

Partial hyperplane activation for generalized intersection cuts

Aleksandr M. Kazachkov^{*1} Selvaprabu Nadarajah² Egon Balas¹
François Margot¹

¹*Tepper School of Business, Carnegie Mellon University, Pittsburgh, PA*

²*College of Business Administration, University of Illinois at Chicago,
Chicago, IL*

Abstract

The generalized intersection cut (GIC) paradigm is a recent framework for generating cutting planes in mixed integer programming with attractive theoretical properties. We investigate this computationally unexplored paradigm and observe that a key hyperplane activation procedure embedded in it is not computationally viable. To overcome this issue, we develop a novel replacement to this procedure called partial hyperplane activation (PHA), introduce a variant of PHA based on a notion of hyperplane tilting, and prove the validity of both algorithms. We propose several implementation strategies and parameter choices for our PHA algorithms and provide supporting theoretical results. We computationally evaluate these ideas in the COIN-OR framework

^{*}akazachk@alumni.cmu.edu

on MIPLIB instances. Our findings shed light on the strengths of the PHA approach as well as suggest properties related to strong cuts that can be targeted in the future.

1 Introduction

An important aspect of modern integer optimization solvers is the use of cutting planes to strengthen a given formulation [1]. Finding methods for generating stronger general-purpose cutting planes has been an active topic of research in the past decade. As part of this trend, Balas and Margot [7] introduced the *generalized intersection cut* (GIC) paradigm with the motivation of finding stronger cutting planes that also possess favorable numerical properties, such as avoiding numerical inaccuracies that arise in traditional cutting plane approaches [33]. Despite offering some theoretical advantages to other modern cutting planes, GICs have remained unexplored computationally. In this paper, we observe that generating GICs as outlined in [7] is computationally intractable due to the exponential size of the linear program used to generate cuts. We offer a solution by extending the notion of a GIC and devising new algorithms to generate such GICs that scale well with the size of the input data. We prove the validity of our algorithms, provide theoretical results that guide our implementation choices, and perform the first computational investigation with GICs. Our investigation identifies properties related to strong cuts that can be targeted in future approaches within the paradigm.

Let $P := \{x \in \mathbb{R}^n : Ax \geq b, x \geq 0\}$ and $P_I := \{x \in P : x_j \in \mathbb{Z} \text{ for all } j \in I\}$ for a set $I \subseteq \{1, \dots, n\}$, where all data is rational. Let C be a relaxation of P defined by a subset of the inequalities defining P . Balas and Margot [7] produce GICs from a collection of intersection points and rays obtained by intersecting the edges of C with the boundary of a convex set S containing no points from P_I in its interior. The simplest GICs are the *standard intersection cuts* (SICs) [5], which use as C a simple polyhedral cone \bar{C} with apex at a vertex \bar{x} of P . GICs generalize SICs by using a tighter relaxation of P , by *activating* additional hyperplanes

of P that are not included in the description of \bar{C} . However, hyperplane activation poses two computational issues: first, maintaining a description of C becomes challenging, and second, the number of points and rays can quickly grow too large for any practical use.

This paper introduces a new method, called *partial hyperplane activation* (PHA), that addresses the issues with the aforementioned *full* hyperplane activation proposed in [7]. The key insight underlying PHA is that we can forgo a complete description of the polyhedron C by always activating hyperplanes on the initial cone \bar{C} , instead of on an iteratively refined relaxation. We show the PHA approach is not only valid, but also generates a collection of intersection points and rays that grows quadratically in size with the number of hyperplane activations, in contrast to the exponential growth exhibited by the full activation procedure.

Activating a hyperplane on \bar{C} creates new vertices lying at most one edge of P away from \bar{x} . Consequently, intersection points are obtained from edges originating at \bar{x} or these distance 1 vertices. Higher-distance PHA methods, with accompanying higher computational cost, can be easily defined by incorporating hyperplane intersections with edges a larger distance from \bar{x} , eventually recovering the full activation procedure. In other words, while we focus in this paper on the details of a distance 1 PHA procedure, its underlying ideas extend to a hierarchy of PHA procedures delivering progressively "stronger" sets of intersection points.

An important observation related to the effectiveness of PHA is that weak intersection points are created whenever a hyperplane being activated intersects rays of \bar{C} that do not intersect the boundary of S . To mitigate this issue, we introduce a class of *tilted* hyperplanes that can be used to avoid such rays. We also show that tilted hyperplanes offer more control over the number of intersection points generated with essentially no additional overhead, but they require a subtle condition to ensure that these points lead to cuts *valid* for P_I , i.e., to inequalities that do not cut off any points of P_I . We use this idea to design a modified PHA algorithm that creates a point-ray collection that grows linearly with the number of hyperplane activations but quadratically in the number of rays of \bar{C} being cut. The notion

of tilting may also have independent merit outside of the GIC paradigm whenever a \mathcal{V} -polyhedral partial description could be useful, i.e., describing a polyhedron by some of its vertices and edges (see, e.g., [34]).

Implementing PHA or its variant with tilting requires several algorithmic design choices, such as the decision of which particular hyperplanes to activate and which objective functions to use when optimizing over the derived collection of intersection points and rays. Our theoretical contributions provide geometric and structural insights into making some of these choices. Our experiments test the strength of GICs with respect to root node gap closed, a commonly used measure of cut strength, based on three hyperplane activation rules and four types of objective directions.

The first hyperplane activation rule we consider is based on the full activation method, which sequentially activates hyperplanes by pivoting to neighbors of \bar{x} . The second rule targets points that lie deep relative to the SIC. The third rule is motivated by a pursuit of a certain class of intersection points called *final*, relating to facet-defining inequalities for the S_k -closure, which is a relaxation of the split closure obtained by refining P through the addition of one simple split disjunction on a variable x_k .

We then discuss the objective directions evaluated in our implementation, which contribute to a better understanding of the process by which a collection of strong cuts can be generated. The first two sets of objective functions, as with the hyperplane rules, are motivated by finding cuts that improve over the SIC, as well as generalizing the usual objective of maximizing violation with respect to \bar{x} . We then give a result motivating the next set of directions, by showing that the optimal objective value over the S_k -closure can be obtained using only final intersection points and describing how to attain this same value using cuts. The fourth objective function type capitalizes on information that can be used from the typical empirical setup in which multiple split disjunctions are applied in parallel.

Finally, we present the first computational results for any algorithm within the GIC paradigm, implemented in the open source COIN-OR framework [29] and using a set of

benchmark instances from MIPLIB [11]. We find that our PHA approach creates a manageable number of intersection points from which we can generate GICs that effectively generalize SICs: one round of GICs generated from our methodology improves the *integrality gap* closed by one round of SICs on 75% of the instances tested, closing on average an additional 5.1% of the gap over SICs, on instances for which either SICs or GICs close any gap. Though our procedure is intended to be nonrecursive, we do test a second round of GICs, the outcome of which is some indication of the lack of tailing off by GICs as compared to SICs. Underlying these results is a detailed analysis of the various hyperplane activation and objective function choices for the linear program used to generate GICs. These experiments build an understanding of structural properties of strong GICs. In particular, we find that the objective functions inspired by final intersection points frequently find the strongest cuts from a given point-ray collection compared to the other objectives tested. However, such final intersection points appear infrequently in the point-ray collections generated by PHA, yielding a concrete direction for future work: to target final intersection points more directly.

Our research on GICs adds to the growing literature that generalizes and extends cutting plane methods based on intersection cuts and Gomory cuts. SICs are equivalent to Gomory cuts [25–27], a relationship surveyed in Conforti et al. [15], and are among the simplest, yet most effective, cuts used in practice. As SICs are generated using information from only one row of the simplex tableau, GICs are a natural generalization by using information from multiple rows. Using such information has been studied extensively in the past decade [4, 10, 16, 18–20, 22, 23, 30]. The focus of these approaches is to obtain stronger cuts by using a better cut-generating set obtained from the optimal solution to the continuous relaxation. GICs instead attempt to strengthen SICs using other rows but the same cut-generating set. The collection of intersection points and rays used to derive GICs can also be seen as defining a linear system for verifying the validity of cutting planes [21].

Broadly, GICs fall into the class of algorithms using the classical theoretical approach

of polarity to generate cuts [6]. This includes similarities to local or target cuts [14], which use the polar of a set of points and rays to generate cuts, and to the procedures in [30, 32], which use row generation to individually certify the validity of each generated cut. Though utilizing some of the same theoretical tools, GICs derive points and rays from a fundamentally different perspective and in a way that immediately guarantees cut validity.

This paper is organized as follows. In Section 2, we revisit full hyperplane activation from [7] and demonstrate the impediment to a practical GIC algorithm inherent in that approach. Section 3 then gives the main contribution of this paper, the PHA scheme. Section 4 shows the idea of tilting and gives an additional algorithm that can be used within the PHA procedure. Implementation choices and supporting theory are discussed in Section 5. Lastly, we give the results of our computational experiments in Section 6. Additional supporting theoretical results are contained in the appendix.

2 Full hyperplane activation

In this section, we show that the full hyperplane activation approach to generate GICs as proposed by Balas and Margot [7] is impractical, which provides the motivation for developing our new methodology in later sections. Given objective vector $c \in \mathbb{Q}^n$, the optimization problem is

$$\min\{c^\top x : x \in P_I\}. \tag{IP}$$

Let (LP) denote the *continuous* or *linear relaxation* of (IP), obtained by removing any integrality restrictions on the variables:

$$\min\{c^\top x : x \in P\}. \tag{LP}$$

Let \bar{x} denote an optimal basic feasible solution to (LP). For simplicity, we assume throughout that P is a full-dimensional pointed polyhedron, but our results extend to the general case with minor modifications.

Let S be a P_I -free convex set, a closed convex set such that its *interior*, denoted $\text{int } S$, contains no points of P_I , and suppose that $\bar{x} \in \text{int } S$. A commonly used such set is formed from a *simple* split disjunction $(x_k \leq \lfloor \bar{x}_k \rfloor) \vee (x_k \geq \lceil \bar{x}_k \rceil)$ on a variable x_k , $k \in I$:

$$S_k := \{x : \lfloor \bar{x}_k \rfloor \leq x_k \leq \lceil \bar{x}_k \rceil\}.$$

Let $\bar{\mathcal{N}} := \mathcal{N}(\bar{x})$ denote the set of nonbasic variables at \bar{x} . Let $\bar{C} := C(\bar{\mathcal{N}})$ be the polyhedral cone with apex at \bar{x} and defined by the n hyperplanes corresponding to the nonbasic variables $\bar{\mathcal{N}}$; denote these hyperplanes by $\bar{\mathcal{H}}$. Let $\text{rays}(\bar{C})$ be the rays of \bar{C} .

GICs generalize SICs, as the cone \bar{C} is but one possible relaxation of P . \bar{C} has n extreme rays that can be intersected with $\text{bd } S$ to obtain a set of intersection points $\bar{\mathcal{P}}$ and a set of rays $\bar{\mathcal{R}}$ of \bar{C} that do not intersect $\text{bd } S$, so that

$$\begin{aligned} \bar{\mathcal{R}} &:= \{r \in \text{rays}(\bar{C}) : r \cap \text{bd } S = \emptyset\}, \\ \bar{\mathcal{P}} &:= \{p^r : p^r := r \cap \text{bd } S, r \in \text{rays}(\bar{C}) \setminus \bar{\mathcal{R}}\}. \end{aligned}$$

The intersection points in $\bar{\mathcal{P}}$ and rays in $\bar{\mathcal{R}}$ uniquely define the SIC, $\alpha_0^\top x \geq \beta_0$, obtained from \bar{C} and S [6].

Letting \mathcal{H} denote the set of hyperplanes of P , \bar{C} can be replaced by a tighter relaxation of P , denoted C , obtained by activating a subset of hyperplanes from $\mathcal{H} \setminus \bar{\mathcal{H}}$, i.e., of those that define P but not \bar{C} . Subsequently, edges of C can be intersected with $\text{bd } S$ to obtain a set of intersection points \mathcal{P} , and edges that do not intersect $\text{bd } S$ yield a set of rays \mathcal{R} . In contrast to the situation when using \bar{C} , the number of intersection points and rays obtained from C intersected with $\text{bd } S$ can be strictly greater than n , meaning the point-ray collection defines not just one cut, but a collection of valid cuts.

Theorem 4 in [7] defines valid GICs as these cuts and propose to generate them as any $(\alpha, \beta) \in \mathbb{R}^n \times \mathbb{R}$ satisfying $\alpha^\top \bar{x} < \beta$ and

$$\begin{aligned} p^\top \alpha &\geq \beta \quad \text{for all } p \in \mathcal{P} \\ r^\top \alpha &\geq 0 \quad \text{for all } r \in \mathcal{R}. \end{aligned} \tag{1}$$

For a fixed β , define $\text{CutRegion}(\beta, \mathcal{P}, \mathcal{R})$ as the vectors α feasible to the above system, i.e., the coefficients for inequalities with lower-bound β that are valid for \mathcal{P} and \mathcal{R} . It suffices to consider $\beta \in \{-1, 0, 1\}$ to obtain all possible cuts.

Unfortunately, the above method requires maintaining the \mathcal{V} -polyhedral description of C . The size of this \mathcal{V} -polyhedral description, as well as the number of rows of (1), as shown in Proposition 1, can grow exponentially large in the number of hyperplanes defining C . Together, these two issues make the full hyperplane activation approach unviable.

Proposition 1. *Let C be formed by adding k_h hyperplanes to the description of \bar{C} , and let $(\mathcal{P}, \mathcal{R})$ denote the point-ray collection obtained from intersecting the edges of C with $\text{bd } S$. The cardinality of \mathcal{P} can grow exponentially large in k_h .*

Proof. Suppose we start with \bar{C} and consider the activation of a halfspace H^+ . Let $\text{rays}(H)$ denote the rays from $\text{rays}(\bar{C})$ that are intersected by H before $\text{bd } S$. For each ray $r \in \text{rays}(H)$ that intersects H , $n - 1 - |\text{rays}(H)|$ new edges are created (not counting the edge back to \bar{x} or the edges between new vertices created by activating H). If $|\text{rays}(H)| \approx n/2$, the size of the new point-ray collection will be $O(n^2)$. The desired result follows inductively. \square

3 Partial hyperplane activation

In this section, we propose PHA, an approximate activation method that overcomes the exponential growth in the size of the point-ray collection when using full hyperplane activation. Specifically, our algorithm generates at most $O(k_r^2 k_h)$ intersection points and rays, where k_r and k_h are both parameters; $k_r \leq n$ is the maximum number of initial intersection

points and rays that we remove via hyperplane activations and k_h is the number of activated hyperplanes.

Showing the validity of PHA requires a strict extension of the tools used in [7] to prove the validity of full hyperplane activation. We give this extension in Section 3.1 and employ it in Section 3.2, in which we describe a concrete variant of PHA called PHA₁ and prove its validity. To facilitate reading, Table 1 gives a summary of some of the most frequently used notation that is used in the paper.

Table 1: Summary of frequently used notation.

