Unique Sink Orientations of Grids ⋆

Bernd Gärtner[1], Walter D. Morris, Jr.[2], and Leo Rüst[3]

[1] Institute of Theoretical Computer Science, ETH Zürich, CH-8092 Zürich, Switzerland, gaertner@inf.ethz.ch
[2] Department of Mathematical Sciences, George Mason University, MS 3F2, US-Fairfax VA 22030-4444, USA, wmorris@gmu.edu
[3] Institute of Theoretical Computer Science, ETH Zürich, CH-8092 Zürich, Switzerland, ruestle@inf.ethz.ch

Abstract. We introduce unique sink orientations of grids as digraph models for many well-studied problems, including linear programming over products of simplices and generalized linear complementarity problems over P-matrices (PGLCP). We investigate the combinatorial structure of such orientations and develop randomized algorithms for finding the sink. We show that the orientations arising from PGLCP satisfy the combinatorial *Holt-Klee* condition known to hold for polytope digraphs, and we give the first expected linear-time algorithms for solving PGLCP with a fixed number of blocks.

1 Introduction

A *grid* is a graph whose vertex set is the Cartesian product of n finite sets, with edges joining all pairs of vertices that differ in exactly one component.

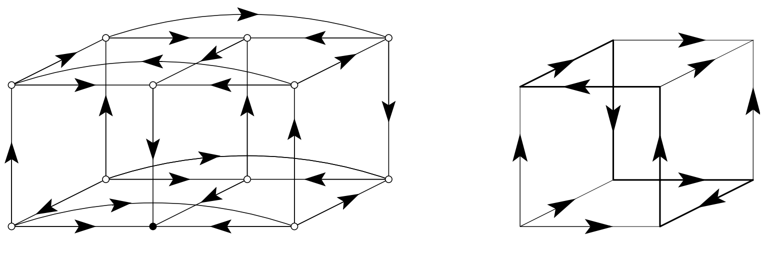

Fig. 1. Left: USO of the $(3 \times 2 \times 2)$-grid. Right: cyclic USO of the 3-cube

If all sets have size two, we get the graph of the n-cube. A *face* or *subgrid* is any induced subgraph spanned by the Cartesian product of subsets of the

⋆ The first and the third author acknowledge support from the Swiss Science Foundation (SNF), Project No. 200021-100316/1. Part of this research was done at the 2004 Barbados Undercurrent Workshop *Polyhedra, Convex Geometry, and Optimization* at Bellairs Research Institute, McGill University. For a full paper see [1].

original sets. An orientation ψ of the grid is called a *unique sink orientation* (USO) if any nonempty face has a unique sink with respect to ψ. Figure 1 (left) depicts a USO of the $(3 \times 2 \times 2)$-grid. In particular, the grid itself must have a unique global sink. Grid USO may contain directed cycles, as the 3-cube in Figure 1 (right) shows.

The significance of USO on grids comes from the fact that they form a simple combinatorial framework subsuming a number of well-studied problems. We show in this paper that the problem of solving a *generalized linear complementarity problem* over a P-matrix (PGLCP), as introduced by Cottle and Dantzig [2], can be recast as the problem of finding the unique sink of an implicitly given grid USO. As special cases, this includes the well-known standard linear complementarity problems over P-matrices (PLCP) [3], linear programming (LP) over products of simplices, and LP over combinatorial cubes. In the LP applications, we get acyclic unique sink orientations (AUSO).

Two major open problems motivate our research. On the one hand, it is unknown whether polynomial-time algorithms exist for PLCP or PGLCP, even though both problems are unlikely to be NP-hard. Megiddo has shown that hardness of PLCP would imply NP = co-NP [4], and his proof extends to PGLCP easily. LP, on the other hand, is solvable in polynomial time (a celebrated result of Khachyian [5]), but a *strongly* polynomial algorithm is not known, even if we are dealing only with LP over combinatorial cubes. Candidates for strongly polynomial algorithms must be *combinatorial* in the sense that the number of arithmetic operations they perform depends only on the combinatorial structure of the LP but not on the actual numbers that encode it. The AUSO approach attempts to extract the combinatorial structure behind LP; in this paper, we generalize to the combinatorics of PGLCP. A polynomial-time algorithm for finding the sink of a grid USO (using an oracle that returns the orientation of a given edge) would solve both problems in strongly polynomial time. It seems unlikely that such an algorithm for general USO will be easier to find than one for PGLCP. Still, the generalization reveals some (algorithmically useful) hidden structure, leading to new results for PGLCP. Ultimately, results obtained along these lines may help to resolve the (combinatorial) complexities of LP and PGLCP.

The AUSO framework also covers a generalization of LP resulting from the replacement of linear objective functions with *abstract objective functions* (AOF) [6, 7], or *completely unimodal numberings* of vertices [8, 9]. On general polytopes, these concepts are dual to the notion of *shellings* [9], and they have successfully been applied to the theory of polytope (di)graphs and linear programming [10, 11].

The case of grid AUSO—equivalently, AOF on products of simplices—has been treated in detail by Björklund et al. as a combinatorial framework for the problem of computing optimal infinite game strategies [12, 13]. The games considered include parity, mean-payoff, and simple stochastic games. Whether parity games (the easiest among the three) can be solved in polynomial time is an important open question [14].

In the planar case ($n = 2$), the combinatorial and algorithmic properties of AUSO have been examined by Tschirschnitz et. al. [15, 16].

