Quantum Computing 4/14/2022

Lecture 21: Phase Estimation

Instructor: Dieter van Melkebeek Scribe: Rishabh Khandelwal

In this lecture, we talk about the quantum Fourier transform and phase estimation, which is a precursor to eigenvalue estimation, a key ingredient in the algorithm we will later develop for finding the order of an integer modulo another integer. We start by discussing Quantum Fourier Transform for the the group $\mathbb{Z}_N, +$ (integers modulo N), where $N = 2^n$. Then we talk about phase estimation - the problem statement, algorithm, and its analysis.

1 Quantum Fourier transform over $\mathbb{Z}_N, +$

1.1 Finding expression for Fourier transform

In previous lectures, we derived the expression for Quantum Fourier Transform as:

$$F|x\rangle = \frac{1}{\sqrt{N}} \sum_{y=0}^{N-1} \exp\left(\frac{2\pi i x y}{N}\right) |y\rangle. \tag{1}$$

Qualitatively, applying the Fourier transform (F) to a basis state, $|x\rangle$ results in a superposition of the basis states $|y\rangle$ where the amplitude for each basis state has absolute value $\frac{1}{\sqrt{N}}$ and the phase is given by $\exp\left(\frac{2\pi i x y}{N}\right)$. Since y here represents an N bit binary number, considering the binary expansion, it can be written as

$$y = \sum_{j=1}^{n} y_j 2^{n-j} \qquad \text{for } \; y_i \in \{0, 1\}. \tag{2}$$

Thus, we can write $|y\rangle$ as:

$$|y\rangle = |y_1\rangle|y_2\rangle|y_3\rangle...|y_n\rangle = |y_1\rangle \otimes |y_2\rangle \otimes |y_3\rangle \otimes ... \otimes |y_n\rangle, \tag{3}$$

where y_n is the lowest order bit and y_1 is the highest order bit. We can further use (2) to simplify the phase term in (1):

$$\exp\left(\frac{2\pi i x y}{N}\right) = \exp\left(\sum_{j=1}^{n} \frac{2\pi i x y_j 2^{n-j}}{2^n}\right) = \prod_{j=1}^{n} \exp\left(\frac{\pi i x y_j}{2^{j-1}}\right). \tag{4}$$

Combining the equations (2), (3) and (4), the overall expression for the Fourier transform can be written as:

$$F|x\rangle = \frac{1}{\sqrt{N}} \sum_{y_1=0}^{1} \sum_{y_2=0}^{1} ... \sum_{y_n=0}^{1} \exp(\pi i x y_1)|y_1\rangle \otimes \exp(\pi i y_2/2)|y_2\rangle \otimes ... \otimes \exp(\pi i x y_n/2^{n-1})|y_n\rangle. \tag{5}$$

In the expression above, the sum over each individual y_j goes from 0 to 1 (from (2)). Further, we have also written $|y\rangle$ as a tensor product on individual qubits. Since the y_i part does not depend on any previous element of the tensor product, we can use distributivity of tensor products over sums to move the sum inside of the tensor product to obtain

$$F|x\rangle = \frac{1}{\sqrt{N}} \sum_{y_1=0}^{1} \exp(\pi i x y_1)|y_1\rangle \otimes \sum_{y_2=0}^{1} \exp(\pi i x y_2/2)|y_2\rangle \otimes ... \otimes \sum_{y_n=0}^{1} \exp\left(\frac{\pi i x y_n}{2^{n-1}}\right)|y_n\rangle. \quad (6)$$

