Violator Spaces: Structure and Algorithms ⋆

B. Gärtner [a], J. Matoušek [b], L. Rüst [a], P. Škovroň [b]

[a] *Institute of Theoretical Computer Science, ETH Zürich, 8092 Zürich, Switzerland*

[b] *Department of Applied Mathematics and Institute of Theoretical Computer Science, Charles University, Malostranské nám. 25, 118 00 Praha 1, Czech Republic*

Abstract

Sharir and Welzl introduced an abstract framework for optimization problems, called *LP-type problems* or also *generalized linear programming problems*, which proved useful in algorithm design. We define a new, and as we believe, simpler and more natural framework: *violator spaces*, which constitute a proper generalization of LP-type problems. We show that Clarkson's randomized algorithms for low-dimensional linear programming work in the context of violator spaces. For example, in this way we obtain the fastest known algorithm for the P-*matrix generalized linear complementarity problem* with a constant number of blocks. We also give two new characterizations of LP-type problems: they are equivalent to *acyclic* violator spaces, as well as to *concrete* LP-type problems (informally, the constraints in a concrete LP-type problem are subsets of a linearly ordered ground set, and the value of a set of constraints is the minimum of its intersection).

Key words: LP-type problem, generalized linear programming, violator space, Clarkson's algorithms, unique sink orientation, generalized linear complementarity problem

⋆ The first and the third author acknowledge support from the Swiss Science Foundation (SNF), Project No. 200021-100316/1. The fourth author acknowledges support from the Czech Science Foundation (GACR), Grant No. 201/05/H014.

Email addresses: `gaertner@inf.ethz.ch` (B. Gärtner), `matousek@kam.mff.cuni.cz` (J. Matoušek), `ruestle@inf.ethz.ch` (L. Rüst), `xofon@kam.mff.cuni.cz` (P. Škovroň).

1 Introduction

The framework of LP-type problems, invented by Sharir and Welzl in 1992 [1], has become a well-established tool in the field of geometric optimization. Its origins are in linear programming: Sharir and Welzl developed a randomized variant of the dual simplex algorithm for linear programming and showed that this algorithm actually works for a more general class of problems they called LP-type problems.

For the theory of linear programming, this algorithm constituted an important progress, since it was later shown to be *subexponential* in the RAM model [2]. Together with a similar result independently obtained by Kalai [3], this was the first linear programming algorithm provably requiring a number of arithmetic operations subexponential in the dimension and number of constraints (independent of the precision of the input numbers).

For many other geometric optimization problems in fixed dimension, the algorithm by Sharir and Welzl was the first to achieve expected linear runtime, simply because these problems could be formulated as LP-type problems. The class of LP-type problems for example includes the problem of computing the minimum-volume ball or ellipsoid enclosing a given point set in $\mathbf{R}^d$, and the problem of finding the distance of two convex polytopes in $\mathbf{R}^d$. Many other problems have been identified as LP-type problems over the years [2,4–7].

Once it is shown that a particular optimization problem is an LP-type problem, and certain algorithmic primitives are implemented for it, several efficient algorithms are immediately at our disposal: the Sharir–Welzl algorithm, two other randomized optimization algorithms due to Clarkson [8] (see [9,10] for a discussion of how it fits the LP-type framework), a deterministic version of it [10], an algorithm for computing the minimum solution that violates at most k of the given n constraints [11], and probably more are to come in the future.

The framework of LP-type problems is not only a prototype for concrete optimization problems, it also serves as a mathematical tool by itself, in algorithmic [12,13] and non-algorithmic contexts [14].

An (abstract) LP-type problem is given by a finite set H of *constraints* and a *value* $w(G)$ for every subset $G \subseteq H$. The values can be real numbers or, for technical convenience, elements of any other linearly ordered set. Intuitively, $w(G)$ is the minimum value of a solution that satisfies all constraints in G. The assignment $G \mapsto w(G)$ has to obey the axioms in the following definition.

Definition 1 *An* abstract LP-type problem *is a quadruple $(H, w, W, \leq)$, where H is a finite set, W is a set linearly ordered by $\leq$, and $w\colon 2^H \to W$ is a mapping satisfying the following two conditions:*

Monotonicity: *for all* $F \subseteq G \subseteq H$ *we have* $w(F) \leq w(G)$*, and*

Locality: *for all* $F \subseteq G \subseteq H$ *and all* $h \in H$ *with* $w(F) = w(G)$ *and* $w(G) < w(G \cup \{h\})$*, we have* $w(F) < w(F \cup \{h\})$*.*

As our running example, we will use the smallest enclosing ball problem, where H is a finite point set in $\mathbf{R}^d$ and $w(G)$ is the radius of the smallest ball that encloses all points of G. In this case monotonicity is obvious, while verifying locality requires the nontrivial but well known geometric result that the smallest enclosing ball is unique for every set.

It seems that the order $\leq$ of subsets is crucial; after all, LP-type problems model *optimization problems*, and indeed, the subexponential algorithm for linear programming and other LP-type problems [2] heavily relies on such an order.

A somewhat deeper look reveals that often, we only care whether two subsets have the *same* value, but not how they compare under the order $\leq$. The following definition is taken from [1]:

Definition 2 *Consider an abstract LP-type problem* $(H, w, W, \leq)$*. We say that* $B \subseteq H$ *is a* basis *if for all proper subsets* $F \subset B$ *we have* $w(F) \neq w(B)$*. For* $G \subseteq H$*, a* basis of G *is a minimal subset* B *of* G *with* $w(B) = w(G)$*.*

We observe that a minimal subset $B \subseteq G$ with $w(B) = w(G)$ is indeed a basis.

Solving an abstract LP-type problem $(H, w, W, \leq)$ means to find a basis of H. In the smallest enclosing ball problem, a basis of H is a minimal set B of points such that the smallest enclosing ball of B has the same radius (and is in fact the same) as the smallest enclosing ball of H, $w(B) = w(H)$.

In defining bases, and in saying what it means to solve an LP-type problem, we therefore do not need the order $\leq$. The main contribution of this paper is that many of the things one can prove about LP-type problems do not require a concept of order.

We formalize this by defining the new framework of *violator spaces.* Intuitively, a violator space is an LP-type problem without order. This generalization of LP-type problems is proper, and we can exactly characterize the violator spaces that "are" LP-type problems. In doing so, we also establish yet another equivalent characterization of LP-type problems that is closer to the applications than the abstract formulation of Definition 1. In a concrete LP-type problem, the constraints are not just elements of a set, but they are associated with subsets of some linearly ordered ground set X, with the minimal elements in the intersections of such subsets corresponding to "solutions". The framework of concrete LP-type problems is similar to the model presented in [5] as a mathematical programming problem, with a few technical differences.

These are our main findings on the structural side. Probably the most surprising insight on the algorithmic side is that Clarkson's algorithms [8] work for violator spaces of fixed dimension, leading to an expected linear-time algorithm for "solving" the violator space. Clarkson's algorithms were originally developed for linear programs with small dimension. They can be generalized for LP-type problems [9,10]. The fact that the scheme also works for violator spaces may come as a surprise since the structure of violator spaces is not acyclic in general (in contrast to LP-type problems). The LP-type algorithm from [2] is also applicable to violator spaces, but its analysis breaks down.

We give an application of Clarkson's algorithms in the more general setting by linking our new violator space framework to well-known abstract and concrete frameworks in combinatorial optimization. For this, we show that any *unique sink orientation* (USO) of the cube [15–24] or the more general grid [16] gives rise to a violator space, but not to an LP-type problem in general. Grid USO capture some important problems like linear programming over products of simplices, *generalized linear complementarity problems* over P-matrices [16] or games like parity, mean-payoff, and simple stochastic games [25–27].

We show that we can find the sink in a unique sink orientation by solving the violator space, for example with Clarkson's algorithms. A concrete new result is obtained by applying this to P-matrix generalized linear complementarity problems. These problems are not known to be polynomial-time solvable, but NP-hardness would imply NP=co-NP [28,16]. Since any P-matrix generalized linear complementarity problem gives rise to a unique sink orientation [16], we may use violator spaces and Clarkson's algorithms to solve the problem in expected linear time in the (polynomially solvable) case of a *fixed* number of *blocks*. This is optimal and beats all previous algorithms.

The rest of the paper is organized as follows. In Section 2, we formally define the frameworks of concrete LP-type problems and violator spaces, along with their essential terminology. Then we state our main structural result.

In Section 3, we prove this result by deriving the equivalence of abstract and concrete LP-type problems, and of *acyclic* violator spaces.

