

TWO NEW BOUNDS FOR THE RANDOM-EDGE SIMPLEX ALGORITHM

BERND GÄRTNER AND VOLKER KAIBEL

ABSTRACT. We prove that the RANDOM-EDGE simplex algorithm requires an expected number of at most $13n/\sqrt{d}$ pivot steps on any simple d -polytope with n vertices. This is the first nontrivial upper bound for general polytopes. We also describe a refined analysis that potentially yields much better bounds for specific classes of polytopes. As one application, we show that for combinatorial d -cubes, the trivial upper bound of 2^d on the performance of RANDOM-EDGE can asymptotically be improved by any desired polynomial factor in d .

1. INTRODUCTION

Dantzig's *simplex method* [8] is a widely used tool for solving linear programs (LP). The feasible region of an LP is a polyhedron; any algorithm implementing the simplex method traverses a sequence of vertices, such that (i) consecutive vertices are equal (the *degenerate* case) or connected by a polyhedron edge, and (ii) the objective function strictly improves along any traversed edge. In both theory and practice, we may assume that some initial vertex is available, and that the optimal solution to the LP is attained at a vertex, if there is an optimum at all. It follows that if the algorithm does not cycle, it will eventually find an optimal solution, or discover that the problem is unbounded (see e.g. Chvátal's book [7] for a comprehensive introduction to the simplex method).

For most (complexity-)theoretic investigations, one can safely assume that the LP's that are considered are bounded as well as both primally and dually non-degenerate [19]. Thus, we will only deal with *simple polytopes*, i.e., bounded d -dimensional polyhedra, where at each vertex exactly d facets meet, and with objective functions that are non-constant along any edge of the polytope.

The distinguishing feature of each simplex-algorithm is the *pivot rule* according to which the next vertex in the sequence is selected in case there is a choice. Many popular pivot rules are efficient in practice, meaning that they induce a short vertex sequence in typical applications. The situation in theory is in sharp contrast to this: Among most of the *deterministic* pivot rules proposed in the literature (including the ones widely used in practice), the simplex algorithm is forced to traverse exponentially (in the number of variables and constraints of the LP) many vertices in the worst case. It is open whether there is a pivot rule that always induces a sequence of polynomial length.

To explain simplex's excellent behavior in practice, the tools of *average case analysis* [5] and *smoothed analysis* [20] have been devised, and to conquer the worst case bounds, research has turned to *randomized* pivot rules. Indeed, Kalai [13, 14] as well as Matoušek, Sharir and Welzl [17] could prove that the expected number of steps taken by the RANDOM-FACET pivot rule is only *subexponential* in the worst case. These results hold under our above assumption that the feasible region of the LP is a simple and full-dimensional polytope.

The first author acknowledges support from the Swiss Science Foundation (SNF), Project No. 200021-100316/1. The second author has been supported by the DFG Research Center MATHEON and by the DFG Research Group *Algorithms, Structure, Randomness*.

Much less is known about another (actually, the most natural) randomized pivot rule: choose the next vertex in the sequence uniformly at random among the neighbors of the current vertex with better objective function value. This rule is called RANDOM-EDGE, and unlike RANDOM-FACET, it has no recursive structure to peg an analysis to. Nontrivial upper bounds on its expected number of pivot steps on general polytopes do not exist. Results are known for 3-polytopes [6, 12], d -polytopes with $d + 2$ facets [9], and for linear assignment problems [21]. Only recently, Pemantle and Balogh solved the long standing problem of finding a tight bound for the expected performance of RANDOM-EDGE on the d -dimensional *Klee-Minty* cube [3]. This polytope is the 'mother' of many worst-case inputs for deterministic pivot rules [15, 2].

None of the existing results exclude the possibility of both RANDOM-FACET and RANDOM-EDGE being the desired (expected) polynomial-time pivot rules. In the more general and well-studied setting of *abstract objective functions* on polytopes [1, 22, 23, 14], superpolynomial lower bounds are known for both rules, where the construction for RANDOM-EDGE [18] is very recent and much more involved than the one for RANDOM-FACET [16]. Both approaches inherently use objective functions (on cubes) that are not linearly induced.

In this paper, we derive the first nontrivial upper bound for the expected performance of RANDOM-EDGE on simple polytopes, with edge orientations induced by abstract objective functions. Even when we restrict to linear objective functions on combinatorial cubes, the result is new. The general bound itself is rather weak and also achieved for example by the deterministic GREATEST-DECREASE rule. The emphasis here is on the fact that we are able to make progress at all, given that RANDOM-EDGE has turned out to be very difficult to attack in the past. Also, our new bound separates RANDOM-EDGE from many deterministic rules (for example, *Dantzig's rule*, *Bland's rule*, or the *shadow vertex rule*) that may visit all vertices in the worst case [2].

