



# Face Dimensions of General-Purpose Cutting Planes for Mixed-Integer Linear Programs

Matthias Walter<sup>(✉)</sup>

Department of Applied Mathematics, University of Twente,  
Enschede, The Netherlands  
[m.walter@utwente.nl](mailto:m.walter@utwente.nl)

**Abstract.** Cutting planes are a key ingredient to successfully solve mixed-integer linear programs. For specific problems, their strength is often theoretically assessed by showing that they are facet-defining for the corresponding mixed-integer hull. In this paper we experimentally investigate the dimensions of faces induced by general-purpose cutting planes generated by a state-of-the-art solver. Therefore, we relate the dimension of each cutting plane to its impact in a branch-and-bound algorithm.

## 1 Introduction

We consider the mixed-integer program

$$\max c^\top x \tag{1a}$$

$$\text{s.t. } Ax \leq b \tag{1b}$$

$$x_i \in \mathbb{Z} \quad \forall i \in I \tag{1c}$$

for a matrix  $A \in \mathbb{R}^{m \times n}$ , vectors  $c \in \mathbb{R}^n$  and  $b \in \mathbb{R}^m$  and a subset  $I \subseteq \{1, 2, \dots, n\}$  of integer variables. Let  $P := \text{conv}\{x \in \mathbb{R}^n : x \text{ satisfies (1b) and (1c)}\}$  denote the corresponding *mixed-integer hull*. A *cutting plane* for (1) is an inequality  $a^\top x \leq \beta$  that is valid for  $P$  (and possibly invalid for some point computed during branch-and-cut). Such a valid inequality *induces the face*  $F := \{x \in P : a^\top x = \beta\}$  of  $P$ . For more background on polyhedra we refer to [12].

To tackle specific problems, one often tries to find *facet-defining* inequalities, which are inequalities whose induced face  $F$  satisfies  $\dim(F) = \dim(P) - 1$ . This is justified by the fact that any system  $Cx \leq d$  with  $P = \{x : Cx \leq d\}$  has to contain an inequality that induces  $F$ . Since the dimension of a face can vary between  $-1$  (no  $x \in P$  satisfies  $a^\top x = \beta$ ) and  $\dim(P)$  (all  $x \in P$  satisfy  $a^\top x = \beta$ ), the following hypothesis is a reasonable generalization of facetness as a strength indicator.

**Hypothesis 1.** *The practical strength of an inequality correlates with the dimension of its induced face.*

There is no unified notion of the practical strength of an inequality, but we will later define one that is related to its impact in a branch-and-bound algorithm. The main goal of this paper is to computationally test this hypothesis for general-purpose cutting planes used in MIP solvers.

*Outline.* In Sect. 2 we present the algorithm we used to compute dimensions. Sect. 3 is dedicated to the score we used to assess a cutting plane's impact, and in Sect. 4 we present our findings, in particular regarding Hypothesis 1.

## 2 Computing the Dimension of a Face

This section is concerned about how to effectively compute the dimension of  $P$  or of one of its faces  $F$  induced by some valid inequality  $a^\top x \leq \beta$ . In our experiments we will consider instances with several hundreds variables, and hence the enumeration of all vertices of  $P$  or  $F$  is typically impossible. Instead, we use an oracle-based approach. An *optimization oracle* for  $P$  is a black-box subroutine that can solve any linear program over the polyhedron  $P$ , i.e., for any given  $w \in \mathbb{R}^n$  it can solve

$$\max w^\top x \text{ s.t. } x \in P. \quad (2)$$

In case (2) is feasible and bounded, the oracle shall return an optimal solution, and in case it is feasible and unbounded, it shall return an unbounded direction. Note that for such an oracle we neither require all vertices of  $P$  nor all valid (irredundant) inequalities. Indeed, we can use any MIP solver and apply it to problem (1) with  $c := w$ . We now discuss how one can compute  $\dim(P)$  by only accessing an optimization oracle for  $P$ . The basic algorithm is known (see Lemma 6.5.3 in [4]), but we provide a slightly improved version that requires at most  $2n$  oracle queries (in contrast to the bound of  $3n$  for the cited one). In fact, one can prove that there is no oracle-based algorithm that always requires less than  $2n$  [13].