Note that x is an integer between 0 and $N-1$, and each y_j can either be 0 or 1. Focusing on each individual sum of the vector product in (6), we note that for $\sum_{y_j=0}^{1} \exp(\pi i x y_j/2^{j-1})\,|y_j\rangle$, the first term ($y_j = 0$) is always equal to $|0\rangle$. Further, for $y_j = 1$, the product $\exp(\pi i x y_j)$ becomes $\exp(\pi i x/2^{j-1})$ (we get a factor of 2 as we go further). Thus, the sum becomes $|0\rangle + \exp(\pi i x/2^{j-1})|1\rangle$. Re-writing our (6) we get:

$$\frac{1}{\sqrt{2}}(|0\rangle + \exp(\pi i x)|1\rangle) \otimes \frac{1}{\sqrt{2}}(|0\rangle + \exp(\pi i x/2)|1\rangle) \otimes ... \otimes \frac{1}{\sqrt{2}}\left(|0\rangle + \exp\left(\frac{\pi i x}{2^{n-1}}\right)|1\rangle\right). \quad (7)$$

Note that the factor $\frac{1}{\sqrt{N}}$ in (6) is broken down into n factors of $\frac{1}{\sqrt{2}}$ in (7). Using this, we can then write our final Fourier transform as

$$F|x\rangle = |z_1\rangle|z_2\rangle...|z_n\rangle,$$

where $|z_k\rangle = \frac{1}{\sqrt{2}}(|0\rangle + \exp(\pi i x/2^{k-1})|1\rangle)$.

1.2 Expressing $|z_j\rangle$ in terms of $|x\rangle$

In this section, we see how we can realize each of these qubits. We start with the simplest case: $|z_1\rangle$. We first note that similar to (2), we can write $x = \sum_{j=1}^{n} x_j 2^{n-j}$, where $x_j \in \{0, 1\}$. Using this, the expression for $|z_1\rangle$ can be written as:

$$|z_1\rangle = \frac{1}{\sqrt{2}}(|0\rangle + \exp(\pi i x_n)|1\rangle). \quad (8)$$

Here, note that the higher order bit in the binary expansion of x do not contribute to the exponential term as they give us a factor of 2 (through the exponent of 2) and $\exp(2\pi i k) = 1$ for any integer k. Thus, for $|z_1\rangle$, only the lowest order bit of x contributes. Next, we analyze the exponential term in z_1. The general distribution of $\exp(i\theta)$ is shown in Fig. 1.2. Here, we see that when $\theta = \pi$, then the value is -1 whereas if $\theta = 0$, the value is 1. Thus, for $x_n = 0$, $|z_1\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ whereas for $x_n = 1$ we have $|z_1\rangle = \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$. In other words, $|z_1\rangle$ can be obtained from $|x_n\rangle$ by application of an Hadamard gate which allows us to rewrite (8) as:

$$|z_1\rangle = \frac{1}{\sqrt{2}}(|0\rangle + \exp(\pi i x_n)|1\rangle) = H|x_n\rangle, \quad (9)$$

which can be represented in circuits as:

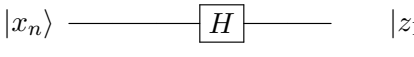

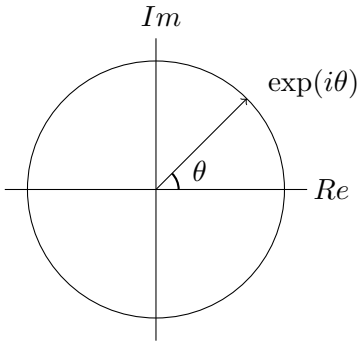

Figure 1: Figure showing the distribution of values for $\exp(i\theta)$.