In a second part, we refine the analysis, with the goal of obtaining better bounds for specific classes of polytopes. Roughly speaking, these are polytopes with large and regular local neighborhoods. Our prime example is the class of combinatorial cubes, for which we improve the general upper bound by any desired polynomial factor in the dimension. As before, this also works for abstract objective functions and thus complements the recent lower bound of Matoušek and Szabó [18] with a first nontrivial upper bound.

2. A BOUND FOR GENERAL POLYTOPES

Throughout this section, P is a d -dimensional simple polytope with a set V of n vertices. A directed graph $D = (V, A)$ is called an *acyclic unique sink orientation* (AUSO) of P if

- (i) its underlying undirected graph is the vertex-edge graph of P ,
- (ii) D is acyclic, and
- (iii) any subgraph of D induced by the vertices of a nonempty face of P has a unique sink.

Any linear function $\varphi : V \rightarrow \mathbb{R}$ that is *generic* (non-constant on edges of P) induces an AUSO in a natural way: there is a directed edge $v \rightarrow w$ between adjacent vertices if and only if $\varphi(v) > \varphi(w)$. The global sink of the AUSO is the unique vertex that minimizes φ over P . If φ is any generic (not necessarily linear) function inducing an AUSO that way, φ is called an *abstract objective function*. For a given AUSO D of P , any function φ that maps vertices to their ranks w.r.t. a fixed topological sorting of D is an abstract objective function that induces D . In general, D need not be induced by a linear function, for example if D fails to satisfy the necessary *Holt-Klee* condition for linear realizability [11]. For the remainder of this section,

we fix an AUSO D of P , an abstract objective function φ that induces D and some vertex $s \in V$.

Let π be the random variable defined as the directed path in D , starting at s and ending at the sink v_{opt} of D , induced by the RANDOM-EDGE pivot rule. From each visited vertex $v \neq v_{\text{opt}}$, π proceeds to a neighbor w of v , along an outgoing edge chosen uniformly at random from all outgoing edges.

For each $v \in V$, denote by

$$\text{out}(v) := \{w \in V : (v, w) \in A\}$$

the set of all smaller (w.r.t. φ) neighbors of v . If $|\text{out}(v)| = k$, then v is called a k -vertex. We denote by V_k the set of all k -vertices.

For every vertex $v \neq v_{\text{opt}}$ on the path π let v' be its successor on π . We denote by

$$S(v) := \{w \in \text{out}(v) : \varphi(v') < \varphi(w)\},$$

the set of neighbors of v that are 'skipped' by π at the step from v to v' . For every $0 \leq k \leq d$ let

$$\eta_k(\pi) := |\{v \in \pi \cap V_k : |S(v)| \geq \lfloor \frac{|\text{out}(v)|}{2} \rfloor\}|$$

be the number of k -vertices on π , where π skips at least $\lfloor \frac{k}{2} \rfloor$ neighbors. (Here, as in the following, we write, depending on the context, ' π ' for the set of vertices on the path π .)

If we denote by $n_k(\pi)$ the total number of k -vertices on the path π , then we obtain

$$(1) \quad \mathbb{E}[\eta_k(\pi)] \geq \frac{1}{2} \mathbb{E}[n_k(\pi)].$$

Indeed, we have

$$\mathbb{E}[\eta_k(\pi)] = \sum_{v \in V_k} \mathbb{P}[v \in \pi \text{ and } |S(v)| \geq \lfloor \frac{k}{2} \rfloor]$$

and

$$\mathbb{E}[n_k(\pi)] = \sum_{v \in V_k} \mathbb{P}[v \in \pi].$$

The claim then follows from

$$\mathbb{P}[|S(v)| \geq \lfloor \frac{|\text{out}(v)|}{2} \rfloor \mid v \in \pi] \geq \frac{1}{2}.$$

Due to $\varphi(v) > \varphi(w) > \varphi(v')$ for all $w \in S(v)$, the sets $S(v)$ are pairwise disjoint. Thus, we obtain (exploiting the linearity of expectation) for the number $\text{length}(\pi)$ of vertices on π

$$\mathbb{E}[\text{length}(\pi)] \leq n - \sum_{k=0}^d \mathbb{E}[\eta_k(\pi)] \lfloor \frac{k}{2} \rfloor \leq n - \sum_{k=0}^d \frac{1}{2} \lfloor \frac{k}{2} \rfloor \mathbb{E}[n_k(\pi)]$$

