

Computing a Fixed Point of Contraction Maps in Polynomial Queries

Xi Chen^{*} Yuhao Li[†] Mihalis Yannakakis[‡]
Columbia University Columbia University Columbia University
xichen@cs.columbia.edu yuhaoli@cs.columbia.edu mihalis@cs.columbia.edu

Abstract

We give an algorithm for finding an ε -fixed point of a contraction map $f : [0, 1]^k \mapsto [0, 1]^k$ under the ℓ_∞ -norm with query complexity $O(k \log(1/\varepsilon))$.

^{*}Supported by NSF grants IIS-1838154, CCF-2106429 and CCF-2107187.

[†]Supported by NSF grants IIS-1838154, CCF-2106429 and CCF-2107187.

[‡]Supported by NSF grants CCF-2107187, CCF-2212233, and CCF-2332922.

1 Introduction

A map $f : \mathcal{M} \mapsto \mathcal{M}$ on a metric space (\mathcal{M}, d) is called a *contraction* map (or a $(1 - \gamma)$ -contraction map) if there exists $\gamma \in (0, 1]$ such that $d(f(x), f(y)) \leq (1 - \gamma) \cdot d(x, y)$ for all points $x, y \in \mathcal{M}$. In 1922, Banach [Ban22] proved a seminal fixed point theorem which states that every contraction map must have a unique fixed point, i.e., there is a unique $x \in \mathcal{M}$ that satisfies $f(x) = x$. Distinct from another renowned fixed point theorem by Brouwer, Banach’s theorem not only guarantees the uniqueness of the fixed point but also provides a method for finding it: iteratively applying the map f starting from any initial point will always converge to the unique fixed point. Over the past century, Banach’s fixed point theorem has found extensive applications in many fields. For example, in mathematics it can be used to prove theorems such as the Picard–Lindelöf (or Cauchy–Lipschitz) theorem on the existence and uniqueness of solutions to differential equations (see e.g. [CL55]), and the Nash embedding theorem [Nas56, Gün89]. In optimization and machine learning, it is used in the convergence and uniqueness analysis of value and policy iteration in Markov decision processes and reinforcement learning [Bel57, How60]. Indeed, as pointed out by Denardo [Den67], contraction mappings underlie many classical dynamic programming (DP) problems and sequential decision processes, including DP models of Bellman, Howard, Blackwell, Karlin and others.

A particularly important metric space to study the problem of finding a Banach’s fixed point is the k -cube $[0, 1]^k$ with respect to the ℓ_∞ -norm, since many important problems can be reduced to that of finding an ε -fixed point (i.e., $x \in [0, 1]^k$ satisfying $\|f(x) - x\|_\infty \leq \varepsilon$) in a $(1 - \gamma)$ -contraction map under the ℓ_∞ -norm. Such problems arise from a variety of fields including stochastic analysis, optimization, verification, semantics, and game theory. For example, the classical dynamic programming models mentioned above (Markov decision processes etc.) involve contraction maps under the ℓ_∞ -norm. Furthermore, the same holds for several well-known open problems that have been studied extensively and are currently not known to be in P. For instance, Condon’s simple stochastic games (SSGs) [Con92] can be reduced to the problem of finding an ε -fixed point in a $(1 - \gamma)$ -contraction map over $[0, 1]^k$ under the ℓ_∞ -norm. A similar reduction from [EY10] extends to an even broader class of games, namely, Shapley’s stochastic games [Sha53], which lay the foundation of multi-agent reinforcement learning [Lit94]. The same holds also of course for other problems known to be subsumed by SSGs, like parity games, which are important in verification (see e.g. [EJ91, CJK⁺22]), and mean payoff games [ZP96]. Crucially, in all these reductions, both the approximation parameter ε and the contraction parameter γ are inversely exponential in the input size. Therefore, efficient algorithms in this context are those with a complexity upper bound that is polynomial in $k, \log(1/\varepsilon)$ and $\log(1/\gamma)$.

In this paper we consider general algorithms that access the contraction map in a black-box manner (as an oracle), and study the *query complexity* of finding an ε -fixed point of a $(1 - \gamma)$ -contraction map over the k -cube $[0, 1]^k$ under the ℓ_∞ -norm (which we denote by $\text{CONTRACTION}_\infty(\varepsilon, \gamma, k)$). An algorithm under this model is given k, ε, γ , and oracle access to an unknown $(1 - \gamma)$ -contraction map f over $[0, 1]^k$. In each round the algorithm can send a point $x \in [0, 1]^k$ to the oracle to reveal its value $f(x)$. The goal of the algorithm is to find an ε -fixed point with as few queries as possible.

Prior work. Despite much ongoing interest on this problem (e.g., [EY10, DP11, DTZ18, FGMS20, Hol21, FGHS23]), progress in understanding the query complexity of $\text{CONTRACTION}_\infty(\varepsilon, \gamma, k)$ has been slow. Banach’s value iteration method needs $\Omega((1/\gamma) \log(1/\varepsilon))$ iterations to converge to an ε -fixed point. For the special case of $k = 2$, [SS02] obtained an $O(\log(1/\varepsilon))$ -query algorithm. Subsequently, [SS03] obtained an $O(\log^k(1/\varepsilon))$ -query algorithm for general k by applying a non-trivial recursive binary search procedure across all k dimensions. (Recently [FGMS20] obtained

similar upper bounds for all ℓ_p -norms with $2 < p < \infty$, though the complexity grows to infinity as $p \rightarrow \infty$.) Note, however, that all known upper bounds so far are exponential in either k or $\log(1/\gamma)$, and this is in sharp contrast with the ℓ_2 -norm case, for which [STW93, HKS99] gave an algorithm with both query and time complexity polynomial in $k, \log(1/\varepsilon)$ and $\log(1/\gamma)$.

Our contribution. We obtain the first algorithm for $\text{CONTRACTION}_\infty(\varepsilon, \gamma, k)$ with polynomial query complexity:

Theorem 1. *There is an $O(k \log(1/\varepsilon))$ -query algorithm for $\text{CONTRACTION}_\infty(\varepsilon, \gamma, k)$.*

The observation below explains why our upper bound does not depend on γ :

Observation 1. *Let $f : [0, 1]^k \mapsto [0, 1]^k$ be a $(1 - \gamma)$ -contraction map under the ℓ_∞ -norm. Consider the map $g : [0, 1]^k \mapsto [0, 1]^k$ defined as $g(x) := (1 - \varepsilon/2)f(x)$. Clearly g is a $(1 - \varepsilon/2)$ -contraction. Let x be any point with $\|g(x) - x\|_\infty \leq \varepsilon/2$. We have*

$$\varepsilon/2 \geq \|g(x) - x\|_\infty = \|(1 - \varepsilon/2)f(x) - x\|_\infty \geq \|f(x) - x\|_\infty - \varepsilon/2.$$

This gives a black-box reduction from $\text{CONTRACTION}_\infty(\varepsilon, \gamma, k)$ to $\text{CONTRACTION}_\infty(\varepsilon/2, \varepsilon/2, k)$, which is both query-efficient and time-efficient.

In Section 4, we give an $O(k \log(1/\varepsilon))$ -query algorithm for $\text{CONTRACTION}_\infty(\varepsilon/2, \varepsilon/2, k)$, from which Theorem 1 follows. Indeed, note that Observation 1 holds even if f is a *non-expansive* map (i.e., f has Lipschitz constant 1: $\|f(x) - f(y)\|_\infty \leq \|x - y\|_\infty$ for all $x, y \in [0, 1]^k$). As a result, the same query upper bound applies to $\text{NONEXP}_\infty(\varepsilon, k)$, the problem of finding an ε -fixed point in a non-expansive map over $[0, 1]^k$ under the ℓ_∞ -norm:

Corollary 1. *There is an $O(k \log(1/\varepsilon))$ -query algorithm for $\text{NONEXP}_\infty(\varepsilon, k)$.*

Another corollary of Theorem 1 is about finding a *strong* ε -fixed point in a contraction map f . We say x is a strong ε -fixed point of f if $\|x - x^*\|_\infty \leq \varepsilon$, where x^* is the unique fixed point of f . The following observation leads to Corollary 2, where $\text{STRCONTRACTION}_\infty(\varepsilon, \gamma, k)$ denotes the problem of finding a strong ε -fixed point:

Observation 2. *Let f be a $(1 - \gamma)$ -contraction map and x^* be its unique fixed point. Let x be any $(\varepsilon\gamma)$ -fixed point of f , i.e., x satisfies $\|f(x) - x\|_\infty \leq \varepsilon\gamma$. Then we have*

$$\|x - x^*\|_\infty \leq \|x - f(x)\|_\infty + \|f(x) - x^*\|_\infty \leq \varepsilon\gamma + (1 - \gamma)\|x - x^*\|_\infty,$$

which implies $\|x - x^\|_\infty \leq \varepsilon$. This gives a black-box reduction from $\text{STRCONTRACTION}_\infty(\varepsilon, \gamma, k)$ to $\text{CONTRACTION}_\infty(\varepsilon\gamma, \gamma, k)$, which is both query-efficient and time-efficient.*

Corollary 2. *There is an $O(k \log(1/(\varepsilon\gamma)))$ -query algorithm for $\text{STRCONTRACTION}_\infty(\varepsilon, \gamma, k)$.*

In sharp contrast with Corollary 2, however, we show that it is impossible to strongly approximate an exact fixed point in a non-expansive map over $[0, 1]^2$ under the ℓ_∞ norm.

Theorem 2. *There is no deterministic or randomized algorithm which, when given oracle access to any non-expansive map $f : [0, 1]^2 \mapsto [0, 1]^2$ under the ℓ_∞ -norm, computes in an expected bounded number of queries a point that is within distance $1/4$ of an exact fixed point of f .*

The impossibility of strong approximation for expansive maps under the ℓ_∞ norm, i.e. maps with Lipschitz constant $L > 1$, was shown previously by [SS03], and it was conjectured that the same fact should hold also for non-expansive maps ($L = 1$). For the ℓ_2 norm, impossibility of strong approximation for non-expansive maps was shown earlier [Sik01].

The problem $\text{CONTRACTION}_\infty(\varepsilon, \gamma, k)$ is a *promise problem*, i.e., it is promised that the function f in the black-box (the oracle) is a $(1 - \gamma)$ -contraction. In the various relevant applications (stochastic games etc.), the corresponding function that is induced is by construction a contraction, thus it is appropriate in these cases to restrict attention to functions that satisfy the contraction promise.

For any promise problem, one can define a corresponding total search problem, where the black-box can be any function f on the domain and the problem is to compute either a solution or a violation of the promise. In our case, the corresponding total search problem, denoted $\text{T-CONTRACTION}_\infty(\varepsilon, \gamma, k)$, is the problem of computing for a given function $f : [0, 1]^k \mapsto [0, 1]^k$ either an ε -fixed point or a violation of the contraction property, i.e. a pair of points $x, y \in [0, 1]^k$ such that $\|f(x) - f(y)\|_\infty > (1 - \gamma)\|x - y\|_\infty$. For any promise problem, the corresponding total search problem is clearly at least as hard as the promise problem. For some problems it can be strictly harder (and it may depend on the type of violation that is desired). However, we show that in our case the two versions have the same query complexity.