We can extend this analysis to $|z_2\rangle$. Note here that only the two lowest order bits of x contibute here as the other bits yield a factor of 1 (similar to above). After applying the simplification, the expression can be written as:

$$|z_2\rangle = \frac{1}{\sqrt{2}}(|0\rangle + \exp(\pi i x_{n-1}) \exp(\pi i x_n/2) |1\rangle). \tag{10}$$

This can also be realized by applying the Hadamard gate to $|x_{n-1}\rangle$, i.e., $H|x_{n-1}\rangle = \frac{1}{\sqrt{2}}(|0\rangle + \exp(\pi i x_{n-1})|1\rangle)$, with an additional phase factor on the component $|1\rangle$. The value of the additional phase factor is 1 if $x_n = 0$ and $\exp(i\pi/2)$ for $x_n = 1$. We can also write this as a conditional rotation, conditioned on the qubit $|x_n\rangle$ being $|1\rangle$. This is expressed as:

$$CR(\pi/2) |x_n\rangle H |x_{n-1}\rangle = |x_n\rangle|z_2\rangle. \tag{11}$$

In terms of circuits, this can be expressed as:

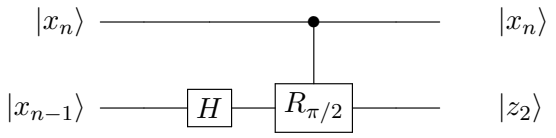

Qualitatively, to obtain $|z_2\rangle$, we first apply the Hadamard gate to $|x_{n-1}\rangle$ and then apply a controlled rotation of $\frac{\pi}{2}$ controlled by qubit $|x_n\rangle$. Note that the order of gates matter here. For example, we cannot swap the Hadamard gate and the controlled rotation. To see this, consider the case when the $|x_{n-1}\rangle$ is in the $|0\rangle$ state. Here, if we apply the rotation first and then the Hadamard gate, the conditional rotation does not have an effect (it only affects the $|1\rangle$ state) and the application of the Hadamard gate results in the final state being $|+\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ (without any additional phase on $|1\rangle$). On the contrary, if we apply the Hadamard first, then the input for the conditional rotation is $|+\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$, resulting in the final state $|+\rangle$ if $x_n = 0$ but $|-\rangle = \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$ if $x_n = 1$. Thus, the order of application of gates are matters.

Extending the procedure to $|z_3\rangle$, we can construct the following circuit:

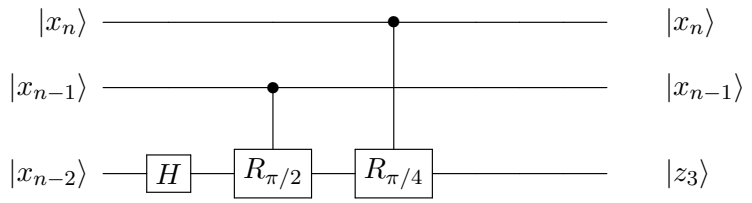

1.3 Putting it all together

We can extend this pattern to get circuits for $|z_k\rangle$, where we transform the qubit $|x_{n-k+1}\rangle$ into $|z_k\rangle$. In order to calculate the overall Fourier transform, we need the jth bit of x to act as a control on $|x_k\rangle$ for all $k < j$, and thus we cannot change $|x_j\rangle$ until all $|x_k\rangle$ for $k < j$ have been transformed. Due to this, we must compute the $|z_k\rangle$ in descending order, computing $|z_n\rangle$ and then $|z_{n-1}\rangle$, until computing $|z_1\rangle$. This results in the following circuit:

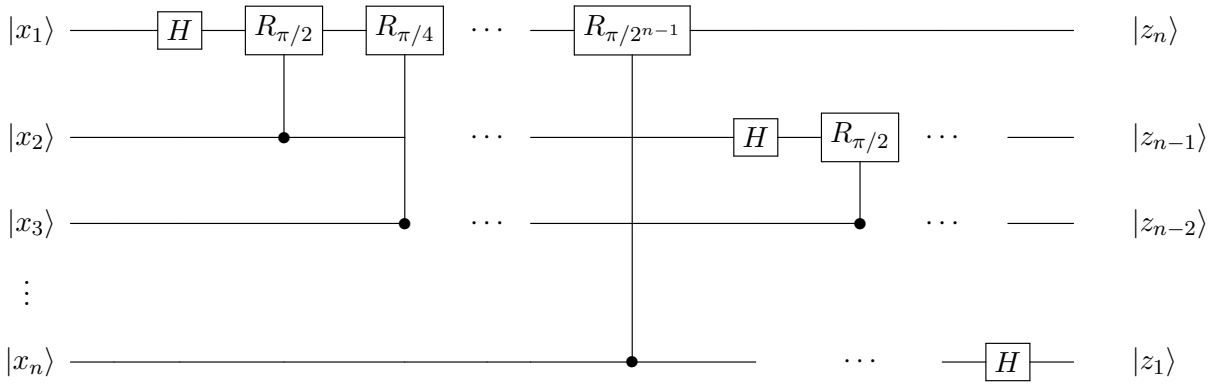

Since this unitary circuit works for all basis states $|x\rangle$, it also works for all possible superpositions $|\psi\rangle$. Note that this circuit reverts the order of the output qubits, which can later be reversed using swaps to get the qubits in order. The resulting circuit consists of $n - j + 1$ gates for each x_j, which makes the total gates to be $O(n^2)$.

An important observation here is that many of the further rotations are extremely small, so to compute this transform approximately, we can omit them and still obtain a good approximation of the whole circuit. More precisely:

Exercise 1. *Dropping rotations $R_{\pi/2^j}$ for $j \geq \log(n/\epsilon)$ yields circuit with $O(n \log(n/\epsilon))$ gates that $O(\epsilon)$ approximates F in 2-norm.*

2 Phase Estimation

Next, we discuss the phase estimation problem. We describe both the classical and quantum versions, with the quantum version described as the subroutine output below.

2.1 Statement

Input: $|\psi\rangle = \frac{1}{\sqrt{N}} \sum_{x=0}^{N-1} \exp(2\pi i \omega x)|x\rangle$ where $\omega \in [0, 1)$ is unknown.

Note that the input expression is very similar to the Quantum Fourier Transform. The key difference being that ω can take any value between 0 and 1. The objective here is to find ω (or a good approximation of it).

Output: ω, or a good approximation of the form y/N for $y \in \mathbb{Z}_N$, such that $|\omega - \frac{y}{N}|_{\mathbb{T}} \leq \delta$, where we can take δ to be another parameter. $\mathbb{T}$ refers to "modulo 1" (explained in 2.1)

Subroutine Output: Pure state $|\tilde{\omega}\rangle$ on n qubits with total weight of good y's at least $1 - \epsilon$ (i.e., it is close to 1). Here, good y's are the ones that satisfy the condition above, i.e., they are close to ωN modulo 1. To obtain the classical output, we can just observe the particular state. It is expected to give us y *close to* ωN with a high probability (at least $1 - \epsilon$).

Definition of $\mathbb{T}$

Consider the diagram below with two points, x and y, close to each other on a circle. Here, we have that $|x - y|_{\mathbb{T}}$ is the distance between x and y on the circle, which can be formally defined by $|x - y|_{\mathbb{T}} = \min_{z \in \mathbb{Z}} |x - y + z|$. We get this metric by "wrapping the interval $[0, 1]$ around on the circle", as can be seen in the Fig. 2.1:

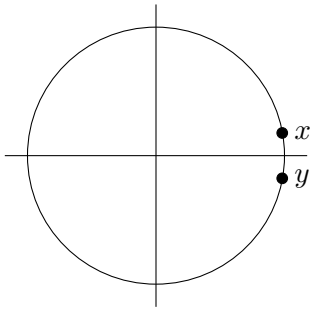

Figure 2: Figure to intuitively explain wrapping the interval $[0, 1]$ around on the circle.

In the above figure, we get x from a very small angle and y from a very large angle, however in this picture x and y are still close. The distance on the circle captures this, making points very close to 1 and very close to 0 close to each other under the $\mathbb{T}$ metric. The letter $\mathbb{T}$ stands for a torus, as a circle is a 1-dimensional torus.

