

Consistency for 0–1 Programming

Danial Davarnia¹ and J. N. Hooker²

¹ Iowa state University
davarnia@iastate.edu

² Carnegie Mellon University
jh38@andrew.cmu.edu

Abstract. Concepts of consistency have long played a key role in constraint programming but never developed in integer programming (IP). Consistency nonetheless plays a role in IP as well. For example, cutting planes can reduce backtracking by achieving various forms of consistency as well as by tightening the linear programming (LP) relaxation. We introduce a type of consistency that is particularly suited for 0-1 programming and develop the associated theory. We define a 0-1 constraint set as LP-consistent when any partial assignment that is consistent with its linear programming relaxation is consistent with the original 0-1 constraint set. We prove basic properties of LP-consistency, including its relationship with Chvátal-Gomory cuts and the integer hull. We show that a weak form of LP-consistency can reduce or eliminate backtracking in a way analogous to k -consistency. This work suggests a new approach to the reduction of backtracking in IP that focuses on cutting off infeasible partial assignments rather than fractional solutions.

Keywords: Consistency · Resolution · Constraint satisfaction · Integer programming · Backtracking · Cutting planes

1 Introduction

Consistency is a fundamental concept of constraint programming (CP) and an essential tool for the reduction of backtracking during search [1]. Curiously, the concept never explicitly developed in mathematical programming, even though solvers rely on a similar type of branching search. In fact, the cutting planes of integer programming can reduce backtracking by achieving various forms of consistency as well as by tightening the linear programming (LP) relaxation.

This suggests that it may be useful to investigate the potential role of consistency concepts in mathematical programming. We do so for 0–1 integer programming in particular. We study how consistency relates to such integer programming ideas as the LP relaxation, Chvátal-Gomory cutting planes [3], and the integer hull, as well as how consistency can be achieved for 0–1 inequalities. Our main contribution is to introduce a type of consistency, *LP-consistency*, that seems particularly relevant to 0–1 programming, and to develop the underlying theory. We show that achieving a form of partial LP-consistency can reduce backtracking in ways that traditional cutting planes do not.

One way to reduce backtracking is to identify partial assignments to the variables that are *inconsistent* with the constraint set, meaning that they cannot occur in a feasible solution of the constraints. Branching decisions that result in such partial assignments can then be avoided, thus removing infeasible subtrees from the search. Unfortunately, it is generally hard to identify inconsistent partial assignments in advance.

The essence of consistency is that it makes it easier to identify inconsistent partial assignments. Full consistency allows one to recognize an inconsistent partial assignment by the fact that it violates a constraint that contains only the variables in the partial assignment. Because full consistency is very hard to achieve, CP solvers rely on *domain consistency* (generalized arc consistency) [1,4,10,11], which reduces variable domains to the point that every value in them occurs in some feasible solution. If domain consistency is obtained at the current node of the search tree, branching on any value in a variable's domain can lead to a feasible solution. Domain consistency is itself hard to achieve for the entire constraint set, but can often be achieved, or partially achieved, for individual global constraints in the CP model, and this reduces backtracking significantly [14].

Our approach is based on the idea that consistency can be defined with respect to a *relaxation* of the constraint set. Specifically, we interpret consistency as making it possible to identify inconsistent partial assignments by checking whether they are consistent with a certain type of relaxation. This perspective allows us to propose alternative types of consistency by using various types of relaxation. For traditional consistency, the relaxation is obtained simply by dropping constraints that contain variables that are not in the partial assignment. We define LP-consistency by replacing this relaxation with the LP relaxation. Thus LP-consistency ensures that any partial assignment that is consistent with the LP relaxation is consistent with the original constraint set. Fortunately, one can easily check consistency with an LP relaxation simply by solving the LP problem that results from adding the partial assignment to the LP relaxation.

This poses the question of whether it is practical to achieve LP-consistency for a 0–1 problem. There is no known practical method for achieving full LP-consistency, but we take a cue from the concept of k -consistency in CP [5,15,16], which is weaker than full consistency but sufficient to avoid backtracking if the constraints are not too tightly coupled by common variables. We define a similar property, *sequential LP k -consistency*, that can avoid some backtracking that traditional cutting planes may permit, because it focuses on identifying inconsistent partial assignments rather than cutting off fractional solutions of the LP relaxation.

A method for obtaining sequential LP k -consistency is suggested by our practice of defining consistency concepts in terms of projection, as proposed in [9]. One can define sequential LP k -consistency, in particular, in terms of the results of lifting a problem from $k - 1$ dimensions to k dimensions, and then projecting it back into $k - 1$ dimensions. A modified form of the well-known lift-and-project technique of IP [2] achieves sequential LP k -consistency.

We begin below by defining and illustrating basic consistency concepts and showing how they can be cast in terms of projection. We also indicate how consistency can eliminate or reduce backtracking. We review some prior work showing that an inference method of propositional logic, resolution, can achieve consistency for 0–1 problems, and that a weak form of resolution, input resolution, can generate all Chvátal-Gomory cuts for a set of logical clauses.

At this point we introduce LP-consistency and show some elementary properties, namely that consistency implies LP-consistency, and a constraint set that describes the integer hull is necessarily LP-consistent. Yet LP-consistency is a concept that does not occur in polyhedral theory, and an LP-consistent constraint set need not describe the integer hull. While the facet-defining inequalities that describe the integer hull are generally regarded as the strongest valid inequalities, we show that they can be weaker than a non-facet-defining inequality that achieves LP-consistency, in the sense that they exclude fewer inconsistent 0–1 (partial) assignments. We further elaborate on connections with cutting plane theory by showing that a 0–1 partial assignment is consistent with the LP relaxation if and only if it violates no logical clause that is a Chvátal-Gomory (C-G) cut, and a 0–1 problem is LP-consistent if and only if all of its implied logical clauses are C-G cuts. We also note that while input resolution derives C-G cuts, it does not achieve LP-consistency.