Theorem 3. *There is an $O(k \log(1/\varepsilon))$ -query algorithm for $\text{T-CONTRACTION}_\infty(\varepsilon, \gamma, k)$.*

Similar results hold for $\text{NONEXP}_\infty(\varepsilon, k)$ and $\text{STRCONTRACTION}_\infty(\varepsilon, \gamma, k)$: the total search versions have the same query complexity as the corresponding promise problems.

Remark. It is also important to note that while our algorithm in Theorem 1 is query-efficient, we do not currently see a way of making it time-efficient. The algorithm guarantees that within polynomial queries we can find an ε -fixed point, but each iteration requires a brute force procedure to *determine the next query point*. We will explain more details of techniques in Section 1.1.

Other Related Work. We have already mentioned the most relevant works addressing the query complexity of computing the fixed point of a contraction map. For continuous functions $f : [0, 1]^k \mapsto [0, 1]^k$ that have Lipschitz constant greater than 1 (i.e. are expansive), there are exponential lower bounds on the query complexity of computing a (weak) approximate fixed point [HPV89, CD08]. In a succinct black-box model, it is known that the query complexity of TFNP problems is polynomial in the description size of the object in the box [KNY19]. Their results do not apply to our setting. For instance, in the $\text{T-CONTRACTION}_\infty(\varepsilon, \gamma, k)$ problem, the objects are essentially all continuous functions over $[0, 1]^k$, whose description size far exceeds an exponential function of k .

$\text{CONTRACTION}_\infty(\varepsilon, \gamma, k)$ when considered in the white-box model¹ can be formulated as a total search problem so that it lies in the class TFNP. In fact, it is one of the motivating problems in [DP11] to define the class CLS for capturing problems that lie in both PLS [JPY88] and PPAD [Pap94]. Later, the problem of computing the exact fixed point of Piecewise-Linear contraction maps, an easier variant of $\text{CONTRACTION}_\infty(\varepsilon, \gamma, k)$, was placed in UEOPL [FGMS20], a subclass of CLS to capture problems with a unique solution. It is not known that $\text{CONTRACTION}_\infty(\varepsilon, \gamma, k)$ is

¹The white-box model refers to the model where the function is explicitly given by a polynomial-size circuit, in contrast to the black-box model where the function can only be accessed via an oracle as we studied in this paper. When we talk about a computational problem under the white-box model, we measure the efficiency by the time complexity.

complete for any TFNP class. Notably, to the best of our knowledge, for the known fixed point problems that are complete for some TFNP class, their query complexity in the black-box model is exponential. Examples of such well-known problems include PPAD-complete problems BROUWER and SPERNER [Pap94, CD09], PPA-complete problems BORSUK–ULAM, TUCKER [Pap94, ABB20] and MÖBIUSSPERNER [DEF⁺21], CLS-complete problems KKT [FGHS23] and METRICBANACH² [DTZ18], and UEOPPL-complete problem OPDC [FGMS20].

However, our results indicate that $\text{CONTRACTION}_\infty(\varepsilon, \gamma, k)$ is dramatically different from all these fixed point problems above in terms of query complexity. Indeed, it is a necessary step towards obtaining a polynomial-time general-purpose (i.e., black-box) algorithm for $\text{CONTRACTION}_\infty(\varepsilon, \gamma, k)$, which, ideally if true, would imply many breakthroughs in the fields of verification, semantics, learning theory, and game theory as we discussed before.

1.1 Sketch of the Main Algorithm

We give a high-level sketch of the main query algorithm for [Theorem 1](#). We start by discretizing the search space. Let $g : [0, n]^k \mapsto [0, n]^k$ with $g(x) := n \cdot f(x/n)$ and $n := \lceil 16/(\gamma\varepsilon) \rceil$. It is easy to show that g remains a $(1 - \gamma)$ -contraction over $[0, n]^k$ and it suffices to find a $(16/\gamma)$ -fixed point of g . Moreover, by rounding the unique fixed point x^* of g to an integer point, we know trivially that there exists at least one point x in $\text{EVEN}(n, k)$ ³, which is the set of integral points in the grid $\{0, 1, \dots, n\}^k$ with all coordinates being even, that satisfies $\|x - x^*\|_\infty \leq 1$. Furthermore, it is easy to show that any such point x must be a $(16/\gamma)$ -fixed point. So our goal is to find a point $x \in \text{EVEN}(n, k)$ that satisfies $\|x - x^*\|_\infty \leq 1$ query-efficiently.

To this end, we use Cand^t to denote the set of $\text{EVEN}(n, k)$ that remains possible to be close to the unknown exact fixed point x^* of g . Starting with Cand^0 set to be the full set $\text{EVEN}(n, k)$, the success of the algorithm relies on whether we can cut down the size of Cand^t efficiently. For this purpose we prove a number of geometric lemmas in [Section 3](#) to give a characterization of the exact fixed point x^* , which lead to the following primitive used by the algorithm repeatedly:

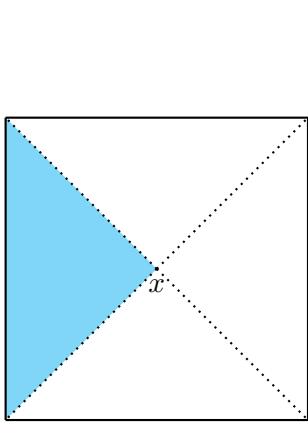
Given $x \in [0, n]^k$, $i \in [k]$ and $\phi \in \{\pm 1\}$, we write $\mathcal{P}_i(x, \phi)$ to denote the set of points $y \in [0, n]^k$ such that $\phi \cdot (y_i - x_i) = \|y - x\|_\infty$, where \mathcal{P} is a shorthand for pyramid; see [Figure 1](#) for illustrations. Then after querying a point $a \in [0, n]^k$, either a was found to be a $(16/\gamma)$ -fixed point (in which case the algorithm is trivially done), or one can find $\phi_i \in \{\pm 1\}$ for each $i \in [k]$ such that no point in $\mathcal{P}_i(a, \phi_i)$ can be close (within ℓ_∞ -distance 1) to x^* (in which case we can update Cand^t by removing all points in $\bigcup_{i \in [k]} \mathcal{P}_i(a, \phi_i)$).

Given this, it suffices to show that for any set of points $T \subseteq \text{EVEN}(n, k)$ (as Cand^t), there exists a point a to be queried such that for any $\phi_i \in \{\pm 1\}$:

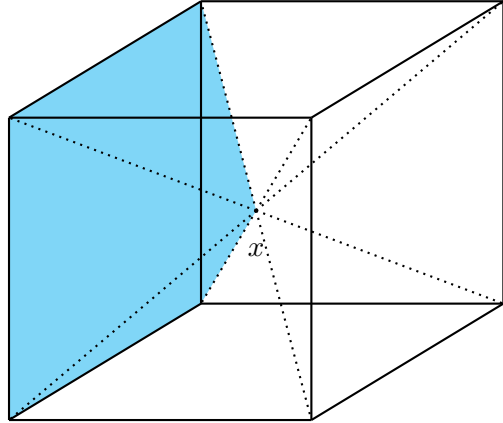
$$\left| T \cap \left(\bigcup_{i \in [k]} \mathcal{P}_i(a, \phi_i) \right) \right|$$

²METRICBANACH refers to the problem of computing an approximate fixed point of a contraction map where the distance function d is also part of the input.

³Looking ahead, the usage of $\text{EVEN}(n, k)$ will guarantee that all queried points lie in the normal grid $\{0, \dots, n\}^k$, which implies that the approximate fixed point that we will find lies in the normal grid $\{0, \dots, n\}^k$ as well.



(a) The light blue region is given by $\mathcal{P}_1(x, -1)$ where x is the center point.



(b) The light blue region is given by $\mathcal{P}_1(x, -1)$ where x is the center point.

Figure 1: Pyramids in two-dimensional and three-dimensional cubes.

is large relative to $|T|$. Quantitatively, we show that there exists a point a such that for any $\phi_i \in \{\pm 1\}$, the set $T \cap \left(\bigcup_{i \in [k]} \mathcal{P}_i(a, \phi_i) \right)$ has size at least half of $|T|$. This bound is clearly sharp and we refer to such point a as *balanced point*.

To prove the existence of a balanced point, we construct an infinite sequence of continuous maps $\{f^t\}$ that can be viewed as relaxed versions of the search for a balanced point. Using Brouwer's fixed point theorem, every map f^t has a fixed point p^t and thus, by the Bolzano–Weierstrass theorem, there must be an infinite subsequence of $\{p^t\}$ that converges. Letting $p^* \in [0, n]^k$ be the point it converges to, we can show that p^* must be a desired balanced point. We could further round p^* to $q^* \in \{0, \dots, n\}^k$ and show that the latter is a balanced point in the grid. While we show such a point always exists, the brute-force search to find $q^* \in \{0, \dots, n\}^k$ is the reason why our algorithm is not time-efficient.

2 Preliminaries

Definition 1 (Contraction). *Let $0 < \gamma < 1$ and (\mathcal{M}, d) be a metric space. A map $f : \mathcal{M} \mapsto \mathcal{M}$ is a $(1 - \gamma)$ -contraction map with respect to (\mathcal{M}, d) if $d(f(x), f(y)) \leq (1 - \gamma) \cdot d(x, y)$ for all $x, y \in \mathcal{M}$.*

A map $f : \mathcal{M} \mapsto \mathcal{M}$ is said to be non-expansive if $d(f(x), f(y)) \leq d(x, y)$ for all $x, y \in \mathcal{M}$.

Every contraction map has a unique fixed point, i.e., x^* with $f(x^*) = x^*$ by Banach's theorem. Every non-expansive map over $[0, 1]^k$ with respect to the infinity norm has a (not necessarily unique) fixed point by Brouwer's theorem. In this paper, we study the query complexity of finding an ε -fixed point of a $(1 - \gamma)$ -contraction map f over the k -cube $[0, 1]^k$ with respect to the infinity norm:

Definition 2 ($\text{CONTRACTION}_\infty(\varepsilon, \gamma, k)$). *In problem $\text{CONTRACTION}_\infty(\varepsilon, \gamma, k)$, we are given oracle access to a $(1 - \gamma)$ -contraction map f over $[0, 1]^k$ with respect to the infinity norm, i.e., f satisfies*

$$\|f(x) - f(y)\|_\infty \leq (1 - \gamma) \cdot \|x - y\|_\infty, \quad \text{for all } x, y \in [0, 1]^k$$

and the goal is to find an ε -fixed point of f , i.e., a point $x \in [0, 1]^k$ such that $\|f(x) - x\|_\infty \leq \varepsilon$.

We also write $\text{STRCONTRACTION}_\infty(\varepsilon, \gamma, k)$ to denote the problem with the same input but the goal is to find a strong ε -fixed point of f , i.e., $x \in [0, 1]^k$ such that $\|x - x^*\|_\infty \leq \varepsilon$, where x^* is the unique fixed point of f .

We define similar problems for non-expansive maps over $[0, 1]^k$:

Definition 3 ($\text{NONEXP}_\infty(\varepsilon, k)$). In problem $\text{NONEXP}_\infty(\varepsilon, k)$, we are given oracle access to a non-expansive map $f : [0, 1]^k \rightarrow [0, 1]^k$ with respect to the infinity norm, i.e., f satisfies

$$\|f(x) - f(y)\|_\infty \leq \|x - y\|_\infty, \quad \text{for all } x, y \in [0, 1]^k$$

and the goal is to find an ε -fixed point of f .