Section 4 shows that Clarkson's algorithms work for (possibly cyclic) violator spaces. Section 5, finally, shows how unique sink orientations induce violator spaces. A unique sink orientation can be cyclic, and a cyclic orientation gives rise to a cyclic violator space. Unique sink orientations are therefore nontrivial examples of possibly cyclic violator spaces.

2 Structural Results

2.1 Concrete LP-type problems.

Although intuitively one thinks about $w(G)$ as the value of an optimal solution of an optimization problem, the solution itself is not explicitly represented in Definition 1. In specific geometric examples, the constraints can usually be interpreted as a subset of some ground set X of points, and the optimal solution for G is the point with the smallest value in the intersection of all constraints in G. For example, in linear programming, the constraints are halfspaces, the value is given by the objective function, and the optimum is the point with minimum value in the admissible region, i.e., the intersection of the halfspaces. In order to have a unique optimum for every set of constraints (which is needed for w to define an LP-type problem), one assumes that the points are linearly ordered by the value; for linear programming, we can always take the lexicographically smallest optimal solution, for instance.

Such an interpretation is possible for the smallest enclosing ball problem too, although it looks a bit artificial. Namely, the "points" of X are all balls in $\mathbf{R}^d$, where the ordering can be an arbitrary linear extension of the partial ordering of balls by radius. The "constraint" for a point $h \in H$ is the set of all balls containing h.

The following definition captures this approach to LP-type problems.

Definition 3 *A* concrete LP-type problem *is a triple* $(X, \preceq, \mathcal{H})$*, where* X *is a set linearly ordered by* $\preceq$*,* $\mathcal{H}$ *is a finite multiset whose elements are subsets of* X*, and for any* $\mathcal{G} \subseteq \mathcal{H}$*, if the intersection* $\bigcap \mathcal{G} := \bigcap_{G \in \mathcal{G}} G$ *is nonempty, then it has a minimum element with respect to* $\preceq$ *(for* $\mathcal{G} = \emptyset$ *we define* $\bigcap \mathcal{G} := X$*).*

The definition allows $\mathcal{H}$ to be a multiset, i.e., a constraint set $A \subseteq X$ may be included several times. For example, in an instance of linear programming, some constraints can be the same, which we can reflect by this. In Subsection 3.4 we provide an example of an abstract LP-type problem, for which the multiplicity comes in handy to represent it as a concrete LP-type problem.

A similar model has been presented in [5] (mathematical programming problem). The slight difference is that it allows several points to have the same value but the constraints form a set rather than a multiset.

Bases are defined analogously to Definition 2.

Definition 4 *Consider a concrete LP-type problem* $(X, \preceq, \mathcal{H})$*. We say that* $\mathcal{B} \subseteq \mathcal{H}$ *is a* basis *if for all proper submultisets* $\mathcal{F} \subset \mathcal{B}$ *we have* $\min(\bigcap \mathcal{F}) \prec$

$\min(\bigcap \mathcal{B})$. *For* $\mathcal{G} \subseteq \mathcal{H}$*, a* basis of $\mathcal{G}$ *is a minimal* $\mathcal{B} \subseteq \mathcal{G}$ *with* $\min(\bigcap \mathcal{B}) = \min(\bigcap \mathcal{G})$.

As before, a minimal $\mathcal{B} \subseteq \mathcal{G}$ with $\min(\bigcap \mathcal{B}) = \min(\bigcap \mathcal{G})$ is indeed a basis.

Given any concrete LP-type problem $\mathcal{P} = (X, \preceq, \mathcal{H})$, we obtain an abstract LP-type problem $P = (\mathcal{H}, w, X, \preceq)$ according to Definition 1 by putting $w(\mathcal{G}) = \min(\bigcap \mathcal{G})$ (or $w(\mathcal{G}) = +\infty$, if $\bigcap \mathcal{G}$ is empty), as is easy to check (proof omitted). It is clear that $\mathcal{B} \subseteq \mathcal{G}$ is a basis of $\mathcal{G}$ in P if and only if $\mathcal{B}$ is a basis of $\mathcal{G}$ in $\mathcal{P}$. We say that P is *basis-equivalent* to $\mathcal{P}$.

Somewhat surprising is the converse, which we prove below in Theorem 8: Any abstract LP-type problem $(H, w, W, \leq)$ has a "concrete representation", that is, a concrete LP-type problem that is basis-equivalent to $(H, w, W, \leq)$.

Strictly speaking, if the multiset $\mathcal{H}$ in the concrete LP-type problem has elements with multiplicity bigger than 1, then $\mathcal{H}$ cannot be used as the set of constraints for the abstract LP-type problem (since it is not a set). However, we can bijectively map $\mathcal{H}$ to a set, i.e., we take any set H with $|H| = |\mathcal{H}|$ and a mapping $f: H \to \mathcal{H}$ such that for any $\bar{h} \in \mathcal{H}$, the number of elements $h \in H$ that map to $\bar{h}$ is equal to the multiplicity of $\bar{h}$. For $G \subseteq H$ we then define $w(G) = \min(\bigcap_{g \in G} f(g))$ which gives us a fair abstract LP-type problem $P = (H, w, X, \preceq)$ basis-equivalent to $\mathcal{P}$. In this case, by basis-equivalence we mean the existence of a suitable mapping f together with the condition that $B \subseteq G$ is a basis of G in P if and only if the multiset $\{f(b): b \in B\}$ is a basis of $\{f(g): g \in G\}$ in $\mathcal{P}$.

2.2 Violator spaces.

Let $(H, w, W, \leq)$ be an abstract LP-type problem. It is natural to define that a constraint $h \in H$ *violates* a set $G \subseteq H$ of constraints if $w(G \cup \{h\}) > w(G)$. For example, in the smallest enclosing ball problem, a point h violates a set G if it lies outside of the smallest ball enclosing G (which is unique).

Definition 5 *The* violator mapping *of* $(H, w, W, \leq)$ *is defined by* $\mathsf{V}(G) = \{h \in H: w(G \cup \{h\}) > w(G)\}$. *Thus,* $\mathsf{V}(G)$ *is the set of all constraints violating* G.

It turns out that the knowledge of $\mathsf{V}(G)$ for all $G \subseteq H$ is enough to describe the "structure" of an LP-type problem. That is, while we cannot reconstruct $W, \leq$, and w from this knowledge, it is natural to consider two LP-type problems with the same mapping $\mathsf{V}: 2^H \to 2^H$ the same (isomorphic). Indeed, the algorithmic primitives needed for implementing the Sharir–Welzl algorithm and the other algorithms for LP-type problems mentioned above can be phrased in terms of testing violation (does $h \in \mathsf{V}(G)$ hold for a certain set $G \subseteq H$?), and they never deal explicitly with the values of w.

We now introduce the notion of *violator space*:

Definition 6 *A* violator space *is a pair* (H, V)*, where* H *is a finite set and* V *is a mapping* $2^H \to 2^H$ *such that*

Consistency: $G \cap \mathsf{V}(G) = \emptyset$ *holds for all* $G \subseteq H$*, and*

Locality: *for all* $F \subseteq G \subseteq H$*, where* $G \cap \mathsf{V}(F) = \emptyset$*, we have* $\mathsf{V}(G) = \mathsf{V}(F)$.

A basis of a violator space is defined in analogy to a basis of an LP-type problem.

Definition 7 *Consider a violator space* (H, V)*. We say that* $B \subseteq H$ *is a* basis *if for all proper subsets* $F \subset B$ *we have* $B \cap \mathsf{V}(F) \neq \emptyset$*. For* $G \subseteq H$*, a* basis of G *is a minimal subset* B *of* G *with* $\mathsf{V}(B) = \mathsf{V}(G)$*.*

Observe that a minimal subset $B \subseteq G$ with $\mathsf{V}(B) = \mathsf{V}(G)$ is indeed a basis: Assume for contradiction that there is a set $F \subset B$ such that $B \cap \mathsf{V}(F) = \emptyset$. Locality then yields $\mathsf{V}(B) = \mathsf{V}(F) = \mathsf{V}(G)$, which contradicts minimality of B.

We will check in Subsection 3.2 that the violator mapping of an abstract LP-type problem satisfies the two axioms above. Consistency is immediate: since $w(G) = w(G \cup \{h\})$ for $h \in G$, no element in G violates G. The locality condition has the following intuitive interpretation: adding only non-violators to a set does not change the value.

We actually show more: given an abstract LP-type problem $(H, w, W, \leq)$, the pair (H, V), with V being the violator mapping, is an *acyclic* violator space. (Acyclicity of a violator space will be defined later in Definition 10.) It turns out in Subsection 3.3 that acyclicity already characterizes the violator spaces obtained from LP-type problems, and thus any acyclic violator space can be represented as an LP-type problem (abstract or concrete). These equivalences are stated in our main theorem.

