Hardness of the Covering Radius Problem (APPROX 2026)

Peter Ly (Joint with Huck Bennett)

August 19, 2026

Lattices

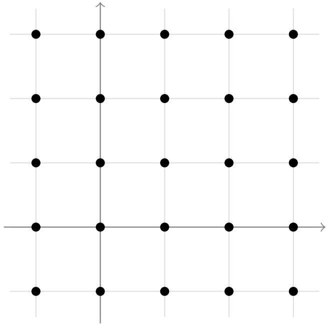

Let $\boldsymbol{b}_1, \boldsymbol{b}_2, \ldots, \boldsymbol{b}_n \in \mathbb{R}^m$ be linearly independent vectors. The lattice generated by $\mathbf{B} := [\boldsymbol{b}_1, \boldsymbol{b}_2, \ldots, \boldsymbol{b}_n]$ is the set of integer linear combinations of $\boldsymbol{b}_i$, i.e., $\mathcal{L}(\mathbf{B}) := \{\mathbf{B}\boldsymbol{z} : \boldsymbol{z} \in \mathbb{Z}^n\}$.

Lattices

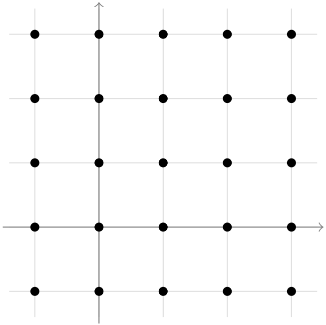

Let $\boldsymbol{b}_1, \boldsymbol{b}_2, \ldots, \boldsymbol{b}_n \in \mathbb{R}^m$ be linearly independent vectors. The lattice generated by $\mathbf{B} := [\boldsymbol{b}_1, \boldsymbol{b}_2, \ldots, \boldsymbol{b}_n]$ is the set of integer linear combinations of $\boldsymbol{b}_i$, i.e., $\mathcal{L}(\mathbf{B}) := \{\mathbf{B}\boldsymbol{z} : \boldsymbol{z} \in \mathbb{Z}^n\}$.

Lattices

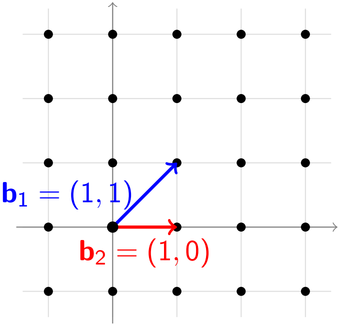

Let $\boldsymbol{b}_1, \boldsymbol{b}_2, \ldots, \boldsymbol{b}_n \in \mathbb{R}^m$ be linearly independent vectors. The lattice generated by $\mathbf{B} := [\boldsymbol{b}_1, \boldsymbol{b}_2, \ldots, \boldsymbol{b}_n]$ is the set of integer linear combinations of $\boldsymbol{b}_i$, i.e., $\mathcal{L}(\mathbf{B}) := \{\mathbf{B}\boldsymbol{z} : \boldsymbol{z} \in \mathbb{Z}^n\}$.

The Covering Radius Problem

Let $\mathcal{L} := \mathcal{L}(\mathbf{B})$ be a lattice. The **covering radius** of $\mathcal{L}$ is

$$\mu_p(\mathcal{L}) := \max_{\boldsymbol{w} \in \mathbb{R}^n} \min_{\boldsymbol{x} \in \mathbb{Z}^n} \|\mathbf{B}(\boldsymbol{w} - \boldsymbol{x})\|_p$$

In other words, covering radius is the farthest distance a point in the ambient space can be from the lattice.

The γ-GapCRP$_p$ problem

Let $p \in [1, \infty]$ and let $\gamma = \gamma(n) \geq 1$.
For inputs a basis $\mathbf{B} \in \mathbb{Q}^{m \times n}$ and a value $r > 0$, the goal of the γ-GapCRP$_p$ is to distinguish between the cases

- (YES instance.) $\mu_p(\mathcal{L}(\mathbf{B})) \leq r$.
- (NO instance.) $\mu_p(\mathcal{L}(\mathbf{B})) > \gamma r$.