We write $\text{STRNONEXP}_\infty(\varepsilon, k)$ to denote the problem with the same input but the goal is to find a strong ε -fixed point of f , namely, a point x such that $\|x - x^*\|_\infty \leq \varepsilon$, where x^* is any fixed point of f .

Let $f : [0, 1]^k \rightarrow [0, 1]^k$ be a $(1 - \gamma)$ -contraction map. For convenience, our main algorithm for $\text{CONTRACTION}_\infty(\varepsilon, \gamma, k)$ will work on $g : [0, n]^k \mapsto [0, n]^k$ with

$$n := \left\lceil \frac{16}{\gamma\varepsilon} \right\rceil \quad \text{and} \quad g(x) := n \cdot f(x/n), \quad \text{for all } x \in [0, n]^k.$$

We record the following simple property about g :

Lemma 1. The map g constructed above is still a $(1 - \gamma)$ -contraction and finding an ε -fixed point of f reduces to finding a $(16/\gamma)$ -fixed point of g .

Proof. For any two points $x, y \in [0, n]^k$, we have

$$\|g(x) - g(y)\|_\infty = \|n \cdot f(x/n) - n \cdot f(y/n)\|_\infty \leq n(1 - \gamma) \cdot \|x/n - y/n\|_\infty = (1 - \gamma)\|x - y\|_\infty.$$

Suppose we found a $(16/\gamma)$ -fixed point a of g . We show that $x := a/n$ is an ε -fixed point of f :

$$\|f(x) - x\|_\infty = \|g(a)/n - a/n\|_\infty = \frac{1}{n} \cdot \|g(a) - a\|_\infty \leq \frac{1}{n} \cdot \frac{16}{\gamma} \leq \varepsilon.$$

This finishes the proof of the lemma. \square

Notation. Given a positive integer m , we use $[m]$ to denote $\{1, \dots, m\}$ and $[0 : m]$ to denote $\{0, \dots, m\}$. For a real number $t \in \mathbb{R}$, we let $\text{sgn}(t) = 1$ if $t > 0$, $\text{sgn}(t) = -1$ if $t < 0$, and $\text{sgn}(t) = 0$ if $t = 0$. Given positive integers n and k , we use $\text{EVEN}(n, k)$ to denote the set of all integer points $x \in [0, n]^k$ such that x_i is even for all $i \in [k]$.

For a point $x \in \mathbb{R}^k$ (not necessarily in $[0, n]^k$), a coordinate $i \in [k]$, and a sign $\phi \in \{\pm 1\}$, we use $\mathcal{P}_i(x, \phi)$ to denote

$$\mathcal{P}_i(x, \phi) := \left\{ y \in [0, n]^k : \phi \cdot (y_i - x_i) = \|y - x\|_\infty \right\},$$

where \mathcal{P} is a shorthand for pyramid.

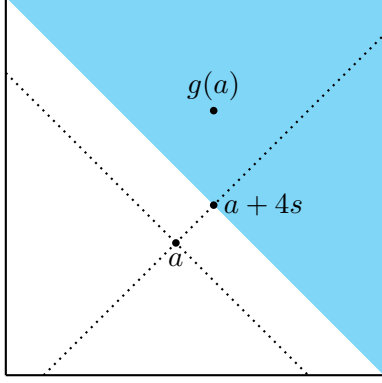
Given a $(1 - \gamma)$ -contraction map g over $[0, n]^k$, we use $\text{Fix}(g)$ to denote the unique fixed point of g . For any point $x \in [0, n]^k$, we use $\text{Around}(x)$ to denote the set

$$\text{Around}(x) := \{y \in [0, n]^k : \|x - y\|_\infty \leq 1\}.$$

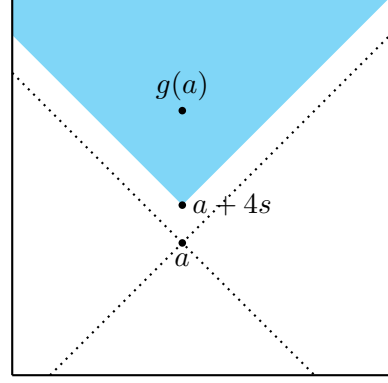
We note that $\|g(y) - y\|_\infty \leq 2$ for all $y \in \text{Around}(\text{Fix}(g))$ and thus, any point in $\text{Around}(\text{Fix}(g))$ is a desired $(16/\gamma)$ -fixed point (given that $\gamma < 1$). To see this, letting $x = \text{Fix}(g)$, we have

$$\|g(y) - y\|_\infty \leq \|g(y) - g(x)\|_\infty + \|x - y\|_\infty \leq (2 - \gamma)\|x - y\|_\infty \leq 2, \quad \text{for any } y \in \text{Around}(x).$$

Note that for any $x \in [0, n]^k$, $\text{Around}(x) \cap \text{EVEN}(n, k)$ is non-empty.



(a) a and $g(a)$ satisfy $s = (+1, +1)$.



(b) a and $g(a)$ satisfy $s = (0, +1)$.

Figure 2: An illustration of [Lemma 2](#). The light blue region is given by $\bigcup_{i:s_i \neq 0} \mathcal{P}_i(a + 4s, s_i)$, which contains $\text{Fix}(g)$ (the unique fixed point of g).

3 Characterizing the Unique Fixed Point

In this section, we prove three geometrical lemmas for characterizing the exact fixed point $\text{Fix}(g)$. The corresponding geometrical illustrations are in [Figure 2](#), [Figure 3](#), and [Figure 4](#).

Lemma 2. *Let $g : [0, n]^k \mapsto [0, n]^k$ be a $(1 - \gamma)$ -contraction map and let $a \in [0, n]^k$ be a point such that $\|g(a) - a\|_\infty > 16/\gamma$ and $s \in \{\pm 1, 0\}^k$ be the sign vector such that $s_i = \text{sgn}(g(a)_i - a_i)$. Then*

$$\text{Fix}(g) \in \bigcup_{i:s_i \neq 0} \mathcal{P}_i(a + 4s, s_i).$$

Proof. Let $c = a + 4s$ and $x^* = \text{Fix}(g)$ be the unique fixed point.

First, we show that $\|x^* - a\|_\infty > 8/\gamma$. Otherwise, we have

$$\|g(a) - a\|_\infty \leq \|g(a) - g(x^*)\|_\infty + \|x^* - a\|_\infty \leq (1 - \gamma + 1) \cdot \|x^* - a\|_\infty \leq \frac{16}{\gamma},$$

which contradicts the assumption that $\|g(a) - a\|_\infty > 16/\gamma$. Now it suffices to show that $g(x) \neq x$ for any point $x \in [0, n]^k$ that satisfies both

$$x \notin \bigcup_{i,s_i \neq 0} \mathcal{P}_i(c, s_i) \quad \text{and} \quad \|x - a\|_\infty > \frac{8}{\gamma}.$$

Let j be a dominating coordinate between x and c , i.e., a $j \in [k]$ such that $|x_j - c_j| = \|x - c\|_\infty$. We divide the proof into two parts.

Part 1: $s_j = 0$. Thus $g(a)_j = a_j$ and $c_j = a_j$. Assume without loss of generality that $x_j \geq c_j$; the case when $x_j \leq c_j$ is symmetric. On the one hand, we have $x_j - c_j = \|x - c\|_\infty$, which gives

$$x_j = c_j + \|x - c\|_\infty. \tag{*}$$

On the other hand, using that g is a $(1 - \gamma)$ -contraction and $g(a)_j = a_j$, we have

$$g(x)_j - a_j = g(x)_j - g(a)_j \leq (1 - \gamma) \cdot \|x - a\|_\infty.$$

Equivalently, $g(x)_j \leq a_j + \|x - a\|_\infty - \gamma\|x - a\|_\infty$. Combining with $\|x - a\|_\infty > 8/\gamma$, this implies

$$g(x)_j < a_j + \|x - a\|_\infty - 8. \quad (\diamond)$$

Putting (\star) and (\diamond) together and the facts that $\|c - a\|_\infty = 4$ and $c_j = a_j$, we have

$$x_j - g(x)_j > \|x - c\|_\infty - \|x - a\|_\infty + 8 \geq -\|c - a\|_\infty + 8 > 0,$$

which implies that $g(x)_j \neq x_j$ and $g(x) \neq x$.

Part 2: $s_j \neq 0$. Assume without loss of generality that $s_j = +1$; the case $s_j = -1$ is symmetric.

Since we are considering points not in $\bigcup_{i:s_i \neq 0} \mathcal{P}_i(c, s_i)$, it must be the case that $x \in \mathcal{P}_j(c, -s_j)$ and thus, $x_j \leq c_j$. Since $\|x - a\|_\infty > 8/\gamma$, we have $x_j \leq a_j$; otherwise $a_j \leq x_j \leq c_j$ and $\|x - c\|_\infty \leq c_j - x_j \leq c_j - a_j \leq 4$ and thus,

$$\|x - a\|_\infty \leq \|x - c\|_\infty + \|a - c\|_\infty \leq 8.$$

Given the $(1 - \gamma)$ -contraction of g , we have $g(a)_j - g(x)_j \leq (1 - \gamma) \cdot \|x - a\|_\infty$, which implies

$$g(x)_j \geq g(a)_j - \|x - a\|_\infty + \gamma\|x - a\|_\infty > g(a)_j - \|x - a\|_\infty + 8. \quad (1)$$

Next, we show an upper bound on $\|x - a\|_\infty$. Recall that $x \in \mathcal{P}_j(c, -s_j)$. Consider $y = x - 4s$. We have $y \in \mathcal{P}_j(a, -s_j)$. So

$$\|x - a\|_\infty \leq \|x - y\|_\infty + \|y - a\|_\infty = 4 + (a_j - y_j) = 4 + (a_j - x_j + 4) = 8 + (a_j - x_j).$$

Given this and plugging the upper bound in [Equation \(1\)](#), we will get

$$g(x)_j > g(a)_j + x_j - a_j.$$

Recall that $s_j = +1$ implies that $g(a)_j > a_j$. So we have $g(x)_j > x_j$ and $g(x) \neq x$. This finishes the proof of [Lemma 2](#). \square

Lemma 3. Let $b \in \mathbb{R}^k$ and $s \in \{\pm 1, 0\}^k$ such that $s \neq 0^k$. Then every $x \in \bigcup_{i:s_i \neq 0} \mathcal{P}_i(b + 2s, s_i)$ must have $\text{Around}(x) \subseteq \bigcup_{i:s_i \neq 0} \mathcal{P}_i(b, s_i)$.

Proof. Let i^* be such that $s_{i^*} \neq 0$ and $x \in \mathcal{P}_{i^*}(b + 2s, s_{i^*})$. Assume without loss of generality that $s_{i^*} = +1$. We have $x_{i^*} - (b_{i^*} + 2) \geq |x_i - (b_i + 2s_i)|$ for every $i \in [k]$.

