The hidden subgroup problem for $\mathbb{Z}^k$ for infinite-index subgroups

Greg Kuperberg

UC Davis

June 13, 2022

(E-print in preparation.)

The hidden subgroup problem

Suppose that

$$\begin{tikzcd} G \arrow[rr, "f"] \arrow[dr, two heads] & & X \\ & G/H \arrow[ur, hook] & \end{tikzcd}$$

where G is a discrete group, X is an unstructured set, f can be computed in polynomial time, and $H \leq G$ is a hidden subgroup. The hidden subgroup problem (HSP) is the computational problem of finding H, given f as functional input or an oracle. More explicitly, f hides H means that $f(x) = f(y)$ if and only if $x = yh$. f must be H-periodic, and otherwise 1-to-1.

The performance of HSP is rated by the bit complexity of the output.

The Shor-Kitaev algorithm

Theorem (Shor-Kitaev) Suppose that $G = \mathbb{Z}^k$ and that $H \leq \mathbb{Z}^k$ has finite index (*i.e.*, max rank k). Then we can calculate H in quantum polynomial time (*i.e.*, in functional BQP), uniformly in k and $\|H\|_{\text{bit}}$.

Corollary (Generalized discrete logarithm) If A is an algorithmic finite abelian group, then an isomorphism

$$\phi : A \xrightarrow{\cong} (\mathbb{Z}/a_1) \times (\mathbb{Z}/a_2) \times \cdots \times (\mathbb{Z}/a_\ell)$$

can be constructed and evaluated in quantum polynomial time.

This corollary has many applications to algorithmic number theory and public-key cryptography.

HSP when G is finite

Most algorithms for HSP other than Shor-Kitaev assume that the ambient group G is finite:

- G is finite and H is normal [Hallgren-Russell-Ta-Shma].
- G finite, almost abelian [Grigni-Schulman-Vazirani-Vazirani].
- G is Heisenberg over $\mathbb{Z}/p$ [Bacon-Childs-van-Dam].
- G is finite and 2-step nilpotent [Ivanyos-Sanselme-Santha].
- G is dihedral [K.,Regev,Peikert].
- Some other cases.

HSP has polynomial quantum query complexity whenever G is finite [Ettinger-Høyer-Knill], but it still looks hard in cases such as $G = S_n$. But that is another topic.

New negative results when G is infinite

Theorem (K.) If $G = (\mathbb{Q}, +)$ with standard encoding of rationals, then HSP is NP-hard.

Theorem (K.) If $G = F_k$ is a non-abelian free group with word encoding of elements, then HSP is NP-hard even for normal subgroups.

Theorem (K.) If $G = \mathbb{Z}^k$ with unary vector encoding (*i.e.*, pseudopolynomial query cost), then HSP is as hard as uSVP (unique short lattice vector).

In this context, I first thought that Shor-Kitaev completely solves $\mathbb{Z}^k$ with standard binary encoding of vectors. Then I noticed the finite-index hypothesis.

A new positive result

Theorem (K.) If $G = \mathbb{Z}^k$ with binary encoding and $H \leq \mathbb{Z}^k$ is a lattice with any rank, then H can be found in quantum polynomial time, uniformly in k and $\|H\|_{\text{bit}}$.

The new algorithm begins the same way as Shor-Kitaev, but it requires new ideas for the classical post-processing stage.

Unlike Shor-Kitaev, I do not know of any challenging instances of hiding functions $f : \mathbb{Z}^k \to X$ for this problem; much less, useful applications to number theory or cryptography. I cheerfully conjecture that applications exist.

The HSP algorithm in $\mathbb{Z}^k$

Suppose that $f : \mathbb{Z}^k \to X$ hides a sublattice $H \leq \mathbb{Z}^k$ of some rank $\ell \leq k$. Given parameters $Q \gg S \gg 1$, we follow a version of the standard quantum opening for this HSP:

1. Prepare an approximate Gaussian state on a cube in $\mathbb{Z}^k$:

$$|\psi_G\rangle \propto \sum_{\substack{\vec{x} \in \mathbb{Z}^k \\ \|\vec{x}\|_\infty < Q/2}} \exp(-\pi \|\vec{x}\|_2^2 / S^2) |\vec{x}\rangle$$

2. Apply the hiding function f to $|\psi_G\rangle$ in unitary form:

$$U_f |\psi_G\rangle \propto \sum_{\vec{x}} \exp(-\pi \|\vec{x}\|_2^2 / S^2) |\vec{x}, f(\vec{x})\rangle$$

   (As usual, U_f must use uncomputation to erase scratch work.) Throw away the output, leaving a partially measured input state $|\psi_{H+\vec{v}}\rangle \in \ell^2((\mathbb{Z}/Q)^k)$.