The Covering Radius Problem

Let $\mathcal{L} := \mathcal{L}(\mathbf{B})$ be a lattice. The **covering radius** of $\mathcal{L}$ is

$$\mu_p(\mathcal{L}) := \max_{\boldsymbol{w} \in \mathbb{R}^n} \min_{\boldsymbol{x} \in \mathbb{Z}^n} \|\mathbf{B}(\boldsymbol{w} - \boldsymbol{x})\|_p$$

In other words, covering radius is the farthest distance a point in the ambient space can be from the lattice.

The γ-GapCRP$_p$ problem

Let $p \in [1, \infty]$ and let $\gamma = \gamma(n) \geq 1$.
For inputs a basis $\mathbf{B} \in \mathbb{Q}^{m \times n}$ and a value $r > 0$, the goal of the γ-GapCRP$_p$ is to distinguish between the cases

- (YES instance.) $\mu_p(\mathcal{L}(\mathbf{B})) \leq r$.
- (NO instance.) $\mu_p(\mathcal{L}(\mathbf{B})) > \gamma r$.

Example

In the ℓ_2 norm, for $\gamma = 3/2$ and $r = \sqrt{2}/2$,

YES instance

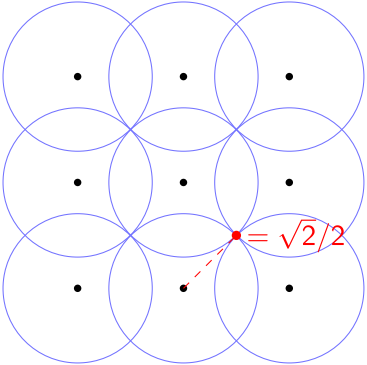

NO instance

Example

In the ℓ_2 norm, for $\gamma = 3/2$ and $r = \sqrt{2}/2$,

YES instance

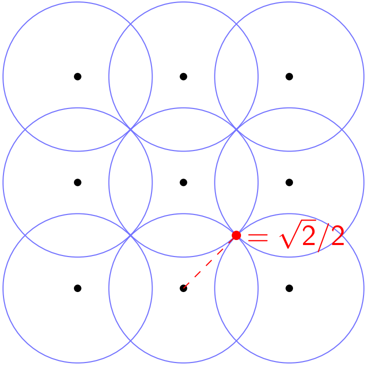

NO instance

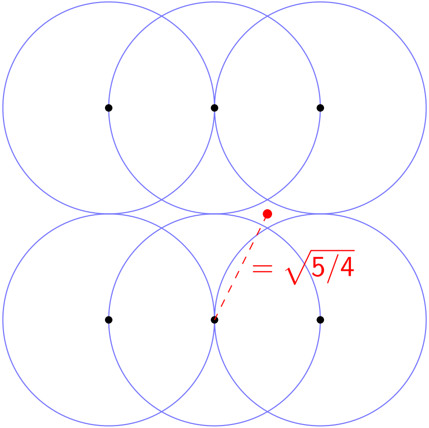

What is known about covering radius

Micciancio
Worst-case to average-case reduction
2004

Guruswami–Micciancio–Regev
Approximation algorithms
2-GapCRP$_p$ $\in$ AM
2005

Haviv–Regev
$3/2$-GapCRP$_\infty$ is Π_2-hard
(Π_2-hardness of approximation for large ℓ_p norms)
2006

Haviv–Lyubashevsky–Regev
$\sqrt{3}$-GapCRP$_2$ $\in$ AM
2009

Things missing from this timeline

- Π_2-hardness of γ-GapCRP$_2$
- NP-hardness of γ-GapCRP$_2$
- NP-hardness of 1-GapCRP$_2$ (i.e. NP-hardness in the exact case)

What is known about covering radius

2004 — **Micciancio**
Worst-case to average-case reduction

2005 — **Guruswami–Micciancio–Regev**
Approximation algorithms
$2\text{-GapCRP}_p \in \mathrm{AM}$

2006 — **Haviv–Regev**
$3/2\text{-GapCRP}_\infty$ is Π_2-hard
(Π_2-hardness of approximation for large ℓ_p norms)

2009 — **Haviv–Lyubashevsky–Regev**
$\sqrt{3}\text{-GapCRP}_2 \in \mathrm{AM}$

Things missing from this timeline

- Π_2-hardness of γ-GapCRP_2
- NP-hardness of γ-GapCRP_2
- NP-hardness of 1-GapCRP_2 (i.e. NP-hardness in the exact case)