Fix an arbitrary $y \in \text{Around}(x)$. We must have $y_{i^*} - b_{i^*} \geq 0$, which follows from

$$0 \leq x_{i^*} - (b_{i^*} + 2) \leq y_{i^*} + 1 - (b_{i^*} + 2)$$

and thus, $y_{i^*} \geq b_{i^*} + 1$. Let

$$j \in \arg \max_{i \in [k]} \{s_i(y_i - b_i) \mid s_i \neq 0\}.$$

Since $s_{i^*} = +1$, we have

$$s_j(y_j - b_j) \geq y_{i^*} - b_{i^*} \geq 0,$$

thus $s_j(y_j - b_j) = |y_j - b_j|$. Our goal is to show $y \in \mathcal{P}_j(b, s_j)$ and it suffices for us to show $|y_j - b_j| \geq |y_i - b_i|$ for all $i \in [k]$.

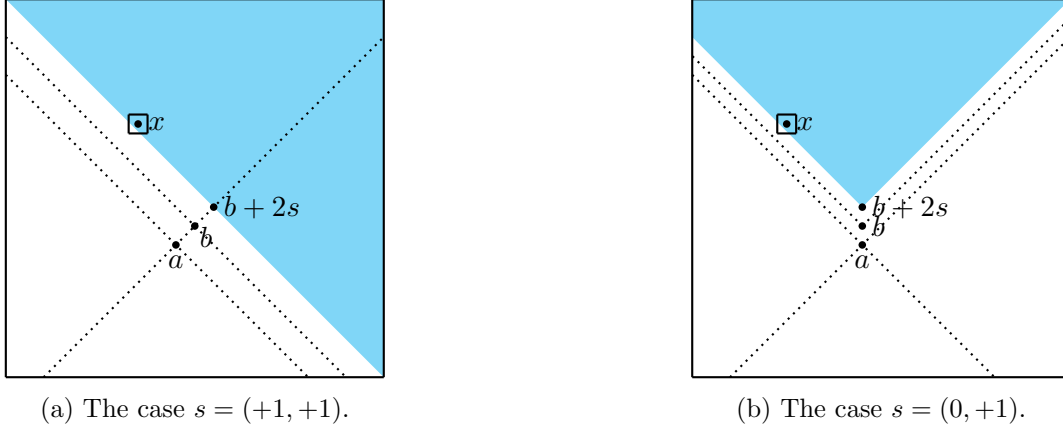


Figure 3: An illustration of **Lemma 3**. The light blue region is given by $\bigcup_{i:s_i \neq 0} \mathcal{P}_i(b+2s, s_i)$, which contains x . $\text{Around}(x)$ is the small square around x , which is contained in $\bigcup_{i:s_i \neq 0} \mathcal{P}_i(b, s_i)$.

Let's consider first an arbitrary $i \in [k]$ with $s_i = 0$. Recall that $x_{i^*} - (b_{i^*} + 2) \geq |x_i - b_i|$. In particular, this implies $|x_{i^*} - b_{i^*}| - 2 \geq |x_i - b_i|$. Putting everything together, we have

$$|y_j - b_j| \geq |y_{i^*} - b_{i^*}| \geq |x_{i^*} - b_{i^*}| - 1 \geq |x_i - b_i| + 1 \geq |y_i - b_i|.$$

Finally let's consider an $i \in [k]$ such that $s_i \neq 0$. If $s_i(y_i - b_i) > 0$, by the definition of how we picked j , we have $|y_j - b_j| \geq |y_i - b_i|$. If $s_i(y_i - b_i) < 0$, then we have

$$|y_i - b_i| = |y_i - (b_i + 2s_i)| - 2 \leq |x_i - (b_i + 2s_i)| - 1 \leq x_{i^*} - b_{i^*} - 3 \leq y_{i^*} - b_{i^*} - 2 < |y_j - b_j|.$$

This finishes the proof of the lemma. \square

Lemma 4. Let $a \in [0, n]^k$ and $s \in \{\pm 1, 0\}^k$ such that $s \neq 0^k$. Then for every $j \in [k]$, there exists $\phi \in \{\pm 1\}$ such that

$$\mathcal{P}_j(a, \phi) \cap \left(\bigcup_{i:s_i \neq 0} \mathcal{P}_i(a + 2s, s_i) \right) = \emptyset.$$

Proof. First we note that

$$\left(\bigcup_{i:s_i \neq 0} \mathcal{P}_i(a, -s_i) \right) \cap \left(\bigcup_{i:s_i \neq 0} \mathcal{P}_i(a + 2s, s_i) \right) = \emptyset.$$

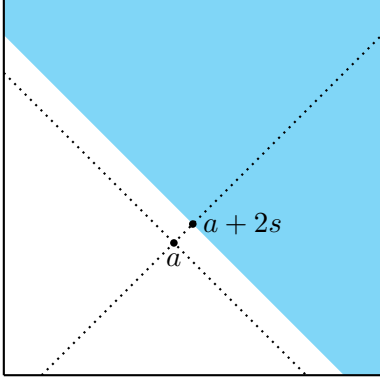
This implies that

$$\mathcal{P}_j(a, \phi) \cap \left(\bigcup_{i:s_i \neq 0} \mathcal{P}_i(a + 2s, s_i) \right) = \emptyset$$

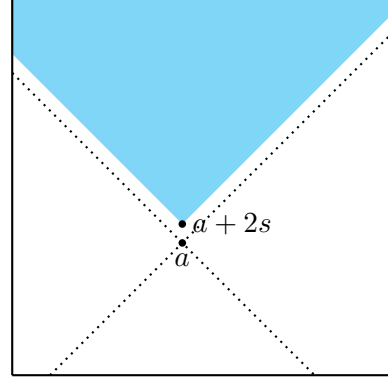
for all j with $s_j \neq 0$ by setting $\phi = -s_j$.

Now consider a j with $s_j = 0$. Under this case, we show in fact that

$$\mathcal{P}_j(a, \phi) \cap \left(\bigcup_{i:s_i \neq 0} \mathcal{P}_i(a + 2s, s_i) \right) = \emptyset$$



(a) The case $s = (+1, +1)$. The light blue region is given by $\bigcup_{i:s_i \neq 0} \mathcal{P}_i(a + 2s, s_i)$. Clearly it does not intersect with $\mathcal{P}_1(a, -1)$ or $\mathcal{P}_2(a, -1)$.



(b) The case $s = (0, +1)$. The light blue region is given by $\bigcup_{i:s_i \neq 0} \mathcal{P}_i(a + 2s, s_i)$. Clearly it does not intersect with $\mathcal{P}_2(a, -1)$, $\mathcal{P}_1(a, -1)$, or $\mathcal{P}_1(a, 1)$.

Figure 4: An illustration of [Lemma 4](#). When $s_i \neq 0$, we can prove that $\bigcup_{i:s_i \neq 0} \mathcal{P}_i(a + 2s, s_i)$ does not intersect with $\mathcal{P}_i(a, -s_i)$; see left figure above. When $s_i = 0$, we can prove $\bigcup_{i:s_i \neq 0} \mathcal{P}_i(a + 2s, s_i)$ does not intersect with $\mathcal{P}_i(a, -1)$ or $\mathcal{P}_i(a, 1)$; see right figure above.

Algorithm 1 Query Algorithm for $\text{CONTRACTION}_\infty(\varepsilon, \gamma, k)$

- 1: Let $\text{Cand}^0 \leftarrow \text{EVEN}(n, k)$.
- 2: **for** $t = 1, 2, \dots$ **do**
- 3: Find and **query** an $a^t \in [0 : n]^k$ such that a^t is a balanced point of Cand^{t-1} .
- 4: **if** $\|g(a^t) - a^t\|_\infty \leq 16/\gamma$ **then return** a^t as a $(16/\gamma)$ -fixed point of g
- 5: Let $s \in \{\pm 1, 0\}^k$ be such that $s_i = \text{sgn}(g(a^t)_i - a_i^t)$ for all $i \in [k]$, $b^t \leftarrow a^t + 2s$, and

$$\text{Cand}^t \leftarrow \text{Cand}^{t-1} \cap \left(\bigcup_{i:s_i \neq 0} \mathcal{P}_i(b^t, s_i) \right). \quad (2)$$

for both $\phi \in \{\pm 1\}$. Consider any point $x \in \bigcup_{i:s_i \neq 0} \mathcal{P}_i(a + 2s, s_i)$ and we show that $x \notin \mathcal{P}_j(a, -1)$ and $x \notin \mathcal{P}_j(a, +1)$. Let $b = a + 2s$. As $x \in \bigcup_{i:s_i \neq 0} \mathcal{P}_i(b, s_i)$, there exists i^* with $s_{i^*} \neq 0$ such that

$$s_{i^*}(x_{i^*} - b_{i^*}) = \|x - b\|_\infty \geq |x_j - b_j|.$$

Note also that $|x_{i^*} - a_{i^*}| = s_{i^*}(x_{i^*} - b_{i^*} + 2s_{i^*}) = s_{i^*}(x_{i^*} - b_{i^*}) + 2$ and $b_j = a_j$, so we have

$$|x_{i^*} - a_{i^*}| = \|x - b\|_\infty + 2 > |x_j - b_j| = |x_j - a_j|.$$

Thus $x \notin \mathcal{P}_j(a, -1)$ and $x \notin \mathcal{P}_j(a, +1)$. This finishes the proof of the lemma. \square

4 The Algorithm

We prove [Theorem 1](#) in this section. Our algorithm for $\text{CONTRACTION}_\infty(\varepsilon, \gamma, k)$ is described in [Algorithm 1](#). Given oracle access to a $(1 - \gamma)$ -contraction map $g : [0, n]^k \rightarrow [0, n]^k$, we show that it can find a $(16/\gamma)$ -fixed point of g within $O(k \log(n))$ many queries. This is sufficient given [Lemma 1](#).

The analysis of [Algorithm 1](#) uses the following theorem which we prove in the next section. In particular, it guarantees the existence of the point a^t to be queried in round t that satisfies that for any $s \in \{\pm 1\}^k$, we have $\left(\bigcup_{i \in [k]} \mathcal{P}_i(a^t, s_i)\right)$ intersects with at least half of \mathbf{Cand}^{t-1} . We will call a point q^* with the property stated in [Theorem 4](#) below a *balanced* point for T .

Theorem 4. *For any $T \subseteq \mathbf{EVEN}(n, k)$, there exist a point $q^* \in [0 : n]^k$ such that*

$$\left| T \cap \left(\bigcup_{i \in [k]} \mathcal{P}_i(q^*, s_i) \right) \right| \geq \frac{1}{2} \cdot |T|, \quad \text{for all } s \in \{\pm 1\}^k.$$

At a high level, [Algorithm 1](#) maintains a subset of grid points $\mathbf{EVEN}(n, k)$ as candidate solutions, which is denoted by \mathbf{Cand}^t after round t . We show the following invariants:

Lemma 5. *For every round $t \geq 1$, either the point a^t queried is a $(16/\gamma)$ -fixed point of g (and the algorithm terminates), or we have both $(\mathbf{Around}(\mathbf{Fix}(g)) \cap \mathbf{EVEN}(n, k)) \subseteq \mathbf{Cand}^t$ and*

$$|\mathbf{Cand}^t| \leq \frac{1}{2} \cdot |\mathbf{Cand}^{t-1}|. \quad (3)$$

Proof of Lemma 5. We start with the proof of (3). Suppose that a^t is a balanced point of \mathbf{Cand}^{t-1} . Then by [Lemma 4](#), for any coordinate $j \in [k]$, there must exist a sign $\phi_j \in \{\pm 1\}$ such that

$$\mathcal{P}_j(a^t, \phi_j) \cap \left(\bigcup_{i: s_i \neq 0} \mathcal{P}_i(b^t, s_i) \right) = \emptyset.$$

Since $\mathbf{Cand}^t := \mathbf{Cand}^{t-1} \cap \left(\bigcup_{i: s_i \neq 0} \mathcal{P}_i(b^t, s_i) \right)$, we know that for every coordinate $j \in [k]$, there exists $\phi_j \in \{\pm 1\}$ such that $\mathbf{Cand}^t \cap \mathcal{P}_j(a^t, \phi_j) = \emptyset$. The inequality (3) follows directly from the definition of the balanced point.