2.2 Algorithm

To obtain the phase value, we can simply apply the inverse Fourier transform $F^{-1}|\psi\rangle$ over $\mathbb{Z}_N, +$. We note that:

- In case $\omega = \frac{y}{N}$ exactly, then we have $|\tilde{\omega}\rangle = |y^*\rangle$. In this case, measuring the final state always yields a y^*. In this scenario, $|\psi\rangle$ is just the fourier transform of $|y^*\rangle$.
- In the general case, the output $|\tilde{\omega}\rangle$ satisfies the aforementioned requirements for any $\delta > 0$ with $\delta \cdot \epsilon = O(1/N)$. Intuitively, we can understand this as $|\tilde{\omega}\rangle$ has most of its weight on $y \in \mathbb{Z}_N$

with $\frac{y}{N}$ close to ω modulo 1. Measuring the final state yields a good y (i.e, $|\omega - \frac{y}{N}|_{\mathbb{T}} \leq \delta$) with probability at least $1 - \epsilon$. We can also see that there is a trade-off between ϵ and δ. This can be understood intuitively as if we want to be very certain in our measurement (low ϵ) then we would have to increase the error margin on ω.

2.3 Analysis

We have that $F^{-1}|x\rangle = \frac{1}{\sqrt{N}} \sum_{y=0}^{N-1} \exp\left(-\frac{2\pi i x y}{N}\right)|y\rangle$. We apply this to the superposition $|\psi\rangle$ that we began with, $|\psi\rangle = \frac{1}{\sqrt{N}} \sum_{x=0}^{N-1} \exp(2\pi i \omega x)|x\rangle$. This gives the equation:

$$F^{-1}|\psi\rangle = \frac{1}{\sqrt{N}} \sum_{x=0}^{N-1} \exp(2\pi i \omega x) \frac{1}{\sqrt{N}} \sum_{y=0}^{N-1} \exp\left(-\frac{2\pi i x y}{N}\right)|y\rangle = \sum_{y=0}^{N-1} \alpha_y |y\rangle,$$

where $\alpha_y = \frac{1}{N} \sum_{x=0}^{N-1} \exp(2\pi i \omega x) \exp\left(-\frac{2\pi i x y}{N}\right)$. We can simplify this expression by using the product properties of exponentials, to get that

$$\alpha_y = \frac{1}{N} \sum_{x=0}^{N-1} \exp(2\pi i \Delta x), \text{ for } \Delta \doteq \omega - \frac{y}{N},$$

which is equal to

$$\frac{1}{N} \sum_{x=0}^{N-1} r^x, \text{ for } r \doteq \exp(2\pi i \Delta).$$

If $r = 1$, then this is an arithmetic sum. We have that $r = 1 \Leftrightarrow \Delta = 0 \Leftrightarrow \omega = \frac{y^*}{N}$ for some $y^* \in \mathbb{Z}_N$, in which case $\alpha_{y^*} = \frac{1}{N} \cdot N = 1$, and we are guaranteed to observe y^*. Otherwise,

$$\alpha_y = \frac{1}{N} \cdot \frac{1 - r^N}{1 - r} = \frac{1}{N} \cdot \frac{1 - \exp(2\pi i \Delta N)}{1 - \exp(2\pi i \Delta)}.$$

Our goal is now to show that the weight of the bad y's is small, i.e. $\sum_{\text{bad } y} |\alpha_y|^2$ is small, where a y is considered bad if $|\omega - \frac{y}{N}|_{\mathbb{T}} > \delta$. We have that

$$\begin{aligned}
|1 - \exp(i\theta)| &= |1 - \cos(\theta) - i\sin(\theta)| \\
&= \sqrt{(1 - \cos(\theta))^2 + \sin^2(\theta)} \\
&= \sqrt{1 - 2\cos(\theta) + \cos^2(\theta) + \sin^2(\theta)} \\
&= \sqrt{1 - 2\cos(\theta) + 1} \\
&= \sqrt{2(1 - \cos(\theta))}.
\end{aligned}$$

