# Paths Beyond Local Search: A Tight Bound for Randomized Fixed-Point Computation* 

Xi Chen ${ }^{\dagger}$<br>Department of Computer Science<br>Tsinghua University

Shang-Hua Teng ${ }^{\ddagger}$<br>Department of Computer Science<br>Boston University


#### Abstract

In 1983, Aldous proved that randomization can speedup local search. For example, it reduces the query complexity of local search over grid $[1: n]^{d}$ from $\Theta\left(n^{d-1}\right)$ to $O\left(d^{1 / 2} n^{d / 2}\right)$. It remains open whether randomization helps fixed-point computation. Inspired by the recent advances on the complexity of equilibrium computation, we solve this open problem by giving an asymptotically tight bound of $(\Omega(n))^{d-1}$ on the randomized query complexity for computing a fixed point of a discrete Brouwer function over grid $[1: n]^{d}$. Our result can be extended to the black-box query model for Sperner's Lemma in any dimension. It also yields a tight bound for the computation of d-dimensional approximate Brouwer fixed points as defined by Scarf and by Hirsch, Papadimitriou, and Vavasis.

Since the randomized query complexity of global optimization over $[1: n]^{d}$ is $\Theta\left(n^{d}\right)$, the randomized query model over $[1: n]^{d}$ strictly separates these three important search problems:


> Global optimization is harder than fixed-point computation, and fixed-point computation is harder than local search.

Our result indeed demonstrates that randomization does not help much in fixed-point computation in the black-box query model. Our randomized lower bound matches the deterministic complexity of this problem, which is $\Theta\left(n^{d-1}\right)$.

[^0]
## 1 Introduction

In this paper, we prove that randomization cannot be used to speedup fixed-point computation (FPC) in the black-box query model. Our result draws a striking contrast between FPC and local search: In 1983, Aldous proved that randomization can significantly speedup local search [2]. Our randomized lower bound is asymptotically tight - it matches the deterministic upper bound $[18,6]$ for this problem. Our result resolves a question that has been open since the 1989 paper of Hirsch, Papadimitriou, and Vavasis [18] which introduced the black-box query model for FPC.

## Motivation

The Simplex Algorithm [13] is an implementation of local search ${ }^{1}$ and finding a Nash equilibrium [25] is an example of FPC. A general approach for local search is Iterative Improvement with Steepest-Descent as its most popular example. It follows a path along which the objective values are improving, in the feasible space. The end of the path is a local optimum. Likewise, algorithms for FPC, such as the Lemke-Howson algorithm [24] and the algorithm for Sperner's Lemma [35], also follow a path whose endpoint is an equilibrium or a fixed-point. But unlike in local search, a path in FPC does not have an obvious locally computable, monotonic ${ }^{2}$ measure-of-progress. Moreover, path following in FPC from an arbitrary point could lead to a cycle while the union of paths in Iterative Improvement is acyclic.

## Do these structural differences have any algorithmic implication?

There has been increasing evidence that local search and FPC are different. First, Aldous [2] showed that randomization can speedup local search (details below). His method crucially utilizes the monotonicity discussed above.

[^1]It remains open whether randomization helps FPC. Second, polynomial-time algorithms have been developed for several non-trivial classes of local search problems. These algorithms include the interior-point algorithm for linear and convex programming [22,26] and edge-insertion algorithms for geometric optimization [16]. However, popular fixed-point problems such as the computation of Nash or market equilibria [3] might be hard for polynomial time $[14,8,12]$. Other than those that can be solved by convex programming, we haven't yet discovered a significantly non-trivial class of equilibrium problems solvable in polynomial time. Third, an approximate local optimum for each PLS (Polynomial Local Search [21]) problem can be found in fully-polynomial time [27]. In contrast, one can apply a fully-polynomial-time approximate Nash equilibrium algorithm to find an exact Nash equilibrium in polynomial time [9]. Same can be said about approximate market equilibria [19]. Fourth, although they all have exponential worst-case complexity [31,23], the smoothed complexity of the Simplex Algorithm and Lemke-Howson Algorithm (or Scarf's market equilibrium algorithm [34]) might be drastically different [36, 9, 19]. This evidence inspires us to ask:

## Is fixed-point computation fundamentally harder than local search?

Naturally, the best way to address this question is to prove "If PPAD is in FP, then PLS is in FP." This question remains open and there are oracles separating PPAD and PLS from each other [4]. Here, we show that in the blackbox query model, FPC is strictly harder than local search.

## Black-Box Query Model for FPC and Local Search

For local search, we will use the model studied by Aldous [2] and Aaronson [1] (see also [30, 39, 37]). The search space is defined over $\mathbb{Z}_{n}^{d}=[1: n]^{d}$ and we are given a black-box function $h: \mathbb{Z}_{n}^{d} \rightarrow \mathbb{R}$. The local search problem is to find a local optimum of $h$, a vector $\mathbf{x} \in \mathbb{Z}_{n}^{d}$ such that $h(\mathbf{x}) \geq h(\mathbf{y}), \forall \mathbf{y}$ with $\|\mathbf{x}-\mathbf{y}\|_{1} \leq 1$. For FPC, we consider the model introduced by Hirsch, Papadimitriou and Vavasis [18]. In the HPV model, the search domain is also $\mathbb{Z}_{n}^{d}=$ $[1: n]^{d}$. We are given a black-box function $F: \mathbb{Z}_{n}^{d} \rightarrow \mathbb{Z}_{n}^{d}$ that satisfies Brouwer's condition [5] - a set of continuity and boundary conditions (see Section 2) - that guarantees the existence of a fixed-point. The FPC problem is to find a fixed-point of $F$, a vector $\mathbf{v} \in \mathbb{Z}_{n}^{d}$ satisfying $F(\mathbf{v})=\mathbf{v}$. The complexity of both problems is measured by the number of queries, of the form "What is $h(\mathbf{x})$ ?" or "What is $F(\mathbf{x})$ ?", needed to find a solution.

The HPV model is in fact a very good approximation for the computation of Brouwer's fixed points, as argued in [18]. To be self-contained, we include a brief discussion here. Note that Brouwer's Fixed Point Theorem [5] - that
any continuous map $f$ from a convex compact body, such as a simplex or a hypercube, to itself has a fixed point - is inherently continuous. In order to study complexity of this continuous problem in the Turing model, some inaccuracy must be introduced to ensure the existence of a solution with finite description [32, 33, 28, 18, 15].

Following Scarf [32], an approximate fixed point of a continuous map $f$ is a point $\mathbf{x}$ in the convex body such that $\|f(\mathbf{x})-\mathbf{x}\| \leq \epsilon$ for a given $\epsilon>0$. In 1928, Sperner [35] discovered a discrete fixed point theorem that led to the most elegant proof of Brouwer's Theorem. Suppose that $\Omega$ is a $d$ dimensional simplex with vertices $v_{1}, v_{2}, \ldots, v_{d+1}$, and that $\mathcal{S}$ is a simplicial decomposition of $\Omega$. Suppose $\Pi$ assigns to each vertex of $\mathcal{S}$ a color from $\{1,2, \ldots, d+1\}$ such that, for every vertex $v$ of $\mathcal{S}, \Pi(v) \neq i$ if the $i^{\text {th }}$ component of the barycentric coordinate of $v$ (the convex combination of $v_{1}, v_{2}, \ldots, v_{d+1}$ to express $v$ ), is 0 . Sperner's Lemma asserts that there exists a cell in $\mathcal{S}$ that contains all colors. Consider a Brouwer map $f$ with Lipschitz constant $L$ over the simplex $\Omega$. Suppose further that the diameter of each simplex cell in $\mathcal{S}$ is at most $\epsilon / L$. Then, one can define a color assignment $\Pi_{f}$ such that each fully-colored simplex in $\left(\mathcal{S}, \Pi_{f}\right)$ must have a vertex $\mathbf{v}$ satisfying $\|f(\mathbf{v})-\mathbf{v}\| \leq \Theta(\epsilon)$. Thus, a fully-colored cell of $\left(\mathcal{S}, \Pi_{f}\right)$ can be viewed as an approximate, discrete fixed point of $f$. The HPV model is an extension of Sperner's Lemma from the simplex to the hypercube. Our randomized lower bound can be extended to the black-box query model for Sperner's Lemma and hence for the computation of approximate fixed points.

There are some similarities between FPC and local search over $\mathbb{Z}_{n}^{d}$. For both, divide-and-conquer has positive but limited success: It can solve both problems using $O\left(n^{d-1}\right)$ queries [6]. An alternative approach to solve them is path-following. When following a short path, it can be faster than divide-and-conquer, but long and winding paths are the cause of inefficiency. The most prominent difference between paths in these two problems is that the values of $h$ along a path to a local optimum are monotonic, serving as a measure-of-progress along the path. Aldous [2] used this fact in a randomized algorithm: Randomly query $d^{1 / 2} n^{d / 2}$ points in $\mathbb{Z}_{n}^{d}$; let s be the sample point with the largest $h$ value; follow a path starting at $\mathbf{s}$. If a path to a local optimum is long, say much longer than $d^{1 / 2} n^{d / 2}$, then with high probability, the random samples intersect the path and partition it into sub-paths, each with expected length $O\left(d^{1 / 2} n^{d / 2}\right)$. As s has the largest $h$ value, its sub-path is the last sub-path of a potentially long path, but with expected length $O\left(d^{1 / 2} n^{d / 2}\right)$. So with randomization, Aldous reduced the expected query complexity to $O\left(d^{1 / 2} n^{d / 2}\right)$. It remains open whether randomization can reduce the query complexity of FPC over $\mathbb{Z}_{n}^{d}$. The lack of a measure-ofprogress along a path makes it impossible for us to directly use Aldous' idea.

## Our Main Result

As the main technical result of this paper, we prove that an expected number of $(\Omega(n))^{d-1}$ queries are indeed needed for FPC. Our lower bound is asymptotically tight ${ }^{3}$ as a function of $n$, since the divide-and-conquer algorithm in [6] can find a fixed point by querying $O\left(n^{d-1}\right)$ vectors.

In contrast to Aldous' result [2], our result demonstrates that randomization does not help much in FPC in the query model. It shows that, in the randomized query model over $\mathbb{Z}_{n}^{d}$, a fixed-point is strictly harder to find than a local optimum! The significant gap between these two problems is revealed only in randomized computation. In the deterministic framework, both have query complexity $\Theta\left(n^{d-1}\right)$.

One can show that the randomized query complexity for finding a global optimum over $\mathbb{Z}_{n}^{d}$ is $\Theta\left(n^{d}\right)$. Thus, the randomized query model over $\mathbb{Z}_{n}^{d}$ strictly separates these three important search problems:

## Global optimization is harder than fixed-point computation, and fixed-point computation is harder than local search.

Although the beautiful question "does PPAD in FP implies PLS is in FP?" remains open, our results uncover some fundamental difficulties in fixed-point and equilibrium computation. We also anticipate that a similar gap can be obtained in the quantum query model.

Our proof has two stages: We first define a string problem (Section 2.2), and reduce it to the FPC problem over $\mathbb{Z}_{n}^{d}$. We then prove a randomized lower bound for the string problem, which implies a same bound for FPC. The lower bound proof is obtained via a hierarchical construction of random long strings. Our reduction exposes the hardness of FPC, and allows us focus on the combinatorially simpler string problem. The idea behind the reduction is illustrated by Figure 1, in which we generate a simple graph $G^{\prime}$ over [3:27] ${ }^{2}$ from a string " 1537 " of integers. Our string-based method systematically generates random, long and winding paths in the grid graph over $\mathbb{Z}_{n}^{d}$. To achieve our nearlytight lower bound, these paths must be much longer than the random paths constructed in [39, 37] for local search. Our paths have expected length $(\Theta(n))^{d-1}$, while those paths for local search have length $\Theta\left(n^{d / 2}\right)$. We also develop new techniques for unknotting a self-intersecting path in $\mathbb{Z}_{n}^{d}$ and for realizing a path with a Brouwer function. These techniques are instrumental to our analysis and could be useful in the future complexity studies of FPC and its applications.

## Related Work

Our work is partially inspired by the results of Aaronson [1], Santha and Szegedy [30], Zhang [39], and Sun and Yao

[^2][37] on the randomized and quantum query complexity of local search over $\mathbb{Z}_{n}^{d}$. It is built on the model of Hirsch, Papadimitriou and Vavasis [18], who proved a tight $\Theta(n)$ deterministic bound for $\mathbb{Z}_{n}^{2}$ and an $\Omega\left(n^{d-2}\right)$ lower bound for $\mathbb{Z}_{n}^{d}$. Chen and Deng [6] improved this bound to $\Theta\left(n^{d-1}\right)$ for $\mathbb{Z}_{n}^{d}$. Friedl, Ivanyos, Santha, and Verhoeven [17] gave an $\Omega\left(n^{1 / 4}\right)$-lower bound on the randomized query complexity of the 2-dimensional Sperner problem. Our method for unknotting self-intersecting paths can be viewed as an extension of the 2D technique of [7] to high dimensions.