Next we prove by induction that $(\mathbf{Around}(\mathbf{Fix}(g)) \cap \mathbf{EVEN}(n, k)) \subseteq \mathbf{Cand}^t$ for every t before the round that the algorithm terminates. The basis is trivial given that \mathbf{Cand}^0 is set to be $\mathbf{EVEN}(n, k)$. For round $t \geq 1$, we assume that $\|g(a^t) - a^t\|_\infty > 16/\gamma$; otherwise a solution is found and the algorithm terminates.

Let $b^t = a^t + 2s$ and $c^t = b^t + 2s$. By [Lemma 2](#), we know that

$$\mathbf{Fix}(g) \in \left(\bigcup_{i: s_i \neq 0} \mathcal{P}_i(c^t, s_i) \right).$$

Since b^t is defined as $c^t - 2s$, by [Lemma 3](#), we know that

$$\mathbf{Around}(\mathbf{Fix}(g)) \subseteq \bigcup_{i: s_i \neq 0} \mathcal{P}_i(b^t, s_i).$$

We finish the proof by using the inductive hypothesis $(\mathbf{Around}(\mathbf{Fix}(g)) \cap \mathbf{EVEN}(n, k)) \subseteq \mathbf{Cand}^{t-1}$ and (2). \square

Recall that $\mathbf{Around}(\mathbf{Fix}(g)) \cap \mathbf{EVEN}(n, k)$ is non-empty. Given [Lemma 5](#) and that $|\mathbf{Cand}^0| \leq n^k$, within at most

$$O\left(\log(n^k)\right) = O\left(k \log\left(\frac{1}{\varepsilon\gamma}\right)\right)$$

many rounds, one of the points a^t queried by the algorithm must be a $(16/\gamma)$ -fixed point of g . This finishes the proof of [Theorem 1](#).

5 Existence of Balanced Point

We prove [Theorem 4](#) in this section, which we restate below:

Theorem 4. *For any $T \subseteq \text{EVEN}(n, k)$, there exist a point $q^* \in [0 : n]^k$ such that*

$$\left| T \cap \left(\bigcup_{i \in [k]} \mathcal{P}_i(q^*, s_i) \right) \right| \geq \frac{1}{2} \cdot |T|, \quad \text{for all } s \in \{\pm 1\}^k.$$

Proof. For each positive integer $t \geq 4$ we let

$$S^t := \bigcup_{x \in T} B(x, 1/t) \subset [-1/4, n + 1/4]^k,$$

where $B(x, 1/t)$ denotes the ℓ_2 -ball of radius $1/t$ centered at x . We write $\text{vol}(S^t)$ to denote the volume of S^t and $\text{vol}(S^t \cap \mathcal{P})$ to denote the volume of the intersection of S^t and some pyramid \mathcal{P} .

We apply Brouwer's fixed point to prove the existence of a balanced (real) point for the balls:

Lemma 6. *For every integer $t \geq 4$, there exist $p^* \in [-1/4, n + 1/4]^k$ such that $\text{vol}(\mathcal{P}_i(p^*, +1) \cap S^t) = \text{vol}(\mathcal{P}_i(p^*, -1) \cap S^t)$ holds for all coordinates $i \in [k]$.*

Proof. We define a continuous map $f : [-1/4, n + 1/4]^k \mapsto [-1/4, n + 1/4]^k$ and apply Brouwer's fixed point theorem on f to find a fixed point p^* of f , and show that p^* satisfies the property above.

We define f as follows: For every $p \in [-1/4, n + 1/4]^k$ and $i \in [k]$, let

$$f_i(p) := p_i + \frac{\text{vol}(\mathcal{P}_i(p, +1) \cap S^t) - \text{vol}(\mathcal{P}_i(p, -1) \cap S^t)}{(n + 0.5)^{k-1}}.$$

It is clear that f is continuous. To see that it is from $[-1/4, n + 1/4]^k$ to itself, we note that

$$0 \leq \frac{\text{vol}(\mathcal{P}_i(p, +1) \cap S^t)}{(n + 0.5)^{k-1}} \leq n + \frac{1}{4} - p_i \quad \text{and} \quad 0 \leq \frac{\text{vol}(\mathcal{P}_i(p, -1) \cap S^t)}{(n + 0.5)^{k-1}} \leq p_i + \frac{1}{4},$$

where the upper bounds above simply follow from the fact that

$$\text{vol}(\mathcal{P}_i(p, +1)) \leq \left(n + \frac{1}{4} - p_i \right) \cdot (n + 0.5)^{k-1} \quad \text{and} \quad \text{vol}(\mathcal{P}_i(p, -1)) \leq \left(p_i + \frac{1}{4} \right) \cdot (n + 0.5)^{k-1}.$$

As a result, one can apply Brouwer's fixed point theorem on f to conclude that there exists a point $p^* \in [-1/4, n + 1/4]^k$ such that $f(p^*) = p^*$, which implies that

$$\text{vol}(\mathcal{P}_i(p^*, +1) \cap S^t) = \text{vol}(\mathcal{P}_i(p^*, -1) \cap S^t)$$

for all $i \in [k]$. This finishes the proof of the lemma. \square

By [Lemma 6](#), we have that for every t , there exists $p^t \in [-1/4, n + 1/4]^k$ such that

$$\text{vol}(\mathcal{P}_i(p^t, +1) \cap S^t) = \text{vol}(\mathcal{P}_i(p^t, -1) \cap S^t)$$

for all $i \in [k]$. Given that $[-1/4, n + 1/4]^k$ is compact, $\{p^t\}_{t \geq 1}$ has an infinite subsequence that converges to a point $p^* \in [-1/4, n + 1/4]^k$. We refer to the subsequence as $\{p^{t_\ell}\}_{\ell \geq 1}$.

In [Lemma 7](#), we show that p^* is already a desired balanced point but it may be off-grid. In [Lemma 8](#), we show how to round p^* to $q^* \in [0 : n]^k$ while making sure that

$$\mathcal{P}_i(p^*, +1) \cap T \subseteq \mathcal{P}_{i^*}(q^*, +1) \cap T \quad \text{and} \quad \mathcal{P}_{i^*}(p^*, -1) \cap T \subseteq \mathcal{P}_{i^*}(q^*, -1) \cap T.$$

Our goal then follows by combining these two lemmas.

Lemma 7. *We have*

$$\left| \left(\bigcup_{i \in [k]} \mathcal{P}_i(p^*, s_i) \right) \cap T \right| \geq \frac{1}{2} \cdot |T|, \quad \text{for all } s \in \{\pm 1\}^k.$$

Proof. We write A to denote the following (potentially empty) set of positive real numbers defined using p^* : $a \in (0, 1)$ is in A if there are $i \neq j \in [k]$ such that either

1. $p_i^* + p_j^*$ is an integer plus a ; or
2. $p_i^* + p_j^*$ is an integer minus a ; or
3. $p_i^* - p_j^*$ is an integer plus a ; or
4. $p_i^* - p_j^*$ is an integer minus a .

Note that it is clear that A is a finite set.

Consider the easier case when A is empty, i.e. $p_i^* + p_j^*$ and $p_i^* - p_j^*$ are integers for all $i \neq j \in [k]$. Let ℓ be a sufficiently large integer such that $1/t_\ell \leq 0.1$ and $\|p^{t_\ell} - p^*\|_\infty \leq 0.1$. Fixing any $s \in \{\pm 1\}^k$, we first show the following claim:

Claim 1. *For every point $x \in T$, we have*

$$\left(\bigcup_{i \in [k]} \mathcal{P}_i(p^{t_\ell}, s_i) \right) \cap B(x, 1/t_\ell) \neq \emptyset \implies x \in \bigcup_{i \in [k]} \mathcal{P}_i(p^*, s_i).$$

Proof. We show a slightly stronger statement in the proof: For any $i \in [k]$, we have

$$\mathcal{P}_i(p^{t_\ell}, s_i) \cap B(x, 1/t_\ell) \neq \emptyset \implies x \in \mathcal{P}_i(p^*, s_i).$$

Let's prove the contrapositive so take any $x \in T$ such that $x \notin \mathcal{P}_i(p^*, s_i)$. Assume without loss of generality that $s_i = +1$. We know that there exists a $j \neq i$ such that either

$$x_i - p_i^* < x_j - p_j^* \quad \text{or} \quad x_i - p_i^* < p_j^* - x_j.$$

For the first case we have $x_i - x_j < p_i^* - p_j^*$. Since both sides are integers we have

$$x_i - x_j \leq p_i^* - p_j^* - 1 \tag{4}$$

so intuitively x is far from $\mathcal{P}_i(p^*, +1)$. From this we can conclude that $B(x, 1/t_\ell) \cap \mathcal{P}_i(p^{t_\ell}, +1)$ is empty. To see this is the case, for any $y \in B(x, 1/t_\ell)$, it follows from $\|x - y\|_\infty \leq \|x - y\|_2 \leq 0.1$ and $\|p^* - p^{t_\ell}\|_\infty \leq 0.1$ and (4) that

$$y_i - y_j \leq p_i^{t_\ell} - p_j^{t_\ell} - (1 - 0.4) < p_i^{t_\ell} - p_j^{t_\ell}$$

and thus, $y \notin \mathcal{P}_i(p^{t_\ell}, +1)$.