Unique sink orientations of cubes (not necessarily acyclic) were first considered by Stickney and Watson as digraph models for PLCP [17]. Most remarkably, Szabó and Welzl gave algorithms for finding the sink of an n-cube USO by looking at only $O(c^n)$ vertices and edges, for some c strictly smaller than 2 [18]. This in particular yields the first combinatorial algorithms for PLCP with nontrivial runtime bounds. Unique sink orientations of the graphs of general polytopes are dual to *exact signings* studied by Kleinschmidt and Onn [19].

In this paper, we generalize results known to hold in some of the above special cases, and we prove new structural and algorithmic results of particular significance for the theory of PLCP and PGLCP. Probably the most surprising fact is that all these results hold even under the presence of directed cycles in the orientation.

We develop two simple *randomized* algorithms whose expected number of calls to the oracle is of the order $f(n)N$, with N being the sum of sizes of the n sets whose product forms the grid, and with $f(n) \approx n!$. If n is fixed, we get *linear-time* algorithms. Specialized to PGLCP, this corresponds to the case in which we have a fixed number of *blocks*; in this situation, we get the first algorithms whose expected complexity is linear in the number of variables—this is optimal.

In the acyclic case, linear-time algorithms for fixed n are known. This follows from the fact that the problem of finding the sink in an AUSO can be formulated as an *LP-type problem* [20] in a natural way [12]. In the LP-type framework, a number of $f(n)N$ oracle calls suffices, even deterministically [21]. The currently best algorithm for the acyclic case combines two randomized algorithms [20, 22] and requires an expected number of $O(Nn + f(n))$ oracle calls, where $f(n) = \exp(O(\sqrt{n \log n}))$ is a *subexponential* function [23]. No subexponential bounds are known for USO that contain cycles.

Unique sink orientations coming from LP over products of simplices satisfy an interesting *geometric* property: consider any subgrid along with its unique sink and unique source.[4] The *Holt-Klee* (HK) condition states that there is a set of as many vertex-disjoint directed paths between source and sink as there are neighbors of the source (equivalently, the sink) in the subgrid [24]. The HK condition does not hold for general AUSO; there are two nonequivalent AUSO of the 3-cube with no set of three vertex-disjoint directed paths from source to sink. The HK condition is important as the only known simple combinatorial condition that can distinguish geometric from abstract situations, and there is an interesting algorithmic scenario in which this distinction becomes apparent [25, 26].

We prove that all grid USO coming from PGLCP do satisfy the HK condition (we also say that they are HK), even in the presence of cycles. This result emphasizes the geometric nature of PGLCP and establishes a new combinatorial way of proving that a given USO cannot be realized as a PGLCP instance. The

[4] The existence of a unique source follows from Theorem 1.

result is new also in the context of PLCP with its wide range of applications [3, 27].

The rest of this paper is organized as follows. After giving some basic results in the next section, Section 3 introduces two problems giving rise to USO of grids. Section 4 analyses two algorithms finding the sink of a grid USO and the final section is devoted to the proof that PGLCP induced orientations fulfill the HK condition.

2 Basics

Throughout this paper, we fix two natural numbers $N \geq n \geq 1$ and an ordered partition

$$\Pi = (\Pi_1, \ldots, \Pi_n)$$

of the set $[N] := \{1, \ldots, N\}$ into n nonempty subsets. We also refer to Π_i as the *block* i.

A subset $J \subseteq [N]$ is called a *Π-vertex* (or simply *vertex*) if $|J \cap \Pi_i| = 1$ for all i. Let V be the set of all vertices. The n-dimensional *grid* spanned by $S \subseteq [N]$ is the undirected graph $G(S) = (V(S), E(S))$, with

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

The vertices of $G(S)$ canonically correspond to the elements of the Cartesian product

$$\prod_{i=1}^{n} S_i, \quad S_i := S \cap \Pi_i.$$

Edges join pairs of vertices J, J' that differ in exactly one coordinate.

A *face* or *subgrid* of $G(S)$ is any graph of the form $G(S')$, for $S' \subseteq S$. Throughout, we abbreviate $G([N])$ as G.

Definition 1. *Let ψ be an orientation of G. ψ is called a* unique sink orientation *(USO) if all nonempty faces of G have unique sinks w.r.t. ψ.*

If ψ induces the directed edge (J, J'), we also write $J \xrightarrow{\psi} J'$.

Outmap and h-vector. Any USO can be specified by associating each vertex J with its outgoing edges. Given J and $j \in [N] \setminus J$, we define $J \rhd j$ to be the unique vertex $J' \subseteq J \cup \{j\}$ which 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.

Given an orientation ψ, the function $s_\psi : V \to 2^{[N]}$, defined via

$$s_\psi(J) := \{j \in [N] \setminus J \mid J \xrightarrow{\psi} J \rhd j\}, \tag{1}$$

is called the *outmap* of ψ. Björklund et. al consider the outmap for acyclic grid USO and call it VID-function (vector of improving directions) [12]. Szabó and

Welzl [18] deal with outmaps for cube USO; formally, these are different from ours, even when we specialize to the cube case, because Szabó and Welzl identify the n-cube vertices with the subsets of some n-element set, while we identify them with certain n-subsets of some $2n$-element set. Still, we can generalize the characterization of USO outmaps by Szabó and Welzl, using one more ingredient.