To keep the presentation simple we assume that  $P$  is bounded, although our implementation can handle unbounded polyhedra as well. The algorithm maintains a set  $X \subseteq P$  of affinely independent points and a system  $Dx = e$  of valid equations, where  $D$  has full row-rank  $r$ . Hence,  $\dim(X) \leq \dim(P) \leq n - r$  holds throughout and the algorithm works by either increasing  $\dim(X)$  or  $r$  in every iteration. The details are provided in Algorithm 1.

**Proposition 1.** *For an optimization oracle for a non-empty polytope  $P \subseteq \mathbb{R}^n$ , Algorithm 1 requires  $2n$  oracle queries to compute a set  $X \subseteq P$  with  $|X| = \dim(P) + 1$  and a system  $Dx = e$  of  $n - \dim(P)$  equations satisfying*

$$\text{aff}(P) = \text{aff}(X) = \{x \in \mathbb{R}^n : Dx = e\}.$$

*Proof.* Since every point  $x \in X$  was computed by the optimization oracle, we have  $X \subseteq P$ . Moreover,  $Dx = e$  is valid for  $P$  since for each equation  $d^\top x = \gamma$ ,  $\min\{d^\top x : x \in P\} = \gamma = \max\{d^\top x : x \in P\}$  holds. We now prove by induction

on the number of iterations that in every iteration of the algorithm, the points  $x \in X$  are affinely independent and that the rows of  $D$  are linearly independent.

Initially, this invariant holds because  $X = \emptyset$  and  $D$  has no rows. Whenever an equation  $d^\top x = d^\top x^+$  is added to  $Dx = e$  in step 7, the vector  $d$  is linearly independent to the rows of  $D$  (see step 3). If in step 10,  $X$  is initialized with a single point, it is clearly affinely independent. Similarly, the two points in step 14 also form an affinely independent set since  $d^\top x^+ > d^\top x^-$  implies  $x^+ \neq x^-$ . Suppose  $X$  is augmented by  $\bar{x} := x^+$  in step 19 or by  $\bar{x} := x^-$  in step 22. Note that by the choice of  $d$  in step 3,  $d^\top x = \gamma$  holds for all  $x \in X$ . Due to  $d^\top \bar{x} \neq \gamma$ ,  $X \cup \{\bar{x}\}$  remains affinely independent.

After the first iteration, we have  $|X| + |\text{rows}(D)| = 2$ , and in every further iteration, this quantity increases by 1. Hence, the algorithm requires  $n$  iterations, each of which performs 2 oracle queries. The result follows.  $\square$

**Algorithm 1:** Affine hull of a polytope via optimization oracle

**Input:** Optimization oracle for a polytope  $\emptyset \neq P \subseteq \mathbb{R}^n$ .

**Output:** Affine basis  $X$  of  $\text{aff}(P)$ , non-redundant system  $Dx = e$  with  $\text{aff}(P) = \{x \in \mathbb{R}^n : Dx = e\}$ .

```

1 Initialize  $X := \emptyset$  and  $Dx = e$  empty.
2 while  $|X| + 1 < n - |\text{rows}(D)|$  do
3   Compute a direction vector  $d := \text{aff}(X)^\perp \setminus \text{span}(\text{rows}(D))$ .
4   Query the oracle to maximize  $d^\top x$  over  $x \in P$  and let
      $x^+ := \arg \max\{d^\top x : x \in P\}$ .
5   Query the oracle to maximize  $-d^\top x$  over  $x \in P$  and let
      $x^- := \arg \min\{d^\top x : x \in P\}$ .
6   if  $d^\top x^+ = d^\top x^-$  then
7     Add equation  $d^\top x = d^\top x^+$  to system  $Dx = e$ .
8     if  $X = \emptyset$  then
9       set  $X := \{x^+\}$ .
10    end
11  end
12  else if  $X = \emptyset$  then
13    set  $X := \{x^+, x^-\}$ .
14  end
15  else
16    Let  $\gamma := d^\top x$  for some  $x \in X$ .
17    if  $d^\top x^+ \neq \gamma$  then
18      set  $X := X \cup \{x^+\}$ .
19    end
20    else
21      set  $X := X \cup \{x^-\}$ .
22    end
23  end
24 end
25 return  $(X, Dx = e)$ .
```

