

Structural Analysis of Branch-and-Cut and the Learnability of Gomory Mixed Integer Cuts

Maria-Florina Balcan* Siddharth Prasad† Tuomas Sandholm‡ Ellen Vitercik§

Abstract

The incorporation of cutting planes within the branch-and-bound algorithm, known as branch-and-cut, forms the backbone of modern integer programming solvers. These solvers are the foremost method for solving discrete optimization problems and thus have a vast array of applications in machine learning, operations research, and many other fields. Choosing cutting planes effectively is a major research topic in the theory and practice of integer programming. We conduct a novel structural analysis of branch-and-cut that pins down how every step of the algorithm is affected by changes in the parameters defining the cutting planes added to the input integer program. Our main application of this analysis is to derive sample complexity guarantees for using machine learning to determine which cutting planes to apply during branch-and-cut. These guarantees apply to infinite families of cutting planes, such as the family of Gomory mixed integer cuts, which are responsible for the main breakthrough speedups of integer programming solvers. We exploit geometric and combinatorial structure of branch-and-cut in our analysis, which provides a key missing piece for the recent generalization theory of branch-and-cut.

1 Introduction

Integer programming (IP) solvers are the most widely-used tools for solving discrete optimization problems. They have numerous applications in machine learning, operations research, and many other fields, including MAP inference [22], combinatorial auctions [33], natural language processing [23], neural network verification [11], interpretable classification [37], training of optimal decision trees [9], and optimal clustering [30], among many others.

Under the hood, IP solvers use the tree-search algorithm branch-and-bound [26] augmented with *cutting planes*, known as *branch-and-cut* (B&C). A cutting plane is a linear constraint that is added to the LP relaxation at any node of the search tree. With a carefully selected cutting plane, the LP guidance can more efficiently lead B&C to the globally optimal integral solution. Cutting planes, specifically the family of *Gomory mixed integer cuts* which we study in this paper, are responsible for breakthrough speedups of modern IP solvers [14].

Successfully employing cutting planes can be challenging because there are infinitely many cuts to choose from and there are still many open questions about which cuts to employ when. A growing body of research has studied the use of machine learning for cut selection [7, 8, 19, 34]. In this paper, we analyze a machine learning setting where there is an unknown distribution over IPs—for

*Computer Science Department, Machine Learning Department, Carnegie Mellon University. ninamf@cs.cmu.edu

†Computer Science Department, Carnegie Mellon University. sprasad2@cs.cmu.edu

‡Computer Science Department, Carnegie Mellon University, Optimized Markets, Inc., Strategic Machine, Inc., Strategy Robot, Inc. sandholm@cs.cmu.edu

§Department of Electrical Engineering and Computer Sciences, UC Berkeley. vitercik@berkeley.edu

example, a distribution over a shipping company’s routing problems. The learner receives a *training set* of IPs sampled from this distribution which it uses to learn cut parameters with strong average performance over the training set (leading, for example, to small search trees). We provide *sample complexity bounds* for this procedure, which bound the number of training instances sufficient to ensure that if a set of cut parameters leads to strong average performance over the training set, it will also lead to strong expected performance on future IPs from the same distribution. These guarantees apply *no matter* what procedure is used to optimize the cut parameters over the training set—optimal or suboptimal, automated or manual. We prove these guarantees by analyzing how the B&C tree varies as a function of the cut parameters on any IP. By bounding the “intrinsic complexity” of this function, we are able to provide our sample complexity bounds.

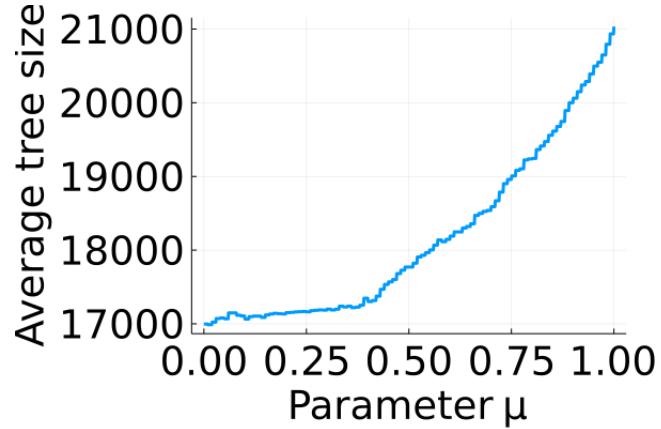


Figure 1: Facility location with 40 locations and 40 clients. Samples generated by perturbing a base facility location IP.

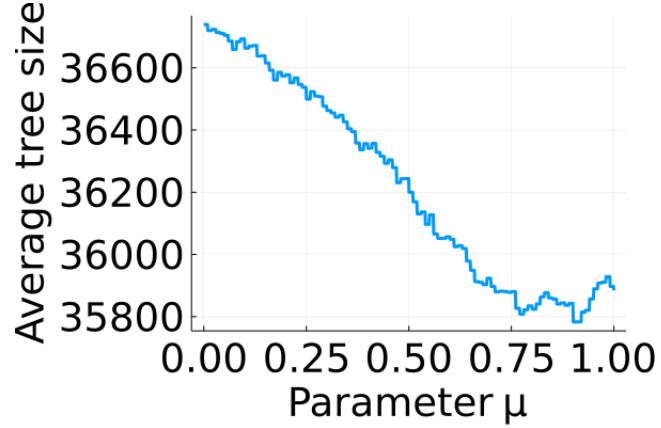


Figure 2: Facility location with 80 locations, 80 clients, and random Euclidean distance costs.

Figures 1 and 2 illustrate the need for distribution-dependent policies for choosing cutting planes. We plot the average number of nodes expanded by B&C as a function of a parameter μ that controls its cut-selection policy, as we detail in Appendix A. In each figure, we draw a training set of facility location integer programs from two different distributions. In Figure 1, we define the distribution by starting with a uniformly random facility location instance and perturbing its costs. In Figure 2, the costs are more structured: the facilities are located along a line and the clients

have uniformly random locations. In Figure 1, a smaller value of μ leads to small search trees, but in Figure 2, a larger value of μ is preferable. These figures illustrate that tuning cut parameters according to the instance distribution at hand can have a large impact on the performance of B&C, and that for one instance distribution, the best parameters for cut evaluation can be very different—in fact opposite—than the optimal parameters for another instance distribution.

The key challenge we face in developing a theory for cutting planes is that a cut added at the root remains in the LP relaxations stored in each node all the way to the leaves, thereby impacting the LP guidance that B&C uses to search throughout the whole tree. Tiny changes to any cut can thus completely change the entire course of B&C. At its core, our analysis therefore involves understanding an intricate interplay between the continuous and discrete components of our problem. The first, continuous component requires us to characterize how the solution the LP relaxation changes as a function of its constraints. This optimal solution will move continuously through space until it jumps from one vertex of the LP tableau to another. We then use this characterization to analyze how the B&C tree—a discrete, combinatorial object—varies as a function of its LP guidance.

1.1 Our contributions

Our first main contribution (Section 3) addresses a fundamental question: how does an LP’s solution change when new constraints are added? As the constraints vary, the solution will jump from vertex to vertex of the LP polytope. We prove that one can partition the set of all possible constraint vectors into a finite number of regions such that within any one region, the LP’s solution has a clean closed form. Moreover, we prove that the boundaries defining this partition have a specific form, defined by degree-2 polynomials.

We build on this result to prove our second main contribution (Section 4), which analyzes how the entire B&C search tree changes as a function of the cuts added at the root. To prove this result, we analyze how every aspect of B&C—the variables branched on, the nodes selected to expand, and the nodes fathomed—changes as a function of the LP relaxations that are computed throughout the search tree. We prove that the set of all possible cuts can be partitioned into a finite number of regions such that within any one region, B&C builds the exact same search tree.

This result allows us to prove sample complexity bounds for learning high-performing cutting planes from the class of *Gomory mixed integer (GMI) cuts*, our third main contribution (Section 5). GMI cuts are one of the most important families of cutting planes in the field of integer programming. Introduced by Gomory [16], they dominate most other families of cutting planes [15], and are perhaps most directly responsible for the realization that a branch-and-cut framework is necessary for the speeds now achievable by modern IP solvers [3]. A historical account of these cuts is provided by Cornuéjols [14]. The structural results from Section 4 allow us to understand the “intrinsic complexity” of B&C’s performance as a function of the GMI cuts it uses. We quantify this notion of intrinsic complexity using *pseudo-dimension* [32], which then implies a sample complexity bound.

1.2 Related research

Learning to cut. This paper helps develop a theory of generalization for cutting plane selection. This line of inquiry began with a paper by Balcan et al. [8], who studied Chvátal-Gomory cuts for (pure) integer programs (IPs). Unlike that work, which exploited the fact that there are only finitely many distinct Chvátal-Gomory cuts for a given IP, our analysis of GMI cuts is far more involved.

The main distinction between our analysis in this paper and the techniques used in previous papers on generalization guarantees for integer programming [5, 7, 8] can be summarized as follows. Let μ be a (potentially multidimensional) parameter controlling some aspect of the IP solver (e.g. a mixture parameter between branching rules or a cutting-plane parameter). In previous works, as μ varied, there were only a finite number of states each node of branch-and-cut could be in. For example, in the case of branching/variable selection, μ controls the additional branching constraint added to the IP at any given node of the search tree. There are only finitely many possible branching constraints, so there are only finitely many possible “child” IPs induced by μ . Similarly, if μ represents the parameterization for Chvátal-Gomory cuts [12, 17], since Balcan et al. [8] showed that there are only finitely many distinct Chvátal-Gomory cuts for a given IP, as μ varies, there are only finitely many possible child IPs induced by μ at any stage of the search tree. However, in many settings, this property does not hold. For example if $\mu = (\alpha, \beta)$ controls the normal vector and offset of an additional feasible constraint $\alpha^T \mathbf{x} \leq \beta$, there are infinitely many possible IPs corresponding to the choice of (α, β) . Similarly, if μ controls the parameterization of a GMI cut, there are infinitely many IPs corresponding to the choice of μ (unlike Chvátal-Gomory cuts). In this paper, we develop a new structural understanding of B&C that is significantly more involved than the structural results in prior work.

This paper ties in to a broader line of research that provides sample complexity bounds for algorithm configuration [e.g., 6, 18]. A chapter by Balcan [4] provides a comprehensive survey.

There have also been several papers that study how to use machine learning for cut selection from an applied perspective [19, 34]. In contrast, the goal of this paper is to provide theoretical guarantees.

Sensitivity analysis of integer and linear programs. A related line of research studied the *sensitivity* of LPs, and to a lesser extent IPs, to changes in their parameters. Mangasarian and Shaiu [28] and Li [27], for example, show that the optimal solution to an LP is a Lipschitz function of the right-hand-side of its constraints but not of its objective. Cook et al. [13] study how the set of optimal solutions to an IP changes as the objective function varies and the right-hand-side of the constraints varies. This paper fits in to this line of research as we study how the solution to an LP varies as new rows are added. This function is not Lipschitz, but we show that it is well-structured.

2 Notation and branch-and-cut background

Integer and linear programs. An *integer program* (IP) is defined by an objective vector $\mathbf{c} \in \mathbb{R}^n$, a constraint matrix $A \in \mathbb{Z}^{m \times n}$, and a constraint vector $\mathbf{b} \in \mathbb{Z}^m$, with the form

$$\max\{\mathbf{c}^T \mathbf{x} : A\mathbf{x} \leq \mathbf{b}, \mathbf{x} \geq \mathbf{0}, \mathbf{x} \in \mathbb{Z}^n\}. \quad (1)$$

The *linear programming (LP) relaxation* is formed by removing the integrality constraints:

$$\max\{\mathbf{c}^T \mathbf{x} : A\mathbf{x} \leq \mathbf{b}, \mathbf{x} \geq \mathbf{0}\}. \quad (2)$$

We denote the optimal solution to (1) by \mathbf{x}_{IP}^* . We denote the optimal solution to (2) by \mathbf{x}_{LP}^* and its objective value by $z_{\text{LP}}^* = \mathbf{c}^T \mathbf{x}_{\text{LP}}^*$. If σ is a set of constraints, we let $\mathbf{x}_{\text{LP}}^*(\sigma)$ denote the LP optimum of (2) subject to these additional constraints (similarly define $z_{\text{LP}}^*(\sigma)$ and $\mathbf{x}_{\text{IP}}^*(\sigma)$).

Polyhedra and polytopes. A set $\mathcal{P} \subseteq \mathbb{R}^n$ is a *polyhedron* if there exists an integer m , $A \in \mathbb{R}^{m \times n}$, and $\mathbf{b} \in \mathbb{R}^m$ such that $\mathcal{P} = \{\mathbf{x} \in \mathbb{R}^n : A\mathbf{x} \leq \mathbf{b}\}$. \mathcal{P} is a *rational polyhedron* if there exists $A \in \mathbb{Z}^{m \times n}$ and $\mathbf{b} \in \mathbb{Z}^m$ such that $\mathcal{P} = \{\mathbf{x} \in \mathbb{R}^n : A\mathbf{x} \leq \mathbf{b}\}$. A bounded polyhedron is called a *polytope*. The feasible regions of all IPs considered in this paper are assumed to be rational polytopes¹. Let $\mathcal{P} = \{\mathbf{x} \in \mathbb{R}^n : \mathbf{a}^i \mathbf{x} \leq b_i, i \in M\}$ be a nonempty polyhedron. For any $I \subseteq M$, the set $F_I := \{\mathbf{x} \in \mathbb{R}^n : \mathbf{a}^i \mathbf{x} = b_i, i \in I, \mathbf{a}^i \mathbf{x} \leq b_i, i \in M \setminus I\}$ is a face of \mathcal{P} . Conversely, if F is a nonempty face of \mathcal{P} , then $F = F_I$ for some $I \subseteq M$. Given a set of constraints σ , let $\mathcal{P}(\sigma)$ denote the polyhedron that is the intersection of \mathcal{P} with all inequalities in σ .

Cutting planes. A *cutting plane* is a linear constraint $\boldsymbol{\alpha}^T \mathbf{x} \leq \beta$. Let \mathcal{P} be the feasible region of the LP relaxation in Equation (2) and $\mathcal{P}_I = \mathcal{P} \cap \mathbb{Z}^n$ be the feasible set of the IP in Equation (1). A cutting plane is *valid* if it is satisfied by every integer-feasible point: $\boldsymbol{\alpha}^T \mathbf{x} \leq \beta$ for all $\mathbf{x} \in \mathcal{P}_I$. A valid cut *separates* a point $\mathbf{x} \in \mathcal{P} \setminus \mathcal{P}_I$ if $\boldsymbol{\alpha}^T \mathbf{x} > \beta$. We interchangeably refer to a cut by its parameters $(\boldsymbol{\alpha}, \beta) \in \mathbb{R}^{n+1}$ and the halfspace $\boldsymbol{\alpha}^T \mathbf{x} \leq \beta$ in \mathbb{R}^n it defines.

An important family of cuts that we study in this paper is the set of *Gomory mixed integer (GMI) cuts*.

Definition 2.1 (Gomory mixed integer cut). Suppose the feasible region of the IP is in equality form $A\mathbf{x} = \mathbf{b}$, $\mathbf{x} \geq \mathbf{0}$ (which can be achieved by adding slack variables). For $\mathbf{u} \in \mathbb{R}^m$, let f_i denote the fractional part of $(\mathbf{u}^T A)_i$ and let f_0 denote the fractional part of $\mathbf{u}^T \mathbf{b}$. That is, $(\mathbf{u}^T A)_i = (\lfloor \mathbf{u}^T A \rfloor)_i + f_i$ and $\mathbf{u}^T \mathbf{b} = \lfloor \mathbf{u}^T \mathbf{b} \rfloor + f_0$. The *Gomory mixed integer (GMI) cut* parameterized by \mathbf{u} is given by

$$\sum_{i: f_i \leq f_0} f_i x_i + \frac{f_0}{1 - f_0} \sum_{i: f_i > f_0} (1 - f_i) x_i \geq f_0.$$

Branch-and-cut. We provide a high-level overview of branch-and-cut (B&C) and refer the reader to the textbook by Nemhauser and Wolsey [31] for more details. Given an IP, B&C searches through the IP’s feasible region by building a binary search tree. B&C solves the LP relaxation of the input IP and then adds any number of cutting planes. It stores this information at the root of its binary search tree. Let $\mathbf{x}_{\text{LP}}^* = (\mathbf{x}_{\text{LP}}^*[1], \dots, \mathbf{x}_{\text{LP}}^*[n])$ be the solution to the LP relaxation with the addition of the cutting planes. B&C next uses a *variable selection policy* to choose a variable x_i to branch on. This means that it splits the IP’s feasible region in two: one set where $x_i \leq \lfloor \mathbf{x}_{\text{LP}}^*[i] \rfloor$ and the other where $x_i \geq \lceil \mathbf{x}_{\text{LP}}^*[i] \rceil$. The left child of the root now corresponds to the IP with a feasible region defined by the first subset and the right child likewise corresponds to the second subset. B&C then chooses a leaf using a *node selection policy* and recurses, adding any number of cutting planes, branching on a variable, and so on. B&C *fathoms* a node—which means that it will never branch on that node—if 1) the LP relaxation at the node is infeasible, 2) the optimal solution to the LP relaxation is integral, or 3) the optimal solution to the LP relaxation is no better than the best integral solution found thus far. Eventually, B&C will fathom every leaf, and it can be verified that it has found the globally optimal integral solution. We assume there is a bound κ on the size of the tree we allow B&C to build before we terminate, as is common in prior research [5, 7, 8, 20, 24, 25].

Every step of B&C—including node and variable selection and the choice of whether or not to fathom—depends crucially on guidance from LP relaxations. To give an example, this is true of the *product scoring rule* [1], a popular variable selection policy that our results apply to.

¹This assumption is not a restrictive one. The Minkowski-Weyl theorem states that any polyhedron can be decomposed as the sum of a polytope and its recession cone. All results in this paper can be derived for rational polyhedra by considering the corresponding polytope in the Minkowski-Weyl decomposition.

Definition 2.2. Let \mathbf{x}_{LP}^* be the solution to the LP relaxation at a node and $z_{\text{LP}}^* = \mathbf{c}^T \mathbf{x}_{\text{LP}}^*$. The *product scoring rule* branches on the variable $i \in [n]$ that maximizes: $\max\{z_{\text{LP}}^* - z_{\text{LP}}^*(x_i \leq \lfloor \mathbf{x}_{\text{LP}}^*[i] \rfloor), 10^{-6}\} \cdot \max\{z_{\text{LP}}^* - z_{\text{LP}}^*(x_i \geq \lceil \mathbf{x}_{\text{LP}}^*[i] \rceil), 10^{-6}\}$.

The tighter the LP relaxation, the more valuable the LP guidance, highlighting the importance of cutting planes.

Polynomial arrangements in Euclidean space. Let $p \in \mathbb{R}[y_1, \dots, y_k]$ be a polynomial of degree at most d . The polynomial p partitions \mathbb{R}^k into connected components that belong to either $\mathbb{R}^k \setminus \{(y_1, \dots, y_k) : p(y_1, \dots, y_k) = 0\}$ or $\{(y_1, \dots, y_k) : p(y_1, \dots, y_k) = 0\}$. When we discuss the connected components of \mathbb{R}^k induced by p , we include connected components in both these sets. We make this distinction because previous work on sample complexity for data-driven algorithm design oftentimes only needed to consider the connected components of the former set. The number of connected components in both sets is $O(d^k)$ [29, 35, 36].

