Chapter 5

Fourier Sampling over Abelian Groups

5.1 Introduction

In this chapter we analyze the relationship between the Fourier transform over the cyclic group $\mathbb{Z}_p$ and the Fourier transform over the cyclic group $\mathbb{Z}_q$, where $q > p$. Since any abelian group is a direct product of cyclic groups, there is a natural extension of the results to general abelian groups. The two main results are a Fourier sampling theorem, showing that Fourier sampling is robust under group change, and a new quantum Fourier transform algorithm that is the most efficient known to date.

The Fourier transform over the group $\mathbb{Z}_p$ is a map from $\mathbb{C}^p$ to $\mathbb{C}^p$. Recall that Fourier sampling is a map from $\mathbb{C}^p$ to $\mathbb{R}^p$, where given a state $|v\rangle$, we get a probability distribution $\mathcal{D}_{|\hat{v}\rangle}$. We wish to study how the Fourier transform and how Fourier sampling,

with respect to $\mathbb{Z}_q$, carried out on the state $|v\rangle$, relates to the respective operations over $\mathbb{Z}_p$. To do this we must first say how to map the state $|v\rangle$ over $\mathbb{Z}_p$ to a state over $\mathbb{Z}_q$. We must also specify how to map back from $\mathbb{Z}_q$ to $\mathbb{Z}_p$ after the Fourier transform. Below we define several maps for each direction. The choice of map is quite crucial in determining how closely the Fourier transforms approximate each other.

We define two maps on the state $|v\rangle$ from $\mathbb{Z}_p$ to $\mathbb{Z}_q$. The first is the inclusion map $\iota : \mathbb{C}^p \to \mathbb{C}^q$ which just identifies $\mathbb{C}^p$ with a subspace of $\mathbb{C}^q$ and maps $|v\rangle$ to itself. The second map R_q repeats the vector to have q coefficients. On a vector $|v\rangle = \sum_{i=0}^{p-1} v_i|i\rangle \in \mathbb{C}^p$, $R_q|v\rangle = c \cdot \sum_{i=0}^{q-1} v_{i \bmod p}|i\rangle$, where c is the appropriate normalization factor.

After we compute the Fourier transform over $\mathbb{Z}_q$, we have three possible ways of mapping back to $\mathbb{Z}_p$ to approximate Fourier sampling. A natural map is to divide $[0, q-1]$ into p equal intervals, and map all integer points in the j^{th} interval to j. This gives us the bucket distribution $\mathcal{D}^{\mathrm{B}}_{|\hat{u}\rangle}(i) = \sum_{t \in T} |\widehat{u}_{s_i+t}|^2$, where $T = \{-\lfloor \frac{q}{2p} \rceil + 1, \ldots, \lfloor \frac{q}{2p} \rceil\}$. Another possibility is to use only the midpoint of each interval (i.e. the closest integer to the midpoint). First we define the map $s(i) = \lfloor \frac{q}{p} i \rceil$. Let $s_i = s(i)$. The inverse of this map, the projection map $P : \mathbb{C}^q \to \mathbb{C}^p$, is then defined as

$$P|j\rangle = \begin{cases} |i\rangle & \text{if } s_i = j \text{ for some } i \in \{0, \ldots, p-1\} \\ 0 & \text{otherwise} \end{cases}$$

Let $|\hat{u}\rangle$ be the result of the Fourier transform over $\mathbb{Z}_q$. The midpoint distribution $\mathcal{D}^{\mathrm{M}}$ is defined by $\mathcal{D}^{\mathrm{M}}_{|\hat{u}\rangle}(i) = c^{-2} \cdot \mathcal{D}_{|\hat{u}\rangle}(s_i)$ where the normalization factor is $c^2 = \sum_{i=0}^{p-1} \mathcal{D}_{|\hat{u}\rangle}(s_i)$.

There are cases where we must Fourier sample and p is unknown, but an upper-bound n is known. In this case, we define the map from $\mathbb{Z}_q$ to $\mathbb{Q}$, such that i is mapped to the best rational approximation of i/q with denominator at most n (this is found by continued

fractions). This induced distribution $\mathcal{D}^{\mathrm{CF}}$, is called the continued fractions distribution.

In Section 5.2 we show that by zero-filling and using the midpoint distribution $\mathcal{D}^{\mathrm{M}}$, we can approximate Fourier sampling over $\mathbb{Z}_p$. The intuitive reason that this works is as follows: the Fourier transform over $\mathbb{Z}_p$ can be obtained by regarding $|v\rangle$ as a vector over $\mathbb{Z}$, taking the inverse Fourier transform with respect to the circle group to get a continuous function, and then restricting to its values at p evenly spaced points. Using the midpoint distribution is wasteful in terms of sample complexity, and it is natural to try to use the bucket distribution. However, we give a counterexample showing that using buckets does not even work for cyclic groups.

In Section 5.3 we show that repeating the vector via R_q and using the bucket distribution gives extremely good approximations to Fourier sampling over $\mathbb{Z}_p$. This is based on some results discovered independently in the context of eigenvalue estimation. Since this technique is more efficient in its use of samples, it extends to abelian groups.

In Section 5.4 we show that by combining the two methods of repeating the vector and then zero-filling, and adding another ingredient to map back to $\mathbb{Z}_p$, we get a very fast quantum algorithm for the Fourier transform over $\mathbb{Z}_p$.

5.2 Zero-Filling

5.2.1 The Main Theorem

The main theorem of this section may be regarded as showing that the diagram

$$\begin{array}{ccc} \mathbb{C}^p & \xrightarrow{F_p} & \mathbb{C}^p \\ \downarrow \iota & & \uparrow \sqrt{\frac{q}{p}} P \\ \mathbb{C}^q & \xrightarrow{F_q} & \mathbb{C}^q \end{array}$$

"commutes approximately" when q is sufficiently large compared to p. By "commutes approximately" we mean that for all $v \in \mathbb{C}^p$, $|F_p|v\rangle - \sqrt{\frac{q}{p}} P F_q \iota |v\rangle| \leq \epsilon$. In the context of quantum computation, this allows us to Fourier sample over $\mathbb{Z}_q$ instead of $\mathbb{Z}_p$.

Theorem 5.1 *For all $s \geq 0$ and $|v\rangle \in \mathbb{C}^p$, if $q \geq (24 \cdot s \cdot \ln p) \cdot p$, then*

$$\| F_p|v\rangle - \sqrt{\tfrac{q}{p}} P F_q \iota |v\rangle \| \leq \frac{\| |v\rangle \|}{8s}$$

The setup is illustrated with an example in Figure 5.1.

5.2.2 Application to the Quantum Case

We will now apply Theorem 5.1 to the quantum case to prove that Fourier sampling over different group sizes works. Let $|v\rangle$ be a quantum state, $|\hat{v}\rangle = F_p|v\rangle$, and $|\hat{u}\rangle = F_q \iota |v\rangle$. Let $\mathcal{D}_{|\hat{v}\rangle} : \{0, \ldots, p-1\} \to [0,1)$ and $\mathcal{D}_{|\hat{u}\rangle} : \{0, \ldots, q-1\} \to [0,1)$ be the probability distributions induced by measuring $|\hat{v}\rangle$ and $|\hat{u}\rangle$, respectively. Let $\mathcal{D}^{\mathrm{M}}_{|\hat{u}\rangle} : \{0, \ldots, p-1\} \to [0,1)$ be the distribution defined by $\mathcal{D}^{\mathrm{M}}_{|\hat{u}\rangle}(i) = c^{-2} \cdot \mathcal{D}_{|\hat{u}\rangle}(s_i)$ where the normalization factor is $c^2 = \sum_{i=0}^{p-1} \mathcal{D}_{|\hat{u}\rangle}(s_i)$. The distribution $\mathcal{D}^{\mathrm{M}}_{|\hat{u}\rangle}$ can be thought of as induced by the experiment where $|\hat{u}\rangle$ is measured, and if the value returned is not s_i for some $i \in \{0, \ldots, p-1\}$,

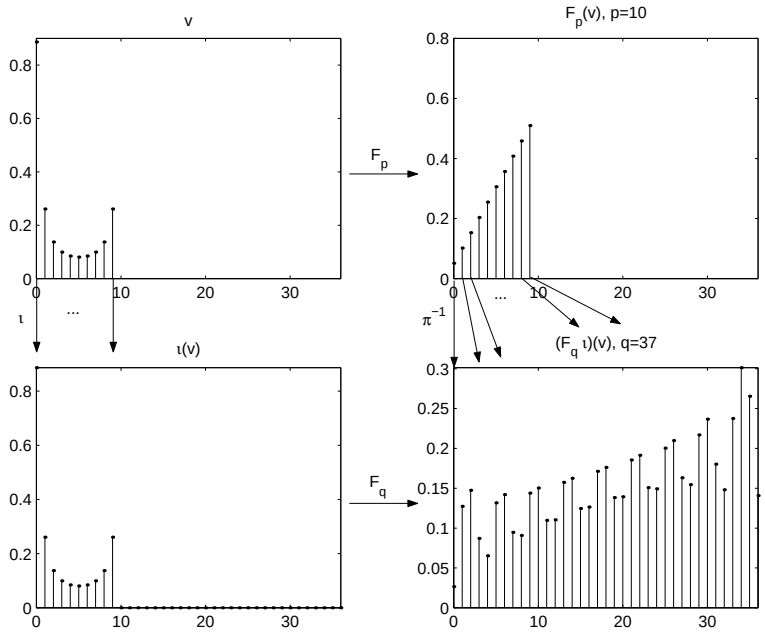

Figure 5.1: Graphs of magnitudes of amplitudes of four vectors.

throw it out and repeat. For this reason it is also important to have a lower bound on the probability of measuring some s_i.