(where we used (1) for the second inequality). Clearly, we have $\mathbb{E}[\text{length}(\pi)] = \sum_{k=0}^d \mathbb{E}[n_k(\pi)]$. Therefore, we obtain (note $\frac{1}{2} \lfloor \frac{k}{2} \rfloor \geq \frac{k-1}{4}$)

$$(2) \quad \mathbb{E}[\text{length}(\pi)] \leq \min \left\{ \sum_{k=0}^d \mathbb{E}[n_k(\pi)], n - \sum_{k=0}^d \frac{k-1}{4} \mathbb{E}[n_k(\pi)] \right\}.$$

If h_k denotes the total number of k -vertices in V , then we clearly have $0 \leq \mathbb{E}[n_k(\pi)] \leq h_k$. Thus, (2) yields

$$(3) \quad \mathbb{E}[\text{length}(\pi)] \leq \max \left\{ \min \left\{ \sum_{k=0}^d x_k, n - \sum_{k=0}^d \frac{k-1}{4} x_k \right\} : 0 \leq x_k \leq h_k \text{ for all } k \right\}.$$

In (3), the maximum must be attained by some $x \in \mathbb{R}^{d+1}$ for which the minimum is attained by both $\sum x_k$ and $n - \sum \frac{k-1}{4} x_k$. Indeed, if $\sum x_k < n - \sum \frac{k-1}{4} x_k$ then not all x_k can be at their respective upper bounds h_k (since $n = \sum h_k$), thus one of them

can slightly be increased in order to increase the minimum. If $\sum x_k > n - \sum \frac{k-1}{4}x_k$ then not all x_k can be zero (since this would yield $0 > n$), so one of them can be decreased in order to increase the minimum. Thus we conclude

$$(4) \quad \mathbb{E}[\text{length}(\pi)] \leq \max \left\{ \sum_{k=0}^d x_k : \sum_{k=0}^d \frac{k+3}{4}x_k = n, 0 \leq x_k \leq h_k \text{ for all } k \right\}.$$

By (weak) linear programming duality (and exploiting $n = \sum_{k=0}^d h_k$ once more), we can derive from (4) the estimate

$$(5) \quad \mathbb{E}[\text{length}(\pi)] \leq \sum_{k=0}^d h_k \cdot \max\{y, 1 - \frac{k-1}{4}y\}$$

for every $y \in \mathbb{R}$.

In the sequel, we need two important results from the theory of convex polytopes. The parameters h_k are independent of the actual acyclic unique sink orientation of the polytope. The h -vector formed by them is a linear transformation of the f -vector of the polytope, storing for each i the number of i -dimensional faces of the polytope.

The first classical result we need are the *Dehn-Sommerville equations*

$$(6) \quad h_k = h_{d-k} \quad \text{for all } 0 \leq k \leq d$$

(see [24, Sect. 8.3]). The second one is the *unimodality of the h -vector*:

$$(7) \quad h_0 \leq h_1 \leq \dots \leq h_{\lfloor d/2 \rfloor}$$

The latter is equivalent to the *nonnegativity of the g -vector*, which is one of the hard parts of the *g -theorem for simplicial polytopes*, see [24, Sect. 8.6].

From (6) and (7) we can derive

$$n = \sum_{k=0}^d h_k \geq (d - 8\sqrt{d})h_{\lfloor 4\sqrt{d} \rfloor},$$

which yields (for $d > 64$)

$$(8) \quad h_{\lfloor 4\sqrt{d} \rfloor} \leq \frac{n}{d - 8\sqrt{d}}.$$

Now we choose $y := 1/\sqrt{d}$ in (5). We have

$$\frac{1}{\sqrt{d}} \geq 1 - \frac{k-1}{4\sqrt{d}} \Leftrightarrow k \geq 4\sqrt{d} - 3.$$

Thus, (5) (with $y = 1/\sqrt{d}$) gives

$$(9) \quad \mathbb{E}[\text{length}(\pi)] \leq \sum_{k=0}^{\lfloor 4\sqrt{d}-3 \rfloor} h_k \left(1 - \frac{k-1}{4\sqrt{d}}\right) + \sum_{k=\lfloor 4\sqrt{d}-3 \rfloor+1}^d \frac{h_k}{\sqrt{d}}.$$

By the unimodality of the h -vector and (8), the first sum in (9) can be estimated by

$$4\sqrt{d} \cdot h_{\lfloor 4\sqrt{d} \rfloor} \leq \frac{4n}{\sqrt{d}-8} \leq \frac{12n}{\sqrt{d}}, \quad d \geq 144.$$

Clearly, the second sum in (9) is bounded by n/\sqrt{d} . The resulting total bound of $13n/\sqrt{d}$ also holds for $d < 144$, because n is a trivial upper bound. Thus we have proved the following result.