Definition 2. *Let $s : V \rightarrow 2^{[N]}$. The vector $h(s) = (h_0(s), \ldots, h_{N-n}(s))$, defined via*

$$h_k(s) = \#\{J \in V \mid |s(J)| = k\}$$

is called the h-vector *of s. If $s = s_\psi$ for some orientation ψ on G, we also refer to $h(s)$ as $h(\psi)$.*

The following result is well-known for the acyclic case [9, 12] but going through its proof, one realizes that only the unique sink property is used [15].

Theorem 1. *$h(\psi) = h(\psi')$ for any two USO ψ, ψ' of G.*

As all USO on a grid have the same h-vector, we denote it by $h(\Pi)$, emphasizing the fact that it depends on the parameters of the grid only. An immediate corollary of this theorem is that the h-vector is symmetric, meaning that $h_k(\Pi) = h_{N-n-k}(\Pi)$ for all $k \in \{0, \ldots, N-n\}$ (reversing all edge orientations of a USO yields a USO again [1]). In particular, w.r.t. any given USO, all nonempty subgrids of G also have unique *sources*.

Here is the characterization of functions s that are of the form s_ψ, for ψ being USO (see [1] for a proof).

Lemma 1. *Let $s : V \rightarrow 2^{[N]}$ satisfy $s(J) \cap J = \emptyset$ for all $J \in V$. Such an s is the outmap of a USO of G if and only if*

(i) $(s(J) \oplus s(J')) \cap (J \oplus J') \neq \emptyset$, *for all* $J \neq J'$, *and*
(ii) $h(s) = h(\Pi)$.

Refined Index. The outmap value $s_\psi(J) \subseteq [N]$ of a vertex J w.r.t. some USO ψ is partitioned according to the dimensions of the grid (the sets Π_i). Accordingly, the outdegree of J can be refined to an n-vector of dimensional outdegrees, as follows.

Definition 3. *Let ψ be a grid USO. The function*

$$r_\psi : V \rightarrow \prod_{i=1}^{n} \{0, \ldots, |\Pi_i| - 1\},$$

with

$$r_\psi(J) = (|s_\psi(J) \cap \Pi_1|, \ldots, |s_\psi(J) \cap \Pi_n|), \quad J \in V$$

is called the refined index *of ψ.*

Unlike the outmap, the refined index is a mapping between two sets of the same size, so it is natural to ask whether this mapping is a bijection. This is true in the cube case (where the refined index is just a different way of writing the outmap) [18], and it also holds for grid AUSO [12], where Björklund et al. use the term *signature* (SIG). The proof of the latter result does not generalize to the case of general USO. We show in [1] that the refined index is a bijection in general.

Theorem 2. *Let ψ be a grid USO. The refined index r_ψ is a bijection.*

In the proof, we make use of the concept of *inherited orientations*. Given a grid orientation, an inherited grid orientation can be obtained by collapsing dimensions of the grid, merging vertices along the collapsed dimensions to hypervertices. Hypervertices correspond to faces of the original grid, and the outmap of a hypervertex is defined as the outmap of the corresponding face sink.

The important fact is that inherited orientations of USO are USO again. This is utilized to get a contradiction in the case of two vertices with the same refined index: collapse the grid until we get a 1-dimensional grid USO where two vertices have the same outdegree, contradicting Theorem 1.

The bijection property of the refined index is useful in the proof of the following theorem [1]. Previously, this was only known to hold for USO of 2-dimensional grids that satisfy the Holt-Klee condition [15, 16].

Theorem 3. *Any unique sink orientation of a* 2*-dimensional grid is acyclic.*

3 Grid LP and Generalized LCP

In this section, we present two geometric models of grid USO arising from *Grid LP* and *Generalized LCP*. Consider a linear program in the variables $x = (x_1, \dots, x_N)^T$, of the form

$$\begin{aligned} \text{minimize } & c^T x \\ \text{subject to } & Ax = b, \\ & x \geq 0, \end{aligned} \tag{2}$$

where $A \in \mathbb{R}^{n \times N}, b \in \mathbb{R}^n, c \in \mathbb{R}^N$. For $S \subseteq [N]$, let A_S denote the submatrix of columns indexed by S. Furthermore assume that every vertex J is a nondegenerate basis in (2), meaning that $A_J^{-1} b > 0$. We say that the LP is Π*-compatible* (or, equivalently, the LP is a *Grid LP*), and we call A_J a *representative submatrix*. We will assume that the ordering of the columns in A_J is compatible with Π, meaning that the i-th column of A_J comes from block $\Pi_i, i \in [n]$.

The feasible region of a Π-compatible LP is combinatorially equivalent to the product of n simplices with a total of N facets. A unique sink orientation of the grid G is obtained from a Π-compatible LP with *generic* objective function vector c, meaning that $c^T x$ is not constant on any edge. In this case, an edge can be directed towards its vertex of lower objective function value. The resulting orientation is a USO, with the unique sink of the face $G(S)$ corresponding to the unique optimal basis of the linear program resulting from (2) by restricting to the variables with indices in S.