The remainder of the paper defines and develops the concept of sequential LP k -consistency. It shows that achieving sequential LP k -consistency for $k = 1, \dots, n$ (where n is the number of variables) avoids backtracking altogether for branching order x_1, x_2, \dots, x_n . In practice, one would achieve sequential LP k -consistency for a few small values of k . We then prove that one step of the lift-and-project procedure [2] achieves sequential k -consistency for a given k . Finally, we illustrate how achieving sequential LP k -consistency even for $k = 2$ can avoid backtracking that is permitted by traditional separating cuts.

2 Consistency and Projection

To define consistency, it is convenient to adopt basic terminology as follows. The *domain* D_j of a variable x_j is the set of values that can be assigned to x_j . A *constraint* C is an object that *contains* some set $\{x_1, \dots, x_k\}$ of variables, such that any given assignment of values to (x_1, \dots, x_k) either *satisfies* or *violates* C . Thus a constraint is satisfied or violated only when all of its variables have been assigned values. An assignment to x satisfies a constraint set \mathcal{S} when it satisfies all the constraints in \mathcal{S} . A list of symbols defined hereafter appears in Table 1.

Let x_J be the tuple containing the variables in $\{x_j \mid j \in J\}$ for $J \subseteq N = \{1, \dots, n\}$. A *partial assignment* to x is an assignment of values to x_J for some $J \subseteq N$. We can now define a consistent partial assignment and a consistent constraint set.

Definition 1. *Given a constraint set \mathcal{S} , a partial assignment $x_J = v_J$ is consistent with \mathcal{S} if $\mathcal{S} \cup \{x_J = v_J\}$ is feasible.*

Since it is hard in general to determine whether $\mathcal{S} \cup \{x_J = v_J\}$ is feasible, it is hard to identify which partial assignments are consistent with \mathcal{S} . Consistent constraint sets are defined so that it is easy to identify which partial assignments are consistent with them.

Definition 2. *A constraint set \mathcal{S} is consistent if every partial assignment to x that violates no constraint in \mathcal{S} is consistent with \mathcal{S} .*

The contrapositive is perhaps more intuitive: \mathcal{S} is consistent when every partial assignment that is inconsistent with \mathcal{S} violates some individual constraint in \mathcal{S} . Thus a consistent constraint set can be viewed as one in which implied constraints are made explicit, in the sense that every inconsistent partial assignment is explicitly ruled out by some constraint in the set.

Since full consistency is generally hard to achieve, the constraint programming community has found various weaker forms of consistency to be more useful. By far the most popular is domain consistency, also known as generalized arc consistency [1,4,10,11].

Definition 3. *A constraint set \mathcal{S} is domain consistent if $x_j = v_j$ is consistent with \mathcal{S} for all $v_j \in D_j$ and all variables x_j .*

That is, every value in the domain of a variable x_j is assigned to x_j in some feasible solution of \mathcal{S} . A consistent constraint set is necessarily domain consistent.

Example 1. Suppose that \mathcal{S} is the constraint set

$$\begin{aligned} x_1 + x_2 &+ x_4 \geq 1 \\ x_1 - x_2 + x_3 &\geq 0 \\ x_1 &- x_4 \geq 0 \\ x_j &\in \{0, 1\}, \text{ all } j \end{aligned}$$

Table 1: List of symbols.

x_J	tuple of variables x_j for $j \in J$
N	the set $\{1, \dots, n\}$
J_k	the set $\{1, \dots, k\}$
D_j	domain of x_j
D_J	cartesian product of D_j for $j \in J$
$D(\mathcal{S})$	set of assignments* to x that satisfy \mathcal{S}
$D_J(\mathcal{S})$	set of assignments* to x_J that are consistent with \mathcal{S}
$D(\mathcal{S}) _J$	projection of $D(\mathcal{S})$ onto x_J
\mathcal{S}_J	set of constraints in \mathcal{S} that contain only variables in x_J
\mathcal{S}_{LP}	LP relaxation of 0–1 constraint set \mathcal{S}
$D_J(\mathcal{S}_{LP})$	set of 0–1 assignments to x_J that are consistent with \mathcal{S}_{LP}
\mathcal{S}_C	set of clausal inequalities implied by individual constraints of \mathcal{S}

*an assignment $x = v$ or $x_J = v_J$ assumes that $v \in D$, $v_J \in D_J$.

The domains are $D_j = \{0, 1\}$ for $j = 1, \dots, 4$. The feasible solutions (x_1, \dots, x_4) of \mathcal{S} are listed below:

$$\begin{array}{ccc} (0, 1, 1, 0) & (1, 0, 1, 0) & (1, 1, 0, 1) \\ (1, 0, 0, 0) & (1, 0, 1, 1) & (1, 1, 1, 0) \\ (1, 0, 0, 1) & (1, 1, 0, 0) & (1, 1, 1, 1) \end{array}$$

Set \mathcal{S} is not consistent because, for instance, the partial assignment $(x_1, x_2) = (0, 0)$ violates no constraint in \mathcal{S} but is inconsistent with \mathcal{S} due to the fact that $(x_1, x_2) = (0, 0)$ in none of the feasible solutions. On the other hand, \mathcal{S} is domain consistent because $x_j = 0$ and $x_j = 1$ occur in some feasible solution for each j .

The various consistency concepts are more easily defined in terms of projection, as proposed in [9]. Let D_J be the cartesian product of D_j for $j \in J$, and let $D = D_N$. When we speak of an assignment $x = v$ or a partial assignment $x_J = v_J$, we assume $v \in D$ and $v_J \in D_J$. Let $D(\mathcal{S})$ be the set of assignments to x that satisfy \mathcal{S} , and let $D_J(\mathcal{S})$ be the set of assignments to x_J that are consistent with \mathcal{S} . Thus

$$D_J(\mathcal{S}) = \{v_J \in D_J \mid \mathcal{S} \cup \{x_J = v_J\} \text{ is feasible}\}$$

The *projection* of $D(\mathcal{S})$ onto x_J , which we may write $D(\mathcal{S})|_J$, is $\{x_J \mid x \in D(\mathcal{S})\}$.