Corollary 5.1 *If* $q \geq (24 \cdot s \cdot \ln p) \cdot p$ *, then* $|\mathcal{D}^M_{|\hat{u}\rangle} - \mathcal{D}_{|\hat{v}\rangle}|_1 \leq \frac{1}{s}$ *and* $\Pr[$*see some* s_i*,* $i \in \{0, \ldots, p-1\}] \geq \frac{p}{q}\left(1 - \frac{1}{s}\right)$.

For the same reason as the lower bound, the probability of seeing some s_i is at most $\frac{p}{q}\left(1 + \frac{1}{s}\right)$. Note that this implies that as q gets larger relative to p, the chance of actually measuring some s_i decreases, and that there is always at least a $\frac{1}{24s\ln p}$ factor loss.

Proof: (Corollary) Let $|\delta\rangle = \sqrt{\frac{q}{p}}P|\hat{u}\rangle - |\hat{v}\rangle$ and $|\delta'\rangle = c \cdot P|\hat{u}\rangle - |\hat{v}\rangle$ (see Figure 5.2). The bound for $|\delta'\rangle$ follows from the fact that $\||\delta'\rangle\| \leq \||\delta\rangle\| + \||\delta'\rangle - |\delta\rangle\| \leq 2\||\delta\rangle\|$, since $|\delta'\rangle - |\delta\rangle$ cannot be longer than $|\delta\rangle$. By Theorem 5.1, $\||\delta\rangle\| \leq \frac{1}{8s}$, so $\||\delta'\rangle\| \leq \frac{1}{4s}$. The bound in the

corollary follows from the following lemma:

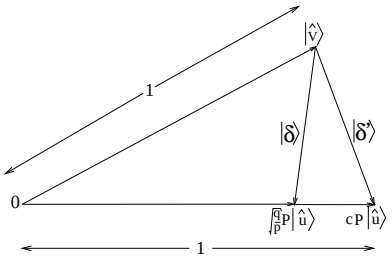

Figure 5.2: The notation used in the proof of Theorem 5.1. $P|\hat{u}\rangle = \sum_{i=0}^{p-1} \hat{u}_{s_i}|i\rangle$, and $c \cdot \hat{u}$ is $P|\hat{u}\rangle$ normalized to unit length.

Lemma 5.1 ([8],Lemma 3.2.6) *Let $|\phi\rangle$ and $|\psi\rangle$ be two unit length vectors with $\||\phi\rangle - |\psi\rangle\| \leq \epsilon$. Then the total variation distance between the probability distributions resulting from measurements of $|\phi\rangle$ and $|\psi\rangle$ is at most 4ϵ.*

Since Pr[see some s_i, $i \in \{0, \ldots, p-1\}$] $= \|P|\hat{u}\rangle\|^2$ and $1 - \|\sqrt{\frac{q}{p}}P|\hat{u}\rangle\| \leq \|\sqrt{\frac{q}{p}}P|\hat{u}\rangle - |\hat{v}\rangle\| \leq \frac{1}{8s}$ by Theorem 5.1, we have $\|P|\hat{u}\rangle\| \geq \sqrt{\frac{q}{p}}(1 - \frac{1}{8s})$, and the bound follows. ■

This method only extends to abelian groups with a constant number of factors. An arbitrary abelian group A is a direct product of cyclic groups $\mathbb{Z}_{n_1} \times \cdots \times \mathbb{Z}_{n_k}$, and we could try a Fourier sampling approach using $\mathcal{D}^{\mathrm{M}}$ by using a larger cyclic group for each factor. That is, we would compute the Fourier transform $F_{A'} = F_{m_1} \bigotimes \cdots \bigotimes F_{m_k}$, where $m_j \geq n_j$ for all j and pick points $s_i = (s_{i_1}^1, \ldots, s_{i_k}^k)$, where s^j rounds according to m_j and n_j. The problem is that we know that for each j, a factor of n_j/m_j is introduced, and this makes the probability of seeing a point s_i too small, unless the group only has a constant number of factors. For example, if n_j/m_j is only $1/2$, then the probability of seeing a point s_i drops below $1/2^k$.

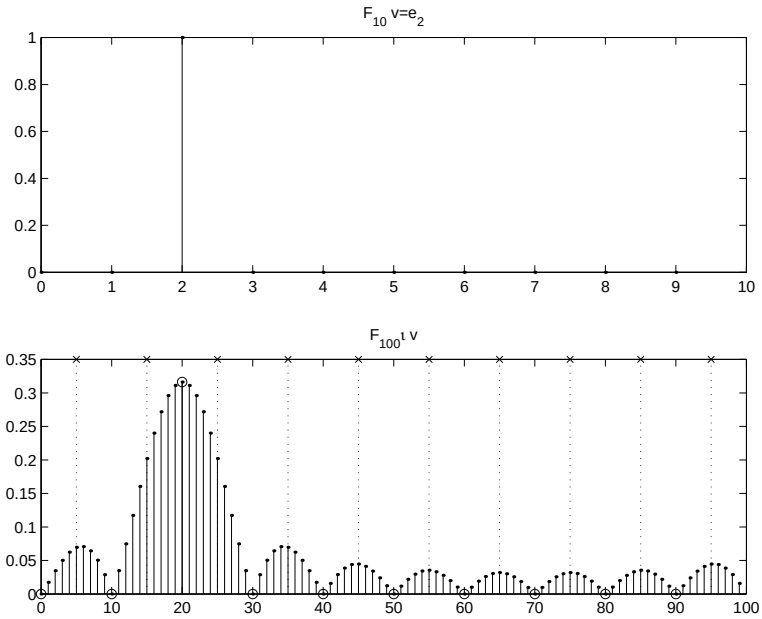

Figure 5.3: $|\hat{v}\rangle = |2\rangle$ and $q = 100$ in $|\hat{u}\rangle$. Buckets are marked off with dotted lines.

5.2.3 The Continuous Picture

In this section we will give the intuition behind the main theorem. We will show that the Fourier transform over $\mathbb{Z}_p$ can be obtained as follows: regard the input vector $|v\rangle$ as a vector over $\mathbb{Z}$, compute the inverse Fourier transform with respect to the circle group to get a continuous function, and then evaluate the function at p evenly spaced points to get the Fourier transform over $\mathbb{Z}_p$.

More formally, the Hilbert space l^2 is the space of complex-valued sequences $\{a_i | i \in \mathbb{Z}\}$ such that $\sum_{i\in\mathbb{Z}} |a_i|^2 < \infty$, with inner product $\langle a, b\rangle = \sum_i a_i \overline{b_i}$. The Hilbert space L^2 is the space of complex-valued functions f on $[0, 1)$ such that $\int_0^1 |f(x)|^2\, dx < \infty$, with inner product $\langle f, g\rangle = \int f(x)\overline{g(x)}\, dx$. The inverse Fourier transform is an isometry (a vector space isomorphism preserving inner products) from l^2 to L^2, defined by $f(x) = \sum_{n=-\infty}^{\infty} a_n \omega^{nx}$,

where $\{a_n\} \in l^2$. See Conway [15] for more details. Terras [47] points out the connection between the continuous and discrete case, as we are discussing here.

Map a vector $|v\rangle \in \mathbb{C}^p$ to the sequence $\{u_i\} \in l^2$, where $u_i = v_i$ for $0 \leq i < p$ and $u_i = 0$ otherwise. The inverse Fourier transform of this sequence is $f(x) = \sum_{i=0}^{p-1} v_i \omega^{ix}$, $x \in [0, 1)$. A coefficient $\hat{v}_i$ of $|\hat{v}\rangle$ is $\hat{v}_i = \frac{1}{\sqrt{p}} f(\frac{i}{p})$ and similarly for $|\hat{u}\rangle$. If q is a multiple of p, $|\hat{u}\rangle$ will contain the exact same p amplitudes of $|\hat{v}\rangle$ at multiples of $\frac{q}{p}$, up to a $\sqrt{\frac{q}{p}}$ factor. When q is not a multiple of p our goal is similar: pick p amplitudes in $|\hat{u}\rangle$ that approximate the amplitudes in $|\hat{v}\rangle$. In particular, we match $\hat{v}_i$ with $\widehat{u}_{s_i}$, for all $i \in \{0, \ldots, p-1\}$. The reason this works can be seen in Figure 5.4. Ideally, we match $\hat{v}_i$ with $\widehat{u}_{\frac{q}{p}i}$, but we must round $\frac{q}{p}i$ to the nearest integer. As q gets larger, the interval becomes more and more populated with evenly spaced points, so $\sqrt{\frac{q}{p}}\widehat{u}_{s_i} = \frac{1}{\sqrt{p}} f(\frac{i}{p} + \frac{\epsilon}{q})$ approaches $\frac{1}{\sqrt{p}} f(\frac{i}{p})$ (in the Figure, scaling $|\hat{u}\rangle$ to the width of $|\hat{v}\rangle$, $\frac{p}{q}s_i = i + \frac{p}{q}\epsilon$ approaches i).