## 2 Three High-Dimensional Search Problems

We define three search problems. The first one concerns FPC. We introduce the last two to help the study of the first one. For each of the three problems, we define its mathematical structure, a query model for accessing this structure, the search problem itself, and its query complexity. Below, let $\mathbb{E}^{d}=\left\{ \pm \mathbf{e}_{1}, \pm \mathbf{e}_{2}, \ldots, \pm \mathbf{e}_{d}\right\}$. Let $\|\cdot\|$ denote $\|\cdot\|_{\infty}$. For two vectors $\mathbf{u} \neq \mathbf{v}$ in $\mathbb{Z}^{d}$, we say $\mathbf{u}<\mathbf{v}$ lexicographically if $u_{i}<v_{i}$ for some $i$, and $u_{j}=v_{j}$ for all $1 \leq j<i$.

### 2.1 Discrete Brouwer Fixed-Points

A function $f: \mathbb{Z}_{n}^{d} \rightarrow\{\mathbf{0}\} \cup \mathbb{E}^{d}$ is bounded if $f(\mathbf{x})+\mathbf{x} \in$ $\mathbb{Z}_{n}^{d}$ for all $\mathbf{x} \in \mathbb{Z}_{n}^{d} ; \mathbf{v} \in \mathbb{Z}_{n}^{d}$ is a zero point of $f$ if $f(\mathbf{v})=\mathbf{0}$. Clearly, if $F(\mathbf{x})=\mathbf{x}+f(\mathbf{x})$, then $\mathbf{v}$ is a fixed point of $F$ iff $\mathbf{v}$ is a zero point of $f$. A function $f: S \rightarrow\{\mathbf{0}\} \cup \mathbb{E}^{d}$, where $S \subset \mathbb{Z}^{d}$, is direction-preserving if $\left\|f\left(\mathbf{r}_{1}\right)-f\left(\mathbf{r}_{2}\right)\right\| \leq 1$ for all pairs $\mathbf{r}_{1}, \mathbf{r}_{2} \in S$ such that $\left\|\mathbf{r}_{1}-\mathbf{r}_{2}\right\| \leq 1$.

Following the discrete fixed-point theorem of [20], we have: For every function $f: \mathbb{Z}_{n}^{d} \rightarrow\{\mathbf{0}\} \cup \mathbb{E}^{d}$, if $f$ is both bounded and direction-preserving, then there exists $\mathbf{v} \in \mathbb{Z}_{n}^{d}$ such that $f(\mathbf{v})=\mathbf{0}$. We refer to a bounded and directionpreserving function $f$ over $\mathbb{Z}_{n}^{d}$ as a Brouwer function over $\mathbb{Z}_{n}^{d}$. In the query model, one can only access $f$ by asking queries of the form: "What is $f(\mathbf{r})$ ?" for a point $\mathbf{r} \in \mathbb{Z}_{n}^{d}$.

The FPC problem $\mathbf{Z P}^{d}$ that we will study is: Given $a$ Brouwer function $f: \mathbb{Z}_{n}^{d} \rightarrow\{\mathbf{0}\} \cup \mathbb{E}^{d}$ in the query model, find a zero point of $f$. Given any randomized algorithm $\Upsilon^{4}$ for solving $\mathrm{ZP}^{d}$, we let $\mathrm{RQ}(\Upsilon, f)$ denote the expected number of queries used by $\Upsilon$ to output a zero point. We let

$$
\mathrm{RQ}_{\mathrm{ZP}}^{d}(n)=\min _{\Upsilon} \max _{f: \text { Brouwer function over } \mathbb{Z}_{n}^{d}} \mathrm{RQ}(\Upsilon, f)
$$

denote the randomized query complexity of $\mathrm{ZP}^{d}$. In this paper, we will prove:
Theorem 2.1 (Main). For all $d \geq 2$, there exists a constant $c$ such that for all sufficiently large $n$,

$$
\mathrm{RQ}_{\mathrm{ZP}}^{d}(n) \geq\left(c^{-d} n\right)^{d-1}
$$

[^3]In contrast, the deterministic query complexity for solving $\mathrm{ZP}^{d}$ is at most $7 n^{d-1}$ [6].

The FPC problem defined here is computationally equivalent to the fixed problems studied in [18, 14, 9, 35]. Thus, our result carries over to these FPC problems.

### 2.2 End-of-a-String

Suppose $\Sigma$ is a finite set. A string $S$ over $\Sigma$ of length $m$ is a sequence $S=a_{1} a_{2} \ldots a_{m}$ with $a_{i} \in \Sigma$.

Definition 2.2 (Non-Repeating-Strings). A string $S=a_{1}$ $a_{2} \ldots a_{m}$ over $\mathbb{Z}_{n}=[1: n]$ is $d$-non-repeating for $d \in[1:$ $m$ ] if (1) each string over $\mathbb{Z}_{n}$ of length d appears in $S$ as a substring at most once; (2) $a_{i}$ is odd if $i$ is a multiple of $d$ and $a_{i}$ is even otherwise; and (3) $m$ is a multiple of $d$.

We define end $_{d}(S)=a_{m-d+1} \ldots a_{m-1} a_{m}$.
Each $d$-non-repeating string $S=a_{1} \ldots a_{m}$ over $\mathbb{Z}_{n}$ defines a query oracle $\mathbb{B}_{S}:[1: n]^{d} \rightarrow\left(\{"\right.$ no" $\left.\} \cup \mathbb{Z}_{n}\right) \times$ $\left(\{" \mathrm{no"}\} \cup \mathbb{Z}_{n}\right)$ : For $S^{\prime}=b_{1} \ldots b_{d} \in[1: n]^{d}$, if $S^{\prime}$ is not a substring of $S$, then $\mathbb{B}_{S}\left(S^{\prime}\right)=$ ("no", "no"); otherwise, there is a unique $k$ such that $a_{k+i-1}=b_{i}, \forall i \in[1: d]$. Then $\mathbb{B}_{S}\left(S^{\prime}\right)=\left("\right.$ no", $\left.a_{d+1}\right)$ when $k=1, \mathbb{B}_{S}\left(S^{\prime}\right)=$ ( $a_{m-d}$, "no") when $k=m-d+1$ (i.e., $S^{\prime}=\operatorname{end}_{d}(S)$ ), and $\mathbb{B}_{S}\left(S^{\prime}\right)=\left(a_{k-1}, a_{k+d}\right)$, otherwise.

Let $\mathrm{ES}^{d}$ denote the search problem: Given a pair $(S, n)$ where $S$ is a d-non-repeating string over $\mathbb{Z}_{n}$ accessible by $\mathbb{B}_{S}{ }^{5}$, and its first d symbols $a_{1} \ldots a_{d-1} a_{d}$ with $a_{d}=1$, find end $(S)$. We let $\mathrm{RQ}_{\mathrm{ES}}^{d}(n)$ denote its randomized query complexity ${ }^{6}$. It is easy to show that $\operatorname{RQ}_{\mathrm{ES}}^{1}(n)=\Theta(n)$. In Section 4, we will prove
Theorem 2.3 (Complexity of ES ${ }^{d}$ ). For all $d \geq 2$ and sufficiently large $n, \mathrm{RQ}_{\mathrm{ES}}^{d}(4 n+4) \geq 2^{-(d+1)}\left(24^{-d} n\right)^{d}$.

In Section 3, we will reduce $\mathrm{ES}^{d-1}$ to $\mathrm{ZP}^{d}$ and prove Theorem 2.1 as a corollary of Theorem 2.3. In the reduction, the two problems are connected via the following problem on graphs.

### 2.3 End-of-a-Path in Grid-PPAD Graphs

The mathematical structure for this search problem is a directed graph $G=(V, E)$. A vertex $v \in V$ satisfies $E u$ ler's condition if $\Delta_{I}(v)=\Delta_{O}(v)$ where $\Delta_{I}(v)$ and $\Delta_{O}(v)$ are the in-degree and out-degree of $v$, respectively. We start with the following definition motivated by Papadimitriou's PPAD class [29].

Definition 2.4 (Generalized PPAD Graphs). A directed graph $G=(V, E)$ is a generalized PPAD graph if (1) there

[^4]exist exactly one vertex $v_{S} \in V$ with $\Delta_{O}\left(v_{S}\right)=\Delta_{I}\left(v_{S}\right)+$ 1 , and exactly one $v_{T} \in V$ with $\Delta_{I}\left(v_{T}\right)=\Delta_{O}\left(v_{T}\right)+1$; (2) all vertices in $V-\left\{v_{S}, v_{T}\right\}$ satisfy Euler's condition; and (3) if $\left(v_{1}, v_{2}\right)$ is a directed edge in $E$, then $\left(v_{2}, v_{1}\right) \notin E$.

We refer to $v_{S}$ and $v_{T}$ as the starting and ending vertices of $G$, respectively. $G$ is a PPAD graph if in addition $\Delta_{I}(v), \Delta_{O}(v) \leq 1$, for all $v \in V$.

Edges of a PPAD graph form a collection of disjoint directed cycles and a directed path from $v_{S}$ to $v_{T}$. In this paper, we are interested in a special family of PPAD graphs over $\mathbb{Z}_{n}^{d}$ : A directed graph $G=\left(\mathbb{Z}_{n}^{d}, E\right)$ is a generalized grid-PPAD graph over $\mathbb{Z}_{n}^{d}$ if it is a generalized PPAD graph and the underlying undirected graph of $G$ is a subgraph of the grid graph defined over $\mathbb{Z}_{n}^{d}$. Moreover, if $G$ is also a PPAD graph, then we say $G$ is a grid-PPAD graph. $G^{*}$ and $G^{\prime}$ in Figure 1 are examples of generalized grid-PPAD graph and grid-PPAD graph, respectively. We now define the query model $\mathbb{B}_{G}$ for accessing a grid-PPAD graph.

Definition 2.5. Given a grid-PPAD graph $G=\left(\mathbb{Z}_{n}^{d}, E\right)$, we define $\mathbb{B}_{G}: \mathbb{Z}_{n}^{d} \rightarrow\left(\{" n o "\} \cup \mathbb{E}^{d}\right) \times\left(\{" n o "\} \cup \mathbb{E}^{d}\right)$ as follows: for $\mathbf{v} \in \mathbb{Z}_{n}^{d}(1) \mathbb{B}_{G}(\mathbf{v})=\left(" n o ", \mathbf{v}_{1}-\mathbf{v}\right)$ if $\mathbf{v}$ is the starting vertex of $G$ and $\left(\mathbf{v}, \mathbf{v}_{1}\right) \in E$; (2) $\mathbb{B}_{G}(\mathbf{v})=(\mathbf{v}-$ $\mathbf{v}_{1}$, "no") if $\mathbf{v}$ is the ending vertex of $G$ and $\left(\mathbf{v}_{1}, \mathbf{v}\right) \in E$; (3) $\mathbb{B}_{G}(\mathbf{v})=\left(\mathbf{v}-\mathbf{v}_{1}, \mathbf{v}_{2}-\mathbf{v}\right)$ if $\left(\mathbf{v}_{1}, \mathbf{v}\right)$ and $\left.\left(\mathbf{v}, \mathbf{v}_{2}\right) \in E\right)$; and $(4) \mathbb{B}_{G}(\mathbf{v})=(" n o ", " n o ")$, otherwise.

In other words, for each $\mathbf{v} \in \mathbb{Z}_{n}^{d}, \mathbb{B}_{G}(\mathbf{v})$ gives the predecessor and successor of $\mathbf{v}$ in $G$. We will use the property that if $\mathbb{B}_{G}(\mathbf{v})=\left(\mathbf{s}_{1}, \mathbf{s}_{2}\right)$ and $\mathbf{s}_{1}, \mathbf{s}_{2} \in \mathbb{E}^{d}$, then $\mathbf{s}_{1}+\mathbf{s}_{2} \neq 0$.

Let $\mathrm{GP}^{d}$ be the search problem: Given a triple $(G, n, \mathbf{u})$ where $G$ is a grid-PPAD graph over grid $\mathbb{Z}_{n}^{d}$ accessible by $\mathbb{B}_{G}$ and $\mathbf{u}$ is the starting vertex of $G$ satisfying $u_{d}=1$, find its ending vertex. We use $\mathrm{RQ}_{\mathrm{GP}}^{d}(n)$ to denote the randomized query complexity for solving this problem.

## 3 Reduction Among Search Problems

In this section, we reduce $E S^{d-1}$ to $\mathrm{ZP}^{d}$ by first reducing $\mathrm{ES}^{d-1}$ to $\mathrm{GP}^{d}$ and then reducing $\mathrm{GP}^{d}$ to $\mathrm{ZP}^{d}$ (Theorem 3.1 and 3.2 below). Theorem 2.1 then follows from Theorem 2.3. In this section, we will prove Theorem 3.1. A proof of Theorem 3.2 can be found in the full version [10].

Theorem 3.1 (From ES ${ }^{d-1}$ to $\mathrm{GP}^{d}$ ). For all $d \geq 2$,

$$
\mathrm{RQ}_{\mathrm{ES}}^{d-1}(n) \leq 4 d \cdot \mathrm{RQ}_{\mathrm{GP}}^{d}(8 n+1)
$$

Theorem 3.2 (From $\mathrm{GP}^{d}$ to $\mathrm{ZP}^{d}$ ). For all $d \geq 2$,

$$
\mathrm{RQ}_{\mathrm{GP}}^{d}(n) \leq \mathrm{RQ}_{\mathrm{ZP}}^{d}(24 n+7)
$$

Proof of Theorem 3.1. We define a map $\mathcal{F}_{d}$ from $\mathbb{Z}^{d-1}$ to $\mathbb{Z}^{d}:$ for $d=2, \mathcal{F}_{2}(a)=(a, a) ;$ and for $d>2, \mathcal{F}_{d}(\mathbf{a})=$ $\left(a_{1}, a_{1}+a_{2}, \ldots, a_{d-2}+a_{d-1}, a_{d-1}\right)$. We will crucially use
the following nice property of $\mathcal{F}_{d}$ : For any $k \in[1: d-1]$ and for any $\mathbf{a} \in \mathbb{Z}^{d-1}$, we can uniquely determine the first $k$ and the last $k$ entries of a, respectively, from the first $k$ and the last $k$ entries of $\mathcal{F}_{d}(\mathbf{a})$.