Notation	Description
\mathcal{H}	Hyperplanes defining P
H_h, H_h^+, H_h^-	$H_h = \{x : a_h^\top x = b_h\}$, $H_h^+ = \{x : a_h^\top x \geq b_h\}$, $H_h^- = \mathbb{R}^n \setminus H_h^+$
\mathcal{N}	Set of nonbasic variables defining \bar{x}
\bar{C}	Simple cone defined by the nonbasic variables at \bar{x}
rays(\bar{C})	Rays of \bar{C}
$\bar{\mathcal{H}}$	Hyperplanes defining \bar{C}
(\mathcal{P}, \mathcal{R})	Set of points on $\text{bd } S$ and rays not intersecting $\text{bd } S$
($\bar{\mathcal{P}}, \bar{\mathcal{R}}$)	Initial point-ray collection from $\bar{C} \cap \text{bd } S$
CutRegion($\bar{\beta}, \mathcal{P}, \mathcal{R}$)	Feasible region of (PRLP)
$\mathcal{R}_{\mathcal{P}}$	$\{p - \bar{x} : p \in \mathcal{P}\}$
$\hat{C}(\mathcal{P}, \mathcal{R})$	$\bar{x} + \text{cone}(\mathcal{R}_{\mathcal{P}} \cup \mathcal{R})$
$\mathcal{G}(\mathcal{P}, \mathcal{R})$	$\text{conv}(\mathcal{P}) + \text{cone}(\mathcal{R})$
$\mathcal{H}(r)$	Hyperplanes defining a ray r
rays(H_h)	Rays from rays(\bar{C}) intersected by H_h before $\text{bd } S$

3.1 Proper point-ray collections

We define a point-ray collection as a pair (\mathcal{P}, \mathcal{R}) of points and rays obtained from intersecting the polyhedron C with $\text{bd } S$, where C is a relaxation of P . Let \mathcal{K} denote the skeleton of the polyhedron C . Let \mathcal{K}' denote the connected component of $\mathcal{K} \cap \text{int } S$ that includes \bar{x} and \mathcal{K}'' denote the union of the other components of $\mathcal{K} \cap \text{int } S$.

Definition 2. *The point-ray collection (\mathcal{P}, \mathcal{R}) is called proper if $\alpha^\top x \geq \beta$ is valid for P_I whenever (α, β) is feasible to (1) and $\alpha^\top v < \beta$ for some $v \in \mathcal{K}'$.*

With this definition, Theorem 3 shows that full hyperplane activation creates a proper point-ray collection, which follows easily from the proof of Theorem 4 in [7].

Theorem 3. *The point-ray collection $(\mathcal{P}, \mathcal{R})$ obtained from intersecting all edges of C with $\text{bd } S$, where C is defined by a subset of the hyperplanes defining P , is proper.*

Given a proper point-ray collection, a GIC is defined as a facet of the convex hull of points and rays, $\mathcal{G}(\mathcal{P}, \mathcal{R}) := \text{conv}(\mathcal{P}) + \text{cone}(\mathcal{R})$, referred to by \mathcal{G} for short, that cuts off a vertex of \mathcal{K}' . Balas and Margot [7] show that these facets can be generated by solving the following linear program formed from the given points and rays and optimized with respect to a given objective direction $w \in \mathbb{R}^n$ and a fixed $\beta \in \{-1, 0, 1\}$:

$$\min\{w^\top \alpha : \alpha \in \text{CutRegion}(\beta, \mathcal{P}, \mathcal{R})\}. \quad (\text{PRLP})$$

In fact, we show something stronger in Lemma 4: the cuts we obtain are not only valid for P_I , but also for $\text{conv}(C \setminus \text{int } S)$. This extends Theorem 3 and will be invoked in proving the results in Section 3.2. For simplicity, we assume that the vertex $v \in \mathcal{K}'$ being cut in Definition 2 is \bar{x} in the rest of the paper unless specified otherwise.

Lemma 4. *Let $(\mathcal{P}, \mathcal{R})$ be the point-ray collection obtained from intersecting all edges of C with $\text{bd } S$. If (α, β) is feasible to (1) and $\alpha^\top \bar{x} < \beta$, then*

$$F := \{x \in C : \alpha^\top x < \beta\} \subseteq \text{int } S.$$

Proof. Assume for the sake of contradiction that $F \setminus \text{int } S$ is nonempty. Then $F \cap \text{bd } S$ is nonempty. We show that this implies an intersection point or ray of C will violate $\alpha^\top x \geq \beta$.

First, note that $\text{cl } F \cap \text{bd } S$ cannot contain any vertex violating the inequality $\alpha^\top x \geq \beta$, as each vertex of $\text{cl } F \cap \text{bd } S$ is an intersection point of C . Hence, as we are also assuming $F \cap \text{bd } S$ is nonempty, the linear program $\min\{\alpha^\top x : x \in \text{cl } F \cap \text{bd } S\}$ must be unbounded. Consider initializing this program at \bar{x} and proceeding by the simplex method using the

steepest edge rule. As the simplex method follows edges of C and it will never increase the objective value, we will either eventually intersect $\text{bd } S$ or end up on a ray of C that is parallel to $\text{bd } S$. This is the desired contradiction: in the former case, this is an intersection point that violates the inequality, and in the latter case, this is a ray of C that does not intersect $\text{bd } S$ and violates the inequality. \square

3.2 Algorithm and validity

We describe the PHA_1 approach for activating a hyperplane partially and prove that it yields a proper point-ray collection. The algorithm proceeds similarly to the full hyperplane activation procedure, in that it activates a hyperplane valid for P , but unlike the procedure from [7], PHA_1 can use any hyperplane valid for P (not necessarily one that defines P), and the activation is performed on the initial relaxation \bar{C} instead of an iteratively refined relaxation that includes previously activated hyperplanes. We show this is not only valid but also computationally advantageous. Part of this advantage comes from that fact that PHA_1 only requires primal simplex pivots in the tableau of the LP relaxation to calculate new intersection points and rays.

The tradeoff for the computational advantages of PHA_1 is that it potentially creates a weaker point-ray collection compared to full activation, in the sense that the cuts generated from this collection may be weaker. This is because all vertices created from the partial activations are restricted to being only one pivot away from \bar{x} , i.e., at (edge) distance 1 from \bar{x} along the skeleton of P . For this reason, we refer to Algorithm 1 as PHA_1 . A stronger version of PHA_1 , such as a valid distance 2 (or higher distance) procedure, is not difficult to create. In other words, one could define a hierarchy of PHA procedures of increasing distance using the principles presented below, eventually leading to the full activation procedure.

Before proceeding, we first provide an illustrative example, shown in Figure 1, of our procedure that guides the intuition for the subsequent formal proof. The polytope in the first panel shows the feasible region of $P := \{x \in \mathbb{R}^3 : -2x_2 + x_3 \leq 0; -2x_1 + x_3 \leq$




Figure 1: This figure illustrates the sequence of activations performed by PHA_1 , where the cut-generating set is $\{x \in \mathbb{R}^3 : 0 \leq x_j \leq 1, j \in \{1, 2\}\}$. The first panel shows P . The second panel depicts \bar{C} . The bottom two panels show the sequential activations of the two hyperplanes not defining \bar{C} .

$0; 12x_1 + 10x_2 - 5x_3 \leq 9; 10x_1 + 12x_2 - 5x_3 \leq 9; x_1 + x_2 + x_3 \leq 1\}$. We denote the hyperplanes defining P by H_1 – H_5 in the order that they are listed in the definition of P . The next panel shows the cone \bar{C} , which is defined by H_1 – H_3 . The rays of \bar{C} are denoted r^1 to r^3 . The cut-generating set S is the box, $\{x \in \mathbb{R}^3 : 0 \leq x_j \leq 1, j \in \{1, 2\}\}$. We use circular and square nodes to represent vertices and intersection points, respectively. The remaining two panels illustrate the full and partial activation of hyperplanes H_4 and H_5 that do not play a role in defining \bar{C} . The third panel shows the full activation of H_4 . This activation cuts off the two intersection points corresponding to r^1 and r^2 intersected with $\text{bd } S$, and replaces them with two intersection points on H_4 . There are also two new vertices that are created — since activation always occurs on \bar{C} , these vertices lie on rays of this cone. The last panel shows H_5 partially activated on \bar{C} . This second activation creates a stronger intersection point along axis x_2 , as well as a *weaker* intersection point along axis x_1 , and it removes one of the intersection points from the activation of H_4 .

We now define some notation used in the algorithm description. For a given $(a_h, b_h) \in$

$\mathbb{R}^n \times \mathbb{R}$, let $H_h := \{x : a_h^\top x = b_h\}$ be a hyperplane such that the corresponding halfspace $H_h^+ := \{x : a_h^\top x \geq b_h\}$ is valid for P . Every element of a proper point-ray collection $(\mathcal{P}, \mathcal{R})$ arises from the intersection of an edge of some relaxation of P with $\text{bd } S$; if the starting vertex of that edge lies on a ray $r^j \in \text{rays}(\bar{C})$, then we say that the point or ray *originates* from r^j . For each $r^j \in \text{rays}(\bar{C})$, let the distance to H_h along the ray r^j starting at \bar{x} be

$$\text{dist}(H_h, r^j) := \begin{cases} (b_h - a_h^\top \bar{x}) / a_h^\top r^j & \text{if } a_h^\top r^j > 0, \\ \infty & \text{otherwise.} \end{cases}$$

We will slightly overload this notation to also refer to the distance along ray r^j to $\text{bd } S$, $\text{dist}(\text{bd } S, r^j)$. This will be the distance along r^j from \bar{x} to the nearest facet of S (or ∞ if no facet of S is intersected). With these definitions, we present Algorithm 1, which takes as an input the polyhedron P , the cut-generating set S , the hyperplane H_h being activated, a set of rays \mathcal{R}_A from $\text{rays}(\bar{C})$ that will be permitted to be cut by H_h (a parameter we will use only later), and any proper point-ray collection to be modified via the activation of H_h . We also assume that along with the point-ray collection, we have kept a history of which edge led to each point and ray to be added to the collection. This is used in the algorithm in step 3 when deciding which points and rays to remove.

Algorithm 1 Distance 1 Partial Hyperplane Activation

Input: Polyhedron P defined by hyperplane set \mathcal{H} ; P_I -free convex set S ; hyperplane H_h valid for P ; set of rays $\mathcal{R}_A \subseteq \text{rays}(\bar{C})$; proper point-ray collection $(\mathcal{P}, \mathcal{R})$ contained in \bar{C} .

- 1: **function** PHA₁($P, S, H_h, \mathcal{R}_A, (\mathcal{P}, \mathcal{R})$)
 - 2: **for all** $r \in \mathcal{R}_A$ such that $\text{dist}(H_h, r) < \text{dist}(\text{bd } S, r)$ **do**
 - 3: Remove all points and rays in $(\mathcal{P}, \mathcal{R})$ that violate H_h^+ and originate from a vertex on r but do not lie on or coincide with any ray from $\text{rays}(\bar{C}) \setminus \mathcal{R}_A$.
 - 4: **for all** new edges e originating at vertex $r \cap H_h$ **do**
 - 5: **if** e (emanating from $r \cap H_h$) does not intersect any $H \in \bar{\mathcal{H}} \setminus \mathcal{H}(r)$ prior to $\text{bd } S$
 - 6: **then**
 - 7: **if** $e \cap \text{bd } S \neq \emptyset$ **then** Add intersection point of e with $\text{bd } S$ to \mathcal{P} .
 - 8: **else** Add the ray e to \mathcal{R} .
 - 9: **return** \mathcal{P} and \mathcal{R} .
-

Theorem 5 states the validity of Algorithm 1 for the special case when all rays of \bar{C} are

permitted to be cut. We show how to relax this restriction in Section 4.

Theorem 5. *Algorithm 1 when $\mathcal{R}_A = \text{rays}(\bar{C})$ outputs a proper point-ray collection.*

In the rest of this section, we will build to a proof of Theorem 5 with the help of several useful intermediate results (Lemmas 6–9). The validity of PHA₁ follows from proving that the inequalities that can be generated from (PRLP) are valid for a relaxation of P_I related only to \bar{x} and the points and rays in the collection. Unlike the proof of full hyperplane activation, this does not rely on any particular relaxation of P used to obtain the points and rays. Given a point-ray collection $(\mathcal{P}, \mathcal{R})$, let $\mathcal{R}_{\mathcal{P}} := \{p - \bar{x} : p \in \mathcal{P}\}$ and $\hat{C}(\mathcal{P}, \mathcal{R}) := \bar{x} + \text{cone}(\mathcal{R}_{\mathcal{P}} \cup \mathcal{R})$. That is, $\hat{C}(\mathcal{P}, \mathcal{R})$ is the polyhedral cone with apex at \bar{x} and rays including \mathcal{R} and a ray from \bar{x} through each intersection point in \mathcal{P} .

Lemma 6. *If $(\mathcal{P}, \mathcal{R})$ is a proper point-ray collection, then $P_I \subseteq \hat{C}(\mathcal{P}, \mathcal{R})$.*

Proof. Otherwise there exists an inequality valid for \mathcal{G} that cuts both \bar{x} and a point of P_I . This contradicts the assumption that $(\mathcal{P}, \mathcal{R})$ is proper. \square

The next two lemmas state inclusion properties for the sets of cuts obtainable from two different point-ray collections.

Lemma 7. *For $i \in \{1, 2\}$, let \mathcal{P}^i denote a set of points on $\text{bd } S$ and \mathcal{R}^i a set of rays. If for every (α, β) such that α is feasible to $\text{CutRegion}(\beta, \mathcal{P}^2, \mathcal{R}^2)$ and $\alpha^\top \bar{x} < \beta$, α is also feasible to $\text{CutRegion}(\beta, \mathcal{P}^1, \mathcal{R}^1)$, then*

$$\hat{C}(\mathcal{P}^1, \mathcal{R}^1) \subseteq \hat{C}(\mathcal{P}^2, \mathcal{R}^2).$$

Proof. Let $\hat{C}^i := \hat{C}(\mathcal{P}^i, \mathcal{R}^i)$, $i \in \{1, 2\}$. Assume for the sake of contradiction that there exists an extreme ray r of \hat{C}^1 that is not in \hat{C}^2 . If r does not intersect $\text{bd } S$, define $q := r$; otherwise, let $q := r \cap \text{bd } S$, and note that $q \notin \hat{C}^2$. Then there exists an inequality $\hat{\alpha}^\top x \geq \hat{\beta}$ that separates \bar{x} and q from $\mathcal{G}(\mathcal{P}^2, \mathcal{R}^2)$. This is a contradiction, as $\hat{\alpha}$ is feasible to $\text{CutRegion}(\hat{\beta}, \mathcal{P}^2, \mathcal{R}^2)$ and $\hat{\alpha}^\top \bar{x} < \hat{\beta}$, but it is not feasible to $\text{CutRegion}(\hat{\beta}, \mathcal{P}^1, \mathcal{R}^1)$, since it

violates the inequality corresponding to q . \square

Lemma 8. *Let P^1 and P^2 be rational convex polyhedra such that $P^1 \subseteq P^2$, and $\bar{x} \in P^2$. Let \mathcal{P}^i and \mathcal{R}^i denote the intersection points and rays generated from the intersection of P^i with $\text{bd } S$, for $i \in \{1, 2\}$. Then, for a given constant β , a vector α feasible to $\text{CutRegion}(\beta, \mathcal{P}^2, \mathcal{R}^2)$ and cutting \bar{x} will also be feasible to $\text{CutRegion}(\beta, \mathcal{P}^1, \mathcal{R}^1)$.*

Proof. Since $P^1 \subseteq P^2$, we have $P^1 \cap \text{bd } S \subseteq P^2 \cap \text{bd } S \subseteq P^2 \setminus \text{int } S$. By Lemma 4, $P^2 \setminus \text{int } S \subseteq \{x : \alpha^\top x \geq \beta\}$, which completes the proof. \square

Next, we remark on one subtlety of Algorithm 1, in step 2. Namely, whenever the hyperplane being activated intersects a ray of \bar{C} outside of $\text{int } S$, the algorithm skips that ray. In Lemma 9, we show this is valid because such intersections would lead to only *redundant* intersection points, in the sense that these points will never be part of a valid inequality that cuts away \bar{x} . We show the result only for points, as rays are never redundant. Recall that \mathcal{K} denotes the skeleton of the polyhedron C , \mathcal{K}' denotes the connected component of $\mathcal{K} \cap \text{int } S$ that includes \bar{x} , and \mathcal{K}'' denotes the union of the other components of $\mathcal{K} \cap \text{int } S$. Let \mathcal{P}' and \mathcal{P}'' denote the intersection points between $\text{bd } S$ and the edges of the closure of \mathcal{K}' and \mathcal{K}'' , respectively.

