Nearly Optimal Embeddings of Flat Tori

Ishan Agarwal Oded Regev Yi Tang

APPROX 2020

Preliminaries: Lattices

Lattice L: set of all *integer* linear combinations of a basis in $\mathbb{R}^n$.

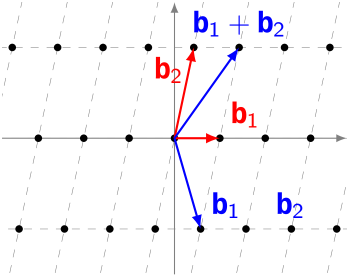

Figure: Example of a 2D lattice, generated by basis $\{\mathbf{b}_1, \mathbf{b}_2\}$.

$_1(L)$: shortest nonzero length in L.

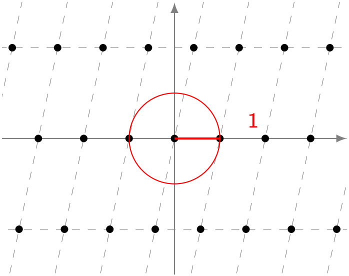

Figure: $_1$ for the 2D example.

We often refer to $_1(L)$ as the “scale” of L.

Preliminaries: Flat Tori

Flat torus R^n=L: quotient space of Euclidean space by lattice; elements: cosets of the form $\mathbf{x} + L$.

Generalizes the standard 2D "torus" R^2=Z^2:

Figure: Transition between the 3D and the quotient representations of R^2=Z^2.

Preliminaries: Flat Tori

Flat torus R^n=L: quotient space of Euclidean space by lattice;
elements: cosets of the form $\mathbf{x} + L$.
Generalizes the standard 2D "torus" R^2=Z^2:

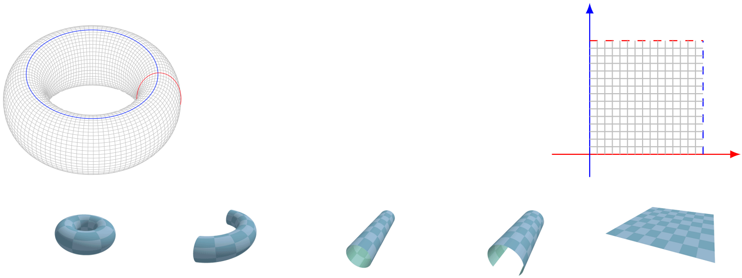

Figure: Transition between the 3D and the quotient representations of R^2=Z^2.

Preliminaries: Distances in Flat Tori

Standard quotient metric on $\mathbb{R}^n/L$:

$\text{dist}_{\mathbb{R}^n/L}(\mathbf{x}+L, \mathbf{y}+L) =$ min distance between $\mathbf{x}+L$ and $\mathbf{y}+L$.

1D Example: What is the distance between 0.2 and 0.8 in $\mathbb{R}/\mathbb{Z}$?

$\text{dist}_{\mathbb{R}/\mathbb{Z}}(0.2+\mathbb{Z}, 0.8+\mathbb{Z}) = 0.4$, the min distance between the cosets:

Preliminaries: Distances in Flat Tori

Standard quotient metric on $\mathrm{R}^n{=}L$:

$\mathrm{dist}_{\mathrm{R}^n=L}(\mathbf{x} + L; \mathbf{y} + L) =$ min distance between $\mathbf{x} + L$ and $\mathbf{y} + L$.

1D Example: What is the distance between 0.2 and 0.8 in R=Z?

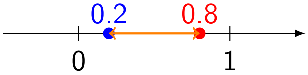

$\mathrm{dist}_{\mathrm{R=Z}}(0{:}2 + \mathrm{Z}; 0{:}8 + \mathrm{Z}) = 0{:}4$, the min distance between the cosets:

Preliminaries: Distances in Flat Tori

Standard quotient metric on $\mathbb{R}^n/L$:

$\mathrm{dist}_{\mathbb{R}^n/L}(\mathbf{x}+L, \mathbf{y}+L)$ = min distance between $\mathbf{x}+L$ and $\mathbf{y}+L$.

1D Example: What is the distance between 0.2 and 0.8 in $\mathbb{R}/\mathbb{Z}$?