*Implementation Details.* We now describe some details of the implementation within the software framework IPO [14]. In the description of Algorithm 1 we did not specify step 3 precisely. While one may enforce the requirement  $d \notin \text{span}(\text{rows}(D))$  via  $d \in \text{rows}(D)^\perp$ , it turned out that this is numerically less stable than first computing a basis of  $\text{aff}(X)^\perp$  and selecting a basis element  $d$  that is not in the span of  $D$ 's rows. Moreover, we can take  $d$ 's sparsity and other numerical properties into account. Sparsity can speed-up the overall computation since very sparse objective vectors are sometimes easier to optimize for a MIP solver. In theory, for all  $x \in X$ , the products  $d^\top x$  have the same value. However, due to floating-point arithmetic the computed values may differ. It turned out that preferring directions  $d$  for which the range of these products is small helps to avoid numerical difficulties.

In our application we compute the dimension of  $P$  and of several of its faces. We can exploit this by caching all points  $x \in P$  returned by an optimization oracle in a set  $\bar{X}$ . Then, for each face  $F$  induced by an inequality  $a^\top x \leq \beta$  we can then compute the set  $\bar{X}_F := \{x \in \bar{X} : a^\top x = \beta\}$ . Before querying the oracle we can then check the set  $\bar{X}_F$  for a point with sufficiently large objective value, which saves two calls to the MIP solvers.

Let  $\bar{D}x = \bar{e}$  be the equation system returned by Algorithm 1 for  $P$ . Now, for a face  $F$  induced by inequality  $a^\top x \leq \beta$  we can initialize  $Dx = e$  in the algorithm by  $\bar{D}x = \bar{e}$ . Moreover, if  $a^\top x = \beta$  is not implied by  $Dx = e$  we add this equation to  $Dx = e$  as well.

Since the algorithm is implemented in floating-point arithmetic, errors can occur which may lead to wrong dimension results. We checked the results using an exact arithmetic implementation of Algorithm 1 for easier instances (with less than 200 variables). For these, the computed dimension varied by at most 2 from the true dimension. We conclude that the relative dimension errors are sufficiently small.

In a first implementation, our code frequently reported dimension  $-1$ , and it turned out that often the right-hand side was only slightly larger than needed to make the inequality supporting. Thus, for each inequality  $a^\top x \leq \beta$ , normalized to  $\|a\|_2 = 1$ , we computed  $\beta^{\text{true}} := \max\{a^\top x : x \in P\}$  by a single oracle query. Whenever we observed  $\beta < \beta^{\text{true}} - 10^{-4}$ , we considered the cut  $a^\top x \leq \beta$  as invalid (indicated by the symbol  $\blacklightning$ ). If  $\beta > \beta^{\text{true}} + 10^{-4}$ , we declare the cut to be non-supporting. Otherwise, we replace  $\beta$  by  $\beta^{\text{true}}$  before running Algorithm 1.

### 3 Measuring the Strength of a Single Inequality

In this section we introduce our *cut impact measure* for indicating, for a given cutting plane  $a^\top x \leq \beta$ , how useful its addition in the context of branch-and-cut is. The main goal of solving an LP at a branch-and-bound node is to determine a dual bound of the current subproblem. If this bound is at most the value of the best feasible solution known so far, then the subproblem can be discarded. Thus, we consider the value of such a bound (after adding a certain cut) in relation to the problem's optimum  $z^* := \max\{c^\top x : Ax \leq b, x_i \in \mathbb{Z} \forall i \in I\}$  and the value  $z^{\text{LP}} := \max\{c^\top x : Ax \leq b\}$  of the LP relaxation.

Our first approach was to just evaluate the dual bound obtained from the LP relaxation  $Ax \leq b$  augmented by  $a^\top x \leq \beta$ . However, adding a single inequality often does not cut off the optimal face of the LP relaxation, which means that the bound does not change. In a second attempt we tried to evaluate the dual bound of the LP relaxation augmented by a random selection of  $k$  cutting planes. However, the variance of the resulting cut impact measure was very large even after averaging over 10.000 such selections.

