The number of unique-sink orientations of the hypercube*

Jiří Matoušek

Department of Applied Mathematics and
Institute of Theoretical Computer Science (ITI)
Charles University
Malostranské nám. 25, 118 00 Praha 1
Czech Republic

Abstract

Let Q_d denote the graph of the d-dimensional cube. A *unique-sink orientation* (USO) is an orientation of Q_d such that every face of Q_n has exactly one sink (vertex of outdegree 0); it does not have to be acyclic. USO have been studied as an abstract model for many geometric optimization problems, such as linear programming, finding the smallest enclosing ball of a given point set, certain classes of convex programming, and certain linear complementarity problems. It is shown that the number of USO is $d^{\Theta(2^d)}$.

1 Introduction

We write $\log x$ for logarithm of x in base 2. The set $\{1, 2, \ldots, n\}$ is denoted by $[n]$.

Let Q_d denote the d-dimensional cube in the graph-theoretical sense. The vertex set is $\{0,1\}^d$, and two vertices u, v are connected by an edge iff they differ in exactly one position. Alternatively, Q_n is the graph of the convex polytope $[0,1]^d$. The notions of face and facet are inherited from the polytope interpretation: Explicitly, a k-dimensional face of Q_n

*This research was partially supported by ETH Zürich and done in part during the workshop "Towards the Peak" in La Claustra, Switzerland and during a visit to ETH Zürich.

is the subgraph induced by a set of the form $F_{v,I} = \{u \in \{0,1\}^d : u_i = v_i \text{ for all } i \in I\}$, where $I \subseteq [d]$ has cardinality $d-k$ and $v \in \{0,1\}^d$. A facet is a $(d-1)$-dimensional face.

A *unique-sink orientation of* Q_n, or USO for short, is an orientation of the graph Q_n such that every face has exactly one sink (a vertex of outdegree 0). The 2-dimensional USO are, up to isomorphism, as in the following picture:

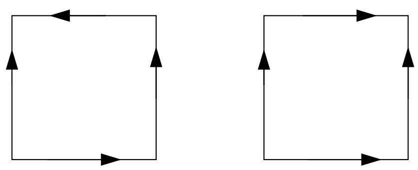

Here are two examples of a 3-dimensional USO:

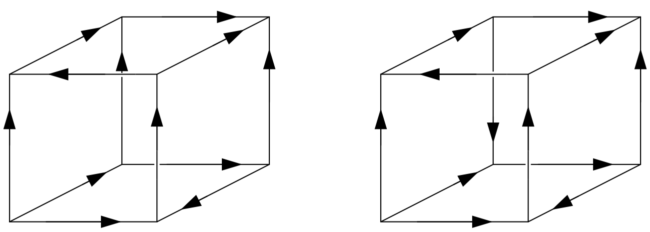

Note that the second example has a directed cycle (of length 6). If a given USO contains no directed cycle, we call it an *acyclic USO*, or briefly AUSO. This notion has been investigated for graphs general convex polytopes, under the names *abstract objective function* (e.g., Kalai [?]) or *completely unimodal numbering* (Hoke [?]).

The motivation for studying USO and AUSO comes mainly from the theory of linear programming and other geometric optimization problems, and it is presented briefly at the end of this section. Here we investigate the number of USO as a function of d.

Let $uso(d)$ denote the number of d-dimensional USO, and similarly for $auso(d)$. Formally, the counting is in the labeled sense, not up to isomorphism; on the other hand, the number of isomorphisms of Q_d, which is $2^d d! = 2^{\Theta(d \log d)}$, is asymptotically negligible compared to the number of (A)USO.

Further, let USO2 denote an orientation of Q_d in which every 2-dimensional face has a unique sink, and let $uso2(d)$ be the corresponding counting function. It is well known that every *acyclic* USO2 is an AUSO, but a general USO2 need not be a USO.

The total number of orientations of Q_d is $2^{d2^{d-1}}$. On the other hand, it is easy to see that $auso(d) \geq 2^{2^{d-1}}$: this follows, for example, from $auso(1) = 2$ and the recurrence $auso(d) \geq auso(d-1)^2$, since taking arbitrary AUSO of two complementary facets of Q_d and orienting all the edges going between them in the same direction yields an AUSO.

Further, it is known that the number of *all* acyclic orientations of Q_d is $d^{\Theta(2^d)}$; see Kahale and Schulman [**?**]. More precise bounds are $(\frac{d}{2} + O(1))^{2^d}$ from below and $(d+1)^{2^d}$ from above.

Therefore, we have $2^{2^{d-1}} \leq auso(d) \leq (d+1)^{2^d}$. The lower bound can be improved a little, but at present I am not aware of any lower bound for $auso(d)$ exceeding $2^{O(2^d)}$.