Let $S$ be a $(d-1)$-non-repeating string over $\mathbb{Z}_{n}$ of length $m(d-1)$ for some $m \geq 2$, whose $(d-1)^{s t}$ symbol is 1 . We view $S$ as a sequence of $m$ points $\mathbf{a}_{1}, \mathbf{a}_{2}, \ldots, \mathbf{a}_{m}$ in $\mathbb{Z}_{n}^{d-1}$, where $\mathbf{a}_{i}=a_{i, 1} \ldots a_{i, d-1}$, such that,

$$
S=a_{1,1} a_{1,2} \ldots a_{1, d-1} \ldots a_{m, 1} a_{m, 2} \ldots a_{m, d-1}
$$

From $S$, we will construct a grid-PPAD graph $G^{\prime}$ in two stages. In the first stage, we construct a generalized gridPPAD graph $G^{*}$ over $\mathbb{Z}_{2 n}^{d}$ such that
(A.1) Its starting vertex is $\mathbf{u}^{*}=\mathcal{F}_{d}\left(\mathbf{a}_{1}\right)$ and its ending vertex is $\mathbf{w}^{*}=\mathcal{F}_{d}\left(\mathbf{a}_{m}\right)$; (A.2) For each directed edge $(\mathbf{u}, \mathbf{v})$ with $\mathbf{u}-\mathbf{v} \in \mathbb{E}^{d}$, at most one query to $\mathbb{B}_{S}$ is needed to determine whether $(\mathbf{u}, \mathbf{v}) \in G^{*}$ or not.

Suppose $\mathbf{u}$ and $\mathbf{v} \in \mathbb{Z}_{2 n}^{d}$ are two vertices that differ in only one coordinate, say the $i^{\text {th }}$ coordinate, and $\mathbf{e}=(\mathbf{v}-$ $\mathbf{u}) /\left|v_{i}-u_{i}\right| \in \mathbb{E}^{d}$. Let $E(\mathbf{u}, \mathbf{v})=\{(\mathbf{u}, \mathbf{u}+\mathbf{e}),(\mathbf{u}+\mathbf{e}, \mathbf{u}+$ $2 \mathbf{e}) \ldots(\mathbf{v}-\mathbf{e}, \mathbf{v})\}$. For $n, m_{1}, m_{2} \in \mathbb{Z}$ and $s \in\{ \pm 1\},(n, s)$ is consistent with $\left(m_{1}, m_{2}\right)$ if either (1) $m_{1} \leq n<m_{2}$ and $s=1$; or (2) $m_{2}<n \leq m_{1}$ and $s=-1$. Recall a directed path is simple if it contains each vertex at most once.

Now consider two consecutive points $\mathbf{a}=\mathbf{a}_{t}$ and $\mathbf{b}=$ $\mathbf{a}_{t+1}$ in the $(d-1)$-non-repeating string $S=\mathbf{a}_{1} \ldots \mathbf{a}_{m}$. We know $\mathbf{a} \neq \mathbf{b}$. We map them to $\mathbf{u}=\mathcal{F}_{d}(\mathbf{a})$ and $\mathbf{w}=\mathcal{F}_{d}(\mathbf{b})$ in $\mathbb{Z}_{2 n}^{d}$, and connect them with a path through a sequence of $(d-1)$ vertices $\mathbf{v}_{0}=\mathbf{u}, \mathbf{v}_{1}, \ldots, \mathbf{v}_{d-1}, \mathbf{v}_{d}=\mathbf{w}$, where $v_{i, j}=u_{j}$ if $i<j$ and $v_{i, j}=w_{j}$ if $i \geq j$.

Note that $\mathbf{v}_{i-1}$ and $\mathbf{v}_{i}$ differ only in the $i^{\text {th }}$ coordinate. Let $P(\mathbf{a}, \mathbf{b})=\cup_{i=0}^{d-1} E\left(\mathbf{v}_{i}, \mathbf{v}_{i+1}\right)$, then $P(\mathbf{a}, \mathbf{b})$ is a simple directed path from $\mathbf{u}=\mathbf{v}_{0}$ to $\mathbf{w}=\mathbf{v}_{d}$.

Construction of $G^{*}$ from $S: G^{*}=\left(\mathbb{Z}_{2 n}^{d}, \cup_{i=1}^{m-1} P\left(\mathbf{a}_{i}\right.\right.$, $\left.\mathbf{a}_{i+1}\right)$ ). See Figure 1 for an example. By Lemma 3.4 below, $G^{*}$ is a generalized grid-PPAD graph. Property (A.2) can be derived from Proposition 3.3.

Proposition 3.3 (Local Characterization of $P(\mathbf{a}, \mathbf{b})$ ). Let $\mathbf{a}, \mathbf{b}$ be two points in $\mathbb{Z}_{n}^{d-1}$. For $\mathbf{v} \in \mathbb{Z}_{2 n}^{d}$ and $s \in\{ \pm 1\}$ :
(1) $\left(\mathbf{v}, \mathbf{v}+s \mathbf{e}_{1}\right) \in P(\mathbf{a}, \mathbf{b})$ iff $\left(v_{1}, s\right)$ is consistent with $\left(a_{1}, b_{1}\right), a_{d-1}=v_{d}$, and $a_{d-i}=v_{d-i+1}-a_{d-i+1}$ for all $i \in[2: d-1]$;
(2) $\left(\mathbf{v}, \mathbf{v}+s \mathbf{e}_{d}\right) \in P(\mathbf{a}, \mathbf{b})$ iff $\left(v_{d}, s\right)$ is consistent with $\left(a_{d-1}, b_{d-1}\right), b_{1}=v_{1}$ and $b_{i}=v_{i}-b_{i-1}, \forall i \in[2: d-1]$;
(3) When $k \in[2: d-1],\left(\mathbf{v}, \mathbf{v}+s \mathbf{e}_{k}\right) \in P(\mathbf{a}, \mathbf{b})$ iff: (3.0) $\left(v_{k}, s\right)$ is consistent with $\left(a_{k-1}+a_{k}, b_{k-1}+b_{k}\right)$; (3.1) $a_{d-1}=v_{d}$ and $a_{d-i}=v_{d-i+1}-a_{d-i+1}, \forall i \in[2: d-k]$; and (3.2) $b_{1}=v_{1}$ and $b_{i}=v_{i}-b_{i-1}, \forall i \in[2: k-1]$.

Lemma 3.4 (Structural Correctness). For all $(d-1)$-nonrepeating string $S=\mathbf{a}_{1} \mathbf{a}_{2} \ldots \mathbf{a}_{m}$ over $\mathbb{Z}_{n}$, if edge $(\mathbf{u}, \mathbf{v}) \in$ $P\left(\mathbf{a}_{i}, \mathbf{a}_{i+1}\right)$ then $(\mathbf{u}, \mathbf{v}),(\mathbf{v}, \mathbf{u}) \notin P\left(\mathbf{a}_{j}, \mathbf{a}_{j+1}\right)$ for all $j \neq i$.


Figure 1. $G^{*}$ and $G^{\prime}$ constructed from 1537
Proof. We only prove the case when $\mathbf{e}=\mathbf{v}-\mathbf{u}=s \mathbf{e}_{k}$ with $1<k<d$ and $s \in\{ \pm 1\}$. The other two cases are similar.

From Proposition 3.3, $(\mathbf{u}, \mathbf{v}) \in P\left(\mathbf{a}_{i}, \mathbf{a}_{i+1}\right)$ implies that $\mathbf{a}_{i}$ and $\mathbf{a}_{i+1}$ satisfy conditions (3.1) and (3.2). If (u,v) or $(\mathbf{v}, \mathbf{u})$ is in $P\left(\mathbf{a}_{j}, \mathbf{a}_{j+1}\right)$, then $\mathbf{a}_{j}$ and $\mathbf{a}_{j+1}$ also satisfy these two conditions. Thus, $a_{i, k} \ldots a_{i, d-1} a_{i+1,1} \ldots a_{i+1, k-1}=a_{j, k}$ $\ldots a_{j, d-1} a_{j+1,1} \ldots a_{j+1, k-1}$, which contradicts with the assumption that $S$ is $(d-1)$-non-repeating.

In the second stage, we build a grid-PPAD graph $G^{\prime}$ over $\mathbb{Z}_{8 n+1}^{d}$ from $G^{*}$. Let map $\Gamma(\mathbf{v})=4 \mathbf{v}-\mathbf{1}$ for all $\mathbf{v} \in \mathbb{Z}_{2 n}^{d}$. Our $G^{\prime}$, see Figure 1 for an example, satisfies:
(B.1) Its starting vertex is $\mathbf{u}^{\prime}=\Gamma\left(\mathbf{u}^{*}\right)-2 \mathbf{e}_{d}$, and its ending vertex $\mathbf{w}^{\prime}$ satisfies $\left\|\mathbf{w}^{\prime}-\Gamma\left(\mathbf{w}^{*}\right)\right\| \leq 1$; (B.2) For each $\mathbf{v} \in \mathbb{Z}_{8 n+1}^{d}$, one can determine $\mathbb{B}_{G^{\prime}}(\mathbf{v})$ from the predecessors and successors of $\mathbf{u}$ in $G^{*}$, where $\mathbf{u} \in$ $\mathbb{Z}_{2 n}^{d}$ is the lexicographically smallest vertex such that $\|\mathbf{v}-\Gamma(\mathbf{u}) \leq 2\|$.
Every vertex in $G^{*}=\left(\mathbb{Z}_{2 n}^{d}, E^{*}\right)$, other than the starting and ending ones, satisfies Euler's condition. Some vertices may have in-degree and out-degree more than one. In this stage, we systematically unknot high-degree intersections of $G^{*}$ using a larger space $\mathbb{Z}_{8 n+1}^{d}$.

Two subsets $H_{1}$ and $H_{2}$ of $\mathbb{E}^{d}$, where $d \geq 2$, form a balanced-non-canceling pair if $\left|H_{1}\right|=\left|H_{2}\right|$ and $\mathbf{s}_{1}+\mathbf{s}_{2}$ $\neq \mathbf{0}$ for all $\mathbf{s}_{1} \in H_{1}$ and $\mathbf{s}_{2} \in H_{2}$. Let $H_{I}(\mathbf{u})=\left\{\mathbf{e} \in \mathbb{E}^{d} \mid\right.$ $\left.(\mathbf{u}-\mathbf{e}, \mathbf{u}) \in E^{*}\right\}$ be the vector differences of $\mathbf{u}$ and its predecessors in $G^{*}$. Similarly, let $H_{O}(\mathbf{u})=\left\{\mathbf{e} \in \mathbb{E}^{d} \mid(\mathbf{u}\right.$, $\left.\mathbf{u}+\mathbf{e}) \in E^{*}\right\}$ be the vector differences of the successors of $\mathbf{u}$ and $\mathbf{u}$. In the construction below, we will use the fact that if $\mathbf{u}$ satisfies Euler's condition then $\left(H_{I}, H_{O}\right)$ is a balanced-non-canceling pair.

Using the procedure of Figure 2, we define a directed graph $G\left[H_{1}, H_{2}\right]=\left(\{-1,0,+1\}^{d}, E\left[H_{1}, H_{2}\right]\right)$ for every balanced-non-canceling pair $H_{1}$ and $H_{2} . G\left[H_{1}, H_{2}\right]$ has the following properties: (1) For every vertex $\mathbf{u}$ of $G\left[H_{1}, H_{2}\right]$, $\Delta_{I}(\mathbf{u}), \Delta_{O}(\mathbf{u}) \leq 1$; (2) A vertex $\mathbf{u}$ has $\Delta_{I}(\mathbf{u})=0$ and $\Delta_{O}(\mathbf{u})=1$ iff there exists an $\mathbf{e} \in H_{1}$ such that $\mathbf{u}=-\mathbf{e}$; (3) A vertex $\mathbf{u}$ has $\Delta_{I}(\mathbf{u})=1$ and $\Delta_{O}(\mathbf{u})=0$ iff there exists an $\mathbf{e} \in H_{2}$ such that $\mathbf{u}=+\mathbf{e}$.