We can use the double angle identity $(1 - \cos(\theta)) = 2\sin^2(\theta/2)$, to get that this is equal to:

$$\begin{aligned}
&= \sqrt{4\sin^2(\theta/2)} \\
&= 2|\sin(\theta/2)|.
\end{aligned}$$

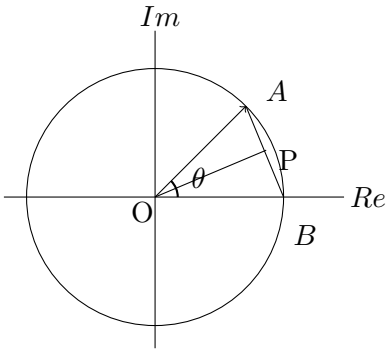

Figure 3: Geometric interpretation of the distance $|A - B|$ where $A = \exp(i\theta)$ and $B = 1$.

Note that we can also compute $|1-\exp(i\theta)|$ geometrically as follows: In Fig. 2.3, $|1-\exp(i\theta)|$ is given by the distance of the line segment AB. To compute this distance, we can draw a perpendicular to the segment of the circle. This perpendicular line makes an angle $\theta/2$ with the x-axis. In the right-triangle OPB, the side PB is given by $\sin(\theta/2)$. This gives us the length of AB as $2\sin(\theta/2)$. Therefore $|1 - \exp(i\theta)| = 2|\sin(\theta/2)|$.

Thus, we have that

$$\alpha_y = \frac{1}{N} \cdot \frac{|\sin(\pi\Delta N)|}{|\sin(\pi\Delta)|}.$$

We now attempt to do two things: we would like to bound $|\alpha_y|$ from above for the bad y, and we would like to give a lower bound for the weight of y^*, where y^*/N is the best approximation of ω.

Bounding weight of optimal y from below: We first start by bounding $\sin(\theta)$ from above and below using the convexity of the sine function.

- For the upper bound, we note that the tangent always stays above, thus we have that $\sin(\theta) \leq \theta$ for all $\theta \in [0, \pi]$. Further, since $\cos(\theta) \leq 1$ we know that θ is increasing faster than $\sin(\theta)$.
- For the lower bound, we note that a chord always stays below the function. Thus, $\sin(\theta) \geq \frac{2}{\pi}\theta$ for all $\theta \in [0, \frac{\pi}{2}]$, as can be seen in Fig. 2.3. The equation of the chord ($y = \frac{2}{\pi}\theta$) is simply the equation of line connecting origin and $\theta = \sin\frac{\pi}{2}$.

We use these bounds for the sine function to get lower bounds for $|\alpha_{y^*}|_{\mathbb{T}}$, where y^* is the best approximation of ω. To get this bound, we take the lower bound of the numerator and the upper bound of the denominator. Specifically, we get a lower bound for the numerator by using $\sin(\theta) \geq \frac{2\theta}{\pi}$ for all $\theta \in [0, \frac{\pi}{2}]$ (lower bound for the sine function). Further, we get an upper bound for the denominator by using $\sin(\theta) \leq \theta$ (upper bound for the sine function). Now, there are N evenly spaced possibilities for the weight of optimal y^* in the interval $[0, 1]$. Whatever ω is, there must be at least one y within $\frac{1}{2N}$ of ω, as otherwise there would be an interval of length $1/N$ around ω without any y. This gives us that the best approximation $\frac{y^*}{N}$ satisfies $|\Delta| \doteq |\omega - \frac{y^*}{N}|_{\mathbb{T}} \leq \frac{1}{2N}$. We can use this, combined with the formula above for $|\alpha_{y^*}|$ by plugging in the lower bound of $\sin(\theta)$ to the numerator and upper bound of $\sin(\theta)$ to the denominator, and get:

$$|\alpha_{y^*}| \geq \frac{1}{N} \cdot \frac{2\Delta N}{\pi\Delta} = \frac{2}{\pi}.$$