Lemma 9. *Any intersection points in $\mathcal{P}'' \setminus \mathcal{P}'$ are redundant.*

Proof. Let α be any feasible solution to $\text{CutRegion}(\beta, \mathcal{P}', \mathcal{R})$ such that $\alpha^\top \bar{x} < \beta$. Suppose for the sake of contradiction that a point $w \in \mathcal{P}'' \setminus \mathcal{P}'$ violates the inequality. Since $w \in \mathcal{K}''$ and $\bar{x} \in \mathcal{K}'$, any path from \bar{x} to w along the skeleton of C must intersect $\text{bd } S$ at some point in \mathcal{P}' . Since C is convex and both \bar{x} and w are cut, there must exist such a path entirely in $\{x : \alpha^\top x < \beta\}$, implying some intersection point in \mathcal{P}' is cut, a contradiction. \square

We now use these results to prove Theorem 5.

Proof of Theorem 5. Let $(\mathcal{P}^*, \mathcal{R}^*)$ be the point-ray collection that Algorithm 1 outputs. Let $\hat{C}^* := \hat{C}(\mathcal{P}^*, \mathcal{R}^*)$. Observe that intersecting the edges of \hat{C}^* with $\text{bd } S$ yields exactly the same

point-ray collection $(\mathcal{P}^*, \mathcal{R}^*)$. Then, using Lemma 4 with \widehat{C}^* , it follows that all obtainable cuts will be valid for $\widehat{C}^* \setminus \text{int } S$. To prove the theorem, it suffices to show the inclusion $P_I \subseteq \widehat{C}^*$, which we do next.

From Lemma 6, we have that $P_I \subseteq \widehat{C} := \widehat{C}(\mathcal{P}, \mathcal{R})$, where $(\mathcal{P}, \mathcal{R})$ is the proper point-ray collection given as an input to the algorithm. We need to show that the point-ray collection after activating H_h is proper. Let \mathcal{P}^c and \mathcal{R}^c denote the set of points and rays from $(\mathcal{P}, \mathcal{R})$ that violate H_h^+ .

First we consider intersection points that lie on H_h . Let $P^1 := \widehat{C} \cap H_h$ and $P^2 := \bar{C} \cap H_h$, with \mathcal{P}^i and \mathcal{R}^i the corresponding intersection points and rays of P^i with $\text{bd } S$, for $i \in \{1, 2\}$. Since \mathcal{P} and \mathcal{R} are contained in \bar{C} , $\widehat{C} \subseteq \bar{C}$. As a result, we can apply Lemma 8, which implies that every α (for a β such that $\alpha^\top \bar{x} < \beta$) that is feasible to $\text{CutRegion}(\beta, \mathcal{P}^2, \mathcal{R}^2)$ will also be feasible to $\text{CutRegion}(\beta, \mathcal{P}^1, \mathcal{R}^1)$.

Now we consider the entire halfspace H_h^+ . The set of intersection points of $\widehat{C} \cap H_h^+$ with $\text{bd } S$ is $\bar{\mathcal{P}} := \mathcal{P}^1 \cup \mathcal{P} \setminus \mathcal{P}^c$, and the set of rays is $\bar{\mathcal{R}} := \mathcal{R}^1 \cup \mathcal{R} \setminus \mathcal{R}^c$. In addition, we have that $\mathcal{P}^* = \mathcal{P}^2 \cup \mathcal{P} \setminus \mathcal{P}^c$, and $\mathcal{R}^* = \mathcal{R}^2 \cup \mathcal{R} \setminus \mathcal{R}^c$. Here we are assuming that all of the points in \mathcal{P}^2 belong to the same part of the skeleton of $\bar{C} \cap H_h^+$ that contains \bar{x} . This is without loss of generality as a result of Lemma 9.

Hence, for any β , suppose α is feasible to $\text{CutRegion}(\beta, \mathcal{P}^*, \mathcal{R}^*)$ and satisfies $\alpha^\top \bar{x} < \beta$. Then α will also be feasible to $\text{CutRegion}(\beta, \bar{\mathcal{P}}, \bar{\mathcal{R}})$. By Lemma 7, $\widehat{C}(\bar{\mathcal{P}}, \bar{\mathcal{R}}) \subseteq \widehat{C}^*$. Moreover, since $P_I \subseteq \widehat{C}$ and $P_I \subseteq H_h^+$, we have that $P_I \subseteq \widehat{C} \cap H_h^+$. The intersection points and rays obtained from intersecting $\widehat{C}(\bar{\mathcal{P}}, \bar{\mathcal{R}})$ with $\text{bd } S$ are precisely $\bar{\mathcal{P}}$ and $\bar{\mathcal{R}}$. As a result, by Lemma 6, $P_I \subseteq \widehat{C}(\bar{\mathcal{P}}, \bar{\mathcal{R}})$. It follows that $P_I \subseteq \widehat{C}^*$, as desired. \square

Performing Algorithm 1 k_h times is dramatically more efficient than the $O(n^{k_h})$ complexity of full hyperplane activation. For every activation using Algorithm 1, letting $k_r := |\mathcal{R}_A|$, the hyperplane H_h can cut $O(k_r)$ rays to create $O(k_r^2)$ new intersection points. Therefore, if k_h hyperplanes are activated, the number of intersection points generated by partial activation is $O(k_r^2 k_h)$.

In the presence of rays in the initial point-ray collection, we next describe the role of the parameter \mathcal{R}_A in PHA_1 , which restricts the set of rays from \bar{C} that can be cut by any activated hyperplane. Edges of C that do not intersect $\text{bd } S$ can be viewed as intersection points infinitely far away, and Balas and Margot [7] show that this view suffices when extending $\text{CutRegion}(\beta, \mathcal{P}, \mathcal{R})$ with only points to a valid system with both points and rays. However, when generating GICs, rays play a fundamentally different role than points. In particular, when activating a new hyperplane cuts some $p \in \bar{\mathcal{P}}$, the new intersection points that are created are either on the SIC or are deeper relative to p (see Theorem 3 in [7]). In contrast, any intersection points created as a result of cutting a ray from $\bar{\mathcal{R}}$ will all lie on the SIC, as shown in Proposition 10.

Proposition 10. *Suppose $r \in \bar{\mathcal{R}}$ is a ray that does not intersect $\text{bd } S$, and H is an activated hyperplane that intersects r at $v := r \cap H$. Denote by Y the set of new rays on H that originate from v . Then the subset of rays Y' in Y that intersect $\text{bd } S$ create intersection points on the hyperplane $\alpha_0^\top x = \beta_0$, where $\alpha_0^\top x \geq \beta_0$ is the SIC from \bar{C} and S . The rays r' in $Y \setminus Y'$ satisfy $\alpha_0^\top r' = 0$.*

Proof. Let $\bar{\mathcal{R}}^c \subseteq \bar{\mathcal{R}}$ denote the set of initial rays cut by H before $\text{bd } S$. When ray r intersects H , it creates $n - 1 - |\bar{\mathcal{R}}^c|$ new rays emanating from v , which we denote by Y . There are $n - 1$ hyperplanes that define r . Any ray $r' \in Y$ lies on $n - 2$ of these hyperplanes, as well as on H . There exists a ray $\hat{r} \in \bar{\mathcal{R}} \setminus \bar{\mathcal{R}}^c$ that also lies on these $n - 2$ facets, and r' can be written as a nontrivial conic combination of r and \hat{r} . Let F denote the two-dimensional face of \bar{C} defined by all conic combinations of r and \hat{r} .

Suppose \hat{r} intersects $\text{bd } S$ and let $\hat{p} := \hat{r} \cap \text{bd } S$. By the remark above, the half-line $L := \hat{p} + \theta r, \theta \geq 0$ is contained in $\{x : \alpha_0^\top x = \beta_0\} \cap \text{bd } S$, since $r \in \bar{\mathcal{R}}$ does not intersect $\text{bd } S$. Further, L is also contained in F , which is parallel to r, \hat{r} , and r' . Since r' is a nontrivial conic combination of \hat{r} and r , it intersects $\text{bd } S$. The resulting intersection point is on L , which lies on the SIC.

Now suppose \hat{r} does not intersect $\text{bd } S$. Then $\alpha_0^\top r = \alpha_0^\top \hat{r} = 0$. Since r' is a nontrivial

conic combination of r and \hat{r} , it too satisfies $\alpha_0^\top r' = 0$. \square

The consequence of Proposition 10 is that activating hyperplanes on rays of \bar{C} that do not intersect $\text{bd } S$ will create weak points, as these points actually lie on the SIC, not deeper as desired. Thus, the input \mathcal{R}_A to Algorithm 1 will, in our implementation, always be chosen from the set of rays of \bar{C} that intersect $\text{bd } S$, i.e., from $\text{rays}(\bar{C}) \setminus \bar{\mathcal{R}}$. Choosing \mathcal{R}_A arbitrarily may in general lead to invalid cuts; the validity of our choice follows from the results in Section 4. The full details of how we implement PHA₁ to generate cuts are given in Algorithm 2.

Algorithm 2 Generalized Intersections Cuts by PHA₁

Input: Polyhedron P defined by a set of hyperplanes \mathcal{H} ; a vertex \bar{x} of P ; indices of fractional integer variables σ ; hyperplane selection criterion \mathcal{SC} ; objectives \mathcal{O} ; number hyperplanes k_h .

- 1: **function** PHA₁CUTGENERATOR($P, \bar{x}, \sigma, \mathcal{SC}, \mathcal{O}, k_h$)
- 2: $\mathcal{C} \leftarrow \emptyset$.
- 3: **for** $k \in \sigma$ **do**
- 4: $S_k \leftarrow \{x : \lfloor \bar{x}_k \rfloor \leq x_k \leq \lceil \bar{x}_k \rceil\}$.
- 5: $\mathcal{R}_k^\parallel \leftarrow \{r \in \text{rays}(\bar{C}) : \text{dist}(\text{bd } S_k, r) = \infty\}$; $\mathcal{R}_k \leftarrow \mathcal{R}_k^\parallel$; $\mathcal{P}_k \leftarrow \{r \cap \text{bd } S_k : r \in \text{rays}(\bar{C}) \setminus \mathcal{R}_k\}$.
- 6: **for** $h \in \{1, \dots, k_h\}$ **do**
- 7: **for** $k \in \sigma$ **do**
- 8: Choose a hyperplane H to activate according to selection criteria \mathcal{SC} .
- 9: $(\mathcal{P}_k, \mathcal{R}_k) \leftarrow \text{PHA}_1(P, S_k, H, \text{rays}(\bar{C}) \setminus \mathcal{R}_k^\parallel, (\mathcal{P}_k, \mathcal{R}_k))$.
- 10: Add to \mathcal{C} valid cuts by solving (PRLP) with objective types \mathcal{O} , always ensuring \bar{x} is cut.
- 11: **return** \mathcal{C} .

4 Partial hyperplane activation with tilted hyperplanes

An outcome of proving the validity of PHA₁ is that the hyperplane given as an input to Algorithm 1 can be an arbitrary valid hyperplane for P , not necessarily one from the hyperplane description of P . One application of this insight is to choose hyperplanes that only cut rays of \bar{C} that intersect $\text{bd } S$, to avoid the weak intersection points that would otherwise be created as shown in Proposition 10. However, it is not practical to search for arbitrary valid




Figure 2: Tilted activations of the two hyperplanes not defining \bar{C} from Figure 1.

hyperplanes that satisfy specific desired properties such as which particular rays of \bar{C} are or are not cut. Instead, we propose to activate the hyperplanes defining P , but only on a subset of the rays that they intersect. The validity of this follows from a geometric argument, in which Algorithm 1 is applied to a *tilted* hyperplane. Crucially, we show that we never need to explicitly compute this tilted hyperplane; rather, the activation can be computed using the original non-tilted hyperplane, and therefore without incurring any additional computational burden. A byproduct of the tilting theory is that choosing \mathcal{R}_A as in Algorithm 2 is valid.

Using the concept of tilting, in Algorithm 3, we provide an alternative, but not mutually exclusive, approach to Algorithm 2 for using PHA₁. This algorithm uses tilting to reduce the size of the point-ray collection when the number of rays of \bar{C} that are cut is large, which is desirable when seeking cuts stronger than the SIC. When cutting k_r rays and activating k_h hyperplanes, the algorithm creates $O(k_r k_h^2)$ points and rays, which has a clear advantage to the $O(k_r^2 k_h)$ points and rays that are created from Algorithm 2 when k_r is large. In our implementation in Section 6, we compare the two ideas when used individually and together.

The essential motivation for this section is that we wish to activate a given hyperplane using a subset of the rays of \bar{C} that the hyperplane intersects, while ignoring its intersections with the remaining rays. However, not picking this subset with care can easily result in invalid cuts. We will show that a simple rule suffices to guarantee validity: if a ray $r \in \text{rays}(\bar{C})$ has been previously cut and a hyperplane being activated intersects r , then r must be included

in the rays that are cut when performing the activation. Our idea is illustrated in Figure 2 on the same example as in Figure 1. The panel on the left shows a partial activation of H_4 , where we choose to intersect the hyperplane with ray r^1 of \bar{C} , but not ray r^2 . This involves only adding the intersection point obtained from intersecting H_4 with r^1 , while keeping the original intersection point from $r^2 \cap \text{bd } S$ intact. The picture illustrates how this corresponds to a tilting of H_4 . In the next panel, we activate hyperplane H_5 . Though we are not obliged to, we select to cut ray r^2 . However, because r^1 was previously cut, our rule states that we *must* cut the ray r^1 when activating H_5 . Ignoring that ray (i.e., not adding the intersection point on the x_1 axis) results in an improper point-ray collection, as shown in Appendix C.

We now describe the tilting operation formally. Let H be a hyperplane defining P . We assume that H is not tight at \bar{x} for ease of exposition; the degenerate case follows similar reasoning and is handled explicitly in Appendix B.

The hyperplane H can be uniquely defined by the n affinely independent points or rays obtained by its intersection with \bar{C} . For $r \in \text{rays}(\bar{C})$, define $d(H, r) := \min\{\text{dist}(H, r), -\text{dist}(H, -r)\}$ as the distance along r (or $-r$) to H . Let $v(H, r)$ denote the intersection of r (or $-r$) with H :

$$v(H, r) := \begin{cases} \bar{x} + d(H, r) \cdot r & \text{if } d(H, r) \neq \infty, \\ r & \text{otherwise.} \end{cases}$$

Then the affine hull of $\{v(H, r) : r \in \text{rays}(\bar{C})\}$ is precisely H . The affine independence of these points and rays is a direct result of the affine independence of the rays of \bar{C} .