We can now define consistency in terms of projection. Let \mathcal{S}_J be the set of constraints in \mathcal{S} whose variables belong to x_J . Then $D_J(\mathcal{S}_J)$ is the set of assignments to x_J that violate no constraints in \mathcal{S} .

Proposition 1. *A constraint set \mathcal{S} is consistent if and only if $D_J(\mathcal{S}_J) = D(\mathcal{S})|_J$ for all $J \subseteq N$. In addition, \mathcal{S} is domain consistent if and only if $D_j = D(\mathcal{S})|_{\{j\}}$ for all $j \in N$.*

3 Consistency and Backtracking

It is well known that consistency is closely related to backtracking. We note first that branching can find a feasible solution for a fully consistent constraint set without backtracking, assuming of course that the constraints have a solution. Suppose we branch on variables x_1, \dots, x_n in that order. Each node in level j of the branching tree corresponds to a partial assignment $(x_1, \dots, x_{j-1}) = (v_1, \dots, v_{j-1})$. We branch on x_j at the node by assigning to x_j each value $v_j \in D_j$ for which the partial assignment $(x_1, \dots, x_j) = (v_1, \dots, v_j)$ violates no constraint in \mathcal{S} . Due to the consistency of \mathcal{S} , this partial assignment is consistent with \mathcal{S} for at least one value $v_j \in D_j$. Thus branching can continue to the bottom of the tree with no need to backtrack.

A weaker form of consistency, k -consistency, avoids backtracking if there is limited coupling of variables [6]. More relevant to our purposes is a still weaker form of consistency that assumes the branching order is given, namely x_1, \dots, x_n . Let $J_k = \{1, \dots, k\}$.

Definition 4. A constraint set \mathcal{S} is sequentially k -consistent if $D_{J_{k-1}}(\mathcal{S}_{J_{k-1}}) = D_{J_k}(\mathcal{S}_{J_k})|_{J_{k-1}}$.

Thus \mathcal{S} is sequentially k -consistent if for every partial assignment $(x_1, \dots, x_{k-1}) = (v_1, \dots, v_{k-1})$ that violates no constraint in \mathcal{S} , there is a value v_k in D_k such that $(x_1, \dots, x_k) = (v_1, \dots, v_k)$ violates no constraint in \mathcal{S} . The following is easy to show.

Proposition 2. If the branching order is x_1, \dots, x_n , constraint set \mathcal{S} can be solved without backtracking if \mathcal{S} is sequentially k -consistent for $k = 1, \dots, n$.

Example 2. Let $\mathcal{S} = \{3x_1 + 2x_2 \geq 1, -x_1 + 2x_2 \geq 0, x \in \{0, 1\}^2\}$. Proposition 2 implies that we can avoid backtracking by branching in the order x_1, x_2 , because \mathcal{S} is sequentially 1-consistent and sequentially 2-consistent.

4 Consistency and Resolution

Previous research has shown that the resolution procedure of propositional logic achieves consistency for a 0–1 constraint set. First, some definitions. A *literal* ℓ_j is a proposition of the form x_j or $\neg x_j$. A logical *clause* is a disjunction $\bigvee_{j \in J} \ell_j$ of literals. A clause C_1 implies C_2 when C_1 *absorbs* C_2 , meaning that all the literals of C_1 are in C_2 . There is a resolution proof of any clause that is logically implied by a clause set \mathcal{C} [12,13].

Now let \mathcal{S} be 0–1 constraint set $\{Ax \geq b, x \in \{0, 1\}^n\}$, where the domains are $D_j = \{0, 1\}$ for all j . \mathcal{S} *logically implies* 0–1 constraint set \mathcal{S}' when all 0–1 points that satisfy \mathcal{S} also satisfy \mathcal{S}' . \mathcal{S} and \mathcal{S}' are *logically equivalent* when they logically imply each other. A logical clause $\bigvee_{j \in J^+} x_j \vee \bigvee_{j \in J^-} \neg x_j$ is *represented* by the 0–1 inequality

$$\sum_{j \in J^+} x_j + \sum_{j \in J^-} (1 - x_j) \geq 1$$

A 0–1 inequality is *clausal* when it represents a clause. It is clear that a 0–1 inequality is logically equivalent to the set of clausal inequalities it implies. Thus if we let \mathcal{S}_C be the set of clausal inequalities that are implied by some inequality in \mathcal{S} , then \mathcal{S} is logically equivalent to \mathcal{S}_C . It is shown in [8] that resolution on clausal inequalities achieves consistency.

The following example illustrates how a traditional cutting plane can serve the dual purpose of tightening the linear programming (LP) relaxation and achieving consistency. Let the LP relaxation of $\mathcal{S} = \{Ax \geq b, x \in \{0, 1\}^n\}$ be $\mathcal{S}_{LP} = \{Ax \geq b, x \in [0, 1]^n\}$.

Example 3. Suppose that \mathcal{S} is the constraint set of Example 1. In this case, \mathcal{S} and \mathcal{S}_C are identical. Resolution yields two additional clausal inequalities, $x_1 + x_2 \geq 1$ and $x_1 + x_3 \geq 1$. Adding these inequalities to \mathcal{S} achieves consistency. These inequalities are also traditional cutting planes for \mathcal{S} , in particular Chvátal-Gomory (C-G) cuts. The first cuts off two fractional vertices $(x_1, \dots, x_4) = (\frac{1}{3}, \frac{1}{3}, 0, \frac{1}{3}), (\frac{1}{2}, 0, 0, \frac{1}{2})$ of the polytope described by \mathcal{S}_{LP} , and the second cuts off

the vertex $(\frac{1}{2}, \frac{1}{2}, 0, 0)$ as well. The inequalities therefore serve the dual purpose of achieving consistency and tightening the LP relaxation. As it happens, adding both resolvents yields an integral polytope, but we will see that a consistent constraint does not in general describe an integral polytope.

A special class of resolution proofs, namely input proofs, derive all clausal C–G cuts for \mathcal{S}_C [7].