Theorem 1. *The expected number of vertices visited by the RANDOM-EDGE simplex algorithm on a d -dimensional simple polytope with $n \geq d+1$ vertices, equipped with an abstract (in particular: a linear) objective function is bounded by*

$$13 \cdot \frac{n}{\sqrt{d}} .$$

A similar analysis reveals that the running-time for the GREATEST-DECREASE-rule is bounded by

$$C' \cdot \frac{n}{\sqrt{d}} .$$

In each step, this rule selects the neighboring vertex with *smallest* φ -value, thus skipping all other neighbors of the current vertex v .

For general simple polytopes, our analysis of the bound for RANDOM-EDGE stated in (3) is essentially best possible. This can be seen through the examples of duals of stacked simplicial polytopes (see, e.g., [4]), which are simple d -polytopes with n vertices, $h_0 = h_d = 1$, and $h_k = \frac{n-2}{d-1}$ for all $1 \leq k \leq d-1$.

3. A BOUND FOR CUBES

The core argument of the analysis presented in Section 2 is the following: For every vertex on the RANDOM-EDGE path π with out-degree k we know that π skips (in expectation) $k/2$ vertices *in the single step from v to its successor*. We then exploited the Dehn-Sommerville equations as well as the unimodality of the h -vector in order to argue that many vertices on π must have large out-degree – unless π is ‘short’ anyway.

For the d -dimensional cube, we have much more information on the h -vector: $h_k = \binom{d}{k}$ for every k . Thus, ‘most’ vertices have out-degree roughly $d/2$ in case of cubes. We will exploit this stronger knowledge in a sharper analysis for cubes, which relies on studying larger structures around vertices than just their out-neighbors. We actually do the analysis for general polytopes and obtain a bound on the expected path length in terms of two specific quantities. Later we bound these quantities for the case of cubes.

3.1. The General Approach. Within this subsection, (as in Section 2), let P be a d -dimensional simple polytope with n vertices V , $D = (V, A)$ an AUSO of P , $\varphi : V \rightarrow \mathbb{R}$ an abstract objective function inducing D , and $s \in V$ a fixed vertex. We denote by $\text{dist}^\rightarrow(v, w)$ the length (number of arcs) of a shortest directed path from v to w ($\text{dist}^\rightarrow(v, w)$ may be ∞ if there is no such path).

Definition 1 (t -reach). *Let $t, k \in \mathbb{N}$ and $v \in V$.*

(1) *We call*

$$R_t(v) := \{w \in V : \text{dist}^\rightarrow(v, w) \leq t\}$$

the t -reach of v . The boundary of $R_t(v)$, denoted by $\partial R_t(v)$, is the set of all $w \in R_t(v)$, for which there is a directed (not necessarily shortest) path of length precisely t from v to w .

(2) *The t -reach $R_t(v)$ is k -good if*

$$|\text{out}(w)| \geq k$$

holds for all $w \in R_t(v)$ with $\text{dist}^\rightarrow(v, w) \leq t-1$.

(3) *A vertex v is (t, k) -good if its t -reach is k -good. The set of all (t, k) -good vertices is denoted by $G(t, k)$.*

In particular, if v is (t, k) -good, the optimal vertex v_{opt} may occur in the boundary of $R_t(v)$, but not in its interior. For $t, k \in \mathbb{N}$, we define

$$g(t, k) := \min\{|\partial R_t(v)| : v \in G(t, k)\} .$$

For every vertex $v \in V$, and some $t \in \mathbb{N}$, denote by (the random variable) $w_t(v)$ the vertex that is reached by the RANDOM-EDGE simplex-algorithm, started at v , after t steps (let $w_t(v) := v_{\text{opt}}$ in case the sink is reached before step t). Generalizing the notion from Section 2, we denote by

$$\tilde{S}_t(v) := \{u \in R_t(v) : \varphi(u) > \varphi(w_t(v))\}$$

the set of vertices in $R_t(v)$ left behind while walking from v to $R_t(v)$.

