Christian Wagner (Autor)

Maximal Lattice-Free Polyhedra in Mixed-Integer Cutting Plane Theory

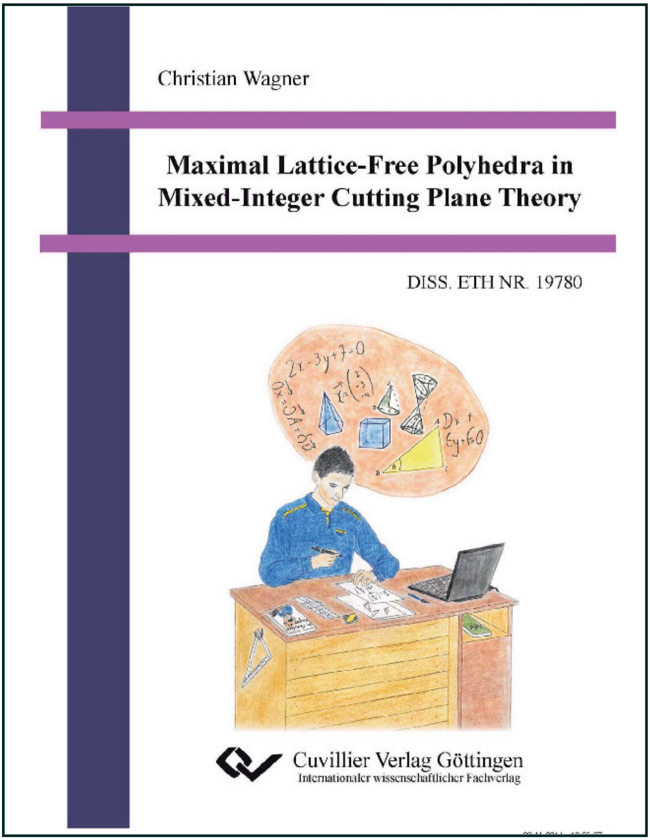

https://cuvillier.de/de/shop/publications/188

CHAPTER 1

INTRODUCTION

The aim of this chapter is to provide an overview of the agenda of this thesis. We introduce the underlying optimization problem and explain step by step our motivation for choosing the research questions that are studied in this thesis.

We assume that a general mixed-integer linear program (MILP) is given in the form

$$\begin{aligned} \max\ c^\top x \quad \text{s.t. } Ax &= b, \\ x &\geq 0, \\ x_i &\in \mathbb{Z} \quad \text{for } i \in \mathcal{I}, \\ x_i &\in \mathbb{R} \quad \text{for } i \in \mathcal{C}, \end{aligned} \tag{1.1}$$

where A, b, and c are rational and $\mathcal{I}$, resp. $\mathcal{C}$, is a set of integer constrained, resp. continuous, variables. The linear programming relaxation of (1.1) is the optimization problem (1.1) where the condition $x_i \in \mathbb{Z}$ is replaced by the weaker condition $x_i \in \mathbb{R}$ for all $i \in \mathcal{I}$. To avoid trivial cases we assume that the feasible region of (1.1) is non-empty and that its linear programming relaxation is bounded. Solving the linear programming relaxation yields an optimal vertex x^* with corresponding sets B and N of basic and non-basic variables which satisfy

$$x_i = f_i + \sum_{j \in N} r_i^j x_j \quad \forall i \in B,$$

where $f_i \in \mathbb{Q}_+$ and $r_i^j \in \mathbb{Q}$ for all $i \in B$ and all $j \in N$. We assume that x^* is not feasible for (1.1), otherwise we have already found an optimal solution.

Our aim is to generate cutting planes (or *cuts* for short) which cut off x^*, i.e. inequalities which are valid for every feasible point of (1.1), but violated by x^*.

Virtually all traditional cutting planes that are used by general-purpose MILP solvers, most notably lift-and-project cuts (see, for instance, [BCC93]), Gomory mixed-integer cuts (see, for instance, [Gom60]), or mixed-integer rounding cuts (see, for instance, [NW90]), are derived by considering only one equation. Normally, the strategy is to generate a linear combination of the original constraints $Ax = b$. Then one applies integrality arguments to the resulting equation. Cuts obtained in this way are *split cuts* (see, for instance, [CKS90]). Unfortunately, an approach that is based on such cuts alone does not give rise to a finite cutting plane algorithm. In [CKS90], an instance in only three variables is presented and it is shown that a cutting plane algorithm based on split cuts does not converge finitely.