To tilt H , we need only to change some of these points; there are many ways to tilt, so we refer to our particular method as *targeted tilting*. As before, let $\text{rays}(H)$ denote the rays of \bar{C} that are intersected by H before $\text{bd } S$, i.e., those rays $r \in \text{rays}(\bar{C})$ for which $\text{dist}(H, r) < \text{dist}(\text{bd } S, r)$. Let $\text{rays}' \subseteq \text{rays}(H)$ be an arbitrary subset of these rays that we select to be cut. Overloading notation again, we will use $v(\text{bd } S, r)$ to denote the intersection point of ray $r \in \text{rays}(\bar{C})$ with $\text{bd } S$, where $v(\text{bd } S, r) = r$ if the ray does not intersect $\text{bd } S$. The *targeted tilting* of H (with respect to rays') is defined as the unique hyperplane \tilde{H} going

through $v(\text{bd } S, r)$ for all $r \in \text{rays}(H) \setminus \text{rays}'$, and $v(H, r)$ otherwise.

The first thing we observe is that \widetilde{H}^+ is valid for P , as it is a relaxation of H^+ . This immediately implies, by Theorem 5, that we can activate \widetilde{H} using Algorithm 1; i.e., given a proper point-ray collection $(\mathcal{P}, \mathcal{R})$, the point-ray collection outputted by $\text{PHA}_1(P, S, \widetilde{H}, \text{rays}(\bar{C}), (\mathcal{P}, \mathcal{R}))$ is proper. This requires knowing \widetilde{H} explicitly. We next show that computing \widetilde{H} is actually unnecessary: we can calculate exactly the same points and rays via $\text{PHA}_1(P, S, H, \text{rays}', (\mathcal{P}, \mathcal{R}))$.

Theorem 11. *Let H be a hyperplane defining P and $\text{rays}^c \subseteq \text{rays}(\bar{C})$ the set of rays previously cut, \widetilde{H} be a hyperplane obtained via a targeted tilting of H on a set of rays $\text{rays}' \subseteq \text{rays}(H)$ such that $\text{rays}^c \cap \text{rays}(H) \subseteq \text{rays}'$. Given a proper point-ray collection $(\mathcal{P}, \mathcal{R})$, the point-ray collection returned by $\text{PHA}_1(P, S, H, \text{rays}', (\mathcal{P}, \mathcal{R}))$ is proper.*

Proof. We show that $\text{PHA}_1(P, S, H, \text{rays}', (\mathcal{P}, \mathcal{R}))$ produces the same point-ray collection as $\text{PHA}_1(P, S, \widetilde{H}, \text{rays}(\bar{C}), (\mathcal{P}, \mathcal{R}))$. In particular, we analyze each iteration of step 2 of Algorithm 1 and show that both algorithms produce the same point-ray collection after each iteration. Observe that the only rays of \bar{C} processed by the algorithm are those that belong to \mathcal{R}_A and are cut by H prior to intersecting $\text{bd } S$. Our targeted tilting method implies that $\text{rays}(\widetilde{H}) = \text{rays}'$, all of which are intersected by both H and \widetilde{H} before $\text{bd } S$. The result is that activation is performed on the same set of rays rays' for both H and \widetilde{H} . Let us fix a ray $r \in \text{rays}'$ being processed. Throughout, the reader may find it useful to refer to Figure 3, which illustrates the two-dimensional face of \bar{C} defined by rays $r \in \mathcal{R}'$ and $r' \in \text{rays}(H) \setminus \mathcal{R}'$, the hyperplane H and its tilted version \widetilde{H} , as well as the result of a prior activation of a hyperplane H' that was not tilted on ray r' .

We first discuss the set of points and rays that are added by Algorithm 1 in each of the cases, starting at step 4, in which the edges emanating from $v(H, r)$ are considered. Note that there is one edge per ray $r' \in \text{rays}(\bar{C}) \setminus \{r\}$, because \bar{C} is a simplicial cone.

Consider an edge corresponding to a ray r' whose intersection with H and \widetilde{H} remains unchanged ($v(H, r') = v(\widetilde{H}, r')$), i.e., for $r' \in \text{rays}' \cup (\text{rays}(\bar{C}) \setminus \text{rays}(H))$. This is an edge that exists in both $H \cap \bar{C}$ and $\widetilde{H} \cap \bar{C}$. Since the edge is the same, the resulting intersection




Figure 3: Reference illustration for the proof of Theorem 11, activating a tilted version of H (after an earlier activation of a hyperplane H').

point or ray will be the same in either case.

Now suppose that $r' \in \text{rays}(H) \setminus \text{rays}'$, as in Figure 3. When activating H , the corresponding edge is skipped in step 5, because $r' \in \text{rays}(H)$ so that $v(H, r') \in \text{int } S$. When activating \widetilde{H} , $v(\widetilde{H}, r') = v(\text{bd } S, r')$ (the intersection point of r' with $\text{bd } S$) instead of $v(H, r')$. Thus, the new intersection point or ray created is precisely $v(\text{bd } S, r')$. However, because $r' \notin \text{rays}'$, it must be the case that $r' \notin \text{rays}^c$, i.e., the ray r' has not been cut, which implies that $v(\text{bd } S, r')$ already belongs to $(\mathcal{P}, \mathcal{R})$. This proves that the same points and rays are added whether H or \widetilde{H} is activated.

Lastly, we show that the same points and rays are removed in step 3. Consider the two-dimensional face of \bar{C} defined by r and r' . As before, when the edge of H and \widetilde{H} is the same on this face, exactly the same points and rays of $(\mathcal{P}, \mathcal{R})$ that lie on this face are cut when activating either hyperplane. Thus, assume that $r' \in \text{rays}(H) \setminus \text{rays}'$. Since \widetilde{H} is weaker than H , anything cut by \widetilde{H} is certainly cut by H .

It remains to show the converse, that any point or ray removed by H is also removed by \widetilde{H} . Observe (as depicted in Figure 3) that the only part of the face defined by r and r' that is cut by H but not by \widetilde{H} is in $F := \text{conv}(v(H, r), v(H, r'), v(\text{bd } S, r'))$.¹ Since $\text{relint}(F) \subseteq \text{int } S$, it contains no points or rays from $(\mathcal{P}, \mathcal{R})$, completing the proof. \square

¹There is one intersection point/ray in F : $v(\text{bd } S, r')$. It is true that, when activating a prior hyperplane H' , $v(H', r')$ could coincide with $v(\text{bd } S, r')$, and in step 3, this would be cut by H but not by \widetilde{H} . However, this point/ray is a duplicate of the intersection point of r' with $\text{bd } S$, which would remain in $(\mathcal{P}, \mathcal{R})$.

One effect of Theorem 11 is that it is valid to simply ignore all intersections of hyperplanes with rays in $\bar{\mathcal{R}}$, as is done in Algorithm 2. The activation performed in step 9 of Algorithm 2 is valid because the rays in $\bar{\mathcal{R}}$ are always ignored, so that the conditions of Theorem 11 apply.

We now operationalize the above theorem in a different way. The motivation for targeted tilting comes not only from Proposition 10, to avoid cutting the rays in $\bar{\mathcal{R}}$, but also from our desire to find GICs strictly stronger than the SIC from the same P_I -free convex set S . This may require cutting many rays of \bar{C} . However, using Algorithm 2 for this purpose could involve unnecessarily cutting some rays of \bar{C} multiple times. This will happen when we have already cut a certain ray of \bar{C} but a hyperplane being subsequently activated intersects it again; not only might this effort be wasted, but also it may create redundant or weak intersection points. As an alternative, we present the targeted tilting algorithm, Algorithm 3, which is tailored to the task of cutting the initial intersection points and rays more efficiently.

Algorithm 3 Generalized Intersection Cuts by PHA₁ and Targeted Tilting

Input: Polyhedron P defined by a set of hyperplanes \mathcal{H} ; a vertex \bar{x} of P ; indices of fractional integer variables σ ; hyperplane selection criterion \mathcal{SC} ; objectives \mathcal{O} ; set $\mathcal{R}_A \subseteq \text{rays}(\bar{C})$ for the rays of \bar{C} to cut.

```

1: function PHA1CUTGENERATORWITHTARGETEDTILTING( $P, \bar{x}, \sigma, \mathcal{SC}, \mathcal{O}, \mathcal{R}_A$ )
2:   for  $k \in \sigma$  do
3:      $S_k \leftarrow \{x : \lfloor \bar{x}_k \rfloor \leq x_k \leq \lceil \bar{x}_k \rceil\}$ .
4:      $\mathcal{R}_k \leftarrow \{r \in \text{rays}(\bar{C}) : \text{dist}(\text{bd } S_k, r) = \infty\}$ ;  $\mathcal{P}_k \leftarrow \{r \cap \text{bd } S_k : r \in \text{rays}(\bar{C}) \setminus \mathcal{R}_k\}$ .
5:      $j \leftarrow 0$ ;  $\text{rays}' \leftarrow \emptyset$ .
6:     for the next ray  $r \in \mathcal{R}_A$  do ▷ Assume sorted by decreasing reduced cost
7:        $j \leftarrow j + 1$ ;  $\text{rays}' \leftarrow \text{rays}' \cup \{r\}$ ;  $\mathcal{R}_A \leftarrow \mathcal{R}_A \setminus \{r\}$ .
8:       for  $k \in \sigma$  do
9:          $\mathcal{H}^j \leftarrow \{H \in \mathcal{H} \setminus \bar{\mathcal{H}} : \text{dist}(H, r) < \text{dist}(\text{bd } S_k, r)\}$ .
10:      if  $\mathcal{H}^j \neq \emptyset$  then
11:        Select a hyperplane  $H_j \in \mathcal{H}^j$  according to selection criterion  $\mathcal{SC}$ .
12:         $(\mathcal{P}_k, \mathcal{R}_k) \leftarrow \text{PHA}_1(P, S_k, H_j, \mathcal{R}', (\mathcal{P}_k, \mathcal{R}_k))$ .
13:      if  $\log(j) \in \mathbb{Z}$  or  $\mathcal{R}_A = \emptyset$  then
14:        Add to  $\mathcal{C}$  valid cuts by solving (PRLP) with objective types  $\mathcal{O}$ , always ensuring  $\bar{x}$ 
is cut.
```

Algorithm 3 focuses on cutting one ray of \bar{C} at a time. The hyperplane chosen for each ray is tilted in step 12 in a way that avoids cutting all other rays, other than the ones that

have been previously cut, as required. However, this may result in cuts that do not get monotonically stronger as more rays of \bar{C} are cut, as might be expected. This is due to our observation that cutting a ray multiple times may lead to the addition of weak intersection points to the collection. We counteract this in step 14 of Algorithm 3, by performing cut generation before all hyperplanes have been chosen and activated. To reduce computational effort, this step is only performed $\lceil \log(n) \rceil$ times.

Although targeted tilting may reduce the number of weak intersection points that result from cutting the same rays of \bar{C} repeatedly, it may also miss improving the point-ray collection in certain directions. This creates an inherent tradeoff when deciding whether to use PHA₁ as part of Algorithm 2 or Algorithm 3. We discuss this as part of our numerical study in Section 6.

5 Implementation choices for PHA₁

Algorithms 2 and 3 both involve steps in which a hyperplane is selected and then, after the points and rays have been collected, a set of objective directions is used for generating GICs from the resulting (PRLP). In this section, we discuss the choices we make for these steps and provide theoretical motivation for the decisions.

5.1 Choosing hyperplanes to activate

We first state the three criteria we consider for choosing hyperplanes to activate as the parameter \mathcal{SC} (used in Algorithm 2 step 8 and Algorithm 3 step 11). For an index $k \in I$, we consider optimizing over the S_k -closure, i.e., $\text{conv}(P \setminus \text{int } S_k)$.

(H1) Choose the hyperplane that first intersects some ray $r \in \text{rays}(\bar{C}) \setminus \bar{\mathcal{R}}$ before $\text{bd } S_k$. This corresponds to pivoting to the nearest neighbor from \bar{x} along the ray r , which would be the same hyperplane selected in the procedure from [7].

(H2) Choose the hyperplane that yields a set of points with highest average depth. Here, *depth* is calculated as the Euclidean distance to the SIC (using the same split set). (H2) seeks a set of points that are far from the cut we are trying to improve upon. The idea is that the resulting cuts will then be deeper as well.

(H3) Choose the hyperplane creating the most final intersection points. A *final* intersection point is defined as follows.

Definition 12. *Suppose $(\mathcal{P}_k, \mathcal{R}_k)$ is a proper point-ray collection. An intersection point in \mathcal{P}_k or a ray in \mathcal{R}_k is final (with respect to S_k) if it belongs to P and $\text{bd } S_k$, meaning it cannot be cut away by any valid hyperplane activations.*

We denote the set of final intersection points by \mathcal{P}_k^F and final rays by \mathcal{R}_k^F .

Proposition 13. *Suppose $C = P$. Then, the intersection points in \mathcal{P}_k^F define vertices of the S_k -closure and the rays \mathcal{R}_k^F define extreme rays of S_k -closure.*

Proof. The S_k -closure is defined as $P_k = \text{conv}(P \setminus \text{int } S_k)$. Therefore, the points $p \in \mathcal{P}^F$ are in the S_k -closure since they belong to $P \cap \text{bd } S_k$. Furthermore, the rays in \mathcal{R}^F are in the S_k -closure since they are extreme rays of P that do not intersect $\text{bd } S_k$.

Consider an arbitrary side of the split disjunction, $S_k^0 := \{x : x_k = \lfloor \bar{x}_k \rfloor\}$. Suppose that a vertex $p = r \cap S_k^0$ in \mathcal{P}_k is not a vertex of the S_k -closure. Then it can be written as a convex combination of other vertices of $P \cap S_k^0$ that are also intersection points created from edges of P . This implies that r , an edge of P , can be written as a conic combination of edges in P that intersect S_k^0 , which is a contradiction. Showing that rays in \mathcal{R}_k are extreme rays of S_k -closure follows similar reasoning. \square

Thus, (H3) targets a set of points that lead to facet-defining inequalities for the set $\text{conv}(P \setminus \text{int } S)$. It turns out to be useful to distinguish between intersection points that are in P and those that are not.

5.2 Choosing objective functions

Even if we know that strong cuts exist from a given point-ray collection, it remains to generate these GICs using (PRLP). To do this, we need to appropriately choose objective function coefficient vectors. We consider the following objective directions:

- (*R*) Ray directions of \bar{C} (i.e., the initial intersection points)
- (*V*) Vertices that are generated from hyperplane activations
- (*T*) Intersection points obtained for S_k (the *tight point* heuristic)
- (*S*) Intersection points from other splits

The first set of objectives, (*R*), is incentivized by the observation that obtaining GICs that strictly dominate the SIC from the same P_I -free convex set requires cutting many of the initial intersection points. This observation is made concrete in the Theorem 17 in Appendix A; we show that finding a strictly dominating cut requires reducing the dimension of the convex hull of the initial intersection points and rays. Another objective typically used for a cut-generating linear program is to minimize $\bar{x}^\top \alpha$, to obtain the most violated cut with respect to \bar{x} . We generalize this with (*V*), by optimizing to find the most violated cuts with respect to each of the vertices created on \bar{C} by hyperplane activations.

The third set of directions, (*T*), is referred to as the tight point heuristic as we are simply trying to find a cut tight on each of the rows of (PRLP) corresponding to points. In other words, we minimize $\alpha^\top p$ for every $p \in \mathcal{P}$ from the current split. Obviously, we will not be able to cut away p ; the purpose of this objective function is to place a cut as close to the chosen intersection point as possible. The fourth set, (*S*), capitalizes on the fact that multiple split sets are considered simultaneously in practice, so that we can share information across the sets to obtain stronger cuts.

The sets of objective directions (*T*) and (*S*) switch the perspective typically used for linear programs that generate cuts. Instead of finding a cut with maximal violation, the cuts

we obtain from these objectives aim to approximate the convex hull of intersection points and rays from every direction, thereby obtaining as many of the facets of that convex hull as possible and a set of cuts with a wider diversity of angles, which is a desirable computational quality.