5.2.4 Limitations of Zero-Filling

Our approach of Fourier sampling over a larger group and looking for points s_i may seem wasteful since we throw out much of the distribution, but unfortunately using all the points will not work. To see this, rather than using $\mathcal{D}^{\mathrm{M}}_{|\hat{u}\rangle}$ we will use the bucket distribution $\mathcal{D}^{\mathrm{B}}_{|\hat{u}\rangle} : \{0, \ldots p-1\} \to [0, 1]$, defined by $\mathcal{D}^{\mathrm{B}}_{|\hat{u}\rangle}(i) = \sum_{t \in T} |\widehat{u}_{s_i+t}|^2$, where $T = \{-\lfloor \frac{q}{2p} \rceil + 1, \ldots, \lfloor \frac{q}{2p} \rceil\}$. This uses all points in the distribution by counting any point in the interval $\{s_i - \lfloor \frac{q}{2p} \rceil + 1, \ldots, s_i + \lfloor \frac{q}{2p} \rceil\}$ as i (see Figure 5.3). Just by looking at Figure 5.4 it does not seem like this should work. The curve f is moving slowly between each integer point (it does for example not hit zero), so the bucket boundary probably represents some function of the two neighboring points. We now give a more formal argument.

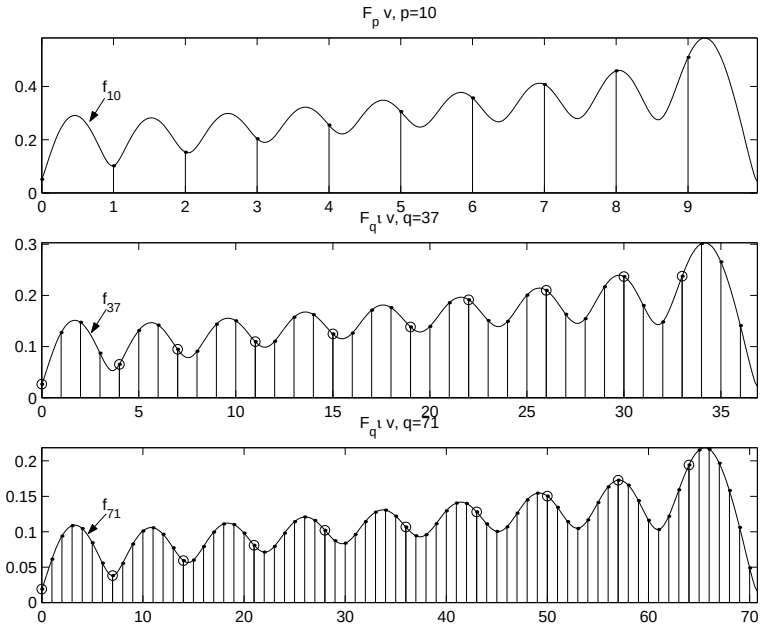

Figure 5.4: Scaled versions of f overlayed on $|\hat{v}\rangle$ (from Figure 5.1), and on $|\hat{u}\rangle$ for two different values of q. $f_p(x) = \frac{1}{\sqrt{p}} f(px)$ and $f_q(x) = \frac{1}{\sqrt{q}} f(qx)$. Points s_i are marked with circles.

We will show that this method does not even work for a delta function. Given any q and any interval $I = [a/q, b/q] \subseteq [0,1]$ with a, $b \in \mathbb{Z}$, we have $\Pr[\text{see } i \text{ s.t. } i/q \in I] = \frac{1}{q} \sum_{i=a}^{b} |f(\frac{i}{q})|^2$. We can now let the point spacing get arbitrarily small while keeping the interval the same, and we have $\lim_{n\to\infty} \frac{1}{nq} \sum_{i=na}^{nb} |f(\frac{i}{nq})|^2 = \int_{a/q}^{b/q} |f(x)|^2 \, dx$. If f is nonzero outside of I, then the integral over I is less than one since f is continuous. If $|\hat{v}\rangle = |0\rangle$, then $f(x) = \frac{1}{\sqrt{p}} \sum_{i=0}^{p-1} \omega^{ix}$, and this is only zero for multiples of $\frac{1}{p}$. So for any q, $|\mathcal{D}_{|\hat{v}\rangle}(0) - \mathcal{D}^{\mathrm{B}}_{|\hat{u}\rangle}(0)|_1 \geq c$, for c a nonzero constant.

In the next section we will take an existing algorithm and show how to convert it into a Fourier sampling theorem where this problem does not occur.

5.2.5 Proof of Theorem 5.1

Proof: (Theorem 5.1). We will use [35] to improve the theorem statement in [31] and give a simpler proof.

Recall that we need to show that $\| |\hat{v}\rangle - \sqrt{\frac{q}{p}} P |\hat{u}\rangle \| \leq \frac{\| |v\rangle \|}{8s}$. Let $|\delta\rangle = \sqrt{\frac{q}{p}} P |\hat{u}\rangle - |\hat{v}\rangle$ be the vector between the two vectors. We will start by computing a coefficient δ_j. For a given j, we can use the linearity of the Fourier transform and write $|\hat{u}\rangle = \hat{v}_j (F_q F_p^{-1} |j\rangle) + \sum_{c \neq j} \hat{v}_c (F_q F_p^{-1} |c\rangle)$. Now we only have $F_q F_p^{-1}$ of delta functions. Figure 5.5 and Lemma 5.2 describe this case.

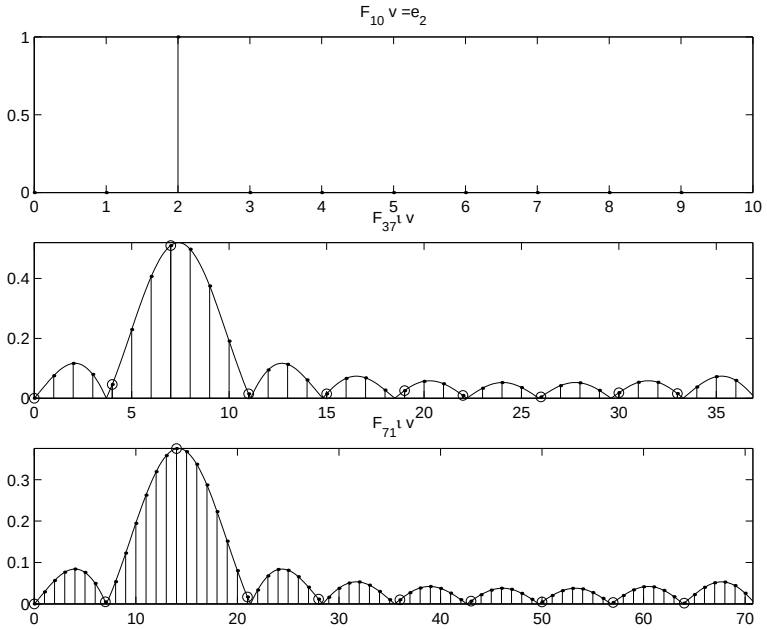

Figure 5.5: $|\hat{v}\rangle = |2\rangle$, $|\hat{u}\rangle = F_q F_p^{-1} |2\rangle$ for two values of q, the underlying continuous function, and points s_i marked with circles.

Lemma 5.2 *For a fixed $j \in \{0, \ldots, p-1\}$, let $|\hat{v}\rangle = |j\rangle$ and $|\hat{u}\rangle = F_q F_p^{-1} |\hat{v}\rangle$. Let $\epsilon = s_j - \frac{q}{p} j$.*

Then:

- $\widehat{u}_{s_j} = \sqrt{\frac{p}{q}}\frac{1}{p}\sum_{i=0}^{p-1}\omega_p^{i\epsilon\frac{p}{q}}$ *and* $Re(\widehat{u}_{s_j}) \geq \sqrt{\frac{p}{q}}\left(1 - 5\frac{p^2}{q^2}\right)$

- *For* $c \neq j$, $\widehat{u}_{s_c} = \sqrt{\frac{p}{q}}\frac{1}{p}\sum_{i=0}^{p-1}\omega_p^{i(j-c+\frac{p}{q}\epsilon)}$ *and* $|\widehat{u}_{s_c}| \leq \sqrt{\frac{p}{q}}\frac{1}{|j-c|_p}\frac{p}{q}$

$$\textit{where } |x|_p = \begin{cases} x \bmod p & \textit{if } 0 \leq x \bmod p \leq p/2 \\ -x \bmod p & \textit{otherwise} \end{cases}$$

The lemma states that when $|\hat{v}\rangle$ is a delta function, $|\hat{u}\rangle$ is relatively large at s_j and falls off as inverse distance at points $s_c,\ c \neq j$, as shown in Figure 5.5.