Theorem 8 *Abstract and concrete LP-type problems as well as acyclic violator spaces are equivalent (in the sense of basis-equivalence).*

The construction is illustrated on simple instances of problems of linear programming and the smallest enclosing ball in Subsection 3.4. Several more results concerning violator spaces have been achieved in the MSc. thesis of the fourth author [29].

3 Equivalence of LP-type Problems and Acyclic Violator Spaces

In this section we prove Theorem 8.

3.1 Preliminaries on Violator Spaces

To show that every acyclic violator space (H, V) originates from some concrete LP-type problem, we need an appropriate linearly ordered set X of "points", and then we will identify the elements of H with certain subsets of X.

What set X will we take? Recall that for smallest enclosing balls, X is the set of all balls, and the subset for $h \in H$ is the subset of balls containing h. It is not hard to see that we may restrict X to smallest enclosing balls of *bases*; in fact, we may choose X as the set of bases, in which case the subset for h becomes the set of bases not violated by h.

This also works for general acyclic violator spaces, with bases suitably ordered. The only blemish is that we may get several minimal bases for $G \subseteq H$; for smallest enclosing balls, this corresponds to the situation in which several bases define the same smallest enclosing ball. To address this, we will declare such bases as equivalent and choose X as the set of all equivalence classes instead.

In the following, we fix a violator space (H, V). The set of all bases in (H, V) will be denoted by $\mathcal{B}$.

Definition 9 *$B, C \in \mathcal{B}$ are* equivalent*, $B \sim C$, if $\mathsf{V}(B) = \mathsf{V}(C)$.*

Clearly, the relation $\sim$ defined on $\mathcal{B}$ is an equivalence relation. The equivalence class containing a basis B will be denoted by $[B]$.

Now we are going to define an ordering of the bases, and we derive from this an ordering of the equivalence classes as well as the notion of acyclicity in violator spaces.

Definition 10 *For $F, G \subseteq H$ in a violator space (H, V), we say that $F \leq_0 G$ (F is* locally smaller *than G) if $F \cap \mathsf{V}(G) = \emptyset$.*

For equivalence classes $[B], [C] \in \mathcal{B}/\!\sim$, we say that $[B] \leq_0 [C]$ if there exist $B' \in [B]$ and $C' \in [C]$ such that $B' \leq_0 C'$.

We define the relation $\leq_1$ on the equivalence classes as the transitive closure of $\leq_0$. The relation $\leq_1$ is clearly reflexive and transitive. If it is antisymmetric,

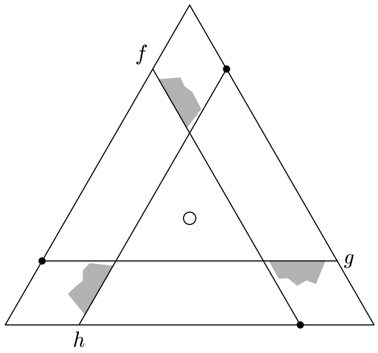

Fig. 1. A cyclic violator space.

we say that the violator space is acyclic*, and we define the relation* $\leq$ *as an arbitrary linear extension of* $\leq_1$*.*

The intuition of the *locally-smaller* notion comes from LP-type problems: if no element of F violates G, then $G \cup F$ has the same value as G (this is formally proved in Lemma 11 below), and monotonicity yields that value-wise, F is smaller than or equal to G.

Note that in the definition of $[B] \leq_0 [C]$ we do not require $B' \leq_0 C'$ to hold for *every* B' and C'. So $[B] \leq_0 [C] \leq_0 [B]$ may hold.

To show that acyclicity does not always hold, we conclude this section with an example of a cyclic violator space.

We begin with an intuitive geometric description; see Figure 1. We consider a triangle without the center point. We say that a point is "locally smaller" if it is farther clockwise with respect to the center. The constraints in our violator space are the three halfplanes f, g, h.

The locally smallest point within each halfplane is marked, and a halfplane violates a set of halfplanes if it does not contain the locally smallest point in their intersection.

Now we specify the corresponding violator space formally. We have $H = \{f, g, h\}$, and V is given by the following table:

G	$\emptyset$	f	g	h	f,g	f,h	g,h	f,g,h
$\mathsf{V}(G)$	f,g,h	h	f	g	h	g	f	$\emptyset$

This (H, V) is really a violator space, since we can easily check both consistency and locality. The bases are $\emptyset$, one-element sets, and H. We have $\{f\} \leq_0 \{h\} \leq_0 \{g\} \leq_0 \{f\}$, but none of the one-element bases are equivalent; i.e., $\leq_1$ is not antisymmetric.

In this subsection, we show that the violator mapping of an abstract LP-type problem is an acyclic violator space. To this end, we need the following two lemmas.

Lemma 11 *Consider an abstract LP-type problem* $(H, w, W, \leq)$ *with violator mapping* V*. Let* $A, B \subseteq H$*, where* B *is not violated by any* $h \in A$ *(*$A \cap \mathsf{V}(B) = \emptyset$*). Then* $w(A \cup B) = w(B)$*.*

PROOF. From monotonicity, we immediately obtain the inequality "$\geq$". The inequality "$\leq$" can be shown by induction on $|A|$. If $|A| = 1$, i.e., $A = \{h\}$, then $w(B \cup \{h\}) > w(B)$ would imply that B is violated by $h \in A$, a contradiction.

Let $|A| > 1$ and $A = A_0 \dot{\cup} \{h\}$ (disjoint union). From the induction hypothesis we have $w(B \cup A_0) = w(B)$. Now, if $w(B \cup A_0) < w(B \cup A_0 \cup \{h\})$, then by locality (for $B \cup A_0$, B and h) we get $w(B) < w(B \cup \{h\})$. This means that $h \in \mathsf{V}(B)$, and since $h \in A$ we have $h \in A \cap \mathsf{V}(B)$, a contradiction. So $w(B) = w(B \cup A_0) \geq w(B \cup A_0 \cup \{h\}) = w(B \cup A)$. We have proved $w(A \cup B) \leq w(B)$. □

Lemma 12 *Consider an abstract LP-type problem* $(H, w, W, \leq)$ *with violator mapping* V*. Then for any* $A, B \subseteq H$ *with* $\mathsf{V}(A) = \mathsf{V}(B)$ *we have* $w(A) = w(B)$*. Conversely,* $w(A) = w(B) = w(A \cup B)$ *implies* $\mathsf{V}(A) = \mathsf{V}(B)$*. In particular, if* $A \subseteq B$ *and* $w(A) = w(B)$*, then* $\mathsf{V}(A) = \mathsf{V}(B)$*.*

Note that the condition $w(A) = w(B)$ generally does not suffice for $\mathsf{V}(A) = \mathsf{V}(B)$. For example, having any H, we can define w by $w(G) = |G|$ for all $G \subseteq H$ (it can be checked that it is an abstract LP-type problem). Then any G's of the same size have the same w, however, $\mathsf{V}(G) = H \setminus G$, and so no distinct G's share the value of V. Roughly speaking, the equality $w(A) = w(B)$ may hold just "by accident". This is one way in which we can see that w by itself does not reflect the combinatorial structure of the problem in a natural way.

PROOF of Lemma 12. Let $w(A) \neq w(B)$. Without loss of generality we assume $w(A) > w(B)$ (note that here we use the linearity of the ordering $\leq$). If $A \cap \mathsf{V}(B) = \emptyset$, from Lemma 11 we would get $w(A \cup B) = w(B)$, which contradicts $w(A \cup B) \geq w(A) > w(B)$. So there necessarily exists $h \in A \cap \mathsf{V}(B)$, but since $h \in A$, we have $h \notin \mathsf{V}(A)$. So $\mathsf{V}(A) \neq \mathsf{V}(B)$.

Conversely, suppose $w(A) = w(B) = w(A \cup B)$. We want to show $\mathsf{V}(A) = \mathsf{V}(B)$, i.e., that $w(A) < w(A \cup \{h\})$ holds iff $w(B) < w(B \cup \{h\})$ holds.