Here we prove that $uso(d)$, as well as $uso2(d)$, have the order of magnitude $d^{\Theta(2^d)}$, which is similar to the number of all acyclic orientations and much smaller than the number of all orientations.

Theorem 1 *There exist constants $C \geq c > 0$ such that for all $d \geq 2$, we have*

$$d^{c2^d} \leq uso(d) \leq uso2(d) \leq d^{C2^d}.$$

Background and references. Unique-sink orientations arise when the edges of a deformed geometric d-dimensional cube (i.e., a convex polytope with the combinatorial structure of a cube) are oriented according to some generic linear function. These orientations are easily seen to be acyclic. Probably the most famous such orientation is the *Klee–Minty cube*. Some relevant references for this direction are Hoke [**?**], Kalai [**?**], Gärtner, Henk, and Zigler [**?**].

A motivation for considering USO possibly containing directed cycles are certain *linear complementary problems*, namely those defined by the so-called *P-matrices*; see, e.g., Morris [**?**] for definitions and references. Cycles may appear in these orientations.

Gärtner, Myiazawa, and Welzl [**?**] identified a class of convex programs that give rise to USOs, again possibly cyclic ones. A primary example is the problem of finding the smallest ball enclosing given $d+1$ affinely independent points in $\mathbf{R}^d$ (also see [**?**]).

In all these cases, it would be extremely interesting to find an efficient algorithm, deterministic or randomized, that finds the sink of a given USO (or AUSO) quickly. The USO is assumed to be specified by an oracle that, given a vertex of Q_d, returns the orientations of all edges incident to that vertex. Alternatively, good lower bounds might be easier to prove in this abstract model than for a concrete geometric optimization problem. For some progress in this direction see Szabó and Welzl [?] and Schurr and Szabó [?].

2 Proof of the lower bound in Theorem 1

Let O_0 denote the "Boolean orientation" of Q_d; that is, every edge is oriented from the vertex with fewer 1s to the one with more 1s. Let M be any matching in Q_d (a set of pairwise vertex-disjoint edges), and let O_M be the orientation obtained from O_0 by reversing the orientation of all edges of M (the idea of this "reversing a matching" goes back to Günter M. Ziegler, as far as I could find out; he suggested it in a somewhat different context).

We claim that O_M is a USO. (Note that it need not be acyclic; the right 3-dimensional cube in the picture in Section 1 is an example.) Clearly, it suffices to verify that the whole cube has a unique sink. Let us consider the vertex $s = (1, 1, \ldots, 1)$ that is the sink of O_0. We note that the neighbors of s have outdegree exactly 1 in O_0. If none of the edges of M is incident to s, then s remains a sink, and no other sink may arise, since neighbors of s retain outdegree 1 and the outdegree of any other vertex changes by at most 1. If $\{u, s\} \in M$ for some neighbor u of s, then s ceases to be a sink and u becomes the new sink, while no other vertex may become a sink for the same reasons as before.

The number of matchings of a graph is a well-studied topic. The famous Van der Waerden conjecture proved by Falikman shows that any k-regular bipartite graph on $n+n$ vertices has at least $\frac{k^n n!}{n^n} \geq (\frac{k}{e})^n$ perfect matchings (see Schrijver [?] for a slightly improved bound and references). For Q_d, we thus obtain the lower bound of $\left(\frac{d}{e}\right)^{2^{d-1}}$. This proves the lower bound in Theorem 1. □

Remark. It might seem that the "matching-reversal' orientations as above could produce many AUSO, if we were able to choose the matchings M carefully enough to avoid directed cycles.

However, an acyclic orientation of any graph is completely determined by the outdegrees of vertices. This is well known and easy to check: Given an unoriented graph and the outdegrees, the acyclic orientation can be reconstructed by fixing the orientations of the edges entering the sinks, removing them from the graph and updating the outdegrees accordingly, and continuing until no edges remain.

The matching-reversal construction changes each outdegree by at most 1. Therefore, the total number of acyclic USO resulting from this construction is at most 2^{2^n}, and we cannot improve the lower bound for $auso(d)$ in this way.

Remark. As was pointed out by Ingo Schurr, the matching-reversal construction can be generalized as follows: One can choose an arbitrary collection of pairwise disjoint faces $F_1, \ldots, F_m$ in Q_d and change the Boolean orientation O_0 on each F_i to an arbitrary USO. The result is still an USO.

3 Proof of the upper bound in Theorem 1

We are going to prove that the number of USO2 is at most $d^{O(2^d)} = 2^{O(2^d \log d)}$. We prove the following more general statement by induction on d and k.