5 LP-consistency

While resolution can always achieve consistency, it is not a practical method for the reduction of backtracking. Resolution proofs tend to explode rapidly in length and complexity. However, the LP relaxation of \mathcal{S} provides an additional tool for this purpose. Specifically, it provides a more useful test for consistency than whether a partial assignment violates a constraint.

Consistency of \mathcal{S} implies that any partial assignment $x_J = v_J$ that is consistent with \mathcal{S}_J (i.e., violates no constraint in \mathcal{S}) is consistent with \mathcal{S} . We want a type of consistency that ensures that any partial assignment consistent with \mathcal{S}_{LP} is consistent with \mathcal{S} . We can achieve this by defining consistency with respect to the LP relaxation \mathcal{S}_{LP} rather than the relaxation \mathcal{S}_J . Recall that classical consistency is defined so that $D_J(\mathcal{S}_J) = D(\mathcal{S})|_J$. We therefore define *LP-consistency* as follows.

Definition 5. A 0–1 constraint set \mathcal{S} is LP-consistent if $D_J(\mathcal{S}_{LP}) = D(\mathcal{S})|_J$ for all $J \subseteq N$.

Note that $D_J(\mathcal{S}_{LP})$ refers to the set of 0–1 assignments to x_J that are consistent with \mathcal{S}_{LP} , since the domains are $D_j = \{0, 1\}$ for all j . Thus \mathcal{S} is *LP-consistent* if $\mathcal{S}_{LP} \cup \{x_J = v_J\}$ is infeasible for any 0–1 partial assignment $x_J = v_J$ that is inconsistent with \mathcal{S} .

Example 4. Consider the 0–1 constraint set $\mathcal{S} = \{4x_1 + 4x_2 \geq 1, 2x_1 - 4x_2 \geq -3, x \in \{0, 1\}^2\}$ (Fig. 1). The partial assignment $x_1 = 0$ is consistent with \mathcal{S}_{LP} but not with \mathcal{S} , because both $(x_1, x_2) = (0, 0)$ and $(x_1, x_2) = (0, 1)$ violate \mathcal{S} . So \mathcal{S} is not LP-consistent.

Two elementary properties of LP-consistency follow.

Proposition 3. A consistent 0–1 constraint set is LP-consistent.

Proof. Consider any 0–1 partial assignment $x_J = v_J$ that is consistent with \mathcal{S}_{LP} . We claim that $x_J = v_J$ is consistent with \mathcal{S} , which suffices to show that \mathcal{S} is LP-consistent. Since $\mathcal{S}_{LP} \cup \{x_J = v_J\}$ is feasible, $x_J = v_J$ violates no constraints in \mathcal{S} . Now since \mathcal{S} is consistent, this means that $x_J = v_J$ is consistent with \mathcal{S} , as claimed. \square

In addition, a 0–1 constraint set that describes the integer hull (the convex hull of feasible 0–1 points) is LP-consistent.




Fig. 1: Illustration of Example 4.

Proposition 4. *Given 0–1 constraint set \mathcal{S} , if \mathcal{S}_{LP} describes the integer hull of $D(\mathcal{S})$, then \mathcal{S} is LP-consistent.*

Proof. Suppose that $\mathcal{S} \cup \{x_J = v_J\}$ is infeasible for a given 0–1 partial assignment $x_J = v_J$. Then $x_J = v_J$ describes a face of the unit hypercube that is disjoint from $D(\mathcal{S})$. This implies that the face is disjoint from the convex hull of $D(\mathcal{S})$, which is described by \mathcal{S}_{LP} . Thus $\mathcal{S}_{\text{LP}} \cup \{x_J = v_J\}$ is infeasible, and it follows that \mathcal{S} is LP-consistent. \square

It is essential to observe that a convex hull model is not necessary to achieve LP-consistency, a fact that will be exploited in later sections. This can be seen in an example.

Example 5. Consider the following two constraint sets (Fig. 2), which have the same feasible set:

$$\begin{aligned} \mathcal{S}^1 &= \{x_1 + x_2 \leq 1, x_2 + x_3 \leq 1, x \in \{0, 1\}^3\} \\ \mathcal{S}^2 &= \{x_1 + 2x_2 + x_3 \leq 2, x \in \{0, 1\}^3\} \end{aligned}$$

The LP relaxation $\mathcal{S}_{\text{LP}}^1$ describes the integer hull of $D(\mathcal{S}^1) = D(\mathcal{S}^2)$, and so \mathcal{S}^1 is LP-consistent by Proposition 4. Yet the constraint set \mathcal{S}^2 is also LP-consistent, even though $\mathcal{S}_{\text{LP}}^2$ does not describe the integer hull, but describes a polytope with fractional extreme points $(x_1, x_2, x_3) = (0, \frac{1}{2}, 1), (1, \frac{1}{2}, 0)$. Interestingly, the inequality $x_1 + 2x_2 + x_3 \geq 2$ in \mathcal{S}^2 is the sum of the two nontrivial facet-defining inequalities in \mathcal{S}^1 and is therefore weaker than either of them from a polyhedral point of view. Yet it cuts off more infeasible 0–1 points than either of the facet-defining inequalities and is therefore stronger in this sense. Indeed, the purpose of achieving LP-consistency is to cut off infeasible 0–1 (partial) assignments, not to cut off fractional vertices of the LP relaxation.

6 Characterizing LP-Consistency

The following result gives a necessary condition for consistency based on clausal inequalities.




Fig. 2: Illustration of Example 5

Proposition 5. *If a constraint set \mathcal{S} is consistent, then all of its implied clausal inequalities are in \mathcal{S}_C .*

Proof. Suppose that \mathcal{S} is consistent, and let C be any clausal inequality implied by \mathcal{S} . Then the assignment $x_J = v_J$ violates C , where x_J are the variables in C and v_j is 1 when x_j is negated in C and 0 otherwise. This means $x_J = v_J$ is inconsistent with \mathcal{S} , which implies by the consistency of \mathcal{S} that $x_J = v_J$ violates an inequality $\alpha x \geq \beta$ in \mathcal{S} . As a result, C must be implied by $\alpha x \geq \beta$, showing that $C \in \mathcal{S}_C$. \square

LP-consistency allows us to derive a stronger argument on the relation between an LP-consistent set and its implied clausal inequalities, as it provides both necessary and sufficient conditions. In particular, a 0–1 constraint set \mathcal{S} is LP-consistent if and only if all of its implied clauses are C-G cuts for \mathcal{S}_{LP} . This is due to the following fact.