The simplex algorithm (Chvátal's book [28] contains an excellent introduction) determines the edge orientations in terms of *reduced cost coefficients*. More precisely, if J is some basis, the row vector

$$\bar{c}(J) := c^T - c_J^T A_J^{-1} A \tag{3}$$

is the *reduced cost vector* associated with J. Note that $\bar{c}(J)_J = 0$, and that $\bar{c}(J)_j \neq 0$ for $j \notin J$, if c is generic. Thus, if J' is adjacent to J, with j being the unique index in $J' \setminus J$, we have

$$J \to J' \quad \Leftrightarrow \quad \bar{c}(J)_j < 0. \tag{4}$$

The existence of a unique sink therefore provides us with a solution to the following feasibility problem.

Definition 4. *Let $A \in \mathbb{R}^{n \times N}$ and $c \in \mathbb{R}^N$ such that A has* property P*, meaning that all determinants of representative submatrices A_J of A have the same nonzero sign. The* P-generalized linear complementarity problem (PGLCP) *defined by (A, c) is the problem of finding a vector $y \in \mathbb{R}^n$ such that $c^T \geq y^T A$, and with the property that for every $i \in [n]$, there is some $j \in \Pi_i$ satisfying $c_j = (y^T A)_j$.*

Intuitively, PGLCP is LP of the form (2) 'without a right-hand side'. It turns out that under this generalization, uniqueness of solution as well as the USO formulation persist. Because already the cyclic cube USO of Figure 1 (right) arises from a PGLCP [17], the generalization is proper.

Theorem 4. *Let $A \in \mathbb{R}^{n \times N}$ and $c \in \mathbb{R}^N$ define a* PGLCP *instance. Then*

(i) there exists a unique solution $y \in \mathbb{R}^n$ to (A, c), and
(ii) if c is generic, the edge orientations given by (4) define a unique sink orientation of the grid G.

Proof. Fix some vertex J, define

$$M^T := A_J^{-1} A,$$
$$q^T := \bar{c}(J) = c^T - c_J^T A_J^{-1} A,$$

and consider the problem of finding $z \in \mathbb{R}^n, w \in \mathbb{R}^N$ such that

$$w - Mz = q, \tag{5}$$

$$\prod_{j \in \Pi_i} w_j = 0, \quad i \in [n], \tag{6}$$

and

$$w, z \geq 0. \tag{7}$$

By definition, every solution must satisfy

$$z_i = w_j, \quad i \in [n], j \in J_i, \tag{8}$$

meaning that (6) is equivalent to

$$z_i \prod_{j \in \Pi_i} w_j = 0. \tag{9}$$

Equations (5), (7) and (9) define the P-generalized linear complementarity problem according to Cottle & Dantzig who show that a feasible solution (w, z) exists if every representative submatrix of M^T is a P-*matrix* [2]. A matrix is a P-matrix if all determinants of principal minors are positive [3]. In our case, M^T satisfies this, which easily follows from the fact that M^T has property P *and* contains a representative identity matrix. Note that the n variables $w_j, j \in J$ of system (5) are redundant by (8) and can be deleted from the problem, along with their corresponding rows.

To prove the existence of y in part (i), it remains to observe that y fulfills the conditions of Definition 4 if and only if $z^T = c_J^T - y^T A_J$ and $w^T = c^T - y^T A$ solve (5), (6) and (7). The uniqueness of y follows from the known uniqueness proofs in the setup of Cottle & Dantzig [29, 30].

Statement (ii) is a corollary of (i), because c being generic implies that the vector y from (i) can be expressed in the form $y^T = c_K^T A_K^{-1}$ for exactly one K. This set K is the unique sink of G in the orientation defined by (4). The fact that this orientation defines a USO easily follows: applying the above arguments to the PGLCP instance (A_S, c_S), we can prove the existence of a unique sink in the subgrid $G(S)$. □

4 Algorithms

In this section, we develop randomized algorithms for finding the sink of a given grid USO, implicitly specified by an *edge evaluation* oracle. The oracle must be able to return the orientation of any given grid edge. Our complexity measure will be the maximum (expected) number of oracle calls needed to find the sink in the worst case. In all concrete instances, this oracle can easily be implemented in polynomial time, meaning that the number of oracle calls is a good measure of complexity. In the case of PGLCP, for example, an edge evaluation must return the sign of a single coefficient of the reduced cost vector (3).

Any USO algorithm which calls the oracle only a polynomial number of times is actually a *strongly* polynomial algorithm for LP and PGLCP. Moreover, the complexity of a single edge evaluation typically only depends on n but not on N (in the PGLCP case, this complexity is $O(n^3)$). Thus, if n is considered to be a constant, any bound on the number of edge evaluations determines the complexity of the algorithm up to a constant factor.

4.1 The Product Algorithm

This algorithm generalizes the *product algorithm* of Szabó and Welzl from cubes [18] to grids, with a slight twist: while in the n-cube, all dimensions are equivalent with respect to their size (which is two), a general grid may have 'heavy' dimensions (with large Π_i) and 'light' dimensions. Our algorithm gives priority to the heavy dimensions. Recall that $S_i := S \cap \Pi_i$. A generic call to Product finds the sink of a nonempty face $G(S)$, see Figure 2 (left).