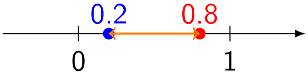

$\mathrm{dist}_{\mathbb{R}/\mathbb{Z}}(0.2+\mathbb{Z}, 0.8+\mathbb{Z}) = 0.4$, the min distance between the cosets:

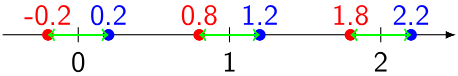

Preliminaries: Distances in Flat Tori cont.

2D Example: What is the distance between the colored points in the following torus (dashed)?

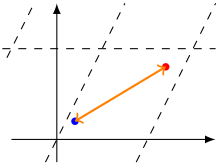

Similarly the distance is not the one within parallelogram, but again the min distance between corresponding cosets.

Preliminaries: Distances in Flat Tori cont.

2D Example: What is the distance between the colored points in the following torus (dashed)?

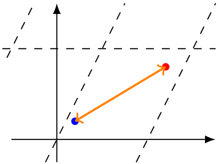

Similarly the distance is not the one within parallelogram, but again the min distance between corresponding cosets.

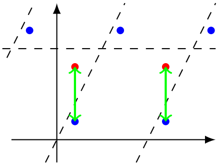

Distortion of metric embedding $f : M_1 \to M_2$:

- the factor by which f changes the distance between two points;
- definition: $\dfrac{\text{expansion factor}}{\text{contraction factor}}$.

Example: embed $\mathbb{R}/\mathbb{Z}$ into Euclidean space $\mathbb{R}^2$.

Figure: Embedding $f : x + \mathbb{Z} \mapsto (\cos(2\pi x), \sin(2\pi x))$, whose distortion is $\pi/2$.

Preliminaries: Distortion

Distortion of metric embedding $f : M_1 \to M_2$:

the factor by which f changes the distance between two points;

definition: $\dfrac{\text{expansion factor}}{\text{contraction factor}}$.

Example: embed $\mathbb{R}/\mathbb{Z}$ into Euclidean space $\mathbb{R}^2$.

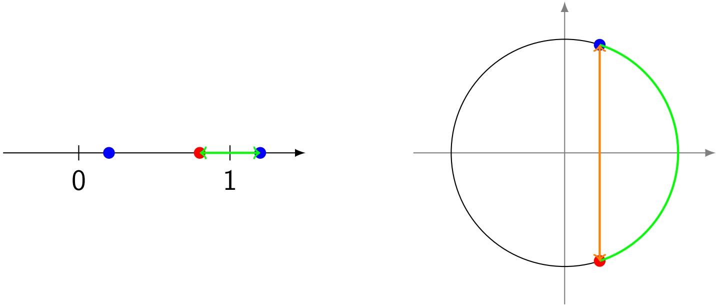

Figure: Embedding $f : x + \mathbb{Z} \mapsto (\cos(2\pi x), \sin(2\pi x))$, whose distortion is $\pi/2$.

Previous Literature

- Question (Khot and Naor, 2005):
  How to embed flat tori into Hilbert space with low distortion?
- Lower bound (KN05): worst case $\Omega(\sqrt{n})$.
- Upper bound (KN05): embedding with distortion $O(n^{3n/2})$.
- Upper bound (Haviv and Regev, 2010):
  - embedding with distortion $O(n\sqrt{\log n})$.
  - embedding with distortion $O(\sqrt{n \log n})$ (under certain condition).

Our Contribution

For any torus $\mathbb{R}^n/L$, we construct a metric embedding of $\mathbb{R}^n/L$ into Hilbert space with distortion $O(\sqrt{n \log n})$.

Previous Literature

- Question (Khot and Naor, 2005):
  How to embed flat tori into Hilbert space with low distortion?
- Lower bound (KN05): worst case $\Omega(\sqrt{n})$.
- Upper bound (KN05): embedding with distortion $O(n^{3n/2})$.
- Upper bound (Haviv and Regev, 2010):
  - embedding with distortion $O(n\sqrt{\log n})$.
  - embedding with distortion $O(\sqrt{n \log n})$ (under certain condition).

Our Contribution

For any torus $\mathbb{R}^n/\mathcal{L}$, we construct a metric embedding of $\mathbb{R}^n/\mathcal{L}$ into Hilbert space with distortion $O(\sqrt{n \log n})$.