3 Linear programming sensitivity

Our main result in this section characterizes how an LP’s optimal solution is affected by the addition of one or more new constraints. In particular, fixing an LP with m constraints and n variables, if $\mathbf{x}_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta) \in \mathbb{R}^n$ denotes the new LP optimum when the constraint $\boldsymbol{\alpha}^T \mathbf{x} \leq \beta$ is added, we pin down a precise characterization of $\mathbf{x}_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta)$ as a function of $\boldsymbol{\alpha}$ and β . We show that $\mathbf{x}_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta)$ has a piece-wise closed form: there are surfaces partitioning \mathbb{R}^{n+1} such that within each connected component induced by these surfaces, $\mathbf{x}_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta)$ has a closed form. While the geometric intuition used to establish this piece-wise structure relies on the basic property that optimal solutions to LPs are achieved at vertices, the surfaces defining the regions are perhaps surprisingly nonlinear: they are defined by multivariate degree-2 polynomials in $\boldsymbol{\alpha}, \beta$. In Appendix B.1 we illustrate these surfaces for an example two-variable LP.

There are two main steps of our proof: (1) tracking the set of edges of the LP polytope intersected by the new constraint, and once that set of edges is fixed, (2) tracking which edge yields the vertex with the highest objective value.

Let $M = [m]$ denote the set of m constraints. For $E \subseteq M$, let $A_E \in \mathbb{R}^{|E| \times n}$ and $\mathbf{b}_E \in \mathbb{R}^{|E|}$ denote the restrictions of A and \mathbf{b} to E . For $\boldsymbol{\alpha} \in \mathbb{R}^n$, $\beta \in \mathbb{R}$, and $E \subseteq M$ with $|E| = n - 1$, let $A_{E, \boldsymbol{\alpha}} \in \mathbb{R}^{n \times n}$ denote the matrix obtained by adding row vector $\boldsymbol{\alpha}$ to A_E and let $A_{E, \boldsymbol{\alpha}, \beta}^i \in \mathbb{R}^{n \times n}$ be the matrix $A_{E, \boldsymbol{\alpha}}$ with the i th column replaced by $(\mathbf{b}_E, \beta)^T$.

Theorem 3.1. Let $(\mathbf{c}, A, \mathbf{b})$ be an LP and let \mathbf{x}_{LP}^* denote the optimal solution. There is a set of at most m^n hyperplanes and at most m^{2n} degree-2 polynomial hypersurfaces partitioning \mathbb{R}^{n+1} into connected components such that for each component C , one of the following holds: either (1) $\mathbf{x}_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta) = \mathbf{x}_{\text{LP}}^*$ or (2) there is a set of constraints $E \subseteq M$ with $|E| = n - 1$ such that

$$\mathbf{x}_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta) = \left(\frac{\det(A_{E, \boldsymbol{\alpha}, \beta}^1)}{\det(A_{E, \boldsymbol{\alpha}})}, \dots, \frac{\det(A_{E, \boldsymbol{\alpha}, \beta}^n)}{\det(A_{E, \boldsymbol{\alpha}})} \right)$$

for all $(\boldsymbol{\alpha}, \beta) \in C$.

Proof. First, if $\boldsymbol{\alpha}^T \mathbf{x} \leq \beta$ does not separate \mathbf{x}_{LP}^* , then $\mathbf{x}_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta) = \mathbf{x}_{\text{LP}}^*$. The set of all such cuts is the halfspace in \mathbb{R}^{n+1} given by $\{(\boldsymbol{\alpha}, \beta) \in \mathbb{R}^{n+1} : \boldsymbol{\alpha}^T \mathbf{x}_{\text{LP}}^* \leq \beta\}$. All other cuts separate \mathbf{x}_{LP}^* and thus pass through $\mathcal{P} = \{\mathbf{x} \in \mathbb{R}^n : A\mathbf{x} \leq \mathbf{b}, \mathbf{x} \geq \mathbf{0}\}$, and the new LP optimum is achieved at

a vertex created by the cut. We consider the new vertices formed by the cut, which lie on edges (faces of dimension 1) of \mathcal{P} . Letting M denote the set of m constraints that define \mathcal{P} , each edge e of \mathcal{P} can be identified with a subset $E \subset M$ of size $n - 1$ such that the edge is precisely the set of all points \mathbf{x} such that

$$\begin{aligned}\mathbf{a}_i^T \mathbf{x} &= b_i & \forall i \in E \\ \mathbf{a}_i^T \mathbf{x} &\leq b_i & \forall i \in M \setminus E,\end{aligned}$$

where \mathbf{a}_i is the i th row of A . Let $A_E \in \mathbb{R}^{n-1 \times n}$ denote the restriction of A to only the rows in E , and let $\mathbf{b}_E \in \mathbb{R}^{|E|}$ denote the entries of \mathbf{b} corresponding to constraints in E . Drop the inequality constraints defining the edge, so the equality constraints define a line in \mathbb{R}^n . The intersection of the cut $\boldsymbol{\alpha}^T \mathbf{x} \leq \beta$ and this line is precisely the solution to the system of n linear equations in n variables: $A_E \mathbf{x} = \mathbf{b}_E, \boldsymbol{\alpha}^T \mathbf{x} = \beta$. By Cramer's rule, the (unique) solution $\mathbf{x} = (x_1, \dots, x_n)$ to this system is given by $x_i = \frac{\det(A_{E, \alpha, \beta}^i)}{\det(A_{E, \alpha})}$. To ensure that the intersection point indeed lies on the edge of the polytope, we simply stipulate that it satisfies the inequality constraints in $M \setminus E$. That is,

$$\sum_{j=1}^n a_{ij} \cdot \frac{\det(A_{E, \alpha, \beta}^j)}{\det(A_{E, \alpha})} \leq b_i \tag{3}$$

for every $i \in M \setminus E$ (note that if $\boldsymbol{\alpha}, \beta$ satisfy any of these constraints, it must be that $\det(A_{E, \alpha}) \neq 0$, which guarantees that $A_E \mathbf{x} = \mathbf{b}_E, \boldsymbol{\alpha}^T \mathbf{x} = \beta$ indeed has a unique solution). Multiplying through by $\det(A_{E, \alpha})$ shows that this constraint is a halfspace in \mathbb{R}^{n+1} , since $\det(A_{E, \alpha})$ and $\det(A_{E, \alpha, \beta}^i)$ are both linear in $\boldsymbol{\alpha}$ and β . The collection of all the hyperplanes defining the boundaries of these halfspaces over all edges of \mathcal{P} induces a partition of \mathbb{R}^{n+1} into connected components such that for all $(\boldsymbol{\alpha}, \beta)$ within a given connected component, the (nonempty) set of edges of \mathcal{P} that the hyperplane $\boldsymbol{\alpha}^T \mathbf{x} = \beta$ intersects is invariant.

Now, consider a single connected component, denoted by C for brevity. Let e_1, \dots, e_k denote the edges intersected by cuts in C , and let $E_1, \dots, E_k \subset M$ denote the sets of constraints that are binding at each of these edges, respectively. For each pair e_p, e_q , consider the surface

$$\sum_{i=1}^n c_i \cdot \frac{\det(A_{E_p, \alpha, \beta}^i)}{\det(A_{E_p, \alpha})} = \sum_{i=1}^n c_i \cdot \frac{\det(A_{E_q, \alpha, \beta}^i)}{\det(A_{E_q, \alpha})}. \tag{4}$$

Clearing the (nonzero) denominators shows this is a degree-2 polynomial hypersurface in $\boldsymbol{\alpha}, \beta$ in \mathbb{R}^{n+1} . This hypersurface is the set of all $(\boldsymbol{\alpha}, \beta)$ for which the LP objective value achieved at the vertex on edge e_p is equal to the LP objective value achieved at the vertex on edge e_q . The collection of these surfaces for each p, q partitions C into further connected components. Within each of these connected components, the edge containing the vertex that maximizes the objective is invariant. If this edge corresponds to binding constraints E , $\mathbf{x}_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta)$ has the closed form $\mathbf{x}_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta)[i] = \frac{\det(A_{E, \alpha, \beta}^i)}{\det(A_{E, \alpha})}$ for all $(\boldsymbol{\alpha}, \beta)$ within this component. We now count the number of surfaces used to obtain our decomposition. \mathcal{P} has at most $\binom{m}{n-1} \leq m^{n-1}$ edges, and for each edge E we first considered at most $|M \setminus E| \leq m$ hyperplanes representing decision boundaries for cuts intersecting that edge (Equation (3)), for a total of at most m^n hyperplanes. We then considered a degree-2 polynomial hypersurface for every pair of edges (Equation (4)), of which there are at most $\binom{m^n}{2} \leq m^{2n}$. \square

In Appendix B.2, we generalize Theorem 3.1 to understand \mathbf{x}_{LP}^* as a function of any K constraints. In this case, we show that the piecewise structure is given by degree- $2K$ multivariate polynomials.

4 Structure and sensitivity of branch-and-cut

We now use Theorem 3.1 to answer a fundamental question about B&C: how does the B&C tree change when cuts are added at the root? Said another way, what is the structure of the B&C tree as a function of the set of cuts? We prove that the set of all possible cuts can be partitioned into a finite number of regions where by employing cuts from any one region, the B&C tree remains exactly the same. Moreover, we prove that the boundaries between regions are defined by constant-degree polynomials. As in the previous section, we focus on a single cut added to the root of the B&C tree. We provide an extension to multiple cuts in Appendix C.2.

We outline the main steps of our analysis:

1. In Lemma 4.2 we use Theorem 3.1 to understand how the LP optimum at any node in the B&C tree behaves as a function of cuts added at the root.
2. In Lemma 4.3, we analyze how the branching decisions of B&C are impacted by variations in the cuts.
3. In Lemma 4.4, we analyze how cuts affect which nodes are fathomed due to the integrality of the LP relaxation.
4. In Theorem 4.5, we analyze how the LP estimates based on cuts can lead to pruning nodes of the B&C tree, which gives us a complete description of when two cutting planes lead to the same B&C tree.

The full proofs from this section are in Appendix C.

Given an IP, let $\tau = \lceil \max_{\mathbf{x} \in \mathcal{P}} \|\mathbf{x}\|_\infty \rceil$ be the maximum magnitude coordinate of any LP-feasible solution, rounded up. The set of all possible branching constraints is contained in $\mathcal{BC} := \{\mathbf{x}[i] \leq \ell, \mathbf{x}[i] \geq \ell\}_{0 \leq \ell \leq \tau, i \in [n]}$ which is a set of size $2n(\tau + 1)$. Naïvely, there are at most $2^{2n(\tau+1)}$ subsets of branching constraints, but the following observation allows us to greatly reduce the number of sets we consider.

Lemma 4.1. *Fix an IP $(\mathbf{c}, \mathbf{A}, \mathbf{b})$. Define an equivalence relation on pairs of branching-constraint sets $\sigma_1, \sigma_2 \subseteq \mathcal{BC}$, by $\sigma_1 \sim \sigma_2 \iff \mathbf{x}_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma_1) = \mathbf{x}_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma_2)$ for all possible cutting planes $\boldsymbol{\alpha}^T \mathbf{x} \leq \beta$. The number of equivalence classes of \sim is at most τ^{3n} .*

By Cramer's rule, $\tau \leq |\det(\tilde{\mathbf{A}})|$, where $\tilde{\mathbf{A}}$ is any square submatrix of \mathbf{A} . This is at most $a^n n^{n/2}$ by Hadamard's inequality, where a is the maximum absolute value of any entry of \mathbf{A} . However, τ can be much smaller in various cases. For example, if \mathbf{A} contains even one row with only positive entries, then $\tau \leq \|\mathbf{b}\|_\infty$.

We will use the following notation in the remainder of this section. Let A_σ and \mathbf{b}_σ denote the augmented constraint matrix and vector when the constraints in $\sigma \subseteq \mathcal{BC}$ are added. For $E \subseteq M \cup \sigma$, let $A_{E,\sigma} \in \mathbb{R}^{|E| \times n}$ and $\mathbf{b}_E \in \mathbb{R}^{|E|}$ denote the restrictions of A_σ and \mathbf{b}_σ to E . For $\boldsymbol{\alpha} \in \mathbb{R}^n, \beta \in \mathbb{R}$ and $E \subseteq M \cup \sigma$ with $|E| = n - 1$, let $A_{E,\boldsymbol{\alpha},\sigma} \in \mathbb{R}^{n \times n}$ denote the matrix obtained by adding row vector $\boldsymbol{\alpha}$ to $A_{E,\sigma}$ and let $A_{E,\boldsymbol{\alpha},\beta,\sigma}^i \in \mathbb{R}^{n \times n}$ be the matrix $A_{E,\boldsymbol{\alpha},\sigma}$ with the i th column replaced by $(\mathbf{b}_{E,\sigma}, \beta)^T$.

Lemma 4.2. *For any LP $(\mathbf{c}, \mathbf{A}, \mathbf{b})$, there are at most $(m + 2n)^n \tau^{3n}$ hyperplanes and at most $(m + 2n)^{2n} \tau^{3n}$ degree-2 polynomial hypersurfaces partitioning \mathbb{R}^{n+1} into connected components such that for each component C and every $\sigma \subset \mathcal{BC}$, either: (1) $\mathbf{x}_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma) = \mathbf{x}_{\text{LP}}^*(\sigma)$ and $z_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma) = z_{\text{LP}}^*(\sigma)$, or (2) there is a set of constraints $E \subseteq M \cup \sigma$ with $|E| = n - 1$ such that $\mathbf{x}_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma)[i] = \frac{\det(A_{E,\boldsymbol{\alpha},\beta,\sigma}^i)}{\det(A_{E,\boldsymbol{\alpha},\sigma})}$ for all $(\boldsymbol{\alpha}, \beta) \in C$.*

Proof sketch. The same reasoning in the proof of Theorem 3.1 yields a partition with the desired properties. \square

Next, we refine the decomposition obtained in Lemma 4.2 so that the branching constraints added at each step of B&C are invariant within a region. Our results apply to the product scoring rule (Def. 2.2), which is used, for example, by the leading open-source solver SCIP [10].

Lemma 4.3. *There are at most $3(m + 2n)^n \tau^{3n}$ hyperplanes, $3(m + 2n)^{3n} \tau^{4n}$ degree-2 polynomial hypersurfaces, and $(m + 2n)^{6n} \tau^{4n}$ degree-5 polynomial hypersurfaces partitioning \mathbb{R}^{n+1} into connected components such that within each component, the branching constraints used at every step of B&C are invariant.*

Proof sketch. Fix a connected component C in the decomposition established in Lemma 4.2. Then, for each σ , either $\mathbf{x}_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma) = \mathbf{x}_{\text{LP}}^*(\sigma)$ or there exists $E \subseteq M \cup \sigma$ such that $\mathbf{x}_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma)[i] = \frac{\det(A_{E, \boldsymbol{\alpha}, \beta, \sigma}^i)}{\det(A_{E, \boldsymbol{\alpha}, \sigma})}$ for all $(\boldsymbol{\alpha}, \beta) \in C$ and all $i \in [n]$. Now, if we are at a stage in the branch-and-cut tree where σ is the list of branching constraints added so far, and the i th variable is being branched on next, the two constraints generated are $x_i \leq \lfloor \mathbf{x}_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma)[i] \rfloor$ and $x_i \geq \lceil \mathbf{x}_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma)[i] \rceil$, respectively. If C is a component where $\mathbf{x}_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma) = \mathbf{x}_{\text{LP}}^*(\sigma)$, then there is nothing more to do, since the branching constraints at that point are trivially invariant over $(\boldsymbol{\alpha}, \beta) \in C$. Otherwise, in order to further decompose C such that the right-hand-side of these constraints are invariant for every σ and every $i = 1, \dots, n$, we add the two decision boundaries given by

$$k \leq \frac{\det(A_{E, \boldsymbol{\alpha}, \beta, \sigma}^i)}{\det(A_{E, \boldsymbol{\alpha}, \sigma})} \leq k + 1$$

for every i , σ , and every integer $k = 0, \dots, \tau - 1$. This ensures that within every connected component of C induced by these boundaries (hyperplanes), $\lfloor \mathbf{x}_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma)[i] \rfloor$ and $\lceil \mathbf{x}_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma)[i] \rceil$ are invariant. A careful analysis of the definition of the product scoring rule provides the appropriate refinement of this partition. \square

We now move to the most critical phase of branch-and-cut: deciding when to fathom a node. One reason a node might be fathomed is if the LP relaxation of the IP at that node has an integral solution. We derive conditions that ensure that nearby cuts have the same effect on the integrality of the original IP at any node in the search tree. Recall that $\mathcal{P}_I = \mathcal{P} \cap \mathbb{Z}^n$ is the set of integer points in \mathcal{P} . Let $\mathcal{V} \subseteq \mathbb{R}^{n+1}$ denote the set of all valid cuts for the input IP $(\mathbf{c}, A, \mathbf{b})$. The set \mathcal{V} is a polyhedron since it can be expressed as

$$\mathcal{V} = \bigcap_{\bar{\mathbf{x}} \in \mathcal{P}_I} \{(\boldsymbol{\alpha}, \beta) \in \mathbb{R}^{n+1} : \boldsymbol{\alpha}^T \bar{\mathbf{x}} \leq \beta\},$$

and \mathcal{P}_I is finite as \mathcal{P} is bounded. For cuts outside \mathcal{V} , we assume the B&C tree takes some special form denoting an invalid cut. Our goal now is to decompose \mathcal{V} into connected components such that $\mathbf{1}[\mathbf{x}_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma) \in \mathbb{Z}^n]$ is invariant for all $(\boldsymbol{\alpha}, \beta)$ in each component.

Lemma 4.4. *For any IP $(\mathbf{c}, A, \mathbf{b})$, there are at most $3(m + 2n)^n \tau^{4n}$ hyperplanes, $3(m + 2n)^{3n} \tau^{4n}$ degree-2 polynomial hypersurfaces, and $(m + 2n)^{6n} \tau^{4n}$ degree-5 polynomial hypersurfaces partitioning \mathbb{R}^{n+1} into connected components such that for each component C and each $\sigma \subseteq \mathcal{BC}$, $\mathbf{1}[\mathbf{x}_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma) \in \mathbb{Z}^n]$ is invariant for all $(\boldsymbol{\alpha}, \beta) \in C$.*

Proof sketch. Fix a connected component C in the decomposition that includes the facets defining \mathcal{V} and the surfaces obtained in Lemma 4.3. For all σ , $\mathbf{x}_l \in \mathcal{P}_l$, and $i \in [n]$, consider the surface

$$\mathbf{x}_{LP}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma)[i] = \mathbf{x}_l[i]. \quad (5)$$

By Lemma 4.2, this surface is a hyperplane. Clearly, within any connected component of C induced by these hyperplanes, for every σ and $\mathbf{x}_l \in \mathcal{P}_l$, $\mathbf{1}[\mathbf{x}_{LP}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma) = \mathbf{x}_l]$ is invariant. Finally, if $\mathbf{x}_{LP}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma) \in \mathbb{Z}^n$ for some cut $\boldsymbol{\alpha}^T \mathbf{x} \leq \beta$ within a given connected component, $\mathbf{x}_{LP}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma) = \mathbf{x}_l$ for some $\mathbf{x}_l \in \mathcal{P}_l(\sigma) \subseteq \mathcal{P}_l$, which means that $\mathbf{x}_{LP}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma) = \mathbf{x}_l \in \mathbb{Z}^n$ for all cuts $\boldsymbol{\alpha}^T \mathbf{x} \leq \beta$ in that connected component. \square

Suppose for a moment that a node is fathomed by B&C if and only if either the LP at that node is infeasible, or the LP optimal solution is integral—that is, the “bounding” of B&C is suppressed. In this case, the partition of \mathbb{R}^{n+1} obtained in Lemma 4.4 guarantees that the tree built by branch-and-cut is invariant within each connected component. Indeed, since the branching constraints at every node are invariant, and for every σ the integrality of $\mathbf{x}_{LP}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma)$ is invariant, the (bounding-suppressed) B&C tree (and the order in which it is built) is invariant within each connected component in our decomposition. Equipped with this observation, we now analyze the full behavior of B&C.