Construction of $G^{\prime}$ from $G^{*}$ : Let $\mathbf{u}^{*}$ be the starting vertex and $\mathbf{w}^{*}$ be the ending vertex of $G^{*}$. We build $G^{\prime}=$ $\left(\mathbb{Z}_{8 n+1}^{d}, E^{\prime}\right)$ by applying the procedure of Figure 2 locally
: set edge set $E\left[H_{1}, H_{2}\right]=\emptyset$
while $H_{1} \neq \emptyset$ do
let $\mathbf{s}_{1}$ be the smallest vector in $H_{1}$ and $\mathbf{s}_{2}$ be the
largest vector in $\mathrm{H}_{2}$ under the lexicographical ordering set $H_{1}=H_{1}-\left\{\mathbf{s}_{1}\right\}$ and $H_{2}=H_{2}-\left\{\mathbf{s}_{2}\right\}$; insert $\left\{\left(-\mathbf{s}_{1},-\mathbf{s}_{1}+\mathbf{s}_{2}\right),\left(-\mathbf{s}_{1}+\mathbf{s}_{2}, \mathbf{s}_{2}\right)\right\}$ into $E\left[H_{1}, H_{2}\right]$

Figure 2. Construction of Graph $G\left[H_{1}, H_{2}\right]$
to every vertex $\mathbf{u} \in \mathbb{Z}_{2 n}^{d}$ of $G^{*}$. When $\mathbf{u} \in\left\{\mathbf{u}^{*}, \mathbf{w}^{*}\right\}$, we use a slight modification of $\left(H_{I}(\mathbf{u}), H_{O}(\mathbf{u})\right)$. Initially, the edge set $E^{\prime}=\emptyset$. Recall $\Gamma(\mathbf{u})=4 \mathbf{u}-\mathbf{1}$.

## [ local embedding of the starting vertex $\mathbf{u}^{*}$ ]

As $u_{d}^{*}=1, \mathbf{e}_{d} \notin H_{I}\left(\mathbf{u}^{*}\right)$ and $-\mathbf{e}_{d} \notin H_{O}\left(\mathbf{u}^{*}\right)$. Let $H_{I}=H_{I}\left(\mathbf{u}^{*}\right) \cup\left\{\mathbf{e}_{d}\right\}$. Add directed edges $\left(\Gamma\left(\mathbf{u}^{*}\right)-\right.$ $\left.2 \mathbf{e}_{d}, \Gamma\left(\mathbf{u}^{*}\right)-\mathbf{e}_{d}\right)$ and $\left(\Gamma\left(\mathbf{u}^{*}\right)+\mathbf{s}_{1}, \Gamma\left(\mathbf{u}^{*}\right)+\mathbf{s}_{2}\right)$ to $E^{\prime}$ for all edges $\left(\mathbf{s}_{1}, \mathbf{s}_{2}\right)$ in $G\left[H_{I}, H_{O}\left(\mathbf{u}^{*}\right)\right]$.
[ local embedding of the ending vertex $\mathbf{w}^{*}$ ]
As $\left|H_{I}\left(\mathbf{w}^{*}\right)\right|=\left|H_{O}\left(\mathbf{w}^{*}\right)\right|+1, H_{I}\left(\mathbf{w}^{*}\right) \neq \emptyset$. Let $\mathbf{e}$ be the smallest vector in $H_{I}\left(\mathbf{w}^{*}\right)$, and $H_{I}=H_{I}\left(\mathbf{w}^{*}\right)-$ $\{\mathbf{e}\}$. Add edges $\left(\Gamma\left(\mathbf{w}^{*}\right)+\mathbf{s}_{1}, \Gamma\left(\mathbf{w}^{*}\right)+\mathbf{s}_{2}\right)$ to $E^{\prime}$ for all edges $\left(\mathbf{s}_{1}, \mathbf{s}_{2}\right)$ in $G\left[H_{I}, H_{O}\left(\mathbf{w}^{*}\right)\right]$.

## [ local embedding of other vertices $u$ ]

For each $\mathbf{u} \in G^{*}$, add $\left(\Gamma(\mathbf{u})+\mathbf{s}_{1}, \Gamma(\mathbf{u})+\mathbf{s}_{2}\right)$ to $E^{\prime}$ for all edges $\left(\mathbf{s}_{1}, \mathbf{s}_{2}\right)$ in $G\left[H_{I}(\mathbf{u}), H_{O}(\mathbf{u})\right]$.
[ connecting local embeddings]
For each $(\mathbf{u}, \mathbf{v}) \in G^{*}$, letting $\mathbf{e}=\mathbf{v}-\mathbf{u} \in \mathbb{E}^{d}$, we add $(\Gamma(\mathbf{u})+\mathbf{e}, \Gamma(\mathbf{u})+2 \mathbf{e}),(\Gamma(\mathbf{u})+2 \mathbf{e}, \Gamma(\mathbf{u})+3 \mathbf{e})$ to $E^{\prime}$.

It is quite mechanical to check that $G^{\prime}$ is a grid-PPAD graph that satisfies (B.1) and (B.2). Theorem 3.1 follows directly from Property (A.1), (A.2), (B.1) and (B.2).

## 4 Randomized Lower Bound for ES ${ }^{d}$

The technical objective of this section is to construct a distribution $\mathcal{S}$ of $d$-non-repeating strings and show that, for a random string $S$ drawn according to $\mathcal{S}$, every deterministic algorithm for $\mathrm{ES}{ }^{d}$ needs expected $(\Omega(n))^{d}$ queries to $\mathbb{B}_{S}$. Thus, by Yao's Minimax Principle [38], we have $\mathrm{RQ}_{\mathrm{ES}}^{d}(n)=(\Omega(n))^{d}$. We will apply random permutations hierarchically to define the distribution $\mathcal{S}$, and ensure that a random string from $\mathcal{S}$ has sufficient entropy that its search problem is expected to be difficult.

### 4.1 Hierarchical Construction of Random d-Non-Repeating Strings

Let $\mathbb{J}_{n}=[2: 2 n+2], \mathbb{O}_{n}=\{3,5,7, \ldots, 2 n+1\}$, and $\mathbb{F}_{n}=\{4,6, \ldots, 2 n+2\}$. Let strings $S_{0}=2, S_{1}=3 \circ 4 \ldots$ $S_{n}=(2 n+1) \circ(2 n+2)$. Each permutation $\pi:[1: n] \rightarrow$


Figure 3. A valid (2,2)-ToC $\mathcal{T}$
$[1: n]$ defines a string $C=S_{0} \circ S_{\pi(1)} \circ \ldots \circ S_{\pi(n)}$ which we refer to as a connector over $\mathbb{J}_{n}$. Let $r[C]=2 \pi(n)+2$, the last symbol of $C$. We use $\phi_{C}(2)$ to denote the right neighbor of 2 . Each $s \in \mathbb{J}_{n}-\left\{2, r\left[C_{\pi}\right]\right\}$ has two neighbors in $C$. The left neighbor of an even $s$ is $s-1$, and we use $\phi_{C}(s)$ to denote its right neighbor; the right neighbor of an odd $s$ is $s+1$ and we use $\phi_{C}(s)$ to denote its left neighbor. Clearly, if $\phi_{C}(s)=t$ then $\phi_{C}(t)=s$.

Our hierarchical framework is built on $T_{n, d}$, the rooted complete- $(2 n+1)$-nary tree of height $d$. In $T_{n, d}$, each internal node $u$ is connected to its $(2 n+1)$ children by edges with distinct labels from $\mathbb{J}_{n}$; if $u$ is connected to $v$ by an edge labeled with $j$, then we call $v$ the $j^{\text {th }}$-successor of $u$. Each node $v$ of $T_{n, d}$ has a natural name, name $(v)$, the concatenation of labels along the path from the root of $T_{n, d}$ to $v$. Let height $(v)$ and level $(v)$ denote the height and level of node $v$ in the tree. For example, the height of the root is $d$ and the level of the root is 0 .

Definition 4.1 (Tree-of-Connectors). $A n(n, d)-\mathrm{ToC} \mathcal{T}$ is a tree $T_{n, d}$ in which each internal node $v$ is associated with a connector $C_{v}$ over $\mathbb{J}_{n}$. The $\left(r\left[C_{v}\right]\right)^{\text {th }}$-successor of $v$ is referred to as the last successor of $v$. The tail of $v$, tail $(v)$, is the leaf reachable from $v$ by last-successor relations. The tail of a leaf is itself. The tail of $\mathcal{T}$, tail $(\mathcal{T})$, is the tail of its root. The head of a leaf $u$, head $(u)$, is the ancestor of $u$ with the largest height such that $u$ is its tail; if no such ancestor exists, then head $(u)=u$.
Definition 4.2 (Valid ToC). $A n(n, d)-\mathrm{ToC} \mathcal{T}$ is valid if for every internal $v$ and for each pair of $s, t \in \mathbb{J}_{n}$ with $\phi_{C_{v}}(s)=t$, name $\left(u_{s}\right)$ and name $\left(u_{t}\right)$ share a common suffix of length height $(v)-1$, where $u_{s}$ and $u_{t}$ are the tails of the $s^{t h}$-successor and $t^{\text {th }}$-successor of $v$, respectively.
Definition $4.3\left(\mathbb{B}_{\mathcal{T}}\right.$ for accessing $\left.\mathcal{T}\right)$. Suppose $\mathcal{T}$ is a valid $(n, d)$-ToC. The input to $\mathbb{B}_{\mathcal{T}}$ is a point $\mathbf{q} \in\left(\mathbb{J}_{n}\right)^{d}$ (defining the name of a leaf $u$ in $\mathcal{T})$. Let $h=$ height $($ head $(u))$. If $u$ is the tail of $\mathcal{T}$, i.e., $h=d$, then $\mathbb{B}_{\mathcal{T}}(\mathbf{q})=\mathcal{T}$. Otherwise, let $v_{1}=$ head $(u)$ and let $v$ be the parent of $v_{1}$. Note that $v_{1}$ is the $\left(q_{d-h}\right)^{\text {th }}$-successor of $v$. Let $\mathcal{T}_{1}$ be the tree rooted at $v_{1}$. As $u \neq$ tail $(v), \phi_{C_{v}}\left(q_{d-h}\right)$ is defined and let $\mathcal{T}_{2}$ be the subtree rooted at the $\left(\phi_{C_{v}}\left(q_{d-h}\right)\right)^{\text {th }}$-successor of $v$. Then, $\mathbb{B}_{\mathcal{T}}(\mathbf{q})=\left(h, \phi_{C_{v}}\left(q_{d-h}\right), \mathcal{T}_{1}, \mathcal{T}_{2}\right)$.

We define our final search problem Name-the-Tail, $\mathrm{NT}^{d}$ as: Given a valid $(n, d)-\operatorname{To} C \mathcal{T}^{*}$ accessible by $\mathbb{B}_{\mathcal{T}^{*}}$, find the name of its tail. Theorem 2.3 follows from the next two theorems. We will prove Theorem 4.4 in Section 4.3.

Theorem 4.4 (Complexity of $\mathrm{NT}^{d}$ ). For all $d \geq 1$ and sufficiently large $n, \mathrm{RQ}_{\mathrm{NT}}^{d}(n) \geq 2^{-(d+1)}\left(24^{-d} n\right)^{d}$.
Theorem 4.5 (From $\mathrm{NT}^{d}$ to $\mathrm{ES}^{d}$ ). For all $d \geq 1$, we have $\mathrm{RQ}_{\mathrm{NT}}^{d}(n) \leq \mathrm{RQ}_{\mathrm{ES}}^{d}(4 n+4)$.
Proof of Theorem 4.5. For two strings $S_{1}=a_{1} \ldots a_{k}$ and $S_{2}=b_{1} \ldots b_{t}$, let $S_{1} \circ S_{2}=a_{1} \ldots a_{k} b_{1} \ldots b_{t}$. For $d \geq 1$, if $a_{k-d+i}=b_{i}$ for all $i \in[1: d]$, then let $S_{1} \circ_{d} S_{2}=$ $a_{1} \ldots a_{k} b_{d+1} \ldots b_{t}$. Given a string $S$ over $\mathbb{Z}$ of length $k \cdot d$, we write $S$ as $\mathbf{u}_{1} \mathbf{u}_{2} \ldots \mathbf{u}_{k}$ with $\mathbf{u}_{i} \in \mathbb{Z}^{d}$. For $t \in \mathbb{Z}$, let insert $_{d}(S, t)=\mathbf{u}_{1} \circ t \circ \mathbf{u}_{2} \circ t \ldots \mathbf{u}_{k-1} \circ t \circ \mathbf{u}_{k}$.

We need to build a $d$-non-repeating string from a valid $(n, d)-\mathrm{ToC} \mathcal{T}$. In fact, we will construct two strings $S[\mathcal{T}]$ and $Q[\mathcal{T}]$ over $\mathbb{Z}_{4 n+4}$. We define $\mathbf{s}_{d} \in \mathbb{Z}^{d}$ to be $\mathbf{s}_{d}=(2$, $\ldots, 2,1)\left(s_{1}=1\right)$, and $\mathcal{F}(\mathbf{p})=\left(2 p_{1}, \ldots, 2 p_{d-1}, 2 p_{d}-1\right)$ for $\mathbf{p} \in \mathbb{Z}^{d}$. Then string $S[\mathcal{T}]$ starts with $\mathbf{s}_{d}$ and ends with $\mathcal{F}($ name $($ tail $(\mathcal{T}))) ; Q[\mathcal{T}]$ starts with $\mathcal{F}($ name $($ tail $(\mathcal{T})))$ and ends with $\mathbf{s}_{d}$.