Previous Literature

- Question (Khot and Naor, 2005):
  How to embed flat tori into Hilbert space with low distortion?
- Lower bound (KN05): worst case $\Omega(\sqrt{n})$.
- Upper bound (KN05): embedding with distortion $O(n^{3n/2})$.
- Upper bound (Haviv and Regev, 2010):
  - embedding with distortion $O(n\sqrt{\log n})$.
  - embedding with distortion $O(\sqrt{n \log n})$ (under certain condition).

Our Contribution

For any torus $\mathbb{R}^n/L$, we construct a metric embedding of $\mathbb{R}^n/L$ into Hilbert space with distortion $O(\sqrt{n \log n})$.

Previous Literature

- Question (Khot and Naor, 2005):
  How to embed flat tori into Hilbert space with low distortion?
- Lower bound (KN05): worst case $\Omega(\sqrt{n})$.
- Upper bound (KN05): embedding with distortion $O(n^{3n/2})$.
- Upper bound (Haviv and Regev, 2010):
  - embedding with distortion $O(n\sqrt{\log n})$.
  - embedding with distortion $O(\sqrt{n \log n})$ (under certain condition).

Our Contribution

For any torus $\mathbb{R}^n/\mathcal{L}$, we construct a metric embedding of $\mathbb{R}^n/\mathcal{L}$ into Hilbert space with distortion $O(\sqrt{n \log n})$.

Previous Literature

- Question (Khot and Naor, 2005):
  How to embed flat tori into Hilbert space with low distortion?
- Lower bound (KN05): worst case $\Omega(\sqrt{n})$.
- Upper bound (KN05): embedding with distortion $O(n^{3n/2})$.
- Upper bound (Haviv and Regev, 2010):
  - embedding with distortion $O(n\sqrt{\log n})$.
  - embedding with distortion $O(\sqrt{n \log n})$ (under certain condition).

Our Contribution

For any torus $\mathbb{R}^n/L$, we construct a metric embedding of $\mathbb{R}^n/L$ into Hilbert space with distortion $O(\sqrt{n \log n})$.

Embed into Hilbert Space using Gaussians

Observation: for embedding $f : \mathbb{R}^n \to L_2(\mathbb{R}^n), \mathbf{x} \mapsto$ Gaussian centered at $\mathbf{x}$, the L_2 distance between two Gaussians $f(\mathbf{x})$ and $f(\mathbf{y})$

- "saturates" at certain distance, while
- approx. $\propto$ distance between $\mathbf{x}$ and $\mathbf{y}$ before saturation.

Figure: L_2 distance between Gaussians as function of difference of points in 1D.

Embed into Hilbert Space using Gaussians

Observation: for embedding $f : \mathrm{R}^n \;\; ! \;\; L_2(\mathrm{R}^n); \mathbf{x}$ 7! Gaussian centered at $\mathbf{x}$, the L_2 distance between two Gaussians $f(\mathbf{x})$ and $f(\mathbf{y})$

- "saturates" at certain distance, while
- approx. / distance between $\mathbf{x}$ and $\mathbf{y}$ before saturation.

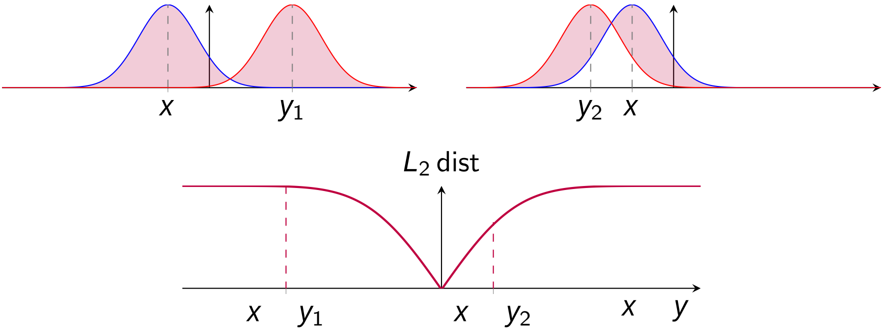

Figure: L_2 distance between Gaussians as function of difference of points in 1D.

The Gaussian Embedding in HR10

The HR10 embedding with distortion $O(\sqrt{n \log n})$:

- Uses Gaussians in the construction.
- Only works if distances to embed never exceed poly(n) $_1$.
- "Saturates" at poly(n) $_1$, like single Gaussian.