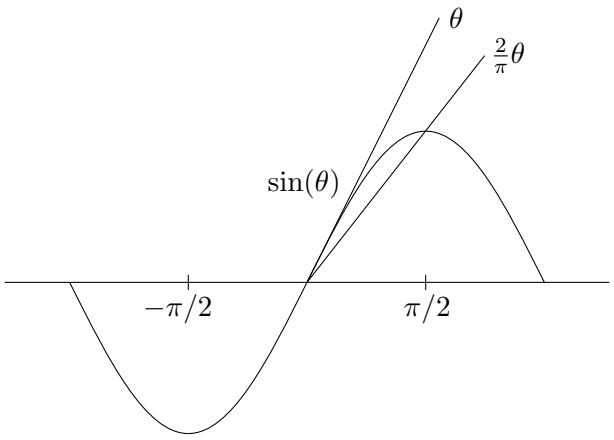

Figure 4: Geometric intuition for bounding $\sin(\theta)$.

Note that α_{y^*} is the amplitude to observe y^* when we measure $\tilde{\omega}$. Thus, the probability of observing the optimal y^* is at least $|\alpha_{y^*}|^2 \geq \left(\frac{2}{\pi}\right)^2 = \frac{4}{\pi^2}$, which is about 40%. We also note that there could be two values that achieve this bound. Due to this, we cannot hope to have probability of observing optimal y^* more than 50%.

Bounding total weight of bad y from above: Bad y's are such that $|\Delta| \geq \delta, |\Delta + 1| \geq \delta, |\Delta - 1| \geq \delta$. We wish to bound $\sum_{\text{bad } y} |\alpha_y|^2$. This expression gives the total weight of bad y's. We have:

$$\sum_{\text{bad } y} |\alpha_y|^2 \leq \frac{1}{N^2} \sum_{\text{bad } y} \frac{1}{\sin^2(\pi\Delta)} \leq \frac{2}{N^2} \sum_{\delta \leq \Delta \leq 1/2} \frac{1}{\sin^2(\pi\Delta)} \leq \frac{2}{N^2} \sum_{\delta \leq \Delta \leq 1/2} \frac{1}{(2\Delta)^2}.$$

Here, in the first step, we have replaced the numerator ($|\sin(\pi\Delta N)|^2$) by 1 as the sine term is always less than or equal to 1. In the next step, we enumerate what these bad y's can be (cases where Δ is at least δ). In these cases, the value of Δ can be positive or negative; here, we only consider the positive values, thus we get a factor of 2. Also note that Δ can not be more than $\frac{1}{2}$, as it is the distance to the closest integer. Finally, since we are trying to get an upper bound of the expression, we need to consider the lower bound for the denominator (the sine term). Thus, we use $\sin(\theta) \leq \theta$ to simplify the expression in the last step.

We can further write this instead as a function of ΔN, pulling in the $\frac{1}{N^2}$ to the summand:

$$\leq \frac{1}{2} \sum_{\delta N \leq \Delta N \leq N/2} \frac{1}{(\Delta N)^2}.$$

But the Δ are going to increase by multiples of $\frac{1}{N}$, so ΔN increases as integers, and we can rewrite the sum above as:

$$\leq \frac{1}{2} \sum_{k \geq \lceil \delta N \rceil} \frac{1}{k^2}.$$

This can in turn be bounded from above by an integral,

$$\leq \frac{1}{2} \int_{\delta N - 1}^{\infty} \frac{1}{x^2} dx = \frac{1}{2} \frac{1}{\delta N - 1} = O(1/(\delta N)).$$

Thus, $\epsilon = O(1/(\delta N))$, and $\epsilon\delta = O(1/N)$, where ϵ is the odds of observing a bad y, which is precisely what we were trying to prove.