By symmetry, it suffices to show only one of the implications. We assume $w(A) < w(A \cup \{h\})$. Then $w(A \cup B) = w(A) < w(A \cup \{h\}) \leq w(A \cup B \cup \{h\})$. Since $B \subseteq A \cup B$ and $w(B) = w(A \cup B)$, we may use locality, which gives $w(B) < w(B \cup \{h\})$. So the desired equivalence holds. $\square$

Proposition 13 *Consider an abstract LP-type problem $(H, w, W, \leq)$, and let V be its violator mapping. Then (H, V) is an acyclic violator space. Moreover, (H, V) is basis-equivalent to $(H, w, W, \leq)$.*

PROOF. Clearly $G \cap \mathsf{V}(G) = \emptyset$, since $w(G \cup \{g\}) = w(G)$ for any $g \in G$, so consistency holds. If $G \cap \mathsf{V}(F) = \emptyset$ for $F \subseteq G$, then by Lemma 11 we get $w(F \cup G) = w(F)$. Since $F \subseteq G$, we have $F \cup G = G$ and so $w(G) = w(F)$. Lemma 12 then yields $\mathsf{V}(G) = \mathsf{V}(F)$, so locality holds.

We proceed to prove acyclicity of (H, V). Fix $[B]$ and $[C]$, $[B] \neq [C]$, with $[B] \leq_0 [C]$, that is $B' \cap \mathsf{V}(C') = \emptyset$ for some $B' \in [B]$ and $C' \in [C]$. Lemma 11 implies $w(C') = w(B' \cup C')$. For contradiction, assume $w(B') \geq w(C')$; then $w(B') \geq w(B' \cup C')$ which with monotonicity yields $w(B') = w(B' \cup C') = w(C')$. Lemma 12 gives $\mathsf{V}(B') = \mathsf{V}(C')$, a contradiction to $[B] \neq [C]$. Thus $[B] \leq_0 [C]$ for $[B] \neq [C]$ implies $w(B') < w(C')$ for some bases B' and C' out of the respective equivalence classes. By Lemma 12, $w(B')$ is the same for all $B' \in [B]$ (because all bases in $[B]$ have the same violators). Therefore, by chaining several $\leq_0$'s we also get $w(B') < w(C')$ for $[B] \leq_1 [C]$. This proves that $\leq_1$ is necessarily antisymmetric (since $\leq$ is an ordering of W).

Finally, observe that by Lemma 12, $B \subseteq G$ is an inclusion-minimal subset of G with $w(B) = w(G)$ if and only if B is an inclusion-minimal subset of G with $\mathsf{V}(B) = \mathsf{V}(G)$. So, $B \subseteq G$ is a basis of G in $(H, w, W, \leq)$ if and only if B is a basis of G in (H, V). Thus (H, V) is basis-equivalent to $(H, w, W, \leq)$. $\square$

At first glance, one might think that for $F \subseteq G$ we should have $\mathsf{V}(F) \supseteq \mathsf{V}(G)$. Unfortunately, this is not the case, as the linear programming example in Figure 2 shows (the y-coordinate is to be minimized).

We put $F = \{h_1, h_2\}$ and $G = \{h_1, h_2, h_3\} \supseteq F$. The point 1 is minimum in the intersection of F, and 2 is minimum in the intersection of G. We have $1 \in h^*$, $2 \notin h^*$, and so $h^* \notin \mathsf{V}(F)$ and $h^* \in \mathsf{V}(G)$.

3.3 Acyclic Violator Spaces yield Concrete LP-type Problems

The following proposition is the last ingredient for Theorem 8.

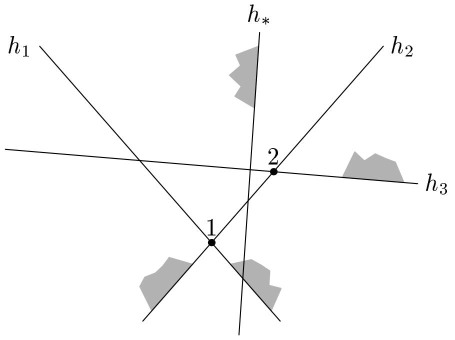

Fig. 2. A linear programming example ($F = \{h_1, h_2\} \subseteq G = \{h_1, h_2, h_3\}$ with $\mathsf{V}(G) \not\subseteq \mathsf{V}(F)$).

Proposition 14 *Every acyclic violator space* (H, V) *can be represented as a concrete LP-type problem that is basis-equivalent to* (H, V)*.*

PROOF. We are given an acyclic violator space (H, V) and we define the mapping $S: H \to 2^{\mathcal{B}/\sim}$ that will act as a "concretization" of the constraints in H:

$$S(h) = \{[B]: B \in \mathcal{B},\ h \notin \mathsf{V}(B)\}.$$

Further, let $\mathcal{H}$ be the image of the mapping S taken as a multiset, i.e.,

$$\mathcal{H} = \{S(h): h \in H\}.$$

Thus, S is a bijection between H and $\mathcal{H}$. By saying that a mapping S is a bijection between a set and a multiset we mean that for any $\bar{h} \in \mathcal{H}$, the number of $h \in H$ that map to $\bar{h}$ is equal to the multiplicity of $\bar{h}$. Note that we cannot use some common properties of set bijections; for instance we have to avoid using the inverse mapping S^{-1}.

Additionally, let σ be the induced bijection of 2^H and $2^{\mathcal{H}}$ defined by $\sigma(G) = \{S(h): h \in G\}$, for $G \subseteq H$.

Now, consider the triple $(\mathcal{B}/\sim, \leq, \mathcal{H})$, where $\leq$ is an arbitrary linear extension of $\leq_1$ (such an extension exists since (H, V) is acyclic and $\leq_1$ therefore antisymmetric). This is a concrete LP-type problem: The only thing to check is the existence of a minimal element of every nonempty intersection $\bigcap \mathcal{G}$ ($\mathcal{G} \subseteq \mathcal{H}$), which is guaranteed by the linearity of $\leq$ (remember from Definition 3 that $\bigcap \mathcal{G} := \bigcap_{G \in \mathcal{G}} G$).

It remains to prove basis-equivalence, which we do with the following two lemmas.

Lemma 15 *If B is an inclusion-minimal subset of G with* $\mathsf{V}(B) = \mathsf{V}(G)$ *in* (H, V) *(that is, B is a basis of G), then* $\min(\bigcap \sigma(B)) = \min(\bigcap \sigma(G))$ *in* $(\mathcal{B}/\!\sim, \leq, \mathcal{H})$*.*

PROOF. It is clear that $[B] \in \bigcap \sigma(G)$. Therefore, showing that there is no other basis in $\bigcap \sigma(G)$ that is locally smaller than $[B]$ proves the lemma, because then $\min(\bigcap \sigma(G)) = [B] = \min(\bigcap \sigma(B))$ (the second equality holds since B is a basis of B; just replace G by B in the following proof). Assume for contradiction that a C with $[C] \neq [B]$, $[C] \in \bigcap \sigma(G)$ and $C \leq_0 [B]$ exists. By $[C] \in \bigcap \sigma(G)$ we have $G \cap \mathsf{V}(C) = \emptyset$, which is equivalent to

$$(G \cup C) \cap \mathsf{V}(C) = \emptyset,$$

and by $C \leq_0 [B]$ we have $C \cap \mathsf{V}(B) = \emptyset$ which is equivalent to (because B is a basis of G)

$$(G \cup C) \cap \mathsf{V}(B) = \emptyset.$$

Applying locality in (H, V) to these two equations tells us that $\mathsf{V}(C) = \mathsf{V}(B)$, a contradiction to $[C] \neq [B]$. $\square$

Lemma 16 *If $\sigma(B)$ is an inclusion-minimal submultiset of $\sigma(G)$ with* $\min(\bigcap \sigma(B)) = \min(\bigcap \sigma(G))$ *in* $(\mathcal{B}/\!\sim, \leq, \mathcal{H})$ *(that is, $\sigma(B)$ is a basis of $\sigma(G)$), then* $\mathsf{V}(B) = \mathsf{V}(G)$ *in* (H, V)*.*

PROOF. Let A be a basis of B, so $\mathsf{V}(A) = \mathsf{V}(B)$. Note that $[A] \in \bigcap \sigma(B)$. Let $[C] = \min(\bigcap \sigma(B))$, thus $B \cap \mathsf{V}(C) = \emptyset$ and therefore also $A \cap \mathsf{V}(C) = \emptyset$. This means that $[A] \leq_0 [C]$ from which we conclude that $[A] = [C]$. From $\min(\bigcap \sigma(G)) = [C]$ we get $G \cap \mathsf{V}(C) = \emptyset$ which is equivalent to

$$G \cap \mathsf{V}(B) = \emptyset.$$

As $\sigma(B) \subseteq \sigma(G)$ if and only if $B \subseteq G$, we can apply locality and derive $\mathsf{V}(B) = \mathsf{V}(G)$ as needed. $\square$

Lemmas 15 and 16 prove that (H, V) and $(\mathcal{B}/\!\sim, \leq, \mathcal{H})$ are basis-equivalent, in the sense that B is a basis of G in (H, V) if and only if $\sigma(B)$ is a basis of $\sigma(G)$ in $(\mathcal{B}/\!\sim, \leq, \mathcal{H})$: Starting with a basis B of G in (H, V), Lemma 15 yields $\min(\bigcap \sigma(B)) = \min(\bigcap \sigma(G))$. This $\sigma(B)$ is inclusion-minimal w.r.t. $\sigma(G)$, since otherwise Lemma 16 would yield a contradiction to the inclusion-minimality of B w.r.t. G (where we again use that $\sigma(B) \subseteq \sigma(G)$ if and only if $B \subseteq G$).