Example 1.1. Consider the following MILP.

$$\begin{aligned} \max\ t \quad \text{s.t.} \qquad -x_1 + t &\leq 0, \\ -x_2 + t &\leq 0, \\ x_1 + x_2 + t &\leq 2, \\ x_1, x_2 &\in \mathbb{Z}, \\ t &\in \mathbb{R}_+. \end{aligned}$$

The cut needed to solve this problem is $t \leq 0$. However, in [CKS90] it is shown that this cut cannot be obtained by applying split cuts only. ◇

In [ALWW07], Andersen et al. initiated a new approach for cutting plane generation by considering two rows of a simplex tableau simultaneously. This approach allows to deduce cutting planes that cannot be obtained by considering one single equation. In particular, the desired cut $t \leq 0$ in Example 1.1 can be derived immediately.

Meanwhile, the two-row case has been analyzed quite exhaustively, most notably due to Andersen et al. [ALWW07], Borozan and Cornuéjols [BC07], Cornuéjols and Margot [CM08], and Basu et al. [BBCM11]. However, the basic idea of the two-row approach can be generalized to the case of multiple rows in a straightforward way. For that, the point of departure is an optimal vertex x^* of the linear programming relaxation of (1.1). We assume that $m := |B \cap \mathcal{I}| \geq 2$ and $f_i \notin \mathbb{Z}$ for at least one $i \in B \cap \mathcal{I}$. We consider the set

$$P_I := \left\{ (x, s) \in \mathbb{Z}^m \times \mathbb{R}^n_+ : x = f + \sum_{j \in N} r^j s_j \right\},$$

where $N := \{1, \dots, n\}$ represents the non-basic variables and f, resp. r^j, is the vector consisting of all f_i's, resp. r^j_i's, such that $i \in B \cap \mathcal{I}$.

The set P_I is the underlying mixed-integer set in this thesis. Our motivation for analyzing P_I is that it can be obtained as a relaxation of the feasible region of a general MILP. Therefore, valid inequalities for P_I give rise to cutting planes for the original mixed-integer set. Consequently, our aim is to derive valid inequalities for P_I, or equivalently, for $\text{conv}(P_I)$.

In Chapter 3, we show that valid inequalities for $\text{conv}(P_I)$ correspond to combinatorial objects in the space of the discrete variables. More precisely, they correspond to *lattice-free* polyhedra, i.e. polyhedra that do not contain an interior integer point. The basic properties of the set $\text{conv}(P_I)$ are summarized in Section 3.1, and the relation between lattice-free polyhedra and the facet-defining inequalities for $\text{conv}(P_I)$ is presented in Section 3.2.

By considering $\text{conv}(P_I)$, the feasible region of the original MILP (1.1) is relaxed in two ways. First, we drop all integrality conditions on the non-basic variables. Second, the non-negativity restrictions on all basic variables are ignored. The latter relaxation has been introduced by Gomory [Gom69] and is known as the classical group relaxation. The first relaxation, however, is the great novelty in the new approach. It preserves much of the complexity of the original model, but keeps it sufficiently simple to analyze it.

The following example illustrates the cutting plane approach that we have in mind.

Example 1.2. Fig. 1.1 exemplifies our intended approach to generate cutting planes. For simplicity, let $m = 2$. The gray regions in Fig.s 1.1(a) and 1.1(b) represent the projection of the linear programming relaxation onto the space of the x-variables. After relaxing the integrality conditions on the non-basic variables and the non-negativity restrictions on the basic variables, we obtain a *corner polyhedron* (see, for instance, [Gom69]). The convex hull of the two dashed half-lines in Fig.s 1.1(a)–1.1(c) is the projection of the corner polyhedron onto the space of the x-variables. Fig. 1.1(b) shows how the solid lattice-free triangle is used to cut off x^*. The intersection points of the triangle and the two dashed half-lines determine the cutting plane. After adding the cutting plane, the feasible region of the linear programming relaxation becomes smaller. Its projection onto the space of the x-variables is the gray region in Fig. 1.1(c). ◇