Lemma 1. *For every $t, k \in \mathbb{N}$ and $v \in G(t, k)$, we have*

$$\mathbb{P}[|\tilde{S}_t(v)| \geq \frac{g(t, k)}{2}] \geq \frac{g(t, k)}{2d^t}.$$

Proof. Let $\partial R_t(v) = \{u_1, \dots, u_q\}$ with $\varphi(u_1) > \dots > \varphi(u_q)$. By construction, there is some i^* with $w_t(v) = u_{i^*}$. Since the outdegree at every vertex is at most d , we have

$$\mathbb{P}[i^* = i] \geq \frac{1}{d^t}$$

for every $1 \leq i \leq q$. Therefore,

$$\mathbb{P}[i^* > q/2] \geq \frac{q}{2d^t}$$

holds. Since $q \geq g(t, k)$ holds and because $i^* > g(t, k)/2$ implies $|\tilde{S}_t(v)| \geq g(t, k)/2$, the claim follows. \square

Now let us consider the path π followed by the RANDOM-EDGE simplex-algorithm started at s (ending in v_{opt}). For $t, k \in \mathbb{N}$ with $t \geq 2$ and $k \geq 1$, we subdivide π into subpaths with the property that every subpath either has length one and starts at a non- (t, k) -good vertex or it has length t (a *long subpath*) and starts at a (t, k) -good vertex. (Such a partitioning is clearly possible.)

Let $n_{t, k}(\pi)$ be the number of long subpaths in our partitioning. We denote the pairs of start and end vertices of these long paths by $(x_1, y_1), \dots, (x_{n_{t, k}(\pi)}, y_{n_{t, k}(\pi)})$. Let

$$S_t(x_i) := \{u \in R_t(x_i) : \varphi(u) > \varphi(y_i)\}$$

and define

$$\eta_{t, k}(\pi) := |\{i \in \{1, \dots, n_{t, k}(\pi)\} : |S_t(x_i)| \geq \frac{g(t, k)}{2}\}|$$

to be the number of those long subpaths which leave behind at least $\frac{g(t, k)}{2}$ vertices from $R_t(x_i)$.

Using Lemma 1 (note that $S_t(x_i)$, conditioned on the event that x_i is the start vertex of a long subpath in the partitioning of π , has the same distribution as $\tilde{S}_t(x_i)$), we can deduce, similarly to our derivation of (1), the following:

$$(10) \quad \mathbb{E}[\eta_{t, k}(\pi)] \geq \frac{g(t, k)}{2d^t} \mathbb{E}[n_{t, k}(\pi)].$$

Also here, the sets $S_t(x_i)$ (for $1 \leq i \leq n_{t, k}(\pi)$) are pairwise disjoint. Thus, for each long subpath (consisting of t arcs) starting at some x_i with $|S_t(x_i)| \geq g(t, k)/2$ we can count at least $g(t, k)/2 - t$ vertices that are not visited by π . Therefore, we can conclude

$$\mathbb{E}[\text{length}(\pi)] \leq n - \left(\frac{g(t, k)}{2} - t\right) \mathbb{E}[\eta_{t, k}(\pi)].$$

Using (10) and defining

$$\tilde{g}(t, k) := \left(\frac{g(t, k)}{2} - t\right) \frac{g(t, k)}{2d^t},$$

this yields

$$(11) \quad \mathbb{E}[\text{length}(\pi)] \leq n - \tilde{g}(t, k) \mathbb{E}[n_{t, k}(\pi)].$$

On the other hand, denote by

$$(12) \quad f(t, k) := |V \setminus G(t, k)|$$

the total number of non- (t, k) -good vertices. From the definition of our path partitioning, we immediately obtain

$$(13) \quad \mathbb{E}[\text{length}(\pi)] \leq f(t, k) + t \cdot \mathbb{E}[n_{t,k}(\pi)] .$$

Adding up nonnegative multiples of (11) and (13) in such a way that $\mathbb{E}[n_{t,k}(\pi)]$ cancels out, one obtains the following bound:

$$\mathbb{E}[\text{length}(\pi)] \leq \frac{tn + \tilde{g}(t, k)(f(t, k))}{\tilde{g}(t, k) + t} \leq \frac{t}{\tilde{g}(t, k)}n + f(t, k)$$

This yields the following estimation.

Lemma 2. *For $t, k \in \mathbb{N}$ with $t \geq 2$ and $k \geq 1$, we have*

$$\mathbb{E}[\text{length}(\pi)] \leq \frac{4td^t}{g(t, k)(g(t, k) - 2t)}n + f(t, k) .$$

A general way to bound the function $f(t, k)$ is as follows.