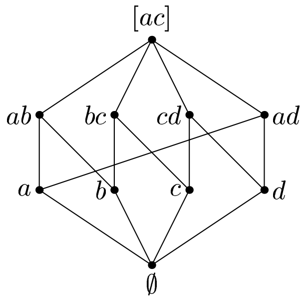

Fig. 3. Hasse diagram from smallest enclosing circle of the vertices of a square.

The reasoning for inclusion-minimality in the other direction is analogous. This concludes the proof of Proposition 14. □

Propositions 13 and 14, together with the fact that every concrete LP-type problem can be transformed into an abstract one (as described below Definition 4), yield Theorem 8.

3.4 Examples

Here we present some particular abstract LP-type problems and we demonstrate the construction (via acyclic violator spaces) of their concrete representations.

Let a, b, c and d be the vertices of a unit square (in the counterclockwise order); let $H = \{a, b, c, d\}$. For $G \subseteq H$ let $w(G)$ be the radius of the smallest circle enclosing all the points of G (for $G = \emptyset$ put $w(G) = -\infty$). The corresponding acyclic violator space is described by the following table:

G	$\emptyset$	a	b	c	d	ab	ac	ad
$\mathsf{V}(G)$	$abcd$	bcd	acd	abd	abc	cd	$\emptyset$	bc
G	bc	bd	cd	abc	abd	acd	bcd	$abcd$
$\mathsf{V}(G)$	ad	$\emptyset$	ab	$\emptyset$	$\emptyset$	$\emptyset$	$\emptyset$	$\emptyset$

The bases are $\emptyset$, a, b, c, d, ab, ac, ad, bc, bd, cd; the only equivalent pair is $ac \sim bd$. There is no inconvenience concerning differences between $\leq_0$ on sets and equivalence classes and $\leq_1$; the ordering $\leq_1$ is given by the Hasse diagram in Figure 3.

As the linear extension $\leq$ of $\leq_1$ we may choose $\emptyset < a < b < c < d < ab <$

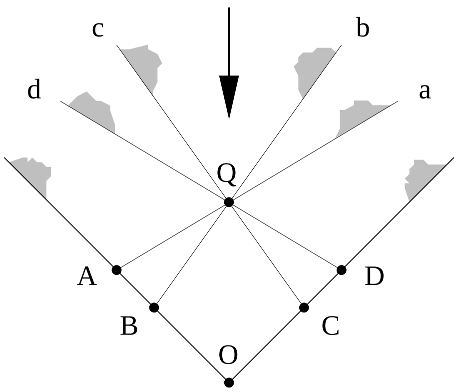

Fig. 4. Illustration example – linear programming.

$bc < cd < ad < [ac]$. Finally, the concrete representation S is as follows:

h	a	b	c	d
$S(h)$	$a, ab, ad, [ac]$	$b, ab, bc, [ac]$	$c, bc, cd, [ac]$	$d, cd, ad, [ac]$

In the geometric view that we have mentioned earlier, $S(a)$ corresponds to the set of all "canonical" (i.e., basic) balls that contain the point a (inside or on the boundary). The same holds for the other points.

As the other example, consider the following LP problem in the positive orthant (rotated by 45 degrees for convenience). Beside the restriction to the positive orthant, the constraints are the four halfplanes depicted in Figure 4. The optimization direction is given by the arrow.

Here the violator space bases are $\emptyset$, a, b, c, d, ac, ad, bc, bd; the equivalence classes are $O = \emptyset$, $A = \{a\}$, $B = \{b\}$, $C = \{c\}$, $D = \{d\}$ and $Q = \{ac\} \sim \{ad\} \sim \{bc\} \sim \{bd\}$. Note that the equivalence classes correspond to the points in the plane. We have $O \leq_1 B \leq_1 A \leq_1 Q$ and $O \leq_1 C \leq_1 D \leq_1 Q$; we choose $\leq$ to be $O < B < A < C < D < Q$. The concrete representation is

h	a	b	c	d
$S(h)$	A, Q	A, B, Q	C, D, Q	D, Q

.

Here we may interpret $S(a)$ as the set of all "canonical" points lying in the halfplane a.

To see why we allow $\mathcal{H}$ in the definition of a concrete LP-type problem to be a multiset, consider the abstract LP-type problem with $H = \{a, b\}$ and $w(G) = 0$ for every $G \subseteq H$. The only basis is $\emptyset$ and it is not violated by any $h \in H$. Thus we have $S(a) = S(b) = \{[\emptyset]\}$. If we do not allow $\mathcal{H}$ to be a multiset, we have $\mathcal{H} = \{S(a), S(b)\} = \{\{[\emptyset]\}\}$ with only one constraint; it seems improper to define this to be basis-equivalent to H. We could alter the construction of $\mathcal{H}$ and get $S(a) = \{0\}$, $S(b) = \{0, 1\}$, which does represent the original abstract LP-type problem with $\mathcal{H}$ being a set; however, we believe that our definition catches the structure in a more straightforward way, although it may seem unusual at first glance.

4 Clarkson's Algorithms

We show that Clarkson's randomized reduction scheme, originally developed for linear programs with many constraints and few variables, actually works for general (possibly cyclic) violator spaces. The two algorithms of Clarkson involved in the reduction have been analyzed for LP and LP-type problems before [8–10]; the analysis we give below is almost identical on the abstract level. Our new contribution is that the combinatorial properties underlying Clarkson's algorithms also hold for violator spaces.

We start off by deriving these combinatorial properties; the analysis of Clarkson's reduction scheme is included for completeness.

4.1 Violator spaces revisited

We recall that an abstract LP-type problem is of the form $(H, w, W, \leq)$. In this subsection we will view a violator space as an "LP-type problem without the order $\leq$", i.e., we will only care whether two subsets F and G, $F \subseteq G \subseteq H$, have the same value (and therefore the same violators, see Lemma 12), but not how they compare under the order $\leq$. It turns out that the order is irrelevant for Clarkson's algorithms.

Even without an order, we can talk about monotonicity in violator spaces:

Lemma 17 *Any violator space* (H, V) *satisfies*

Monotonicity: $\mathsf{V}(F) = \mathsf{V}(G)$ *implies* $\mathsf{V}(E) = \mathsf{V}(F) = \mathsf{V}(G)$,
for all sets $F \subseteq E \subseteq G \subseteq H$.

PROOF. Assume $\mathsf{V}(E) \neq \mathsf{V}(F), \mathsf{V}(G)$. Then locality yields $\emptyset \neq E \cap \mathsf{V}(F) = E \cap \mathsf{V}(G)$ which contradicts consistency. $\square$

Recall Definition 7: A basis is a set B satisfying $B \cap \mathsf{V}(F) \neq \emptyset$ for all proper subsets F of B. A basis of G is an inclusion-minimal subset of G with the same violators. This can be used to prove the following observation, well-known to hold for LP-type problems [9].

Observation 18 *Let (H, V) be a violator space. For $R \subseteq H$ and all $h \in H$, we have*

(i) $\mathsf{V}(R) \neq \mathsf{V}(R \cup \{h\})$ *if and only if* $h \in \mathsf{V}(R)$, *and*
(ii) $\mathsf{V}(R) \neq \mathsf{V}(R \setminus \{h\})$ *if and only if* h *is contained in every basis of* R.

An element h such that (ii) holds is called extreme *in R.*

PROOF. (i) If $h \notin \mathsf{V}(R)$, we get $\mathsf{V}(R) = \mathsf{V}(R \cup \{h\})$ by Lemma 11. If $h \in \mathsf{V}(R)$, then $\mathsf{V}(R) \neq \mathsf{V}(R \cup \{h\})$ is a consequence of consistency applied to $G = R \cup \{h\}$. (ii) if $\mathsf{V}(R) = \mathsf{V}(R \setminus \{h\})$, there is a basis B of $R \setminus \{h\}$, and this basis is also a basis of R not containing h. Conversely, if there is some basis B of R not containing h, then $\mathsf{V}(R) = \mathsf{V}(R \setminus \{h\})$ follows from monotonicity. □

We are particularly interested in violator spaces with small bases.