The other case where $x_i - p_i^* < p_j^* - x_j$ is similar. □

Notice that $B(x, 1/t_\ell)$ are all disjoint from each other, so for every $x \in T$, we have

$$\frac{\text{vol}\left(\left(\bigcup_{i \in [k]} \mathcal{P}_i(p^{t_\ell}, s_i)\right) \cap B(x, 1/t_\ell)\right)}{\text{vol}(S^{t_\ell})} \leq \frac{\text{vol}(B(x, 1/t_\ell))}{\text{vol}(S^{t_\ell})} = \frac{1}{|T|}.$$

Thus, Claim 1 implies that

$$\frac{\left|\left(\bigcup_{i \in [k]} \mathcal{P}_i(p^*, s_i)\right) \cap T\right|}{|T|} \geq \frac{\text{vol}\left(\left(\bigcup_{i \in [k]} \mathcal{P}_i(p^{t_\ell}, s_i)\right) \cap S^{t_\ell}\right)}{\text{vol}(S^{t_\ell})}. \quad (5)$$

It only remains to note that

$$\text{vol}\left(\left(\bigcup_{i \in [k]} \mathcal{P}_i(p^{t_\ell}, s_i)\right) \cap S^{t_\ell}\right) = \sum_{i \in [k]} \text{vol}(\mathcal{P}_i(p^{t_\ell}, s_i) \cap S^{t_\ell})$$

and

$$\text{vol}(\mathcal{P}_i(p^{t_\ell}, +1) \cap S^{t_\ell}) = \text{vol}(\mathcal{P}_i(p^{t_\ell}, -1) \cap S^{t_\ell})$$

for all $i \in [k]$. The first equation holds because the intersection of any two different pyramids $\mathcal{P}_i(p^{t_\ell}, s_i)$ lies on a hyperplane and thus has 0 volume, and the second equation follows from the choice of p^{t_ℓ} . So Equation (5) implies Lemma 7 since

$$\frac{\left|\left(\bigcup_{i \in [k]} \mathcal{P}_i(p^*, s_i)\right) \cap T\right|}{|T|} \geq \frac{\text{vol}\left(\left(\bigcup_{i \in [k]} \mathcal{P}_i(p^{t_\ell}, s_i)\right) \cap S^{t_\ell}\right)}{\text{vol}(S^{t_\ell})} = \sum_{i \in [k]} \frac{\text{vol}(\mathcal{P}_i(p^{t_\ell}, s_i) \cap S^{t_\ell})}{\text{vol}(S^{t_\ell})} = \frac{1}{2}.$$

Now we consider the general case when A is not empty and let $\alpha > 0$ be the smallest value in A ; note that $\alpha \leq 1/2$. In this case we let ℓ be a sufficiently large integer such that $1/t_\ell \leq 0.1\alpha$ and $\|p^{t_\ell} - p^*\|_\infty \leq 0.1\alpha$. Similarly it suffices to show that for every point $x \in T$ and $i \in [k]$, the following property holds

$$\mathcal{P}_i(p^{t_\ell}, s_i) \cap B(x, 1/t_\ell) \neq \emptyset \implies x \in \mathcal{P}_i(p^*, s_i).$$

Let's prove the contrapositive so take any $x \in T$ and $i \in [k]$ such that $x \notin \mathcal{P}_i(p^*, s_i)$. Assume without loss of generality that $s_i = +1$. There exists a $j \neq i$ such that either

$$x_i - p_i^* < x_j - p_j^* \quad \text{or} \quad x_i - p_i^* < p_j^* - x_j.$$

For the first case we have $x_i - x_j < p_i^* - p_j^*$. Since $p_i^* - p_j^*$ is either an integer or an integer \pm something that is between α and $1 - \alpha$, we have

$$x_i - x_j \leq p_i^* - p_j^* - \alpha.$$

The rest of the proof is similar. □

Given p^* , we round it to an integer point $q^* \in [0 : n]^k$ as follows: Let $q_i^* \in [0, n]$ be an integer such that $|p_i^* - q_i^*| \leq 1/2$ for every $i \in [k]$ (in the case that q_i^* is not unique, we break ties by choosing $q_i^* \in \{p_i^* - 1/2, p_i^* + 1/2\}$ such that q_i^* is odd). It is clear that $q^* \in [0 : n]^k$.

Lemma 8. *The point q^* satisfies $q^* \in [0 : n]^k$ and*

$$\mathcal{P}_i(p^*, +1) \cap T \subseteq \mathcal{P}_i(q^*, +1) \cap T \quad \text{and} \quad \mathcal{P}_i(p^*, -1) \cap T \subseteq \mathcal{P}_i(q^*, -1) \cap T \quad \text{for all } i \in [k].$$

Proof. Fix an arbitrary $i \in [k]$. We show $\mathcal{P}_i(p^*, +1) \cap T \subseteq \mathcal{P}_i(q^*, +1) \cap T$ below, and the other part $\mathcal{P}_i(p^*, -1) \cap T \subseteq \mathcal{P}_i(q^*, -1) \cap T$ can be handled in the same way.

Let $x \in \text{EVEN}(n, k)$ be a point in $\mathcal{P}_i(p^*, +1) \cap T$. So for every $j \neq i$ we have

$$x_i - p_i^* \geq x_j - p_j^* \quad \text{and} \quad x_i - p_i^* \geq p_j^* - x_j,$$

or equivalently

$$x_i - x_j \geq p_i^* - p_j^* \quad \text{and} \quad x_i + x_j \geq p_i^* + p_j^*.$$

It suffices to show that

$$x_i - x_j \geq q_i^* - q_j^* \quad \text{and} \quad x_i + x_j \geq q_i^* + q_j^*.$$

Note that

$$p_i^* \geq q_i^* - 1/2 \quad \text{and} \quad p_j^* \leq q_j^* + 1/2.$$

So we have

$$x_i - x_j \geq p_i^* - p_j^* \geq (q_i^* - 1/2) - (q_j^* + 1/2) = (q_i^* - q_j^*) - 1.$$

Given the rounding rule for obtaining q^* and the fact that $x \in \text{EVEN}(n, k)$, we will argue that it must be the case where $x_i - x_j \geq q_i^* - q_j^*$. Consider two cases below.

If $p_i^* > q_i^* - 1/2$ or $p_j^* < q_j^* + 1/2$, then we know that $x_i - x_j > q_i^* - q_j^* - 1$. Since both $x_i - x_j$ and $q_i^* - q_j^*$ are integer numbers, this implies $x_i - x_j \geq q_i^* - q_j^*$.

If $p_i^* = q_i^* - 1/2$ and $p_j^* = q_j^* + 1/2$, then we know that both q_i^* and q_j^* are odd numbers, which means the parity of $|q_i^* - q_j^*|$ is even. Since the parity of $|x_i - x_j|$ is also even, $x_i - x_j \geq q_i^* - q_j^* - 1$ implies $x_i - x_j \geq q_i^* - q_j^*$.

The other part $x_i + x_j \geq q_i^* + q_j^*$ can be proved similarly. □

This finishes the proof of [Theorem 4](#). □

Remark on [Theorem 4](#). We note that the (possibly off-grid point) p^* in the proof already satisfies the desired property and our algorithm can proceed by querying p^* . However, p^* as defined here is the limit of fixed points found in an infinite sequence of maps. In contrast, [Lemma 8](#) shows that, after rounding, a grid balanced point always exists, which can be found by brute-force enumeration.

6 Impossibility of Strong Approximation under Non-expansion

We consider functions f on the plane with bounded domain and range, e.g. the unit square, that are non-expansive under the ℓ_∞ metric. We will show the following impossibility result.

Theorem 2. *There is no deterministic or randomized algorithm which, when given oracle access to any non-expansive map $f : [0, 1]^2 \mapsto [0, 1]^2$ under the ℓ_∞ -norm, computes in an expected bounded number of queries a point that is within distance $1/4$ of an exact fixed point of f .*

In the proof it will be more convenient to use as the domain a square that is tilted by 45° . We call a rectangle whose sides are at 45° and -45° , a *diamond*. Let D be the diamond whose vertices are the midpoints of the sides of the unit square. Any function g over D can be extended to a function g' over the unit square, by defining for every point $p \in [0, 1]^2$ the value of the function as

$g'(p) = g(\pi(p))$, where $\pi(p)$ is the projection of p onto D , namely, the point obtained by translating p along lines perpendicular to the sides of D until it intersects the boundary of the diamond. Clearly, for any two points $p, q \in [0, 1]^2$, $\|\pi(p) - \pi(q)\|_\infty \leq \|p - q\|_\infty$, hence if the function g over D is non-expansive, then so is the function g' over $[0, 1]^2$. Furthermore, the fixed points of g' are exactly the fixed points of g .

We will prove the statement of the theorem for the domain D . The claim then follows for the unit square. To see this, restrict attention to the non-expansive functions g' on the unit square that are extensions of functions g on the diamond D . If we have an algorithm for the unit square, then we can use the algorithm also for the diamond D : when the algorithm queries a point $p \in [0, 1]^2$ then we query instead its projection $\pi(p) \in D$. If the algorithm outputs at the end a point that is close to a fixed point of g' , then its projection on D is a valid output for g .

For any $\delta \in (0, 1/2)$ and any point s on the SW or NE side of the diamond D that is at least at Euclidean distance δ from the vertices of D , we will define a non-expansive function $f_{\delta,s}$ with unique fixed point s . The function is defined as follows. Draw the line l_0 through s at 45° and let t be the point of intersection with the opposite side of D . Let l_1 and l_2 be the two lines parallel to l_0 that are left and right of l_0 respectively at Euclidean distance δ , and let D_0 be the strip of D that is strictly between the lines l_1 and l_2 . Let $D' = D \setminus D_0$. Every point $p \in D'$ is mapped by $f_{\delta,s}$ to the point that is at Euclidean distance δ towards the line l_0 ; i.e., if $p = (p_1, p_2)$ is left and above l_0 then $f_{\delta,s}(p) = (p_1 + \delta/\sqrt{2}, p_2 - \delta/\sqrt{2})$, and if p is right and below l_0 then $f_{\delta,s}(p) = (p_1 - \delta/\sqrt{2}, p_2 + \delta/\sqrt{2})$.

For a point p in D_0 we define $f_{\delta,s}(p)$ as follows. Let p' be the projection of p onto the line l_0 . Then $f_{\delta,s}(p)$ is the point on l_0 that is at Euclidean distance $(\delta - |pp'|) \cdot |p's|$ from p' in the direction of s , where $|pp'|$, $|p's|$ are the (Euclidean) lengths of the segments pp' and $p's$. Thus for example, if $p = s$ then $p' = s$ and $f_{\delta,s}(s) = s$. If $p = t$ then $p' = t$ and $(\delta - |pp'|) \cdot |p's| = \delta/\sqrt{2}$, so t moves along l_0 distance $\delta/\sqrt{2}$ towards s . Note that if p is on line l_1 or l_2 (i.e. on the boundary sides between D_0 and D'), then $(\delta - |pp'|) \cdot |p's| = 0$, since $|pp'| = \delta$, thus p is mapped to p' whether we treat p as a member of D' or as a member of D_0 . It follows that $f_{\delta,s}(p)$ is continuous over D .

As we noted above, s is a fixed point of $f_{\delta,s}(p)$. We claim that it is the only fixed point. Clearly, any fixed point p must be in D_0 and must lie on the line l_0 , thus $p = p'$. It must satisfy also $(\delta - |pp'|) \cdot |p's| = 0$, hence $|p's| = 0$, and thus, $p = s$.

We will show now that $f_{\delta,s}$ is a non-expansive function, i.e. that $\|f_{\delta,s}(p) - f_{\delta,s}(q)\|_\infty \leq \|p - q\|_\infty$ for all $p, q \in D$. We show first that it suffices to check pairs p, q that are *diagonal* to each other, i.e. such that the line connecting them is at 45° or -45° . Note that such points have the property that the L_∞ distance is tight in both coordinates, $\|p - q\|_\infty = |p_1 - q_1| = |p_2 - q_2|$.

