. Capacitated network design—polyhedral structure and computation. *Inform Journal on Computing*, 8(3):243–259, 1996. doi:10.1287/ijoc.8.3.243.
- 10 Suresh Bolusani, Mathieu Besançon, Ksenia Bestuzheva, Antonia Chmiela, João Dionísio, Tim Donkiewicz, Jasper van Doornmalen, Leon Eifler, Mohammed Ghannam, Ambros Gleixner, Christoph Graczyk, Katrin Halbig, Ivo Hedtke, Alexander Hoen, Christopher Hojny, Rolf van der Hulst, Dominik Kamp, Thorsten Koch, Kevin Kofler, Jurgen Lentz, Julian Manns, Gioni Mexi, Erik Mühmer, Marc E. Pfetsch, Franziska Schlösser, Felipe Serrano, Yuji Shinano, Mark Turner, Stefan Vigerske, Dieter Weninger, and Liding Xu. The SCIP Optimization Suite 9.0, 2024. [arXiv:2402.17702](https://arxiv.org/abs/2402.17702).
- 11 Pierre Bonami, Gérard Cornuéjols, Sanjeeb Dash, Matteo Fischetti, and Andrea Lodi. Projected Chvátal–Gomory cuts for mixed integer linear programs. *Mathematical Programming*, 113(2):241–257, 2008. doi:10.1007/s10107-006-0051-y.
- 12 Emmanuel J Candès, Justin Romberg, and Terence Tao. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Transactions on Information Theory*, 52(2):489–509, 2006. doi:10.1109/tit.2005.862083.
- 13 Emmanuel J Candes and Terence Tao. Decoding by linear programming. *IEEE Transactions on Information Theory*, 51(12):4203–4215, 2005. doi:10.1109/TIT.2005.858979.

- 14 Emmanuel J Candes, Michael B Wakin, and Stephen P Boyd. Enhancing sparsity by reweighted ℓ_1 minimization. *Journal of Fourier Analysis and Applications*, 14:877–905, 2008. doi:10.21236/ada528514.
- 15 Alberto Caprara and Matteo Fischetti. $\{0, 1/2\}$ -Chvátal-Gomory cuts. *Mathematical Programming*, 74:221–235, 1996. doi:10.1007/BF02592196.
- 16 Alberto Caprara, Matteo Fischetti, and Adam N Letchford. On the separation of maximally violated mod-k cuts. *Mathematical Programming*, 87:37–56, 2000. doi:10.1007/s101079900107.
- 17 Antonin Chambolle and Thomas Pock. A first-order primal-dual algorithm for convex problems with applications to imaging. *Journal of Mathematical Imaging and Vision*, 40:120–145, 2011. doi:10.1007/s10851-010-0251-1.
- 18 Scott Shaobing Chen, David L Donoho, and Michael A Saunders. Atomic decomposition by basis pursuit. *SIAM Review*, 43(1):129–159, 2001. doi:10.1137/s003614450037906x.
- 19 Philipp M Christophel. *Separation algorithms for cutting planes based on mixed integer row relaxations*. PhD thesis, Paderborn, Univ., Diss., 2009, 2009.
- 20 Vasek Chvátal. Edmonds polytopes and a hierarchy of combinatorial problems. *Discrete Mathematics*, 4(4):305–337, 1973. doi:10.1016/0012-365x(73)90167-2.
- 21 William Cook, Ravindran Kannan, and Alexander Schrijver. Chvátal closures for mixed integer programming problems. *Mathematical Programming*, 47(1):155–174, 1990. doi:10.1007/bf01580858.
- 22 Gérard Cornuéjols. Valid inequalities for mixed integer linear programs. *Mathematical Programming*, 112(1):3–44, 2008. doi:10.1007/s10107-006-0086-0.
- 23 Sanjeeb Dash, Oktay Günlük, and Andrea Lodi. Mir closures of polyhedral sets. *Mathematical Programming*, 121:33–60, 2010. doi:10.1007/s10107-008-0225-x.
- 24 Norman R Draper and Harry Smith. *Applied regression analysis*. McGraw-Hill. Inc, 1998.
- 25 Matteo Fischetti and Andrea Lodi. Optimizing over the first Chvátal closure. *Mathematical Programming*, 110(1):3–20, 2007. doi:10.1007/s10107-006-0054-8.
- 26 Ricardo Fukasawa, Laurent Poirrier, and Shenghao Yang. Split cuts from sparse disjunctions. *Mathematical Programming Computation*, 12:295–335, 2020. doi:10.1007/s12532-020-00180-9.