Next, we discuss the theoretical motivation for (T). Consider optimizing over the S_k -closure: $\min\{c^\top x : x \in \text{conv}(P \setminus \text{int } S_k)\}$. We show that intersection points can be used to find an optimal solution to this problem without using (PRLP), and we address how the heuristic (T) aims to obtain the bound implied by this optimal solution.

Theorem 14 shows what bounds can be computed on the optimal value over the S_k -closure using the point-ray collection when $P \subseteq C$, i.e., not all hyperplanes may have been activated. We define $\bar{p}^k \in \arg \min_{p \in \mathcal{P}_k^F} c^\top p$ (when \mathcal{P}_k^F is nonempty) and $\underline{p}^k \in \arg \min_{p \in \mathcal{P}_k} c^\top p$. Let $\bar{z} = c^\top \bar{p}^k$ (defined to be $+\infty$ when \mathcal{P}_k^F is empty) and $\underline{z} = c^\top \underline{p}^k$.

Theorem 14. *Suppose \mathcal{P}_k^F is nonempty. If $\underline{z} < \bar{z}$, then \underline{z} is a lower bound and \bar{z} is an upper bound on the minimum over the S_k -closure. Otherwise, $\bar{z} \leq \underline{z}$ and \bar{z} is the minimum over the S_k -closure.*

Proof. By Proposition 13, \bar{p}^k is in the S_k -closure. Therefore, \bar{z} always provides an upper bound on $c^\top p^*$, where p^* is a minimum over the S_k -closure.

When $\underline{z} < \bar{z}$, $\underline{z} = \min\{c^\top x : x \in C \cap \text{bd } S_k\}$ provides a valid lower bound on $c^\top p^*$, since $P \cap \text{bd } S_k \subseteq C \cap \text{bd } S_k$. On the other hand, when $\bar{z} \leq \underline{z}$, we have $\bar{z} = \min\{c^\top x : x \in C \cap \text{bd } S_k\}$, which implies that \bar{z} provides a lower bound on $c^\top p^*$. Therefore, it must be a minimum over the S_k -closure. \square

Corollary 15 follows immediately for the special case when $C = P$.

Corollary 15. *If $C = P$, then \bar{p}^k is an optimal solution over the S_k -closure.*

Thus, \underline{p}^k is readily available from \mathcal{P}_k and implies a lower bound on the value of the optimal solution over the S_k -closure. The same bound \underline{z} implied by \underline{p}^k can be obtained through GICs, but it may require many cuts generated from (PRLP). We may be able to

obtain one inequality tight at \underline{p}^k by using \underline{p}^k itself as the objective coefficient vector, since $\alpha^\top \underline{p}^k \geq \beta$ for all (α, β) feasible to (1), and there will exist a facet-defining inequality $\bar{\alpha}^\top x \geq \bar{\beta}$ for \mathcal{G} satisfying $\bar{\alpha}^\top \underline{p}^k = \bar{\beta}$. Note that for validity of this inequality to be guaranteed, we must additionally verify that the inequality cuts a point $v \in \mathcal{K}'$. However, even if this is satisfied, it is unlikely that the one cut will imply the bound $c^\top x \geq \underline{z}$ on the objective value. In the worst case, n facets of \mathcal{G} that are tight at \underline{p}^k may be required to obtain this bound via cuts. To get these other facets tight at \underline{p}^k , we can use points in \mathcal{P}_k that lie close to \underline{p}^k as the objective vectors. This is precisely what we do in approach (T), though we do not only use \underline{p}^k and points in its vicinity, but also other points from \mathcal{P}_k to encourage diversity of the cut collection.

Finally, we give more details and motivation for the objective directions (S). When using a single split set S_k , no intersection point generated from that split can be cut by any inequality generated through (PRLP) from the point-ray collection for S_k . However, different splits (more generally, different P_I -free convex sets) give rise to different point-ray collections, and the buildup of intersection points and rays can and should be done in parallel for several splits. The first reason for this is computational efficiency: one can intersect an edge with the boundaries of more than one split set and store these intersection points or rays separately. The second reason is that, when using multiple split disjunctions, an intersection point on the boundary of one split set may lie in the interior of another and hence can be cut away by a facet of the point-ray collection from this second split. With this in mind, let σ be the indices of a set of integer variables that are fractional at \bar{x} , i.e., $\sigma \subseteq \{j \in I : \bar{x}_j \notin \mathbb{Z}\}$. For every intersection point p generated from intersecting an edge of C with the boundary of some split disjunction, we follow the edge to find the *last* split disjunction S_k , $k \in \sigma$, that this edge intersects. If p lies in the interior of S_k , then we will use p as an objective for (PRLP) with feasible region determined by the point-ray collection from S_k .

6 Computational results

This section contains the results from computational experiments with PHA_1 as used in Algorithm 2 from Section 3 and Algorithm 3 from Section 4. The purpose of these experiments is exploratory; that is, we seek conditions and implementation choices that lead to strong GICs. To do this, we evaluate the effect of a variety of parameters, including those discussed in Section 5, and identify structural properties of instances that can be taken advantage of to find stronger cuts. In particular, our results indicate that it is beneficial to use objective functions targeting GICs that are tight on the points and rays in the point-ray collection. Another observation from our experiments is that PHA_1 leads to few final points (as defined in Section 5.1), which, in combination with the intuitive importance of these points, indicates a limitation of our procedure that can perhaps be a target for improvement via future research.

We test Algorithm 2 and Algorithm 3 when used independently, as well as in combination. The purpose of this is to test the strength of targeted tilting relative to non-tilted hyperplanes as in Algorithm 2. When used together, we first perform Algorithm 3 and afterwards perform non-tilted activations. This is equivalent to inserting the steps of Algorithm 3 before step 6 of Algorithm 2.

6.1 Experimental setup

Parameters. Several parameters are fixed throughout the experiments, while others are varied, summarized in Table 2 and elaborated on below.

Table 2: Parameters that are varied and the values considered.

Parameter	Values Considered
Hyperplane scoring function, \mathcal{SC}	{H1, H2, H3}
Objective functions, \mathcal{O}	{ (R) , (V) , (T) , (S) }
Targeted tilting	{On, Off}
# non-tilted hyperplanes, k_h	{0, 1, 2, 3, 4}

The fixed parameters are as follows. The sets used for cut generation are all the simple splits. The time limits are one hour per set of parameters and at most five seconds per objective function for each time (PRLP) is resolved. The instances are not preprocessed, and presolve is turned off for (PRLP), as using presolve mildly reduced the average gap closed and led to erratic solver behavior, such as occasional crashes. At most 1,000 GICs are generated per instance (a limit seldom attained). As the cardinality of \mathcal{P} may be large, using all the intersection points as objective directions for (PRLP) may be prohibitively expensive; as a result, we limit the number of points used for objective functions to 1,000 per instance. The overall zero tolerance is set to 10^{-7} , while the fractionality tolerance is set to 10^{-3} ; i.e., a component is considered fractional if it is at least 10^{-3} from the nearest integer. To avoid numerical instability, cuts with dynamism higher than 10^6 are rejected, where dynamism is the ratio of the largest and smallest cut coefficients. A one hour time limit is used per instance.

The parameters we modify are used to evaluate the effect of the choice of hyperplanes to be activated and objective functions to be used for (PRLP). In Algorithm 2, we vary k_h from 0 to 4, where $k_h = 0$ implies we do not perform any non-tilted activations. Section 6.3 contains computational results relating to the hyperplane activation rules described in Section 5.1. Section 6.4 presents the results for the objective functions discussed in Section 5.2.

Cut generation. The indices σ used for defining the split sets for cut generation are those of the fractional integer variables at \bar{x} . First, we generate one SIC for each S_k , $k \in \sigma$. We then generate GICs using Algorithms 2 and 3, as described above.

We generate cuts in the nonbasic space. The vertex \bar{x} of \bar{C} in this space corresponds to the zero vector, and the j th ray of \bar{C} in the nonbasic space has a single nonzero j th entry. As a result, the intersection points or rays defining SICs (one pivot from \bar{x}) and GICs from PHA on \bar{C} (two pivots from \bar{x}) have one and two nonzero coordinates, respectively. This feature leads to a sparse constraint coefficient matrix of (PRLP), which improves the

time complexity of implementing Algorithm 1. In addition, it implies that it is sufficient to consider cuts with $\beta = 1$ to obtain all valid inequalities that cut \bar{x} .

The objective functions tested are those discussed in Section 5.2. Assuming all objective functions are utilized, the order in which these are used in the code is (S) , (T) , (R) , then (V) .

As our procedure is not recursive, all the generated cuts have split rank 1 with respect to P , and in particular, they all valid for the first split closure. Neither the SICs nor the GICs are strengthened through the integral nonbasic variables, as conventional strengthening approaches require information not available through (PRLP).

Environment. All algorithms are implemented in C++ in the COIN-OR framework [29], using `Clp version 1.16` as the linear programming solver. The machine used is a shared 64-bit `PowerEdge R515` with 48GB of memory and twelve `AMD Opteron 4176` processors clocked at 2.4GHz. The operating system is `Red Hat Enterprise Linux 7.6` and the compiler is `g++ version 4.8.5 20150623 (Red Hat 4.8.5-36)`. Experiments with instance `mod008` and the second round of cuts were run on a separate shared `SuperServer 6019P-WTR` machine with 512GB of memory running `Oracle Linux 7.5` on 32 `Intel Xeon Gold 6142 2.6GHz` processors.

Cut evaluation. The metric we use to evaluate the cuts obtained from a specific set of parameters is *percent gap closed*. The gap is defined as the difference between the optimal values of the integer program and its linear programming relaxation. Denoting the optimal value of the integer program by z_{IP} , of its linear relaxation by z_{LP} , and of the linear relaxation with cuts added by z' , we have

$$\% \text{ gap closed} := 100 \times (z' - z_{LP}) / (z_{IP} - z_{LP}).$$

The baseline we use is the percent gap closed by SICs, which are the simplest GICs. We then

add GICs along with the SICs to assess what additional effect GICs have on the percent gap closed in the presence of the SICs.

Instance selection. We test forty instances selected from MIPLIB [2, 11, 12, 28] (all versions) based on the following criteria meant to identify small problems so that many different parameter settings can be tested (over 200 in our experiments): (1) The number of rows and number of columns must be no more than 500 each. (2) The instance has to be integer-feasible with $z_{LP} < z_{IP}$, and the gap closed by SICs is not 100%. (3) There must be at least one non-final intersection point created from intersecting the rays of \bar{C} with $\text{bd } S$. (4) The instance must not be known to have 0% gap closed from split cuts based on previous experiments [8, 17]. Criterion 3 exists because we do not cut rays of \bar{C} that do not intersect $\text{bd } S$, so if all intersection points are final, no hyperplanes will be activated.

We modify the `stein15`, `stein27`, and `stein45` instances to reduce symmetry, by replacing the objective $\sum_{j=1}^n x_j$ with $\sum_{j=1}^n jx_j$. We also remove the cardinality constraint $\sum_{j=1}^n x_j \geq (n-1)/2$, as this is not present in the initial formulation of these instances [24]. In addition, we exclude instances `go19` and `pp08aCUTS`: the former because its continuous relaxation solves exceptionally slowly, and the latter because it is simply a strengthened version of `pp08a`.

6.2 Point-ray collection statistics

The following statistics are averaged across all the instances and all the splits used for each instance. First, of the rays of the initial cone \bar{C} , 55% belong to $\bar{\mathcal{R}}$, i.e., do not intersect $\text{bd } S$. This demonstrates the impact of Proposition 10 and our resulting decision to always set the parameter \mathcal{R}_A of Algorithm 1 to a subset of the rays of \bar{C} that intersect $\text{bd } S$. An additional 5% of the initial rays on average lead to final intersection points (with a range of 0% to 34%), leaving possibly few rays that can be cut for some instances. For six instances, less than 10% of the rays are able to be cut.

We can also provide an idea of the number of rows in (PRLP) in practice, i.e., the number of points and rays we generate, for the best combination of parameters for each instance. On average across all splits, each (PRLP) contains about 2,000 points 200 rays. There exist instances with split sets leading to as few as 2 and as many as 46,000 points, and as few as 0 and as many as 4,000 rays. On average, fewer than 10% of the points and rays in the collection are final (with a range of 0.1% to 75%). Figure 4 plots the number of generated points across instances (for the best parameter combination for each instance) as a function of the number of rays of \bar{C} that can be cut; we observe a quadratic relationship, as predicted by our analysis in Section 3.

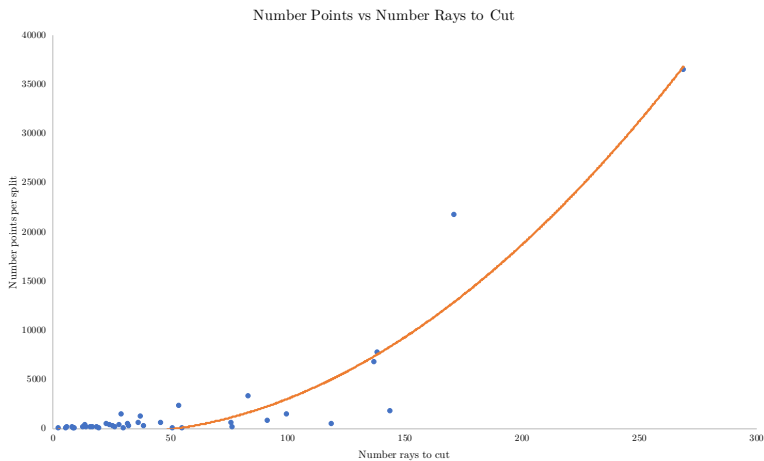


Figure 4: Number generated points versus number of rays that can be cut.

6.3 Effect of hyperplanes activated

We now analyze how the choice of hyperplanes to activate, as described in Section 5.1, affects the percent gap closed by GICs. In addition, in Section 4, we discussed that, theoretically, neither Algorithm 2 nor targeted tilting is strictly stronger than the other. As a result, the computational experiments we perform are necessary to assess the practical strength of each approach.

Table 3: Best percent gap closed by hyperplane activation choice for the 30 instances with any improvement over SICs across all settings. We highlight the cells that achieve the best result, but only if it is the first time this result is attained using that particular combination of Algorithms 2 and 3. Some values differ in the thousandths digit.