Definition 19 *Let (H, V) be a violator space. The size of a largest basis is called the* combinatorial dimension $\delta = \delta(H, \mathsf{V})$ *of* (H, V).

Observation 18 implies that in a violator space of combinatorial dimension δ, every set has at most δ extreme elements. This in turn yields a bound for the *expected* number of violators of a random subset of constraints, using the *sampling lemma* [12].

Lemma 20 [12] *Consider a triple (H, w, W), where w is a function mapping subsets of the set H to the set W (not necessarily ordered). For $R \subseteq H$, we define*

$$
\begin{aligned}
V(R) &:= \{h \in H \setminus R : w(R) \neq w(R \cup \{h\}), \\
X(R) &:= \{h \in R : w(R) \neq w(R \setminus \{h\}).
\end{aligned}
$$

For $0 \leq r \leq |H|$, let v_r be the expected value of $|V(R)|$, for R chosen uniformly at random among all subsets of H with r elements. x_r is defined similarly as the expected value of $|X(R)|$. Then for $0 \leq r < n$, the following equality holds.

$$
\frac{v_r}{n-r} = \frac{x_{r+1}}{r+1}.
$$

To apply this in our situation, we fix a set $W \subseteq H$, and we define $w(R) = \mathsf{V}(W \cup R)$. Since then $|X(R)| \leq \delta$ for all R, the following Corollary is obtained.

Corollary 21 *Let* (H, V) *be a violator space of combinatorial dimension* δ *and* $W \subseteq H$ *some fixed set. Let* v_r *be the expected number of violators of the set* $W \cup R$*, where* $R \subseteq H$ *is a random subset of size* $r < n = |H|$*. Then*

$$v_r \leq \delta \frac{n-r}{r+1}.$$

4.2 The Trivial Algorithm

Given a violator space (H, V) of combinatorial dimension δ, the goal is to find a basis of H. For this, we assume availability of the following primitive.

Primitive 22 *Given* $G \subseteq H$ *and* $h \in H \setminus G$*, decide whether* $h \in \mathsf{V}(G)$*.*

Given this primitive, the problem can be solved in a brute-force manner by going through all sets of size $\leq \delta$, testing each of them for being a basis of H. More generally, $B \subseteq G$ is a basis of G if and only if

$$\begin{aligned} &h \in \mathsf{V}(B \setminus \{h\}), \quad && \forall h \in B, \\ &h \notin \mathsf{V}(B), && \forall h \in G \setminus B. \end{aligned}$$

Consequently, the number of times the primitive needs to be invoked in order to find a basis of H is at most

$$n \sum_{i=0}^{\delta} \binom{n}{i} = O(n^{\delta+1}).$$

The next two subsections show that this can be substantially improved.

4.3 Clarkson's First Algorithm

Fix a violator space (H, V) of combinatorial dimension δ, implicitly specified through Primitive 22. Clarkson's first algorithm calls Clarkson's second algorithm (`Basis2`) as a subroutine. Given $G \subseteq H$, both algorithms compute a basis B of G.

`Basis1`(G):
  (* computes a basis B of G *)
  `IF` $|G| \leq 9\delta^2$ `THEN`
    `RETURN Basis2`(G)
  `ELSE`

```
    r := ⌊δ√|G|⌋
    W := ∅
    REPEAT
       choose R to be a random r-element subset of G, R ∈ (G choose r)
       C := Basis2(W ∪ R)
       V := {h ∈ G \ C: h ∈ V(C)}
       IF |V| ≤ 2√|G| THEN
          W := W ∪ V
       END
    UNTIL V = ∅
    RETURN C
END
```

Assuming `Basis2` is correct, this algorithm is correct as well: if B is a basis of $W \cup R \subseteq G$ that in addition has no violators in G, B is a basis of G. Moreover, the algorithm augments the working set W at most δ times, which is guaranteed by the following observation.

Observation 23 *If $C \subseteq G$ and $G \cap \mathsf{V}(C) \neq \emptyset$, then $G \cap \mathsf{V}(C)$ contains at least one element from every basis of G.*

PROOF. Let B be a basis of G. Assuming

$$\emptyset = B \cap G \cap \mathsf{V}(C) = B \cap \mathsf{V}(C),$$

consistency yields $C \cap \mathsf{V}(C) = \emptyset$, implying $(B \cup C) \cap \mathsf{V}(C) = \emptyset$. From locality and monotonicity (Lemma 17), we get

$$\mathsf{V}(C) = \mathsf{V}(B \cup C) = \mathsf{V}(G),$$

meaning that $G \cap \mathsf{V}(G) = G \cap \mathsf{V}(C) = \emptyset$, a contradiction. $\square$

It is also clear that `Basis2` is called only with sets of size at most $3\delta\sqrt{|G|}$. Finally, the expected number of iterations through the `REPEAT` loop is bounded by 2δ: by Corollary 21 (applied to $(G, \mathsf{V}|_G)$) and the Markov inequality, the expected number of calls to `Basis2` before we next augment W is bounded by 2.

Lemma 24 *Algorithm* `Basis1` *computes a basis of G with an expected number of at most $2\delta|G|$ calls to Primitive 22, and an expected number of at most 2δ calls to* `Basis2`*, with sets of size at most $3\delta\sqrt{|G|}$.*

This algorithm calls the trivial algorithm as a subroutine. Instead of adding violated constraints to a working set, it gives them larger probability of being selected in further iterations. Technically, this is done by maintaining G as a multiset, where $\mu(h)$ denotes the multiplicity of h (we set $\mu(F) := \sum_{h\in F}\mu(h)$). Sampling from G is done as before, imagining that G contains $\mu(h)$ copies of the element h.

```
Basis2(G):
  (* computes a basis B of G *)
  IF |G| ≤ 6δ² THEN
    RETURN Trivial(G)
  ELSE
    r := 6δ²
    REPEAT
      choose random R ∈ (G choose r)
      C := Trivial(R)
      V := {h ∈ G \ C: h ∈ V(C)}
      IF μ(V) ≤ μ(G)/3δ THEN
        μ(h) := 2μ(h),    h ∈ V
      END
    UNTIL V = ∅
    RETURN C
  END
```

Invoking Corollary 21 again (which also applies to multisets as we use them), we see that the expected number of calls to `Trivial` before we next reweight elements (a *successful* iteration), is bounded by 2. It remains to bound the number of successful iterations.

Lemma 25 *Let k be a positive integer. After $k\delta$ successful iterations, we have*

$$2^k \leq \mu(B) \leq |G|e^{k/3},$$

for every basis B of G. In particular, $k < 3\ln|G|$.

PROOF. Every successful iteration multiplies the total weight of elements in G by at most $(1 + 1/3\delta)$, which gives the upper bound (not only for $\mu(B)$ but actually for $\mu(G)$). For the lower bound, we use Observation 23 again to argue that each successful iteration doubles the weight of some element in B, meaning that after $k\delta$ iterations, one element has been doubled at least k

times. Because the lower bound exceeds the upper bound for $k \geq 3 \ln |G|$, the bound on k follows. $\square$

Summarizing, we get the following lemma.

Lemma 26 *Algorithm* `Basis2` *computes a basis of G with an expected number of at most $6\delta|G| \ln |G|$ calls to Primitive 22, and expected number of at most $6\delta \ln |G|$ calls to* `Trivial`*, with sets of size $6\delta^2$.*

4.5 Combining the Algorithms

Theorem 27 *Using a combination of the above two algorithms, a basis of H in a violator space (H, V) can be found calling Primitive 22 expected*

$$O\left(\delta n + \delta^{O(\delta)}\right)$$

many times.

PROOF. Using the above bound for the trivial algorithm, `Basis2` can be implemented to require an expected number of at most

$$O\left(\delta \log |G|(|G| + \delta^{O(\delta)})\right)$$

calls to the primitive. Applying this as a subroutine in `Basis1`(H) with $|H| = n$, $|G|$ is bounded by $3\delta\sqrt{n}$, and we get an overall expected complexity of

$$O\left(\delta n + \delta^2(\log n(\delta\sqrt{n} + \delta^{O(\delta)}))\right)$$

in terms of the number of calls to Primitive 22. The terms $\delta^2 \log n\ \delta\sqrt{n}$ and $\delta^2 \log n\ \delta^{O(\delta)}$ are asymptotically dominated by either δn or $\delta^{O(\delta)}$, and we get the simplified bound of $O\left(\delta n + \delta^{O(\delta)}\right)$. $\square$

5 Grid USO as Models for Violator Spaces

We show in this section that the problem of finding the sink in a δ-dimensional *grid unique sink orientation* [16] can be reduced to the problem of finding the (unique) basis of a violator space of combinatorial dimension δ.