Lemma 3. *For $t, k \in \mathbb{N}$, we have*

$$f(t, k) \leq \frac{d^t - 1}{d - 1} h_{<k} ,$$

where $h_{<k} := \sum_{j=0}^{k-1} h_j$ is the number of vertices with outdegree less than k .

Proof. If $v \in V \setminus G(t, k)$, then there is some $w \in R_{t-1}(v)$ with $|\text{out}(w)| < k$. On the other hand, each w is contained in at most $\sum_{i=0}^{t-1} d^i = \frac{d^t - 1}{d - 1}$ $(t - 1)$ -reaches (since the undirected graph is d -regular). The claim follows. \square

The following describes a way of bounding the function $g(t, k)$ by studying the undirected graph of the polytope.

Definition 2 ((t, k) -neighborhood, $\gamma(t, k)$). *Let $t, k \in \mathbb{N}$.*

- (1) *A subset $N \subset V$ is called a (t, k) -neighborhood of $v \in V$ if $N = \{v\}$ in case of $t = 0$, or, if $t \geq 1$, there are k neighbors w_1, \dots, w_k of v in the graph of P together with $(t - 1, k)$ -neighborhoods N_1, \dots, N_k of w_1, \dots, w_k , respectively, such that $N = \bigcup_{i=1}^k N_i$.*
- (2) *We define $\gamma(t, k)$ as the minimum cardinality of $\{w \in N : \text{dist}(v, w) = t\}$, taken over all $v \in V$ and all (t, k) -neighborhoods of v . (Here, $\text{dist}(v, w)$ denotes the graph-theoretical distance between v and w in the undirected graph of P .)*

If v is (t, k) -good, then it follows right from the definitions that the boundary $\partial R_t(v)$ of its t -reach contains a (t, k) -neighborhood N of v . In particular, all vertices $w \in N$ with $\text{dist}(v, w) = t$ are in $\partial R_t(v)$, and these are the ones of use to us.

Lemma 4. *For $t, k \in \mathbb{N}$ with $t \geq 2$, we have*

$$g(t, k) \geq \gamma(t, k) .$$

3.2. Specialization to Cubes. In order to obtain from Lemma 2 an explicit bound for the expected number of vertices visited by the RANDOM-EDGE simplex-algorithm on the d -cube, we will derive estimates on the functions $f(t, k)$ and $g(t, k)$ for $k = \lfloor \frac{d}{4} \rfloor$.

Lemma 5. *There is a constant $0 < \alpha < 1$ such that*

$$f(t, \lfloor \frac{d}{4} \rfloor) \leq 2^{\alpha d + o(d)}$$

holds for all $t \in \mathbb{N}$ (where f is the function defined in (12) for the case of the d -cube, and with $k = \lfloor \frac{d}{4} \rfloor$).

Proof. In the case of a d -cube and $k = \lfloor \frac{d}{4} \rfloor$, we have

$$h_{<k} = \sum_{i=0}^{\lfloor \frac{d}{4} \rfloor - 1} \binom{d}{i} = 2^{h(\frac{1}{4})d + o(d)},$$

where $h(x) = x \log \frac{1}{x} + (1-x) \log \frac{1}{1-x}$ is the binary entropy function (see, e.g., [10, Chap. 9, Ex. 42]). By Lemma 3 this implies the claimed bound (with the $o(d)$ term depending on t). \square

The final building block of our bound for the special case of cubes is the following. Here, we denote by $a^{\underline{b}}$ (*falling factorial power*) the product $a(a-1)\cdots(a-b+1)$ (for $a, b \in \mathbb{N}$).

Lemma 6. *Let $t, k \in \mathbb{N}$ with $1 \leq t, k \leq d$. If the polytope P considered in Section 3.1 is a d -cube, then the following is true:*

- (1) $\gamma(t, k) \geq \frac{k^t}{t!} - \sum_{i=1}^{t-1} \frac{k^i}{t^i} \binom{d-1}{t-i-1}$.
- (2) *If t is a constant, then $\gamma(t, \lfloor \frac{d}{4} \rfloor) = \Omega(d^t)$.*

Proof. Part (2) follows immediately from part (1), since the sum becomes a polynomial in d of degree $t-1$ for $k = \frac{d}{4}$ (and constant t).

Let us prove (1) for each fixed k , by induction on t , where the case $t=1$ holds due to $\gamma(1, k) = k$. Thus, let us consider the case $t \geq 2$.