Theorem 4.5. *Given an IP $(\mathbf{c}, A, \mathbf{b})$, there is a set of at most $O(14^n(m+2n)^{3n^2}\tau^{5n^2})$ polynomial hypersurfaces of degree ≤ 5 partitioning \mathbb{R}^{n+1} into connected components such that the branch-and-cut tree built after adding the cut $\boldsymbol{\alpha}^T \mathbf{x} \leq \beta$ at the root is invariant over all $(\boldsymbol{\alpha}, \beta)$ within a given component.*

Proof sketch. Fix a connected component C in the decomposition induced by the set of hyperplanes and degree-2 hypersurfaces established in Lemma 4.4. Let

$$Q_1, \dots, Q_{i_1}, I_1, Q_{i_1+1}, \dots, Q_{i_2}, I_2, Q_{i_2+1}, \dots \quad (6)$$

denote the nodes of the tree branch-and-cut creates, in order of exploration, under the assumption that a node is pruned if and only if either the LP at that node is infeasible or the LP optimal solution is integral (so the “bounding” of branch-and-bound is suppressed). Here, a node is identified by the list σ of branching constraints added to the input IP. Nodes labeled by Q are either infeasible or have fractional LP optimal solutions. Nodes labeled by I have integral LP optimal solutions and are candidates for the incumbent integral solution at the point they are encountered. (The nodes are functions of $\boldsymbol{\alpha}$ and β , as are the indices i_1, i_2, \dots) By Lemma 4.4 and the observation following it, this ordered list of nodes is invariant over all $(\boldsymbol{\alpha}, \beta) \in C$.

Now, given an node index ℓ , let $I(\ell)$ denote the incumbent node with the highest objective value encountered up until the ℓ th node searched by B&C, and let $z(I(\ell))$ denote its objective value. For each node Q_ℓ , let σ_ℓ denote the branching constraints added to arrive at node Q_ℓ . The hyperplane

$$z_{LP}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma_\ell) = z(I(\ell)) \quad (7)$$

(which is a hyperplane due to Lemma 4.2) partitions C into two subregions. In one subregion, $z_{LP}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma_\ell) \leq z(I(\ell))$, that is, the objective value of the LP optimal solution is no greater than the objective value of the current incumbent integer solution, and so the subtree rooted at Q_ℓ is pruned. In the other subregion, $z_{LP}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma_\ell) > z(I(\ell))$, and Q_ℓ is branched on further. Therefore, within each connected component of C induced by all hyperplanes given by Equation 16 for all ℓ , the set of node within the list (15) that are pruned is invariant. Combined with the surfaces established in Lemma 4.4, these hyperplanes partition \mathbb{R}^{n+1} into connected components such that as $(\boldsymbol{\alpha}, \beta)$ varies within a given component, the tree built by branch-and-cut is invariant. \square

5 Sample complexity bounds for B&C

In this section, we show how the results from the previous section can be used to provide sample complexity bounds for configuring B&C. Our results will apply to families of cuts parameterized by vectors \mathbf{u} from a set \mathcal{U} , such as the family of GMI cuts from Definition 2.1. We assume there is an unknown, application-specific distribution \mathcal{D} over IPs. The learner receives a *training set* $\mathcal{S} \sim \mathcal{D}^N$ of N IPs sampled from this distribution. A *sample complexity guarantee* bounds the number of samples N sufficient to ensure that for any parameter setting $\mathbf{u} \in \mathcal{U}$, the B&C tree size on average over the training set \mathcal{S} is close to the expected B&C tree size. More formally, let $g_{\mathbf{u}}(\mathbf{c}, A, \mathbf{b})$ be the size of the tree B&C builds given the input $(\mathbf{c}, A, \mathbf{b})$ after applying the cut defined by \mathbf{u} at the root. Given $\epsilon > 0$ and $\delta \in (0, 1)$, a sample complexity guarantee bounds the number of samples N sufficient to ensure that with probability $1 - \delta$ over the draw $\mathcal{S} \sim \mathcal{D}^N$, for every parameter setting $\mathbf{u} \in \mathcal{U}$,

$$\left| \frac{1}{N} \sum_{(\mathbf{c}, A, \mathbf{b}) \in \mathcal{S}} g_{\mathbf{u}}(\mathbf{c}, A, \mathbf{b}) - \mathbb{E}[g_{\mathbf{u}}(\mathbf{c}, A, \mathbf{b})] \right| \leq \epsilon. \quad (8)$$

To derive our sample complexity guarantee, we use the notion of *pseudo-dimension* [32]. Let $\mathcal{G} = \{g_{\mathbf{u}} : \mathbf{u} \in \mathcal{U}\}$. The *pseudo-dimension* of \mathcal{G} , denoted $\text{Pdim}(\mathcal{G})$, is the largest integer N for which there exist N IPs $(\mathbf{c}_1, A_1, \mathbf{b}_1), \dots, (\mathbf{c}_N, A_N, \mathbf{b}_N)$ and N thresholds $r_1, \dots, r_N \in \mathbb{R}$ such that for every binary vector $(\sigma_1, \dots, \sigma_N) \in \{0, 1\}^N$, there exists $g_{\mathbf{u}} \in \mathcal{G}$ such that $g_{\mathbf{u}}(\mathbf{c}_i, A_i, \mathbf{b}_i) \geq r_i$ if and only if $\sigma_i = 1$. The number of samples sufficient to ensure that Equation (8) holds is $N = O(\frac{\kappa^2}{\epsilon^2}(\text{Pdim}(\mathcal{G}) + \log \frac{1}{\delta}))$ [32]. Equivalently, for a given number of samples N , the left-hand-side of Equation (8) can be bounded by $\kappa \sqrt{\frac{1}{N}(\text{Pdim}(\mathcal{G}) + \log \frac{1}{\delta})}$.

So far, α, β are parameters that do not depend on the input instance $\mathbf{c}, A, \mathbf{b}$. Suppose now that they do: α, β are functions of $\mathbf{c}, A, \mathbf{b}$ and a parameter vector \mathbf{u} (as they are for GMI cuts). Despite the structure established in the previous section, if α, β can depend on $(\mathbf{c}, A, \mathbf{b})$ in arbitrary ways, one cannot even hope for a finite sample complexity, illustrated by the following impossibility result. The full proofs of all results from this section are in Appendix D.

Theorem 5.1. *There exist functions $\alpha_{\mathbf{c}, A, \mathbf{b}} : \mathcal{U} \rightarrow \mathbb{R}^n$ and $\beta_{\mathbf{c}, A, \mathbf{b}} : \mathcal{U} \rightarrow \mathbb{R}$ such that*

$$\text{Pdim}(\{g_{\mathbf{u}} : \mathbf{u} \in \mathcal{U}\}) = \infty,$$

where \mathcal{U} is any set with $|\mathcal{U}| = |\mathbb{R}|$.

However, in the case of GMI cuts, we show that the cutting plane coefficients parameterized by \mathbf{u} are highly structured. Combining this structure with our analysis of B&C allows us to derive polynomial sample complexity bounds.

Lemma 5.2. *Consider the family of GMI cuts parameterized by $\mathbf{u} \in [-U, U]^m$. There is a set of at most $O(nU^2 \|\mathbf{A}\|_1 \|\mathbf{b}\|_1)$ hyperplanes partitioning $[-U, U]^m$ into connected components such that $\lfloor \mathbf{u}^T \mathbf{a}_i \rfloor$, $\lfloor \mathbf{u}^T \mathbf{b} \rfloor$, and $\mathbf{1}[f_i \leq f_0]$ are invariant, for every i , within each component.*

Proof sketch. We have $f_i = \mathbf{u}^T \mathbf{a}_i - \lfloor \mathbf{u}^T \mathbf{a}_i \rfloor$, $f_0 = \mathbf{u}^T \mathbf{b} - \lfloor \mathbf{u}^T \mathbf{b} \rfloor$, and since $\mathbf{u} \in [-U, U]^m$, $\lfloor \mathbf{u}^T \mathbf{a}_i \rfloor \in [-U \|\mathbf{a}_i\|_1, U \|\mathbf{a}_i\|_1]$ and $\lfloor \mathbf{u}^T \mathbf{b} \rfloor \in [-U \|\mathbf{b}\|_1, U \|\mathbf{b}\|_1]$. For all i , $k_i \in [-U \|\mathbf{a}_i\|_1, U \|\mathbf{a}_i\|_1] \cap \mathbb{Z}$ and $k_0 \in [-U \|\mathbf{b}\|_1, U \|\mathbf{b}\|_1] \cap \mathbb{Z}$, hyperplanes define the two regions

$$\lfloor \mathbf{u}^T \mathbf{a}_i \rfloor = k_i \iff k_i \leq \mathbf{u}^T \mathbf{a}_i < k_i + 1$$

and the hyperplanes defining the two halfspaces

$$\lfloor \mathbf{u}^T \mathbf{b} \rfloor = k_0 \iff k_0 \leq \mathbf{u}^T \mathbf{b} < k_0 + 1.$$

In addition, for each i , consider the hyperplane

$$\mathbf{u}^T \mathbf{a}_i - k_i = \mathbf{u}^T \mathbf{b} - k_0. \quad (9)$$

Within any connected component of \mathbb{R}^m determined by these hyperplanes, $\lfloor \mathbf{u}^T \mathbf{a}_i \rfloor$ and $\lfloor \mathbf{u}^T \mathbf{b} \rfloor$ are constant. Also, $\mathbf{1}[f_i \leq f_0]$ is invariant within each component, since if $\lfloor \mathbf{u}^T \mathbf{a}_i \rfloor = k_i$ and $\lfloor \mathbf{u}^T \mathbf{b} \rfloor = k_0$, $f_i \leq f_0 \iff \mathbf{u}^T \mathbf{a}_i - k_i \leq \mathbf{u}^T \mathbf{b} - k_0$, which is the hyperplane from Equation 9. The lemma follows by counting the hyperplanes. \square

Let $\boldsymbol{\alpha} : [-U, U]^m \rightarrow \mathbb{R}^n$ denote the function taking GMI cut parameters \mathbf{u} to the corresponding vector of coefficients determining the resulting cutting plane, and let $\beta : [-U, U]^m \rightarrow \mathbb{R}$ denote the offset of the resulting cutting plane. So (after multiplying through by $1 - f_0$),

$$\boldsymbol{\alpha}(\mathbf{u})[i] = \begin{cases} f_i(1 - f_0) & \text{if } f_i \leq f_0 \\ f_0(1 - f_i) & \text{if } f_i > f_0 \end{cases}$$

and $\beta(\mathbf{u}) = f_0(1 - f_0)$ (of course f_0 and each f_i are functions of \mathbf{u} , but we suppress this dependence for readability).

The next lemma allows us to transfer the polynomial partition of \mathbb{R}^{n+1} from Theorem 4.5 to a polynomial partition of $[-U, U]^m$, incurring only a factor 2 increase in degree.

Lemma 5.3. *Let $p \in \mathbb{R}[y_1, \dots, y_{n+1}]$ be a polynomial of degree d . Let $D \subseteq [-U, U]^m$ be a connected component from Lemma 5.2. Define $q : D \rightarrow \mathbb{R}$ by $q(\mathbf{u}) = p(\boldsymbol{\alpha}(\mathbf{u}), \beta(\mathbf{u}))$. Then q is a polynomial in \mathbf{u} of degree $2d$.*

Proof. By Lemma 5.2, there are integers k_0, k_i for $i \in [n]$ such that $\lfloor \mathbf{u}^T \mathbf{a}_i \rfloor = k_i$ and $\lfloor \mathbf{u}^T \mathbf{b} \rfloor = k_0$ for all $\mathbf{u} \in D$. Also, the set $S = \{i : f_i \leq f_0\}$ is fixed over all $\mathbf{u} \in D$.

A degree- d polynomial p in variables y_1, \dots, y_{n+1} can be written as $\sum_{T \subseteq [n+1], |T| \leq d} \lambda_T \prod_{i \in T} y_i$ for some coefficients $\lambda_T \in \mathbb{R}$, where $T \subseteq [n+1]$ means that T is a multiset of $[n+1]$. Evaluating at $(\boldsymbol{\alpha}(\mathbf{u}), \beta(\mathbf{u}))$, we get

$$\sum_{|T| \leq d} \lambda_T \prod_{\substack{i \in T \cap S \\ i \neq n+1}} f_i(1 - f_0) \prod_{\substack{i \in T \setminus S \\ i \neq n+1}} f_0(1 - f_i) \prod_{i=n+1}^{i \in T} f_0(1 - f_0).$$

Now, $f_i = \mathbf{u}^T \mathbf{a}_i - k_i$ and $f_0 = \mathbf{u}^T \mathbf{b} - k_0$ are linear in \mathbf{u} . The sum is over all multisets of size at most d , so each monomial consists of the product of at most d degree-2 terms of the form $f_i(1 - f_0)$, $f_0(1 - f_i)$, or $f_0(1 - f_0)$. Thus, $\deg(q) \leq 2d$, as desired. \square

Applying Lemma 5.3 to every polynomial hypersurface in the partition of \mathbb{R}^{n+1} established in Theorem 4.5 yields our main structural result for GMI cuts.

Lemma 5.4. *Consider the family of GMI cuts parameterized by $\mathbf{u} \in [-U, U]^m$. For any IP $(\mathbf{c}, \mathbf{A}, \mathbf{b})$, there are at most $O(nU^2 \|\mathbf{A}\|_1 \|\mathbf{b}\|_1)$ hyperplanes and $2^{O(n^2)}(m + 2n)^{O(n^3)}\tau^{O(n^3)}$ degree-10 polynomial hypersurfaces partitioning $[-U, U]^m$ into connected components such that the B&C tree built after adding the GMI cut defined by \mathbf{u} is invariant over all \mathbf{u} within a single component.*

Bounding the pseudo-dimension of the class of tree-size functions $\{g_{\mathbf{u}} : \mathbf{u} \in [-U, U]^m\}$ is a direct application of the main theorem of Balcan et al. [6] along with standard results bounding the VC dimension of polynomial boundaries [2].

Theorem 5.5. *The pseudo-dimension of the class of tree-size functions $\{g_{\mathbf{u}} : \mathbf{u} \in [-U, U]^m\}$ on the domain of IPs with $\|A\|_1 \leq a$ and $\|\mathbf{b}\|_1 \leq b$ is*

$$O(m \log(abU) + mn^3 \log(m+n) + mn^3 \log \tau).$$

We generalize the analysis of this section to multiple GMI cuts at the root of the B&C tree in Appendix D. The analysis there is more involved since GMI cuts can be applied in sequence, re-solving the LP relaxation after each cut. In particular, GMI cuts applied in sequence have one more parameter than the next, so the hyperplane defined by each GMI cut depends (polynomially) on the parameters defining all GMI cuts before it. We show that if K GMI cuts are sequentially applied at the root, the resulting partition of the parameter space is induced by polynomials of degree $O(K^2)$.

6 Conclusions

In this paper, we investigated fundamental questions about linear and integer programs: given an integer program, how many possible branch-and-cut trees are there if one or more additional feasible constraints can be added? Even more specifically, what is the structure of the branch-and-cut tree as a function of a set of additional constraints? Through a detailed geometric and combinatorial analysis of how additional constraints affect the LP relaxation’s optimal solution, we showed that the branch-and-cut tree is piecewise constant and precisely bounded the number of pieces. We showed that the structural understandings that we developed could be used to prove sample complexity bounds for configuring branch-and-cut.

Acknowledgements

This material is based on work supported by the National Science Foundation under grants CCF-1733556, CCF-1910321, IIS-1901403, and SES-1919453, the ARO under award W911NF2010081, the Defense Advanced Research Projects Agency under cooperative agreement HR00112020003, a Simons Investigator Award, an AWS Machine Learning Research Award, an Amazon Research Award, a Bloomberg Research Grant, and a Microsoft Research Faculty Fellowship.

References

- [1] Tobias Achterberg. *Constraint Integer Programming*. PhD thesis, Technische Universität Berlin, 2007.
- [2] Martin Anthony and Peter Bartlett. *Neural Network Learning: Theoretical Foundations*. Cambridge University Press, 2009.
- [3] Egon Balas, Sebastian Ceria, Gérard Cornuéjols, and N Natraj. Gomory cuts revisited. *Operations Research Letters*, 19(1):1–9, 1996.
- [4] Maria-Florina Balcan. Data-driven algorithm design. In Tim Roughgarden, editor, *Beyond Worst Case Analysis of Algorithms*. Cambridge University Press, 2020.
- [5] Maria-Florina Balcan, Travis Dick, Tuomas Sandholm, and Ellen Vitercik. Learning to branch. In *International Conference on Machine Learning (ICML)*, 2018.

- [6] Maria-Florina Balcan, Dan DeBlasio, Travis Dick, Carl Kingsford, Tuomas Sandholm, and Ellen Vitercik. How much data is sufficient to learn high-performing algorithms? Generalization guarantees for data-driven algorithm design. In *Annual Symposium on Theory of Computing (STOC)*, 2021.
- [7] Maria-Florina Balcan, Siddharth Prasad, Tuomas Sandholm, and Ellen Vitercik. Improved learning bounds for branch-and-cut. *arXiv preprint arXiv:2111.11207*, 2021.
- [8] Maria-Florina Balcan, Siddharth Prasad, Tuomas Sandholm, and Ellen Vitercik. Sample complexity of tree search configuration: Cutting planes and beyond. In *Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2021.
- [9] Dimitris Bertsimas and Jack Dunn. Optimal classification trees. *Machine Learning*, 106(7):1039–1082, 2017.
- [10] Ksenia Bestuzheva, Mathieu Besançon, Wei-Kun Chen, Antonia Chmiela, Tim Donkiewicz, Jasper van Doornmalen, Leon Eifler, Oliver Gaul, Gerald Gamrath, Ambros Gleixner, Leona Gottwald, Christoph Graczyk, Katrin Halbig, Alexander Hoen, Christopher Hojny, Rolf van der Hulst, Thorsten Koch, Marco Lübecke, Stephen J. Maher, Frederic Matterer, Erik Mühmer, Benjamin Müller, Marc E. Pfetsch, Daniel Rehfeldt, Steffan Schlein, Franziska Schlösser, Felipe Serrano, Yuji Shinano, Boro Sofranac, Mark Turner, Stefan Vigerske, Fabian Wegscheider, Philipp Wellner, Dieter Weninger, and Jakob Witzig. The SCIP Optimization Suite 8.0. Technical report, Optimization Online, December 2021. URL http://www.optimization-online.org/DB_HTML/2021/12/8728.html.
- [11] Rudy Bunel, Ilker Turkaslan, Philip H.S. Torr, Pushmeet Kohli, and M. Pawan Kumar. A unified view of piecewise linear neural network verification. In *Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2018.
- [12] Vašek Chvátal. Edmonds polytopes and a hierarchy of combinatorial problems. *Discrete mathematics*, 4(4):305–337, 1973.
- [13] William Cook, Albertus MH Gerards, Alexander Schrijver, and Éva Tardos. Sensitivity theorems in integer linear programming. *Mathematical Programming*, 34(3):251–264, 1986.
- [14] Gérard Cornuéjols. Revival of the Gomory cuts in the 1990’s. *Annals of Operations Research*, 149(1):63–66, 2007.
- [15] Gérard Cornuéjols and Yanjun Li. Elementary closures for integer programs. *Operations Research Letters*, 28(1):1–8, 2001.
- [16] Ralph Gomory. An algorithm for the mixed integer problem. resreport RM-2597, The Rand Corporation, 1960.
- [17] Ralph E. Gomory. Outline of an algorithm for integer solutions to linear programs. *Bulletin of the American Mathematical Society*, 64(5):275 – 278, 1958.
- [18] Rishi Gupta and Tim Roughgarden. A PAC approach to application-specific algorithm selection. *SIAM Journal on Computing*, 46(3):992–1017, 2017.
- [19] Zeren Huang, Kerong Wang, Furui Liu, Hui-Ling Zhen, Weinan Zhang, Mingxuan Yuan, Jianye Hao, Yong Yu, and Jun Wang. Learning to select cuts for efficient mixed-integer programming. *Pattern Recognition*, 123:108353, 2022.