We use the following recursive procedure. Let $r$ be the root of $\mathcal{T}$ and $C_{r}=a_{1} \ldots a_{2 n+1}$ be the connector at $r$. When $d=1, S[\mathcal{T}]=1 b_{1} b_{2} \ldots b_{2 n+1}$ and $Q[\mathcal{T}]=b_{2 n+1} \ldots b_{2} b_{1} 1$, where $b_{i}=2 a_{i}-1$. When $d \geq 2$,

1. let $\mathcal{T}_{i}$ be the subtree of $\mathcal{T}$ rooted at the $\left(a_{i}\right)^{t h}$-successor of $r$ and let $\mathbf{p}_{i} \in\left(\mathbb{F}_{n}\right)^{d-1}$ be the name of the tail of $\mathcal{T}_{i}$ given by $\mathcal{T}_{i}$ (not by $\mathcal{T}$ ).
2. for every odd integer $i \in[1: 2 n+1]$, set string $S_{i}^{\prime}=$ insert $_{d-1}\left(S\left[\mathcal{T}_{i}\right], 2 a_{i}\right)$ which starts with $\mathbf{s}_{d-1}$ and ends with $\mathcal{F}\left(\mathbf{p}_{i}\right)$; for every even $i \in[1: 2 n+1]$, set string $S_{i}^{\prime}=$ insert $_{d-1}\left(Q\left[\mathcal{T}_{i}\right], 2 a_{i}\right)$ which starts with $\mathcal{F}\left(\mathbf{p}_{i}\right)$ and ends with $\mathbf{s}_{d-1} ; S[\mathcal{T}]=\mathbf{s}_{d} \circ_{d-1} S_{1}^{\prime} \circ_{d-1} S_{2}^{\prime} \circ_{d-1}$ $S_{3}^{\prime} \circ_{d-1} \ldots \circ_{d-1} S_{2 n}^{\prime} \circ_{d-1} S_{2 n+1}^{\prime} ;$
3. for every odd integer $i \in[1: 2 n+1]$, we set string $Q_{i}^{\prime}=$ insert $_{d-1}\left(Q\left[\mathcal{I}_{i}\right], 2 a_{i}\right)$, which starts with $\mathcal{F}\left(\mathbf{p}_{i}\right)$ and ends with $\mathbf{s}_{d-1}$; for every even $i \in[1: 2 n+1]$, $Q_{i}^{\prime}=$ insert $_{d-1}\left(S\left[\mathcal{T}_{i}\right], 2 a_{i}\right)$, which starts with $\mathbf{s}_{d-1}$ and ends with $\mathcal{F}\left(\mathbf{p}_{i}\right) ; Q[\mathcal{T}]=\left(2 a_{2 n+1}\right) \circ Q_{2 n+1}^{\prime} \circ_{d-1}$ $Q_{2 n}^{\prime} \circ{ }_{d-1} \cdots \circ_{d-1} Q_{2}^{\prime} \circ{ }_{d-1} Q_{1}^{\prime} \circ \mathbf{s}_{d}$.

The two strings for the example in Figure 3 are:

$$
\begin{aligned}
& S[\mathcal{T}]={ }^{2} 1^{4} 3^{4} 5^{4} 7^{4} 9^{4} 11^{10} 9^{10} 7^{10} 5^{10} 3^{10} 1^{12} 3^{12}{ }_{9}{ }^{12} 11 \\
& { }^{12} 5^{12} 7^{6} 5^{6} 11^{6} 9^{6} 3^{6} 1^{8} 3^{8} 5_{5}^{8} 7^{8} 9^{8} 11 \\
& Q[\mathcal{T}]={ }^{8} 11{ }^{8} 9^{8} 7^{8} 5_{5}^{8} 3^{8} 1^{6} 3_{3}{ }_{9}{ }^{6}{ }_{11} 6_{5}{ }^{6} 7_{7}{ }^{12} 5^{12}{ }_{11}{ }^{12}{ }_{9}{ }^{12} 3 \\
& { }^{12} 1^{10} 3^{10} 5^{10} 7^{10} 9{ }_{9}{ }^{10} 11^{4} 9^{4} 7_{5}^{4} 3_{3} 1^{2} 1
\end{aligned}
$$

The correctness of our construction (and thus, Theorem 4.5) can be established using the next two lemmas [10].

Lemma 4.6 (Non-Repeating). If $\mathcal{T}$ is a valid $(n, d)$-ToC, then both $S[\mathcal{T}]$ and $Q[\mathcal{T}]$ are d-non-repeating.

Lemma 4.7 (Asking $\mathbb{B}_{\mathcal{T}}$ ). Suppose $\mathcal{T}$ is a valid $(n, d)$-ToC, $S=S[\mathcal{T}]$ and $Q=Q[\mathcal{T}]$. For any $\mathbf{u} \in \mathbb{Z}_{4 n+4}^{d}$, we can compute $\mathbb{B}_{S}(\mathbf{u})$ and $\mathbb{B}_{Q}(\mathbf{u})$ by querying $\mathbb{B}_{\mathcal{T}}$ at most once.

### 4.2 Knowledge Representation in $\mathrm{NT}^{d}$

An algorithm for $\mathrm{NT}^{d}$ tries to learn about the connectors in $\mathcal{T}^{*}$ by repeatedly querying its leaves. To capture its intermediate knowledge about this $\mathcal{T}^{*}$, we introduce a notion of partial connectors.

Let $\sigma=[\sigma(1), \ldots, \sigma(k)]$ be an array of distinct elements from $\{0,1, \ldots n\}$. Then, $\sigma$ defines a string $S_{\sigma(1)} \circ \ldots \circ S_{\sigma(k)}$, referred to as a connecting segment. Recall $S_{0}=2, S_{1}=$ $3 \circ 4, \ldots S_{n}=(2 n+1) \circ(2 n+2)$. A partial connector over $\mathbb{J}_{n}$ is then a set $\mathcal{C}$ of connecting segments such that (1) each $j \in \mathbb{J}_{n}$ is contained in exactly one segment in $\mathcal{C}$; and (2) 2 is the first element of the segment containing it. If $\mathcal{C}$ has $n+1$ segments, that is, $\mathcal{C}=\{2,3 \circ 4, \ldots,(2 n+1) \circ(2 n+2)\}$, then $\mathcal{C}$ is called an empty connector. We say a connector $C$ is consistent with a partial connector $\mathcal{C}$ if every segment in $\mathcal{C}$ is a substring of $C$.

Let $r[\mathcal{C}]$ be the last symbol of the segment in $\mathcal{C}$ that starts with 2. Let $L[\mathcal{C}]$ and $R[\mathcal{C}]$, respectively, be the set of first and the last symbols of other segments in $\mathcal{C}$. So, $r[\mathcal{C}] \in \mathbb{F}_{n} \cup$ $\{2\}, L[\mathcal{C}] \subset \mathbb{O}_{n}$, and $R[\mathcal{C}] \subset \mathbb{F}_{n}$. Also, $|L[\mathcal{C}]|=|R[\mathcal{C}]|$. If $2 \neq r[\mathcal{C}]$, we use $\phi_{\mathcal{C}}(2)$ to denote its right neighbor. Note that each $s \in \mathbb{J}_{n}-L[\mathcal{C}] \cup R[\mathcal{C}] \cup\{r[\mathcal{C}], 2\}$ has two neighbors in $\mathcal{C}$. If $s$ is even, we let $\phi_{\mathcal{C}}(s)$ denote its right neighbor and if $s$ is odd, we let $\phi_{\mathcal{C}}(s)$ denote its left neighbor.

Initially, the knowledge of an algorithm for $\mathrm{NT}^{d}$ can be viewed as a tree $\mathcal{T}$ of empty connectors. At each round, the algorithm chooses a query point $\mathbf{q}$ (according to its knowledge about $\mathcal{T}^{*}$ ) and asks for $\mathbb{B}_{\mathcal{T}^{*}}(\mathbf{q})$, which may connect some segments in the partial connectors, and $\mathcal{T}$ is updated. The algorithm succeeds when $\mathcal{T}$ grows into $\mathcal{T}^{*}$.

At intermediate steps, the knowledge of the algorithm can be expressed by a tree $\mathcal{T}$ of partial connectors. To simplify our proof, we relax the oracle $\mathbb{B}_{\mathcal{T}^{*}}$ to sometime provide more information to the algorithm than being asked (which will be explained in Section 4.3) so that $\mathcal{T}$ always satisfies the conditions of the following definition:
Definition 4.8 (Valid Tree of Partial Connectors). $\mathcal{C}$ is a $\beta$ partial connector for $0<\beta<1$ if the number of segments in $\mathcal{C}$ is at least $(1-\beta) n+1$.
$\mathcal{T}$ is a valid $(n, d, \beta)-\mathrm{ToPC}$ if the root of $T_{n, d}$ is associated with a $\beta$-partial connector. Moreover, for each internal node $v \in T_{n, d}$ whose children are not leaves, if $v$ has a $\beta$ partial connector $\mathcal{C}_{v}$, then: (1) The $s^{\text {th }}$-successor of $v$, for each $s \in L[\mathcal{C}] \cup R[\mathcal{C}] \cup\{r[\mathcal{C}]\}$, has a $\beta$-partial connector; (2) For all $s, t \in \mathbb{J}_{n}$ with $\phi_{\mathcal{C}_{v}}(s)=t$, the tree $\mathcal{T}_{s}$ rooted at the $s^{\text {th }}$-successor $v_{s}$ and the tree $\mathcal{T}_{t}$ rooted at the $t^{\text {th }}$ successor $v_{t}$ of $v$ are both valid ToCs, and name (tail $\left(v_{s}\right)$ ) in $\mathcal{T}_{s}$ is the same as name (tail $\left.\left(v_{t}\right)\right)$ in $\mathcal{T}_{t}$.

A valid $(n, d)-\operatorname{ToC} \mathcal{T}^{*}$ is consistent with a valid ( $n, d$, $\beta)-\operatorname{ToPC} \mathcal{T}$, denoted by $\mathcal{T} \models \mathcal{T}^{*}$, if for each internal node of $T_{n, d}$, its connector in $\mathcal{T}^{*}$ is consistent with its (partial) connector in $\mathcal{T}$.

The definition implies: Given a valid $(n, d, \beta)$-ToPC $\mathcal{T}$, for each internal node $v$ of $T_{n, d}$, it is associated with either a connector or a $\beta$-partial connector, and moreover, the subtree of $\mathcal{T}$ rooted at $v$ is either a valid $(n$, height $(v))$-ToC or a valid ( $n$, height $(v), \beta$ )-ToPC.

A key to our analysis is Lemma 4.9 below, stating that every valid $(n, d, \beta)$-ToPC has a large number of consistent valid $(n, d)$-ToCs. Moreover, the tails of these valid ToCs are nearly-uniformly distributed. We let set $\mathcal{F}[\mathcal{T}]=$ \{name (tail $\left.\left.\left(\mathcal{T}^{*}\right)\right) \mid \mathcal{T} \models \mathcal{T}^{*}\right\}$. Also, for each $\mathbf{p} \in\left(\mathbb{F}_{n}\right)^{d}$, let $N[\mathcal{T}, \mathbf{p}]=\mid\left\{\mathcal{T}^{*} \mid \mathcal{T} \models \mathcal{T}^{*}\right.$, name $\left(\right.$ tail $\left.\left.\left(\mathcal{T}^{*}\right)\right)=\mathbf{p}\right\} \mid$.

Lemma 4.9 (Key Lemma). For $d \geq 1$ and $\beta \in\left[0,24^{-d}\right]$, $|\mathcal{F}[\mathcal{T}]| \geq((1-\beta) n)^{d}$ for each valid $(n, d, \beta)$-ToPC $\mathcal{T}$. Moreover, for all $\mathbf{p}_{1}, \mathbf{p}_{2} \in \mathcal{F}[\mathcal{T}]$,

$$
\begin{gather*}
\frac{1}{\alpha_{d}(\beta)} \leq \frac{N\left[\mathcal{T}, \mathbf{p}_{1}\right]}{N\left[\mathcal{T}, \mathbf{p}_{2}\right]} \leq \alpha_{d}(\beta), \text { where } \alpha_{1}(\beta)=1  \tag{1}\\
\text { and } \alpha_{d}(\beta)=\frac{\left(\alpha_{d-1}(\beta)\right)^{7}}{\left(2(1-\beta)^{d-1}-1\right)^{3}}, \text { for } d \geq 2 \tag{2}
\end{gather*}
$$

Proof. When $d=1$, let $\mathcal{C}$ be the only partial connector in $\mathcal{T}$. Clearly, $\mathcal{F}[\mathcal{T}]=R[\mathcal{C}]$. Thus, in this case the lemma is true. We will use this case as the base of the induction.