Idea about next step:

- first partition torus into direct sum of tori, each representing a different scale;
- then apply the HR10 embedding to each scale separately;
- even if saturation happens at certain scale, a larger scale might compensate.

The HR10 embedding with distortion $O(\sqrt{n\log n})$:

- Uses Gaussians in the construction.
- Only works if distances to embed never exceed $\mathrm{poly}(n)\ {}_1$.
- "Saturates" at $\mathrm{poly}(n)\ {}_1$, like single Gaussian.

Idea about next step:

- first partition torus into direct sum of tori, each representing a different scale;
- then apply the HR10 embedding to each scale separately;
- even if saturation happens at certain scale, a larger scale might compensate.

The Gaussian Embedding in HR10

The HR10 embedding with distortion $O(\sqrt{n \log n})$:

- Uses Gaussians in the construction.
- Only works if distances to embed never exceed $\mathrm{poly}(n)^{-1}$.
- "Saturates" at $\mathrm{poly}(n)^{-1}$, like single Gaussian.

Idea about next step:

- first partition torus into direct sum of tori, each representing a different scale;
- then apply the HR10 embedding to each scale separately;
- even if saturation happens at certain scale, a larger scale might compensate.

2D Example of Saturation

Consider two points in the following torus:

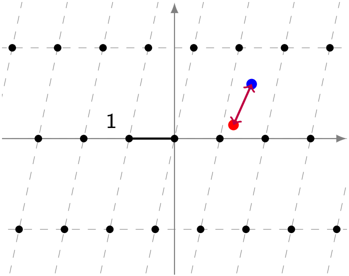

Note that $_1$ is the scale in x-direction, while imagine the scale in y-direction could be arbitrarily large.
Saturation happens and distances in y-direction are not captured.

Besides the original torus, also consider its projection on the y-axis:

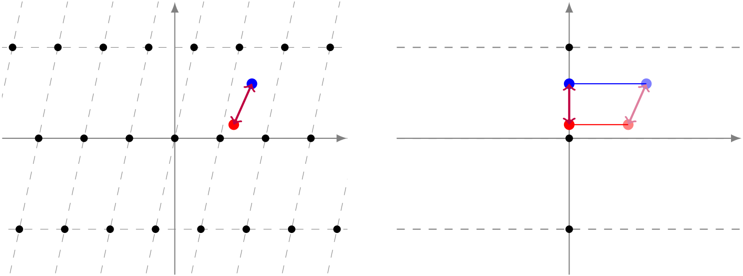

$_1$ of the projection is the scale in y-direction, and HR10 embedding of the projection would capture distances in y-direction.

Partition of Scales in the 2D Example

Besides the original torus, also consider its projection on the y-axis:

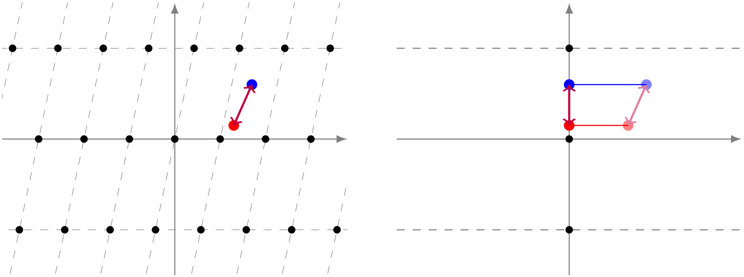

$_1$ of the projection is the scale in y-direction, and HR10 embedding of the projection would capture distances in y-direction.

Partition of Scales in General

General technique:

- Take a chain of *sublattices* $\{\mathbf{0}\} = L_0 \quad L_1 \quad \quad L_m = L$;
- Take projections of $\mathbb{R}^n/L$ orthogonal to L_i for all $i = 0, 1, \ldots, m - 1$;
- Proper choice of sublattices can partition the scales well.

For the 2D example:

- $m = 2$, $L_1 = L \cap x$-axis.
- To project orthogonally to L_1 is to project onto y-axis.

Partition of Scales in General

General technique:

- Take a chain of *sublattices* $\{\mathbf{0}\} = L_0 \subset L_1 \subset \cdots \subset L_m = L$;
- Take projections of $\mathbb{R}^n/L$ orthogonal to L_i for all $i = 0, 1, \ldots, m-1$;
- Proper choice of sublattices can partition the scales well.