[20] Frank Hutter, Holger H Hoos, Kevin Leyton-Brown, and Thomas Stützle. ParamILS: An automatic algorithm configuration framework. *Journal of Artificial Intelligence Research*, 36(1):267–306, 2009. ISSN 1076-9757.

[21] Robert G Jeroslow. Trivial integer programs unsolvable by branch-and-bound. *Mathematical Programming*, 6(1):105–109, 1974.

[22] Jörg Hendrik Kappes, Markus Speth, Gerhard Reinelt, and Christoph Schnörr. Towards efficient and exact map-inference for large scale discrete computer vision problems via combinatorial optimization. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1752–1758. IEEE, 2013.

[23] Daniel Khashabi, Tushar Khot, Ashish Sabharwal, Peter Clark, Oren Etzioni, and Dan Roth. Question answering via integer programming over semi-structured knowledge. In *International Joint Conference on Artificial Intelligence (IJCAI)*, 2016.

[24] Robert Kleinberg, Kevin Leyton-Brown, and Brendan Lucier. Efficiency through procrastination: Approximately optimal algorithm configuration with runtime guarantees. In *International Joint Conference on Artificial Intelligence (IJCAI)*, 2017.

[25] Robert Kleinberg, Kevin Leyton-Brown, Brendan Lucier, and Devon Graham. Procrastinating with confidence: Near-optimal, anytime, adaptive algorithm configuration. *Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2019.

[26] Ailsa H Land and Alison G Doig. An automatic method of solving discrete programming problems. *Econometrica*, pages 497–520, 1960.

[27] Wu Li. The sharp lipschitz constants for feasible and optimal solutions of a perturbed linear program. *Linear algebra and its applications*, 187:15–40, 1993.

[28] Olvi L Mangasarian and T-H Shiau. Lipschitz continuity of solutions of linear inequalities, programs and complementarity problems. *SIAM Journal on Control and Optimization*, 25(3):583–595, 1987.

[29] John Milnor. On the Betti numbers of real varieties. *Proceedings of the American Mathematical Society*, 15(2):275–280, 1964.

[30] Atsushi Miyauchi, Tomohiro Sonobe, and Noriyoshi Sukegawa. Exact clustering via integer programming and maximum satisfiability. In *AAAI Conference on Artificial Intelligence*, 2018.

[31] George Nemhauser and Laurence Wolsey. *Integer and Combinatorial Optimization*. John Wiley & Sons, 1999.

[32] David Pollard. *Convergence of Stochastic Processes*. Springer, 1984.

[33] Tuomas Sandholm. Very-large-scale generalized combinatorial multi-attribute auctions: Lessons from conducting \$60 billion of sourcing. In Zvika Neeman, Alvin Roth, and Nir Vulkan, editors, *Handbook of Market Design*. Oxford University Press, 2013.

[34] Yunhao Tang, Shipra Agrawal, and Yuri Faenza. Reinforcement learning for integer programming: Learning to cut. *International Conference on Machine Learning (ICML)*, 2020.

[35] René Thom. Sur l’homologie des variétés algébriques réelles. In *Differential and combinatorial topology*, pages 255–265. Princeton University Press, 1965.

- [36] Hugh E Warren. Lower bounds for approximation by nonlinear manifolds. *Transactions of the American Mathematical Society*, 133(1):167–178, 1968.
- [37] Jiaming Zeng, Berk Ustun, and Cynthia Rudin. Interpretable classification models for recidivism prediction. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 180(3):689–722, 2017.

A Further details about plots

The version of the *facility location* problem we study involves a set of locations J and a set of clients C . Facilities are to be constructed at some subset of the locations, and the clients in C are served by these facilities. Each location $j \in J$ has a cost f_j of being the site of a facility, and a cost $s_{c,j}$ of serving client $c \in C$. Finally, each location j has a capacity κ_j which is a limit on the number of clients j can serve. The goal of the facility location problem is to arrive at a feasible set of locations for facilities and a feasible assignment of clients to these locations that minimizes the overall cost incurred.

The facility location problem can be formulated as the following 0,1 IP:

$$\begin{aligned} \text{minimize} \quad & \sum_{j \in J} f_j x_j + \sum_{j \in J} \sum_{c \in C} s_{c,j} y_{c,j} \\ \text{subject to} \quad & \sum_{j \in J} y_{c,j} = 1 \quad \forall c \in C \\ & \sum_{c \in C} y_{c,j} \leq \kappa_j x_j \quad \forall j \in J \\ & y_{c,j} \in \{0, 1\} \quad \forall c \in C, j \in J \\ & x_j \in \{0, 1\} \quad \forall j \in J \end{aligned}$$

We consider the following two distributions over facility location IPs.

First distribution Facility location IPs are generated by perturbing the costs and capacities of a base facility location IP. We generated the base IP with 40 locations and 40 clients by choosing the location costs and client-location costs uniformly at random from $[0, 100]$ and the capacities uniformly at random from $\{0, \dots, 39\}$. To sample from the distribution, we perturb this base IP by adding independent Gaussian noise with mean 0 and standard deviation 10 to the cost of each location, the cost of each client-location pair, and the capacity of each location.

Second distribution Facility location IPs are generated by placing 80 evenly-spaced locations along the line segment connecting the points $(0, 1/2)$ and $(1, 1/2)$ in the Cartesian plane. The location costs are all uniformly set to 1. Then, 80 clients are placed uniformly at random in the unit square $[0, 1]^2$. The cost $s_{c,j}$ of serving client c from location j is the distance between j and c . Location capacities are chosen uniformly at random from $\{0, \dots, 43\}$.

In our experiments, we add five cuts at the root of the B&C tree. These five cuts come from the set of Chvátal-Gomory and Gomory mixed integer cuts derived from the optimal simplex tableau of the LP relaxation. The five cuts added are chosen to maximize a weighting of cutting-plane scores:

$$\mu \cdot \mathbf{score}_1 + (1 - \mu) \cdot \mathbf{score}_2. \quad (10)$$

\mathbf{score}_1 is the *parallelism* of a cut, which intuitively measures the angle formed by the objective vector and the normal vector of the cutting plane—promoting cutting planes that are nearly parallel with the objective direction. \mathbf{score}_2 is the *efficacy*, or depth, of a cut, which measures the perpendicular distance from the LP optimum to the cut—promoting cutting planes that are “deeper”, as measured with respect to the LP optimum. More details about these scoring rules can be found in Balcan et al. [8] and references therein. Given an IP, for each $\mu \in [0, 1]$ (discretized at steps of 0.01) we choose the five cuts among the set of Chvátal-Gomory and Gomory mixed integer cuts that maximize (10). Figures 1 and 2 display the average tree size over 1000 samples drawn from the respective distribution for each value of μ used to choose cuts at the root. We ran our experiments using the C API of IBM ILOG CPLEX 20.1.0, with default cut generation disabled.

B Omitted results and proofs from Section 3

B.1 Example in two dimensions

Consider the LP

$$\max\{x + y : x \leq 1, y \geq 0, y \leq x\}.$$

The optimum is at $(x^*, y^*) = (1, 1)$. Consider adding an additional constraint $\alpha_1 x + \alpha_2 y \leq 1$. Let h denote the hyperplane $\alpha_1 x + \alpha_2 y = 1$. We derive a description of the set of parameters (α_1, α_2) such that h intersects the hyperplanes $x = 1$ and $y = x$. The intersection of h and $x = 1$ is given by

$$(x, y) = \left(1, \frac{1 - \alpha_1}{\alpha_2}\right),$$

which exists if and only if $\alpha_2 \neq 0$. This intersection point is in the LP feasible region if and only if $0 \leq \frac{1 - \alpha_1}{\alpha_2} \leq 1$ (which additionally ensures that $\alpha_2 \neq 0$). Similarly, h intersects $y = x$ at

$$(x, y) = \left(\frac{1}{\alpha_1 + \alpha_2}, \frac{1}{\alpha_1 + \alpha_2}\right),$$

which exists if and only if $\alpha_1 + \alpha_2 \neq 0$. This intersection point is in the LP feasible region if and only if $0 \leq \frac{1}{\alpha_1 + \alpha_2} \leq 1$. Now, we put down an ‘‘indifference’’ curve in (α_1, α_2) -space that represents the set of (α_1, α_2) such that the value of the objective achieved at the two aforementioned intersection points is equal. This surface is given by

$$\frac{2}{\alpha_1 + \alpha_2} = 1 + \frac{1 - \alpha_1}{\alpha_2}.$$

Since $\alpha_1 + \alpha_2 \neq 0$ and $\alpha_2 \neq 0$ (for the relevant α_1, α_2 in consideration), this is equivalent to $\alpha_1^2 - \alpha_2^2 - \alpha_1 + \alpha_2 = 0$, which is a degree-2 curve in α_1, α_2 . The left-hand-side can be factored to write this as $(\alpha_1 - \alpha_2)(\alpha_1 + \alpha_2 - 1) = 0$. Therefore, this curve is given by the two lines $\alpha_1 = \alpha_2$ and $\alpha_1 + \alpha_2 = 1$. Figure 3 illustrates the resulting partition of (α_1, α_2) -space.

It turns out that when $n = 2$ the indifference curve can always be factored into a product of linear terms. Let the objective of the LP be (c_1, c_2) , and let $s_1 x + s_2 y = u_1$ and $t_1 x + t_2 y = v_1$ be two intersecting edges of the LP feasible region. Let $\alpha_1 x + \alpha_2 y = \beta$ be an additional constraint. The intersection points of this constraint with the two lines, if they exist, are given by

$$\left(\frac{s_2 \beta - u \alpha_2}{s_2 \alpha_1 - s_1 \alpha_2}, \frac{s_1 \beta - u \alpha_1}{s_1 \alpha_2 - s_2 \alpha_1}\right) \text{ and } \left(\frac{t_2 \beta - v \alpha_2}{t_2 \alpha_1 - t_1 \alpha_2}, \frac{t_1 \beta - v \alpha_1}{t_1 \alpha_2 - t_2 \alpha_1}\right).$$

The indifference surface is thus given by

$$c_1 \frac{s_2 \beta - u \alpha_2}{s_2 \alpha_1 - s_1 \alpha_2} + c_2 \frac{s_1 \beta - u \alpha_1}{s_1 \alpha_2 - s_2 \alpha_1} = c_1 \frac{t_2 \beta - v \alpha_2}{t_2 \alpha_1 - t_1 \alpha_2} + c_2 \frac{t_1 \beta - v \alpha_1}{t_1 \alpha_2 - t_2 \alpha_1}.$$

For α_1, α_2 such that $s_2 \alpha_1 - s_1 \alpha_2 \neq 0$ and $t_2 \alpha_1 - t_1 \alpha_2 \neq 0$, clearing denominators and some manipulation yields

$$(c_1 \alpha_2 - c_2 \alpha_1)((u t_1 - v s_1) \alpha_2 - (u t_2 - v s_2) \alpha_1 + (s_2 t_2 - t_1 s_2) \beta) = 0.$$

This curve consists of the two planes $c_1 \alpha_2 - c_2 \alpha_1 = 0$ and $(u t_1 - v s_1) \alpha_2 - (u t_2 - v s_2) \alpha_1 + (s_2 t_2 - t_1 s_2) \beta = 0$.

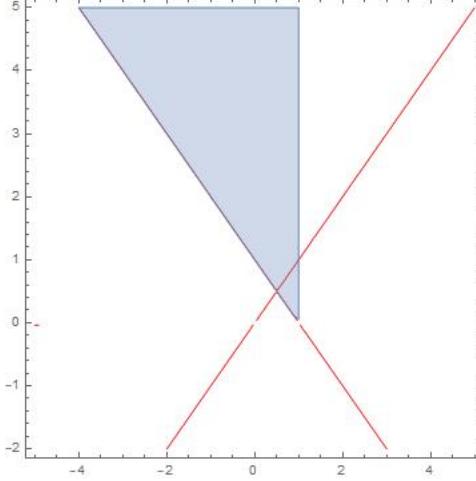


Figure 3: Decomposition of the parameter space: the blue region contains the set of (α_1, α_2) such that the constraint intersects the feasible region at $x = 1$ and $x = y$. The red lines consist of all (α_1, α_2) such that the objective value is equal at these intersection points. The red lines partition the blue region into two components: one where the new optimum is achieved at the intersection of h and $x = y$, and one where the new optimum is achieved at the intersection of h and $x = 1$.

This is however not true if $n > 2$. For example, consider an LP in three variables x, y, z with the constraints $x + y \leq 1, x + z \leq 1, x \leq 1, z \leq 1$. Writing out the indifference surface (assuming the objective is $\mathbf{c} = (1, 1, 1)^T$) for the vertex on the intersection of $\{x + y = 1, x = 1\}$ and the vertex on $\{x + z = 1, z = 1\}$ yields

$$\alpha_1\alpha_2 - \alpha_2\beta - \alpha_3^2 + \alpha_3\beta = 0.$$

Setting $\beta = 1$, we can plot the resulting surface in $\alpha_1, \alpha_2, \alpha_3$ (Figure 4).

B.2 Linear programming sensitivity for multiple constraints

Lemma B.1. *Let $(\mathbf{c}, A, \mathbf{b})$ be an LP and let M denote the set of its m constraints. Let \mathbf{x}_{LP}^* and z_{LP}^* denote the optimal solution and its objective value, respectively. For $F \subseteq M$, let $A_F \in \mathbb{R}^{|F| \times n}$ and $\mathbf{b}_F \in \mathbb{R}^{|F|}$ denote the restrictions of A and \mathbf{b} to F . For $k \leq n$, $\alpha_1, \dots, \alpha_k \in \mathbb{R}^n$, $\beta_1, \dots, \beta_k \in \mathbb{R}$, and $F \subseteq M$ with $|F| = n - k$, let $A_{F, \alpha_1, \dots, \alpha_k} \in \mathbb{R}^{n \times n}$ denote the matrix obtained by adding row vectors $\alpha_1, \dots, \alpha_k$ to A_F and let $A_{F, \alpha_1, \dots, \alpha_k, \beta_k}^i \in \mathbb{R}^{n \times n}$ be the matrix $A_{F, \alpha_1, \dots, \alpha_k} \in \mathbb{R}^{n \times n}$ with the i th column replaced by $[\mathbf{b}_F \ \beta_1 \ \dots \ \beta_k]^T$. There is a set of at most K hyperplanes, $nK^n m^n$ degree- K polynomial hypersurfaces, and $nK^n m^{2n}$ degree- $2K$ polynomial hypersurfaces partitioning $\mathbb{R}^{K(n+1)}$ into connected components such that for each component C , one of the following holds: either (1) $\mathbf{x}_{\text{LP}}^*(\alpha_1^T \mathbf{x} \leq \beta_1, \dots, \alpha_K^T \mathbf{x} \leq \beta_K) = \mathbf{x}_{\text{LP}}^*$, or (2) there is a subset of cuts indexed by $\ell_1, \dots, \ell_k \in [K]$ and a set of constraints $F \subseteq M$ with $|F| = n - k$ such that*

$$\mathbf{x}_{\text{LP}}^*(\alpha_1^T \mathbf{x} \leq \beta_1, \dots, \alpha_K^T \mathbf{x} \leq \beta_K) = \left(\frac{\det(A_{F, \alpha_{\ell_1}, \dots, \alpha_{\ell_k}, \beta_{\ell_k}}^1)}{\det(A_{F, \alpha_1, \dots, \alpha_k})}, \dots, \frac{\det(A_{F, \alpha_{\ell_1}, \dots, \alpha_{\ell_k}, \beta_{\ell_k}}^n)}{\det(A_{F, \alpha_1, \dots, \alpha_k})} \right),$$

for all $(\alpha_1, \beta_1, \dots, \alpha_K, \beta_K) \in C$.

Proof. First, if none of $\alpha_1^T \mathbf{x} \leq \beta_1, \dots, \alpha_K^T \mathbf{x} \leq \beta_K$ separate \mathbf{x}_{LP}^* , then $\mathbf{x}_{\text{LP}}^*(\alpha_1^T \mathbf{x} \leq \beta_1, \dots, \alpha_K^T \mathbf{x} \leq \beta_K) = \mathbf{x}_{\text{LP}}^*$ and $z_{\text{LP}}^*(\alpha_1^T \mathbf{x} \leq \beta_1, \dots, \alpha_K^T \mathbf{x} \leq \beta_K) = z_{\text{LP}}^*$. The set of all such cuts is given by the

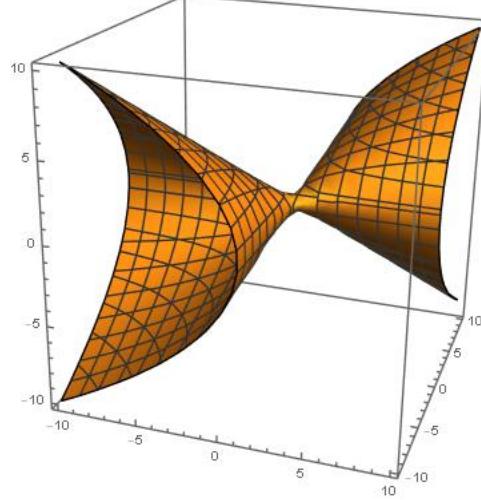


Figure 4: Indifference surface for two edges of the feasible region of an LP in three variables.

intersection of halfspaces in $\mathbb{R}^{K(n+1)}$ given by

$$\bigcap_{j=1}^K \left\{ (\boldsymbol{\alpha}_1, \beta_1, \dots, \boldsymbol{\alpha}_k, \beta_k) \in \mathbb{R}^{K(n+1)} : \boldsymbol{\alpha}_j^T \mathbf{x}_{\text{LP}}^* \leq \beta_j \right\}. \quad (11)$$

All other vectors of K cuts contain at least one cut that separates \mathbf{x}_{LP}^* , and those cuts therefore pass through $\mathcal{P} = \{\mathbf{x} \in \mathbb{R}^n : A\mathbf{x} \leq \mathbf{b}, \mathbf{x} \geq \mathbf{0}\}$. The new LP optimum is thus achieved at a vertex created by the cuts that separate \mathbf{x}_{LP}^* . As in the proof of Theorem 3.1, we consider all possible new vertices formed by our set of K cuts. In the case of a single cut, these new vertices necessarily were on edges of \mathcal{P} , but now they may lie on higher dimensional faces.