We will prove by induction on $d$ that both (1) and (**) $|\mathcal{F}[\mathcal{T}]| \geq((1-\beta) n)^{d}$ are true for all $d \geq 1$. When $d \geq 2$, let $\mathcal{C}=\left\{Y_{0}, Y_{1}, Y_{2}, \ldots, Y_{m}\right\}$ be the $\beta$-partial connector at the root of $\mathcal{T}$; assume $Y_{0}$ is the segment starting with 2 . We use $r_{i}$ and $t_{i}$, respectively, to denote the ending and starting symbols of $Y_{i}$. For each $k \in\left\{r_{0}, r_{1} \ldots, r_{m}, t_{1}, \ldots, t_{m}\right\}$, let $\mathcal{T}_{k}$ denote the $(n, d-1, \beta)$-ToPC at the $k^{t h}$-successor of the root. For each $(i, j) \in[0: m] \times[1: m]$ with $i \neq j$, we let

$$
\begin{equation*}
N_{i, j}=\sum_{\mathbf{p} \in \mathcal{F}\left[\mathcal{T}_{r_{i}}\right] \cap \mathcal{F}\left[\mathcal{T}_{t_{j}}\right]}\left(N\left[\mathcal{T}_{r_{i}}, \mathbf{p}\right] \cdot N\left[\mathcal{T}_{t_{j}}, \mathbf{p}\right]\right) \tag{3}
\end{equation*}
$$

Inductively, (1) and ( ${ }^{* *}$ ) hold for $d-1$. Since $\beta \leq 24^{-d}$ $<24^{-(d-1)}$, we have

$$
\begin{aligned}
& \left|\mathcal{F}\left[\mathcal{T}_{r_{i}}\right] \cap \mathcal{F}\left[\mathcal{T}_{t_{j}}\right]\right|=\left|\mathcal{F}\left[\mathcal{T}_{r_{i}}\right]\right|+\left|\mathcal{F}\left[\mathcal{T}_{t_{j}}\right]\right| \\
& \quad-\left|\mathcal{F}\left[\mathcal{T}_{r_{i}}\right] \cup \mathcal{F}\left[\mathcal{T}_{t_{j}}\right]\right| \geq\left(2(1-\beta)^{d-1}-1\right) n^{d-1}>0
\end{aligned}
$$

so $N_{i, j}>0$ for all $(i, j) \in[0: m] \times[1: m]$ with $i \neq j$.
To show (**), it suffices to prove $\mathcal{F}[\mathcal{T}]=\cup_{k \in R[\mathcal{C}]}(k \circ$ $\left.\mathcal{F}\left[\mathcal{T}_{k}\right]\right)$. Clearly, for any $\mathbf{p} \notin \cup_{k \in R[\mathcal{C}]}\left(k \circ \mathcal{F}\left[\mathcal{T}_{k}\right]\right), N[\mathcal{T}, \mathbf{p}]$ $=0$. So, let us consider a point $\mathbf{p} \in \cup_{k \in R[\mathcal{C}]}\left(k \circ \mathcal{F}\left[\mathcal{T}_{k}\right]\right)$. Since $p_{1} \in R[\mathcal{C}]$, WLOG, we assume $p_{1}=r_{m}$. We let $\mathcal{P}$ denote the set of permutations $s_{0} s_{1} \ldots s_{m-1}$ over $[0: m-1]$ with $s_{0}=0$. Then

$$
\begin{align*}
& N[\mathcal{T}, \mathbf{p}]=\sum_{s_{0} \ldots s_{m-1} \in \mathcal{P}}\left(\left(\prod_{i=0}^{m-2} N_{s_{i}, s_{i+1}}\right)\right. \\
& \left.\quad \cdot N_{s_{m-1}, m} \cdot N\left[\mathcal{T}_{r_{m}},\left(p_{2}, p_{3}, \ldots, p_{d}\right)\right]\right)>0 \tag{4}
\end{align*}
$$

and thus, $\mathcal{F}[\mathcal{T}]=\cup_{k \in R[\mathcal{C}]}\left(k \circ \mathcal{F}\left[\mathcal{T}_{k}\right]\right)$.

To prove (1), consider $\mathbf{p}_{1}$ and $\mathbf{p}_{2} \in \mathcal{F}[\mathcal{T}]$. There are two basic cases. When $p_{1,1}=p_{2,1}$, Eqn. (1) follows directly from (4) and the inductive hypothesis. When $p_{1,1} \neq p_{2,1}$, WLOG, we assume $p_{1,1}=r_{m}$ and $p_{2,1}=r_{m-1}$.

Let $\mathcal{P}_{1}$ denote the set of permutations over $\{0,1, \ldots, m-$ $2, m-1\}$ with $s_{0}=0$, and $\mathcal{P}_{2}$ denote the set of permutations over $\{0,1, \ldots, m-2, m\}$ with $s_{0}=0$. For each $P=s_{0} s_{1} \ldots s_{m-1} \in \mathcal{P}_{1}$, let $\Pi(P) \in \mathcal{P}_{2}$ be the permutation obtained from $P$ by replacing $m-1$ by $m$. Clearly $\Pi$ is a bijection from $\mathcal{P}_{1}$ to $\mathcal{P}_{2}$.

Now by mimicking Eqn. (4), we can write $N\left[\mathcal{T}, \mathbf{p}_{1}\right]$ and $N\left[\mathcal{T}, \mathbf{p}_{2}\right]$ as two summations: $N\left[\mathcal{T}, \mathbf{p}_{1}\right]=\sum_{P \in \mathcal{P}_{1}} N_{1}(P)$, and $N\left[\mathcal{T}, \mathbf{p}_{2}\right]=\sum_{P \in \mathcal{P}_{1}} N_{2}(\Pi(P))$, where $N_{1}(P)$ and $N_{2}(\Pi(P))$ are given by similar terms as in (4).

We prove for all $P \in \mathcal{P}_{1},\left(N_{1}(P) / N_{2}(\Pi(P)) \leq \alpha_{d}(\beta)\right.$, from which Eqn. (1) follows. Let $P=s_{0} s_{1} \ldots s_{m-1}$ where $s_{k}=m-1$ for some $1 \leq k \leq m-1$. For the case when $k<m-1$, after expanding $N_{1}(P)$ and $N_{2}(\Pi(P))$, and canceling out the common terms, we get

$$
\begin{aligned}
\frac{N_{1}(P)}{N_{2}(\Pi(P))} & =\frac{N_{s_{k-1}, m-1} \cdot N_{m-1, s_{k+1}}}{N_{s_{k-1}, m} \cdot N_{m, s_{k+1}}} \\
& \times \frac{N_{s_{m-1}, m} \cdot N\left[\mathcal{T}_{r_{m}},\left(p_{1,2}, p_{1,3}, \ldots, p_{1, d}\right)\right]}{N_{s_{m-1}, m-1} \cdot N\left[\mathcal{T}_{r_{m-1}},\left(p_{2,2}, p_{2,3}, \ldots, p_{2, d}\right)\right]}
\end{aligned}
$$

We then expand each term $N_{i, j}$ using the definition (Eqn. (3)). It follows from the application of the inductive hypothesis that $N_{1}(P) / N_{2}(\Pi(P)) \leq \alpha_{d}(\beta)$. Similarly, we can establish the same bound for the case $k=m-1$.

### 4.3 The Randomized Query Complexity

By querying every leaf, one can solve any instance of $\mathrm{NT}^{d}$ with $n^{d}$ queries. Below, we prove Theorem 4.4 by showing $\operatorname{RQ}_{\mathrm{NT}}^{d}(n)=(\Omega(n))^{d}$. We first relax $\mathbb{B}_{\mathcal{T}^{*}}$ by extending its domain to $\left(\mathbb{J}_{n}\right)^{m}$, for all $m \in[1: d]$.

Definition 4.10 (Relaxation of $\mathbb{B}_{\mathcal{T}^{*}}$ ). Suppose $\mathcal{T}^{*}$ is a valid $(n, d)-T o C$ and $\mathbf{q} \in\left(\mathbb{J}_{n}\right)^{m}$. Let $v$ be the node such that name $(v)=q_{1} \ldots q_{m} . \operatorname{Let} \mathbf{q}^{\prime}=$ name $($ tail $(v)) \in\left(\mathbb{J}_{n}\right)^{d}$ (in $\left.\mathcal{T}^{*}\right)$. Then, $\mathbb{B}_{\mathcal{T}^{*}}(\mathbf{q})=\mathbb{B}_{\mathcal{T}^{*}}\left(\mathbf{q}^{\prime}\right)$.

Proof (of Theorem 4.4). We consider the distribution $\mathcal{D}$ in which each valid $(n, d)$-ToC $\mathcal{T}^{*}$ is chosen with the same probability. We will prove that the expected query complexity of any deterministic algorithm $\mathcal{A}$ for $\mathrm{NT}^{d}$ over $\mathcal{D}$ is $(\Omega(n))^{d}$. Let $\beta_{d}=24^{-d}$. To simplify the proof, we assume $n$ is a multiple of $2 \cdot 24^{d}$ so that $\beta_{d} n / 2 \in \mathbb{Z}$.

Suppose, at a particular step, the knowledge of $\mathcal{A}$ can be expressed by a valid $\left(n, d, \beta_{d}\right)$-ToPC $\mathcal{T} \models \mathcal{T}^{*}$, which is clearly true initially, and $\mathcal{A}$ wants to query $\mathbf{q} \in\left(\mathbb{J}_{n}\right)^{d}$. Let $u_{0}$ be the root of $\mathcal{T}$ and $u_{i}$ be the node with name $\left(u_{i}\right)=$ $q_{1} \ldots q_{i}$. Let $\mathcal{C}_{i}$ be the connector or $\beta_{d}$-partial connector at $u_{i}$ in $\mathcal{T}$ and $\mathcal{T}_{i}$ be the subtree of $\mathcal{T}$ rooted at $u_{i}$. There are

```
Query-and-Update \((\mathcal{T}, \mathbf{q})\), where \(\mathbf{q} \in\left(\mathbb{J}_{n}\right)^{d}\)
if \(\exists 0 \leq i \leq d-1:\left|R\left[\mathcal{C}_{i}\right]\right|=\left(1-\beta_{d}\right) n\) then
    set \(m\) be the smallest of such \(i(m \in[0: d-1])\)
else set \(m=d\)
if \(m=0\) then \(\operatorname{set} \mathcal{T}=\mathcal{T}^{*} \quad\{\operatorname{set} I=1\}\)
else Update \(\left(\mathcal{T},\left(q_{1}, q_{2}, \ldots, q_{m}\right), m\right)\)
```

$\operatorname{Update}(\mathcal{T}, \mathbf{q}, m)$, where $\mathbf{q} \in\left(\mathbb{J}_{n}\right)^{m}$ and $1 \leq m \leq d$
fetch $\mathbb{B}_{\mathcal{T}^{*}}(\mathbf{q})$
$\left\{\right.$ set $\left.A_{m}=A_{m}+1, \mathcal{B}_{m}\left[A_{m}\right]=0, \mathcal{B}_{m, k}\left[A_{m}\right]=0\right\}$
if $\mathbb{B}_{\mathcal{T}^{*}}(\mathbf{q})=\mathcal{T}^{*}$ then set $\mathcal{T}=\mathcal{T}^{*}\left\{\right.$ set $\left.\mathcal{B}_{m}\left[A_{m}\right]=1\right\}$
else $\left[\right.$ let $\left.\mathbb{B}_{\mathcal{T}^{*}}(\mathbf{q})=\left(h, r, \mathcal{T}^{\prime}, \mathcal{T}^{\prime \prime}\right), m^{\prime}=d-h-1\right]$
$\exists Y_{1}, Y_{2} \in \mathcal{C}_{m^{\prime}}$ such that, $\{$ the ending symbol
of $Y_{1}$, the starting symbol of $\left.Y_{2}\right\}=\left\{q_{m^{\prime}+1}, r\right\}$
10: replace $Y_{1}$ and $Y_{2}$ in $\mathcal{C}_{m^{\prime}}$ by the concatenation
of $Y_{1}$ and $Y_{2}\left\{\right.$ set $\left.\mathcal{B}_{m, m^{\prime}}\left[A_{m}\right]=1\right\}$
replace the subtree of $\mathcal{T}$ rooted at $u_{m^{\prime}+1}$ with $\mathcal{T}^{\prime}$;
replace the subtree of $\mathcal{T}$ rooted at the $r$-successor
of $u_{m^{\prime}}$ with $\mathcal{T}^{\prime \prime}$

Figure 4. The Query-and-Update procedure
two cases (1) $\forall i \in[0, d-1], \mathcal{C}_{i}$ is a $\beta_{d}$-partial connector and $q_{i+1} \in L\left[\mathcal{C}_{i}\right] \cup R\left[\mathcal{C}_{i}\right] \cup\left\{r\left[\mathcal{C}_{i}\right]\right\} ;(2)$ otherwise. By the definition of $\mathbb{B}_{\mathcal{T}^{*}}$, we can show that in case $(2), \mathbb{B}_{\mathcal{T}^{*}}(\mathbf{q})$ can be answered based on $\mathcal{T}$ only. So, WLOG, we assume $\mathcal{A}$ never asks queries of the second type.

In case (1), $\mathcal{C}_{i}$ is a $\beta_{d}$-partial connector for all $i \in[0$ : $d-1]$. Let $h=$ height $(\operatorname{head}(\mathbf{q}))$. If $h=d$, then $\mathcal{A}$ gets $\mathcal{T}^{*}$. Otherwise, the knowledge gained by querying $\mathbb{B}_{\mathcal{T}^{*}}(\mathbf{q})$ connects two segments in $\mathcal{C}_{d-h-1}$ and replaces the two involved subtrees by the corresponding ones in the 4 -tuple $\mathbb{B}_{\mathcal{T}^{*}}(\mathbf{q})$. The resulting tree $\mathcal{T}$, however, may no longer be a valid $\left(n, d, \beta_{d}\right)-T o P C$, if $\left|R\left[\mathcal{C}_{d-h-1}\right]\right|=\left(1-\beta_{d}\right) n$ (or the number of segments in $\mathcal{C}_{d-h-1}$ is exactly $\left.\left(1-\beta_{d}\right) n+1\right)$ before the query. When this happens, we will relax $\mathbb{B}_{\mathcal{T}^{*}}$ to provide $\mathcal{A}$ more information to ensure that the resulting $\mathcal{T}$ remains a valid $\left(n, d, \beta_{d}\right)$-ToPC.