Since, by assumption, x^* is not feasible for (1.1), we aim at generating cutting planes that are violated by the basic solution $x_i^* = f_i$ for all $i \in B$ and $x_j^* = 0$ for all $j \in N$. For that, we look for valid inequalities for $\text{conv}(P_I)$ which cut off the point (f, o). It turns out that the non-trivial facet-defining inequalities for $\text{conv}(P_I)$, i.e. the strongest inequalities that we can derive from our relaxation, do perform this task: all of them are violated by (f, o). This implies that we can focus our attention on non-trivial facet-defining inequalities for $\text{conv}(P_I)$. At this point the enormous power of the applied

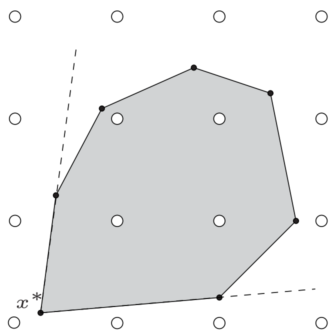

(a) The point x^* must be cut off from the integer points in the feasible region.

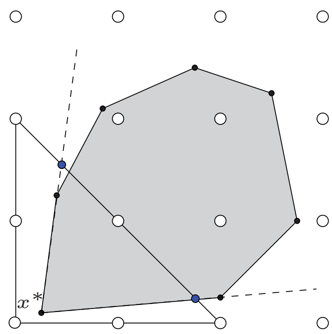

(b) The solid triangle contains x^* in its interior, but none of the feasible integer points.

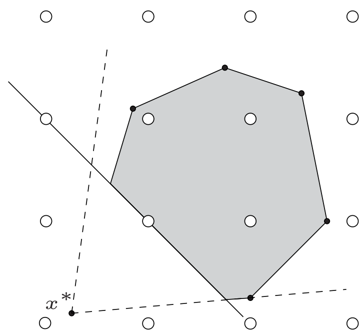

(c) The solid line separates x^* from the integer points in the feasible region.

Figure 1.1: Derivation of a cutting plane.

relaxation comes into play, because all these non-trivial facet-defining inequalities correspond to lattice-free polyhedra which possess beautiful geometrical properties. The exact relation is stated in Theorem 3.9 where we show that every non-trivial facet-defining inequality for $\operatorname{conv}(P_I)$ can be derived from a lattice-free polyhedron which has a representation as the sum of a polytope and a linear space. It follows that strongest cutting planes are associated with *maximal lattice-free polyhedra*, i.e. lattice-free polyhedra which are not properly contained in another lattice-free polyhedron. Structural properties of maximal lattice-free polyhedra entail information on the corresponding cutting planes and therefore, instead of analyzing cutting planes, we can equivalently analyze maximal lattice-free polyhedra. As a result of this, several questions related to these polyhedra and their associated cutting planes arise.

Certainly, the aim in cutting plane generation should not be to produce a bulk of cuts which just cut off the current optimal linear programming solution, but rather to identify a (preferably small) set of well-chosen cuts which are "important" in some sense. Here, "important" is difficult to define. There are several approaches to evaluate cutting planes, for instance with respect to the volume which is cut off, a comparison of the cut coefficients, or the improvement of the objective function value after adding a cut or a set of cuts. The choice of the measure is highly dependent on the particular structure of the problem. Since we start from a general MILP it is simply not possible to say which measure is most suitable. In this thesis, we use a strength measure of Goemans [Goe95] to evaluate non-trivial facet-defining inequalities for $\text{conv}(P_I)$. Every such inequality corresponds to a maximal lattice-free polyhedron in the m-dimensional space of the x-variables. And each such polyhedron P can be represented as $P = \mathcal{P} + \mathcal{L}$, where $\mathcal{P}$ is a polytope and $\mathcal{L}$ is a linear space. The codimension of $\mathcal{L}$ is called the *split-dimension of* P. In turn, the split-dimension of a non-trivial facet-defining inequality I for $\text{conv}(P_I)$ is defined to be the smallest split-dimension of a maximal lattice-free polyhedron P such that P can be used to derive an inequality for $\text{conv}(P_I)$ which is equal to or which dominates the inequality I.