As a consequence, we discarded cut impact measures based on the combined impact of several cutting planes. Instead we carried out the following steps for given cutting planes  $a_1^\top x \leq \beta_1, a_2^\top x \leq \beta_2, \dots, a_k^\top x \leq \beta_k$ :

1. Compute the optimum  $z^*$  and optimum solution  $x^* \in P$ .
2. Compute  $z^{\text{LP}}$ .
3. For  $i = 0, 1, 2, \dots, k$ , solve

$$\max c^\top x \text{ s.t. } Ax \leq b, x_j \in \mathbb{Z} \forall j \in I \text{ and } a_i^\top x \leq \beta_i \text{ if } i \geq 1$$

with  $x^*$  as an initial incumbent, with presolve, cutting planes and heuristics disabled and where a time limit of 60 s is enforced.

4. Let  $N$  denote the minimum number of branch-and-bound nodes used in any these  $k + 1$  runs.
5. For  $i = 0, 1, \dots, k$ , let  $z^i$  be the dual bound obtained in run  $i$  when stopping after  $N$  branch-and-bound nodes. For  $i \geq 1$ , the *closed gap of cut  $i$*  is defined as  $(z^i - z^{\text{LP}})/(z^* - z^{\text{LP}})$ . For  $i = 0$ , this yields the *closed gap without cuts*.

*Remarks.* The closed gap is essentially the dual bound, normalized such that a value of 0 means no bound improvement over the root LP bound without cuts and a value of 1 means that the instance was solved to optimality. The effective limit of  $N$  branch-and-bound nodes was introduced such that all runs reach this limit. This circumvents the question of how to compare runs in which the problem was solved to optimality with those that could not solve it. For all runs, presolve and domain propagation were disabled due to our focus on the branch-and-bound algorithm itself. To avoid interaction with heuristics, the latter were disabled, but an optimal solution  $x^*$  was provided.

We are aware that it is debatable how meaningful this cut measure is with respect to the actual importance of a single inequality in a practical setting. However, we think that it is highly nontrivial to design a cut measure that is meaningful, robust and computable with reasonable effort. Therefore, we see our proposed measure as one contribution and hope that further research leads to the development of more measures, which in turn can lead to more reliable statements.

## 4 Computational Study

In order to test Hypothesis 1 we considered the 65 instances from the MIPLIB 3 [10]. We did not use the more recent MIPLIB 2017 [11] because

we had to restrict ourselves to problems with a decent number of variables. For each of them we computed the dimension of the mixed-integer hull  $P$ , imposing a time limit of 10 min<sup>1</sup>. Moreover, we ran the state-of-the-art solver SCIP [3] and collected all cutting planes generated in the root node, including those that were discarded by SCIP’s cut selection routine<sup>2</sup>. For each of the cuts we computed the dimension of its induced face (see Sect. 2) as well as the closed gap after processing  $N$  (as defined in Sect. 3) branch-and-bound nodes.

While for instances `air03`, `air05`, `nw04`, no cuts were generated by SCIP, our implementation of Algorithm 1 ran into numerical difficulties during the computation of  $\dim(P)$  for `set1ch`. Moreover,  $\dim(P)$  could not be computed within 10 min for instances `air04`, `arki001`, `cap6000`, `dano3mip`, `danooint`, `dsbmip`, `fast0507`, `gesa2_0`, `gesa3_0`, `l152lav`, `mas74`, `misc06`, `mitre`, `mkc`, `mod010`, `mod011`, `pk1`, `pp08aCUTS`, `pp08a`, `qnet1`, `rentacar`, `rout` and `swath`.

We evaluated the remaining 37 instances, whose characteristic data is shown in Table 1. We distinguish how many cuts SCIP found by which cut separation method, in particular to investigate whether certain separation routines tend to generate cuts with low- or high-dimensional faces.

We verified some of the invalid cutting planes manually, i.e., checked that these cuts are generated by SCIP and that there exists a feasible solution that is indeed cut off. The most likely reason for their occurrence is that SCIP performed dual reductions although we disabled them via corresponding parameters<sup>3</sup>.

Since the results on face dimensions and cut strength turned out to be very instance-specific, we created one plot per instance. We omit the ones for `p2756` (too many failures, see Table 1), `blend2` (only 7 cuts analyzed), `enigma` ( $\dim(P) = 3$  is very small), and for `markshare1`, `markshare2` and `noswot` (all cuts were ineffective). Moreover, we present several plots in the Appendix A since these are similar to those of other instances.