For the 2D example:

- $m = 2$, $L_1 = L \cap x$-axis.
- To project orthogonally to L_1 is to project onto y-axis.

Partition of Scales in General

General technique:

- Take a chain of *sublattices* $\{\mathbf{0}\} = L_0 \subseteq L_1 \subseteq \cdots \subseteq L_m = L$;
- Take projections of $\mathbb{R}^n/L$ orthogonal to L_i for all $i = 0, 1, \ldots, m-1$;
- Proper choice of sublattices can partition the scales well.

For the 2D example:

- $m = 2$, $L_1 = L \cap x$-axis.
- To project orthogonally to L_1 is to project onto y-axis.

Partition of Scales in General

General technique:

- Take a chain of *sublattices* $\{\mathbf{0}\} = L_0 \subset L_1 \subset \cdots \subset L_m = L$;
- Take projections of $\mathbb{R}^n/L$ orthogonal to L_i for all $i = 0, 1, \ldots, m-1$;
- Proper choice of sublattices can partition the scales well.

For the 2D example:

- $m = 2$, $L_1 = L \cap x$-axis.
- To project orthogonally to L_1 is to project onto y-axis.

Pitfall of HR10

- HR10 indeed uses the orthogonal projections.
- They only get $O(n\sqrt{\log n})$, losing a factor of $\sqrt{n}$.
- Issue: distances in projected tori with large i get counted repeatedly. For the last one the repetition is m, which at worst could be n.
- For the 2D example, (small) distances in the y-direction get counted twice — both in the original torus and the projection.

Pitfall of HR10

- HR10 indeed uses the orthogonal projections.
- They only get $O(n\sqrt{\log n})$, losing a factor of $\sqrt{n}$.
- Issue: distances in projected tori with large i get counted repeatedly. For the last one the repetition is m, which at worst could be n.
- For the 2D example, (small) distances in the y-direction get counted twice — both in the original torus and the projection.

Pitfall of HR10

- HR10 indeed uses the orthogonal projections.
- They only get $O(n\sqrt{\log n})$, losing a factor of $\sqrt{n}$.
- Issue: distances in projected tori with large i get counted repeatedly. For the last one the repetition is m, which at worst could be n.
- For the 2D example, (small) distances in the y-direction get counted twice — both in the original torus and the projection.

Pitfall of HR10 cont.

Distances in projected tori with large i get counted repeatedly:

- Consider the subspaces spanned by projected tori — they have reversed subspace relationship from the end.
- Orthogonally decompose the entire space according to the subspace relationship, then the repeated counting is like:

(columns: orthogonal decomposition of the entire space,
rows: subspaces spanned by projected tori)

Pitfall of HR10 cont.

Distances in projected tori with large i get counted repeatedly:

- Consider the subspaces spanned by projected tori — they have reversed subspace relationship from the end.
- Orthogonally decompose the entire space according to the subspace relationship, then the repeated counting is like:

1	1	1	$\cdots$	1
	1	1	$\cdots$	1
		1	$\cdots$	1
			$\ddots$	$\vdots$
				1

(columns: orthogonal decomposition of the entire space,
rows: subspaces spanned by projected tori)

Impossible to Isolate Block

Ideally want projection of L in each block (orthogonal component) of the entire space separately to avoid repeated counting.

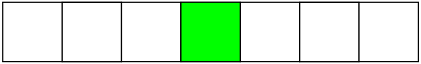

However projecting into arbitrary block is invalid for general lattices.

For the 2D Example: projection onto x-axis is dense.

Impossible to Isolate Block

Ideally want projection of L in each block (orthogonal component) of the entire space separately to avoid repeated counting.

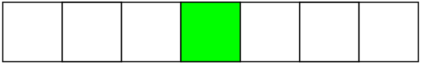

However projecting into arbitrary block is invalid for general lattices.
For the 2D Example: projection onto x-axis is dense.

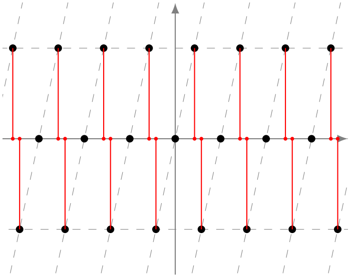

Impossible to Isolate Block

Ideally want projection of L in each block (orthogonal component) of the entire space separately to avoid repeated counting.