Proposition 6. *Given a 0–1 constraint set \mathcal{S} , a 0–1 partial assignment is consistent with \mathcal{S}_{LP} if and only if the assignment violates no clausal C-G cut for \mathcal{S}_{LP} .*

Proof. It suffices to show that a given 0–1 partial assignment $x_J = v_J$ violates a clausal C-G for \mathcal{S}_{LP} if and only if $\mathcal{S}_{LP} \cup \{x_J = v_J\}$ is infeasible. Suppose first that $x_J = v_J$ violates a clausal inequality $ax \geq \beta$ that is a C-G cut for \mathcal{S}_{LP} , where \mathcal{S}_{LP} is the system $Ax \geq b$. Since $x_J = v_J$ violates $ax \geq \beta$, we can write the inequality as $a_J x_J \geq \beta$, where $a_J v_J \leq \beta - 1$. Now since $ax \geq \beta$ is a C-G cut, there is a tuple $u \geq 0$ of multipliers such that $uA = a$ and $\beta - 1 < ub \leq \beta$. We therefore have $(uA)_J v_J = a_J v_J \leq \beta - 1 < ub$. This implies that $x_J = v_J$ violates $uAx \geq ub$, and so $\mathcal{S}_{LP} \cup \{x_J = v_J\}$ must be infeasible.

For the converse, suppose that $\mathcal{S}_{LP} \cup \{x_J = v_J\}$ is infeasible, which means that the face of the unit hypercube defined by $x_J = v_J$ lies outside the polytope defined by \mathcal{S}_{LP} . Let $J^+ = \{j \in J \mid v_j = 0\}$ and $J^- = \{j \in J \mid v_j = 1\}$.

Then some inequality of the form $\sum_{j \in J^+} x_j + \sum_{j \in J^-} (1 - x_j) \geq \bar{\pi}$ for some $\bar{\pi} > 0$ separates the face just mentioned from the polytope; i.e., $x_J = v_J$ violates this inequality. Since this inequality is valid for \mathcal{S}_{LP} , it is dominated by some surrogate of $Ax \geq b$. That is there exists a tuple $u \geq 0$ of multipliers such that $uA \geq ub$ is of the form

$$\sum_{j \in J^+} x_j + \sum_{j \in J^-} (1 - x_j) \geq \pi \quad (1)$$

where $\pi \geq \bar{\pi}$, and $\pi \leq |J|$ because \mathcal{S} is feasible. Now pick any subset $\hat{J} \subseteq J$ with $|\hat{J}| = \lceil \pi \rceil - 1$, let $\hat{J}^+ = J^+ \cap \hat{J}$, and let $\hat{J}^- = J^- \cap \hat{J}$. Take the sum of (1) with $-x_j \geq -1$ for $j \in \hat{J}^+$ and $x_j \geq 0$ for $j \in \hat{J}^-$. This yields a clausal inequality that is a surrogate of $Ax \geq b$:

$$\sum_{j \in J^+ \setminus \hat{J}^+} x_j + \sum_{j \in J^- \setminus \hat{J}^-} (1 - x_j) \geq 1 + \pi - \lceil \pi \rceil$$

Rounding up the right-hand side (if necessary) yields a clausal C-G cut violated by $x_J = v_J$. Thus $x_J = v_J$ violates a clausal C-G cut for \mathcal{S}_{LP} , as claimed. \square

Example 6. Consider again the constraint set \mathcal{S} of Example 3. The partial assignment $(x_1, x_3) = (0, 0)$ is inconsistent with \mathcal{S}_{LP} and violates a clausal C-G cut, namely $x_1 + x_3 \geq 1$. The cut is obtained by assigning multipliers $\frac{1}{4}, \frac{1}{2}, \frac{1}{4}, \frac{1}{4}, \frac{1}{2}$ to the three constraints of \mathcal{S} , $x_2 \geq 0$, and $x_3 \geq 0$, respectively. The partial assignment $(x_1, x_3) = (0, 1)$ is consistent with \mathcal{S}_{LP} and therefore violates no clausal C-G cut.

Corollary 1. *A constraint set \mathcal{S} is LP-consistent if and only if all of its implied clausal inequalities are C-G cuts for \mathcal{S}_{LP} .*

Proof. Suppose first that \mathcal{S} is LP-consistent, and let C be any clausal inequality implied by \mathcal{S} . Then the assignment $x_J = v_J$ violates C , where x_j are the variables in C and v_j is 1 when x_j is negated in C and 0 otherwise. This means $x_J = v_J$ is inconsistent with \mathcal{S} , which implies by the LP-consistency of \mathcal{S} that $x_J = v_J$ is inconsistent with \mathcal{S}_{LP} . By Proposition 6, $x_J = v_J$ violates some clausal C-G cut C' of \mathcal{S}_{LP} . Then C' must absorb C , which means C is likewise a C-G cut of \mathcal{S}_{LP} .

Conversely, suppose all clausal inequalities implied by \mathcal{S} are C-G cuts for \mathcal{S}_{LP} , and consider any partial assignment $x_J = v_J$ that is consistent with \mathcal{S}_{LP} . By Proposition 6, $x_J = v_J$ violates no clausal C-G cut of \mathcal{S}_{LP} . This means that it violates no clause implied by \mathcal{S} , which means that $x_J = v_J$ is consistent with \mathcal{S} , as desired. \square

Example 7. The constraint set \mathcal{S} of Example 1 is LP-consistent because its implied clausal inequalities are all implied by the inequalities in the set $\mathcal{S} \cup \{x_1 + x_2 \geq 1, x_1 + x_3 \geq 1\}$, and these are all C-G cuts for \mathcal{S}_{LP} .