The plots show the dimension of the cuts (horizontal axis, rounded to 19 groups) together with their closed gap (vertical axis, 14 groups) according to Sect. 3. Each circle corresponds to a nonempty set of cuts, where the segments depict the respective cut classes (see Table 1) and their color depicts the number  $k$  of cuts, where the largest occurring number  $L$  is specified in the caption. The colors are red ( $0.9L < k \leq L$ ), orange ( $0.7L < k \leq 0.9L$ ), yellow ( $0.5L < k \leq 0.7L$ ), green ( $0.3L < k \leq 0.5L$ ), turquoise ( $0.1L < k \leq 0.3L$ ) and blue ( $1 \leq k < 0.1L$ ). For instance, the circle for `fixnet6` containing a red and a turquoise segment subsumes cutting planes with face dimensions between 730 and 777, and closed gap of approximately 0.45. As the legend next to the plot indicates, this circle represents  $k \in [109, 121]$  c-MIR cuts and  $k' \in [13, 36]$  multi-commodity flow cuts. The dashed horizontal line indicates the closed gap without cuts (see Sect. 3).




<sup>1</sup> All experiments were carried out on a single core of an Intel Core i3 CPU running at 2.10 GHz with 8 GB RAM.

<sup>2</sup> We disabled presolve, domain propagation, dual reductions, symmetry, and restarts.



<sup>3</sup> We set `misc/allowweakdualreds` and `misc/allowstrongdualreds` to `false`.

**Table 1.** Characteristics of the relevant 37 instances with number of successfully analyzed cuts, failures (numerical problems 0/0, timeouts  $\odot$ , invalid cuts  $\lightning$ ), dimension of the mixed-integer hull, and number  $N$  of branch-and-bound nodes (see Sect. 3).  $\odot$  – Lifted extended weight inequalities [9, 15, 16]  $\ominus$  – Complemented Mixed-Integer Rounding (c-MIR) inequalities [8, 16]  $\ominus$  –  $\{0, 1/2\}$ -Chvátal-Gomory inequalities [2, 6]  $\oplus$  – Strengthened Chvátal-Gomory inequalities [7]  $\ominus$  – Lifted flow-cover inequalities [5]  $\ominus$  – Multi-commodity flow inequalities [1]

Instance	Cuts total	Analyzed by class						Failed			Dim. of $P$	B& B nodes
		$\odot$	$\ominus$	$\ominus$	$\oplus$	$\ominus$	$\ominus$	0/0	$\odot$	$\lightning$		
bell13a	25		25								121	53075
bell15	53		36				1	11		5	97	7815
blend2	8		7							1	245	597
dcmulti	172		137					14		21	467	773
egout	125		97			1		24		3	41	1
enigma	82		59	1	17	5					3	1
fiber	533	127	139	7	4	9	207	33	6	1	946	23395
fixnet6	772		612				83	42		35	779	53900
flugpl	42		42								9	1753
gen	28	17	8					3			540	21
gesa2	470	16	419	2	6			27			1176	21114
gesa3	259		194	2	6	15		38	4		1104	729
gt2	143		92	3	39	8			1		188	1
harp2	1028	659	210	9	4	112		5	28	1	1300	13748
khb05250	122		78					5		39	1229	425
lseu	125	35	85	2	3						89	3495
markshare1	107		104			3					50	318260
markshare2	84		79			3			2		60	286347
mas76	225		204						1	20	151	211381
misc03	634	3	283	33	311				4		116	13
misc07	808		286	34	440			36	12		204	22839
mod008	441	119	272						3	47	319	1158
modglob	268		186					3	1	78	327	112516
noswot	164		151		11			2			120	160344
p0033	94	14	40		10			30			27	127
p0201	263	12	152	14	78			7			139	2
p0282	904	368	441	6	15	1		4		69	282	57
p0548	577	159	213	10	16	62		117			520	921
p2756	1000	82	96	22	48	55		26	671		2716	15149
qiu	63		51					2	10		709	10406
qnet1	162		95	10		47		6	4		1233	5
rgn	278		168			65		45			160	1691
seymour	6246		656	53	5528			9	0		1255	9
stein27	886		517	9	360						27	3673
stein45	1613		1221	10	382						45	45371
vpml	281		129			32		117		3	288	27881
vpml2	353		251	1		30		63		8	286	172271