Lemma 9. *If a function f on the diamond D satisfies $\|f(p) - f(q)\|_\infty \leq \|p - q\|_\infty$ for all diagonal pairs of points p, q , then f is non-expansive.*

Proof. Let x, y be any two points that are not diagonal. Consider the diamond with opposite vertices x, y , i.e. draw the lines through x, y at 45° and -45° and consider the rectangle enclosed by them. Let z, w be the other two vertices of this diamond. Suppose without loss of generality that $\|x - y\|_\infty = x_1 - y_1 > |x_2 - y_2|$. Then $x_1 > z_1 > y_1$, and similarly for w . We have $\|x - y\|_\infty = x_1 - y_1 = (x_1 - z_1) + (z_1 - y_1) = \|x - z\|_\infty + \|z - y\|_\infty$. Since f is non-expansive on diagonal pairs, $\|f(x) - f(z)\|_\infty \leq \|x - z\|_\infty$ and $\|f(z) - f(y)\|_\infty \leq \|z - y\|_\infty$. Therefore, $\|f(x) - f(y)\|_\infty \leq \|f(x) - f(z)\|_\infty + \|f(z) - f(y)\|_\infty \leq \|x - z\|_\infty + \|z - y\|_\infty = \|x - y\|_\infty$. \square

Remark. The lemma can be shown to hold more generally in any dimension. That is, if $f : [0, 1]^k \mapsto [0, 1]^k$ has the property that $\|f(p) - f(q)\|_\infty \leq \|p - q\|_\infty$ for all diagonal pairs of points p, q (i.e. such that $|p_i - q_i| = \|p - q\|_\infty$ for all $i \in [k]$), then f is non-expansive.

Lemma 10. *The function $f_{\delta,s}$ is non-expansive.*

Proof. The function $f_{\delta,s}$ was defined according to which region of the domain D a point lies in. There are three regions: the part of D' left of l_1 , the middle region D_0 , and the part of D' right of l_2 . It suffices to check the non-expansiveness for diagonal pairs of points p, q that lie in the same region. If p, q are both in the region left of l_1 , or if they are both right of l_2 , then from the definition we have $\|f(p) - f(q)\|_\infty = \|p - q\|_\infty$.

So suppose p, q are both in D_0 . Assume first that the line pq has angle 45° , i.e. pq is parallel to the line l_0 . Then $\|p - q\|_\infty = |pq|/\sqrt{2}$. Let p', q' be the projections of p, q on l_0 , and let $p'' = f_{\delta,s}(p)$, $q'' = f_{\delta,s}(q)$. Then $|p'p''| = (\delta - |pp'|) \cdot |p's|$, $|q'q''| = (\delta - |qq'|) \cdot |q's|$. Since pq is parallel to l_0 , $(\delta - |pp'|) = (\delta - |qq'|)$ and $|pq| = |p'q'|$. Assume without loss of generality that $|p's| > |q's|$. Then $|p''q''| = |p'q'| - (\delta - |pp'|)(|p's| - |q's|) \leq |p'q'| = |pq|$. Since $\|p - q\|_\infty = |pq|/\sqrt{2}$ and $\|f(p) - f(q)\|_\infty = |p''q''|/\sqrt{2}$, it follows that $\|f(p) - f(q)\|_\infty \leq \|p - q\|_\infty$.

Assume now that the line pq has angle -45° , i.e., pq is perpendicular to l_0 . Again $\|p - q\|_\infty = |pq|/\sqrt{2}$. Now p and q have the same projection $p' = q'$ on l_0 . Let $p'' = f_{\delta,s}(p)$, $q'' = f_{\delta,s}(q)$. We have $|p'p''| = (\delta - |pp'|) \cdot |p's|$, and $|q'q''| = |p'q''| = (\delta - |qp'|) \cdot |p's|$. Therefore, $|p''q''| = |(|pp'| - |qp'|)| \cdot |p's|$. If p, q are on the same side of l_0 then $|(|pp'| - |qp'|)| = |pq|$. If p, q are on opposite sides of l_0 then $|(|pp'| - |qp'|)| < |pq|$. In either case, we have $|p''q''| \leq |pq| \cdot |p's| < |pq|$, since $|p's| \leq |st| = 1/\sqrt{2}$. Again, since $\|p - q\|_\infty = |pq|/\sqrt{2}$ and $\|f(p) - f(q)\|_\infty = |p''q''|/\sqrt{2}$, it follows that $\|f(p) - f(q)\|_\infty \leq \|p - q\|_\infty$. \square

We are ready now to prove the theorem. Intuitively, if the given function is $f_{\delta,s}$ for some s on the NE or SW side of D and some small δ , then for an algorithm (deterministic or randomized) to find a point that is within L_∞ distance $1/4$ of s , it must ask a query within the central region D_0 around s , because otherwise it cannot know whether the fixed point s is on the NE or the SW side of D_0 .

Proof of Theorem 2. Recall that binary search is an optimal algorithm for searching for an unknown item in a sorted array A , both among deterministic and randomized algorithms. If the array has size N , then any randomized comparison-based algorithm requires expected time at least $\log N - 1$ to look up an item in the array whose location is not known.

Suppose there is a (randomized) algorithm B that computes a point that is within $1/4$ of a fixed point of a non-expansive function f over the domain D within a finite expected number n of queries (the expectation is over the random choices of the algorithm). We will show how to solve faster the array search problem. Partition the diamond D into $N = 2^{2n}$ strips by drawing $N - 1$ parallel lines at 45° , spaced at distance $1/(N\sqrt{2})$ from each other, between the NW and SE side of D . Let S_1, \dots, S_N be the N strips. Fix a $\delta < 1/(N2\sqrt{2})$. For each $x \in [N]$, let s_x be the point on the SW side of S_x at Euclidean distance δ from the S vertex of S_x , and t_x the point on the NE side of S_x at Euclidean distance δ from the N vertex. Note that $\|s_x - t_x\|_\infty > 1/2$. Let \mathcal{F} be the family of non-expansive functions $\{f_{\delta,s_x}, f_{\delta,t_x} \mid x \in [N]\}$.

Consider the execution of algorithm B for a function $f \in \mathcal{F}$. Note that the central region D_0 for the functions f_{δ,s_x} and f_{δ,t_x} is contained in the strip S_x . If B queries a point p in another strip S_j , the answer $f(p)$ only conveys the information whether $j < x$ or $j > x$. If some execution of B returns a point q without ever having queried any point in S_x , then all the answers in the execution are consistent with both f_{δ,s_x} and f_{δ,t_x} . Since $\|s_x - t_x\|_\infty > 1/2$, either $\|q - s_x\|_\infty > 1/4$ or $\|q - t_x\|_\infty > 1/4$. Therefore, a correct algorithm B cannot terminate before querying a point in the strip S_x that contains the fixed point of the function.

We can map now the algorithm B to an algorithm B' for the problem of searching for an item in a sorted array A of size N . A choice of an index x in the array A corresponds to a choice of the strip S_x that contains the fixed point of the function $f \in \mathcal{F}$, i.e. choosing one of $f_{\delta, s_x}, f_{\delta, t_x}$. Since B terminates in expected number n of queries, it asks within n steps a query within the strip S_x of the fixed point, hence the expected time of the algorithm B' is at most $n = \log N/2$, a contradiction. \square

7 Promise Problem versus Total Search Version

The problem $\text{CONTRACTION}_\infty(\varepsilon, \gamma, k)$ is a *promise problem*, where we want to compute an ε -fixed point of a given function f with promise that f is a $(1 - \gamma)$ -contraction. For a promise problem, one can define its total search version by asking to find a desired solution as in the promise problem, or a short violation certificate indicating that the given function doesn't satisfy the promise. The total search version of $\text{CONTRACTION}_\infty(\varepsilon, \gamma, k)$, denoted $\text{T-CONTRACTION}_\infty(\varepsilon, \gamma, k)$ is naturally defined as the following search problem.

Definition 4 (Total search version $\text{T-CONTRACTION}_\infty(\varepsilon, \gamma, k)$). *Given a function $f : [0, 1]^k \mapsto [0, 1]^k$, find one of the following:*

- a point $x \in [0, 1]^k$ such that $\|f(x) - x\|_\infty \leq \varepsilon$;
- two points $x, y \in [0, 1]^k$ such that $\|f(x) - f(y)\|_\infty > (1 - \gamma)\|x - y\|_\infty$.

In the black-box setting, the function f is given by an oracle access.

Our theorem in this section shows $\text{T-CONTRACTION}_\infty(\varepsilon, \gamma, k)$ admits the same query bounds as $\text{CONTRACTION}_\infty(\varepsilon, \gamma, k)$.

Theorem 3. *There is an $O(k \log(1/\varepsilon))$ -query algorithm for $\text{T-CONTRACTION}_\infty(\varepsilon, \gamma, k)$.*

Theorem 3 follows from **Lemma 11** below.

Lemma 11. *Let $\{q^1, \dots, q^m\}$ be a set of points in $[0, 1]^k$ and $\{a^1, \dots, a^m\}$ be the corresponding answers from the black-box oracle. There is a $(1 - \gamma)$ -contraction f that is consistent with all the answers if and only if there is no pair t_1, t_2 such that $\|a^{t_1} - a^{t_2}\|_\infty > (1 - \gamma)\|q^{t_1} - q^{t_2}\|_\infty$.*

Proof. If there is some pair t_1, t_2 such that $\|a^{t_1} - a^{t_2}\|_\infty > (1 - \gamma)\|q^{t_1} - q^{t_2}\|_\infty$, then obviously there is no $(1 - \gamma)$ contraction that is consistent with the answers.

Now suppose that no such pair exists. We define a function $f : [0, 1]^k \mapsto [0, 1]^k$ as follows: For every point $x \in [0, 1]^k$ and coordinate $i \in [k]$, we let $f(x)_i = \min_{t \in [m]} \{(1 - \gamma)\|x - q^t\|_\infty + a_i^t\}$; if the minimal value of this set is larger than 1, then we set $f(x)_i = 1$.

We show first that f is consistent with all the query answers, i.e. $f(q^j) = a^j$ for all $j \in [m]$. Since the query points satisfy the contraction property, we have $\|a^j - a^t\|_\infty \leq (1 - \gamma)\|q^j - q^t\|_\infty$ for all $t \neq j$. Therefore, for every coordinate $i \in [k]$, $a_i^j \leq (1 - \gamma)\|q^j - q^t\|_\infty + a_i^t$. Hence, $f(q^j)_i = \min_{t \in [m]} \{(1 - \gamma)\|q^j - q^t\|_\infty + a_i^t\} = a_i^j$. Thus, $f(q^j) = a^j$.