Algorithm 5
Product(S):
  IF is_vertex(S) THEN
    RETURN S
  ELSE
    choose a heaviest $S_i \in S$
    choose $j \in S_i$ at random
    $K :=$ Product($S \setminus \{j\}$)
    $K' := K \triangleright j$
    IF $K' \stackrel{\psi}{\to} K$ THEN
      RETURN K
    ELSE
      RETURN
      Product($(S \setminus S_i) \cup \{j\}$)
    END
  END

Algorithm 6
RandomFacet(J, S):
  IF $S = J$ THEN
    RETURN J
  ELSE
    choose a heaviest $S_i \in S$
    choose $j \in S_i \setminus J_i$ at random
    $K :=$ RandomFacet($J, S \setminus \{j\}$)
    $K' := K \triangleright j$
    IF $K' \stackrel{\psi}{\to} K$ THEN
      RETURN K
    ELSE
      RETURN
      RandomFacet($K', (S \setminus S_i) \cup \{j\}$)
    END
  END

Fig. 2. The Algorithms Product and RandomFacet

The algorithm recursively computes the sink K of the subgrid $G(S \setminus \{j\})$. If the edge incident to K in direction j is incoming, we have already found the global sink, otherwise we need to search the lower-dimensional 'facet' $G((S \setminus S_i) \cup \{j\})$ recursively. To prepare the analysis of the algorithm, let $S \subseteq [N]$ be a set containing a vertex and let $z = \max_{i=1}^{n} |S_i|$ be the size of a heaviest dimension. The $(z-1)$-vector $(a_2, \ldots, a_z)$, defined through

$$a_t = |\{i \in [n] \mid |S_i| = t\}|$$

is called the *characteristic* of S. If $z = 1$ (meaning that S is a vertex itself), the characteristic is the empty vector ().

It can easily be shown by induction that the expected number of edge evaluations in Product(S) only depends on the characteristic of S but (maybe surprisingly) not on the input USO ψ. We can even compute the exact expectation.

Theorem 5. *For $z \geq 1$, let $T_e(a_2, \ldots, a_z)$ denote the expected number of edge evaluations in a call to* PRODUCT(S)*, where S has characteristic $(a_2, \ldots, a_z)$. Then*

$$\begin{aligned} T_e(a_2, \ldots, a_z) = T_e(a_2, \ldots, a_{z-1}+1, a_z-1) + 1 \\ + T_e(a_2, \ldots, a_{z-1}, a_z-1)/z \end{aligned} \tag{10}$$

for $z > 1$, with $T_e(a_2, \ldots, a_{z-1}, 0) := T_e(a_2, \ldots, a_{z-1})$ and $T_e() = 0$. The solution to this recurrence is

$$T_e(a_2, \ldots, a_z) = \prod_{k=2}^{z} H_k^{a_k} + \sum_{k=2}^{z} \prod_{\ell=k}^{z} (H_\ell - H_{k-1} + 1)^{a_\ell} - z. \tag{11}$$

Here, H_k is the k-th Harmonic number.

Proof. It is clear that $(a_2, \ldots, a_{z-1}+1, a_z-1)$ and $(a_2, \ldots, a_{z-1}, a_z-1)$ are the characteristics of the grids handled in the recursive calls. Moreover, the second recursive call is executed if and only if the global sink contains the chosen element j. This happens with probability $1/z$. The recurrence follows. The closed form (11) can be checked by induction. □

For fixed z, (11) is maximized if $a_z = n$. This corresponds to the characteristic $(0, \ldots, 0, n)$ of the $(z \times \cdots \times z)$-grid. It follows that

$$T_e(a_2, \ldots, a_z) \leq H_z^n + \sum_{k=2}^{z} (H_z - H_{k-1} + 1)^n - z. \tag{12}$$

The middle term

$$f(n, z) := \sum_{k=2}^{z} (H_z - H_{k-1} + 1)^n \tag{13}$$

asymptotically dominates the bound in (12), and an estimate of $f(n, z) \leq (z-1)H_z^n \approx z \ln^n z$ immediately follows. The next result shows that the bound is actually *linear* in z.

Lemma 2.

$$T_e(a_2, \ldots, a_z) \leq (\lfloor en! \rfloor - 1)\, z + H_z^n.$$

Proof. Using the estimate

$$\sum_{t=\ell}^{u} g(t) \leq \int_{\ell-1}^{u} g(x)\; dx \tag{14}$$

for any decreasing function g such that the integral exists, we can bound (13) as follows.

$$\begin{aligned} f(n, z) &= \sum_{k=2}^{z} (H_z - H_{k-1} + 1)^n \overset{(14)}{\leq} \sum_{k=2}^{z} (\ln z - \ln(k-1) + 1)^n \\ &= \sum_{k=1}^{z-1} \left(1 - \ln \frac{k}{z}\right)^n < \sum_{k=1}^{z} \left(1 - \ln \frac{k}{z}\right)^n \\ &\overset{(14)}{\leq} \int_0^z \left(1 - \ln \frac{k}{z}\right)^n dk = z \int_0^1 (1 - \ln x)^n\; dx := z I_n. \end{aligned}$$

Integration by parts yields the recurrence relation

$$I_n = 1 + nI_{n-1},$$

with $I_0 = 1$. This solves to $I_n = \lfloor en! \rfloor$ for $n > 0$ [31]. The statement follows. □

This means, algorithm Product solves any PGLCP instance with a fixed number n of blocks in expected time $O(z) = O(N)$ which is asymptotically optimal.

Algorithm 5 is a close relative of algorithms due to Seidel (for linear programming with n variables and N constraints) [32] and Welzl (for finding the smallest enclosing ball of a set of N points in dimension n) [33].