Consider a subset of $k \leq n$ cuts that separate \mathbf{x}_{LP}^* . Without loss of generality, denote these cuts by $\boldsymbol{\alpha}_1^T \mathbf{x} \leq \beta_1, \dots, \boldsymbol{\alpha}_k^T \mathbf{x} \leq \beta_k$. We now establish conditions for these k cuts to “jointly” form a new vertex of \mathcal{P} . Any vertex created by these cuts must lie on a face f of \mathcal{P} with $\dim(f) = k$ (in the case that $k = n$, the relevant face f with $\dim(f) = n$ is \mathcal{P} itself). Letting M denote the set of m constraints that define \mathcal{P} , each dimension- k face f of \mathcal{P} can be identified with a (potentially empty) subset $F \subset M$ of size $n - k$ such that f is precisely the set of all points \mathbf{x} such that

$$\begin{aligned} \mathbf{a}_i^T \mathbf{x} &= b_i & \forall i \in F \\ \mathbf{a}_i^T \mathbf{x} &\leq b_i & \forall i \in M \setminus F, \end{aligned}$$

where \mathbf{a}_i is the i th row of A . Let $A_F \in \mathbb{R}^{n-k \times n}$ denote the restriction of A to only the rows in F , and let $\mathbf{b}_F \in \mathbb{R}^{n-k}$ denote the entries of \mathbf{b} corresponding to the constraints in F . Consider removing the inequality constraints defining the face. The intersection of the cuts $\boldsymbol{\alpha}_1^T \mathbf{x} \leq \beta_1, \dots, \boldsymbol{\alpha}_k^T \mathbf{x} \leq \beta_k$ and this unbounded surface (if it exists) is precisely the solution to the system of n linear equations

$$A_F \mathbf{x} = \mathbf{b}_F$$

$$\boldsymbol{\alpha}_1^T \mathbf{x} = \beta_1$$

$$\vdots$$

$$\boldsymbol{\alpha}_k^T \mathbf{x} = \beta_k.$$

Let $A_{F, \alpha_1, \dots, \alpha_k} \in \mathbb{R}^{n \times n}$ denote the matrix obtained by adding row vectors $\alpha_1, \dots, \alpha_k$ to A_F , and let $A_{F, \alpha_1, \beta_1, \dots, \alpha_k, \beta_k}^i \in \mathbb{R}^{n \times n}$ denote the matrix $A_{F, \alpha_1, \dots, \alpha_k}$ where the i th column is replaced by

$$\begin{bmatrix} \mathbf{b}_F \\ \beta_1 \\ \vdots \\ \beta_k \end{bmatrix} \in \mathbb{R}^n.$$

By Cramer's rule, the solution to this system is given by

$$\mathbf{x} = \left(\frac{\det(A_{F, \alpha_1, \beta_1, \dots, \alpha_k, \beta_k}^1)}{\det(A_{F, \alpha_1, \dots, \alpha_k})}, \dots, \frac{\det(A_{F, \alpha_1, \beta_1, \dots, \alpha_k, \beta_k}^n)}{\det(A_{F, \alpha_1, \dots, \alpha_k})} \right),$$

and the value of the objective at this point is

$$\mathbf{c}^T \mathbf{x} = \sum_{i=1}^n c_i \cdot \frac{\det(A_{F, \alpha_1, \beta_1, \dots, \alpha_k, \beta_k}^i)}{\det(A_{F, \alpha_1, \dots, \alpha_k})}.$$

Now, to ensure that the unique intersection point \mathbf{x} (1) exists and (2) actually lies on f (or simply lies in \mathcal{P} , in the case that $F = \emptyset$), we stipulate that it satisfies the inequality constraints in $M \setminus F$. That is,

$$\sum_{j=1}^n a_{ij} \frac{\det(A_{F, \alpha_1, \beta_1, \dots, \alpha_k, \beta_k}^1)}{\det(A_{F, \alpha_1, \dots, \alpha_k})} \leq b_i \quad (12)$$

for every $i \in M \setminus F$. If $\alpha_1, \beta_1, \dots, \alpha_k, \beta_k$ satisfies any of these constraints, it must be that $\det(A_{F, \alpha_1, \dots, \alpha_k}) \neq 0$, which guarantees that $A_F \mathbf{x} = \mathbf{b}_F, \alpha_1^T \mathbf{x} = \beta_1, \dots, \alpha_k^T \mathbf{x} = \beta_k$ indeed has a unique solution. Now, $\det(A_{F, \alpha_1, \dots, \alpha_k})$ is a polynomial in $\alpha_1, \dots, \alpha_k$ of degree $\leq k$, since it is multilinear in each coefficient of each α_ℓ , $\ell = 1, \dots, k$. Similarly, $\det(A_{F, \alpha_1, \beta_1, \dots, \alpha_k, \beta_k}^1)$ is a polynomial in $\alpha_1, \beta_1, \dots, \alpha_k, \beta_k$ of degree $\leq k$, again because it is multilinear in each cut parameter. Hence, the boundary each constraint of the form given by Equation 12 is a polynomial of degree at most k .

The collection of these polynomials for every k , every subset of $\{\alpha_1^T \mathbf{x} \leq \beta_1, \dots, \alpha_K^T \mathbf{x} \leq \beta_K\}$ of size k , and every face of \mathcal{P} of dimension k , along with the hyperplanes determining separation constraints (Equation 11), partition $\mathbb{R}^{K(n+1)}$ into connected components such that for all $(\alpha_1, \beta_1, \dots, \alpha_K, \beta_K)$ within a given connected component, there is a fixed subset of K and a fixed set of faces of \mathcal{P} such that the cuts with indices in that subset intersect every face in the set at a common vertex.

Now, consider a single connected component, denoted by C . Let f_1, \dots, f_ℓ denote the faces intersected by vectors of cuts in C , and let (without loss of generality) $1, \dots, k$ denote the subset of cuts that intersect these faces. Let $F_1, \dots, F_\ell \subset M$ denote the sets of constraints that are binding at each of these faces, respectively. For each pair f_p, f_q , consider the surface

$$\sum_{i=1}^n c_i \cdot \frac{\det(A_{F_p, \alpha_1, \beta_1, \dots, \alpha_k, \beta_k}^i)}{\det(A_{F_p, \alpha_1, \dots, \alpha_k})} = \sum_{i=1}^n c_i \cdot \frac{\det(A_{F_q, \alpha_1, \beta_1, \dots, \alpha_k, \beta_k}^i)}{\det(A_{F_q, \alpha_1, \dots, \alpha_k})},$$

which can be equivalently written as

$$\sum_{i=1}^n c_i \cdot \det(A_{F_p, \alpha_1, \beta_1, \dots, \alpha_k, \beta_k}^i) \det(A_{F_q, \alpha_1, \dots, \alpha_k}) = \sum_{i=1}^n c_i \cdot \det(A_{F_q, \alpha_1, \beta_1, \dots, \alpha_k, \beta_k}^i) \det(A_{F_p, \alpha_1, \dots, \alpha_k}). \quad (13)$$

This is a degree- $2k$ polynomial hypersurface in $(\boldsymbol{\alpha}_1, \beta_1, \dots, \boldsymbol{\alpha}_K, \beta_K) \in \mathbb{R}^{K(n+1)}$. This hypersurface is precisely the set of all cut vectors for which the LP objective achieved at the vertex on face f_p is equal to the LP objective value achieved at the vertex on face f_q . The collection of these surfaces for each p, q partitions C into further connected components. Within each of these connected components, the face containing the vertex that maximizes the objective is invariant, and the subset of cuts passing through that vertex is invariant. If $F \subseteq M$ is the set of binding constraints representing this face, and $\ell_1, \dots, \ell_k \in [K]$ represent the subset of cuts intersecting this face, $\mathbf{x}_{\text{LP}}^*(\boldsymbol{\alpha}_1^T \mathbf{x} \leq \beta_1, \dots, \boldsymbol{\alpha}_K^T \mathbf{x} \leq \beta_K)$ and $z_{\text{LP}}^*(\boldsymbol{\alpha}_1^T \mathbf{x} \leq \beta_1, \dots, \boldsymbol{\alpha}_K^T \mathbf{x} \leq \beta_K)$ have the closed forms:

$$\mathbf{x}_{\text{LP}}^*(\boldsymbol{\alpha}_1^T \mathbf{x} \leq \beta_1, \dots, \boldsymbol{\alpha}_K^T \mathbf{x} \leq \beta_K) = \left(\frac{\det(A_{F, \boldsymbol{\alpha}_{\ell_1}, \beta_1, \dots, \boldsymbol{\alpha}_{\ell_k}, \beta_{\ell_k}}^1)}{\det(A_{F, \boldsymbol{\alpha}_{\ell_1}, \dots, \boldsymbol{\alpha}_{\ell_k}})}, \dots, \frac{\det(A_{F, \boldsymbol{\alpha}_{\ell_1}, \beta_1, \dots, \boldsymbol{\alpha}_{\ell_k}, \beta_{\ell_k}}^n)}{\det(A_{F, \boldsymbol{\alpha}_{\ell_1}, \dots, \boldsymbol{\alpha}_{\ell_k}})} \right),$$

and

$$z_{\text{LP}}^*(\boldsymbol{\alpha}_1^T \mathbf{x} \leq \beta_1, \dots, \boldsymbol{\alpha}_K^T \mathbf{x} \leq \beta_K) = \sum_{i=1}^n c_i \cdot \frac{\det(A_{F, \boldsymbol{\alpha}_{\ell_1}, \beta_1, \dots, \boldsymbol{\alpha}_{\ell_k}, \beta_{\ell_k}}^i)}{\det(A_{F, \boldsymbol{\alpha}_{\ell_1}, \dots, \boldsymbol{\alpha}_{\ell_k}})}.$$

for all $(\boldsymbol{\alpha}_1, \beta_1, \dots, \boldsymbol{\alpha}_K, \beta_K)$ within this component. We now count the number of surfaces used to obtain our decomposition. First, we added K hyperplanes encoding separation constraints for each of the K cuts (Equation 11). Then, for every subset $S \subseteq K$ of size $\leq n$, and for every face F of \mathcal{P} with $\dim(F) = |S|$, we first considered at most $|M \setminus F| \leq m$ degree- $\leq K$ polynomial hypersurfaces representing decision boundaries for when cuts in S intersected that face (Equation 12). The number of k -dimensional faces of \mathcal{P} is at most $\binom{m}{n-k} \leq m^{n-k} \leq m^{n-1}$, so the total number of these hypersurfaces is at most $((\binom{K}{0} + \dots + \binom{K}{n})m^n \leq nK^n m^n)$. Finally, we considered a degree- $2K$ polynomial hypersurface for every subset of cuts and every pair of faces with degree equal to the size of the subset, of which there are at most $nK^n \binom{m^n}{2} \leq nK^n m^{2n}$. \square

C Omitted results and proofs from Section 4

Proof of Lemma 4.1. Consider as an example $\sigma_1 = \{\mathbf{x}[1] \leq 1, \mathbf{x}[1] \leq 5\}$ and $\sigma_2 = \{\mathbf{x}[1] \leq 1\}$. We have $\mathbf{x}_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma_1) = \mathbf{x}_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma_2)$ for any cut $\boldsymbol{\alpha}^T \mathbf{x} \leq \beta$, because the constraint $\mathbf{x}[1] \leq 5$ is redundant in σ_1 . More generally, any $\sigma \subseteq \mathcal{BC}$ can be reduced by preserving only the tightest \leq constraint and tightest \geq constraint without affecting the resulting LP optimal solutions. The number of such unique reduced sets is at most $((\tau + 2)^2)^n < \tau^{3n}$ (for each variable, there are $\tau + 2$ possibilities for the tightest \leq constraint: no constraint or one of $\mathbf{x}[i] \leq 0, \dots, \mathbf{x}[i] \leq \tau$, and similarly $\tau + 2$ possibilities for the \geq constraint). \square

Proof of Lemma 4.2. We carry out the same reasoning in the proof of Theorem 3.1 for each reduced σ . The number of edges of $\mathcal{P}(\sigma)$ is at most $\binom{m+|\sigma|}{n-1} \leq (m+|\sigma|)^{n-1}$. For each edge E , we considered at most $|(M \cup \sigma) \setminus E| \leq m + |\sigma|$ hyperplanes, for a total of at most $(m + |\sigma|)^n$ halfspaces. Then, we had a degree-2 polynomial hypersurface for every pair of edges, of which there are at most $\binom{m+|\sigma|}{2}^n \leq (m + |\sigma|)^{2n}$. Summing over all reduced σ (of which there are at most τ^{3n}), combined with the fact that if σ is reduced then $|\sigma| \leq 2n$, we get a total of at most $(m + 2n)^n \tau^{3n}$ hyperplanes and at most $(m + 2n)^{2n} \tau^{3n}$ degree-2 hypersurfaces, as desired. \square

Proof of Lemma 4.4. Fix a connected component C in the decomposition that includes the facets defining \mathcal{V} and the surfaces obtained in Lemma 4.3. For all $\sigma \in \mathcal{BC}$, $\mathbf{x}_1 \in \mathcal{P}_1$, and $i = 1, \dots, n$, consider the surface

$$\mathbf{x}_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma)[i] = \mathbf{x}_1[i]. \quad (14)$$

This surface is a hyperplane, since by Lemma 4.2, either $\mathbf{x}_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma)[i] = \mathbf{x}_{\text{LP}}^*(\sigma)[i]$ or $\mathbf{x}_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma)[i] = \frac{\det(A_{E, \boldsymbol{\alpha}, \beta, \sigma}^i)}{\det(A_{E, \boldsymbol{\alpha}, \sigma})}$, where $E \subseteq M \cup \sigma$ is the subset of constraints corresponding to σ and C . Clearly, within any connected component of C induced by these hyperplanes, for every σ and $\mathbf{x}_I \in \mathcal{P}_I$, $\mathbf{1}[\mathbf{x}_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma) = \mathbf{x}_I]$ is invariant. Finally, if $\mathbf{x}_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma) \in \mathbb{Z}^n$ for some cut $\boldsymbol{\alpha}^T \mathbf{x} \leq \beta$ within a given connected component, $\mathbf{x}_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma) = \mathbf{x}_I$ for some $\mathbf{x}_I \in \mathcal{P}_{\text{IH}}(\sigma) \subseteq \mathcal{P}_I$, which means that $\mathbf{x}_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma) = \mathbf{x}_I \in \mathbb{Z}^n$ for all cuts $\boldsymbol{\alpha}^T \mathbf{x} \leq \beta$ in that connected component.

We now count the number of hyperplanes given by Equation 14. For each σ , there are $\binom{m+|\sigma|}{n-1} \leq (m+2n)^{n-1}$ binding edge constraints $E \subseteq M \cup \sigma$ defining the formula of Lemma 4.2, and we have $n|\mathcal{P}_I|$ hyperplanes for each E . Since $\tau = \max_{\mathbf{x} \in \mathcal{P}_I} \|\mathbf{x}\|_\infty$, $|\mathcal{P}_I| \leq \tau^n$. So the total number of hyperplanes given by Equation 14 is at most $\tau^{3n}(m+2n)^{n-1}n\tau^n \leq (m+2n)^n\tau^{4n}$. The number of facets defining \mathcal{V} is at most $|\mathcal{P}_{\text{IH}}| \leq |\mathcal{P}_I| \leq \tau^n$. Adding these to the counts obtained in Lemma 4.3 yields the final tallies in the lemma statement. \square

Proof of Theorem 4.5. Fix a connected component C in the decomposition induced by the set of hyperplanes and degree-2 hypersurfaces established in Lemma 4.4. Let

$$Q_1, \dots, Q_{i_1}, I_1, Q_{i_1+1}, \dots, Q_{i_2}, I_2, Q_{i_2+1}, \dots \quad (15)$$

denote the nodes of the tree branch-and-cut creates, in order of exploration, under the assumption that a node is pruned if and only if either the LP at that node is infeasible or the LP optimal solution is integral (so the “bounding” of branch-and-bound is suppressed). Here, a node is identified by the list σ of branching constraints added to the input IP. Nodes labeled by Q are either infeasible or have fractional LP optimal solutions. Nodes labeled by I have integral LP optimal solutions and are candidates for the incumbent integral solution at the point they are encountered. (The nodes are functions of $\boldsymbol{\alpha}$ and β , as are the indices i_1, i_2, \dots) By Lemma 4.4 and the observation following it, this ordered list of nodes is invariant over all $(\boldsymbol{\alpha}, \beta) \in C$.

Now, given an node index ℓ , let $I(\ell)$ denote the incumbent node with the highest objective value encountered up until the ℓ th node searched by B&C, and let $z(I(\ell))$ denote its objective value. For each node Q_ℓ , let σ_ℓ denote the branching constraints added to arrive at node Q_ℓ . The hyperplane

$$z_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma_\ell) = z(I(\ell)) \quad (16)$$

(which is a hyperplane due to Lemma 4.2) partitions C into two subregions. In one subregion, $z_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma_\ell) \leq z(I(\ell))$, that is, the objective value of the LP optimal solution is no greater than the objective value of the current incumbent integer solution, and so the subtree rooted at Q_ℓ is pruned. In the other subregion, $z_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma_\ell) > z(I(\ell))$, and Q_ℓ is branched on further. Therefore, within each connected component of C induced by all hyperplanes given by Equation 16 for all ℓ , the set of node within the list (15) that are pruned is invariant. Combined with the surfaces established in Lemma 4.4, these hyperplanes partition \mathbb{R}^{n+1} into connected components such that as $(\boldsymbol{\alpha}, \beta)$ varies within a given component, the tree built by branch-and-cut is invariant.

Finally, we count the total number of surfaces inducing this partition. Unlike the counting stages of the previous lemmas, we will first have to count the number of connected components induced by the surfaces established in Lemma 4.4. This is because the ordered list of nodes explored by branch-and-cut (15) can be different across each component, and the hyperplanes given by Equation 16 depend on this list. From Lemma 4.4 we have $3(m+2n)^n\tau^{4n}$ hyperplanes, $3(m+2n)^{3n}\tau^{4n}$ degree-2 polynomial hypersurfaces, and $(m+2n)^{6n}\tau^{4n}$ degree-5 polynomial hypersurfaces. To determine the connected components of \mathbb{R}^{n+1} induced by the zero sets of these polynomials, it suffices to consider the zero set of the product of all polynomials defining these surfaces. Denote this product

polynomial by p . The degree of the product polynomial is the sum of the degrees of $3(m+2n)^n\tau^{4n}$ degree-1 polynomials, $3(m+2n)^{3n}\tau^{4n}$ degree-2 polynomials, and $(m+2n)^{6n}\tau^{4n}$ degree-5 polynomials, which is at most $3(m+2n)^n\tau^{4n} + 2 \cdot 3(m+2n)^{3n}\tau^{4n} + 5 \cdot (m+2n)^{6n}\tau^{4n} < 14(m+2n)^{3n}\tau^{4n}$. By Warren's theorem, the number of connected components of $\mathbb{R}^{n+1} \setminus \{(\boldsymbol{\alpha}, \beta) : p(\boldsymbol{\alpha}, \beta) = 0\}$ is $O((14(m+2n)^{3n}\tau^{4n})^{n-1})$, and by the Milnor-Thom theorem, the number of connected components of $\{(\boldsymbol{\alpha}, \beta) : p(\boldsymbol{\alpha}, \beta) = 0\}$ is $O((14(m+2n)^{3n}\tau^{4n})^{n-1})$ as well. So, the number of connected components induced by the surfaces in Lemma 4.4 is $O(14^n(m+2n)^{3n^2}\tau^{4n^2})$. For every connected component C in Lemma 4.4, the closed form of $z_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma_\ell)$ is already determined due to Lemma 4.2, and so the number of hyperplanes given by Equation 16 is at most the number of possible $\sigma \subseteq \mathcal{BC}$, which is at most τ^{3n} . So across all connected components C , the total number of hyperplanes given by Equation 16 is $O(14^n(m+2n)^{3n^2}\tau^{5n^2})$. Finally, adding this to the surface-counts established in Lemma 4.4 yields the lemma statement. \square

C.1 Product scoring rule for variable selection

Let σ be the set of branching constraints added thus far. The product scoring rule branches on the variable $i \in [n]$ that maximizes:

$$\max\{z_{\text{LP}}^*(\sigma) - z_{\text{LP}}^*(x_i \leq \lfloor x_{\text{LP}}^*(\sigma)[i] \rfloor, \sigma), \gamma\} \cdot \max\{z_{\text{LP}}^*(\sigma) - z_{\text{LP}}^*(x_i \geq \lceil x_{\text{LP}}^*(\sigma)[i] \rceil, \sigma), \gamma\},$$

where $\gamma = 10^{-6}$.