Unique sink orientations of grids arise from various problems, including linear programming over products of simplices and generalized linear complementarity problems (GLCP) over P-matrices [16]. The GLCP has been introduced by Cottle and Dantzig [30] as a generalization of the well known LCP [31]. There are also applications in game theory; for instance [25–27] show how parity, mean-payoff, and simple stochastic games are related to grid USO.

5.1 Grid USO

Fix a partition

$$\Pi = (\Pi_1, \ldots, \Pi_\delta)$$

of the set $H := \{1, \ldots, n\}$ into δ nonempty subsets, where we refer to Π_i as the *block* i. A subset $J \subseteq H$ is called a *vertex* if $|J \cap \Pi_i| = 1$ for all i. The vertices naturally correspond to the Cartesian product of the Π_i. Let $\mathcal{V}$ be the set of all vertices.

In the following definition, we introduce the *grid spanned* by subsets Π'_i whose union is $G \subseteq H$. The vertex set of this grid contains all vertices $J \subseteq G$ ($J \in \mathcal{V}$), with two vertices being adjacent whenever they differ in exactly two elements.

Definition 28 *The δ-dimensional* grid *spanned by $G \subseteq H$ is the undirected graph $\mathcal{G}(G) = (\mathcal{V}(G), \mathcal{E}(G))$, with*

$$\mathcal{V}(G) := \{J \in \mathcal{V} \colon J \subseteq G\}, \quad \mathcal{E}(G) := \{\{J, J'\} \subseteq \mathcal{V}(G) \colon |J \oplus J'| = 2\}.$$

Here, $\oplus$ is the symmetric difference of sets.

$\mathcal{V}(G)$ is in one-to-one correspondence with the Cartesian product

$$\prod_{i=1}^{\delta} G_i, \quad G_i := G \cap \Pi_i,$$

and the edges in $\mathcal{E}(G)$ connect vertices in $\mathcal{V}(G)$ whose corresponding tuples differ in exactly one coordinate. See Figure 5 left for an example of a grid.

Note that $\mathcal{G}(G)$ is the empty graph whenever $G_i = G \cap \Pi_i = \emptyset$ for some i. We say that such a G is *not* Π-*valid*, and it is Π-*valid* otherwise.

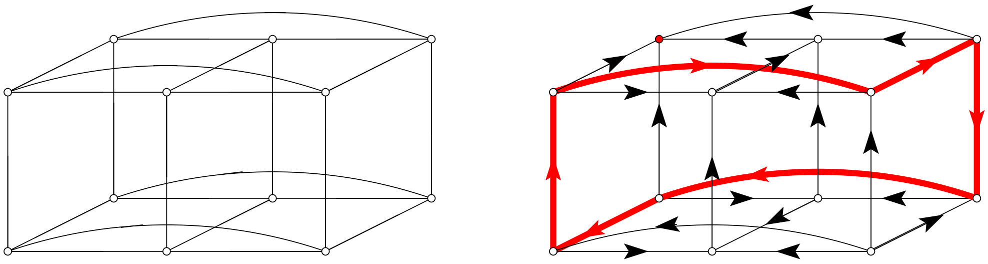

Fig. 5. A 3-dimensional grid $\mathcal{G}(H)$ with $H = \{1,\ldots,7\}$ where $\Pi = (\{1,2,3\},\{4,5\},\{6,7\})$ and a USO of it.

A *subgrid* of $\mathcal{G}(G)$ is any graph of the form $\mathcal{G}(G')$, for $G' \subseteq G$.

Definition 29 *An orientation ψ of the graph $\mathcal{G} := \mathcal{G}(H)$ is called a* unique sink orientation *(USO) if all nonempty subgrids of $\mathcal{G}$ have unique sinks w.r.t. ψ.*

We are interested in finding the sink in a USO of $\mathcal{G}$ as fast as possible, since the sink corresponds to the solution of the underlying problem (the P-matrix GLCP, for example). Our measure of complexity will be the expected number of *edge evaluations*, see [16]. An edge evaluation returns the orientation of the considered edge and can typically be implemented to run in polynomial time (depending on the underlying problem). In the remainder of this paper, we derive the following theorem.

Theorem 30 *The sink of a unique sink grid orientation can be found by evaluating expected $O\left(\delta n + \delta^{O(\delta)}\right)$ edges.*

Note that a USO ψ can be cyclic (see the thick edges in Figure 5 right). If ψ induces the directed edge (J, J'), we also write $J \xrightarrow{\psi} J'$. Any USO can be specified by associating each vertex J with its outgoing edges. Given J and $j \in H \setminus J$, we define $J \triangleright j$ to be the unique vertex $J' \subseteq J \cup \{j\}$ that is different from J, and we call J' the *neighbor* of J *in direction* j. Note that J is a neighbor of J' in some direction different from j.

Definition 31 *Given an orientation ψ of $\mathcal{G}$, the function $s_\psi : \mathcal{V} \to 2^H$, defined by*

$$s_\psi(J) := \{j \in H \setminus J \colon J \xrightarrow{\psi} J \triangleright j\}, \tag{1}$$

is called the outmap *of ψ.*

By this definition, any sink w.r.t. ψ has empty outmap value.

Let us fix a unique sink orientation ψ of $\mathcal{G}$. Given a Π-valid subset $G \subseteq H$, we define $\text{sink}(G) \in \mathcal{V}(G)$ to be the unique sink vertex in $\mathcal{G}(G)$. For a subset G that is not Π-valid, let

$$\bar{G} := \bigcup_{i:\, G_i = \emptyset} \Pi_i.$$

Thus $\bar{G}$ is the set of elements occurring in blocks of Π disjoint from G.

Definition 32 *For $G \subseteq H$, define*

$$\mathsf{V}(G) = \begin{cases} s_\psi(\text{sink}(G)), & \text{if } G \text{ is } \Pi\text{-valid} \\ \bar{G}, & \text{if } G \text{ is not } \Pi\text{-valid.} \end{cases}$$

Theorem 33 *The pair (H, V) from Definition 32 is a violator space of combinatorial dimension δ. Moreover, for all Π-valid $G \subseteq H$, the unique sink of the subgrid $\mathcal{G}(G)$ corresponds to the unique basis of G in (H, V).*

PROOF. For every $G \subseteq H$, consistency holds by definition of $\text{sink}(G), s_\psi(J)$ and $\bar{G}$. In order to prove locality for $F \subseteq G \subseteq H$, we look at three different cases.

G is not Π-valid. Then, $F \subseteq G$ is not Π-valid either. The condition $\emptyset = G \cap \mathsf{V}(F) = G \cap \bar{F}$ means that F is disjoint from the same blocks as G. This implies $\bar{G} = \bar{F}$, hence $\mathsf{V}(G) = \mathsf{V}(F)$.

G and F are both Π-valid. Then $\mathcal{G}(F)$ is a nonempty subgrid of $\mathcal{G}(G)$, and $G \cap \mathsf{V}(F) = \emptyset$ means that the sink of $\mathcal{G}(F)$ has no outgoing edges into $\mathcal{G}(G)$. Thus the unique sink of $\mathcal{G}(F)$ is also a sink of $\mathcal{G}(G)$ and therefore the unique one. This means that $\text{sink}(G) = \text{sink}(F)$, from which $\mathsf{V}(G) = \mathsf{V}(F)$ follows.

G is Π-valid, F is not Π-valid. Then the condition $G \cap \mathsf{V}(F) = \emptyset$ can never be satisfied since $\mathsf{V}(F) = \bar{F}$ contains at least one full block Π_i, and $G_i = G \cap \Pi_i \neq \emptyset$.

Next we prove that a largest basis in (H, V) has at most δ elements. For this, let $G \subseteq H$ be a set of size larger than δ. If G is Π-valid, we have

$$\mathsf{V}(G) := s_\psi(\mathrm{sink}(G)) = s_\psi(\mathrm{sink}(\mathrm{sink}(G))) =: \mathsf{V}(\mathrm{sink}(G)),$$

since $J = \mathrm{sink}(J)$ for any vertex J. This means that G has a subset of size δ with the same violators, so G is not a basis.

If G is not Π-valid, we consider some subset B that contains exactly one element from every block intersected by G. By definition, we have $\bar{G} = \bar{B}$ and $\mathsf{V}(G) = \mathsf{V}(B)$. Since B has less than δ elements, G cannot be a basis in this case, either.