Fourier measurement and dual samples

3. Apply the quantum Fourier operator $F_{(\mathbb{Z}/Q)^k}$ to $|\psi_{H+\vec{v}}\rangle$ and measure a Fourier mode $\vec{y}_0 \in (\mathbb{Z}/Q)^k$. Rescale $\vec{y}_0$ to obtain:

$$\vec{y}_1 = \frac{\vec{y}_0}{Q} \in (\mathbb{R}/\mathbb{Z})^k$$

The vector $\vec{y}_1$ is approximately a randomly chosen element of the dual group

$$H^\# = \widehat{\mathbb{Z}^k/H} \le (\mathbb{R}/\mathbb{Z})^k,$$

Explicitly, $H^\#$ consists of those $\vec{y}$ such that $\vec{x} \cdot \vec{y} \in \mathbb{Z}$ for all $\vec{x} \in H$.

The sample $\vec{y}_1$ also has noise due to both Gaussian blur and discretization. This noise is exponentially small, but so is the feature scale of $H^\#$.

Examples of $H^\#$

Here are two examples of $H^\#$ and a noisy sample $\vec{y}_1 \underset{\sim}{\in} H^\#$.

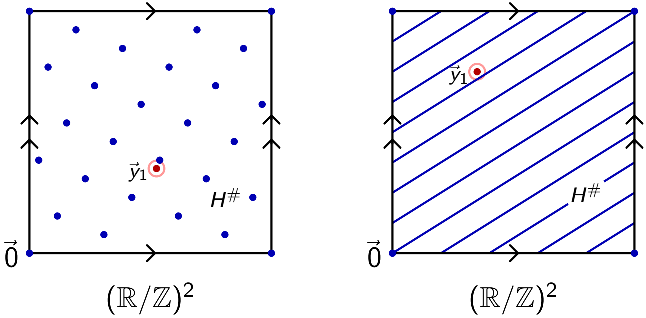

On the left, H has full rank and $H^\#$ is a finite group. On the right, when H has lower rank, $H^\#$ a striped pattern whose connected subgroup $H_1^\#$ is a complicated torus.

Solving for $H^{\#}$ from random samples

The easy case

Goal: Find $H^{\#} \leq (\mathbb{R}/\mathbb{Z})^k$ from noisy random samples $\vec{y}_1 \underset{\sim}{\in} H^{\#}$.

Shor-Kitaev: If H has full rank and $H^{\#}$ is finite, then we can find rational approximations to the coordinates of $\vec{y}_1$ using the continued fraction algorithm. In this case, $O(\log|H^{\#}|)$ samples generate $H^{\#}$ with high probability. (For instance, when $H = h\mathbb{Z} \leq \mathbb{Z}$, then $H^{\#} = \frac{1}{h}\mathbb{Z}/h \leq \mathbb{R}/\mathbb{Z}$, and we can succeed with one or two samples.)

New: If H has rank $\ell < k$, then $\dim H^{\#} = k - \ell$. Any one coordinate of $\vec{y}_1$ is uniformly random in $\mathbb{R}/\mathbb{Z}$. Rational approximation of the coordinates does not work. Happily, the LLL lattice algorithm works, even in high dimensions.

Solving for $H^{\#}$ from random samples

The hard case

A randomly chosen $\vec{y}_0 \in H^{\#}$ almost surely densely generates the connected subgroup $H_1^{\#}$. We look for multiples of $\vec{y}_1 \mathrel{\underset{\sim}{\in}} H^{\#}$ near $\vec{0}$ by lifting the dense orbit to a lattice one dimension higher.