To this end, we consider two subcases when $\mathcal{A}$ queries $\mathbf{q} \in\left(\mathbb{J}_{n}\right)^{d}:$ Case (1.a) $\forall i \in[0: d-1],\left|R\left[\mathcal{C}_{i}\right]\right|>\left(1-\beta_{d}\right) n$, then $\mathcal{A}$ receives $\mathbb{B}_{\mathcal{T}^{*}}(\mathbf{q})$ as it requested; Case (1.b) $\exists i \in$ $[0: d-1]$ such that $\left|R\left[\mathcal{C}_{i}\right]\right|=\left(1-\beta_{d}\right) n$, then we assume $m=\min \left\{i:\left|R\left[\mathcal{C}_{i}\right]\right|=\left(1-\beta_{d}\right) n\right\}$. Let $\mathbf{q}^{\prime}=\left(q_{1}, \ldots, q_{m}\right)$. Instead of getting $\mathbb{B}_{\mathcal{T}^{*}}(\mathbf{q}), \mathcal{A}$ gets $\mathbb{B}_{\mathcal{T}^{*}}\left(\mathbf{q}^{\prime}\right)$. In this way, the resulting $\mathcal{T}$ remains a valid $\left(n, d, \beta_{d}\right)$-ToPC (or grows into $\mathcal{T}^{*}$ and $\mathcal{A}$ stops) after the query. Details of the query-andupdate procedure can be found in Figure 4.

We introduce some "analysis variables" to aid our analysis. These variables include: (1) $I \in\{0,1\}$ : Initially, $I$ $=0$. (2) For each $m \in[1: d], A_{m} \in \mathbb{Z}$, and a set of bi-
nary sequences $\mathcal{B}_{m}[\ldots]$ and $\mathcal{B}_{m, k}[\ldots]$ of length $A_{m}, \forall k \in$ $[0: m-1]$. Initially, $A_{m}=0$, and $\mathcal{B}_{m}, \mathcal{B}_{m, k}$ are empty.

In each step, the variables are updated as follows: (I) If $m=0$ in case (1.b), then we set $I=1$; (II) Otherwise, (to unify the discussion in this part, if we have case (1.a), then set $m=d$ and $\mathbf{q}^{\prime}=\mathbf{q}$ ) we increase $A_{m}$ by 1 ; (II.1) if $\mathbb{B}_{\mathcal{T}^{*}}\left(\mathbf{q}^{\prime}\right)=\mathcal{T}^{*}$, we set $\mathcal{B}_{m}\left[A_{m}\right]=1$ and $\mathcal{B}_{m, k}\left[A_{m}\right]=0$, $\forall k \in[0: m-1]$; (II.2) otherwise, if the first component of $\mathbb{B}_{\mathcal{T}^{*}}\left(\mathbf{q}^{\prime}\right)$ is $d-l$, for $l \in[1: m]$, then $\mathcal{B}_{m, l-1}\left[A_{m}\right]=1$, $\mathcal{B}_{m}\left[A_{m}\right]=\mathcal{B}_{m, k}\left[A_{m}\right]=0, \forall k: 0 \leq k \neq l-1 \leq m-1$. Details can also be found in Figure 4 (highlighted in $\}$ ).

Let $M_{i}=\left(\beta_{d} n / 2\right)^{i}, \forall i \in[1: d]$. Given a valid $(n, d)-$ $\operatorname{ToC} \mathcal{T}^{*}$, if $\mathcal{A}$ stops before making $M_{d}$ queries, then we let $\left\{I, A_{m}, \mathcal{B}_{m}, \mathcal{B}_{m, k}\right\}$ be the variables assigned when $\mathcal{A}$ stops; otherwise, $\left\{I, A_{m}, \mathcal{B}_{m}, \mathcal{B}_{m, k}\right\}$ is assigned after $\mathcal{A}$ makes $M_{d}$ queries. Furthermore, we define a set of binary strings $\left\{\overline{\mathcal{B}}_{m}\left[1 \ldots M_{m}\right], \overline{\mathcal{B}}_{m, k}\left[1 \ldots M_{m}\right], 1 \leq m \leq d, 0 \leq k \leq\right.$ $m-1\}$ from $\left\{\mathcal{B}_{m}, \mathcal{B}_{m, k}\right\}$ : for each $1 \leq i \leq M_{m}, \overline{\mathcal{B}}_{m}[i]=$ $\mathcal{B}_{m}[i]$ and $\overline{\mathcal{B}}_{m, k}[i]=\mathcal{B}_{m, k}[i]$ when $i \leq \min \left(A_{m}, M_{m}\right)$; $\overline{\mathcal{B}}_{m}[i]=\overline{\mathcal{B}}_{m, k}[i]=0$ when $A_{m}<i \leq M_{m}$.

As $\mathcal{T}^{*}$ is chosen randomly from valid $(n, d)-\mathrm{ToCs},\{I$, $\left.A_{m}, \overline{\mathcal{B}}_{m}, \overline{\mathcal{B}}_{m, k}\right\}$ is a set of random variables defined by the deterministic algorithm $\mathcal{A}$. To assist the analysis, we introduce the following definition. Lemma 4.12 is an important step in our analysis.

Definition 4.11 (c-Biased Distributions). Suppose we have a probabilistic distribution over strings $\{0,1\}^{m}$. For every binary string $S$ of length at most $m$, we define set $U_{S}=$ $\left\{S^{\prime} \in\{0,1\}^{m} \mid S\right.$ is a prefix of $\left.S^{\prime}\right\}$. For $0 \leq c \leq 1$, the distribution is said to be $c$-biased if we have $\operatorname{Pr}\left[U_{1}\right] \leq c$ and $\operatorname{Pr}\left[U_{S \circ 1}\right] \leq c \cdot \operatorname{Pr}\left[U_{S}\right]$ for all $S$ with $1 \leq|S| \leq m-1$.
Lemma 4.12. For all $1 \leq m \leq d$, the distribution over $\overline{\mathcal{B}}_{m}$ is $2 / n^{m}$-biased; For $2 \leq m \leq d$ and $0 \leq k \leq m-2$, the distribution over $\overline{\mathcal{B}}_{m, k}$ is $2 / n^{m-k-1}$-biased.

The lemma follows directly from Corollary 4.13 below of our Key Lemma (4.9).
Corollary 4.13. For $d \geq 1$ and $\beta \in\left[0,24^{-d}\right]$, let $\mathcal{T}$ be a valid $(n, d, \beta)-T o P C$, and $N=\sum_{\mathbf{p} \in \mathcal{F}[\mathcal{T}]} N[\mathcal{T}, \mathbf{p}]$ be the number of consistent ToCs.

For $\mathbf{q} \in\left(\mathbb{J}_{n}\right)^{m}$ where $m \in[1: d]$, let $u_{0}$ be the root of $\mathcal{T}$ and $u_{i}$ be the node with name $\left(u_{i}\right)=q_{1} \ldots q_{i}$. Let $\mathcal{C}_{i}$ be the connector or $\beta_{d}$-partial connector at $u_{i}$ and $\mathcal{T}_{i}$ be the subtree of $\mathcal{T}$ rooted at $u_{i}$. If $\forall i \in[0, m-1], \mathcal{C}_{i}$ is a $\beta_{d}$-partial connector and $q_{i+1} \in L\left[\mathcal{C}_{i}\right] \cup R\left[\mathcal{C}_{i}\right] \cup\left\{r\left[\mathcal{C}_{i}\right]\right\}$, then (1) $\left(N^{*} / N\right) \leq\left(2 / n^{m}\right)$ where $N^{*}=\mid\left\{\mathcal{T}^{\prime} \mid \mathbb{B}_{\mathcal{T}^{\prime}}(\mathbf{q})=\right.$ $\mathcal{T}^{\prime}$ and $\left.\mathcal{T} \models \mathcal{T}^{\prime}\right\} \mid$; and (2) $\left(N_{k} / N\right) \leq 2 / n^{m-k-1}$ where for $0 \leq k \leq m-2, N_{k}$ is the number of consistent ToCs $\mathcal{T}^{\prime}$ such that the first component of $\mathbb{B}_{\mathcal{T}^{\prime}}(\mathbf{q})$ is $d-k-1$.
Proof. For each $k \in[0: m-1]$, let $W_{k}=\left\{\mathbf{p} \in \mathcal{F}\left[\mathcal{T}_{k}\right] \subset\right.$ $\left(\mathbb{F}_{n}\right)^{d-k}$ such that $\left.p_{i}=q_{k+i}, \forall i \in[1: m-k]\right\}$. Clearly, $\left|W_{k}\right| \leq n^{d-m}$. By Lemma 4.9, we have

$$
\frac{N^{*}}{N}=\frac{\sum_{\mathbf{p} \in W_{0}} N[\mathcal{T}, \mathbf{p}]}{\sum_{\mathbf{p} \in \mathcal{F}[\mathcal{T}]} N[\mathcal{T}, \mathbf{p}]} \leq \frac{\alpha_{d}(\beta) \cdot\left|W_{0}\right|}{|\mathcal{F}[\mathcal{T}]|} \leq \frac{2}{n^{m}}
$$

The third inequality uses Proposition A. 3 in Appendix A.
To prove the second statement, for $k \in[0: m-2]$, we consider any connector $C^{*}$ over $\mathbb{J}_{n}$ that is consistent with $\mathcal{C}_{k}$ and satisfies $\phi_{C^{*}}\left(q_{k+1}\right) \neq$ "no" (that is, $\left.r\left[C^{*}\right] \neq q_{k+1}\right)$. Assume $\phi_{C^{*}}\left(q_{k+1}\right)=r$. We let $\mathcal{T}^{\prime}$ denote the subtree of $\mathcal{T}$ rooted at the $r^{t h}$-successor of $u_{k}$. Since $q_{k+1} \in L\left[\mathcal{C}_{k}\right] \cup$ $R\left[\mathcal{C}_{k}\right] \cup\left\{r\left[\mathcal{C}_{k}\right]\right\}$, both $\mathcal{T}_{k+1}$ and $\mathcal{T}^{\prime}$ are $(n, d-k-1, \beta)-$ ToPCs. Therefore, we have

$$
\begin{aligned}
& \frac{\sum_{\mathbf{p} \in W_{k+1} \cap \mathcal{F}\left[\mathcal{T}_{k+1}\right] \cap \mathcal{F}\left[\mathcal{T}^{\prime}\right]} N\left[\mathcal{T}_{k+1}, \mathbf{p}\right] \cdot N\left[\mathcal{T}^{\prime}, \mathbf{p}\right]}{\sum_{\mathbf{p} \in \mathcal{F}\left[\mathcal{T}_{k+1}\right] \cap \mathcal{F}\left[\mathcal{T}^{\prime}\right]} N\left[\mathcal{T}_{k+1}, \mathbf{p}\right] \cdot N\left[\mathcal{T}^{\prime}, \mathbf{p}\right]} \\
& \leq \frac{\left(\alpha_{d-k-1}(\beta)\right)^{2} \cdot n^{d-m}}{\left(2(1-\beta)^{d-k-1}-1\right) \cdot n^{d-k-1}} \leq \frac{2}{n^{m-k-1}}
\end{aligned}
$$

Let $B_{m}, B_{m, k}, \bar{B}_{m}$ and $\bar{B}_{m, k}$ denote the number of 1's in $\mathcal{B}_{m}, \mathcal{B}_{m, k}, \overline{\mathcal{B}}_{m}$ and $\overline{\mathcal{B}}_{m, k}$, respectively. Then,

Lemma 4.14. For all $m \in[1: d]$ and $k \in[0: m-2]$, we have $\operatorname{Pr}_{\mathcal{D}}\left[\bar{B}_{m}>0\right]<1 /\left(2 d^{2}\right)$ and

$$
\operatorname{Pr}_{\mathcal{D}}\left[\bar{B}_{m, k}>\frac{16 \cdot M_{m}}{n^{m-k-1}}\right]<\frac{1}{2 d^{2}}
$$

Proof. We will use the following fact: Let $\mathcal{D}_{I N D}^{m}$ denote the distribution over $\{0,1\}^{m}$ where each bit of the string is chosen independently and is equal to 1 with probability $c$. For any $c$-biased distribution $\mathcal{D}^{m}$ over $\{0,1\}^{m}$ and integer $k \in[1: m]$, we have $\operatorname{Pr}_{S \leftarrow \mathcal{D}^{m}}[S$ has at least $k 1$ 's $] \leq$ $\operatorname{Pr}_{S \leftarrow \mathcal{D}_{I N D}^{m}}[S$ has at least $k 1$ 's $]$.

By Lemma 4.12, $\operatorname{Pr}_{\mathcal{D}}\left[\bar{B}_{m}>0\right] \leq 1-\left(1-2 n^{-m}\right)^{M_{m}} \leq$ $4\left(\beta_{d} / 2\right)^{m} \leq 1 / 2 d^{2}$. The second inequality uses Propositions A. 2 and A.1, and the last inequality uses $\beta_{d}=24^{-d}$ and the fact $m \geq 1$. We can apply Lemma 4.12 and the Chernoff bound [11] to prove the second statement.