What is known about covering radius

Micciancio
Worst-case to average-case reduction
2004

Guruswami–Micciancio–Regev
Approximation algorithms
$2\text{-GapCRP}_p \in \mathsf{AM}$
2005

Haviv–Regev
$3/2\text{-GapCRP}_\infty$ is Π_2-hard
(Π_2-hardness of approximation for large ℓ_p norms)
2006

Haviv–Lyubashevsky–Regev
$\sqrt{3}\text{-GapCRP}_2 \in \mathsf{AM}$
2009

Things missing from this timeline

- Π_2-hardness of $\gamma\text{-GapCRP}_2$
- NP-hardness of $\gamma\text{-GapCRP}_2$
- NP-hardness of 1-GapCRP_2 (i.e. NP-hardness in the exact case)

What is known about covering radius

2004 — **Micciancio**
Worst-case to average-case reduction

2005 — **Guruswami–Micciancio–Regev**
Approximation algorithms
2-GapCRP$_p$ $\in$ AM

2006 — **Haviv–Regev**
$3/2$-GapCRP$_\infty$ is Π_2-hard
(Π_2-hardness of approximation for large ℓ_p norms)

2009 — **Haviv–Lyubashevsky–Regev**
$\sqrt{3}$-GapCRP$_2$ $\in$ AM

Things missing from this timeline

- Π_2-hardness of γ-GapCRP$_2$
- NP-hardness of γ-GapCRP$_2$
- NP-hardness of 1-GapCRP$_2$ (i.e. NP-hardness in the exact case)

Takeaway 1

Theorem

For $p > p_0 \approx 35$, there exists $\gamma(p)$ such that for all constants $\varepsilon > 0$, $(\gamma(p) - \varepsilon)$-GapCRP$_p$ is NP-hard.

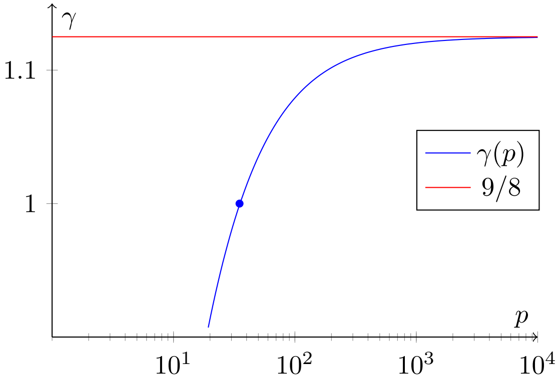

The blue dot is placed at $(p_0, 1)$

Linear discrepancy

Informal definition

The linear discrepancy of a matrix $\mathbf{M} \in \mathbb{Q}^{m \times n}$ is the covering radius problem in the ℓ_∞ norm where the ambient space is replaced with $\mathbf{M} \cdot [0,1]^n$ and the lattice points are replaced with $\mathbf{M} \cdot \{0,1\}^n$.

YES instance

NO instance

At a high-level, we adapt a reduction due to Manurangsi (2021) for the Linear Discrepancy problem to obtain hardness of approximation of GapCRP.

Linear discrepancy

Informal definition

The linear discrepancy of a matrix $\mathbf{M} \in \mathbb{Q}^{m \times n}$ is the covering radius problem in the ℓ_∞ norm where the ambient space is replaced with $\mathbf{M} \cdot [0,1]^n$ and the lattice points are replaced with $\mathbf{M} \cdot \{0,1\}^n$.

YES instance

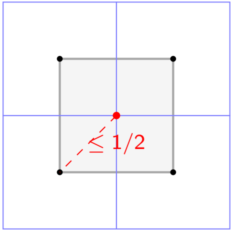

NO instance

At a high-level, we adapt a reduction due to Manurangsi (2021) for the Linear Discrepancy problem to obtain hardness of approximation of GapCRP.

Linear discrepancy

Informal definition

The linear discrepancy of a matrix $\mathbf{M} \in \mathbb{Q}^{m \times n}$ is the covering radius problem in the ℓ_∞ norm where the ambient space is replaced with $\mathbf{M} \cdot [0,1]^n$ and the lattice points are replaced with $\mathbf{M} \cdot \{0,1\}^n$.