Claim 2 *Let $V \subseteq V(Q_d)$ be a set consisting of k vertices of the d-cube, $d \geq 2$, let $E_V = E(Q_d) \cap \binom{V}{2}$ be the set of edges of the subgraph of Q_d induced by V, and let O be an arbitrary orientation of the edges of $E(Q_d) \setminus E_V$. Then the number of* USO2 *of Q_d that extend O (i.e., agree with the O on all the edges of $E(Q_d) \setminus E_V$) is at most $2^{f_d(k)}$, where $f_d(k) = 0$ for $k \leq 1$ and*

$$f_d(k) \leq Ck(\log\log k + \log d),$$

C a sufficiently large constant.

For convenience, let us refer to the facet of Q_d whose vertices have the last component 0 as the *bottom facet*, while the complementary facet is the *top facet*, and the edges connecting the top and bottom facets are the *vertical edges*.

In the proof, we need the following (probably well-known) lemma, essentially stating that the densest subgraphs of Q_d are subcubes:

Lemma 3 *Let G be an m-vertex subgraph of some Q_d. Then $|E(G)| \leq \frac{1}{2} m \log m$.*

Proof. We proceed by induction on m. Let V_0 and V_1 be the parts of V in the bottom and top facets of Q_d, respectively, and let $m_i = |V_i|$. We may assume $m_0 \geq m_1 \geq 1$. The number of vertical edges between V_0 and V_1 is at most m_1, and by the inductive hypothesis, the total number of edges on V is at most $\frac{1}{2}(m_0 \log m_0 + m_1 \log m_1) + m_1$. This is a well-known recurrence and it is not hard to check that it implies the desired bound. $\square$

Proof of Claim 2. Let V and O as in the claim be fixed and let $\bar{O}$ be a USO2 extending O. To each such $\bar{O}$, we assign a string (encoding) of at most $f_d(k)$ bits in such a way that distinct $\bar{O}$ receive distinct strings.

Let V_0 and V_1 be the parts of V in the bottom and top facets of Q_d, and let $k_i = |V_i|$. For the purposes of the following description, we assume that $k_0 \geq k_1 \geq 1$. If $k_1 > k_0$, we exchange the roles of the top and bottom facets, and for $k_1 = 0$, we proceed directly to the bottom facet. (Note that we consider V, and thus k_0 and k_1, as given.)

Let $I \subseteq V_1$ consist of those vertices in V_1 whose neighbor in the bottom facet lies in V_0.

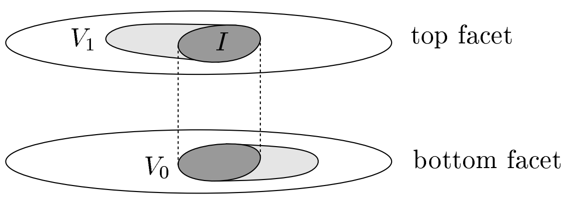

We define a subgraph $G = G(\bar{O})$ of the top facet. An edge $e = \{u, v\}$ of the top facet belongs to G if its orientation in $\bar{O}$ is opposite to the orientation of the corresponding edge $e' = \{u', v'\}$ in the bottom facet. The key observation is this: If $e = \{u, v\} \in E(G)$, then the vertical edges $\{v, v'\}$ and $\{u, u'\}$ have the same orientation (both upwards or both downwards) in $\bar{O}$, for otherwise, the 2-dimensional face $\{u, v, u', v'\}$ would not have a unique sink.

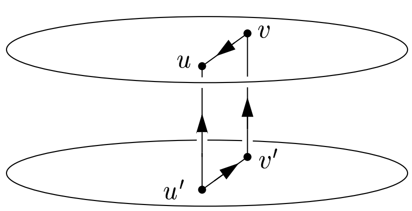

Consequently, if K is the vertex set of a connected component of G then all the vertical edges emanating from K have the same orientation.

Let us call a component of G *good* if it is completely contained in I and has at most $\log k$ vertices.

Now we are ready to describe the encoding of $\bar{O}$ given V and O.

1. We recursively encode the restriction of $\bar{O}$ on the bottom facet. By the inductive hypothesis, this requires at most

$$f_{d-1}(k_0)$$

bits.

2. Let $S \subseteq I$ be the union of the vertex sets of the good components of G, let $R = V_1 \setminus S$, and let $k_S = |S|$, $k_R = |R|$. We encode k_S and the set S; since there are $\binom{k_R+k_S}{k_S}$ possible subsets of V_1 of size k_S, we need at most

$$\log k_S + \log \binom{k_R + k_S}{k_S} + 3$$

bits (the extra 3 accounts for the case $k_S = 0$ and for rounding the logarithms up).