So for a given j, $|\hat{u}\rangle = \hat{v}_j(F_q F_p^{-1}|j\rangle) + \sum_{c\neq j}\hat{v}_c(F_q F_p^{-1}|c\rangle)$, and by Lemma 5.2 we have

$$\widehat{u}_{s_j} = \hat{v}_j\sqrt{\frac{p}{q}}\frac{1}{p}\sum_{i=0}^{p-1}\omega_p^{i\epsilon\frac{p}{q}} + \sum_{c\neq j}\sqrt{\frac{p}{q}}\frac{1}{p}\sum_{i=0}^{p-1}\omega_p^{i(j-k+\frac{p}{q}\epsilon)}.$$

Now we bound δ_j:

$$\begin{aligned} |\delta_j| &= \left|\sqrt{\frac{q}{p}}\widehat{u}_{s_j} - \hat{v}_j\right| = \left|\hat{v}_j\left(\frac{1}{p}\sum_{i=0}^{p-1}\omega_p^{i\epsilon\frac{p}{q}} - 1\right) + \sum_{c\neq j}\hat{v}_c\left(\frac{1}{p}\sum_{i=0}^{p-1}\omega_p^{i(j-c+\frac{p}{q}\epsilon)}\right)\right| \\ &\leq |\hat{v}_j|\left|\frac{1}{p}\sum_{i=0}^{p-1}\omega_p^{i\epsilon\frac{p}{q}} - 1\right| + \sum_{c\neq j}|\hat{v}_c|\left|\frac{1}{p}\sum_{i=0}^{p-1}\omega_p^{i(j-c+\frac{p}{q}\epsilon)}\right| \leq |\hat{v}_j|4\frac{p}{q} + \sum_{c\neq j}|\hat{v}_c|\frac{1}{|j-c|_p}\frac{p}{q} \end{aligned}$$

where if $\alpha = \frac{1}{p}\sum_{i=0}^{p-1}\omega_p^{i\epsilon\frac{p}{q}}$, the last inequality follows from $|1-\alpha|^2 = (1-\mathrm{Re}(\alpha))^2 + \mathrm{Im}(\alpha)^2 = 2(1-\mathrm{Re}(\alpha))$ and Lemma 5.2.

We can now bound the length of $|\delta\rangle$:

$$\begin{aligned} \||\delta\rangle\|^2 &\leq \sum_j\left|\hat{v}_j|4\frac{p}{q} + \sum_{c\neq j}|\hat{v}_c|\frac{1}{|j-c|_p}\frac{p}{q}\right|^2 \leq \frac{\||\hat{v}\rangle\|^2}{p}\sum_j\left|4\frac{p}{q} + \sum_{c\neq j}\frac{1}{|j-c|_p}\frac{p}{q}\right|^2 \\ &\leq \frac{\||\hat{v}\rangle\|^2}{p}\sum_j\left|4\frac{p}{q} + \frac{p}{q}2\ln p\right|^2 \leq \frac{\||\hat{v}\rangle\|^2}{p}\left(3\frac{p}{q}\ln p\right)^2\sum_j 1 = \left(3\frac{p}{q}\||\hat{v}\rangle\|\ln p\right)^2 \end{aligned}$$

The second inequality follows from the following observation. The left hand side can be viewed as the squared length of a matrix-vector product Mv, where the matrix and vector satisfy $M_{jj} = 4\frac{p}{q}$, $M_{jc} = \frac{1}{|j-c|_p}\frac{p}{q}$, and $v_c = |\hat{v}_c|$. An upper bound is the operator norm of the matrix times the length of $|\hat{v}\rangle$. Since the matrix is symmetric, the norm is the magnitude of the largest eigenvalue. As the matrix is circulant, the eigenvalues are all of the form $\sum_i \omega_p^{ij} M_{ij}$ for some j. Since the values of M are positive and real this is maximized when $j = 0$. Therefore the left hand side is maximized when $|\hat{v}_j| = \frac{\||\hat{v}\rangle\|}{\sqrt{p}}$ for all j. ■

We end this section by proving Lemma 5.2.

Proof: **(Lemma 5.2)** The first bound is established as follows:

$$\widehat{u}_{s_j} = \frac{1}{\sqrt{q}}\sum_{i=0}^{p-1}\frac{1}{\sqrt{p}}\omega_p^{-ij}\omega_q^{i(jq/p+\epsilon)} = \frac{1}{\sqrt{q}}\sum_{i=0}^{p-1}\frac{1}{\sqrt{p}}\omega_p^{-ij}\omega_p^{ij}\omega_q^{i\epsilon} = \sqrt{\frac{p}{q}}\frac{1}{p}\sum_{i=0}^{p-1}\omega_q^{i\epsilon}$$

Since $|\epsilon| \leq \frac{1}{2}$, for s_j we have $\mathrm{Re}\left(\frac{1}{p}\sum_{i=0}^{p-1}\omega_p^{i\epsilon p/q}\right) \geq \cos(2\pi\epsilon p/q) \geq 1 - \frac{(2\pi\epsilon p/q)^2}{2} \geq 1 - 5(p/q)^2$.

For the second bound we use the fact that the sum is a geometric series:

$$\widehat{u}_{s_k} = \sqrt{\frac{p}{q}}\frac{1}{p}\sum_{i=0}^{p-1}\omega_p^{i(k-j+\epsilon p/q)} = \sqrt{\frac{p}{q}}\frac{1}{p}\frac{1-\omega^{\epsilon\frac{p}{q}}}{1-\omega_p^{k-j+\epsilon\frac{p}{q}}}$$

In the numerator we have $\left|1-\omega^{\epsilon\frac{p}{q}}\right|^2 = \left|2\left(1-\cos\left(2\pi\epsilon\frac{p}{q}\right)\right)\right| \leq \left(2\pi\epsilon\frac{p}{q}\right)^2 \leq \left(\pi\frac{p}{q}\right)^2$. In the denominator we have:

$$\begin{aligned}
\left|1-\omega_p^{k-j+\epsilon\frac{p}{q}}\right| &= \sqrt{2\left(1-\cos\left(2\pi\left(\frac{k-j}{p}+\frac{\epsilon}{q}\right)\right)\right)} = \left|2\sin\left(\pi\left(\frac{k-j}{p}+\frac{\epsilon}{q}\right)\right)\right| \\
&\geq 2\left|\sin\left(\pi\frac{k-j}{p}\right)\right| - \frac{4\pi\epsilon}{q} \geq 2\left|\sin\left(\pi\frac{k-j}{p}\right)\right| - \frac{2\pi}{q} \\
&\geq \frac{4}{p}|k-j|_p - \frac{2\pi}{q}
\end{aligned}$$

The last inequality follows from the fact that $\sin(x)$ is at least $\frac{x}{\pi/2}$ when $x \in [0, \pi/2]$ and is at least $2 - \frac{x}{\pi/2}$ when $x \in [\pi/2, \pi]$.

So

$$\sqrt{\frac{p}{q}}\left|\frac{1}{p}\sum_{i=0}^{p-1}\omega_p^{i(k-j+\epsilon p/q)}\right| = \sqrt{\frac{p}{q}}\frac{1}{p}\frac{\left|1-\omega^{\epsilon\frac{p}{q}}\right|}{\left|1-\omega_p^{k-j+\epsilon\frac{p}{q}}\right|} \leq \sqrt{\frac{p}{q}}\frac{\pi p/q}{4|k-j|_p - \frac{\pi p}{q}} \leq \sqrt{\frac{p}{q}}\frac{1}{|k-j|_p}\frac{p}{q}$$

■

5.3 Repeating the Vector

In this section we cast results of Kitaev [36] and Cleve, Ekert, Macchiavello and Mosca [14] into our framework and derive another solution to Fourier sampling. Kitaev gives an algorithm for estimating the eigenvalue of a vector. In [14] they show that the quantum circuits for this algorithm and Shor's algorithm are the same. Using this we can take the power from Kitaev's algorithm and use it for our problem. In particular, from our point of view, the power essentially comes from taking the given superposition and repeating it. Given a superposition $|v\rangle = \sum_{i=0}^{p-1} v_i|i\rangle$, first compute $|u\rangle = R_q|v\rangle$. Then we have:

Theorem 5.2 *If* $q = \Omega(p^3/\epsilon)$ *then* $|\mathcal{D}^B_{|\hat{u}\rangle} - \mathcal{D}_{|\hat{v}\rangle}|_1 \leq \epsilon$.

This method immediately works for any abelian group by taking ϵ small enough. In the next section we will give a Fourier sampling theorem where q is smaller.

Proof: As shown in [36] and [14], given an eigenvector $\sum_{i=0}^{q-1}\omega^{-\phi i}|i\rangle$ (for the map $|i\rangle \longrightarrow |i+1\rangle$), computing the Fourier transform results in a state that is highly concentrated near the basis state $|\tilde{\phi}\rangle$, where $\tilde{\phi}$ is translated as the bits in the binary expansion of ϕ. We have

Lemma 5.3 ([39]) $\Pr[|\frac{i}{q} - \phi| \leq \frac{k}{q}] \geq 1 - \frac{1}{2k-1}$.