- 27 Ambros Gleixner, Gregor Hendel, Gerald Gamrath, Tobias Achterberg, Michael Bastubbe, Timo Berthold, Philipp Christophel, Kati Jarck, Thorsten Koch, Jeff Linderoth, et al. MIPLIB 2017: data-driven compilation of the 6th mixed-integer programming library. *Mathematical Programming Computation*, 13(3):443–490, 2021. doi:10.1007/s12532-020-00194-3.
- 28 Ralph Edward Gomory. An algorithm for the mixed integer problem. *Report No. P-1885, The Rand Corporation, Santa Monica, CA.*, 1960.
- 29 Joao PM Gonçalves and Laszlo Ladanyi. An implementation of a separation procedure for mixed integer rounding inequalities. *IBM Res. Report RC23686 (W0508-022)*, IBM, 2005.
- 30 Zonghao Gu, George L Nemhauser, and Martin WP Savelsbergh. Lifted flow cover inequalities for mixed 0-1 integer programs. *Mathematical Programming*, 85:439–467, 1999. doi:10.1007/s101070050067.
- 31 Oscar Guaje, Arnaud Deza, Aleksandr M Kazachkov, and Elias B Khalil. Machine learning for optimization-based separation: the case of mixed-integer rounding cuts. *arXiv preprint arXiv:2408.08449*, 2024.
- 32 Trevor Hastie, Robert Tibshirani, and Ryan J Tibshirani. Extended comparisons of best subset selection, forward stepwise selection, and the lasso. *arXiv preprint arXiv:1707.08692*, 2017.
- 33 Clifford M Hurvich and Chih—Ling Tsai. The impact of model selection on inference in linear regression. *The American Statistician*, 44(3):214–217, 1990. doi:10.1080/00031305.1990.10475722.
- 34 Adam N Letchford and Georgia Souli. On lifted cover inequalities: A new lifting procedure with unusual properties. *Operations Research Letters*, 47(2):83–87, 2019. doi:10.1016/j.orl.2018.12.005.

- 35 Andrea Lodi and Andrea Tramontani. Performance variability in mixed-integer programming. In *Theory driven by influential applications*, pages 1–12. INFORMS, 2013.
- 36 Hugues Marchand and Laurence A Wolsey. The 0-1 knapsack problem with a single continuous variable. *Mathematical Programming*, 85:15–33, 1999. doi:10.1007/s101070050044.
- 37 Hugues Marchand and Laurence A Wolsey. Aggregation and mixed integer rounding to solve MIPs. *Operations Research*, 49(3):363–371, 2001. doi:10.1287/opre.49.3.363.11211.
- 38 Gioni Mexi, Felipe Serrano, Timo Berthold, Ambros Gleixner, and Jakob Nordström. Cut-based conflict analysis in mixed integer programming. *arXiv preprint arXiv:2410.15110*, 2024.
- 39 Sangnam Nam, Mike E Davies, Michael Elad, and Rémi Gribonval. The cosparsity analysis model and algorithms. *Applied and Computational Harmonic Analysis*, 34(1):30–56, 2013. doi:10.1016/j.acha.2012.03.006.
- 40 George L Nemhauser and Laurence A Wolsey. A recursive procedure to generate all cuts for 0-1 mixed integer programs. *Mathematical Programming*, 46(1):379–390, 1990. doi:10.1007/bf01585752.
- 41 Ellen B Roecker. Prediction error and its estimation for subset-selected models. *Technometrics*, 33(4):459–468, 1991. doi:10.2307/1269417.
- 42 Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 58(1):267–288, 1996. doi:10.1111/j.2517-6161.1996.tb02080.x.
- 43 Andreas M Tillmann, Daniel Bienstock, Andrea Lodi, and Alexandra Schwartz. Cardinality minimization, constraints, and regularization: a survey. *SIAM Review*, 66(3):403–477, 2024. doi:10.1137/21m142770x.
- 44 Kati Wolter. Implementation of cutting plane separators for mixed integer programs. *Diploma thesis, Technische Universität Berlin*, 2006. URL: https://www.zib.de/userpage/groetschel/students/Diplom-Wolter-2006.pdf&ved=2ahUKEwimkM2zhPGMAxW2g_0HHfJoGTIQFnoECBoQAQ&usg=AOvVaw2Gan2dr7aRWTh4BA4g7Wdf.
- 45 Liding Xu. Codes for aggregations-based cut separator. Software, version 1.0., Research Campus MODAL, swhId: swh:1:dir:a7313b858eba5b70d7d70e5041dad8041700ae60 (visited on 2025-06-23). URL: https://github.com/lidingxu/aggregation_lp, doi:10.4230/artifacts.23664.