7 LP-Consistency and Backtracking

Like full consistency in CP, full LP-consistency is difficult to achieve. We therefore follow the lead of the CP community and consider a weaker form of consistency, namely an analog of k -consistency. While even k -consistency is hard to achieve in practice, and the CP community focuses on domain consistency instead, a form of LP-consistency analogous to sequential k -consistency may offer possibilities to 0–1 programming.

Recall that \mathcal{S} is sequentially k -consistent if $D_{J_{k-1}}(\mathcal{S}_{J_{k-1}}) = D_{J_k}(\mathcal{S}_{J_k})|_{J_{k-1}}$, and that sequential k -consistency for $k = 1, \dots, n$ suffices to avoid backtracking when the branching order is x_1, \dots, x_n . A parallel definition that relates to linear programming is as follows.

Definition 6. *A 0–1 constraint set \mathcal{S} is sequentially LP k -consistent if $D_{J_{k-1}}(\mathcal{S}_{\text{LP}}) = D_{J_k}(\mathcal{S}_{\text{LP}})|_{J_{k-1}}$.*

Equivalently, we can say that \mathcal{S} is sequentially LP k -consistent if for every 0–1 partial assignment $x_{J_{k-1}} = v_{J_{k-1}}$ that is consistent with \mathcal{S}_{LP} , there is a 0–1 assignment $x_k = v_k$ for which $x_{J_k} = v_{J_k}$ is consistent with \mathcal{S}_{LP} . Thus sequential LP k -consistency is analogous to sequential k -consistency but based on the \mathcal{S}_{LP} relaxation rather than the $\mathcal{S}_{J_{k-1}}$ relaxation.

This form of consistency can also allow us to avoid backtracking, if we are willing to solve appropriate LP problems. Specifically, suppose that at a given node in the branching tree, prior branching has fixed $(x_1, \dots, x_{k-1}) = (v_1, \dots, v_{k-1})$. For the next branch, we select a value $v_k \in \{0, 1\}$ for which the partial assignment $(x_1, \dots, x_k) = (v_1, \dots, v_k)$ is consistent with \mathcal{S}_{LP} ; that is, for which the LP problem $\mathcal{S}_{\text{LP}} \cup \{(x_1, \dots, x_k) = (v_1, \dots, v_k)\}$ is feasible. We then set $x_k = v_k$ and continue to the next level of the tree. The following theorem guarantees that the LP problem will be feasible for at least one value of v_k , and that this process avoids backtracking.

Proposition 7. *If \mathcal{S} is a feasible 0–1 constraint set over x and the branching order is x_1, \dots, x_n , achieving sequential LP k -consistency for $k = 1, \dots, n$ suffices to solve \mathcal{S} without backtracking.*

Proof. Since \mathcal{S} is feasible, \mathcal{S}_{LP} is feasible at the root node of the branching tree, and so the empty assignment is consistent with \mathcal{S}_{LP} . Arguing inductively, suppose the partial assignment $(x_1, \dots, x_{k-1}) = (v_1, \dots, v_{k-1})$ that reflects the branching decisions down to the node at level k is consistent with \mathcal{S}_{LP} . Since \mathcal{S} is sequentially LP k -consistent, there exists a 0–1 value v_k of x_k for which the partial assignment $(x_1, \dots, x_k) = (v_1, \dots, v_k)$ is consistent with \mathcal{S}_{LP} . By induction, $\mathcal{S}_{\text{LP}} \cup \{(x_1, \dots, x_n) = (v_1, \dots, v_n)\}$ is feasible at the terminal node of the tree for some tuple (v_1, \dots, v_n) of 0–1 values. But in this case, $(x_1, \dots, x_n) = (v_1, \dots, v_n)$ satisfies \mathcal{S} , and we have solved the problem without backtracking. \square

Example 8. Consider the constraint set \mathcal{S} of Example 4. \mathcal{S} is not sequentially LP 2-consistent because $x_1 = 0$ is consistent with \mathcal{S}_{LP} , but neither $(x_1, x_2) =$

$(0, 0)$ nor $(x_1, x_2) = (0, 1)$ is consistent with \mathcal{S}_{LP} . Also, backtracking is possible, because if we set $x_1 = 0$ at the root node because $x_1 = 0$ is consistent with \mathcal{S}_{LP} , we cannot find a consistent value for x_2 at the child node and must backtrack. Now suppose we add the clause $x_1 + x_2 \geq 1$ to \mathcal{S} to obtain a constraint set \mathcal{S}' that is sequentially LP 2-consistent (Fig. ??). At the root node we must branch on $x_1 = 1$, because $x_1 = 0$ is not consistent with \mathcal{S}'_{LP} . At the child node, branching on $x_2 = 1$ yields an assignment $(x_1, x_2) = (1, 1)$ that is consistent with \mathcal{S}_{LP} and, in fact, solves \mathcal{S} without backtracking.

8 Achieving LP Consistency

We can achieve sequential LP k -consistency by using one step of a modified lift-and-project method [2]. Given $\mathcal{S} = \{Ax \geq b, x \in \{0, 1\}^n\}$ where $0 \leq x_i \leq 1$ is included in $Ax \geq b$, we generate the nonlinear system

$$\begin{aligned} (Ax - b)x_k &\geq 0 \\ (Ax - b)(1 - x_k) &\geq 0 \end{aligned}$$

We next linearize the system by replacing each x_k^2 with x_k , and each product $x_i x_k$ with y_{ik} . Let the resulting system be $R_k(\mathcal{S}_{\text{LP}})$. Adding the constraints in this system to \mathcal{S}_{LP} yields a sequentially LP k -consistent constraint set.