As usual, we are interested in understanding the delta function case $|\hat{v}\rangle = |c\rangle$ and then using the linearity of the Fourier transform. So we must examine $F_q R_q F_p^{-1}|c\rangle$, where R_q repeats the state up to q. We now apply the theorem with $\phi = c/p$ and $k = \frac{q}{2p}$, and we have $\mathcal{D}^{\mathrm{B}}_{|c\rangle}(c) \geq 1 - \frac{p}{q}$, when $q \geq 2p$. We now use the linearity of the Fourier transform to show the same holds for any $|\hat{v}\rangle$. First examine the value of $\mathcal{D}_{|\hat{u}\rangle}(c)$ for some c. It has probability at least $1 - p/q$ minus the p/q that can interfere from each of the other $p-1$ delta functions. So there is at most p^2/q probability loss at each delta function, and we have $\sum_i |\mathcal{D}^{\mathrm{B}}_{|\hat{u}\rangle}(i) - \mathcal{D}_{|\hat{v}\rangle}(i)| \leq p^3/q$, and choosing $q = \Omega(p^3/\epsilon)$ proves the theorem. ■

Repeating the given vector first in this way causes the superposition to be much more concentrated in the buckets. In fact, when q is an exact multiple of p it is easy to see the exact effect by looking at the underlying continuous function. In particular, if the original superposition is $\sum_{i=0}^{p-1} v_i|i\rangle$, then repeating first results in the state $\frac{1}{\sqrt{s}}\sum_{j=0}^{s-1}\sum_{i=0}^{p-1} v_i|jp+i\rangle$, and the Fourier transform has the associated continuous function

$$f(x) = \frac{1}{\sqrt{s}}\sum_{j=0}^{s-1}\sum_{i=0}^{p-1}\hat{v}_i\omega^{(jp+i)x} = \left(\frac{1}{\sqrt{s}}\sum_{j=0}^{s-1}\omega^{jpx}\right)\left(\sum_{i=0}^{p-1}\hat{v}_i\omega^{ix}\right).$$

Notice that the right factor is exactly the continuous function for the Fourier transform of the original state $|v\rangle$, and the effect of repeating can be factored out into the left scaling factor $g(x) = (\frac{1}{\sqrt{s}}\sum_{j=0}^{s-1}\omega^{jpx})$. Figure 5.6 shows how the functions $g(x)$ works for different values of s. Note that if x is a multiple of $1/p$, then $g(x) = \sqrt{s}$, and if x is not a multiple of $1/p$ but is a multiple of $1/ps$, then $g(x) = 0$.

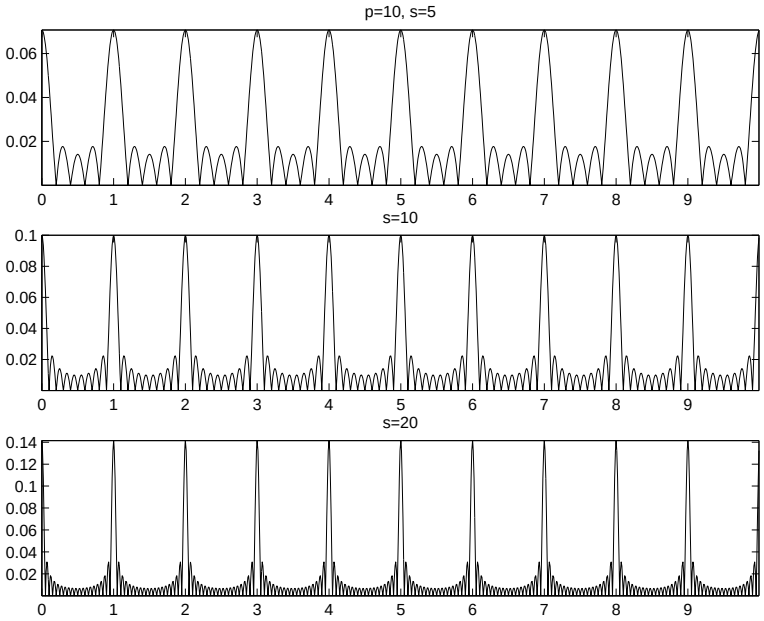

Figure 5.6: A graph of the function $g(x) = \frac{1}{\sqrt{s}} \sum_{i=0}^{s-1} \omega^{px}$.

5.4 Repeating and Zero-Filling

5.4.1 The Main Theorem

In this section we will combine zero-filling and repeating the vector to get a better Fourier transform algorithm and a better Fourier sampling theorem. The main theorem of this section may be regarded as showing that the diagram

$$\begin{array}{ccc} \mathbb{C}^p & \xrightarrow{F_p} & \mathbb{C}^p \\ \downarrow R_{ps} & & \downarrow S \\ \mathbb{C}^{ps+r} & \xrightarrow{F_q} & \mathbb{C}^q \end{array}$$

"commutes approximately" when s is sufficiently large compared to p, and $q = ps + r$ is sufficiently large compared to ps, where s is the number of copies of $\mathbb{C}^p$: $\mathbb{C}^{ps} = \bigoplus_{i=0}^{s-1} \mathbb{C}^p$,

$R_{ps} = \sum_{i=0}^{s-1} \iota_i$ is the repeat map, where ι_i is the inclusion map of $\mathbb{C}^p$ into the i^{th} copy of $\mathbb{C}^p$ in $\mathbb{C}^{ps}$, and S is a map similar to s, and will be described below. By "commutes approximately" we mean that for all $v \in \mathbb{C}^p$, $|SF_p|v\rangle - F_q R_{ps}|v\rangle| \leq \epsilon$. We will then apply this to the quantum situation to achieve a new Fourier transform algorithm and a new Fourier sampling theorem.

Define $T = \{-\lfloor q/2p \rceil + 1, \ldots, \lfloor q/2p \rceil\}$ to index the consecutive basis vectors, and define the map $S : \mathbb{C}^p \to \mathbb{C}^q$ by $S|i\rangle = \sum_{t \in T} c_t |s_i + t\rangle$. We will spell out the normalization constants c_t in the body of the proof – for now, we just point out that these constants are highly concentrated at the center of the interval.

Notice what this map does on an arbitrary vector: we have

$$\begin{aligned} S|\hat{v}\rangle &= \sum_{i=0}^{p-1} \hat{v}_i S|i\rangle = \sum_{i=0}^{p-1} \hat{v}_i \sum_{t \in T} c_t |s_i + t\rangle \\ &= \sum_{t \in T} c_t \left(\sum_{i=0}^{p-1} \hat{v}_i |s_i + t\rangle \right) \end{aligned}$$

So $S|\hat{v}\rangle$ contains $|T|$ copies of $|\hat{v}\rangle$, each weighted by some c_t. Furthermore, each copy is more or less evenly spread out in the interval $\{0, \ldots, q-1\}$. We can now state the main theorem of this section.

Theorem 5.3 *For any* $|v\rangle \in \mathbb{C}^p$,

$$\| F_q R_{ps} |v\rangle - S F_p |v\rangle \| \leq \left(\frac{4ps}{q} + \frac{4 \ln p}{\sqrt{s}} \right) \| |v\rangle \|$$

We can now apply this to the quantum case to get an algorithm for the Fourier transform.

Algorithm 5.1 (Approximate Fourier Transform)

Input: $|v\rangle = \sum_{i=0}^{p-1} v_i|i\rangle$

Output: *An approximation of* $|\hat{v}\rangle = F_p|v\rangle$.

1. *Repeat the state s times:*

$$|v\rangle \xrightarrow{F_s} \frac{1}{\sqrt{s}} \sum_{j=0}^{s-1} |j\rangle \sum_{i=0}^{p-1} v_i|i\rangle \longrightarrow \frac{1}{\sqrt{s}} \sum_{j=0}^{s-1} \sum_{i=0}^{p-1} v_i|jp+i\rangle$$

2. *Compute the Fourier transform over* $\mathbb{Z}_q$*:*

$$\frac{1}{\sqrt{s}} \sum_{j=0}^{s-1} \sum_{i=0}^{p-1} v_i|jp+i\rangle \xrightarrow{F_q} \frac{1}{\sqrt{qs}} \sum_{d=0}^{q-1} \left(\sum_{j=0}^{s-1} \sum_{i=0}^{p-1} v_i \omega_q^{(jp+i)d} \right) |d\rangle$$

3. *Identify which copy of* $|\hat{v}\rangle$ *each d is in:*

$$|d\rangle \longrightarrow |c,t\rangle,$$

where $d = \lfloor c\frac{q}{p} \rceil + t$. *Call the resulting superposition* $\sum_{t\in T} \sqrt{p_t}|\widehat{u}_t\rangle|t\rangle$, *where* p_t *is the probability of measuring t.*

4. *Measure the second register resulting in* $|\widehat{u}_t\rangle|t\rangle$ *with probability* p_t.

Corollary 5.2 *Let* $|\widehat{u}_t\rangle$ *be as in Algorithm 5.1. For any* $\epsilon > 0$, *if* $s = \Omega\left(\frac{\log^2 p}{\epsilon^4}\right)$ *and* $q = \Omega\left(\frac{ps}{\epsilon^2}\right)$, *then with probability* $1 - \epsilon^2$ *over choices of t,* $\||\hat{v}\rangle - |\widehat{u}_t\rangle\| \leq \epsilon$.

Proof: We apply a standard averaging argument to conclude the bound. By Theorem 5.3 we have $\sum_t \|\sqrt{p_t}|\widehat{u}_t\rangle|t\rangle - c_t|\hat{v}\rangle|t\rangle\|^2 \leq \epsilon^4$. Let B denote the set of bad t, where $t \in B$ if $\||\widehat{u}_t\rangle - c_t/\sqrt{p_t}|\hat{v}\rangle\|^2 > \epsilon^2$. Suppose the probability of getting a bad t is at least ϵ^2. Then we have $\sum_{t\in B} \|\sqrt{p_t}|\widehat{u}_t\rangle|t\rangle - c_t|\hat{v}\rangle|t\rangle\|^2 = \sum_{t\in B} p_t\||\widehat{u}_t\rangle - c_t/\sqrt{p_t}|\hat{v}\rangle\|^2 > \epsilon^2 \sum_{t\in B} p_t > \epsilon^4$, contradicting Theorem 5.3.