The expected number of edge evaluations depends on the rule for choosing i, and a bad rule can lead to a complexity which is asymptotically worse than what we found in Theorem 5. For instance, always choosing i to be a lightest dimension yields an expected number of

$$\sum_{k=1}^{z-1} k \left(\prod_{\ell=2}^{k} H_\ell^{a_\ell} \right) \sum_{m=0}^{a_{k+1}-1} H_{k+1}^m \overset{a_z:=n}{=} (z-1) \sum_{m=0}^{n-1} H_z^m = \frac{z-1}{H_z - 1}(H_z^n - 1),$$

edge evaluations, which is superlinear in z. We believe (although we cannot prove it formally) that our choice of i in Algorithm 5 leads to the smallest possible expected number of edge evaluations.

4.2 The Algorithm RandomFacet

The RandomFacet algorithm shares its basic idea with the Product algorithm. In addition to the current set S, it maintains a current vertex $J \in V(S)$ which may be replaced at some point by a neighbor of J along an outgoing edge. This means, we get a *path-following* algorithm. In order to guarantee the invariant $J \in V(S)$, the element j which gets removed from S for the first recursive call must not be in J, see Figure 2 (right).

As in the case of the Product algorithm, we can derive an explicit bound on the runtime, but here it is an *upper* bound instead of an exact bound. We get that the expected number of edge evaluations for a grid of characteristic $(a_2, \ldots, a_z)$ is at most

$$\sum_{k=1}^{z-1} \prod_{\ell=k}^{z-1} (H_\ell - H_{k-1} + 1)^{a_{\ell+1}} - z + 1.$$

The complexity is again maximized if $a_z = n$, and an upper bound of

$$\sum_{k=1}^{z-1} (H_{z-1} - H_{k-1} + 1)^n - z + 1 = (H_{z-1} + 1)^n + \sum_{k=2}^{z} (H_{z-1} - H_{k-1} + 1)^n - z$$

holds. Comparing this with (12), we see that for large z, both algorithms have approximately the same expected worst-case complexity which we have shown to be linear in z. However, only RANDOMFACET has the potential of being faster than the upper bound in practice, for example if the starting vertex is already close to the sink, or if paths tend to be short in the USO under consideration.

The algorithm RANDOMFACET is a close relative of an algorithm by Matoušek, Sharir and Welzl for *LP-type problems* [20].

5 The Holt-Klee Condition

A grid USO is said to be *Holt-Klee* (HK), if there exists a set of N-n vertex-disjoint paths from source to sink and if in addition, every nonempty subgrid is HK.

In the following, we prove that USO coming from PGLCP are HK. We actually prove that a larger class of digraphs, those defined by *complete pointed fans* in $\mathbb{R}^d$, has the Holt-Klee property. Here, we follow the notation of Ziegler [34].

Definition 7. *A fan in $\mathbb{R}^d$ is a family $\mathcal{F} = \{C_1, C_2, \ldots, C_t\}$ of nonempty polyhedral cones, so that*

(i) Every nonempty face of a cone in $\mathcal{F}$ is also a cone in $\mathcal{F}$.
(ii) The intersection of any two cones in $\mathcal{F}$ is a face of both.

A fan $\mathcal{F}$ is *complete* if the union of its cones is $\mathbb{R}^d$. It is *pointed* if the zero vector is one of its cones. The *dual graph* $G_{\mathcal{F}}$ of a fan has as its vertex set the set of d-dimensional cones of $\mathcal{F}$, with two cones joined by an edge if their intersection is a $(d-1)$-dimensional face of $\mathcal{F}$.

In our PGLCP application (see Section 3), we use $d = N-n$, and we consider a matrix $\hat{A}$ that has as its row space the complement of the row space of A in $\mathbb{R}^N$ and say that a set of columns of $\hat{A}$ generates a cone of $\mathcal{F}$ if and only if the set does not contain all of the columns indexed by Π_i for any i. We show in [1] that property P of the matrix A ensures that the family $\mathcal{F}$ satisfies (i) and (ii). The dual graph of such a fan is a grid graph.

A vector $q \in \mathbb{R}^d$ is said to be in general position with respect to $\mathcal{F}$ if it is not contained in any hyperplane that contains a $(d-1)$-dimensional cone of $\mathcal{F}$. If a vector q is in general position with respect to a fan $\mathcal{F}$, we can define an orientation $\Gamma_{q,\mathcal{F}}$, in which an edge joining cones C and C' is oriented from C to C' if $C \backslash C'$ and q are on opposite sides of the hyperplane containing $C \cap C'$. The digraph $\Gamma_{q,\mathcal{F}}$ has a unique sink and source, which are the faces of $\mathcal{F}$ that contain q and $-q$. The interior of a directed path in a digraph is the set of vertices in the path other than the first and the last vertex.

Theorem 6. *There is a set of d directed paths in $\Gamma_{q,\mathcal{F}}$ from the source to the sink that have pairwise disjoint interiors.*

We need the following lemma for the proof.