Proposition 8. *Given a 0–1 constraint set \mathcal{S} , augmenting \mathcal{S} with the constraints in $R_k(\mathcal{S}_{\text{LP}})$ yields a constraint set that is sequentially LP k -consistent.*

Proof. For a given 0–1 partial assignment $x_J = v_J$, suppose that $\mathcal{S}_{\text{LP}} \cup \{(x_{J_k}) = (v_{J_k})\}$ is infeasible for $v_k = 0, 1$. It suffices to show that $R_k(\mathcal{S}_{\text{LP}})|_N \cup \{x_{J_{k-1}} = v_{J_{k-1}}\}$ is infeasible. It follows from Theorem 2.1 in [2] that $R_k(\mathcal{S}_{\text{LP}})|_N$ describes the convex hull of the union of $D(\mathcal{S}_{\text{LP}} \cup \{x_k = v_k\})$ over $v_k = 0, 1$. We claim that $x_{J_{k-1}} = v_{J_{k-1}}$ does not satisfy $R_k(\mathcal{S}_{\text{LP}})|_N$. Assume to the contrary. Then there exists a point $w = (v_{J_{k-1}}, \tilde{v}_k, \tilde{v}_K, \tilde{y})$ that satisfies $R_k(\mathcal{S}_{\text{LP}})$, where $K = N \setminus J_k$. This point must be representable as a convex combination of two points of the form $(v_{J_{k-1}}, 0, \tilde{v}_K, \tilde{y})$ and $(v_{J_{k-1}}, 1, \tilde{v}_K, \tilde{y})$, since the components of $v_{J_{k-1}}$ are integral and cannot be represented as the convex combination of other points. However, by assumption such points do not exist because $\mathcal{S}_{\text{LP}} \cup \{x_{J_k} = v_{J_k}\}$ is infeasible for $v_k = 0, 1$. This yields the desired contradiction. \square

If desired, $R_k(\mathcal{S}_{\text{LP}})$ can be projected onto x , before adding its constraints to \mathcal{S}_{LP} , to obtain a sequentially LP k -consistent system in the space of original variables. Alternatively, $R_k(\mathcal{S}_{\text{LP}})$ can be projected onto $x_{J_{k-1}}$ to obtain sparse cuts that are nonetheless sufficient to achieve sequential LP k -consistency. Every partial assignment $x_{J_{k-1}} = v_{J_{k-1}}$ that is inconsistent with \mathcal{S}_{LP} violates some individual cut in $R_k(\mathcal{S}_{\text{LP}})|_{J_{k-1}}$.

Example 9. Consider again Example 4, in which

$$\mathcal{S} = \{2x_1 - 4x_2 \geq -3, 4x_1 + 4x_2 \geq 1, x_1, x_2 \in \{0, 1\}\}$$

Recall that \mathcal{S} is not LP 2-consistent because $x_1 = 0$ is consistent with \mathcal{S}_{LP} and $(x_1, x_2) = (0, v_2)$ is inconsistent with \mathcal{S}_{LP} for $v_2 = 0, 1$. We wish to achieve sequential LP 2-consistency by applying the modified lift-and-project procedure. First generate the constraints

$$\begin{aligned} (2x_1 - 4x_2 + 3)x_2 &\geq 0 & x_1x_2 &\geq 0 \\ (2x_1 - 4x_2 + 3)(1 - x_2) &\geq 0 & x_1(1 - x_2) &\geq 0 \\ (4x_1 + 4x_2 - 1)x_2 &\geq 0 & (1 - x_1)x_2 &\geq 0 \\ (4x_1 + 4x_2 - 1)(1 - x_2) &\geq 0 & (1 - x_1)(1 - x_2) &\geq 0 \end{aligned}$$

After linearizing and writing y_{12} simply as y , we obtain the system $R_2(\mathcal{S}_{\text{LP}})$:

$$\begin{aligned} -x_2 + 2y &\geq 0 & y &\geq 0 \\ 2x_1 - 3x_2 - 2y + 3 &\geq 0 & x_1 - y &\geq 0 \\ 3x_2 + 4y &\geq 0 & x_2 - y &\geq 0 \\ 4x_1 + x_2 - 4y - 1 &\geq 0 & -x_1 - x_2 + y + 1 &\geq 0 \end{aligned} \tag{2}$$

The third constraint on the left can be omitted because it is implied by $x_2, y \geq 0$. Adding the constraints in (2) to \mathcal{S}_{LP} yields a sequentially LP 2-consistent set, and it is clear on inspection that $x_1 = 0$ is inconsistent with (2). If we wish to obtain a consistent constraint set in the original variables, we can project (2) onto (x_1, x_2) . This yields a constraint $4x_1 - x_2 \geq 1$ that can be added to \mathcal{S}_{LP} to obtain a sequentially LP 2-consistent set, as illustrated in Fig. 3(a). It is evident in the figure that $x_1 = 0$ is inconsistent with this set. Finally, we can obtain a sparse cut that achieves sequential LP 2-consistency by projecting (2) onto x_1 . This yields the cut $x_1 \geq \frac{1}{4}$, which likewise excludes $x_1 = 0$.

An advantage of sequential LP-consistency is that it can avoid branching that traditional cutting planes do not avoid, because it focuses on excluding inconsistent partial assignments, rather than on tightening the LP relaxation by




Fig. 3: Illustration of Examples 9 and 10.

cutting off fractional points. This can be illustrated in a very simple context as follows.

Example 10. Suppose we wish to maximize $3x_2 - x_1$ subject to the constraint set \mathcal{S} in the previous example. We first apply a traditional branch-and-cut procedure that generates separating lift-and-project cuts at the root node (Fig. 4(a)). The solution of the LP relaxation at the root node is $(x_1, x_2) = (\frac{1}{2}, 1)$. Lift and project yields the cuts $x_1 - 4x_2 \geq -3$ and $x_1 + 4x_2 \geq 1$ (corresponding to the disjunction $x_1 = 0 \vee x_1 = 1$) as illustrated in Fig. 3(b), and the cut $4x_1 - x_2 \geq 1$ (corresponding to $x_2 = 0 \vee x_2 = 1$) as illustrated in Fig. 3(a). Only the first cut is generated, because only it cuts off the fractional solution $(\frac{1}{2}, 1)$. This results in a new LP solution $(x_1, x_2) = (0, \frac{3}{4})$. The procedure then branches on the fractional variable x_2 . The $x_2 = 0$ branch yields the fractional LP solution $(x_1, x_2) = (\frac{1}{2}, 0)$, and it is necessary to branch on x_1 . The $x_2 = 1$ branch yields the integer LP solution $(x_1, x_2) = (1, 1)$, which solves the problem. The resulting search tree has 5 nodes.