Now assume we are in the good case when $\||\widehat{u}_t\rangle - c_t/\sqrt{p_t}|\hat{v}\rangle\|^2 \leq \epsilon^2$. Then by the triangle inequality we have $\||\widehat{u}_t\rangle - |\hat{v}\rangle\| \leq \||\widehat{u}_t\rangle - c_t/\sqrt{p_t}|\hat{v}\rangle\| + \|c_t/\sqrt{p_t}|\hat{v}\rangle - |\hat{v}\rangle\| \leq \epsilon + |1 - c_t/\sqrt{p_t}| \leq 2\epsilon$. The bound on $1 - c_t/\sqrt{p_t}$ follows from the fact that $c_t/\sqrt{p_t}|\hat{v}\rangle$ is at most ϵ away from the unit vector $|\widehat{u}_t\rangle$. ■

Corollary 5.3 *For any ϵ and positive integer p, let $n = \log p$. Then Algorithm 5.1 ϵ-approximates the quantum Fourier transform over $\mathbb{Z}_p$ and runs in time $O(n \log n \log\log n + \log^2 \frac{1}{\epsilon})$.*

The previous algorithm [36] for computing the Fourier transform for an arbitrary p takes time $O(n^2)$, so Algorithm 5.1 is strictly faster for polynomial approximations.

Proof: Coppersmith [16] gives an algorithm to ϵ-approximate the Fourier transform over $\mathbb{Z}_{2^n}$ in time $O(n \log(\frac{n}{\epsilon}))$, so choose s and q to be powers of two. Step 3 requires multiplying two n bit numbers which takes time $n \log n \log\log n$. ■

5.4.2 A Better Fourier Sampling Theorem

Theorem 5.3 can also be applied to achieve a better Fourier sampling theorem than in the last section.

Algorithm 5.2 (Fourier Sampling, p known)

Input: $|v\rangle = \sum_{i=0}^{p-1} v_i |i\rangle$

Output: *A sample from a distribution that closely approximates $\mathcal{D}_{|\hat{v}\rangle}$*

1. *Repeat the state s times:*

$$|v\rangle \longrightarrow \frac{1}{\sqrt{s}} \sum_{j=0}^{s-1} \sum_{i=0}^{p-1} v_i |jp + i\rangle$$

2. Compute the Fourier transform over $\mathbb{Z}_q$:

$$\frac{1}{\sqrt{s}}\sum_{j=0}^{s-1}\sum_{i=0}^{p-1} v_i|jp+i\rangle \xrightarrow{F_q} \frac{1}{\sqrt{qs}}\sum_{d=0}^{q-1}\left(\sum_{j=0}^{s-1}\sum_{i=0}^{p-1} v_i\omega_q^{(jp+i)d}\right)|d\rangle \stackrel{def}{=} |\hat{u}\rangle$$

3. Measure.

4. Round according to $\mathcal{D}^B_{|\hat{u}\rangle}$.

Recall from the last chapter that $\mathcal{D}_{|\hat{v}\rangle}$ is the distribution induced by measuring $|\hat{v}\rangle$, and the bucket distribution $\mathcal{D}^{\mathrm{B}}_{|\hat{u}\rangle} : \{0,\ldots p-1\} \to [0,1]$ is defined by $\mathcal{D}^{\mathrm{B}}_{|\hat{u}\rangle}(i) = \sum_{t\in T}|\widehat{u}_{s_i+t}|^2$, where $T = \{-\lfloor\frac{q}{2p}\rceil+1,\ldots,\lfloor\frac{q}{2p}\rceil\}$. This uses all points in the distribution by counting any point in the interval $\{s_i - \lfloor\frac{q}{2p}\rceil+1,\ldots,s_i+\lfloor\frac{q}{2p}\rceil\}$ as i (see Figure 5.3).

Corollary 5.4 *Let $|\hat{u}\rangle$ be as in Algorithm 5.2. If $s = \Omega\left(\frac{\log^2 p}{\epsilon^2}\right)$ and $q = \Omega\left(\frac{ps}{\epsilon}\right)$ then*

$$|\mathcal{D}_{|\hat{v}\rangle} - \mathcal{D}^B_{|\hat{u}\rangle}|_1 \leq \epsilon.$$

Proof: Follows from Lemma 5.1 and Theorem 5.3. ∎

In some of the most interesting quantum algorithms which use Fourier sampling neither p nor $|v\rangle$ are known. However, it is often possible to generate a superposition of a chosen length q that consists of repetitions of $|v\rangle$, where q is not necessarily a multiple of p. That is, for an arbitrary integer q we assume we are given the superposition

$$c\sum_{i=0}^{l-1} v_{(i \bmod p)}|i\rangle,$$

where c is a normalization constant and $\lceil l/p\rceil = s$. Given this state as input, we will try to approximate $\mathcal{D}_{|\hat{v}\rangle}$ as in Algorithm 5.2. This is not exactly possible though because we use the continued fractions algorithm, and that algorithm returns fractions in reduced terms. We will now define the new distribution we will be close to.

Let p be the unknown period of $|v\rangle$. Define the reduced fractions distribution $\mathcal{D}^{\mathrm{RF}}_{|\hat{v}\rangle}$: $\{0,\ldots,p-1\}^2 \to [0,1]$ by $\mathcal{D}^{\mathrm{RF}}_{|\hat{v}\rangle}(c,r) = \mathcal{D}_{|\hat{v}\rangle}(cl)$ if $rl = p$. This distribution can be thought of as sampling some i from $\mathcal{D}_{|\hat{v}\rangle}$ and making c/r the output, where c/r $(= i/p)$ is i/p in reduced terms. For example, if $p = 6$ and $\mathcal{D}_{|\hat{v}\rangle}$ is the uniform distribution over $\{0,1,2,3,4,5\}$, then $\mathcal{D}^{\mathrm{RF}}_{|\hat{v}\rangle}$ is a uniform distribution over $\{0,(1,6),(1,3),(1,2),(2,3),(5,6)\}$. The distributions are basically the same, except some information is lost because the fractions are reduced. Of course, if p is prime, then nothing reduces and the distributions are just the same information-wise. The following algorithm approximates $\mathcal{D}^{\mathrm{RF}}_{|\hat{v}\rangle}$.

Algorithm 5.3 (Fourier Sampling, p unknown)

Input: $|v\rangle = \sum_{i=0}^{q-1} v_{(i \bmod p)}|i\rangle$, *an upper bound* t *on* p

Output: *A sample from a distribution that closely approximates* $\mathcal{D}^{RF}_{|\hat{v}\rangle}$

1. *Compute the Fourier transform over* $\mathbb{Z}_q$*:*

$$\sum_{i=0}^{l-1} v_{(i \bmod p)}|i\rangle \xrightarrow{F_q} \frac{1}{\sqrt{q}} \sum_{d=0}^{q-1} \left(\sum_{i=0}^{l-1} v_{(i \bmod p)} \omega_q^{id} \right) |d\rangle \stackrel{def}{=} |\hat{u}\rangle$$

2. *Measure, see* d*.*

3. *Use continued fractions on* d/q *and get* c/r*, which is* i/p *in reduced terms.*

Given an upper bound t on p, let the continued fractions distribution $\mathcal{D}^{\mathrm{CF}}_{|\hat{u}\rangle}$: $\{0,\ldots,t-1\}^2 \to [0,1]$ be defined by $\mathcal{D}^{\mathrm{CF}}_{|\hat{u}\rangle}(c,r) = \Pr[c/r$ is output by Algorithm 5.3$]$. The continued fractions algorithm used in Step 3 takes integers q and d with $d < q$, and a guarantee that $|d - c\frac{q}{p}| \leq \frac{q}{2t^2}$. In return it gives the unique fraction $\frac{i}{p}$, where $p < t$, in reduced terms in polynomial time.

Corollary 5.5 *Let $|\hat{u}\rangle$ be as in Algorithm 5.3 and let $\lceil l/p \rceil = s$ be the number of repetitions. If $s = \Omega\left(\frac{t^2}{\epsilon^2 p}\right)$ and $q = \Omega\left(\frac{ps}{\epsilon}\right)$, then*

$$|\mathcal{D}^{RF}_{|\hat{v}\rangle} - \mathcal{D}^{CF}_{|\hat{u}\rangle}|_1 \leq \epsilon.$$

We may get a bogus value back from the continued fractions algorithm if the guarantee is not met, but notice that this corollary implies that the probability of this happening is less than ϵ.

Proof: The notation here is defined at the beginning of the proof of Theorem 5.3. In addition let $|\delta_0^{\mathrm{IB}}\rangle$ be the inner bucket, i.e. $|\delta_0^{\mathrm{B}}\rangle$ restricted to integers between $-q/2t^2$ and $q/2t^2$. We will show that $\||\delta_0^{\mathrm{B}}\rangle - |\delta_0^{\mathrm{IB}}\rangle\|^2 \leq \frac{t^2}{ps}$, i.e. almost all the weight is close to the midpoint of the bucket. Then using Theorem 5.3 and the triangle inequality we have $\||\hat{u}\rangle - \sum_{i=0}^{p-1} \hat{v}_i |\delta_{0,i}^{\mathrm{IB}}\rangle\| \leq \||\hat{u}\rangle - \sum_{i=0}^{p-1} \hat{v}_i |\delta_{0,i}^{\mathrm{B}}\rangle\| + \|\sum_{i=0}^{p-1} \hat{v}_i |\delta_{0,i}^{\mathrm{B}}\rangle - \sum_{i=0}^{p-1} \hat{v}_i |\delta_{0,i}^{\mathrm{IB}}\rangle\| \leq \frac{4ps}{q} + \frac{4 \ln p}{\sqrt{s}} + \frac{t}{\sqrt{ps}}$. The corollary then follows from Lemma 5.1.