Lemma 3. *Let $\mathcal{K} = \{K_1, K_2, \ldots, K_{d-1}\}$ be d-dimensional cones of $\mathcal{F}$, and suppose that none of these cones contains q or $-q$. Then there exists a vector w orthogonal to q so that the segments from q to w and from $-q$ to w both have empty intersection with each of the cones of $\mathcal{K}$.*

Proof. Assume without loss of generality that q is the d^{th} unit vector $(0, \ldots, 0, 1)$. Let $K \in \mathcal{K}$ and let $\hat{A}_K$ be the submatrix of $\hat{A}$ containing the generators of K. Now let $\overline{A}_K$ be the matrix obtained from $\hat{A}_K$ by deleting the last row. The columns of $\overline{A}_K$ are the projections of the columns of $\hat{A}_K$ onto the hyperplane H_q orthogonal to q. Because neither q nor $-q$ is in K, the systems $\hat{A}_K x = q, x \geq 0$ and $\hat{A}_K x = -q, x \geq 0$ have no solution. It follows that the system $\overline{A}_K x = 0, x \geq 0, x \neq 0$ has no solution. By Gordan's Theorem [35] there exists a vector z_K so that $z_K^T \overline{A}_K > 0$. For such a z_K, any sufficiently small perturbation of it will also satisfy the inequality. Therefore there we can find a linearly independent set $\{z_{K_1}, z_{K_2}, \ldots, z_{K_{d-1}}\}$ so that for each i, $z_{K_i}^T \overline{A}_{K_i} > 0$. Let Z be a $(d-1) \times (d-1)$ matrix that has as its rows the vectors $z_{K_i}^T, i = 1, \ldots, d-1$. Stiemke's Theorem [36] says that the system $Zw \leq 0, Zw \neq 0$ has a solution if and only if the system $y^T Z = 0, y > 0$ has no solution. But the matrix Z is nonsingular, so the second system has no solution. Therefore there must be a vector $w \in H_q$ that is not in any of the cones $\overline{A}_K$ for $K \in \mathcal{K}$. This w is nonzero, because Zw is nonzero. The vector w satisfies the requirements of the lemma, because any intersection of a cone of $\mathcal{K}$ with the segment from q to w or the segment from $-q$ to w would project to the cone generated by w. It should also be noted that since the cones $\overline{A}_K$ are closed, there is an open set of such w. □

With this, we are ready for the proof.

Proof (of Theorem 6). The directed vertex version of Menger's theorem states that there will be d disjoint directed paths from the source of $\Gamma_{q,\mathcal{F}}$ to the sink if and only if there do not exist $d-1$ vertices of the graph other than the source and the sink that cover all directed paths from the source to the sink. A set of $d-1$ vertices of $\Gamma_{q,\mathcal{F}}$ other than the source and the sink corresponds to a set $\mathcal{K} = \{K_1, K_2, \ldots, K_{d-1}\}$ as in the lemma. Because the set of w satisfying the conditions of the lemma is open, we can choose a w for which the segments from w to q and from w to $-q$ do not meet any cones of $\mathcal{F}$ of dimension less than $d-1$. We claim that the sequence of d-dimensional cones met by the directed segment from $-q$ to w, followed by the sequence of d-dimensional cones met by the directed segment from w to q, corresponds to a directed path from the source to the sink of $\Gamma_{q,\mathcal{F}}$. Suppose C_i and C_j are two d-dimensional cones of $\mathcal{F}$, and that the directed segment from $-q$ to w crosses, in order, $(C_i, C_i \cap C_j, C_j)$. Then $C_i \backslash (C_i \cap C_j)$ and q are on opposite sides of the hyperplane spanned by $C_i \cap C_j$, so the edge of $\Gamma_{q,\mathcal{F}}$ connecting C_i and C_j is oriented from C_i to C_j. Similarly, the edges connecting cones met by the directed segment from w to q are oriented consistently with the direction of the segment. □

The grid orientation defined by $\Gamma_{q,\mathcal{F}}$ is the same as that defined in Section 3 through the reduced costs in Equation (3) [1]. This yields the desired result.

Corollary 1. *Any* PGLCP*-induced grid USO ψ satisfies the Holt-Klee condition.*

Reorienting all edges along a fixed dimension of the grid preserves the property of being PGLCP-induced [1]. It follows that a grid USO can only be PGLCP-induced if all 2^n reorientations satisfy the Holt-Klee condition.

The special case of the Theorem 6 in which $\mathcal{F}$ is the normal fan of a polytope was shown by Holt and Klee [24]. Our proof uses Menger's theorem, which was used by Holt and Klee, but does not use the geometry of a polytope. Our approach may be seen as an alternate way to prove the Holt - Klee Theorem. The graph $\Gamma_{q,\mathcal{F}}$ has been used by Kleinschmidt and Onn to prove that fans are *signable* [19]. Restricted to PGLCP-induced fans, their result simply says that the undirected grid graph Γ_F underlying $\Gamma_{q,\mathcal{F}}$ has a unique sink orientation, namely $\Gamma_{q,\mathcal{F}}$. Theorem 6 strengthens the result of Kleinschmidt and Onn by showing that the signings they produce have an interesting additional property.

6 Acknowledgment

We thank an IPCO referee for useful comments, Falk Tschirschnitz for many helpful discussions and David Loeffler for suggesting the simple proof of Lemma 2.