It remains to prove that for G being Π-valid, the vertex $\mathrm{sink}(G)$ is the unique basis of G in (H, V). We have already shown that $\mathsf{V}(G) = \mathsf{V}(\mathrm{sink}(G))$ must hold in this case. Moreover, $\mathsf{V}(\mathrm{sink}(G))$ contains no full block Π_i. On the other hand, any proper subset F of $\mathrm{sink}(G)$ is not Π-valid, so its violator set *does* contain at least one full block. It follows that $V(F) \neq V(\mathrm{sink}(G))$, so $\mathrm{sink}(G)$ is a basis of G. The argument is complete when we can prove that no other vertex $J \subseteq G$ is a basis of G. Indeed, such a vertex J is not a sink in $\mathcal{G}(G)$, meaning that $G \cap \mathsf{V}(J) \neq \emptyset$. This implies $\mathsf{V}(J) \neq \mathsf{V}(G)$. $\square$

Note that the global sink of the grid USO corresponds to the unique δ-element (and Π-valid) set B with $\mathsf{V}(B) = \emptyset$. This is exactly the set output by the call `Basis1`(H) of Clarkson's algorithms, when we apply it to the violator space constructed in Definition 32.

Primitive 22 corresponds to one edge evaluation in the USO setting. With Theorem 27, we therefore have proved Theorem 30. For small δ, the running time given in the theorem is faster than the one from the *Product Algorithm* [16] which needs expected $O(\delta! n + H_n^\delta)$ edge evaluations, where H_n is the n-th harmonic number.

6 Conclusions

We introduced violator spaces as a new framework for optimization problems and showed that acyclic violator spaces are equivalent to abstract and concrete LP-type problems. It turned out that the explicit ordering inherent to LP-type problems is not necessary in order to capture the structure of the underlying optimization problem. Violator spaces are more general than LP-type problems, yet Clarkson's algorithms still work on them.

The Sharir-Welzl algorithm is also applicable for violator spaces in a straightforward way. However, the most obvious translation of this algorithm to the setting of violator spaces is not even guaranteed to finish, since for a general violator space it may run in a cycle and the subexponential analysis thus breaks down.

We have seen that unique sink orientations are models for possibly cyclic violator spaces, and with Clarkson's algorithms we therefore have a fast scheme to solve fixed dimensional USO problems like the generalized linear complementarity problem with a P-matrix. The GLCP with a P-matrix has in general a cyclic structure and therefore gives rise to a cyclic USO. A violator space obtained from a cyclic USO is again cyclic. It is interesting that there are no cycles in a 2-dimensional grid USO [16]. Whether the same is true for violator spaces of combinatorial dimension 2 is an open question.

Acknowledgment

The second author would like to thank Nina Amenta for discussions concerning LP-type problems, possibly already forgotten by her as they took place many years ago, but nevertheless helpful for reaching the results in this paper.

References

[1] M. Sharir, E. Welzl, A combinatorial bound for linear programming and related problems, in: Proc. 9th Symposium on Theoretical Aspects of Computer Science (STACS), Vol. 577 of Lecture Notes in Computer Science, Springer-Verlag, 1992, pp. 569–579.

[2] J. Matoušek, M. Sharir, E. Welzl, A subexponential bound for linear programming, Algorithmica 16 (1996) 498–516.

[3] G. Kalai, A subexponential randomized simplex algorithm, in: Proc. 24th Annual ACM Symposium on Theory of Computing (STOC), 1992, pp. 475–482.

[4] N. Amenta, Bounded boxes, Hausdorff distance, and a new proof of an interesting Helly-type theorem, in: Proc. 10th Annual Symposium on Computational Geometry (SCG), ACM Press, 1994, pp. 340–347.

[5] N. Amenta, Helly theorems and generalized linear programming, Discrete and Computational Geometry 12 (1994) 241–261.

[6] H. Björklund, S. Sandberg, S. Vorobyov, A discrete subexponential algorithm for parity games, in: Proc. 20th Annual Symposium on Theoretical Aspects of Computer Science (STACS), Springer-Verlag, 2003, pp. 663–674.

[7] N. Halman, Discrete and lexicographic Helly theorems and their relations to LP-type problems, Ph.D. thesis, Tel-Aviv University (2004).

[8] K. L. Clarkson, Las Vegas algorithms for linear and integer programming, Journal of the ACM 42 (1995) 488–499.

[9] B. Gärtner, E. Welzl, Linear programming - randomization and abstract frameworks, in: Proc. 13th Annual Symposium on Theoretical Aspects of Computer Science (STACS), Springer-Verlag, London, UK, 1996, pp. 669–687.

[10] B. Chazelle, J. Matoušek, On linear-time deterministic algorithms for optimization problems in fixed dimension, Journal of Algorithms 21 (1996) 579–597.

[11] J. Matoušek, On geometric optimization with few violated constraints, Discrete and Computational Geometry 14 (1995) 365–384.

[12] B. Gärtner, E. Welzl, A simple sampling lemma - analysis and applications in geometric optimization, Discrete and Computational Geometry 25 (4) (2001) 569–590.

[13] T. Chan, An optimal randomized algorithm for maximum Tukey depth, in: Proc. 15th ACM-SIAM Symposium on Discrete Algorithms (SODA), 2004, pp. 423–429.

[14] N. Amenta, A short proof of an interesting Helly-type theorem, Discrete and Computational Geometry 15 (1996) 423–427.

[15] T. Szabó, E. Welzl, Unique sink orientations of cubes, in: Proc. 42nd IEEE Symposium on Foundations of Computer Science (FOCS), 2000, pp. 547–555.

[16] B. Gärtner, W. D. Morris, Jr., L. Rüst, Unique sink orientations of grids, in: Proc. 11th Conference on Integer Programming and Combinatorial Optimization (IPCO), Vol. 3509 of Lecture Notes in Computer Science, Springer-Verlag, 2005, pp. 210–224.

[17] W. D. Morris, Jr., Randomized principal pivot algorithms for P-matrix linear complementarity problems, Mathematical Programming, Series A 92 (2002) 285–296.

[18] J. Matoušek, The number of unique sink orientations of the hypercube, Combinatorica, to appear (2006).

[19] W. D. Morris, Jr., Distinguishing cube orientations arising from linear programs, Manuscript (2002).

[20] M. Develin, LP-orientations of cubes and crosspolytopes, Advances in Geometry 4 (2004) 459–468.

[21] J. Matoušek, T. Szabó, Random Edge can be exponential on abstract cubes, in: Proc. 45th Annual IEEE Symposium on Foundations of Computer Science (FOCS), 2004, pp. 92–100.

[22] I. Schurr, T. Szabó, Finding the sink takes some time, Discrete and Computational Geometry 31 (2004) 627–642.

[23] I. Schurr, T. Szabó, Jumping doesn't help in abstract cubes, in: Proc. 11th Conference on Integer Programming and Combinatorial Optimization (IPCO), Vol. 3509 of Lecture Notes in Computer Science, Springer-Verlag, 2005, pp. 225–235.

[24] B. Gärtner, I. Schurr, Linear programming and unique sink orientations, in: Proc. 17th Annual Symposium on Discrete Algorithms (SODA), 2006, pp. 749–757.

[25] H. Björklund, S. Vorobyov, Combinatorial structure and randomized subexponential algorithms for infinite games, Theoretical Computer Science (in press).

[26] H. Björklund, S. Sandberg, S. Vorobyov, A combinatorial strongly subexponential strategy improvement algorithm for mean payoff games, in: Proc. 29th International Symposium on Mathematical Foundations of Computer Science (MFCS), Vol. 3153 of Lecture Notes in Computer Science, Springer-Verlag, 2004, pp. 673–685.

[27] B. Gärtner, L. Rüst, Simple stochastic games and P-matrix generalized linear complementarity problems, in: Proc. 15th International Symposium on Fundamentals of Computation Theory (FCT), Vol. 3623 of Lecture Notes in Computer Science, Springer-Verlag, 2005, pp. 209–220.

[28] N. Megiddo, A note on the complexity of P-matrix LCP and computing an equilibrium, Tech. rep., IBM Almaden Research Center, San Jose (1988).

[29] P. Škovroň, Generalized linear programming, Master's thesis, Charles University, Prague, Faculty of Mathematics and Physics (2002).
URL `http://kam.mff.cuni.cz/~xofon/diplomka/`

[30] R. W. Cottle, G. B. Dantzig, A generalization of the linear complementarity problem, Journal on Combinatorial Theory 8 (1970) 79–90.

[31] R. W. Cottle, J. Pang, R. E. Stone, The Linear Complementarity Problem, Academic Press, 1992.