There is one technical point, which is that $|\hat{u}\rangle$ is not exactly as stated in Theorem 5.3 because it is not necessarily an exact multiple of p. This does not change the bound, though. The distance between $|\hat{u}\rangle$ and the same state repeated to the next multiple of p is at most $1/\sqrt{s}$, and adding a $1/\sqrt{s}$ to the sum $\frac{4ps}{q} + \frac{4 \ln p}{\sqrt{s}} + \frac{t}{\sqrt{ps}}$ does not change the bound.

We now bound $\||\delta_0^{\mathrm{B}}\rangle - |\delta_0^{\mathrm{IB}}\rangle\|$. Let $a = \lceil q/2t^2 \rceil$ and $b = \lfloor q/2p \rfloor$. By Equation 5.5 we have

$$\||\delta_0^{\mathrm{B}}\rangle - |\delta_0^{\mathrm{IB}}\rangle\|^2 \leq 2 \sum_{i=a}^{b} |(|\delta_0\rangle)_i|^2 \leq \sum_{i=a}^{b} \frac{q}{2ps} \frac{1}{i^2} \leq \frac{q}{2ps} \int_a^{b+1} \frac{1}{x^2}\, dx \leq \frac{q}{2ps} \frac{2t^2}{q} \leq \frac{t^2}{ps},$$

so $\||\delta_0^{\mathrm{B}}\rangle - |\delta_0^{\mathrm{IB}}\rangle\| \leq \frac{t}{\sqrt{ps}}$. ∎

Notice that this is easily applied to the factoring algorithm. A function is given with period p along with an upper bound $t = N$, the number trying to be factored. In this

case some order r of a periodic function must be found. Since this corresponds to $|v\rangle = |c\rangle$ for some $c \in \{0, \ldots, r-1\}$, the distribution $\mathcal{D}_{|\hat{v}\rangle}$ will be uniform. Using Algorithm 5.3 allows us to sample from $\mathcal{D}^{\mathrm{CF}}_{|\hat{u}\rangle}$, which is ϵ-close to $\mathcal{D}^{\mathrm{RF}}_{|\hat{v}\rangle}$. As soon as we measure some i relatively prime to r, we are done. We will give another application in Chapter 7.

5.4.3 Application: Many-to-One Periodic Functions

We will now give an application of the Fourier sampling theorem due to Hales [32]. We wish to generalize the period finding example by allowing a function to be many-to-one in its fundamental domain. That is, if the function has period p, it is not necessary that the function have distinct values in the domain $\{0, \ldots, p-1\}$. By Corollary 5.5, we know that a starting point is to understand the distribution $\mathcal{D}^{\mathrm{RF}}_{|\hat{v}\rangle}$, however as we will see in the algorithm below, there will be one additional complication.

We will first define a class of functions $C_{1/d(n)}$ that contains a function f if one must change at least a $1/d(n)$ fraction of its values to reduce its period. Such an f can be viewed as being $1/d(n)$ robust with respect to its period. We will then show that for a given polynomial $d(n)$ there is an efficient algorithm finding the period of any $f \in C_{1/d(n)}$.

More formally, let $f : \mathbb{Z} \to \{0, \ldots, 2^n - 1\}$ and $g : \mathbb{Z} \to \{0, \ldots 2^n - 1\}$ be functions with periods p_f and p_g, respectively, each of length at most 2^n. Define $D(f, g)$ to be the fraction of points in $\{0, \ldots, p_f p_g - 1\}$ where f and g differ. Let $C_{1/d(n)} = \{f | \forall g \text{ with } p_g < p_f, D(f, g) > 1/d(n)\}$.

Algorithm 5.4 (Many-to-One Period Finding)

Input: f

Output: *The period p of f.*

1. *Repeat the following $k = poly(\log p)$ times:*

   (a) *Create the periodic state, measuring the second register:*

   $$\frac{1}{\sqrt{q}} \sum_{i=0}^{q-1} |i, f(i)\rangle \longrightarrow \frac{1}{\sqrt{N}} \sum_{i:f(i)=a} |i\rangle |a\rangle$$

   (b) *Run Algorithm 5.3 to get a value c_i/r_i.*

2. *Output the least common multiple of the denominators $r_1, \ldots, r_k$.*

By the definition of f and g we have an upper bound on the period as required in Algorithm 5.3. Define the skewed distribution $\mathcal{D}_f^S : \{0, \ldots, p-1\}^2 \to [0,1]$ by $\mathcal{D}_f^S(c,r) = \Pr[c/r$ is output by Algorithm 5.4]. This distribution is similar to $\mathcal{D}_{|\hat{u}\rangle}^{\mathrm{CF}}$, except that the repeated version of $|v\rangle$ changes depending on the outcome of Step 1a.

Theorem 5.4 *Let $\epsilon = 1/d^3(n)$, $s = \Omega(\frac{2^{2n}}{p\epsilon^2})$, $q = \Omega(\frac{ps}{\epsilon})$, and by $k = 20d^2(n)\log n$. Then Algorithm 5.4 outputs the period of f with probability at least $3/4$.*

Proof: The main problem is that we do not have a uniform superposition over $\{0, \ldots, p-1\}$, so we may never get a value i/p with $\gcd(i,p) = 1$ (i.e. we may never get a value $r = p$). However, even though each denominator returned might never be p, the lcm of the denominators will be. The only way the lcm is missing some factor l of p is if every value returned is of the form cl/rl, where $rl = p$. The next lemma says that this does not happen with high probability.

Lemma 5.4 ([32]) *Suppose that $f \in C_{1/d(n)}$ has period p. Then for every prime l dividing p,*

$$\Pr[l \text{ divides } j] < 1 - \frac{1}{8d^2(n)},$$

where the j are distributed according to $\mathcal{D}_f^S$.

By Corollary 5.3 and by choosing $\epsilon = 1/d^3(n)$ we know that with probability at least $1 - t\epsilon \geq 1 - 20/d(n)$ the condition for the continued fraction algorithm is met, so all of the fractions have a valid form (i.e. r_i divides p). By Lemma 5.4, $\Pr[l \text{ divides } j] \leq 1 - 1/8d^2(n) + \epsilon \leq 1 - 1/9d^2(n)$. The probability that some prime divisor l of p occurs in every fraction i/p measured is at most $n(1 - 1/9d^2(n))^k \leq 1/7$, since there are at most n different possible primes dividing p. The lcm of the denominators therefore returns p with probability at least $3/4$. ■

The probability of success can be amplified because the answer can be tested. This restriction to functions in $C_{1/d(n)}$, for $d(n)$ a polynomial is also necessary:

Theorem 5.5 ([32]) *Let $d(n) = o(2^n)$ be given. Suppose that A is a quantum algorithm which correctly computes the period of any $f \in C_{1/d(n)}$ with probability at least $3/4$. Then A has worst-case run-time $\Omega(\sqrt[4]{d(n)})$.*

We finish this chapter by proving Theorem 5.3.

5.4.4 Proof of Theorem 5.3

Proof: We wish to show that $\|S|\hat{v}\rangle - |\hat{u}\rangle\| \leq \epsilon$. Recall that $|\hat{u}\rangle = F_q R_{ps}|v\rangle$, so $|\hat{u}\rangle = F_q R_{ps} F_p^{-1}|\hat{v}\rangle$. We will use this fact to relate the two vectors. First we examine what

happens when $|\hat{v}\rangle$ is some delta function $|i\rangle$, and then we combine the results using the linearity of the maps. We will define S in terms of the case when $|\hat{v}\rangle = |0\rangle$.