Let $[A]$ denote the event that $A$ is true. Let NOT-YET$\operatorname{FOUND}\left(\mathcal{T}^{*}\right)$ denote the event that algorithm $\mathcal{A}$ has not found the tail of $\mathcal{T}^{*}$ after making $M_{d}=\left(\beta_{d} n / 2\right)^{d}$ queries. Then [NOT-YET-FOUND $\left.\left(\mathcal{T}^{*}\right)\right] \Longleftrightarrow\left[I=0\right.$ and $B_{m}=0$, $\forall m \in[1: d]]$. The theorem directly follows from Lemma 4.15 below.

Lemma 4.15. Let $A$ denote the following event:

$$
\begin{gathered}
\quad\left(\bar{B}_{m}=0 \text { and } \bar{B}_{m, m-1} \leq M_{m}\right. \text { and } \\
\left.\bar{B}_{m, k} \leq \frac{16 \cdot M_{m}}{n^{m-k-1}}, \forall m \in[1: d], k \in[0, m-2]\right),
\end{gathered}
$$

then we have (E.1) $[A]$ implies [NOT-YET-FOUND $\left(\mathcal{T}^{*}\right)$ ] and (E.2) $\operatorname{Pr}_{\mathcal{D}}[A] \geq 1 / 2$.

Proof. To prove (E.1), we need the following inequalities which follow from the definition of our analysis variables.

$$
\begin{gather*}
A_{m} \leq \frac{1}{\beta_{d} n} \sum_{i=m+1}^{d} B_{i, m}, \forall 1 \leq m \leq d-1  \tag{5}\\
I=1 \Longrightarrow \sum_{i=1}^{d} B_{i, 0} \geq \beta_{d} n \tag{6}
\end{gather*}
$$

To prove (E.1), it suffices to show that $[A] \Rightarrow[I=0]$ and $[A] \Rightarrow\left[B_{m}=0, \forall m \in[1: d]\right]$. Here we use $[A] \Rightarrow$ $[B]$ to denote if event $A$ is true then event $B$ is true. It follows immediately from the definitions of $B_{m}$ and $\bar{B}_{m}$, that if $A_{m} \leq M_{m}$, then $B_{m}=\bar{B}_{m}$. So, we first inductively prove that $[A] \Rightarrow\left[A_{d-m} \leq M_{d-m}, \forall m \in[0: d-1]\right]$.

The base case when $m=0$ is trivial, since $A_{d}$ is at most $M_{d}$, the total number of queries. We now consider the case when $m \geq 1$, and assume inductively, that $A_{i} \leq M_{i}$ for all $i \in[d-m+1: d]$. Consequently, for all $i \in[d-m+1: d]$ and $j \in[0, i-1], \bar{B}_{i}=B_{i}$ and $\bar{B}_{i, j}=B_{i, j}$. By Eqn. (5),

$$
\begin{aligned}
& A_{d-m} \leq \sum_{i=d-m+1}^{d} \frac{B_{i, d-m}}{\beta_{d} n}=\sum_{i=d-m+1}^{d} \frac{\bar{B}_{i, d-m}}{\beta_{d} n} \\
& \leq \frac{1}{\beta_{d} n}\left(M_{d-m+1}+\sum_{i=d-m+2}^{d}\left(\frac{16 \cdot M_{i}}{n^{i-d+m-1}}\right)\right) \\
& \leq M_{d-m}\left(\frac{1}{2}+8 \sum_{i=d-m+2}^{d}\left(\frac{\beta_{d}}{2}\right)^{i-d+m-1}\right)<M_{d-m}
\end{aligned}
$$

Thus, $[A] \Rightarrow\left[B_{m}=0, \forall m \in[1: d]\right]$. Now we prove $[A]$ implies $[I=0]$. By Eqn. (5), it suffices to show $[A]$ implies $\left[\sum_{m=1}^{d} B_{m, 0}<\beta_{d} n\right]$. Assume $[A]$, then

$$
\begin{aligned}
\sum_{i=m}^{d} B_{m, 0} & =\bar{B}_{1,0}+\sum_{m=2}^{d} \bar{B}_{m, 0} \leq M_{1}+\sum_{m=2}^{d} \frac{16 \cdot M_{m}}{n^{m-1}} \\
& =\beta_{d} n\left(\frac{1}{2}+8 \sum_{m=2}^{d}\left(\frac{\beta_{d}}{2}\right)^{m-1}\right)<\beta_{d} n
\end{aligned}
$$

The first equation follows from $[A] \Rightarrow\left[A_{d-m} \leq M_{d-m}\right.$, $\forall m \in[0: d-1]]$ and the first inequality uses $\bar{B}_{m, m-1} \leq$ $M_{m}$ for all $m \in[1: d]$. Finally, by Lemma 4.14, we have

$$
\begin{aligned}
& \operatorname{Pr}_{\mathcal{D}}[A] \geq 1-\left(\sum_{m=1}^{d} \operatorname{Pr}_{\mathcal{D}}\left[\bar{B}_{m}>0\right]\right. \\
& \left.\quad+\sum_{m=1}^{d} \sum_{k=1}^{m-2} \operatorname{Pr}_{\mathcal{D}}\left[\bar{B}_{m, k}>\frac{16 \cdot M_{m}}{n^{m-k-1}}\right]\right) \geq \frac{1}{2}
\end{aligned}
$$

## 5 Open Problems

We conjecture that FPC is also strictly harder than local search in the quantum query model over $[1: n]^{d}$.

Our current lower bound does not apply to the case when $n$ is small and $d$ is large. For example, one can define the End-OF-A-PATH problem over hypercubes $\{0,1\}^{n}$. Both the randomized and quantum query complexity of this problem remains open.

We conclude this paper with the following conjecture: If PPAD is in FP, then PLS is in FP.

## 6 Acknowledgments

We would like to thank Dan Spielman, Xiaoming Sun, Xiaotie Deng, Kyle Burke and Stan Sclaroff for discussions that are invaluable to this work.

## References

[1] S. Aaronson. Lower bounds for local search by quantum arguments. In STOC 2004, pages 465-474.
[2] D. Aldous. Minimization algorithms and random walk on the d-cube. Annals of Probability, 11(2):403-413, 1983.
[3] K. Arrow and G. Debreu. Existence of an equilibrium for a competitive economy. Econometrica, 22(3):265-290, 1954.
[4] P. Beame, S. Cook, J. Edmonds, R. Impagliazzo, and T. Pitassi. The relative complexity of NP search problems. J. Comput. Syst. Sci., 57:3-19, 1998.
[5] L. Brouwer. Über abbildung von mannigfaltigkeiten. Mathematische Annalen, 71:97-115, 1910.
[6] X. Chen and X. Deng. On algorithms for discrete and approximate Brouwer fixed points. In STOC 2005.
[7] X. Chen and X. Deng. On the complexity of 2d discrete fixed point problem. In ICALP 2006, pages 489-500.
[8] X. Chen and X. Deng. Settling the complexity of two-player Nash equilibrium. In FOCS 2006.
[9] X. Chen, X. Deng, and S.-H. Teng. Computing Nash equilibria: Approximation and smoothed complexity. In FOCS 2006.
[10] X. Chen and S.-H. Teng. Paths beyond local search: A nearly tight bound for randomized fixed-point computation. arXiv:cs/0702088, 2007.
[11] H. Chernoff. Asymptotic efficiency for tests based on the sum of observations. Ann. Math. Stat., 23:493-507, 1952.
[12] B. Codenotti, A. Saberi, K. Varadarajan, and Y. Ye. Leontief economies encode nonzero sum two-player games. ECCC, TR05-055, 2005.
[13] G. Dantzig. Maximization of linear function of variables subject to linear inequalities. In T. Koopmans, editor, Activity Analysis of Production and Allocation, pages 339-347.
[14] C. Daskalakis, P. Goldberg, and C. Papadimitriou. The complexity of computing a Nash equilibrium. In STOC 2006.
[15] X. Deng, C. Papadimitriou, and S. Safra. On the complexity of price equilibria. Journal of Computer and System Sciences, 67(2):311-324, 2003.
[16] H. Edelsbrunner. Algorithms in Combinatorial Geometry. Springer-Verlag New York, Inc., 1987.
[17] K. Friedl, G. Ivanyos, M. Santha, and F. Verhoeven. On the black-box complexity of Sperner's lemma. In 15th FCT, pages 245-257, 2005.
[18] M. Hirsch, C. Papadimitriou, and S. Vavasis. Exponential lower bounds for finding Brouwer fixed points. Journal of Complexity, 5:379-416, 1989.
[19] L.-S. Huang and S.-H. Teng. On the approximation and smoothed complexity of Leontief market equilibria. ECCC, TR06-031, 2006.
[20] T. Iimura, K. Murota, and A. Tamura. Discrete fixed point theorem reconsidered. Journal of Mathematical Economics, 41:1030-1036, 2005.
[21] D. Johnson, C. Papadimtriou, and M. Yannakakis. How easy is local search? J. Comput. Syst. Sci., 37(1):79-100, 1988.
[22] N. Karmarkar. A new polynomial time algorithm for linear programming. Combinatorica, 4:373-395, 1984.
[23] V. Klee and G. Minty. How good is the simplex algorithm? In O. Shisha, editor, Inequalities - III, pages 159-175. Academic Press, 1972.
[24] C. Lemke and J. J.T. Howson. Equilibrium points of bimatrix games. J. Soc. Indust. Appl. Math., 12:413-423, 1964.
[25] J. Nash. Equilibrium point in n-person games. Porceedings of the National Academy of the USA, 36(1):48-49, 1950.
[26] Y. Nesterov and A. Nemirovskii. Interior Point Polynomial Algorithms in Convex Programming, volume 13 of Studies in Applied Mathematics. SIAM, 1993.
[27] J. Orlin, A. Punnen, and A. Schulz. Approximate local search in combinatorial optimization. In SODA 2004.
[28] C. Papadimitriou. On inefficient proofs of existence and complexity classes. In Proceedings of the 4th Czechoslovakian Symposium on Combinatorics, 1991.
[29] C. Papadimitriou. On the complexity of the parity argument and other inefficient proofs of existence. Journal of Computer and System Sciences, pages 498-532, 1994.
[30] M. Santha and M. Szegedy. Quantum and classical query complexities of local search are polynomially related. In STOC 2004, pages 494-501.
[31] R. Savani and B. von Stengel. Exponentially many steps for finding a nash equilibrium in a bimatrix game. In FOCS 2004, pages 258-267.
[32] H. Scarf. The approximation of fixed points of a continuous mapping. SIAM Journal on Applied Mathematics, 15:9971007, 1967.
[33] H. Scarf. On the computation of equilibrium prices. In W. Fellner, editor, Ten Economic Studies in the Tradition of Irving Fisher. New York: John Wiley \& Sons, 1967.
[34] H. Scarf. The Computation of Economic Equilibria. Yale University Press, 1973.
[35] E. Sperner. Neuer beweis fur die invarianz der dimensionszahl und des gebietes. Abhandlungen aus dem Mathematischen Seminar Universitat Hamburg, 6:265-272, 1928.
[36] D. Spielman and S.-H. Teng. Smoothed analysis of algorithms: Why the simplex algorithm usually takes polynomial time. Journal of ACM, 51(3):385-463, 2004.
[37] X. Sun and A.-C. Yao. On the quantum query complexity of local search in two and three dimensions. In FOCS 2006, pages 429-438.
[38] A.-C. Yao. Probabilistic computations: Towards a unified measure of complexity. In FOCS 1977, pages 222-227.
[39] S. Zhang. New upper and lower bounds for randomized and quantum local search. In STOC 2006, pages 634-643.

## A Inequalities

Proposition A.1. For all $\beta \geq 0,1-\beta \leq e^{-\beta}$.
Proposition A.2. For all $0 \leq \beta \leq 1 / 3,1-\beta \geq e^{-2 \beta}$.
Lemma A.3. For all $d \geq 1$ and $\beta \in\left[0,24^{-d}\right], \alpha_{d}(\beta) \leq$ $e^{2 \cdot 24^{d-1} \beta}$.

Proof. This lemma can be proved by induction on $d$.


[^0]:    *This research is supported mostly by the NSF ITR grant CCR0325630.
    ${ }^{\dagger}$ Part of this work was done while visiting Computer Science Department at Boston University. In part supported by the National Natural Science Foundation of China Grant 60553001, the National Basic Research Program of China Grant (2007CB807900, 2007CB807901), and the Chinese National Key Foundation R \& D Plan (2003CB317807, 2004CB318108).
    ${ }^{\ddagger}$ Part of this work was done while visiting Tsinghua University and Microsoft Research Asia Lab.

[^1]:    ${ }^{1}$ In linear programming, each local optimum is also a global optimum.
    ${ }^{2}$ Each path has a globally computable monotonic measure, the number of hops from the start of the path to a node.

[^2]:    ${ }^{3}$ The constant in $\Omega$ in our lower bound depends exponentially on $d$.

[^3]:    ${ }^{4}$ We assume that $\Upsilon$ does not stop until a zero point of $f$ is found.

[^4]:    ${ }^{5}$ Here the length of $S$ is unknown to the algorithm.
    ${ }^{6}$ The parameter $n$ is not the length of $S$, but the size of the alphabet.