We may assume that the vertex v and its neighbors w_1, \dots, w_k , for which the minimum $\gamma(t, k)$ is attained, are $v = \mathbf{0}$ and $w_i = \mathbf{e}_i$ ($1 \leq i \leq k$). For each i , the $(t-1, k)$ neighborhood N_i of \mathbf{e}_i has at least $\gamma(t-1, k)$ vertices w with $\text{dist}(\mathbf{e}_i, w) = t-1$, by definition. All of them have distance $t-2$ or t from $\mathbf{0}$. The former may be the case at most $\binom{d-1}{t-2}$ times (these vertices cannot have a one at position i). Therefore, we have

$$|\{w \in N_i : \text{dist}(\mathbf{0}, w) = t\}| \geq \gamma(t-1, k) - \binom{d-1}{t-2}.$$

On the other hand, every vertex $w \in N_i$ with $\text{dist}(\mathbf{0}, w) = t$ needs to have a one at position i (otherwise, $\text{dist}(\mathbf{e}_i, w) = t+1$). Hence, every vertex w with $\text{dist}(\mathbf{0}, w) = t$ can be contained in at most t of the neighborhoods N_1, \dots, N_k . Thus, we conclude (for $t \geq 2$)

$$\gamma(t, k) \geq \frac{k(\gamma(t-1, k) - \binom{d-1}{t-2})}{t},$$

and thus,

$$(14) \quad \gamma(t, k) \geq \frac{k}{t} \gamma(t-1, k) - \frac{k \binom{d-1}{t-2}}{t}.$$

Using the induction hypothesis and (14) we derive

$$\begin{aligned} \gamma(t, k) &\geq \frac{k}{t} \left(\frac{k^{t-1}}{(t-1)!} - \sum_{i=1}^{t-2} \frac{k^i}{(t-1)^i} \binom{d-1}{t-i-2} \right) - \frac{k}{t} \binom{d-1}{t-2} \\ &= \frac{k^t}{t!} - \sum_{i=1}^{t-2} \frac{k^{i+1}}{t^{i+1}} \binom{d-1}{t-i-2} - \frac{k^{0+1}}{t^{0+1}} \binom{d-1}{t-0-2} \\ &= \frac{k^t}{t!} - \sum_{i=0}^{t-2} \frac{k^{i+1}}{t^{i+1}} \binom{d-1}{t-i-2}, \end{aligned}$$

which, after an index shift in the sum, yields the claim. \square

Now we can prove our main result:

Theorem 2. *For every fixed $t \in \mathbb{N}$, there is a constant $C_t \in \mathbb{R}$ (depending on t), such that the expected number of vertices visited by the RANDOM-EDGE simplex-algorithm on a d -dimensional cube, equipped with an abstract (in particular: a linear) objective function, is bounded by*

$$C_t \cdot \frac{2^d}{d^t} .$$

Proof. Let π be the (random) path (for some arbitrary start vertex) defined by the RANDOM-EDGE simplex-algorithm on a d -cube equipped with an acyclic unique sink orientation. By Lemma 2, we have, with $d' := \lfloor \frac{d}{4} \rfloor$,

$$(15) \quad \mathbb{E}[\text{length}(\pi)] \leq \frac{4td^t}{g(t, d')(g(t, d') - 2t)} 2^d + f(t, d') .$$

From Lemma 5 we know that there is some constant $0 < \alpha < 1$ with

$$(16) \quad f(t, d') \leq 2^{\alpha d + o(d)} .$$

Finally, by Lemmas 4 and 6 (2) there is some constant $\beta > 0$ such that

$$(17) \quad g(t, d') \geq \beta d^t .$$

Putting (15), (16), and (17) together, we obtain

$$\mathbb{E}[\text{length}(\pi)] \leq \frac{4td^t}{\beta^2 d^{2t} - 2t\beta d^t} 2^d + 2^{\alpha d + o(d)} ,$$

which implies the claim. \square

4. CONCLUSION

Probably one can extend the methods we have used for analyzing RANDOM-EDGE on cubes to other classes of polytopes (e.g., general products of simplices). However, it seems to us that it would be more interesting to find a way of sharpening our bounds by enhancing our approach with some new ideas. As mentioned at the end of Section 2, the analysis of our approach is sharp in the general setting. We suspect that one cannot prove a subexponential bound for RANDOM-EDGE on cubes with our methods. Therefore, it would be most interesting to find a way of combining our kind of analysis with some other ideas.