Let $|\delta_i\rangle = F_q R_{ps} F_p^{-1}|i\rangle$, let the bucket part of the vector be $|\delta_i^{\mathrm{B}}\rangle = \sum_{t\in T} \widehat{u}_{s_i+t}|s_i + t\rangle$, let the tail part of the vector be $|\delta_i^{\mathrm{T}}\rangle = \sum_{t\notin T} \widehat{u}_{s_i+t}|s_i + t\rangle$, which implies that $|\delta_i\rangle = |\delta_i^{\mathrm{B}}\rangle + |\delta_i^{\mathrm{T}}\rangle$. Define $|\delta_{0,i}^{\mathrm{B}}\rangle$ to be $|\delta_0^{\mathrm{B}}\rangle$ shifted to s_i, and define S by $S|i\rangle = |\delta_{0,i}^{\mathrm{B}}\rangle$. Writing the two vectors we are interested in in terms of these functions we have $S|\hat{v}\rangle = \sum_{c=0}^{p-1} \hat{v}_c S|c\rangle = \sum_{c=0}^{p-1} \hat{v}_c|\delta_{0,c}^{\mathrm{B}}\rangle$ and $|\hat{u}\rangle = \sum_{c=0}^{p-1} \hat{v}_c F_q R_{ps} F_p^{-1}|c\rangle = \sum_{c=0}^{p-1} \hat{v}_c|\delta_c\rangle = \sum_{c=0}^{p-1} \hat{v}_c|\delta_c^{\mathrm{B}}\rangle + \sum_{c=0}^{p-1} \hat{v}_c|\delta_c^{\mathrm{T}}\rangle$. Using these we have:

$$\begin{aligned}
\|S|\hat{v}\rangle - |\hat{u}\rangle\| &= \left\| \sum_{c=0}^{p-1} \hat{v}_c(|\delta_{0,c}^{\mathrm{B}}\rangle - |\delta_c^{\mathrm{B}}\rangle) - \sum_{c=0}^{p-1} \hat{v}_c|\delta_c^{\mathrm{T}}\rangle \right\| && (5.1)\\
&\leq \left\| \sum_{c=0}^{p-1} \hat{v}_c(|\delta_{0,c}^{\mathrm{B}}\rangle - |\delta_c^{\mathrm{B}}\rangle) \right\| + \left\| \sum_{c=0}^{p-1} \hat{v}_c|\delta_c^{\mathrm{T}}\rangle \right\| && (5.2)\\
&\leq \frac{4ps}{q}\||v\rangle\| + \frac{4\ln p}{\sqrt{s}}\||v\rangle\|. && (5.3)\\
& && (5.4)
\end{aligned}$$

The left summand of Equation 5.3 follows from Lemma 5.5 and the fact that the vectors in the set $\{|\delta_c^{\mathrm{B}}\rangle\}$ are pairwise orthogonal. The right summand of Equation 5.3 follows from Lemma 5.6.

The following lemma expresses the fact that excluding the tails, the vectors are close.

Lemma 5.5 $\|(|\delta_{0,c}^{B}\rangle - |\delta_c^{B}\rangle)\| \leq \frac{4ps}{q}$

Proof: We will give an upper bound for $\||\delta_{0,c}\rangle - |\delta_c\rangle\|$, which is at least $\||\delta_{0,c}^{\mathrm{B}}\rangle - |\delta_c^{\mathrm{B}}\rangle\|$. For convenience we will compute the difference in the Fourier basis.

$$\begin{aligned}\||\delta_{0,c}\rangle - |\delta_c\rangle\|^2 &= \left\|F_q^{-1}|\delta_{0,c}\rangle - F_q^{-1}|\delta_c\rangle\right\|^2 = \frac{1}{ps}\sum_{i=0}^{ps-1}\left|(\omega_q^{-icq/p-i\epsilon} - \omega_p^{-ic})\right|^2 \\ &= \frac{1}{ps}\sum_{i=0}^{ps-1}\left|\omega_p^{-ic}(\omega_q^{-i\epsilon}-1)\right|^2 = \frac{1}{ps}\sum_{i=0}^{ps-1}\left|1-\omega_q^{-i\epsilon}\right|^2 \le \left|1-\omega_q^{ps/2}\right|^2 \le \left(\frac{4ps}{q}\right)^2\end{aligned}$$

■

The following lemma expresses the fact that the tails are small and so will not interfere too much with the bound in the theorem.

Lemma 5.6 $\left\|\sum_{c=0}^{p-1}\hat{v}_c|\delta_c^T\rangle\right\| \le \frac{4\ln p}{\sqrt{s}}\||v\rangle\|$

Proof: Recall that

$$|\delta_c\rangle = F_q R_{ps} F_p^{-1}|c\rangle = \frac{1}{\sqrt{qps}}\sum_{d=0}^{q-1}\sum_{i=0}^{ps-1}\omega^{i(\frac{d}{q}-\frac{c}{p})}|d\rangle.$$

We have

$$|(|\delta_j\rangle)_d| = \left|\frac{1}{\sqrt{qps}}\sum_{i=0}^{ps-1}\omega^{i(\frac{d}{q}-\frac{j}{p})}\right| \le \frac{1}{\sqrt{qps}}\left|\frac{1-\omega^{psd/q}}{1-\omega^{(\frac{d}{q}-\frac{j}{p})}}\right| \le \sqrt{\frac{q}{ps}}\frac{2}{4|d-j\frac{q}{p}|_q}, \qquad (5.5)$$

since $|1-\omega^{(\frac{d}{q}-\frac{j}{p})}| = 2|\sin(\pi\frac{1}{q}(d-j\frac{q}{p}))| \ge \frac{4}{q}|d-j\frac{q}{p}|_q$, which can be seen using the half angle formula and the fact that $\sin(x) \ge 2x/\pi$, when $x \in [0, 1/2]$.

Expanding the expression we have

$$\left\|\sum_{c=0}^{p-1}\hat{v}_c|\delta_c^{\mathrm{T}}\rangle\right\|^2 = \sum_{d=0}^{q-1}\left|\sum_{\substack{j=0\\ j\neq\lfloor dp/q\rceil}}^{p-1}\hat{v}_j(|\delta_j^{\mathrm{T}}\rangle)_d\right|^2 \le \frac{q}{4ps}\sum_{d=0}^{q-1}\left|\sum_{\substack{j=0\\ j\neq\lfloor dp/q\rceil}}^{p-1}|\hat{v}_j|\frac{1}{|d-j\frac{q}{p}|_q}\right|^2$$

We apply Lemma 5.7 to pull $\hat{v}_i$ out of the sum. The bound on the inner sum follows from the fact that $|d-jq/p|_q \ge q/2p$. We have

$$\sum_{\substack{j=0\\ \neq\lfloor dp/q\rceil}}^{p-1}\frac{1}{|d-j\frac{q}{p}|_q} \le 2\sum_{j=0}^{p-1}\frac{1}{j\frac{q}{p}+\frac{q}{2p}} \le \frac{2p}{q}\sum_{j=0}^{p-1}\frac{1}{j+\frac{1}{2}} \le \frac{2p}{q}\sum_{j=1}^{p}\frac{1}{j} \le \frac{2p}{q}(\ln p+1) \le \frac{4p\ln p}{q}$$

Therefore,

$$\left\| \sum_{c=0}^{p-1} \hat{v}_c |\delta_c^{\mathrm{T}}\rangle \right\|^2 \leq \frac{q}{p^2 s} \||v\rangle\|^2 \sum_{d=0}^{q-1} \left| \frac{4p \ln p}{q} \right|^2 \leq \frac{16 \ln^2 p}{s} \||v\rangle\|^2$$

■

Lemma 5.7 *For any vector* $x \in \mathbb{R}^p_{\geq 0}$,

$$\sum_{d=0}^{q-1} \left(\sum_{\substack{i=0 \\ i \neq \lfloor dp/q \rceil}}^{p-1} \frac{x_i}{|d - \frac{q}{p} i|_q} \right)^2 \leq \frac{4\|x\|^2}{p} \sum_{d=0}^{q-1} \left(\sum_{\substack{i=0 \\ i \neq \lfloor dp/q \rceil}}^{p-1} \frac{1}{|d - \frac{q}{p} i|_q} \right)^2$$

Proof: Let C be the matrix such that $\|Cx\|^2 = \sum_{d=0}^{q-1} \left(\sum_{i=0, i \neq \lfloor dp/q \rceil}^{p-1} \frac{x_i}{|d - \frac{q}{p} i|_q} \right)^2$, i.e. $C_{d,i} = \frac{1}{|d - \frac{q}{p} i|_q}$. Our goal is to find the vector x_0 that maximizes the sum and factor it out. If C was circulant, we could use the usual argument (which we will use below) and conclude that the maximizing vector is parallel to the vector x where $x_i = 1$ for all i. The matrix C is close to being circulant, however, and we will choose another matrix M and use it to achieve the bound.

We will compare $\|Cx\|^2$ with the sum

$$\|Mx\|^2 \stackrel{\text{def}}{=} \sum_{t=0}^{\lceil q/p \rceil - 1} \|M_t x\|^2 = \sum_{t=0}^{\lceil q/p \rceil - 1} \sum_{k=0}^{p-1} \left(\sum_{\substack{i=0 \\ i \neq \lfloor dp/q \rceil}}^{p-1} \frac{x_i}{|\frac{q}{p} k + t - \frac{q}{p} i|_q} \right)^2$$

First we will show that the vector parallel to the vector satisfying $x_i = 1$ for all i maximizes the length of the vector Mx. Notice that for a fixed t, M_t has entries $M_{t,k,i} = \frac{1}{|t + \frac{q}{p}(k-i)|_q}$ and is circulant, i.e., $M_{t,k,i} = M_{t,k+1,i+1}$. As the matrix is circulant, all its eigenvalues are of the form $\sum_i \omega_p^{ij} M_{t,k,i}$ for some j. Since the values of M_t are positive and real this is maximized when $j = 0$. Therefore Mx is maximized when $|x_j| = \frac{\|x\|}{\sqrt{p}}$ for all j.

Now observe that $\frac{1}{2} \leq \frac{\|Cx\|^2}{\|Mx\|^2} \leq 2$. This follows from the fact that for some sequence we have $\sum_i \frac{1}{2l_i} \leq \sum_i \frac{1}{l_i+1}$, which implies $\sum_i \frac{1}{l_i} \leq 2 \sum_i \frac{1}{l_i+1}$.

Let x_0 maximize $\|Mx\|^2$ over all vectors x. Then for any x we have $\|Cx\|^2 \leq 2\|Mx\|^2 \leq 2\|Mx_0\|^2 \leq 4\|Cx_0\|^2$. ■

■