Instance	SIC	Best	Only Alg. 3			Alg. 3 with Alg. 2				Only Alg. 2			
			H1	H2	H3	T+0	T+1	T+2	T+3	T+4	+1	+2	+3
bell3a	44.74	59.52	58.94	58.49	58.94	58.94	58.94	58.94	58.94	59.07	59.52	59.52	59.52
bell3b	44.57	60.34	51.39	53.45	51.39	53.45	53.45	53.45	53.45	60.29	60.30	60.34	60.34
bell4	23.37	26.54	23.82	26.35	23.83	26.36	26.35	26.54	26.54	24.67	24.67	25.26	25.26
bell5	14.53	85.37	17.58	19.99	17.58	19.99	19.99	19.99	19.99	85.37	85.37	85.37	85.37
blend2	16.04	19.78	19.78	17.23	19.78	19.78	19.78	19.78	19.78	16.33	17.36	17.36	17.36
bm23	5.92	11.06	8.79	9.78	8.92	9.78	9.78	9.78	9.78	8.00	10.65	11.06	11.06
egout	51.57	52.11	51.59	51.60	51.59	51.60	51.60	51.60	51.60	52.11	52.11	52.11	52.11
gt2	83.13	84.26	83.13	83.13	83.13	83.13	83.13	83.13	83.13	84.26	84.26	84.26	84.26
k16x240	7.56	7.71	7.56	7.65	7.56	7.65	7.65	7.65	7.65	7.71	7.71	7.71	7.71
lseu	4.57	4.65	4.65	4.57	4.65	4.65	4.65	4.65	4.65	4.57	4.65	4.65	4.65
mas74	3.30	4.31	4.12	3.90	4.06	4.12	4.12	4.12	4.12	4.30	4.31	4.31	4.31
mas76	2.37	2.49	2.37	2.37	2.37	2.37	2.37	2.37	2.37	2.38	2.49	2.49	2.49
mas284	0.38	0.51	0.45	0.51	0.45	0.51	0.51	0.51	0.51	0.39	0.39	0.39	0.39
misc05	3.60	3.62	3.60	3.60	3.60	3.60	3.60	3.60	3.60	3.60	3.60	3.60	3.62
mod008	1.30	1.37	1.31	1.30	1.31	1.31	1.31	1.31	1.31	1.31	1.37	1.37	1.37
mod013	4.41	7.37	4.72	7.28	4.72	7.28	7.37	7.37	7.37	4.41	4.72	4.72	7.28
modglob	9.59	14.02	10.00	10.26	10.00	10.26	10.26	10.26	10.26	10.00	14.02	14.02	14.02
p0033	1.83	5.19	2.59	5.19	2.59	5.19	5.19	5.19	5.19	1.83	2.59	2.59	2.59
p0282	3.67	5.12	4.09	4.50	4.09	4.50	4.77	4.77	4.77	4.65	4.81	4.81	5.12
p0291	27.78	40.12	27.93	34.21	27.93	34.21	34.21	34.21	34.21	39.80	40.12	40.12	40.12
pipex	0.81	1.43	1.43	1.43	1.43	1.43	1.43	1.43	1.43	1.43	1.43	1.43	1.43
pp08a	51.44	54.46	53.42	53.89	53.42	53.89	53.89	53.89	53.89	52.47	54.41	54.46	54.46
probportfolio	25.14	25.28	25.19	25.28	25.19	25.28	25.28	25.28	25.28	25.14	25.14	25.14	25.16
sample2	5.86	13.14	5.86	13.14	5.86	13.14	13.14	13.14	13.14	5.86	5.86	5.86	5.86
sentoy	10.38	14.00	12.19	12.45	12.39	12.45	12.45	12.45	12.45	10.38	13.89	13.89	14.00
stein15_nosym	50.00	58.33	55.78	50.00	55.78	55.78	57.17	57.17	57.17	50.00	50.00	58.33	58.33
stein27_nosym	7.41	8.78	8.78	8.61	8.78	8.78	8.78	8.78	8.78	7.41	7.86	7.86	7.86
stein45_nosym	7.10	7.55	7.36	7.55	7.36	7.55	7.55	7.55	7.55	7.51	7.51	7.51	7.51
vpm1	10.00	10.18	10.00	10.18	10.00	10.18	10.18	10.18	10.18	10.00	10.00	10.00	10.00
vpm2	10.18	11.71	11.18	11.37	11.18	11.37	11.38	11.38	11.38	11.04	11.19	11.19	11.71
Average	17.75	23.34	19.32	19.98	19.33	20.28	20.33	20.35	20.35	21.88	22.41	22.72	22.84
Wins			3	7	3	9	11	12	12	4	10	15	19

Table 3 shows the effect of hyperplane activation choices on gap closed for the 30 instances in which GICs close additional gap over SICs. All possible combinations of objective functions (discussed in Section 5.2) for (PRLP) are used. The best percent gap closed is shown per instance in: column 2 for SICs; column 3 for GICs across all parameter settings; columns 4 through 6 for Algorithm 3 (targeted tilting) using each of the hyperplane scoring functions (H1), (H2), (H3); columns 7 through 11 for Algorithm 3 used in conjunction with Algorithm 2 (for $k_h \in \{0, 1, 2, 3, 4\}$, indicated by the column headers “T+ k_h ”); and columns 12 through 15 for Algorithm 2 (with $k_h \in \{1, 2, 3, 4\}$, indicated by column headers “+ k_h ”). Note that column 7 (T+0) is simply the best result from columns 4 through 6. For each parameter setting, the last two rows give the average percent gap closed and the number of instances for which that parameter setting achieves the best percent gap closed.

The results show that Algorithm 2 (with $k_h = 4$) closes 98% of the best percent gap closed, and that it achieves the best result across all settings for 63% of the instances without using any targeted tilting. Using $k_h = 1$ already achieves 96% of the best percent gap closed, but this comes from strong results from just a few instances. The marginal impact of activating more hyperplanes diminishes as k_h increases.

The table also shows that the effect of targeted tilting is mixed. Targeted tilting alone wins 9 times, meaning it achieves the best result for 30% of the instances. On 8 of these instances, Algorithm 2 on its own does not attain the best result and only closes 0.2% of the integrality gap averaged across the 8 instances, compared to 1.6% from column “T+0”. On the other hand, targeted tilting may lead to weaker cuts, as reflected in the fact that the best average percent gap closed is achieved by Algorithm 2 on its own, and that Algorithm 2 seems to have a relatively insignificant impact on improving the results once targeted tilting has been used. We see that for the 18 instances on which Algorithm 2 uniquely wins (with $k_h = 4$, from column “+4”), Algorithm 3 performs poorly. Thus, it appears as though the two algorithms are typically not both effective on the same instance.

In our experiments, we find that though final intersection points are targeted, they rep-

resent a small percent (less than 16% on average for these 30 instances) of all generated points. This does not necessarily imply bad cuts: for example, GICs close an additional 8% of the gap for `stein15_nosym` over SICs alone, but only 2 of the 168 points created are final. Nevertheless, this suggests that one promising future direction to obtain stronger cuts is to develop a procedure targeting more final intersection points, as this may lead to more facet-defining inequalities for $\text{conv}(P \setminus \text{int } S)$.

6.4 Evaluating objective function choices

Next, we examine which objective functions for (PRLP) lead to the strongest cuts. Table 4 shows the effect of the possible values of the parameter \mathcal{O} when all other parameters are allowed to vary freely. The second column gives the best percent gap closed across all parameter settings. The next four columns provides the best percent gap closed when each of the objective functions is used alone. The following six columns show the best percent gap closed for all pairs of objective functions, while the last four columns give the values for any three objective functions. The highlighted cells are those that achieve the best percent gap closed, but only when that combination of objectives was required to achieve this; by this, we mean that the same result must not have been achieved with any subset of these objectives. The last row for each column shows the number of instances for which that set of objective functions attains the best percent gap closed.

In our experiments, the objective functions (T) (targeting the points in the point-ray collection) on their own work extremely well, closing nearly all of the possible gap on average, and achieving the best result for 19/30 of the instances. Using (R) (cutting along the directions of the rays of \bar{C}) and (T) together achieves the best result across all parameter settings for 22 of the 30 instances. Moreover, 67% of the cuts that are active at the post-cut optimum come from procedures (R) and (T), on average. Note that the number of possible objectives (R) is n , while the number of objectives (T) depends on the size of the point-ray

Table 4: Best percent gap closed by objective function used for the 30 instances with any improvement over SICs across all settings. Highlighted cells indicate objectives (or combinations of objectives) achieving best result. For the combinations, the cell is highlighted only if using the combination is actually necessary to achieve the result. Some values differ in the thousandths digit.

Instance	Best	R	S	T	V	R+S	R+T	R+V	S+T	S+V	T+V	R+S+T	R+S+V	R+T+V	S+T+V
bell3a	59.52	59.52	57.14	59.52	59.50	59.52	59.52	59.52	59.52	59.52	59.52	59.52	59.52	59.52	59.52
bell3b	60.34	60.34	52.71	60.34	60.34	60.34	60.34	60.34	60.34	60.34	60.34	60.34	60.34	60.34	60.34
bell4	26.54	26.34	26.19	26.54	25.90	26.34	26.54	26.34	26.54	26.27	26.54	26.54	26.34	26.54	26.54
bell5	85.37	85.37	65.85	85.37	85.37	85.37	85.37	85.37	85.37	85.37	85.37	85.37	85.37	85.37	85.37
blend2	19.78	16.56	17.70	17.36	19.35	19.71	17.36	19.35	17.70	19.78	19.35	19.71	19.78	19.35	19.78
bm23	11.06	9.48	10.67	11.04	10.21	10.68	11.04	10.43	11.06	10.68	11.04	11.06	10.68	11.04	11.06
egout	52.11	52.11	51.57	51.60	52.11	52.11	52.11	52.11	51.60	52.11	52.11	52.11	52.11	52.11	52.11
gt2	84.26	84.26	83.13	84.26	84.26	84.26	84.26	84.26	84.26	84.26	84.26	84.26	84.26	84.26	84.26
k16x240	7.71	7.71	7.71	7.71	7.71	7.71	7.71	7.71	7.71	7.71	7.71	7.71	7.71	7.71	7.71
lseu	4.65	4.65	4.61	4.65	4.65	4.65	4.65	4.65	4.65	4.65	4.65	4.65	4.65	4.65	4.65
mas74	4.31	4.30	4.31	4.13	4.30	4.31	4.30	4.30	4.31	4.31	4.30	4.31	4.31	4.30	4.31
mas76	2.49	2.49	2.49	2.38	2.48	2.49	2.49	2.49	2.49	2.49	2.48	2.49	2.49	2.49	2.49
mas284	0.51	0.42	0.51	0.39	0.38	0.51	0.42	0.42	0.51	0.51	0.39	0.51	0.51	0.42	0.51
misc05	3.62	3.60	3.60	3.62	3.60	3.60	3.62	3.60	3.62	3.62	3.62	3.62	3.62	3.62	3.62
mod008	1.37	1.37	1.37	1.34	1.33	1.37	1.37	1.37	1.37	1.37	1.34	1.37	1.37	1.37	1.37
mod013	7.37	7.37	7.37	7.37	7.37	7.37	7.37	7.37	7.37	7.37	7.37	7.37	7.37	7.37	7.37
modglob	14.02	14.02	13.79	14.02	14.02	14.02	14.02	14.02	14.02	14.02	14.02	14.02	14.02	14.02	14.02
p0033	5.19	2.59	5.19	5.19	2.59	5.19	5.19	2.59	5.19	5.19	5.19	5.19	5.19	5.19	5.19
p0282	5.12	4.81	4.50	5.12	4.50	5.10	5.12	4.84	5.12	4.74	5.12	5.12	5.10	5.12	5.12
p0291	40.12	39.96	40.11	40.12	32.91	40.11	40.12	39.96	40.12	40.11	40.12	40.12	40.11	40.12	40.12
pipeX	1.43	1.43	1.43	1.43	1.43	1.43	1.43	1.43	1.43	1.43	1.43	1.43	1.43	1.43	1.43
pp08a	54.46	53.50	51.81	54.46	54.45	54.41	54.46	54.45	54.46	54.45	54.46	54.46	54.45	54.46	54.46
probportfolio	25.28	25.28	25.24	25.26	25.20	25.28	25.28	25.28	25.26	25.24	25.26	25.28	25.28	25.28	25.26
sample2	13.14	9.07	8.59	13.14	5.86	10.64	13.14	9.07	13.14	13.14	13.14	13.14	13.14	13.14	13.14
sentoy	14.00	13.62	13.89	13.89	13.92	13.89	13.89	14.00	13.89	13.92	13.92	13.89	14.00	14.00	13.92
stein15_nosym	58.33	58.33	54.67	58.33	54.17	58.33	58.33	58.33	58.33	55.14	58.33	58.33	58.33	58.33	58.33
stein27_nosym	8.78	8.61	7.86	8.78	8.60	8.61	8.78	8.62	8.78	8.60	8.78	8.78	8.62	8.78	8.78
stein45_nosym	7.55	7.51	7.45	7.51	7.49	7.51	7.51	7.55	7.51	7.49	7.51	7.51	7.55	7.55	7.51
vpm1	10.18	10.04	10.04	10.18	10.18	10.18	10.18	10.18	10.18	10.18	10.18	10.18	10.18	10.18	10.18
vpm2	11.71	11.71	11.18	11.60	11.71	11.71	11.71	11.71	11.60	11.71	11.71	11.71	11.71	11.71	11.71
Average	23.34	22.88	21.76	23.22	22.53	23.23	23.25	23.06	23.25	23.19	23.32	23.34	23.32	23.33	23.34
Wins		11	7	19	10	18	22	15	23	19	21	27	24	24	26

collection, which may be far larger than n .

We also tested a bilinear program that finds a maximally violated cut with respect to a vertex of P with all SICs added, but it did not yield additional strong cuts.

6.5 Strength of GICs

Finally, we look at the best percent gap closed across all parameter settings. We compare the percent gap closed by using GICs and SICs together to using SICs on their own. When the GICs close an additional amount of the integrality gap, it is clear that we are getting a tighter relaxation of the integral hull. The converse is not true; that the extra gap closed by GICs is zero does not imply that the cuts have no effect on tightening the relaxation.

Table 5 shows the best result for percent gap closed for each instance, across all parameter settings. The percent of the integrality gap closed by SICs and by GICs is given in columns 2 and 3. Column 4 shows the difference between columns 3 and 2. Columns 5 and 6 show the number of SICs generated as well as how many of the SICs are active at the optimum of the LP with all cuts (both SICs and GICs) added. Columns 7 and 8 show the same for GICs. The top part of the table shows those instances for which GICs or SICs close any gap over the LP relaxation, and the bottom part of the table contains the remaining instances.

GICs close extra gap over SICs for 75% of the instances. The average extra percent gap closed across all instances is around 4.2%, around 5.1% on those instances in which SICs have any effect, and around 5.6% on those instances with nonzero extra gap closed by GICs (representing a 31% improvement over using SICs alone).

Another possible way to assess strength of cuts is how many of the cuts are active at the optimum of the continuous relaxation after the addition of all cuts. On average, over 70% of the active cuts are GICs. Moreover, one can select for the strong cuts with reasonable success. We adopt a common cut selection criterion (see, e.g., [3]) that sorts cuts based on a combination of their *efficacy* (the Euclidean distance by which they cut \bar{x}) and their

Table 5: Best percent gap closed and number of cuts.