YES instance

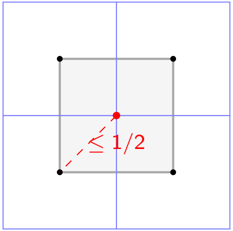

NO instance

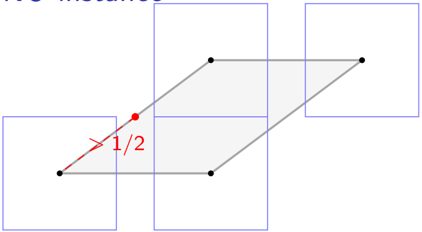

At a high-level, we adapt a reduction due to Manurangsi (2021) for the Linear Discrepancy problem to obtain hardness of approximation of GapCRP.

Linear discrepancy

Informal definition

The linear discrepancy of a matrix $\mathbf{M} \in \mathbb{Q}^{m \times n}$ is the covering radius problem in the ℓ_∞ norm where the ambient space is replaced with $\mathbf{M} \cdot [0,1]^n$ and the lattice points are replaced with $\mathbf{M} \cdot \{0,1\}^n$.

YES instance NO instance

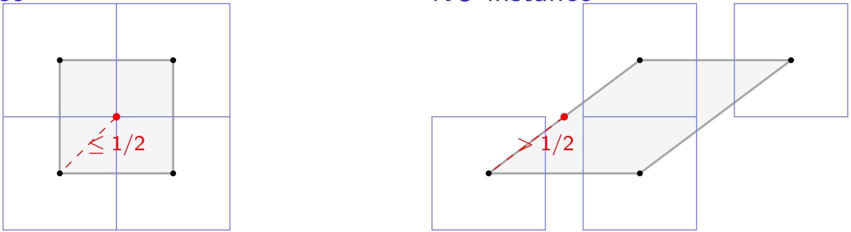

At a high-level, we adapt a reduction due to Manurangsi (2021) for the Linear Discrepancy problem to obtain hardness of approximation of GapCRP.

Our approach

We show hardness of approximation through a hybrid problem BinGapCRP that mixes YES and NO instances of the linear discrepancy problem and the covering radius problem.

Informally $\mathrm{BinGapCRP}_p$ is the problem of determining which of the following two cases hold for a matrix $\mathbf{B}$ encoding a lattice:

- (YES) ℓ_p-analogue of linear discrepancy of $\mathbf{B}$ is small
- (NO) ℓ_p covering radius of $\mathcal{L}(\mathbf{B})$ is large

Our approach

We show hardness of approximation through a hybrid problem BinGapCRP that mixes YES and NO instances of the linear discrepancy problem and the covering radius problem.

Informally BinGapCRP_p is the problem of determining which of the following two cases hold for a matrix $\mathbf{B}$ encoding a lattice:

- (YES) ℓ_p-analogue of linear discrepancy of $\mathbf{B}$ is small
- (NO) ℓ_p covering radius of $\mathcal{L}(\mathbf{B})$ is large

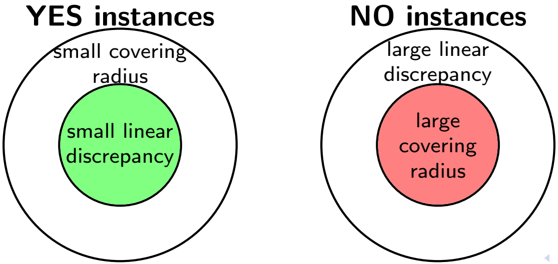

Our approach

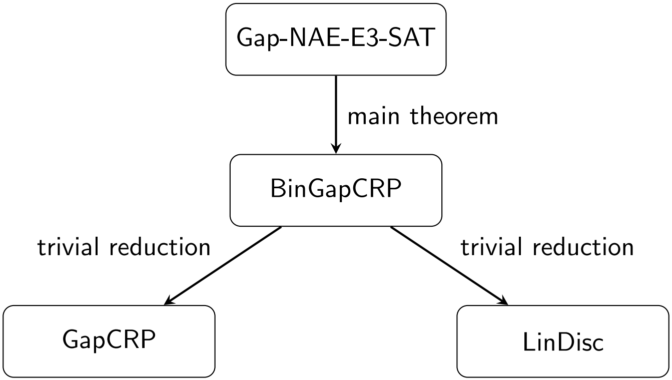

The analysis is more intricate because we have a different NO case to analyze and working with ℓ_p norms allows each clause of the boolean formula to contribute to the ℓ_p norm (rather than only the "worst" ones).