Lemma C.1. *There is a set of at most $3(m+2n)^n\tau^{3n}$ hyperplanes and $(m+2n)^{2n}\tau^{3n}$ degree-2 polynomial hypersurfaces partitioning \mathbb{R}^{n+1} into connected components such that for any connected component C and any σ , the set of branching constraints $\{x_i \leq \lfloor \mathbf{x}_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma)[i] \rfloor, x_i \geq \lceil \mathbf{x}_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma)[i] \rceil \mid i \in [n]\}$ is invariant across all $(\boldsymbol{\alpha}, \beta) \in C$.*

Proof. Fix a connected component C in the decomposition established in Lemma 4.2. By Lemma 4.2, for each σ , either $\mathbf{x}_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma) = \mathbf{x}_{\text{LP}}^*(\sigma)$ or there exists $E \subseteq M \cup \sigma$ such that $\mathbf{x}_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma)[i] = \frac{\det(A_{E, \boldsymbol{\alpha}, \beta, \sigma}^i)}{\det(A_{E, \boldsymbol{\alpha}, \sigma})}$ for all $(\boldsymbol{\alpha}, \beta) \in C$. Fix a variable $i \in [n]$, which corresponds to two branching constraints

$$x_i \leq \lfloor \mathbf{x}_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma)[i] \rfloor \text{ and } x_i \geq \lceil \mathbf{x}_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma)[i] \rceil. \quad (17)$$

If C is a component where $\mathbf{x}_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma) = \mathbf{x}_{\text{LP}}^*(\sigma)$, then these two branching constraints are trivially invariant over $(\boldsymbol{\alpha}, \beta) \in C$. Otherwise, in order to further decompose C such that the right-hand-sides of these constraints are invariant for every σ , we add the two decision boundaries given by

$$k \leq \frac{\det(A_{E, \boldsymbol{\alpha}, \beta, \sigma}^i)}{\det(A_{E, \boldsymbol{\alpha}, \sigma})} \leq k + 1$$

for every i, σ , and every integer $k = 0, \dots, \tau - 1$, where $\tau = \max_{\mathbf{x} \in \mathcal{P} \cap \mathbb{Z}^n} \|\mathbf{x}\|_\infty$. This ensures that within every connected component of C induced by these boundaries (hyperplanes),

$$\lfloor \mathbf{x}_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma)[i] \rfloor = \left\lfloor \frac{\det(A_{E, \boldsymbol{\alpha}, \beta, \sigma}^i)}{\det(A_{E, \boldsymbol{\alpha}, \sigma})} \right\rfloor \text{ and } \lceil \mathbf{x}_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma)[i] \rceil = \left\lceil \frac{\det(A_{E, \boldsymbol{\alpha}, \beta, \sigma}^i)}{\det(A_{E, \boldsymbol{\alpha}, \sigma})} \right\rceil$$

are invariant, so the branching constraints from Equation (17) are invariant. For a fixed σ , there are two hyperplanes for every $E \subseteq M \cup \sigma$ corresponding to an edge of $\mathcal{P}(\sigma)$ and $i = 1, \dots, n$, for a total of at most $2n \binom{m+|\sigma|}{n-1} \leq 2n(m+|\sigma|)^{n-1}$ hyperplanes. Summing over all reduced σ , we get a total of $2n(m+2n)^{n-1}\tau^{3n} < 2(m+2n)^n\tau^{3n}$ hyperplanes. Adding these hyperplanes to the set of hyperplanes established in Lemma 4.2 yields the lemma statement. \square

Proof of Lemma 4.3. Fix a connected component C in the decomposition established in Lemma C.1. We know that for each set of branching constraints σ :

- By Lemma 4.2, either $\mathbf{x}_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma) = \mathbf{x}_{\text{LP}}^*(\sigma)$ or there exists $E \subseteq M \cup \sigma$ such that $\mathbf{x}_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma)[i] = \frac{\det(A_{E, \boldsymbol{\alpha}, \beta, \sigma}^i)}{\det(A_{E, \boldsymbol{\alpha}, \sigma}^i)}$ for all $(\boldsymbol{\alpha}, \beta) \in C$ and all $i \in [n]$, and
- The set of branching constraints $\{x_i \leq \lfloor \mathbf{x}_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma)[i] \rfloor, x_i \geq \lceil \mathbf{x}_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma)[i] \rceil \mid i \in [n]\}$ is invariant across all $(\boldsymbol{\alpha}, \beta) \in C$.

Suppose that σ is the list of branching constraints added so far. For any variable $k \in [n]$, let

$$\sigma_k^- = (x_k \leq \lfloor \mathbf{x}_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma)[k] \rfloor, \sigma) \text{ and } \sigma_k^+ = (x_k \geq \lceil \mathbf{x}_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma)[k] \rceil, \sigma).$$

So long as $(\boldsymbol{\alpha}, \beta) \in C$, σ_k^- and σ_k^+ are fixed. With this notation, we can write the product scoring rule as

$$\max\{z_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma) - z_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma_k^-), \gamma\} \cdot \max\{z_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma) - z_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma_k^+), \gamma\},$$

where $\gamma = 10^{-6}$.

By Lemma 4.2, we know that across all $(\boldsymbol{\alpha}, \beta) \in C$, either $z_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma_k^+) = z_{\text{LP}}^*(\sigma_k^+)$ or there exists $E_k^+ \subseteq M \cup \sigma_k^+$ such that

$$z_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma_k^+) = \sum_{i=1}^n c_i \cdot \frac{\det(A_{E_k^+, \boldsymbol{\alpha}, \beta, \sigma_k^+}^i)}{\det(A_{E_k^+, \boldsymbol{\alpha}, \sigma_k^+}^i)},$$

and similarly for σ_k^- , defined according to some edge set $E_k^- \subseteq M \cup \sigma_k^-$. Therefore, for each $k \in [n]$, there is a single degree-2 polynomial hypersurface partitioning C into connected components such that within each connected component, either

$$z_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma) - z_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma_k^-) \geq \gamma \tag{18}$$

or vice versa, and similarly for σ_k^+ . In particular, the former hypersurface will have one of four forms:

1. $z_{\text{LP}}^*(\sigma) - z_{\text{LP}}^*(\sigma_k^-) \geq \gamma$, which is uniformly satisfied or not satisfied across all $(\boldsymbol{\alpha}, \beta) \in C$,
2. $z_{\text{LP}}^*(\sigma) - \sum_{i=1}^n c_i \cdot \frac{\det(A_{E_k^-, \boldsymbol{\alpha}, \beta, \sigma_k^-}^i)}{\det(A_{E_k^-, \boldsymbol{\alpha}, \sigma_k^-}^i)} \geq \gamma$, which is a hyperplane,
3. $\sum_{i=1}^n c_i \cdot \frac{\det(A_{E, \boldsymbol{\alpha}, \beta, \sigma}^i)}{\det(A_{E, \boldsymbol{\alpha}, \sigma}^i)} - z_{\text{LP}}^*(\sigma_k^-) \geq \gamma$, which is a hyperplane, or
4. $\sum_{i=1}^n c_i \left(\frac{\det(A_{E, \boldsymbol{\alpha}, \beta, \sigma}^i)}{\det(A_{E, \boldsymbol{\alpha}, \sigma}^i)} - \frac{\det(A_{E_k^-, \boldsymbol{\alpha}, \beta, \sigma_k^-}^i)}{\det(A_{E_k^-, \boldsymbol{\alpha}, \sigma_k^-}^i)} \right) \geq \gamma$, which is a degree-2 polynomial hypersurface.

Simply said, these are all degree-2 polynomial hypersurfaces.

Within any region induced by these hypersurfaces, the comparison between any two variables x_k and x_j will have the form

$$\max\{z_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma) - z_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma_k^-), \gamma\} \cdot \max\{z_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma) - z_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma_k^+), \gamma\}$$

$$\geq \max\{z_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma) - z_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma_j^-), \gamma\} \cdot \max\{z_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma) - z_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma_j^+), \gamma\}$$

which at its most complex will equal

$$\begin{aligned} & \sum_{i=1}^n c_i \left(\frac{\det(A_{E, \boldsymbol{\alpha}, \beta, \sigma}^i)}{\det(A_{E, \boldsymbol{\alpha}, \sigma})} - \frac{\det(A_{E_k^-, \boldsymbol{\alpha}, \beta, \sigma_k^-}^i)}{\det(A_{E_k^-, \boldsymbol{\alpha}, \sigma_k^-})} \right) \cdot \sum_{i=1}^n c_i \left(\frac{\det(A_{E, \boldsymbol{\alpha}, \beta, \sigma}^i)}{\det(A_{E, \boldsymbol{\alpha}, \sigma})} - \frac{\det(A_{E_k^+, \boldsymbol{\alpha}, \beta, \sigma_k^+}^i)}{\det(A_{E_k^+, \boldsymbol{\alpha}, \sigma_k^+})} \right) \\ & \geq \sum_{i=1}^n c_i \left(\frac{\det(A_{E, \boldsymbol{\alpha}, \beta, \sigma}^i)}{\det(A_{E, \boldsymbol{\alpha}, \sigma})} - \frac{\det(A_{E_j^-, \boldsymbol{\alpha}, \beta, \sigma_j^-}^i)}{\det(A_{E_j^-, \boldsymbol{\alpha}, \sigma_j^-})} \right) \cdot \sum_{i=1}^n c_i \left(\frac{\det(A_{E, \boldsymbol{\alpha}, \beta, \sigma}^i)}{\det(A_{E, \boldsymbol{\alpha}, \sigma})} - \frac{\det(A_{E_j^+, \boldsymbol{\alpha}, \beta, \sigma_j^+}^i)}{\det(A_{E_j^+, \boldsymbol{\alpha}, \sigma_j^+})} \right). \end{aligned} \quad (19)$$

This inequality can be written as a degree-5 polynomial hypersurface. In any region induced by these hypersurfaces, the variable that branch-and-cut branches on will be fixed.

We now count the total number of hypersurfaces. First, we count the number of degree-2 polynomial hypersurfaces from Equation (18): there is a hypersurface defined by each variable x_k , set of branching constraints σ , cutoff $t \in [\tau]$ such that $\sigma_k^- = (x_k \leq t, \sigma)$, set $E \subseteq M \cup \sigma$ corresponding to an edge of $\mathcal{P}(\sigma)$, and set $E_k^- \subseteq M \cup \sigma_k^-$ (and similarly for σ_k^+ and E_k^+). For a fixed σ , this amounts to $2n\tau \binom{m+|\sigma|}{n-1} \binom{m+|\sigma|+1}{n-1} \leq 2n\tau(m+|\sigma|+1)^{2(n-1)}$ hypersurfaces. Summing over all τ^{3n} reduced σ , we have $2n\tau^{3n+1}(m+2n+1)^{2(n-1)}$ degree-2 polynomial hypersurfaces.

Next, we count the number of degree-5 polynomial hypersurfaces from Equation (19): there is a hypersurface defined by each pair of variables x_k, x_j , set of branching constraints σ , cutoffs $t_k, t_j \in [\tau]$ such that $\sigma_k^- = (x_k \leq t_k, \sigma)$ and $\sigma_j^- = (x_j \leq t_j, \sigma)$, and sets $E, E_k^-, E_k^+, E_j^-, E_j^+$ corresponding to edges of $\mathcal{P}(\sigma), \mathcal{P}(\sigma_k^-), \mathcal{P}(\sigma_k^+), \mathcal{P}(\sigma_j^-), \mathcal{P}(\sigma_j^+)$. For a fixed σ , this amounts to $n^2\tau^2 \binom{m+|\sigma|}{n-1} \binom{m+|\sigma|+1}{n-1}^4 \leq n^2\tau^2(m+|\sigma|+1)^{5(n-1)}$ hypersurfaces. Summing over all τ^{3n} reduced σ , we have $n^2\tau^{3n+2}(m+2n+1)^{5(n-1)}$ degree-5 polynomial hypersurfaces.

Adding these hypersurfaces to those from Lemma C.1, we get the lemma statement. \square

C.2 Extension to multiple cutting planes

We can similarly derive a multi-cut version of Lemma 4.2 that controls $\mathbf{x}_{\text{LP}}^*(\boldsymbol{\alpha}_1^T \mathbf{x} \leq \beta_1, \dots, \boldsymbol{\alpha}_K^T \mathbf{x} \leq \beta_K, \sigma)$ for any set of branching constraints. We use the following notation. Let $(\mathbf{c}, A, \mathbf{b})$ be an LP and let M denote the set of its m constraints. For $F \subseteq M \cup \sigma$, let $A_{F, \sigma} \in \mathbb{R}^{|F| \times n}$ and $\mathbf{b}_{F, \sigma} \in \mathbb{R}^{|F|}$ denote the restrictions of A_σ and \mathbf{b}_σ to F . For $\boldsymbol{\alpha}_1, \dots, \boldsymbol{\alpha}_k \in \mathbb{R}^n$, $\beta_1, \dots, \beta_k \in \mathbb{R}$, and $F \subseteq M \cup \sigma$ with $|F| = n - k$, let $A_{F, \boldsymbol{\alpha}_1, \dots, \boldsymbol{\alpha}_k, \sigma} \in \mathbb{R}^{n \times n}$ denote the matrix obtained by adding row vectors $\boldsymbol{\alpha}_1, \dots, \boldsymbol{\alpha}_k$ to $A_{F, \sigma}$ and let $A_{F, \boldsymbol{\alpha}_1, \dots, \boldsymbol{\alpha}_k, \beta_1, \dots, \beta_k, \sigma}^i \in \mathbb{R}^{n \times n}$ be the matrix $A_{F, \boldsymbol{\alpha}_1, \dots, \boldsymbol{\alpha}_k, \sigma} \in \mathbb{R}^{n \times n}$ with the i th column replaced by $[\mathbf{b}_{F, \sigma} \ \beta_1 \ \dots \ \beta_k]^T$.

Corollary C.2. *Fix an IP $(\mathbf{c}, A, \mathbf{b})$. There is a set of at most K hyperplanes, $nK^n(m+2n)^n\tau^{3n}$ degree- K polynomial hypersurfaces, and $nK^n(m+2n)^{2n}\tau^{3n}$ degree- $2K$ polynomial hypersurfaces partitioning $\mathbb{R}^{K(n+1)}$ into connected components such that for each component C and every $\sigma \subseteq \mathcal{BC}$, one of the following holds: either (1) $\mathbf{x}_{\text{LP}}^*(\boldsymbol{\alpha}_1^T \mathbf{x} \leq \beta_1, \dots, \boldsymbol{\alpha}_K^T \mathbf{x} \leq \beta_K, \sigma) = \mathbf{x}_{\text{LP}}^*(\sigma)$, or (2) there is a subset of cuts indexed by $\ell_1, \dots, \ell_k \in [K]$ and a set of constraints $F \subseteq M \cup \sigma$ with $|F| = n - k$ such that*

$$\mathbf{x}_{\text{LP}}^*(\boldsymbol{\alpha}_1^T \mathbf{x} \leq \beta_1, \dots, \boldsymbol{\alpha}_K^T \mathbf{x} \leq \beta_K, \sigma) = \left(\frac{\det(A_{F, \boldsymbol{\alpha}_{\ell_1}, \beta_{\ell_1}, \dots, \boldsymbol{\alpha}_{\ell_k}, \beta_{\ell_k}, \sigma}^1)}{\det(A_{F, \boldsymbol{\alpha}_{\ell_1}, \dots, \boldsymbol{\alpha}_{\ell_k}, \sigma})}, \dots, \frac{\det(A_{F, \boldsymbol{\alpha}_{\ell_1}, \beta_{\ell_1}, \dots, \boldsymbol{\alpha}_{\ell_k}, \beta_{\ell_k}, \sigma}^n)}{\det(A_{F, \boldsymbol{\alpha}_{\ell_1}, \dots, \boldsymbol{\alpha}_{\ell_k}, \sigma})} \right),$$

for all $(\boldsymbol{\alpha}_1, \beta_1, \dots, \boldsymbol{\alpha}_K, \beta_K) \in C$.

Proof. The exact same reasoning in the proof of Lemma B.1 applies. We still have K hyperplanes. Now, for each σ , for each subset $S \subseteq K$ with $|S| \leq n$, and for every face F of $\mathcal{P}(\sigma)$ with $\dim(F) = |S|$, we have at most m degree- K polynomial hypersurfaces. The number of k -dimensional faces of $\mathcal{P}(\sigma)$ is at most $\binom{m+|\sigma|}{n-k} \leq (m+2n)^{n-1}$, so the total number of these hypersurfaces is at most $nK^n(m+2n)^n\tau^{3n}$. Finally, for every σ , we considered a degree- $2K$ polynomial hypersurfaces for every subset of cuts and every pair of faces with degree equal to the size of the subset, of which there are at most $nK^n(m+2n)^{2n}\tau^{3n}$, as desired. \square

We now refine the decomposition obtained in Lemma 4.2 so that the branching constraints added at each step of branch-and-cut are invariant within a region. For ease of exposition, we assume that branch-and-cut uses a lexicographic variable selection policy. This means that the variable branched on at each node of the search tree is fixed and given by the lexicographic ordering x_1, \dots, x_n . Generalizing the argument to work for other policies, such as the product scoring rule, can be done as in the single-cut case.