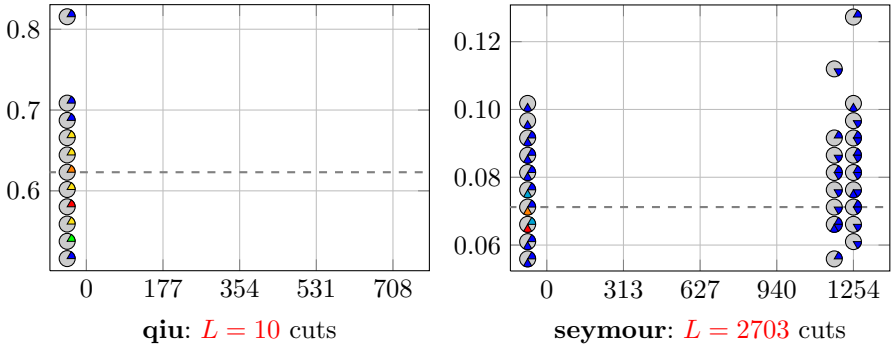
The first three plots already highlight that the results are very heterogeneous: while the faces of the strongest cuts in `fixnet6` have a high dimension, the strongest ones for `misc03` are not even supporting. Even when considering non-supporting cuts as outliers, the dimension does not indicate practical strength, as the plot for `gt2` shows. A quick look at the other plots lets us conclude that Hypothesis 1 is false—at least for the strength measure from Sect. 3.



The dashed line for `harp2` shows that adding a single cut does not necessarily help in branch-and-bound, which may be due to side-effects such as different branching decisions. For some instances, such as `mod008`, the cuts' face dimensions are well distributed. In contrast to this, some instances exhibit only very few distinct dimension values, e.g., only non-supporting cuts for `qiu`. Interest-




ingly, the 6246 cuts for `seymour` induce only empty faces as well as faces with dimensions between 1218 and 1254. This can partially be explained via the cut classes. On the one hand, all generated strengthened Chvátal-Gomory cuts are non-supporting. On the other hand, some of the  $c$ -MIR cuts and  $\{0, 1/2\}$ -cuts are non-supporting while others induce faces of very high dimension.




In general we do not see an indication that cuts from certain classes induce higher dimensional faces than others. At first glance, such a pattern is apparent for `misc07` (at dimensions 130–160), however one has to keep in mind that these blue segments constitute a minority of the cuts. In line with that, the majority of the cuts for `fiber` is concentrated around dimension 900 with a closed gap similar to that without cuts.

Despite the heterogeneity of the results, one observation is common to many instances: the distribution of the face dimensions is biased towards  $-1$  and high-dimensions, i.e., not many cuts inducing low dimensional faces are generated.



A corresponding histogram is depicted in Fig. 1. We conjecture that the high dimensions occur because lifting and strengthening techniques for cutting planes are quite evolved.




**Fig. 1.** Distribution of relative dimensions over the 37 instances from Table 1 except for p2756 (avoiding a bias due to many failures dimension computations). A cut of dimension  $k$  in an instance having  $\ell$  cuts and with  $\dim(P) = d$  contributes  $1/(36\ell)$  to its bar. For  $k = -1$  (resp.  $k = d$ ), this is  $\emptyset$  (resp.  $\infty$ ), and it is  $k/(d - 1)$  otherwise.

The first bar in Fig. 1 indicates that for an instance chosen uniformly at random among the ones we considered and then a randomly chosen cut for this instance, this cut is non-supporting with probability greater than 35%. This is remarkably high and thus we conclude this paper by proposing to investigate means to (heuristically) test for such a situation, with the goal of strengthening a non-supporting cutting plane by a reduction of its right-hand side.


**Acknowledgments.** We thank R. Hoeksma and M. Uetz as well as the SCIP development team, in particular A. Gleixner, C. Hojny and M. Pfetsch for valuable suggestions on the computational experiments and their presentation. Finally, we thank three anonymous referees for raising interesting discussion points, especially regarding Sect. 3.