In Chapter 4, we investigate which of the non-trivial facet-defining inequalities for $\text{conv}(P_I)$ are needed to approximate $\text{conv}(P_I)$ sufficiently well with respect to the strength measure of Goemans. In Theorem 4.4, we show that, in general, good approximations for $\text{conv}(P_I)$ can be expected only by having available all the non-trivial facet-defining inequalities for $\text{conv}(P_I)$ of split-dimension m. This result is clearly unsatisfactory since the complexity of the corresponding maximal lattice-free polyhedra increases with increasing split-dimension. Consequently, inequalities of split-dimension m are difficult to generate. In contrast to this negative result on the strength, in Theorem 4.7, we show that by restricting the size of the data, inequalities of split-dimension m can be approximated using inequalities of split-dimension one (i.e. split cuts). This is a positive message since split cuts are the easiest objects in terms of complexity. In particular, we show that, given the dimension m of the x-variable space, the fractionality of the current optimal solution (f, o), and the *max-facet-width* of a lattice-free polyhedron P of split-dimension m, then the inequality corresponding to P can be approximated to within a constant factor which involves only these three quantities. For the special case where P is a *regular lattice-free simplex* (RLS), in Theorem 4.8, we even state a constant which involves only the dimension m. This raises hope that cuts with low split-dimension perform well in practice.

In Chapter 6, we address the case $m = 2$ in order to obtain deeper results on the approximability of inequalities of split-dimension two by split cuts. As pointed out, the non-trivial facet-defining inequalities for $\text{conv}(P_I)$ are associated with maximal lattice-free polyhedra. In dimension two, these polyhedra can be partitioned into five types which are shown in Fig. 1.2 (see Proposition 5.3 on p. 37 for the precise definition of each type).

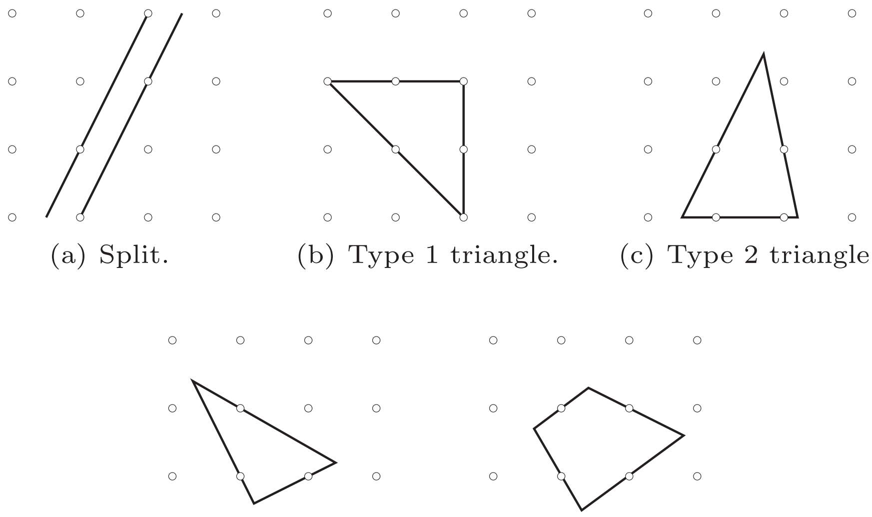

(a) Split. (b) Type 1 triangle. (c) Type 2 triangle.

(d) Type 3 triangle. (e) Quadrilateral.

Figure 1.2: All types of two-dimensional maximal lattice-free polyhedra.

Since every non-trivial facet-defining inequality for $\text{conv}(P_I)$ corresponds to one of the above maximal lattice-free sets, they are called split, type 1, type 2, type 3, or quadrilateral inequalities. In [BBCM11] it has been shown that the closures of split and type 1 inequalities may produce an arbitrarily bad approximation of $\text{conv}(P_I)$, whereas the closures of type 2 or type 3 or quadrilateral inequalities deliver good approximations in terms of the strength measure of Goemans. More concretely, in [BBCM11] sequences of examples are constructed in which cuts from triangles of types 2 and 3, and quadrilaterals cannot be approximated to within a constant factor by using split and type 1 inequalities only. The approximation becomes worse as the triangles and quadrilaterals converge towards a split. We think that this is geometrically counterintuitive. Therefore, in Chapter 6, we refine the argument by taking into consideration the probability that such a situation emerges when f is uniformly distributed in the interior of a given maximal lattice-free triangle of type 2, type 3, or quadrilateral. The precise model is explained in Section 6.2. Our main result of the probabilistic analysis in

Chapter 6 is stated in Theorem 6.2, where we show that the addition of a single type 2 inequality to the split closure becomes less likely to be beneficial the closer the type 2 triangle looks like a split. Our analysis in Chapter 6 suggests that this is true for type 3 and quadrilateral inequalities as well.