REFERENCES

- [1] Ilan Adler and Romesh Saigal. Long monotone paths in abstract polytopes. *Math. Operations Research*, 1(1):89–95, 1976.
- [2] Nina Amenta and Günter Ziegler. Deformed products and maximal shadows of polytopes. In J. Chazelle, J.E. Goodman, and R. Pollack, editors, *Advances in Discrete and Computational Geometry*, volume 223 of *Contemporary Mathematics*, pages 57–90. Amer. Math. Soc., 1999.
- [3] Jozsef Balogh and Robin Pemantle. The Klee-Minty random edge chain moves with linear speed. Manuscript, 2004.
- [4] Margaret M. Bayer and Carl W. Lee. Combinatorial aspects of convex polytopes. In *Handbook of convex geometry, Vol. A, B*, pages 485–534. North-Holland, Amsterdam, 1993.
- [5] Karl Heinz Borgwardt. *The Simplex Method: A Probabilistic Analysis*, volume 1 of *Algorithms and Combinatorics*. Springer-Verlag, New York, 1987.
- [6] Andrei Z. Broder, Martin E. Dyer, Alan M. Frieze, Prabhakar Raghavan, and Eli Upfal. The worst-case running time of the random simplex algorithm is exponential in the height. *Information Processing Letters*, 56(2):79–81, 1995.
- [7] Vášek Chvátal. *Linear Programming*. W. H. Freeman, New York, NY, 1983.
- [8] George B. Dantzig. *Linear Programming and Extensions*. Princeton University Press, Princeton, New Jersey, 1963.
- [9] Bernd Gärtner, József Solymosi, Falk Tschirschnitz, Pavel Valtr, and Emo Welzl. One line and n points. *Random Structures & Algorithms*, 23(4):453–471, 2003.
- [10] Ronald L. Graham, Donald E. Knuth, and Oren Patashnik. *Concrete Mathematics*. Addison-Wesley, Reading, MA, second edition, 1994.

- [11] Fred B. Holt and Victor Klee. A proof of the strict monotone 4-step conjecture. In J. Chazelle, J.E. Goodman, and R. Pollack, editors, *Advances in Discrete and Computational Geometry*, volume 223 of *Contemporary Mathematics*, pages 201–216. Amer. Math. Soc., 1998.
- [12] Volker Kaibel, Rafael Mechtel, Micha Sharir, and Günter M. Ziegler. The simplex algorithm in dimension three. *SIAM J. Computing*, to appear.
- [13] Gil Kalai. A subexponential randomized simplex algorithm. In *Proc. 24th Annu. ACM Sympos. Theory Comput.*, pages 475–482, 1992.
- [14] Gil Kalai. Linear programming, the simplex algorithm and simple polytopes. *Mathematical Programming Ser. B*, 79:217–233, 1997.
- [15] Victor Klee and George J. Minty. How good is the simplex algorithm? In O. Shisha, editor, *Inequalities III*, pages 159–175. Academic Press, 1972.
- [16] Jiří Matoušek. Lower bounds for a subexponential optimization algorithm. *Random Structures & Algorithms*, 5(4):591–607, 1994.
- [17] Jiří Matoušek, Micha Sharir, and Emo Welzl. A subexponential bound for linear programming. *Algorithmica*, 16(4/5):498–516, October/November 1996.
- [18] Jiří Matoušek and Tibor Szabó. RANDOM EDGE can be exponential on abstract cubes. In *Proc. 45th Ann. IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 92–100, 2004.
- [19] Alexander Schrijver. *Theory of Linear and Integer Programming*. Wiley-Interscience, New York, NY, 1986.
- [20] Daniel Spielman and Shang-Hua Teng. Smoothed analysis of algorithms: why the simplex algorithm usually takes polynomial time. In ACM, editor, *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing*, pages 296–305, New York, NY, USA, 2001. ACM Press.
- [21] Craig A. Tovey. Low order polynomial bounds on the expected performance of local improvement algorithms. *Math. Programming*, 35(2):193–224, 1986.
- [22] Doug Wiedemann. Unimodal set-functions. *Congressus Numerantium*, 50:165–169, 1985.
- [23] Kathy Williamson Hoke. Completely unimodal numberings of a simple polytope. *Discrete Applied Mathematics*, 20:69–81, 1988.
- [24] Günter M. Ziegler. *Lectures on Polytopes*, volume 152 of *Graduate Texts in Mathematics*. Springer-Verlag, Heidelberg, 1994.

INSTITUT FÜR THEORETISCHE INFORMATIK, ETH ZÜRICH, HALDENEGGSTEIG 4, CH-8092 ZÜRICH
E-mail address: `gaertner@inf.ethz.ch`

ZUSE-INSTITUT BERLIN, TAKUSTR. 7, D-14195 BERLIN
E-mail address: `kaibel@zib.de`