## A Additional Plots


Here we provide additional instance-specific plots. This underlines the conclusions drawn in Sect. 4 and allows inspection of results for instances with certain characteristics (see Table 1).




**bell3a:**  $L = 5$  cuts




**bell5:**  $L = 3$  cuts




**dcmulti:**  $L = 38$  cuts



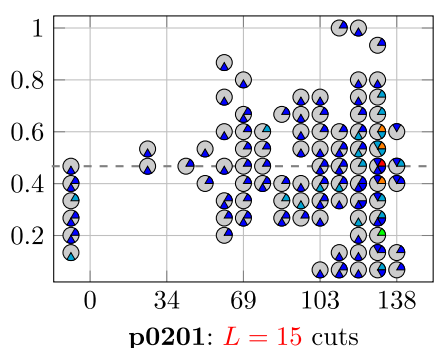
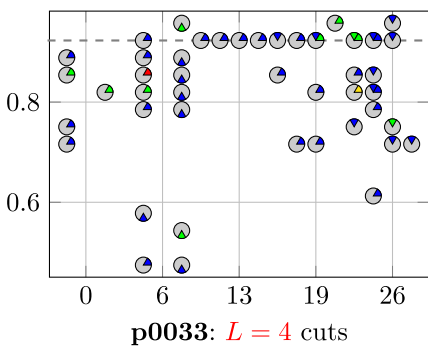
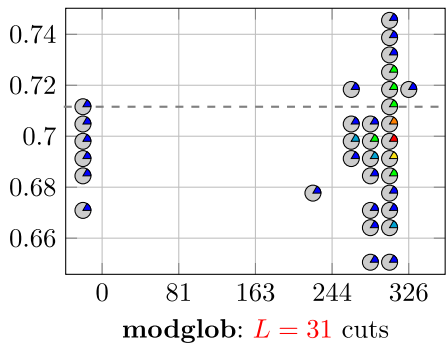
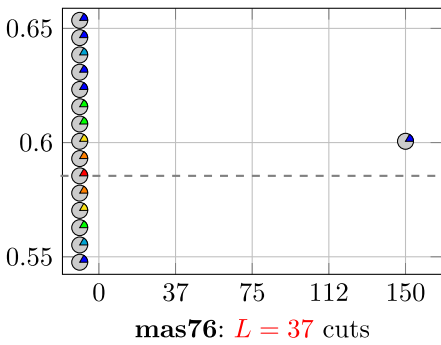

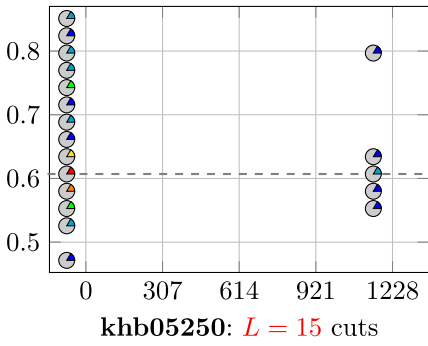
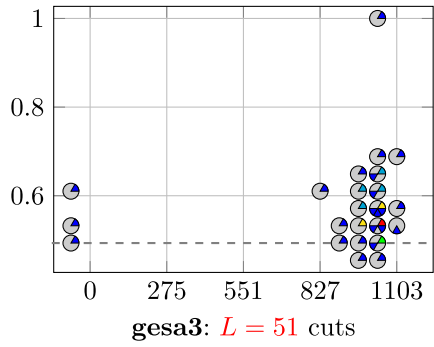