Our approach

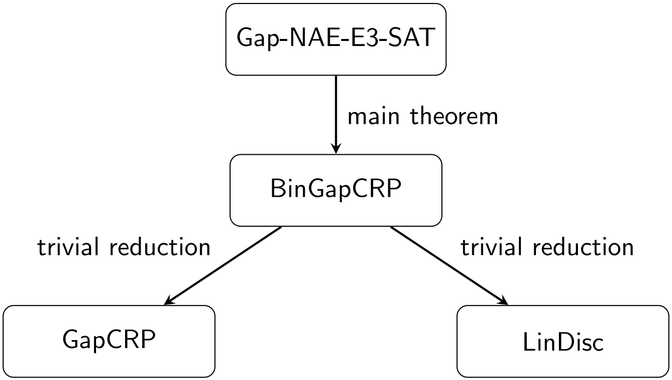

The analysis is more intricate because we have a different NO case to analyze and working with ℓ_p norms allows each clause of the boolean formula to contribute to the ℓ_p norm (rather than only the "worst" ones).

Ingredients

More specifically, the result follows from

1. Using *hardness of approximation* for NAE-E3-SAT with soundness $15/16 + \varepsilon$ and perfect completeness.
2. Appropriately weighting variables in the formula to account for the contributions of each variable to the ℓ_p norm.
3. Extending the analysis from Manurangsi (2021) for a particular gadget matrix **G**.

Ingredients

More specifically, the result follows from

1. Using *hardness of approximation* for NAE-E3-SAT with soundness $15/16 + \varepsilon$ and perfect completeness.
2. Appropriately weighting variables in the formula to account for the contributions of each variable to the ℓ_p norm.
3. Extending the analysis from Manurangsi (2021) for a particular gadget matrix $\mathbf{G}$.

Ingredients

More specifically, the result follows from

1. Using *hardness of approximation* for NAE-E3-SAT with soundness $15/16 + \varepsilon$ and perfect completeness.
2. Appropriately weighting variables in the formula to account for the contributions of each variable to the ℓ_p norm.
3. Extending the analysis from Manurangsi (2021) for a particular gadget matrix $\mathbf{G}$.

Key pieces of the reduction

Let ϕ be a CNF formula with m clauses $(C_1, \ldots, C_m)$ and n variables $(x_1, \ldots, x_n)$. The matrix $\mathbf{B} := \mathbf{B}(\phi)$ with entries

$$B_{i,j} = \begin{cases} 1 & \text{Literal } x_j \text{ appears in clause } C_i \\ -1 & \text{Literal } \neg x_j \text{ appears in clause } C_i \\ 0 & \text{Otherwise} \end{cases}$$

encodes ϕ.

The diagonal matrix $\mathbf{D}_p := \mathbf{D}_p(\phi)$ counts the number of clauses each variable appears in i.e.

$$(D_p)_{i,i} = |\{C_j : C_j \text{ contains } x_i \text{ or its negation}\}|^{1/p}$$

Key pieces of the reduction

Let ϕ be a CNF formula with m clauses $(C_1, \dots, C_m)$ and n variables $(x_1, \dots, x_n)$. The matrix $\mathbf{B} := \mathbf{B}(\phi)$ with entries

$$B_{i,j} = \begin{cases} 1 & \text{Literal } x_j \text{ appears in clause } C_i \\ -1 & \text{Literal } \neg x_j \text{ appears in clause } C_i \\ 0 & \text{Otherwise} \end{cases}$$

encodes ϕ.

The diagonal matrix $\mathbf{D}_p := \mathbf{D}_p(\phi)$ counts the number of clauses each variable appears in i.e.

$$(D_p)_{i,i} = |\{C_j : C_j \text{ contains } x_i \text{ or its negation}\}|^{1/p}$$

Reduction

G is as defined in Manurangsi (2021) i.e.,

$$\mathbf{G} := \begin{pmatrix} 1 & 1 & -1 \\ 1 & -1 & 1 \\ -1 & 1 & 1 \end{pmatrix}$$

The full reduction outputs

$$\mathbf{A} := \mathbf{A}_p(\phi) = \begin{pmatrix} \frac{1}{3}\mathbf{B} & \frac{1}{3}\mathbf{B} & \frac{1}{3}\mathbf{B} \\ & \mathbf{G} \otimes \mathbf{D}_p & \end{pmatrix}$$