Instance	Best % gap closed			# cuts			
	SIC	GIC	Diff	SICs	Active SICs	GICs	Active GICs
bell3a	44.74	59.52	14.77	32	15	95	36
bell3b	44.57	60.34	15.76	35	21	80	28
bell4	23.37	26.54	3.17	46	20	968	21
bell5	14.53	85.37	70.85	25	12	62	38
blend2	16.04	19.78	3.75	6	1	563	9
bm23	5.92	11.06	5.14	6	0	1000	9
egout	51.57	52.11	0.54	38	36	45	41
flugpl	11.74	11.74	0.00	10	6	5	3
gt2	83.13	84.26	1.13	11	11	17	13
k16x240	7.56	7.71	0.15	14	6	25	12
lseu	4.57	4.65	0.08	12	6	87	9
mas74	3.30	4.31	1.01	12	1	1000	48
mas76	2.37	2.49	0.13	11	2	1000	29
mas284	0.38	0.51	0.13	20	1	1000	29
misc05	3.60	3.62	0.02	11	3	644	14
mod008	1.30	1.37	0.07	5	0	865	2
mod013	4.41	7.37	2.96	5	2	57	12
modglob	9.59	14.02	4.43	29	14	321	71
p0033	1.83	5.19	3.35	6	4	31	7
p0040	6.65	6.65	0.00	4	4	7	7
p0282	3.67	5.12	1.45	26	4	975	26
p0291	27.78	40.12	12.34	10	2	105	9
pipex	0.81	1.43	0.62	6	2	218	18
pp08a	51.44	54.46	3.02	53	43	395	75
probportfolio	25.14	25.28	0.15	125	69	1000	143
sample2	5.86	13.14	7.28	12	2	227	8
sentoy	10.38	14.00	3.62	8	0	547	11
stein15_nosym	50.00	58.33	8.33	5	2	41	9
stein27_nosym	7.41	8.78	1.37	27	3	1000	5
stein45_nosym	7.10	7.55	0.45	45	2	1000	6
timtab1	17.54	17.54	0.00	136	38	1	1
vpm1	10.00	10.18	0.18	15	10	98	37
vpm2	10.18	11.71	1.53	31	13	413	49
Average	17.23	22.31	5.08				
glass4	0	0	0	72	36	421	81
misc01	0	0	0	12	4	47	6
misc02	0	0	0	8	8	7	7
misc03	0	0	0	22	8	74	11
misc07	0	0	0	26	12	603	220
p0201	0	0	0	40	19	47	17
rgn	0	0	0	19	11	15	9

orthogonality to cuts that have already been selected. Adding the cuts in this way and setting a cut limit of the same number of GICs as there are SICs, we can close 20.94% of the integrality gap, compared to the 22.31% gap closed using all the GICs. With a limit of five times as many GICs as SICs, we close 22.26% of the integrality gap, that is, over 99% of the improvement from adding all GICs; for comparison, on average across the instances, our cut limit of 1,000 results in about 30 times as many GICs as there are SICs.

Although our procedure is intended to be nonrecursive, we ran a small experiment to test the effect of using GICs with two rounds. These results are shown in Table 6. The second and third columns display the percent gap closed after one round of SICs and one round of GICs (added on top of the SICs, as before). The fourth and fifth columns show the percent gap closed after two rounds of SICs and a second round of GICs (added on top of the first round of SICs and first round of GICs). The next two columns give the number of SICs and GICs generated in the second round. The last column provides the number of rounds of SICs needed to match the gap closed by the second round of GICs. The three instances marked with asterisks in this last column (`bell15`, `flugpl1`, and `mas74`) do not achieve the same gap closed as two rounds of GICs. The first two instances are terminated when as many SICs are generated as GICs. For example, for `bell15`, the 5 rounds of SICs correspond to 97 SICs that together close only 25% of the integrality gap (this is not shown in the table), compared to 96 GICs that close nearly 87%. For `mas74`, SIC generation tails off, with the objective value improving less than 10^{-3} in the last five rounds of SICs (for reference, matching the result of two GIC rounds would require an additional 0.15 improvement in the objective value).

Lastly, although GICs can be generated from any convex cut-generating set, in our experiments, we limit ourselves to variable split disjunctions. As a result, it may be useful to compare GICs to lift-and-project cuts [9], which can generate the facets of each of the sets $\text{conv}(P \setminus \text{int } S_k)$. This comparison is given in Table 7, where the gap closed by lift-and-project cuts is taken from the computational experiments by Bonami [13], for those instances appearing in both test sets. The table indicates that there is a high degree of correlation

between the instances on which lift-and-project cuts are effective, and those on which GICs are effective.

6.6 Summary

The GICs generated from our PHA approach close a significant percentage of the integrality gap with respect to SICs. This validates the premise of PHA, that stronger intersection cuts can be obtained from collecting a manageable number of intersection points and rays. We examine the tradeoff in using PHA_1 with and without the targeted tilting algorithm. Using tilting permits more hyperplanes to be activated, but it can create weaker points that are avoided by the procedure without tilting. Moreover, our experiments identify one aspect of PHA that can be improved to lead to stronger cuts. Only a small percentage of the points we generate are final, which indicates that the GICs generated from our approach are far from the split closure and motivates future work on new methods targeting such final intersection points. We also evaluate several different objective functions that can be used in (PRLP) and determine one (the tight point heuristic) that seems particularly effective, which has implications for any future GIC computational experiments. One of the motivations of the GIC procedure is the ability to generate diverse cuts, which we find is indeed possible. However, we do not wish to add a large number of cuts to our LP relaxation. We show that in fact a small set of diverse GICs achieves nearly the same result as using all the GICs.

References

- [1] Tobias Achterberg and Roland Wunderling. Mixed integer programming: analyzing 12 years of progress. In *Facets of Combinatorial Optimization*, pages 449–481. Springer, Heidelberg, 2013.
- [2] Tobias Achterberg, Thorsten Koch, and Alexander Martin. MIPLIB 2003. *Oper. Res. Lett.*, 34(4):361–372, 2006.

Table 6: Gap closed over two rounds of cuts, compared between SICs and GICs. The first round is suffixed with a 1, and the second round is suffixed with a 2. The last column is the number of rounds of SICs needed to match the percent gap closed by two rounds of GICs. The two instances marked with asterisks do not achieve the same gap closed as GICs.

Instance	% gap closed				# cuts		SIC rds
	SIC1	GIC1	SIC2	GIC2	SIC2	GIC2	
bell3a	44.74	59.52	63.65	63.26	55	206	2
bell3b	44.57	60.34	59.49	67.02	64	327	4
bell4	23.37	26.54	45.21	37.24	81	1968	2
bell5	14.53	85.37	17.93	86.65	43	96	5*
blend2	16.04	19.78	21.11	21.18	16	1564	3
bm23	5.92	11.06	10.40	13.45	12	1958	4
egout	51.57	52.11	90.89	92.48	60	94	3
flugpl	11.74	11.74	13.30	13.00	18	10	1*
gt2	83.13	84.26	95.14	94.54	27	29	2
k16x240	7.56	7.71	13.30	11.50	29	109	2
lseu	4.57	4.65	17.10	31.25	23	342	3
mas74	3.30	4.31	4.06	4.65	27	2000	50*
mas76	2.37	2.49	2.76	2.78	27	2000	3
mas284	0.38	0.51	0.94	0.95	41	2000	3
misc05	3.60	3.62	6.75	6.17	32	1644	2
mod008	1.30	1.37	3.03	2.79	13	1170	2
mod013	4.41	7.37	18.89	9.64	12	264	2
modglob	9.59	14.02	33.07	38.62	67	1321	3
p0033	1.83	5.19	7.30	10.50	18	96	3
p0040	6.65	6.65	8.04	8.55	8	21	3
p0282	3.67	5.12	8.56	6.52	47	1342	2
p0291	27.78	40.12	62.24	47.22	21	822	2
pipex	0.81	1.43	4.21	2.18	12	1218	2
pp08a	51.44	54.46	66.88	70.93	102	1395	3
probportfolio	25.14	25.28	26.07	25.83	218	2000	2
sample2	5.86	13.14	17.58	24.63	26	1227	6
sentoy	10.38	14.00	11.97	15.19	20	770	11
stein15_nosym	50.00	58.33	58.79	76.87	9	97	7
stein27_nosym	7.41	8.78	32.41	34.84	52	1818	3
stein45_nosym	7.10	7.55	9.32	11.21	90	2000	3
timtab1	17.54	17.54	24.86	17.54	255	1	1
vpm1	10.00	10.18	13.91	13.99	34	831	3
vpm2	10.18	11.71	19.32	17.72	63	1228	2
Average	17.23	22.31	26.92	29.72			4.5

Table 7: Comparison of percent gap closed by lift-and-project cuts (taken from the literature) with the gap closed by generalized intersection cuts.

Instance	% gap closed	
	L&PC	GIC
bell3a	64.56	59.52
bell5	86.25	85.37
blend2	21.82	19.78
egout	93.85	52.11
flugpl	11.72	11.74
gt2	92.38	84.26
lseu	16.58	4.65
mas74	5.47	4.31
mas76	3.68	2.49
mod008	9.02	1.37
modglob	57.09	14.02
p0033	8.19	5.19
p0282	93.9	5.12
pp08a	79.29	54.46
timtab1	26.99	17.54
vpm1	31.42	10.18
vpm2	54.29	11.71
Average	44.50	26.11

- [3] Tobias Achterberg, Timo Berthold, Thorsten Koch, and Kati Wolter. Constraint integer programming: A new approach to integrate CP and MIP. In Laurent Perron and Michael A. Trick, editors, *Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems: 5th International Conference, CPAIOR 2008 Paris, France, May 20-23, 2008 Proceedings*, pages 6–20. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008.
- [4] Kent Andersen, Quentin Louveaux, Robert Weismantel, and Laurence A. Wolsey. Inequalities from two rows of a simplex tableau. In *Integer Programming and Combinatorial Optimization*, volume 4513 of *Lecture Notes in Comput. Sci.*, pages 1–15. Springer, Berlin, 2007.
- [5] Egon Balas. Intersection cuts—a new type of cutting planes for integer programming. *Oper. Res.*, 19(1):19–39, 1971.
- [6] Egon Balas. Disjunctive programming. *Ann. Discrete Math.*, 5:3–51, 1979.
- [7] Egon Balas and François Margot. Generalized intersection cuts and a new cut generating paradigm. *Math. Program.*, 137(1-2, Ser. A):19–35, 2013.
- [8] Egon Balas and Anureet Saxena. Optimizing over the split closure. *Math. Program.*, 113(2, Ser. A):219–240, 2008.
- [9] Egon Balas, Sebastián Ceria, and Gérard Cornuéjols. A lift-and-project cutting plane algorithm for mixed 0-1 programs. *Math. Program.*, 58(3, Ser. A):295–324, 1993.
- [10] Amitabh Basu, Pierre Bonami, Gérard Cornuéjols, and François Margot. Experiments with two-row cuts from degenerate tableaux. *INFORMS J. Comput.*, 23(4):578–590, 2011.
- [11] R. E. Bixby, E. A. Boyd, and R. R. Indovina. MIPLIB: A test set of mixed integer programming problems. *SIAM News*, 25:16, 1992.

- [12] R. E. Bixby, S. Ceria, C. M. McZeal, and M. W. P Savelsbergh. An updated mixed integer programming library: MIPLIB 3.0. *Optima*, 58:12–15, June 1998.
- [13] Pierre Bonami. On optimizing over lift-and-project closures. *Math. Program. Comput.*, 4(2):151–179, 2012.
- [14] Christoph Buchheim, Frauke Liers, and Marcus Oswald. Local cuts revisited. *Oper. Res. Lett.*, 36(4):430–433, 2008.
- [15] Michele Conforti, Gérard Cornuéjols, and Giacomo Zambelli. Corner polyhedron and intersection cuts. *Surveys in Operations Research and Management Science*, 16(2):105 – 120, 2011.
- [16] Sanjeeb Dash and Oktay Günlük. On t -branch split cuts for mixed-integer programs. *Math. Program.*, 141(1-2, Ser. A):591–599, 2013.
- [17] Sanjeeb Dash, Oktay Günlük, and Andrea Lodi. MIR closures of polyhedral sets. *Math. Program.*, 121(1, Ser. A):33–60, 2010.
- [18] Sanjeeb Dash, Oktay Günlük, and Juan Pablo Vielma. Computational experiments with cross and crooked cross cuts. *INFORMS J. Comput.*, 26(4):780–797, 2014.
- [19] Sanjeeb Dash, Oktay Günlük, and Marco Molinaro. On the relative strength of different generalizations of split cuts. *Discrete Optim.*, 16:36–50, 2015.
- [20] Sanjeeb Dash, Oktay Günlük, and Diego A. Morán R. On the polyhedrality of cross and quadrilateral closures. *Math. Program.*, pages 1–26, 2016.
- [21] Santanu S. Dey and Sebastian Pokutta. Design and verify: a new scheme for generating cutting-planes. *Math. Program.*, 145(1-2, Ser. A):199–222, 2014.
- [22] Santanu S. Dey, Andrea Lodi, Andrea Tramontani, and Laurence A. Wolsey. On the practical strength of two-row tableau cuts. *INFORMS J. Comput.*, 26(2):222–237, 2014.

- [23] Daniel G. Espinoza. Computing with multi-row Gomory cuts. *Oper. Res. Lett.*, 38(2): 115–120, 2010.
- [24] D. R. Fulkerson, G. L. Nemhauser, and L. E. Trotter. Two computationally difficult set covering problems that arise in computing the 1-width of incidence matrices of Steiner triple systems. *Math. Program. Stud.*, (2):72–81, 1974.
- [25] Ralph E. Gomory. Some polyhedra related to combinatorial problems. *Linear Algebra Appl.*, 2(4):451–558, 1969.
- [26] Ralph E. Gomory and Ellis L. Johnson. Some continuous functions related to corner polyhedra. *Math. Program.*, 3(1):23–85, 1972.
- [27] Ralph E. Gomory and Ellis L. Johnson. Some continuous functions related to corner polyhedra. II. *Math. Program.*, 3(1):359–389, 1972.
- [28] Thorsten Koch, Tobias Achterberg, Erling Andersen, Oliver Bastert, Timo Berthold, Robert E. Bixby, Emilie Danna, Gerald Gamrath, Ambros M. Gleixner, Stefan Heinz, Andrea Lodi, Hans Mittelmann, Ted Ralphs, Domenico Salvagnin, Daniel E. Steffy, and Kati Wolter. MIPLIB 2010: mixed integer programming library version 5. *Math. Program. Comput.*, 3(2):103–163, 2011.
- [29] Robin Lougee-Heimer. The Common Optimization INterface for Operations Research: Promoting open-source software in the operations research community. *IBM Journal of Research and Development*, 47, 2003.
- [30] Quentin Louveaux, Laurent Poirrier, and Domenico Salvagnin. The strength of multi-row models. *Math. Program. Comput.*, 7(2):113–148, 2015.
- [31] Olvi L. Mangasarian. *Nonlinear programming*, volume 10 of *Classics in Applied Mathematics*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1994. Corrected reprint of the 1969 original.

- [32] Michael Perregaard and Egon Balas. Generating cuts from multiple-term disjunctions. In *Integer Programming and Combinatorial Optimization*, volume 2081 of *Lecture Notes in Comput. Sci.*, pages 348–360. Springer, Berlin, 2001.
- [33] Arrigo Zanette, Matteo Fischetti, and Egon Balas. Lexicography and degeneracy: can a pure cutting plane algorithm work? *Math. Program.*, 130(1, Ser. A):153–176, 2011.
- [34] Günter M. Ziegler. *Lectures on Polytopes*, volume 152 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1995.

A Additional theory for PHA

A.1 Existence of strictly dominating GICs

In this section, we provide some theoretical motivation for Algorithm 3 by giving necessary and sufficient conditions for the existence of a GIC that strictly dominates the SIC after activation of a single hyperplane.

Definition 16 ([7]). *Consider two inequalities that are valid for P_I but not necessarily P . Inequality 2 dominates 1 on P if for every $x \in P$, the fact that x satisfies Inequality 2 implies that x satisfies Inequality 1. Inequality 2 strictly dominates 1 if, in addition, there exists $x \in P$ such that x violates Inequality 2 but satisfies Inequality 1.*

The theorem proved in this section strengthens Theorem 5 in Balas and Margot [7] for the case when S is a split disjunction. Theorem 5 of the aforementioned paper gives sufficient conditions for a GIC to strictly dominate the SIC, given that dominance holds. We show that this condition is also necessary for strict dominance when S is a split disjunction. For ease of exposition, Theorem 17 assumes that all rays of C intersect $\text{bd } S$ because Proposition 10 shows that intersecting rays cannot lead to deeper points.