Lemma C.3. *Suppose branch-and-cut uses a lexicographic variable selection policy. Then, there is a set of at most K hyperplanes, $3n^2K^n(m+2n)^n\tau^{3n}$ degree- K polynomial hypersurfaces, and $nK^n(m+2n)^{2n}\tau^{3n}$ degree- $2K$ polynomial hypersurfaces partitioning \mathbb{R}^{n+1} into connected components such that within each connected component, the branching constraints used at every step of branch-and-cut are invariant.*

Proof. Fix a connected component C in the decomposition established in Corollary C.2. Then, by Corollary C.2, for each σ , either $\mathbf{x}_{LP}^*(\boldsymbol{\alpha}_1^T \mathbf{x} \leq \beta_1, \dots, \boldsymbol{\alpha}_K^T \mathbf{x} \leq \beta_K, \sigma) = \mathbf{x}_{LP}^*(\sigma)$ or there exists cuts (without loss of generality) labeled by indices $1, \dots, k \in [K]$ and there exists $F \subseteq M \cup \sigma$ such that

$$\mathbf{x}_{LP}^*(\boldsymbol{\alpha}_1^T \mathbf{x} \leq \beta_1, \dots, \boldsymbol{\alpha}_K^T \mathbf{x} \leq \beta_K, \sigma)[i] = \frac{\det(A_{F, \boldsymbol{\alpha}_1, \dots, \boldsymbol{\alpha}_k, \beta_1, \dots, \beta_K, \sigma}^i)}{\det(A_{F, \boldsymbol{\alpha}_1, \dots, \boldsymbol{\alpha}_k, \sigma})}$$

for all $(\boldsymbol{\alpha}, \beta) \in C$ and all $i \in [n]$. Now, if we are at a stage in the branch-and-cut tree where σ is the list of branching constraints added so far, and the i th variable is being branched on next, the two constraints generated are

$$x_i \leq \lfloor \mathbf{x}_{LP}^*(\boldsymbol{\alpha}_1^T \mathbf{x} \leq \beta_1, \dots, \boldsymbol{\alpha}_K^T \mathbf{x} \leq \beta_K, \sigma)[i] \rfloor \text{ and } x_i \geq \lceil \mathbf{x}_{LP}^*(\boldsymbol{\alpha}_1^T \mathbf{x} \leq \beta_1, \dots, \boldsymbol{\alpha}_K^T \mathbf{x} \leq \beta_K, \sigma)[i] \rceil,$$

respectively. If C is a component where $\mathbf{x}_{LP}^*(\boldsymbol{\alpha}_1^T \mathbf{x} \leq \beta_1, \dots, \boldsymbol{\alpha}_K^T \mathbf{x} \leq \beta_K, \sigma) = \mathbf{x}_{LP}^*(\sigma)$, then there is nothing more to do, since the branching constraints at that point are trivially invariant over $(\boldsymbol{\alpha}_1, \beta_1, \dots, \boldsymbol{\alpha}_K, \beta_K) \in C$. Otherwise, in order to further decompose C such that the right-hand-side of these constraints are invariant for every σ and every $i = 1, \dots, n$, we add the two decision boundaries given by

$$k \leq \frac{\det(A_{F, \boldsymbol{\alpha}_1, \beta_1, \dots, \boldsymbol{\alpha}_k, \beta_k, \sigma}^i)}{\det(A_{F, \boldsymbol{\alpha}_1, \dots, \boldsymbol{\alpha}_k, \sigma})} \leq k+1$$

for every i , σ , and every integer $k = 0, \dots, \tau - 1$, where $\tau = \lceil \max_{\mathbf{x} \in \mathcal{P}} \|\mathbf{x}\|_\infty \rceil$. This ensures that within every connected component of C induced by these boundaries (degree- K polynomial hypersurfaces),

$$\lfloor \mathbf{x}_{LP}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma)[i] \rfloor = \left\lfloor \frac{\det(A_{F, \boldsymbol{\alpha}_1, \beta_1, \dots, \boldsymbol{\alpha}_k, \beta_k, \sigma}^i)}{\det(A_{F, \boldsymbol{\alpha}_1, \dots, \boldsymbol{\alpha}_k, \sigma})} \right\rfloor$$

and

$$\lceil \mathbf{x}_{\text{LP}}^*(\boldsymbol{\alpha}^T \mathbf{x} \leq \beta, \sigma)[i] \rceil = \left\lceil \frac{\det(A_{F, \boldsymbol{\alpha}_1, \beta_1, \dots, \boldsymbol{\alpha}_k, \beta_k, \sigma}^i)}{\det(A_{F, \boldsymbol{\alpha}_1, \dots, \boldsymbol{\alpha}_k, \sigma})} \right\rceil$$

are invariant, so the branching constraints added by, for example, a lexicographic branching rule, are invariant. For a fixed σ , there are two hypersurfaces for every subset $S \subseteq [K]$, every $F \subseteq M \cup \sigma$ corresponding to a $|S|$ -dimensional face of $\mathcal{P}(\sigma)$, and every $i = 1, \dots, n$, for a total of at most $2n^2 K^n \binom{m+|\sigma|}{|S|} \leq 2n^2 K^n (m+2n)^n$. Summing over all reduced σ , we get a total of $2n^2 K^n (m+2n)^n \tau^{3n}$ hypersurfaces. Adding these hypersurfaces to the set of hypersurfaces established in Corollary C.2 yields the lemma statement. \square

Now, as in the single-cut case, we consider the constraints that ensure that all cuts are valid. Let $\mathcal{V} \subseteq \mathbb{R}^{K(n+1)}$ denote the set of all vectors of valid K cuts. As before, \mathcal{V} is a polyhedron, since we may write

$$\mathcal{V} = \bigcap_{k=1}^K \bigcap_{\mathbf{x}_{\text{IH}} \in \mathcal{P}_{\text{IH}}} \left\{ (\boldsymbol{\alpha}_1, \beta_1, \dots, \boldsymbol{\alpha}_K, \beta_K) \in \mathbb{R}^{K(n+1)} : \boldsymbol{\alpha}_k^T \mathbf{x}_{\text{IH}} \leq \beta_k \right\}.$$

We now refine our decomposition further to control the integrality of the various LP solutions at each node of branch-and-cut.

Lemma C.4. *Given an IP $(\mathbf{c}, A, \mathbf{b})$, there is a set of at most $2K\tau^n$ hyperplanes, $4n^2 K^n (m+2n)^n \tau^{4n}$ degree- K polynomial hypersurfaces, and $nK^n (m+2n)^{2n} \tau^{3n}$ degree- $2K$ polynomial hypersurfaces partitioning $\mathbb{R}^{K(n+1)}$ into connected components such that for each component C , and each $\sigma \subseteq \mathcal{BC}$,*

$$\mathbf{1} [\mathbf{x}_{\text{LP}}^* (\boldsymbol{\alpha}_1^T \mathbf{x} \leq \beta_1, \dots, \boldsymbol{\alpha}_K^T \mathbf{x} \leq \beta_K, \sigma) \in \mathbb{Z}^n]$$

is invariant for all $(\boldsymbol{\alpha}_1, \beta_1, \dots, \boldsymbol{\alpha}_K, \beta_K) \in C$.

Proof. Fix a connected component C in the decomposition that includes the facets defining \mathcal{V} and the surfaces obtained in Lemma C.3. For all $\sigma \in \mathcal{BC}$, $\mathbf{x}_{\text{I}} \in \mathcal{P}_{\text{I}}$, and $i = 1, \dots, n$, consider the surface

$$\mathbf{x}_{\text{LP}}^* (\boldsymbol{\alpha}_1^T \mathbf{x} \leq \beta_1, \dots, \boldsymbol{\alpha}_K^T \mathbf{x} \leq \beta_K, \sigma) [i] = \mathbf{x}_{\text{I}}[i]. \quad (20)$$

This surface is a polynomial hypersurface of degree at most K , due to Corollary C.2. Clearly, within any connected component of C induced by these hyperplanes, for every σ and $\mathbf{x}_{\text{I}} \in \mathcal{P}_{\text{I}}$, $\mathbf{1} [\mathbf{x}_{\text{LP}}^* (\boldsymbol{\alpha}_1^T \mathbf{x} \leq \beta_1, \dots, \boldsymbol{\alpha}_K^T \mathbf{x} \leq \beta_K, \sigma) = \mathbf{x}_{\text{I}}] = 1$ is invariant. Finally, if $\mathbf{x}_{\text{LP}}^* (\boldsymbol{\alpha}_1^T \mathbf{x} \leq \beta_1, \dots, \boldsymbol{\alpha}_K^T \mathbf{x} \leq \beta_K, \sigma) \in \mathbb{Z}^n$ for some K cuts $\boldsymbol{\alpha}_1^T \mathbf{x} \leq \beta_1, \dots, \boldsymbol{\alpha}_K^T \mathbf{x} \leq \beta_K$ within a given connected component, $\mathbf{x}_{\text{LP}}^* (\boldsymbol{\alpha}_1^T \mathbf{x} \leq \beta_1, \dots, \boldsymbol{\alpha}_K^T \mathbf{x} \leq \beta_K, \sigma) = \mathbf{x}_{\text{I}}$ for some $\mathbf{x}_{\text{I}} \in \mathcal{P}_{\text{IH}}(\sigma) \subseteq \mathcal{P}_{\text{I}}$, which means that $\mathbf{x}_{\text{LP}}^* (\boldsymbol{\alpha}_1^T \mathbf{x} \leq \beta_1, \dots, \boldsymbol{\alpha}_K^T \mathbf{x} \leq \beta_K, \sigma) = \mathbf{x}_{\text{I}} \in \mathbb{Z}^n$ for all vectors of K cuts $\boldsymbol{\alpha}_1^T \mathbf{x} \leq \beta_1, \dots, \boldsymbol{\alpha}_K^T \mathbf{x} \leq \beta_K$ in that connected component.

We now count the number of hyperplanes given by Equation 20. For each σ , there are nK^n possible subsets of cut indices and at most $(m+2n)^{n-1}$ binding face constraints $F \subseteq M \cup \sigma$ defining the formula of Corollary C.2. For each subset-face pair, there are $n|\mathcal{P}_{\text{I}}| \leq n\tau^n$ degree- K polynomial hypersurfaces given by Equation 20. So the total number of such hypersurfaces over all σ is at most $\tau^{3n} n^2 K^n (m+2n)^{n-1} \tau^n$. The number of facets defining \mathcal{V} is at most $K|\mathcal{P}_{\text{I}}| \leq K\tau^n$. Adding these to the counts obtained in Lemma C.3 yields the final tallies in the lemma statement. \square

At this point, as in the single-cut case, if the bounding aspect of branch-and-cut is suppressed, our decomposition yields connected components over which the branch-and-cut tree built is invariant. We now prove our main structural theorem for B&C as a function of multiple cutting planes at the root.

Theorem C.5. *Given an IP $(\mathbf{c}, A, \mathbf{b})$, there is a set of at most $O(12^n n^{2n} K^{2n^2} (m+2n)^{2n^2} \tau^{5n^2})$ polynomial hypersurfaces of degree at most $2K$ partitioning $\mathbb{R}^{K(n+1)}$ into connected components such that the branch-and-cut tree built after adding the K cuts $\alpha_1^T \mathbf{x} \leq \beta_1, \dots, \alpha_k^T \mathbf{x} \leq \beta_k$ at the root is invariant over all $(\alpha_1, \beta_1, \dots, \alpha_K, \beta_K)$ within a given component. In particular, $f_{\mathbf{c}, A, \mathbf{b}}(\alpha_1, \beta_1, \dots, \alpha_K, \beta_K)$ is invariant over each connected component.*

Proof. Fix a connected component C in the decomposition induced by the set of hyperplanes, degree- K hypersurfaces, and degree- $2K$ hypersurfaces established in Lemma C.4. Let

$$Q_1, \dots, Q_{i_1}, I_1, Q_{i_1+1}, \dots, Q_{i_2}, I_2, Q_{i_2+1}, \dots \quad (21)$$

denote the nodes of the tree branch-and-cut creates, in order of exploration, under the assumption that a node is pruned if and only if either the LP at that node is infeasible or the LP optimal solution is integral (so the “bounding” of branch-and-bound is suppressed). Here, a node is identified by the list σ of branching constraints added to the input IP. Nodes labeled by Q are either infeasible or have fractional LP optimal solutions. Nodes labeled by I have integral LP optimal solutions and are candidates for the incumbent integral solution at the point they are encountered. (The nodes are functions of $\alpha_1, \beta_1, \dots, \alpha_K, \beta_K$, as are the indices i_1, i_2, \dots) By Lemma C.4, this ordered list of nodes is invariant for all $(\alpha_1, \beta_1, \dots, \alpha_K, \beta_K) \in C$.

Now, given an node index ℓ , let $I(\ell)$ denote the incumbent node with the highest objective value encountered up until the ℓ th node searched by B&C, and let $z(I(\ell))$ denote its objective value. For each node Q_ℓ , let σ_ℓ denote the branching constraints added to arrive at node Q_ℓ . The hyperplane

$$z_{\text{LP}}^*(\alpha_1^T \mathbf{x} \leq \beta_1, \dots, \alpha_K^T \mathbf{x} \leq \beta_K, \sigma_\ell) = z(I(\ell)) \quad (22)$$

(which is a hyperplane due to Corollary C.2) partitions C into two subregions. In one subregion, $z_{\text{LP}}^*(\alpha_1^T \mathbf{x} \leq \beta_1, \dots, \alpha_K^T \mathbf{x} \leq \beta_K, \sigma_\ell) \leq z(I(\ell))$, that is, the objective value of the LP optimal solution is no greater than the objective value of the current incumbent integer solution, and so the subtree rooted at Q_ℓ is pruned. In the other subregion, $z_{\text{LP}}^*(\alpha_1^T \mathbf{x} \leq \beta_1, \dots, \alpha_K^T \mathbf{x} \leq \beta_K, \sigma_\ell) > z(I(\ell))$, and Q_ℓ is branched on further. Therefore, within each connected component of C induced by all hyperplanes given by Equation 22 for all ℓ , the set of node within the list (21) that are pruned is invariant. Combined with the surfaces established in Lemma C.4, these hyperplanes partition $\mathbb{R}^{K(n+1)}$ into connected components such that as $(\alpha_1, \beta_1, \dots, \alpha_K, \beta_K)$ varies within a given component, the tree built by branch-and-cut is invariant.

Finally, we count the total number of surfaces inducing this partition. Unlike the counting stages of the previous lemmas, we will first have to count the number of connected components induced by the surfaces established in Lemma C.4. This is because the ordered list of nodes explored by branch-and-cut (21) can be different across each component, and the hyperplanes given by Equation 22 depend on this list. From Lemma C.4 we have $6n^2 K^n (m+2n)^{2n} \tau^{4n}$ polynomial hypersurfaces of degree $\leq 2K$. The set of all $(\alpha_1, \beta_1, \dots, \alpha_K, \beta_K) \in \mathbb{R}^{K(n+1)}$ such that $(\alpha_1, \beta_1, \dots, \alpha_K, \beta_K)$ lies on the boundary of any of these surfaces is precisely the zero set of the product of all polynomials defining these surfaces. Denote this product polynomial by p . The degree of the product polynomial is the sum of the degrees of $6n^2 K^n (m+2n)^{2n} \tau^{4n}$ polynomials of degree $\leq 2K$, which is at most $2K \cdot 6Kn^2 K^n (m+2n)^{2n} \tau^{4n} = 12n^2 K^{n+2} (m+2n)^{2n} \tau^{4n}$. By Warren’s theorem, the number of connected components of $\mathbb{R}^{n+1} \setminus \{(\alpha, \beta) : p(\alpha, \beta) = 0\}$ is $O((12n^2 K^{n+2} (m+2n)^{2n} \tau^{4n})^{n-1})$, and by the Milnor-Thom theorem, the number of connected components of $\{(\alpha, \beta) : p(\alpha, \beta) = 0\}$ is $O((12n^2 K^{n+2} (m+2n)^{2n} \tau^{4n})^{n-1})$ as well. So, the number of connected components induced by the surfaces in Lemma C.4 is $O(12^n n^{2n} K^{2n^2} (m+2n)^{2n^2} \tau^{4n^2})$. For every connected component C in Lemma C.4, the closed form of $z_{\text{LP}}^*(\alpha^T \mathbf{x} \leq \beta, \sigma_\ell)$ is already determined due to Corollary C.2,

and so the number of hyperplanes given by Equation 22 is at most the number of possible $\sigma \subseteq \mathcal{BC}$, which is at most τ^{3n} . So across all connected components C , the total number of hyperplanes given by Equation 22 is $O(12^n n^{2n} K^{2n^2} (m+2n)^{2n^2} \tau^{5n^2})$. Finally, adding this to the surface-counts established in Lemma C.4 yields the theorem statement. \square

D Omitted results from Section 5

Proof of Theorem 5.1. For a set \mathcal{X} , $\mathcal{X}^{<\mathbb{N}}$ denotes the set of finite sequences of elements from \mathcal{X} . There is a bijection between the set of IPs $(\mathbf{c}, A, \mathbf{b}) \in \mathcal{I} := \mathbb{R}^n \times \mathbb{Z}^{m \times n} \times \mathbb{Z}^m$ and \mathbb{R} , so IPs can be uniquely represented as real numbers (and vice versa). Now, consider the set of all finite sequences of pairs of IPs and ± 1 labels of the form $((\mathbf{c}_1, A_1, \mathbf{b}_1), \varepsilon_1), \dots, ((\mathbf{c}_N, A_N, \mathbf{b}_N), \varepsilon_N)$, $\varepsilon_1, \dots, \varepsilon_N \in \{-1, 1\}$, that is, the set $(\mathcal{I} \times \{-1, 1\})^{<\mathbb{N}}$. There is a bijection between this set and $(\mathbb{R} \times \{-1, 1\})^{<\mathbb{N}}$, and in turn there is a bijection between $(\mathbb{R} \times \{-1, 1\})^{<\mathbb{N}}$ and \mathbb{R} . Hence, there exists a bijection between \mathcal{U} and $(\mathcal{I} \times \{-1, 1\})^{<\mathbb{N}}$. Fix such a bijection $\varphi : \mathcal{U} \rightarrow (\mathcal{I} \times \{-1, 1\})^{<\mathbb{N}}$, and let $\varphi^{-1} : (\mathcal{I} \times \{-1, 1\})^{<\mathbb{N}} \rightarrow \mathcal{U}$ denote the inverse of φ , which is well defined and also a bijection.

Let n be odd. For $c \in \mathbb{R}$, let $\text{IP}_c \in \mathcal{I}$ denote the IP

$$\begin{aligned} & \text{maximize} && c \\ & \text{subject to} && 2x_1 + \dots + 2x_n = n \\ & && \mathbf{x} \in \{0, 1\}^n. \end{aligned} \tag{23}$$

Since n is odd, IP_c is infeasible, independent of c . Jeroslow [21] showed that without the use of cutting planes or heuristics, branch-and-bound builds a tree of size $2^{(n-1)/2}$ before determining infeasibility and terminating. The objective c is irrelevant, but is important in generating distinct IPs with this property. Consider the cut $x_1 + \dots + x_n \leq \lfloor n/2 \rfloor$, which is a valid cut for IP_c (this is in fact a Chvátal-Gomory cut [8]). In particular, since n is odd, $x_1 + \dots + x_n \leq \lfloor n/2 \rfloor \implies x_1 + \dots + x_n \leq (n-1)/2 < n/2$, so the equality constraint of IP_c is violated by this cut. Thus, the feasible region of the LP relaxation after adding this cut is empty, and branch-and-bound will terminate immediately at the root (building a tree of size 1). Denote this cut by $(\boldsymbol{\alpha}^{(-1)}, \beta^{(-1)}) = (\mathbf{1}, \lfloor n/2 \rfloor)$. On the other hand, let $(\boldsymbol{\alpha}^{(1)}, \beta^{(1)}) = (\mathbf{0}, 0)$ be the trivial cut $0 \leq 0$. Adding this cut to the IP constraints does not change the feasible region, so branch-and-bound will build a tree of size $2^{(n-1)/2}$.