4. Using a single sample $\vec{y}_1$, make a lattice $L \leq \mathbb{R}^{k+1}$ with basis:

$$\vec{e}_1, \vec{e}_2, \ldots, \vec{e}_k, (\widetilde{\vec{y}}_1, \frac{1}{T})$$

Here $S \gg T \gg R$, and $1/R$ is a lower bound for the feature scale of $H^{\#}$. Then calculate an LLL basis of short vectors of L:

$$\vec{b}_1, \vec{b}_2, \ldots, \vec{b}_{k+1} \in L \leq \mathbb{R}^{k+1}$$

The first $k - \ell + 1$ vectors approximately span $T_{\vec{0}}(H^{\#} \oplus \mathbb{R})$.

Lifting a dense orbit to a lattice

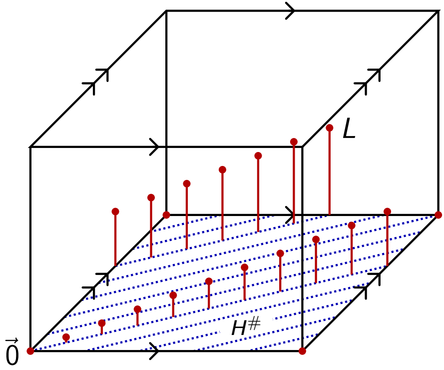

We lift a dense orbit (approximately) in $H^{\#} \leq (\mathbb{R}/\mathbb{Z})^k$ to an anisotropic lattice $L \leq \mathbb{R}^{k+1}$ in the next dimension.

Denoising the data

5. Put the matrix

$$B = [\vec{b}_1, \vec{b}_2, \ldots, \vec{b}_{k-\ell+1}]$$

in CREF form by inverting an appropriate square submatrix A. When $k, k-\ell >> 1$, we can find a good choice with a greedy algorithm based on the Cauchy-Binet formula

$$\det(B^T B) = \sum_A (\det A)^2.$$

6. The entries of $A^{-1}B$ are approximate rational numbers that can be denoised with the continued fraction algorithm. This yields a rational basis for

$$T_{\vec{0}}(H^{\#} \oplus \mathbb{R}) = H_{\mathbb{R}}^{\perp} \leq \mathbb{R}^{k+1},$$

and thus a rational basis for $H_{\mathbb{R}} = H \otimes \mathbb{R}$.

The last step

Where we stand: Using the quantum part of the algorithm, we obtained a noisy sample $\vec{y}_1 \underset{\sim}{\in} H^\# \leq (\mathbb{R}/\mathbb{Z})^k$, where $H \leq \mathbb{Z}^k$ is the hidden subgroup. We then used $\vec{y}_1$ to define a lattice $L \leq \mathbb{R}^{k+1}$. We can apply the LLL algorithm to L and denoise the result to find a rational basis for $H_\mathbb{R} = H \otimes \mathbb{R}$.

To finish the algorithm:

7. We can use the Smith normal form algorithm to convert a rational basis for $H_\mathbb{R}$ to an integral basis for $H_1 = H_\mathbb{R} \cap \mathbb{Z}^k$. Since the original $H \leq \mathbb{Z}^k$ has finite index in its rational closure H_1, we can use the standard Shor-Kitaev algorithm to find H.

Open problems

- Especially when the ambient dimension k is large, it is more efficient to find H using $m > 1$ samples $\vec{y}_1 \underset{\sim}{\in} H^{\#} \leq (\mathbb{R}/\mathbb{Z})^k$ and apply LLL to a lattice $L \leq \mathbb{R}^{k+m}$. This leads to a tradeoff between classical and quantum resources, that also depends on the complexity of the hiding function $f : \mathbb{Z}^k \to X$.
- Is there a challenging hiding function $f : \mathbb{Z}^k \to X$ which is H-periodic and otherwise injective, and H has lower rank $\ell < k$? Challenging here means that the algorithm to compute f does not reveal H directly, nor with the aid of an efficient companion classical algorithm.
- There should be a mutual generalization of this algorithm and the Eisenträger-Hallgren-Kitaev-Song algorithm for $H \leq \mathbb{R}^k$.