References

1. Gärtner, B., Morris, Jr., W.D., Rüst, L.: Unique sink orientations of grids. Technical Report 472, Institute of Theoretical Computer Science (2005)
2. Cottle, R.W., Dantzig, G.B.: A generalization of the linear complementarity problem. Journal on Combinatorial Theory **8** (1970) 79–90
3. Cottle, R.W., Pang, J., Stone, R.E.: The Linear Complementarity Problem. Academic Press (1992)
4. Megiddo, N.: A note on the complexity of P-matrix LCP and computing an equilibrium. Technical report, IBM Almaden Research Center, San Jose (1988)
5. Khachiyan, L.G.: Polynomial algorithms in linear programming. U.S.S.R. Comput. Math. and Math. Phys **20** (1980) 53–72
6. Adler, I., Saigal, R.: Long monotone paths in abstract polytopes. Math. Operations Research **1** (1976) 89–95
7. Kalai, G.: Linear programming, the simplex algorithm and simple polytopes. Math. Programming **79** (1997) 217–233
8. Wiedemann, D.: Unimodal set-functions. Congressus Numerantium **50** (1985) 165–169
9. Hoke, K.W.: Completely unimodal numberings of a simple polytope. Discrete Appl. Math. **20** (1988) 69–81
10. Kalai, G., Kleitman, D.J.: A quasi-polynomial bound for the diameter of graphs of polyhedra. Bulletin Amer. Math. Soc. **26** (1992) 315–316
11. Kalai, G.: A simple way to tell a simple polytope from its graph. J. Combinatorial Theory **49** (1988) 381–383
12. Björklund, H., Sandberg, S., Vorobyov, S.: Randomized subexponential algorithms for parity games. Technical Report TR-2003-019, Department of Information Technology (2003)
13. Björklund, H., Sandberg, S., Vorobyov, S.: A combinatorial strongly subexponential strategy improvement algorithm for mean payoff games. In: Proc. 29th International Symposium on Mathematical Foundations of Computer Science (MFCS). Lecture Notes in Computer Science, Springer-Verlag (2004, to appear)
14. Grädel, E., Thomas, W., (Eds.), T.W.: Automata, Logics, and Infinite Games. Volume 2500 of Lecture Notes in Computer Science. Springer-Verlag (2002)

15. Tschirschnitz, F.: LP-related Properties of Polytopes with Few Facets. PhD thesis, ETH Zürich (2003)
16. Felsner, S., Gärtner, B., Tschirschnitz, F.: Grid orientations, $(d, d+2)$-polytopes, and arrangements of pseudolines. Discrete and Computational Geometry (to appear)
17. Stickney, A., Watson, L.: Digraph models of Bard-type algorithms for the linear complementarity problem. Mathematics of Operations Research **3** (1978) 322–333
18. Szabó, T., Welzl, E.: Unique sink orientations of cubes. In: Proc. 42^{nd} IEEE Symp. on Foundations of Comput. Sci. (2000) 547–555
19. Kleinschmidt, P., Onn, S.: Signable posets and partitionable simplicial complexes. Discrete Comput. Geom. **15** (1996) 443–466
20. Matoušek, J., Sharir, M., Welzl, E.: A subexponential bound for linear programming. Algorithmica **16** (1996) 498–516
21. Chazelle, B., Matoušek, J.: On linear-time deterministic algorithms for optimization problems in fixed dimension. J. Algorithms **21** (1996) 579–597
22. Clarkson, K.L.: Las Vegas algorithms for linear and integer programming. J. ACM **42** (1995) 488–499
23. Gärtner, B., Welzl, E.: Linear programming – randomization and abstract frameworks. In: Proc. 13th Ann. ACM Symp. Theoretical Aspects of Computer Science. Volume 1046 of Lecture Notes Comput. Sci., Springer-Verlag (1996) 669–687
24. Holt, F., Klee, V.: A proof of the strict monotone 4-step conjecture. In Chazelle, J., Goodman, J.B., Pollack, R., eds.: Advances in Discrete and Computational geometry. Contemporary Mathematics. Amer. Math. Soc. (1998)
25. Matoušek, J.: Lower bounds for a subexponential optimization algorithm. Random Structures & Algorithms **5** (1994) 591–607
26. Gärtner, B.: The Random-Facet simplex algorithm on combinatorial cubes. Random Structures & Algorithms **20** (2002)
27. Schäfer, U.: A linear complementarity problem with a P-matrix. SIAM Review **46** (2004) 189–201
28. Chvátal, V.: Linear Programming. W. H. Freeman, New York, NY (1983)
29. Szanc, B.P.: The Generalized Complementarity Problem. PhD thesis, Rensselaer Polytechnic Institute, Troy, NY (1989)
30. Gowda, M.S., Sznajder, R.: The generalized order linear complementarity problem. SIAM J. Matrix Anal. Appl. **15** (1994) 779–795
31. Sloane, N.J.A.: Sequence A000522. The On-Line Encyclopedia of Integer Sequences, `http://www.research.att.com/~njas/sequences/` (2004)
32. Seidel, R.: Small-dimensional linear programming and convex hulls made easy. Discrete Comput. Geom. **6** (1991) 423–434
33. Welzl, E.: Smallest enclosing disks (balls and ellipsoids). In Maurer, H., ed.: New Results and New Trends in Computer Science. Volume 555 of Lecture Notes Comput. Sci. Springer-Verlag (1991) 359–370
34. Ziegler, G.M.: Lectures on Polytopes. Springer-Verlag (1995)
35. Gordan, P.: Über die Auflösung linearer Gleichungen mit reellen Coefficienten. Mathematische Annalen **6** (1873) 23–28
36. Stiemke, E.: Über positive Lösungen homogener linearer Gleichungen. Journal für die reine und angewandte Mathematik **76** (1915) 340–342