We now define $\boldsymbol{\alpha}_{\mathbf{c}, A, \mathbf{b}}$ and $\beta_{\mathbf{c}, A, \mathbf{b}}$. Let

$$(\boldsymbol{\alpha}_{\mathbf{c}, A, \mathbf{b}}(\mathbf{u}), \beta_{\mathbf{c}, A, \mathbf{b}}(\mathbf{u})) = \begin{cases} (\boldsymbol{\alpha}^{(1)}, \beta^{(1)}) & \text{if } ((\mathbf{c}, A, \mathbf{b}), 1) \in \varphi(\mathbf{u}) \text{ and } ((\mathbf{c}, A, \mathbf{b}), -1) \notin \varphi(\mathbf{u}) \\ (\boldsymbol{\alpha}^{(-1)}, \beta^{(-1)}) & \text{if } ((\mathbf{c}, A, \mathbf{b}), -1) \in \varphi(\mathbf{u}) \text{ and } ((\mathbf{c}, A, \mathbf{b}), 1) \notin \varphi(\mathbf{u}) \\ (\mathbf{0}, 0) & \text{otherwise} \end{cases}.$$

The choice to use $(\mathbf{0}, 0)$ in the case that either $((\mathbf{c}, A, \mathbf{b}), \varepsilon) \notin \varphi(\mathbf{u})$ for each $\varepsilon \in \{-1, 1\}$, or $((\mathbf{c}, A, \mathbf{b}), -1) \in \varphi(\mathbf{u})$ and $((\mathbf{c}, A, \mathbf{b}), 1) \in \varphi(\mathbf{u})$ is arbitrary and unimportant. Now, for any integer $N > 0$, constructing a set of N IPs and N thresholds that is shattered is almost immediate. Let $c_1, \dots, c_N \in \mathbb{R}$ be distinct reals, and let $1 < r_1, \dots, r_N < 2^{(n-1)/2}$. Then, the set $\{(\text{IP}_{c_1}, r_1), \dots, (\text{IP}_{c_N}, r_N)\}$ can be shattered. Indeed, given a sign pattern $(\varepsilon_1, \dots, \varepsilon_N) \in \{-1, 1\}^N$, let

$$\mathbf{u} = \varphi^{-1}((\text{IP}_{c_1}, \varepsilon_1), \dots, (\text{IP}_{c_N}, \varepsilon_N)).$$

Then, if $\varepsilon_i = 1$, $(\boldsymbol{\alpha}_{\text{IP}_{c_i}}(\mathbf{u}), \beta_{\text{IP}_{c_i}}(\mathbf{u})) = (\boldsymbol{\alpha}^{(1)}, \beta^{(1)})$, so $g_{\mathbf{u}}(\text{IP}_{c_i}) = 2^{(n-1)/2}$ and $\text{sign}(g_{\mathbf{u}}(\text{IP}_{c_i}) - r_i) = 1$. If $\varepsilon_i = -1$, $(\boldsymbol{\alpha}_{\text{IP}_{c_i}}(\mathbf{u}), \beta_{\text{IP}_{c_i}}(\mathbf{u})) = (\boldsymbol{\alpha}^{(-1)}, \beta^{(-1)})$, so $g_{\mathbf{u}}(\text{IP}_{c_i}) = 1$ and $\text{sign}(g_{\mathbf{u}}(\text{IP}_{c_i}) - r_i) = -1$. So

for any N there is a set of IPs and thresholds that can be shattered, which yields the theorem statement. \square

Proof of Lemma 5.2. We have $f_i = \mathbf{u}^T \mathbf{a}_i - \lfloor \mathbf{u}^T \mathbf{a}_i \rfloor$, $f_0 = \mathbf{u}^T \mathbf{b} - \lfloor \mathbf{u}^T \mathbf{b} \rfloor$, and since $\mathbf{u} \in [-U, U]^m$, $\lfloor \mathbf{u}^T \mathbf{a}_i \rfloor \in [-U \|\mathbf{a}_i\|_1, U \|\mathbf{a}_i\|_1]$ and $\lfloor \mathbf{u}^T \mathbf{b} \rfloor \in [-U \|\mathbf{b}\|_1, U \|\mathbf{b}\|_1]$. Now, for all i , $k_i \in [-U \|\mathbf{a}_i\|_1, U \|\mathbf{a}_i\|_1] \cap \mathbb{Z}$ and $k_0 \in [-U \|\mathbf{b}\|_1, U \|\mathbf{b}\|_1] \cap \mathbb{Z}$, put down the hyperplanes defining the two halfspaces

$$\lfloor \mathbf{u}^T \mathbf{a}_i \rfloor = k_i \iff k_i \leq \mathbf{u}^T \mathbf{a}_i < k_i + 1 \quad (24)$$

and the hyperplanes defining the two halfspaces

$$\lfloor \mathbf{u}^T \mathbf{b} \rfloor = k_0 \iff k_0 \leq \mathbf{u}^T \mathbf{b} < k_0 + 1. \quad (25)$$

In addition, consider the hyperplane

$$\mathbf{u}^T \mathbf{a}_i - k_i = \mathbf{u}^T \mathbf{b} - k_0 \quad (26)$$

for each i . Within any connected component of \mathbb{R}^m determined by these hyperplanes, $\lfloor \mathbf{u}^T \mathbf{a}_i \rfloor$ and $\lfloor \mathbf{u}^T \mathbf{b} \rfloor$ are constant. Furthermore, $\mathbf{1}[f_i \leq f_0]$ is invariant within each connected component, since if $\lfloor \mathbf{u}^T \mathbf{a}_i \rfloor = k_i$ and $\lfloor \mathbf{u}^T \mathbf{b} \rfloor = k_0$, $f_i \leq f_0 \iff \mathbf{u}^T \mathbf{a}_i - k_i \leq \mathbf{u}^T \mathbf{b} - k_0$, which is the hyperplane given by Equation 26. The total number of hyperplanes of type 24 is $O(nU \|\mathbf{A}\|_1)$, the total number of hyperplanes of type 25 is $O(U \|\mathbf{b}\|_1)$, and the total number of hyperplanes of type 26 is $nU^2 \|\mathbf{A}\|_1 \|\mathbf{b}\|_1$. Summing yields the lemma statement. \square

Proof of Lemma 5.4. Let $C \subseteq \mathbb{R}^{n+1}$ be a connected component in the partition established in Theorem 4.5, so C can be written as the intersection of at most $14^n(m + 2n)^{3n^2} \tau^{5n^2}$ polynomial constraints of degree at most 5. Let $D \subseteq [-U, U]^m$ be a connected component in the partition established in Lemma 5.2. By Lemma 5.3, there are at most $14^n(m + 2n)^{3n^2} \tau^{5n^2}$ polynomials of degree at most 10 partitioning D into connected components such that within each component, $\mathbf{1}[(\alpha(\mathbf{u}), \beta(\mathbf{u})) \in C]$ is invariant. If we consider the overlay of these polynomial surfaces over all components C , we will get a partition of $[-U, U]^m$ such that for every C , $\mathbf{1}[(\alpha(\mathbf{u}), \beta(\mathbf{u})) \in C]$ is invariant over each connected component of $[-U, U]^m$. Once we have this we are done, since all \mathbf{u} in the same connected component of $[-U, U]^m$ will be sent to the same connected component of \mathbb{R}^{n+1} by $(\alpha(\mathbf{u}), \beta(\mathbf{u}))$, and thus by Theorem 4.5 the behavior of branch-and-cut will be invariant.

We now tally up the total number of surfaces. The number of connected components C was given by Warren's theorem and the Milnor-Thom theorem to be $O(14^{n(n+1)}(m + 2n)^{3n^2(n+1)} \tau^{5n^2(n+1)})$, so the total number of degree-10 hypersurfaces is $14^n(m + 2n)^{3n^2} \tau^{5n^2}$ times this quantity, which yields the lemma statement. \square

D.1 Multiple GMI cuts at the root

In this section we extend our results to allow for multiple GMI cuts at the root of the B&C tree. These cuts can be added simultaneously, sequentially, or in rounds. If GMI cuts $\mathbf{u}_1, \mathbf{u}_2$ are added simultaneously, both of them have the same dimension and are defined in the usual way. If GMI cuts $\mathbf{u}_1, \mathbf{u}_2$ are added sequentially, \mathbf{u}_2 has one more entry than \mathbf{u}_1 . This is because when cuts are added sequentially, the LP relaxation is re-solved after the addition of the first cut, and the second cut has a multiplier for all original constraints as well as for the first cut (this ensures that the second cut can be chosen in a more informed manner). If K cuts are made at the root, they can be added in sequential rounds of simultaneous cuts. In the following discussion, we focus on the case where all K cuts are added sequentially—the other cases can be viewed as instantiations of this. We refer the reader to the discussion in Balcan et al. [8] for more details.

To prove an analogous result for multiple GMI cuts (in sequence, that is, each successive GMI cut has one more parameter than the previous), we combine the reasoning used in the single-GMI-cut case with some technical observations in Balcan et al. [8].

Lemma D.1. *Consider the family of K sequential GMI cuts parameterized by $\mathbf{u}_1 \in [-U, U]^m, \mathbf{u}_2 \in [-U, U]^{m+1}, \dots, \mathbf{u}_K \in [-U, U]^{m+K-1}$. For any IP $(\mathbf{c}, A, \mathbf{b})$, there are at most*

$$O(nK(1+U)^{2K} \|A\|_1 \|\mathbf{b}\|_1)$$

degree- K polynomial hypersurfaces and

$$2^{O(n^2)} K^{O(n^3)} (m+2n)^{O(n^3)} \tau^{O(n^3)}$$

degree- $4K^2$ polynomial hypersurfaces partitioning $[-U, U]^m \times \dots \times [-U, U]^{m+K-1}$ connected components such that the B&C tree built after sequentially adding the GMI cuts defined by $\mathbf{u}_1, \dots, \mathbf{u}_K$ is invariant over all $(\mathbf{u}_1, \dots, \mathbf{u}_K)$ within a single component.

Proof. We start with the setup used by Balcan et al. [8] to prove similar results for sequential Chvátal-Gomory cuts. Let $\mathbf{a}_1, \dots, \mathbf{a}_n \in \mathbb{R}^m$ be the columns of A . We define the following augmented columns $\tilde{\mathbf{a}}_i^1 \in \mathbb{R}^m, \dots, \tilde{\mathbf{a}}_i^K \in \mathbb{R}^{m+K-1}$ for each $i \in [n]$, and the augmented constraint vectors $\tilde{\mathbf{b}}^1 \in \mathbb{R}^m, \dots, \tilde{\mathbf{b}}^K \in \mathbb{R}^{m+K-1}$ via the following recurrences:

$$\begin{aligned} \tilde{\mathbf{a}}_i^1 &= \mathbf{a}_i \\ \tilde{\mathbf{a}}_i^k &= \begin{bmatrix} \tilde{\mathbf{a}}_i^{k-1} \\ \mathbf{u}_{k-1}^T \tilde{\mathbf{a}}_i^{k-1} \end{bmatrix} \end{aligned}$$

and

$$\begin{aligned} \tilde{\mathbf{b}}^1 &= \mathbf{b} \\ \tilde{\mathbf{b}}^k &= \begin{bmatrix} \tilde{\mathbf{b}}^{k-1} \\ \mathbf{u}_{k-1}^T \tilde{\mathbf{b}}^{k-1} \end{bmatrix} \end{aligned}$$

for $k = 2, \dots, K$. In other words, $\tilde{\mathbf{a}}_i^k$ is the i th column of the constraint matrix of the IP and $\tilde{\mathbf{b}}^k$ is the constraint vector after applying cuts $\mathbf{u}_1, \dots, \mathbf{u}_{k-1}$. An identical induction argument to that of Balcan et al. [8] shows that for each $k \in [K]$,

$$\lfloor \mathbf{u}_k^T \tilde{\mathbf{a}}_i^k \rfloor \in \left[-(1+U)^k \|\mathbf{a}_i\|_1, (1+U)^k \|\mathbf{a}_i\|_1 \right]$$

and

$$\lfloor \mathbf{u}_k^T \tilde{\mathbf{b}}^k \rfloor \in \left[-(1+U)^k \|\mathbf{b}\|_1, (1+U)^k \|\mathbf{b}\|_1 \right].$$

Now, as in the single-GMI-cut setting, consider the surfaces

$$\lfloor \mathbf{u}_k^T \tilde{\mathbf{a}}_i^k \rfloor = \ell_i \iff \ell_i \leq \mathbf{u}_k^T \tilde{\mathbf{a}}_i^k < \ell_i + 1 \quad (27)$$

and

$$\lfloor \mathbf{u}_k^T \tilde{\mathbf{b}}^k \rfloor = \ell_0 \iff \ell_0 \leq \mathbf{u}_k^T \tilde{\mathbf{b}}^k < \ell_0 + 1 \quad (28)$$

for every i, k , and every integer $\ell_i \in [-(1+U)^k \|\mathbf{a}_i\|_1, (1+U)^k \|\mathbf{a}_i\|_1] \cap \mathbb{Z}$ and every integer $\ell_0 \in [-(1+U)^k \|\mathbf{b}\|_1, (1+U)^k \|\mathbf{b}\|_1] \cap \mathbb{Z}$. In addition, consider the surfaces

$$\mathbf{u}_k^T \tilde{\mathbf{a}}_i^k - \ell_i = \mathbf{u}_k^T \tilde{\mathbf{b}}^k - \ell_0 \quad (29)$$

for each i, k, ℓ_i, ℓ_0 . As observed by Balcan et al. [8], $\mathbf{u}_k^T \tilde{\mathbf{a}}_i^k$ is a polynomial in $\mathbf{u}_1[1], \dots, \mathbf{u}_1[m], \mathbf{u}_2[1], \dots, \mathbf{u}_2[m+1], \dots, \mathbf{u}_k[1], \dots, \mathbf{u}_k[m+k-1]$ of degree at most k (as is $\mathbf{u}_k^T \tilde{\mathbf{b}}^k$), so surfaces 27, 28, and 29 are all degree- K polynomial hypersurfaces for all i, k . Within any connected component of $[-U, U]^m \times \dots \times [-U, U]^{m+K-1}$ induced by these hypersurfaces, $\lfloor \mathbf{u}_k^T \tilde{\mathbf{a}}_i^k \rfloor$ and $\lfloor \mathbf{u}_k^T \tilde{\mathbf{b}}^k \rfloor$ are constant. Furthermore $\mathbf{1}[f_i^k \leq f_0^k]$ is invariant for every i, k , where $f_i^k = \mathbf{u}_k^T \tilde{\mathbf{a}}_i^k - \lfloor \mathbf{u}_k^T \tilde{\mathbf{a}}_i^k \rfloor$ and $f_0^k = \mathbf{u}_k^T \tilde{\mathbf{b}}^k - \lfloor \mathbf{u}_k^T \tilde{\mathbf{b}}^k \rfloor$.

Now, fix a connected component $D \subseteq [-U, U]^m \times \dots \times [-U, U]^{m+K-1}$ induced by the above hypersurfaces, and let $C \subseteq \mathbb{R}^{K(n+1)}$ be the intersection of q polynomial inequalities of degree at most d . Consider a single degree- d polynomial inequality in $K(n+1)$ variables $y_1, \dots, y_{K(n+1)}$, which can be written as

$$\sum_{\substack{T \subseteq [K(n+1)] \\ |T| \leq d}} \lambda_T \prod_{j \in T} y_j = \sum_{\substack{T_1, \dots, T_K \subseteq [n+1] \\ |T_1| + \dots + |T_K| \leq d}} \lambda_{T_1, \dots, T_K} \prod_{j_1 \in T_1} y_{j_1} \dots \prod_{j_K \in T_K} y_{j_K} \leq \gamma.$$

Now, the sets S_1, \dots, S_K defined by $S_k = \{i : f_i^k \leq f_0^k\}$ are fixed within D , so we can write this as

$$\sum_{\substack{T_1, \dots, T_K \subseteq [n+1] \\ |T_1| + \dots + |T_K| \leq d}} \lambda_{T_1, \dots, T_K} \prod_{k=1}^K \left[\prod_{\substack{j \in T_k \cap S_k \\ j \neq n+1}} f_j^k (1 - f_0^k) \prod_{\substack{j \in T_k \setminus S_k \\ j \neq n+1}} f_0^k (1 - f_j^k) \prod_{j \in T_k, j=n+1} f_0^k (1 - f_0^k) \right] \leq \gamma.$$

We have that f_j^k and f_0^k are degree- k polynomials in $\mathbf{u}_1, \dots, \mathbf{u}_k$. Since the sum is over all multisets T_1, \dots, T_K such that $|T_1| + \dots + |T_K| \leq d$, there are at most d terms across the products, each of the form $f_j^k (1 - f_0^k)^k$, $f_0^k (1 - f_j^k)$, or $f_0^k (1 - f_0^k)^k$. Therefore, the left-hand-side is a polynomial of degree at most $2dK$, and if $C \subseteq \mathbb{R}^{K(n+1)}$ is the intersection of q polynomial inequalities each of degree at most d , the set

$$\{(\mathbf{u}_1, \dots, \mathbf{u}_K) \in D : (\alpha(\mathbf{u}_1, \dots, \mathbf{u}_K), \beta(\mathbf{u}_1, \dots, \mathbf{u}_K)) \in C\} \subseteq [-U, U]^m \times \dots \times [-U, U]^{m+K-1}$$

can be expressed as the intersection of q degree- $2dK$ polynomial inequalities.

To finish, we run this process for every connected component $C \subseteq \mathbb{R}^{K(n+1)}$ in the partition established by Theorem C.5. This partition consists of $O(12^n n^{2n} K^{2n^2} (m+2n)^{2n^2} \tau^{5n^2})$ degree- $2K$ polynomials over $\mathbb{R}^{K(n+1)}$. By Warren's theorem and the Milnor-Thom theorem, these polynomials partition $\mathbb{R}^{K(n+1)}$ into $O(12^{n(n+1)} n^{2n(n+1)} K^{2n^2(n+1)} (m+2n)^{2n^2(n+1)} \tau^{5n^2(n+1)})$ connected components. Running the above argument for each of these connected components of $\mathbb{R}^{K(n+1)}$ yields a total of $O(12^{n(n+1)} n^{2n(n+1)} K^{2n^2(n+1)} (m+2n)^{2n^2(n+1)} \tau^{5n^2(n+1)}) \cdot O(12^n n^{2n} K^{2n^2} (m+2n)^{2n^2} \tau^{5n^2}) = 2^{O(n^2)} K^{O(n^3)} (m+2n)^{O(n^3)} \tau^{O(n^3)}$ polynomials of degree $4K^2$. Finally, we count the surfaces of the form (27), (28), and (29). The total number of degree- K polynomials of type 27 is at most $O(nK(1+U)^K \|A\|_1)$, the total number of degree- k polynomials of type 28 is $O(K(1+U)^K \|\mathbf{b}\|_1)$, and the total number of degree- K polynomials of type 29 is $O(nK(1+U)^{2K} \|A\|_1 \|\mathbf{b}\|)$. Summing these counts yields the desired number of surfaces in the lemma statement.

In any connected component of $[-U, U]^m$ determined by these surfaces, $\mathbf{1}[(\alpha(\mathbf{u}), \beta(\mathbf{u})) \in C]$ is invariant for every connected component $C \subseteq \mathbb{R}^{K(n+1)}$ in the partition of $\mathbb{R}^{K(n+1)}$ established in Theorem C.5. This means that the tree built by branch-and-cut is invariant, which concludes the proof. \square

Finally, applying the main result of Balcan et al. [6] to Lemma D.1, we get the following pseudo-dimension bound for the class of K sequential GMI cuts at the root of the B&C tree.

Theorem D.2. For $\mathbf{u}_1 \in [-U, U]^m$, $\mathbf{u}_2 \in [-U, U]^{m+1}, \dots, \mathbf{u}_K \in [-U, U]^{m+K-1}$, let $g_{\mathbf{u}_1, \dots, \mathbf{u}_K}(\mathbf{c}, A, \mathbf{b})$ denote the number of nodes in the tree $B\mathcal{E}C$ builds given the input $(\mathbf{c}, A, \mathbf{b})$ after sequentially applying the GMI cuts defined by $\mathbf{u}_1, \dots, \mathbf{u}_K$ at the root. The pseudo-dimension of the set of functions $\{g_{\mathbf{u}_1, \dots, \mathbf{u}_K} : (\mathbf{u}_1, \dots, \mathbf{u}_K) \in [-U, U]^m \times \dots \times [-U, U]^{m+K-1}\}$ on the domain of IPs with $\|A\|_1 \leq a$ and $\|\mathbf{b}\|_1 \leq b$ is

$$O(mK^3 \log U + mn^3 K^2 \log(mnK\tau) + mK^2 \log(ab)).$$