3. Now we specify the partition of S into vertex sets of the good components. To this end, we consider the graph G restricted to S and we choose some spanning forest F for it. This F has fewer than k_S edges. We consider the vertices of S one by one in the lexicographic order, and for each $v \in S$ we specify its neighbors in F. A neighbor u of v is given by the direction of the edge $\{u, v\}$, whose encoding needs at most $\log d$ bits, and one extra bit per edge can be used to distinguish when one should pass to the next vertex. The number of bits for encoding F is no more than

$$k_S(\log d + 2).$$

4. Knowing the vertex sets of the good components and the orientation of the bottom facet, we can reconstruct the orientation of all edges of Q_d on S connecting distinct components, and also of the edges between S and R. Now we consider a good component with vertex set K. By Lemma 3, the restriction of Q_d to K has at most $|K| \log |K| \leq |K| \log \log k$ edges. Summing over all good components, we have at

most $k_S \log\log k$ edges on S whose orientation we need to encode, and the required number of bits is at most

$$k_S \log\log k.$$

5. From the encoding produced up until now, we can reconstruct the orientation of all edges but those within R. By induction, the number of bits for encoding the orientation on R is no more than

$$f_{d-1}(k_R).$$

6. It remains to specify the orientation of all the vertical edges emanating from I. Since we know the orientation in both the top and bottom facets, we can reconstruct G completely. For each good component, we specify the orientation by one bit. If a component is not good, the it has at least one vertex v outside I, in which case we already know the orientation of the vertical edge incident to v, or it has at least $\log k$ vertices. The number of components of the latter type is no more than $\frac{k_R}{\log k}$. Altogether, we need at most

$$k_S + \frac{k_R}{\log k}$$

bits for specifying the orientation of the vertical edges. This finishes the description of the encoding.

We now collect all these contributions and obtain a recurrence for $f_d(k)$. For $k_S = 0$, we have only the terms $f_{d-1}(k_0) + f_{d-1}(k_1) + k_R/\log k + 3$.

Let us now assume $k_S \geq 1$. We have $\binom{k_R+k_S}{k_S} \leq \left(\frac{e(k_R+k_S)}{k_S}\right)^{k_S}$, and so $\log\binom{k_R+k_S}{k_S} \leq k_S(\log(1+\frac{k_R}{k_S})+2)$. If $1+\frac{k_R}{k_S} \leq \log^2 k$, then this expression is at most $2k_S(\log\log k+1)$, and otherwise, it does not exceed $k_S(\log k+2) \leq 2k_S + \frac{k_R}{\log k}$. The resulting recurrence is thus

$$\begin{aligned} f_d(k) \quad \leq \quad \max\Big\{ & f_{d-1}(k_0) + f_{d-1}(k_R) + \frac{2k_R}{\log k} + 3k_S\log\log k + k_S\log d + 6k_S + 3 : \\ & k_0 + k_R + k_S = k, k_0 \geq \tfrac{1}{2}k \Big\}. \end{aligned}$$

We may also assume that k is larger than a suitable constant, since $f_d(k)$ is always bounded (for example) by $\frac{1}{2}k\log k$ according to Lemma 3.

A simple calculation shows that the recurrence indeed implies the bound in Claim 2. The border case $k_R \le 1$ is easy and we omit it. We thus assume $k_R \ge 2$.

Let us consider the case $k_S \ge 1$. The right-hand side of the recurrence can be estimated by

$$f_{d-1}(k_0) + f_{d-1}(k_R) + \frac{2k_R}{\log k} + Bk_S(\log\log k + \log d) \qquad (1)$$

for a constant B. We note that since $k_R \le \frac{1}{2}k$, we have

$$\begin{aligned} k_R \log\log k_R &\le k_R \log(\log k - 1) = k_R\left(\log\log k + \log\left(1 - \frac{1}{\log k}\right)\right) \\ &\le k_R \log\log k - \frac{k_R \log e}{\log k} \end{aligned}$$

(using $\ln(1+x) \le x$). So, by induction, (1) can be estimated by

$$\begin{aligned} &Ck_0(\log\log k + \log d) + Ck_R(\log\log k + \log d) - \frac{Ck_R \log e}{\log k} + \frac{2k_R}{\log k} \\ &\quad + Bk_S(\log\log k + \log d) \\ &\le (\log\log k + \log d)(Ck_0 + Ck_R + Bk_S) \le Ck(\log\log k + \log d) \end{aligned}$$

as required. For $k_S = 0$, the calculation is similar but simpler; the $+3$ term can be absorbed into $-\frac{Ck_R \log e}{\log k}$. This finishes the proof of Claim 2 and thus Theorem 1, too, is proved. □

Acknowledgment

I would like to thank Uli Wagner, Tibor Szabó, and Ingo Schurr for extensive and fruitful discussions, and to Bernd Gärtner and Emo Welzl for organizing an excellent workshop and inviting me to it.