Suppose $S = \{x : 0 \leq x_k \leq 1\}$ is a split disjunction on a variable x_k . Let $S^0 := \{x : x_k = 0\}$, and $S^1 := \{x : x_k = 1\}$. We partition the intersection point set \mathcal{P} into \mathcal{P}^0 and \mathcal{P}^1 , where

$\mathcal{P}^0 := \mathcal{P} \cap S^0$, and $\mathcal{P}^1 := \mathcal{P} \cap S^1$. Recall that $\bar{\mathcal{P}}$ and $\bar{\mathcal{R}}$ are the points and rays obtained from intersecting \bar{C} with $\text{bd } S$. We also partition $\bar{\mathcal{P}}$ into $\bar{\mathcal{P}}^0 := \bar{\mathcal{P}} \cap S^0$ and $\bar{\mathcal{P}}^1 := \bar{\mathcal{P}} \cap S^1$. Intuitively, the theorem shows that a strictly dominating cut with respect to the SIC must reduce the dimension of $\text{conv}(\bar{\mathcal{P}}^0)$ or $\text{conv}(\bar{\mathcal{P}}^1)$.

Theorem 17. *Suppose that $\bar{\mathcal{R}} = \emptyset$, $\bar{\mathcal{P}}^0 \neq \emptyset$ and $\bar{\mathcal{P}}^1 \neq \emptyset$, and $(\mathcal{P}, \mathcal{R})$ is a proper point-ray collection obtained from activating a single hyperplane H valid for P . There exists a basic feasible solution to (1) corresponding to a cut strictly dominating $\alpha_0^\top x \geq \beta_0$ if and only if $\text{relint}(H^+) \cap \bar{\mathcal{P}}^t = \emptyset$ and $H^- \cap \bar{\mathcal{P}}^t \neq \emptyset$, for at least one side t of the split disjunction, $t \in \{0, 1\}$.*

Proof. For the “if” direction of the proof, suppose without loss of generality that $\text{relint}(H^+) \cap \bar{\mathcal{P}}^0 = \emptyset$ and $H^- \cap \bar{\mathcal{P}}^0 \neq \emptyset$. Any point in \mathcal{P}^0 lying on the SIC is in $\text{conv}(\bar{\mathcal{P}}^0)$. Because $\text{relint}(H^+) \cap \bar{\mathcal{P}}^0 = \emptyset$, it holds that $(\mathcal{P}^0 \setminus \bar{\mathcal{P}}^0) \cap \text{conv}(\bar{\mathcal{P}}^0) = \emptyset$. This implies that any point p in $\mathcal{P}^0 \setminus \bar{\mathcal{P}}^0$ satisfies $\alpha_0^\top p > \beta_0$. Recall that $|\bar{\mathcal{P}}^0| + |\bar{\mathcal{P}}^1| = n$. Since some point of $\bar{\mathcal{P}}^0$ lies in H^- , $|\bar{\mathcal{P}}^0 \cap H^+| \leq |\bar{\mathcal{P}}^0| - 1$. It follows that at most $n - 1$ intersection points from $\bar{\mathcal{P}}$ remain in \mathcal{P} . This added degree of freedom and the aforementioned depth of points in $\mathcal{P}^0 \setminus \bar{\mathcal{P}}^0$ allows the SIC to be tilted to obtain a GIC that strictly dominates the SIC. The “only if” direction follows from Theorem 5 in Balas and Margot [7]. \square

The above result shows that any single hyperplane is unlikely to directly lead to a strictly dominating cut. Instead of looking for one such hyperplane, in our implementation we focus on activating a set of hyperplanes that together cut away large parts of $\text{conv}(\bar{\mathcal{P}}^0)$ and $\text{conv}(\bar{\mathcal{P}}^1)$. We do this by targeting each of the intersection points in $\bar{\mathcal{P}}$ one at a time in step 6 of Algorithm 3.

Although strict dominance is difficult to attain, our next result shows that activating hyperplanes is monotonic in the sense that the lower bound implied by the point-ray collection can only be improved by activating hyperplanes. This complements Theorem 3 of Balas and Margot [7], in which it is shown that activating hyperplanes increases the depth of points

with respect to the SIC. This leaves open the question of whether the lower bound on the objective value (implied by the points) improves after activating hyperplanes, which Proposition 18 resolves. Using the notation from Section 5.2, we show that when the ray creating the least cost intersection point \underline{p}^k is cut by a hyperplane, the objective value implied by the new intersection points is greater than or equal to \underline{z} .

Proposition 18. *Let r be the edge of C that leads to \underline{p}^k , i.e., $\underline{p}^k = r \cap \text{bd } S_k$, H be a hyperplane intersected by r before $\text{bd } S_k$, and \mathcal{P}' denote the set of intersection points originating at $r \cap H$, obtained by activating H . Then $\min\{c^\top p : p \in \mathcal{P}'\} \geq \underline{z}$.*

Proof. Suppose S_k^0 is the facet of S_k containing \underline{p}^k , and let S_k^1 be the opposite facet. We have that $\underline{z} = \min\{c^\top x : x \in C \cap S_k^0\} \leq \min\{c^\top x : x \in C \cap S_k^1\}$. Since each of the points $p \in \mathcal{P}'$ is either (possibly strictly) in $C \cap S_k^0$ or $C \cap S_k^1$, the result follows. \square

A.2 Characterizing bounded objective functions for (PRLP)

We turn to an analysis of (PRLP). It is possible for the optimal solution to (PRLP) to be unbounded, a behavior we have in fact observed in our numerical implementation. To better understand this, in this section we present some structural properties of $\text{CutRegion}(\bar{\beta}, \mathcal{P}, \mathcal{R})$, the feasible region to (PRLP) for a fixed right-hand size $\bar{\beta}$, that characterize the objective function choices leading to unboundedness.

We begin by studying when the system $\text{CutRegion}(\bar{\beta}, \mathcal{P}, \mathcal{R})$ has valid cuts for a given proper point-ray collection. Recall that \mathcal{K}' denotes the connected component of the skeleton of P that includes $\bar{x} \cap \text{int } S$, and any inequality feasible to (1) that cuts a point $v \in \mathcal{K}'$ is valid. We will consider the system

$$\mathcal{G}^\# := \{\alpha : \alpha \in \text{CutRegion}(\bar{\beta}, \mathcal{P}, \mathcal{R}); v^\top \alpha < \bar{\beta}\}.$$

Theorem 19. *Let $(\mathcal{P}, \mathcal{R})$ be a proper point-ray collection and let $v \in \mathcal{K}'$. The system $\text{CutRegion}(\bar{\beta}, \mathcal{P}, \mathcal{R})$ has valid cuts as feasible solutions in the following cases: (1) for $\bar{\beta} = 1$*

if and only if $0 \notin \mathcal{G}$ and $v \notin \text{conv}(\mathcal{P}) + \text{cone}(\mathcal{P} \cup \mathcal{R}) = \mathcal{G} + \text{cone}(\mathcal{P})$, (2) for $\bar{\beta} = -1$ if and only if $v \notin \text{conv}(\mathcal{G} \cup \{0\})$, and (3) for $\bar{\beta} = 0$ if and only if $v \notin \text{cone}(\mathcal{P} \cup \mathcal{R})$.

Proof. Let Q be the $|\mathcal{P}| \times n$ matrix containing the intersection points in \mathcal{P} as its rows, and R be the $|\mathcal{R}| \times n$ matrix with rows comprised of the rays in \mathcal{R} . Let e denote the n -vector of all ones. Using the nonhomogeneous Farkas' lemma [31], $\mathcal{G}^\#$ has a feasible solution if and only if the following two systems are infeasible:

$$\left\{ \begin{array}{l} \lambda, \mu \geq 0 : \\ Q^\top \lambda + R^\top \mu = v \\ \bar{\beta} e^\top \lambda \geq \bar{\beta} \end{array} \right\} \quad \left\{ \begin{array}{l} \lambda, \mu \geq 0 : \\ Q^\top \lambda + R^\top \mu = 0 \\ \bar{\beta} e^\top \lambda > 0 \end{array} \right\}$$

When $\bar{\beta} = 1$, the first system is infeasible if and only if $v \notin \mathcal{G} + \text{cone}(\mathcal{P})$, and the second system is infeasible if and only if $0 \notin \mathcal{G}$, since the existence of a solution (λ, μ) implies $(\lambda/e^\top \lambda, \mu)$ is also feasible. When $\bar{\beta} = -1$, the first system is infeasible if and only if $v \notin \text{conv}(\mathcal{G} \cup \{0\})$, and the second system is always infeasible, since $\lambda \geq 0$. When $\bar{\beta} = 0$, the first system is infeasible if and only if $v \notin \text{cone}(\mathcal{P} \cup \mathcal{R})$, and the second system is again always infeasible. \square

The feasible region to (PRLP) is $\text{CutRegion}(\bar{\beta}, \mathcal{P}, \mathcal{R})$, not $\mathcal{G}^\#$. However, if we assume that v is used as the objective to (PRLP), then Theorem 19 can be used to show when (PRLP) has a finite solution. Observe that (PRLP) implicitly ranks valid inequalities and picks the most violated cut with respect to v . If there exists a homogeneous inequality valid for \mathcal{G} that cuts off v , this ranking breaks down, since all homogeneous inequalities can be scaled to have arbitrarily large violation and hence are unbounded directions in (PRLP).

From the $\bar{\beta} = 0$ case in Theorem 19, it follows that the linear program (PRLP) is bounded if and only if v belongs to $\text{cone}(\mathcal{P} \cup \mathcal{R})$. Corollary 20 characterizes the two open objective function sets within $\text{cone}(\mathcal{P} \cup \mathcal{R})$ that admit valid cuts of only one type, either with right-hand side 1 or -1 .

Corollary 20. *The system (PRLP) has valid inequalities that cut off a point v only for*

1. $\bar{\beta} = 1$ if and only if $0 \notin \mathcal{G}$ and $v \in \text{conv}(\mathcal{G} \cup \{0\}) \setminus \mathcal{G}$.

2. $\bar{\beta} = -1$ if and only if $v \in (\mathcal{G} + \text{cone}(\mathcal{P})) \setminus \mathcal{G}$.

Proof. Notice that

$$\text{conv}(\mathcal{G} \cup \{0\}) \setminus \mathcal{G} \subseteq \text{conv}(\mathcal{G} \cup \{0\}) \subseteq \text{cone}(\mathcal{P} \cup \mathcal{R})$$

and $\{\text{conv}(\mathcal{G} \cup \{0\}) \setminus \mathcal{G}\} \cap \{\mathcal{G} + \text{cone}(\mathcal{P})\} = \emptyset$. Therefore, the result in part 1 follows from Theorem 19. The proof of the second part is similar. \square

B Tilting for degenerate hyperplanes

We previously showed how to tilt a hyperplane H defining P that is not degenerate, i.e., \bar{x} does not lie on H . We defined H using n affinely independent points obtained by intersecting the n affinely independent rays of \bar{C} with H . These points all lie on one-dimensional faces of \bar{C} . In the case that H is degenerate, we will instead use two-dimensional faces of \bar{C} to define the hyperplane, which we will then modify to define a targeted tilting.

When a degenerate hyperplane H is activated on \bar{C} , each of the extreme rays of the new cone $\bar{C} \cap H^+$ that lie on H can be defined by H and $n - 2$ hyperplanes of \bar{C} that are not redundant for $\bar{C} \cap H^+$. Thus, each ray of the new cone lying on H is on a two-dimensional face of \bar{C} . Let \mathcal{R}^{H^*} denote an affinely independent set of $n - 1$ of these rays lying on H . Let $\text{rays}(\bar{C})^c$ be the rays of \bar{C} that are cut by H^+ . For each $r \in \mathcal{R}^{H^*}$, since it lies on a two-dimensional face of \bar{C} , we can define $(\bar{r}^1, \bar{r}^2) \in \text{rays}(\bar{C})^c \times (\text{rays}(\bar{C}) \setminus \text{rays}(\bar{C})^c)$ such that r can be expressed as a convex combination of \bar{r}^1 (which is cut by H) and \bar{r}^2 (which is not cut by H); in particular, let λ_r^H be the multiplier such that $r = \lambda_r^H \bar{r}^1 + (1 - \lambda_r^H) \bar{r}^2$. Figure 5 depicts this construction.

If we know λ_r^H and the rays \bar{r}^1 and \bar{r}^2 for every $r \in \mathcal{R}^{H^*}$, then we can use this to give an alternate definition of H . To define a targeted tilting of H , we can modify the values λ_r^H for




Figure 5: Illustration of targeted tilting construction for a degenerate hyperplane.

each $r \in \mathcal{R}^{H^*}$ using some δ_r . In order to coordinate with the definition of a targeted tilting that we gave in Section 4, the intersection points and rays created by activating the tilted hyperplane should either be identical to those obtained from activating H or coincide with some initial intersection point or ray. It is not difficult to see that this means the allowed values for δ_r are 0 and $-\lambda_r^H$. With such a tilting, Theorem 11 will still hold, i.e., activations can be performed by using H and \mathcal{R}_A without a need for explicitly computing the tilted hyperplane.

C Tilting example

In this section, we demonstrate an example in which a valid tilting combined with the implicit computation used in Theorem 11 leads to invalid cuts. The example additionally provides intuition for the targeted tilting rule that if a hyperplane cuts a ray of \bar{C} that has already been cut, then we should not tilt the hyperplane along that ray. This is sufficient, as we have shown, to allow us to apply the implicit computation scheme.

The left panel of Figure 6 shows the feasible region of $P := \{x \in \mathbb{R}^3 : -2x_2 + x_3 \leq 0; -2x_1 + x_3 \leq 0; 12x_1 + 10x_2 - 5x_3 \leq 9; 10x_1 + 12x_2 - 5x_3 \leq 9; x_1 + x_2 + x_3 \leq 1\}$, and the right panel of the same figure shows the cone \bar{C} . The cut-generating set S is the unit box, $\{x \in \mathbb{R}^3 : 0 \leq x_1 \leq 1; 0 \leq x_2 \leq 1\}$. The hyperplane activations are of H_4 and then H_5 . We tilt H_4 along r^2 so that the intersection of this ray with \widetilde{H}_4 coincides $r^2 \cap \text{bd } S$, and H_5 will similarly be tilted along r^1 so that the intersection of r^1 with \widetilde{H}_5 is at the point $r^1 \cap \text{bd } S$.




Figure 6: Feasible region of P and \bar{C} .

The tilted hyperplanes are shown in the top panel of Figure 7.

The tilting defined above is clearly valid. It also satisfies all but one of the conditions of being a targeted tilting; it does not meet the requirement that H_5 and \widetilde{H}_5 must intersect r^1 at the same point, as a result of the activation of H_4 on r^1 prior to the activation of H_5 .

However, as shown in the bottom panel of Figure 7, when the implicit computation algorithm is applied to these tilted hyperplanes, a point of $\text{conv}(P \setminus \text{int } S)$ is cut. This is because \widetilde{H}_5 intersects r^1 outside of the interior of S . Hence, the intersection point $\widetilde{H}_5 \cap r^1$ is not added to the point collection, as a result of step 5 of Algorithm 1. This intersection point is the same as $H_4 \cap r^1$, which has already been removed from the point collection during the activation of \widetilde{H}_4 .

There may be many approaches in which a tilted hyperplane activation can be computed implicitly using only information from the non-tilted hyperplane. Our method prevents the situation in this example from occurring by requiring that r^1 intersects \widetilde{H}_5 at the same point as it intersects H_5 . An example of an alternative would be to add the intersection point $\widetilde{H}_5 \cap r^1 = H_4 \cap r^1$ back into the point collection when activating \widetilde{H}_5 , but then the intersection points obtained from activating H_4 on r^1 would be redundant.




Figure 7: Hyperplane activations leading to cutting a point in $\text{conv}(P \setminus \text{int } S)$.