The performance of cuts may be evaluated in two different ways. In dimension two, if one considers only *one round* of cuts, then – using the strength measure of Goemans – split and type 1 inequalities can be arbitrarily bad in approximating $\text{conv}(P_I)$. On the other hand, within a cutting plane framework where *several rounds* of cuts are considered, it is enough to add split and type 1 inequalities iteratively, in order to terminate with an optimal mixed-integer point after a finite number of applied rounds (see [DL09] and also [BCM11] and [DPW11] for a generalization of the results in [DL09]). Using the correspondence between the non-trivial facet-defining inequalities for $\text{conv}(P_I)$ and maximal lattice-free polyhedra this insight leads to a natural question: Which maximal lattice-free polyhedra are important in a cutting plane framework? Admittedly, this question is too general to be answered completely within this thesis. Nevertheless, the answer must have to do with the integer points on the boundary of the maximal lattice-free polyhedra. Since rationality of the input data is assumed we only need to consider maximal lattice-free rational polyhedra.

In Chapter 7, we show in Theorem 7.2 that, given the dimension (i.e. the number of simplex tableau rows from which a non-trivial facet-defining inequality for $\text{conv}(P_I)$ is derived) and the rationality of a corresponding maximal lattice-free polyhedron P, then only finitely many different shapes are possible for P, provided we identify any two polyhedra which coincide up to a transformation which preserves the integer lattice. Unfortunately, "finitely many" does not mean "few". Indeed, in Section 7.3 we provide an upper bound on the volume of such a polytope. Our bound is by far not best possible, but suggests that the number of potential shapes may explode dramatically with increasing dimension. This makes clear that there is no chance to enumerate all shapes based on a computer code, even for small dimensions.

The fact that in dimension two only split and type 1 inequalities are needed within a cutting plane framework is not just coincidence, but rather has to do with the integer points on the boundary. The *X-body*[1] of a lattice-free polyhedron is the convex hull of the integer points on its boundary. In particular, a lattice-free polyhedron coincides with its X-body if and only if it is an *integral polyhedron* in the sense that every minimal (non-empty) face contains integer points. In Fig. 1.2, only the split and the type 1 triangle

[1]The notion "X-body" is quite unintuitive, but we use it for historical reasons.

are integral polyhedra. Recently, it has been proved by Del Pia and Weismantel [DPW11] that the X-body of a lattice-free polyhedron is connected with the importance of the polyhedron in a cutting plane procedure. To be precise, within a cutting plane framework, only *lattice-free integral polyhedra* are needed. Thus, a characterization of maximal lattice-free integral polyhedra is desired. In dimensions one and two, all shapes of maximal lattice-free integral polyhedra are known. On the other hand, their number is expected to be huge in dimensions beyond three.

In Chapter 8, we classify all three-dimensional maximal lattice-free integral polyhedra. We first show that we can restrict our attention to polytopes. Then, in Theorem 8.1, we enumerate all three-dimensional maximal lattice-free polytopes with integer vertices.

Theorems 6.2 and 8.1 are proved by intensively using two-dimensional tools which cannot be deduced offhand. Therefore, we dedicate an extra chapter to the two-dimensional relation between the area and the *lattice width* of lattice-free convex sets. In Chapter 5, we prove several inequalities which involve the area and the lattice width in the plane. In Theorem 5.6, we present our results for arbitrary lattice-free convex sets and in Theorem 5.9 we present our results for centrally symmetric ones. We further characterize the extreme lattice-free convex sets and relate our results to the *covering minima* introduced in [KL86]. Moreover, in Theorem 5.10 we rectify a result of [KL88] with a new proof.