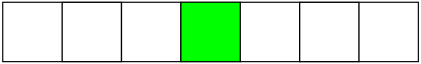

However projecting into arbitrary block is invalid for general lattices.
For the 2D Example: projection onto x-axis is dense.

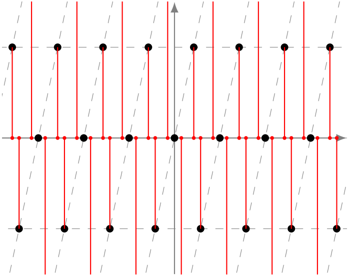

Impossible to Isolate Block

Ideally want projection of L in each block (orthogonal component) of the entire space separately to avoid repeated counting.

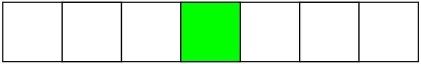

However projecting into arbitrary block is invalid for general lattices.
For the 2D Example: projection onto x-axis is dense.

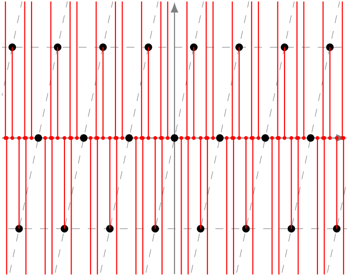

Impossible to Isolate Block

Ideally want projection of L in each block (orthogonal component) of the entire space separately to avoid repeated counting.

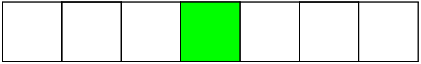

However projecting into arbitrary block is invalid for general lattices.
For the 2D Example: projection onto x-axis is dense.

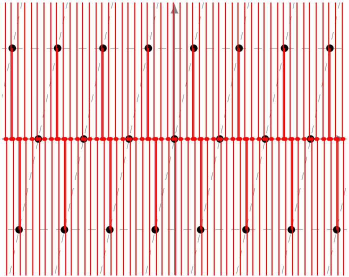

Impossible to Isolate Block

Ideally want projection of L in each block (orthogonal component) of the entire space separately to avoid repeated counting.

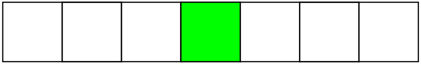

However projecting into arbitrary block is invalid for general lattices.
For the 2D Example: projection onto x-axis is dense.

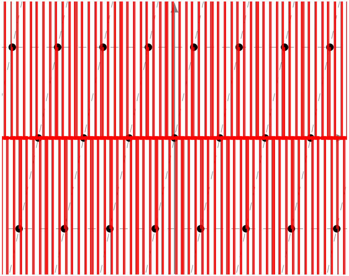

Novel Contribution: Compressed Projections

- Use *compressed projections* that compress the projected tori recursively by a factor 0 <   < 1.
- The repeated counting becomes:

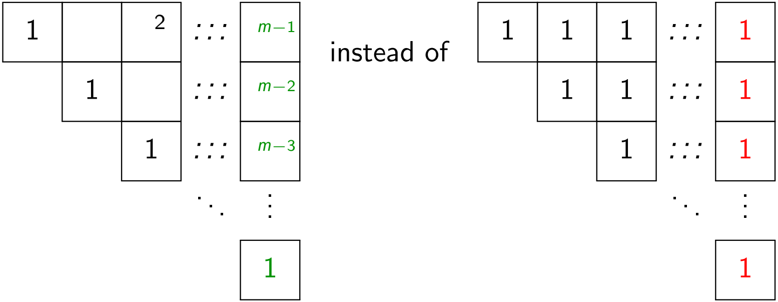

- The geometric series suppresses the $\sqrt{n}$ extra factor into constant.