Suppose now that we achieve sequential LP 2-consistency as described in Example 9 by generating the inequality $4x_1 - x_2 \geq 1$, even though it does not cut off the fractional LP solution (Fig. 4(b)). Since the partial assignment $x_1 = 0$ is inconsistent with the LP relaxation, we immediately branch on $x_1 = 1$, which yields the integer LP solution $(x_1, x_2) = (1, 1)$. The problem is solved with only 2 nodes in the search tree, even though we used no traditional separating cuts at all.




Fig. 4: Illustration of Example 10

One could, in principle, avoid backtracking altogether by applying lift-and-project repeatedly to achieve sequential LP k -consistency for $k = 1, \dots, n$. This is impractical, however, because the resulting constraint set explodes in size. An alternative is to achieve k -consistency for a few small values of k . This can be

accomplished in three ways. (a) Obtain k -consistency by applying the lift-and-project step to the set \mathcal{S}_{LP} obtained from computing $(k - 1)$ -consistency in the previous step, where the procedure is modified to account for the y -variables in \mathcal{S}_{LP} . This causes the number of variables to double in each step, but no projection is required. (b) Project $R_k(\mathcal{S}_{\text{LP}})$ onto x before moving to the next step. The number of variables remains constant, but a time-consuming projection operation must be carried out. (c) Project $R_k(\mathcal{S}_{\text{LP}})$ onto $x_{J_{k-1}}$ before moving to the next step. This adds only sparse cuts to the constraint set but requires more computation to carry out the projection.

These methods become computationally prohibitive in an IP solver as k increases, unless a heuristic is used to identify generated inequalities that are likely to play a role in achieving sequential LP k -consistency—much as separation algorithms are used to identify useful cutting planes. This remains an issue for future research.

9 Conclusion

We provided a theoretical foundation for a new type of consistency, LP-consistency, that is particularly suited to 0–1 programming. It is based on the idea that consistency can, in general, be defined with respect to a type of relaxation. LP-consistency is obtained by replacing the relaxation used for traditional consistency concepts with the LP relaxation. It brings a novel approach to 0–1 programming by identifying cuts that exclude infeasible partial assignments rather than fractional solutions. To our knowledge, no such concept has been proposed in the IP literature, even though it is directly relevant to the amount of backtracking that occurs. We also showed how a non-facet-defining inequality can be stronger than a facet-defining inequality in an interesting sense, and how traditional cutting planes can reduce branching even if no LP relaxation is used, because they can help achieve consistency.

We also defined sequential LP k -consistency, a weaker form of LP-consistency that nonetheless reduces backtracking. Sequential LP k -consistency for a given k can be obtained by one step of the lift-and-project process of integer programming. We showed that achieving even sequential LP 2-consistency can avoid backtracking that traditional separating cuts allow.

This work points to at least three further research programs. One is to extend the concepts introduced here to general mixed integer/linear programming (MILP), which appears to be straightforward. A second is to investigate the computational usefulness of sequential LP k -consistency for MILP solvers, in particular by achieving sequential LP k -consistency for small k near the top of the search tree. A third is to conduct a systematic study of the ability of traditional cutting planes to achieve consistency, both traditional forms and LP-consistency, in an MILP problem. This could allow one to make better use of known cutting planes by generating cuts that do not separate fractional solutions but enhance the consistency properties of the constraint set.

References

1. Apt, K.R.: Principles of Constraint Programming. Cambridge University Press, Cambridge, UK (2003)
2. Balas, E., Ceria, S., Cornuéjols, G.: A lift-and-project cutting plane algorithm for mixed 0-1 programs. *Mathematical Programming* **58**, 295–324 (1993)
3. Chvátal, V.: Edmonds polytopes and a hierarchy of combinatorial problems. *Discrete Mathematics* **4**, 305–337 (1973)
4. Davis, E.: Constraint propagation with intervals labels. *Artificial Intelligence* **32**, 281–331 (1987)
5. Freuder, E.C.: Synthesizing constraint expressions. *Communications of the ACM* **21**, 958–966 (1978)
6. Freuder, E.C.: A sufficient condition for backtrack-free search. *Communications of the ACM* **29**, 24–32 (1982)
7. Hooker, J.N.: Input proofs and rank one cutting planes. *ORSA Journal on Computing* **1**, 137–145 (1989)
8. Hooker, J.N.: *Integrated Methods for Optimization*, 2nd ed. Springer (2012)
9. Hooker, J.N.: Projection, consistency, and George Boole. *Constraints* **21**, 59–76 (2016)
10. Mackworth, A.: Consistency in networks of relations. *Artificial Intelligence* **8**, 99–118 (1977)
11. Montanari, U.: Networks of constraints: Fundamental properties and applications to picture processing. *Information Science* **7**, 95–132 (1974)
12. Quine, W.V.: The problem of simplifying truth functions. *American Mathematical Monthly* **59**, 521–531 (1952)
13. Quine, W.V.: A way to simplify truth functions. *American Mathematical Monthly* **62**, 627–631 (1955)
14. Régin, J.C.: Global constraints: A survey. In: Milano, M., Van Hentenryck, P. (eds.) *Hybrid Optimization: The Ten Years of CPAIOR*, pp. 63–134. Springer, New York (2010)
15. Tsang, E.: *Foundations of Constraint Satisfaction*. Academic Press, London (1983)
16. Van Hentenryck, P.: *Constraint Satisfaction in Logic Programming*. MIT Press, Cambridge, MA (1989)