We show now that the function f constructed above is a $(1 - \gamma)$ -contraction. Consider any two points $x, y \in [0, 1]^k$ and a coordinate $i \in [k]$. Suppose without loss of generality that $f(y)_i \leq f(x)_i$. If $f(y)_i = 1$, then also $f(x)_i = 1$ and $|f(x)_i - f(y)_i| = 0 \leq (1 - \gamma)\|x - y\|_\infty$. So suppose $f(y)_i = (1 - \gamma)\|y - q^t\|_\infty + a_i^t$ for some $t \in [m]$. By the triangle inequality, $\|x - q^t\|_\infty \leq \|x - y\|_\infty + \|y - q^t\|_\infty$. Hence $f(x)_i \leq (1 - \gamma)\|x - q^t\|_\infty + a_i^t \leq (1 - \gamma)(\|x - y\|_\infty + \|y - q^t\|_\infty) + a_i^t = (1 - \gamma)\|x - y\|_\infty + f(y)_i$. Therefore, $0 \leq f(x)_i - f(y)_i \leq (1 - \gamma)\|x - y\|_\infty$. Thus, $\|f(x) - f(y)\|_\infty \leq (1 - \gamma)\|x - y\|_\infty$. \square

It follows from [Lemma 11](#) that we can use any algorithm that can solve the promise problem $\text{CONTRACTION}_\infty(\varepsilon, \gamma, k)$ to solve the total search version $\text{T-CONTRACTION}_\infty(\varepsilon, \gamma, k)$ within the same number of queries (without loss of generality, we can assume that the algorithm for the promise problem $\text{CONTRACTION}_\infty(\varepsilon, \gamma, k)$ always queries the approximate fixed point it finds). To see it, consider the sequence of query points q^1, \dots, q^m generated by the algorithm and the corresponding answers a^1, \dots, a^m of the black-box oracle. If $\|q^i - a^i\|_\infty \leq \epsilon$ for some i , then the algorithm will return the ϵ -approximate fixed point q^i and terminate. If there is a pair of generated queries that violate the contraction property, then the algorithm can return the violating pair and terminate. Suppose that neither of these events happens; since all pairs of queries satisfy the contraction property, there is a $(1 - \gamma)$ -contraction map f that is consistent with all the queries. Then the algorithm for the promise problem $\text{CONTRACTION}_\infty(\varepsilon, \gamma, k)$ fails to find an ϵ -approximate fixed point for the black-box oracle function f within the prescribed time bound, a contradiction. Therefore, given any black-box oracle, the algorithm will find either an ϵ -approximate fixed point or a violating pair within the same number of queries $m = O(k \log(1/\epsilon))$ as the promise problem.

8 Conclusions and Discussions

We gave an algorithm for finding an ε -fixed point of a contraction (or non-expansive) map $f : [0, 1]^k \mapsto [0, 1]^k$ under the ℓ_∞ norm in polynomial query complexity. Contraction maps under the ℓ_∞ norm are especially important because several longstanding open problems from various fields can be cast in this framework. The main open question is whether there is a general-purpose (black-box) algorithm for contraction maps that finds an approximate fixed point in polynomial time. Resolving positively this question would have tremendous implications.

Our results provide the necessary first step towards the goal above since, in the black-box model, any polynomial-time algorithm must imply polynomial-query complexity. Nevertheless, we currently see obstacles to making our algorithm time-efficient. Namely, it is not clear how to efficiently find (or even verify) an approximate balanced point, whose existence guarantee is based on non-constructive arguments.

On the other hand, it is possible that finding an approximate fixed point of contraction maps is computationally hard (e.g., CLS-hard or UEOP-*hard*). However, our result restricts what approaches could work in this direction. At a high level, such reduction must work in a white-box manner. To convey the interpretation more concretely, take the canonical CLS-complete problem $\text{CONTINUOUS-LOCALOPT}$ as an example. The problem takes two continuous functions, p and g , as inputs and asks for an approximate local optimum of p with respect to g ; see its formal definition in [\[FGHS23\]](#). It could be the case where the reduction takes the polynomial-size arithmetic circuits that represent p and g as inputs and simulates their gates to generate a contraction map f such that any approximate fixed point of f encodes an approximate local optimum of p with respect to g . In contrast, it cannot be the case where the reduction treats the arithmetic circuits for p and g as evaluation oracles. Specifically, it cannot produce a contraction map f where the evaluation of f requires only a polynomial number of calls to the oracles for p and g , since there are exponential query lower bounds for the $\text{CONTINUOUS-LOCALOPT}$ problem [\[HY20\]](#).

An open question regarding the query complexity of $\text{CONTRACTION}_\infty(\varepsilon, \gamma, k)$ is whether the upper bound $O(k \log(1/\varepsilon))$ is tight. Currently, we are not aware of lower bounds that are stronger than the standard $\Omega(\log(1/\varepsilon))$, which holds even when $k = 1$. In an earlier version of this article, we asked the question of whether similar query bounds can be obtained for contraction maps under the ℓ_1 -norm or norms ℓ_p with $p > 2$. A very recent unpublished paper [\[HLSW25\]](#) answered

this question affirmatively by showing $O(k^2 \log(1/\varepsilon\gamma))$ query upper bounds for all norms ℓ_p with $p \in [1, \infty) \cup \{\infty\}$.

Acknowledgements

The authors are very grateful to the editor and anonymous reviewers for their helpful feedback on improving the presentation.

References

- [ABB20] James Aisenberg, Maria Luisa Bonet, and Sam Buss. 2-d Tucker is PPA-complete. *Journal of Computer and System Sciences*, 108:92–103, 2020. [1](#)
- [Ban22] Stefan Banach. Sur les opérations dans les ensembles abstraits et leur application aux équations intégrales. *Fundamenta mathematicae*, 3(1):133–181, 1922. [1](#)
- [Bel57] Richard Bellman. A Markovian decision process. *Journal of mathematics and mechanics*, pages 679–684, 1957. [1](#)
- [CD08] Xi Chen and Xiaotie Deng. Matching algorithmic bounds for finding a Brouwer fixed point. *Journal of the ACM (JACM)*, 55(3):1–26, 2008. [1](#)
- [CD09] Xi Chen and Xiaotie Deng. On the complexity of 2D discrete fixed point problem. *Theoretical Computer Science*, 410(44):4448–4456, 2009. [1](#)
- [CJK⁺22] Cristian S. Calude, Sanjay Jain, Bakhadyr Khoussainov, Wei Li, and Frank Stephan. Deciding parity games in quasi-polynomial time. *SIAM J. Comput.*, 51(2):17–152, 2022. [1](#)
- [CL55] E. A. Coddington and N. Levinson. *Theory of Ordinary Differential Equations*. McGraw Hill, 1955. [1](#)
- [Con92] Anne Condon. The complexity of stochastic games. *Information and Computation*, 96(2):203–224, 1992. [1](#)
- [DEF⁺21] Xiaotie Deng, Jack R Edmonds, Zhe Feng, Zhengyang Liu, Qi Qi, and Zeying Xu. Understanding PPA-completeness. *Journal of Computer and System Sciences*, 115:146–168, 2021. [1](#)
- [Den67] Eic V. Denardo. Contraction mappings underlying dynamic programming. *SIAM Review*, 9(2):165–177, 1967. [1](#)
- [DP11] Constantinos Daskalakis and Christos H. Papadimitriou. Continuous local search. In Dana Randall, editor, *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2011, San Francisco, California, USA, January 23-25, 2011*, pages 790–804. SIAM, 2011. [1](#), [1](#)
- [DTZ18] Constantinos Daskalakis, Christos Tzamos, and Manolis Zampetakis. A converse to Banach’s fixed point theorem and its CLS-completeness. In Ilias Diakonikolas, David Kempe, and Monika Henzinger, editors, *Proceedings of the 50th Annual ACM*

- SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*, pages 44–50. ACM, 2018. 1, 1
- [EJ91] E. A. Emerson and C. Jutla. Tree automata, μ -calculus and determinacy. In *Proceedings IEEE Symp. on Foundations of Computer Science*, pages 368–377, 1991. 1
- [EY10] Kousha Etessami and Mihalis Yannakakis. On the complexity of Nash equilibria and other fixed points. *SIAM J. Comput.*, 39(6):2531–2597, 2010. 1
- [FGHS23] John Fearnley, Paul Goldberg, Alexandros Hollender, and Rahul Savani. The complexity of gradient descent: $\text{CLS} = \text{PPAD} \cap \text{PLS}$. *J. ACM*, 70(1):7:1–7:74, 2023. 1, 1, 8
- [FGMS20] John Fearnley, Spencer Gordon, Ruta Mehta, and Rahul Savani. Unique end of potential line. *J. Comput. Syst. Sci.*, 114:1–35, 2020. 1, 1
- [Gün89] Matthias Günther. Zum einbettungssatz von J. Nash. *Mathematische Nachrichten*, 144(1):165–187, 1989. 1
- [HKS99] Z. Huang, Leonid G. Khachiyan, and Christopher (Krzysztof) Sikorski. Approximating fixed points of weakly contracting mappings. *J. Complex.*, 15(2):200–213, 1999. 1
- [HLSW25] Sebastian Haslebacher, Jonas Lill, Patrick Schnider, and Simon Weber. Query-Efficient Fixpoints of ℓ_p -Contractions. *arXiv preprint arXiv:2503.16089*, 2025. 8
- [Hol21] Alexandros Hollender. *Structural results for total search complexity classes with applications to game theory and optimisation*. PhD thesis, University of Oxford, UK, 2021. 1
- [How60] Ronald A Howard. *Dynamic programming and Markov processes*. John Wiley, 1960. 1
- [HPV89] Michael D. Hirsch, Christos H. Papadimitriou, and Stephen A. Vavasis. Exponential lower bounds for finding Brouwer fixpoints. *J. Complex.*, 5(4):379–416, 1989. 1
- [HY20] Pavel Hubáček and Eylon Yogev. Hardness of continuous local search: Query complexity and cryptographic lower bounds. *SIAM Journal on Computing*, 49(6):1128–1172, 2020. 8
- [JPY88] David S. Johnson, Christos H. Papadimitriou, and Mihalis Yannakakis. How easy is local search? *J. Comput. Syst. Sci.*, 37(1):79–100, 1988. 1
- [KNY19] Ilan Komargodski, Moni Naor, and Eylon Yogev. White-box vs. black-box complexity of search problems: Ramsey and graph property testing. *Journal of the ACM (JACM)*, 66(5):1–28, 2019. 1
- [Lit94] Michael L Littman. Markov games as a framework for multi-agent reinforcement learning. In *Machine learning proceedings 1994*, pages 157–163. Elsevier, 1994. 1
- [Nas56] John Nash. The imbedding problem for Riemannian manifolds. *Annals of mathematics*, 63(1):20–63, 1956. 1

- [Pap94] Christos H Papadimitriou. On the complexity of the parity argument and other inefficient proofs of existence. *Journal of Computer and system Sciences*, 48(3):498–532, 1994. 1
- [Sha53] L. Shapley. Stochastic games. *Proc. Nat. Acad. Sci.*, 39(10):1095–1100, 1953. 1
- [Sik01] Krzysztof A. Sikorski. *Optimal solution of nonlinear equations*. Clarendon Press, 2001. 1
- [SS02] Spencer D. Shellman and Kris Sikorski. A two-dimensional bisection envelope algorithm for fixed points. *J. Complex.*, 18(2):641–659, 2002. 1
- [SS03] Spencer D. Shellman and Christopher (Krzysztof) Sikorski. A recursive algorithm for the infinity-norm fixed point problem. *J. Complex.*, 19(6):799–834, 2003. 1, 1
- [STW93] Christopher (Krzysztof) Sikorski, Chey-Woei Tsay, and Henryk Wozniakowski. An ellipsoid algorithm for the computation of fixed points. *J. Complex.*, 9(1):181–200, 1993. 1
- [ZP96] Uri Zwick and Mike Paterson. The complexity of mean payoff games on graphs. *Theor. Comput. Sci.*, 158(1&2):343–359, 1996. 1