Using Compressed Projection in the 2D Example

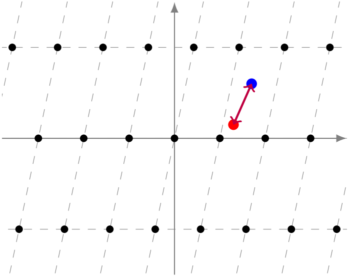

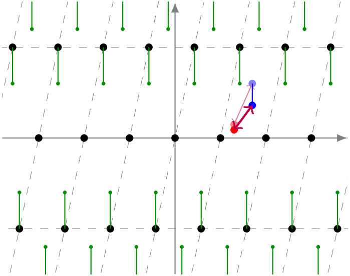

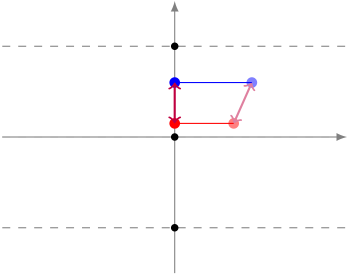

Proof Overview

- Our entire embedding: first partition the scales using compressed projections, and then apply the HR10 embedding to each scale.

- Main technical lemma: the partition step has constant distortion.
- Recall the HR10 embedding has distortion $O(\sqrt{n \log n})$.
- Hence our entire embedding has distortion $O(1) \cdot O(\sqrt{n \log n})$.

Proof Overview

Our entire embedding: rst partition the scales using compressed projections, and then apply the HR10 embedding to each scale.

Main technical lemma: the partition step has constant distortion.

Recall the HR10 embedding has distortion $O(\sqrt{n \log n})$.

Hence our entire embedding has distortion $O(1)$ $O(\sqrt{n \log n})$.

Main Technical Lemma

Main technical lemma: the partition step has constant distortion (even after taking care of the saturation issue of the HR10 embedding afterwards).

- Proving constant expansion of partition step: easy part, due to the geometric series in compressed projections.
- Proving constant contraction of partition step:
  - the main technical part, nontrivial while still elementary;
  - di culty 1: compressed projections worsen the contraction | the geometric series introduces some exponentially small factors;
  - di culty 2: need to prove stronger contraction property that takes care of the saturation issue so that the HR10 embedding works afterwards

Main Technical Lemma

Main technical lemma: the partition step has constant distortion (even after taking care of the saturation issue of the HR10 embedding afterwards).

- Proving constant expansion of partition step: easy part, due to the geometric series in compressed projections.
- Proving constant contraction of partition step:
  - the main technical part, nontrivial while still elementary;
  - di culty 1: compressed projections worsen the contraction | the geometric series introduces some exponentially small factors;
  - di culty 2: need to prove stronger contraction property that takes care of the saturation issue so that the HR10 embedding works afterwards

Main Technical Lemma

Main technical lemma: the partition step has constant distortion (even after taking care of the saturation issue of the HR10 embedding afterwards).

- Proving constant expansion of partition step: easy part, due to the geometric series in compressed projections.
- Proving constant contraction of partition step:
  - the main technical part, nontrivial while still elementary;
  - di culty 1: compressed projections worsen the contraction | the geometric series introduces some exponentially small factors;
  - di culty 2: need to prove stronger contraction property that takes care of the saturation issue so that the HR10 embedding works afterwards

Main Technical Lemma

Main technical lemma: the partition step has constant distortion (even after taking care of the saturation issue of the HR10 embedding afterwards).

- Proving constant expansion of partition step: easy part, due to the geometric series in compressed projections.
- Proving constant contraction of partition step:
  - the main technical part, nontrivial while still elementary;
  - difficulty 1: compressed projections worsen the contraction — the geometric series introduces some exponentially small factors;
  - difficulty 2: need to prove stronger contraction property that takes care of the saturation issue so that the HR10 embedding works afterwards.

Main Technical Lemma

Main technical lemma: the partition step has constant distortion (even after taking care of the saturation issue of the HR10 embedding afterwards).

- Proving constant expansion of partition step: easy part, due to the geometric series in compressed projections.
- Proving constant contraction of partition step:
  - the main technical part, nontrivial while still elementary;
  - difficulty 1: compressed projections worsen the contraction — the geometric series introduces some exponentially small factors;
  - difficulty 2: need to prove stronger contraction property that takes care of the saturation issue so that the HR10 embedding works afterwards.

Open Questions

- Our result: $O(\sqrt{n \log n})$, lower bound: *worst case* $\Omega(\sqrt{n})$.
- KN05, HR10 both give lattice-specific lower bounds as well.
- Tighter bounds for every lattice in terms of lattice parameters?

Thank you!