**egout:**  $L = 31$  cuts

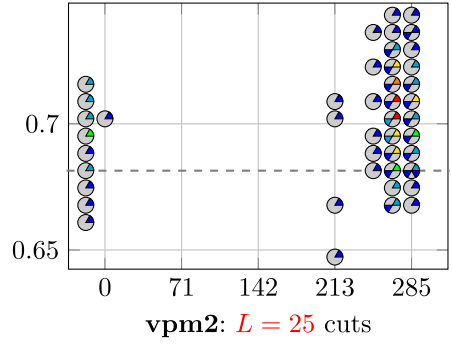
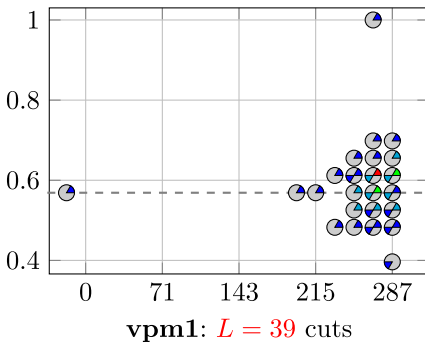
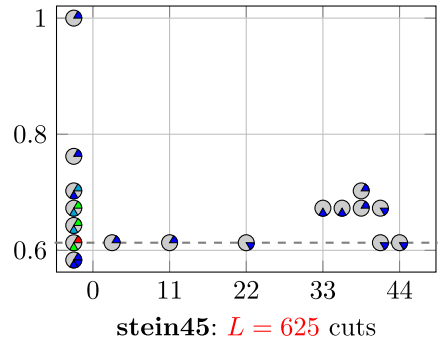
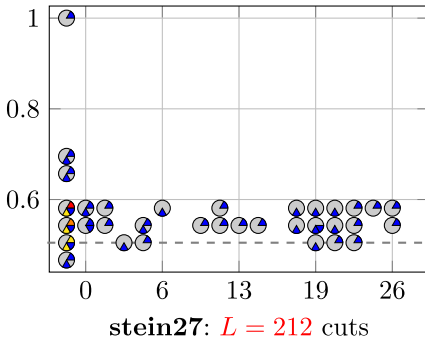






**flugpl:**  $L = 11$  cuts



**gen:**  $L = 5$  cuts





## References

1. Achterberg, T., Raack, C.: The MCF-separator: detecting and exploiting multi-commodity flow structures in MIPs. *Math. Program. Comput.* **2**(2), 125–165 (2010)
2. Caprara, A., Fischetti, M.:  $\{0, 1/2\}$ -Chvátal-Gomory cuts. *Math. Program.* **74**(3), 221–235 (1996)
3. Gamrath, G., et al.: The SCIP Optimization Suite 7.0. Technical report, Optimization Online, March 2020
4. Martin, G., Lovász, L., Schrijver, A.: *Geometric Algorithms and Combinatorial Optimizations*. Springer-Verlag (1993)
5. Gu, Z., Nemhauser, G.L., Savelsbergh, M.W.P.: Lifted flow cover inequalities for mixed 0–1 integer programs. *Math. Program.* **85**(3), 439–467 (1999)
6. Koster, A.M.C.A., Zymolka, A., Kutschka, M.: Algorithms to separate  $\{0, \frac{1}{2}\}$ -Chvátal-Gomory cuts. *Algorithmica* **55**(2), 375–391 (2009)
7. Letchford, A.N., Lodi, A.: Strengthening Chvátal-Gomory cuts and Gomory fractional cuts. *Oper. Res. Lett.* **30**(2), 74–82 (2002)
8. Marchand, H., Wolsey, L.A.: Aggregation and mixed integer rounding to solve MIPs. *Oper. Res.* **49**(3), 363–371 (2001)
9. Martin, A.: *Integer programs with block structure* (1999)
10. Mczeal, C.M., Savelsbergh, M.W.P., Bixby, R.E.: An updated mixed integer programming library: MIPLIB 3.0. *Optima*, **58**, 12–15 (1998)
11. Gleixner, A., et al.: MIPLIB 2017: data-driven compilation of the 6th mixed-integer programming library. *Math. Program. Comput.* (2020). <https://doi.org/10.1007/s12532-020-00194-3>
12. Schrijver, A.: *Theory of Linear and Integer Programming*. John Wiley & Sons Inc, New York, NY, USA (1986)
13. Walter, M.: *Investigating Polyhedra by Oracles and Analyzing Simple Extensions of Polytopes*. PhD thesis, Otto-von-Guericke-Universität Magdeburg (2016)
14. Walter, M.: *IPO - Investigating Polyhedra by Oracles* (2016). Software available at: [bitbucket.org/matthias-walter/ipo/](http://bitbucket.org/matthias-walter/ipo/)
15. Weismantel, R.: On the 0/1 knapsack polytope. *Math. Program.* **77**(3), 49–68 (1997)
16. Wolter, K.: *Implementation of Cutting Plane Separators for Mixed Integer Programs*. Master’s thesis, Technische Universität Berlin (2006)