Reduction

$\mathbf{G}$ is as defined in Manurangsi (2021) i.e.,

$$\mathbf{G} := \begin{pmatrix} 1 & 1 & -1 \\ 1 & -1 & 1 \\ -1 & 1 & 1 \end{pmatrix}$$

The full reduction outputs

$$\mathbf{A} := \mathbf{A}_p(\phi) = \begin{pmatrix} \frac{1}{3}\mathbf{B} & \frac{1}{3}\mathbf{B} & \frac{1}{3}\mathbf{B} \\ & \mathbf{G} \otimes \mathbf{D}_p & \end{pmatrix}$$

Takeaway 2

Theorem

For $p > p_0 \approx 35$, there exists $\gamma(p)$ such that for all constants $\varepsilon > 0$, $(\gamma(p) - \varepsilon)$-BinGapCRP$_p$ is NP-hard.

Theorem

For any constant $\varepsilon > 0$, $(9/8 - \varepsilon)$-BinGapCRP$_\infty$ is Π_2-hard.

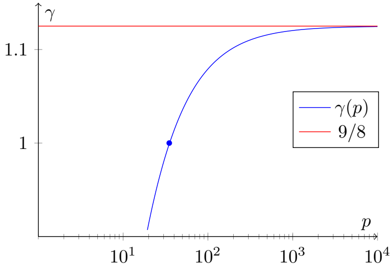

Open questions

- Can we show that $\mathrm{BinGapCRP}_2$ is NP-hard to approximate?
- Can the techniques be refined to show that $\mathrm{BinGapCRP}_p$ is Π_2-hard to approximate?
  - Showing that a gap variant of $\forall\exists$-NAE-3-SAT is Π_2-hard to approximate should suffice here.
  - Can the techniques be improved to show factor c hardness where c agrees with interactive proof bounds? (E.g. Π_2-hardness of approximation for $(\sqrt{3} - \varepsilon)$-GapCRP_2)
- Can we adapt a separate set of techniques that show NP-hardness of approximation for any constant factor?

Open questions

- Can we show that $\mathrm{BinGapCRP}_2$ is NP-hard to approximate?
- Can the techniques be refined to show that $\mathrm{BinGapCRP}_p$ is Π_2-hard to approximate?
  - Showing that a gap variant of $\forall\exists$-NAE-3-SAT is Π_2-hard to approximate should suffice here.
  - Can the techniques be improved to show factor c hardness where c agrees with interactive proof bounds? (E.g. Π_2-hardness of approximation for $(\sqrt{3}-\varepsilon)$-GapCRP_2)
- Can we adapt a separate set of techniques that show NP-hardness of approximation for any constant factor?

Open questions

- Can we show that $\mathrm{BinGapCRP}_2$ is NP-hard to approximate?
- Can the techniques be refined to show that $\mathrm{BinGapCRP}_p$ is Π_2-hard to approximate?
  - Showing that a gap variant of $\forall\exists$-NAE-3-SAT is Π_2-hard to approximate should suffice here.
  - Can the techniques be improved to show factor c hardness where c agrees with interactive proof bounds? (E.g. Π_2-hardness of approximation for $(\sqrt{3} - \varepsilon)$-GapCRP_2)
- Can we adapt a separate set of techniques that show NP-hardness of approximation for any constant factor?

Open questions

- Can we show that $\mathrm{BinGapCRP}_2$ is NP-hard to approximate?
- Can the techniques be refined to show that $\mathrm{BinGapCRP}_p$ is Π_2-hard to approximate?
  - Showing that a gap variant of $\forall\exists$-NAE-3-SAT is Π_2-hard to approximate should suffice here.
  - Can the techniques be improved to show factor c hardness where c agrees with interactive proof bounds? (E.g. Π_2-hardness of approximation for $(\sqrt{3} - \varepsilon)$-GapCRP_2)
- Can we adapt a separate set of techniques that show NP-hardness of approximation for any constant factor?

Summary

Main problem

Obtaining a better understanding of covering radius problem

Theorem (Main)

There exists a function $\gamma(p)$ such that for all $p \geq p_0 \approx 35$ and constants $\varepsilon > 0$, $(\gamma(p) - \varepsilon)$-BinGapCRP$_p$ is NP*-hard.*

Main ideas

- Adapt techniques showing hardness of approximation for linear discrepancy
- Techniques transfer nicely but require more intricate analysis because of more general NO case in decision problem