

Quantum algorithms (CO 781/CS 867/QIC 823, Winter 2011)

Andrew Childs, University of Waterloo

LECTURE 9: Fourier sampling

In this lecture, we will see how the Fourier transform can be used to simplify the structure of the states obtained in the standard approach to the hidden subgroup problem. In particular, we will see how *weak Fourier sampling* is sufficient to identify any normal hidden subgroup (generalizing the solution of the abelian HSP). We will also briefly discuss the potential of *strong Fourier sampling* to go beyond the limitations of weak Fourier sampling.

Weak Fourier sampling

Recall that the standard approach to the HSP allows us to produce a *coset state*

$$|gH\rangle := \frac{1}{\sqrt{|H|}} \sum_{h \in H} |gh\rangle \quad (1)$$

where each $g \in G$ occurs uniformly at random; or equivalently, the *hidden subgroup state*

$$\rho_H := \frac{1}{|G|} \sum_{g \in G} |gH\rangle\langle gH|. \quad (2)$$

The symmetry of such a state can be exploited using the quantum Fourier transform. In particular, we have

$$|gH\rangle = \frac{1}{\sqrt{|H|}} \sum_{h \in H} R(h)|g\rangle \quad (3)$$

where R is the right regular representation of G . Thus the hidden subgroup state can be written

$$\rho_H = \frac{1}{|G| \cdot |H|} \sum_{g \in G} \sum_{h, h' \in H} R(h)|g\rangle\langle g|R(h')^\dagger \quad (4)$$

$$= \frac{1}{|G| \cdot |H|} \sum_{h, h' \in H} R(hh'^{-1}) \quad (5)$$

$$= \frac{1}{|G|} \sum_{h \in H} R(h). \quad (6)$$

Since the right regular representation is block-diagonal in the Fourier basis, the same is true of ρ_H . In particular, we have

$$\hat{\rho}_H := F_G \rho_H F_G^\dagger \quad (7)$$

$$= \frac{1}{|G|} \bigoplus_{\sigma \in \hat{G}} (I_{d_\sigma} \otimes \sigma(H)^*) \quad (8)$$

where

$$\sigma(H) := \sum_{h \in H} \sigma(h). \quad (9)$$

Since $\hat{\rho}_H$ is block diagonal, with blocks labeled by irreducible representations, we may now measure the irrep label without loss of information. This procedure is referred to as *weak Fourier sampling*. The probability of observing representation $\sigma \in \hat{G}$ under weak Fourier sampling is

$$\Pr(\sigma) = \frac{1}{|G|} \operatorname{tr} (I_{d_\sigma} \otimes \sigma(H)^*) \quad (10)$$

$$= \frac{d_\sigma}{|G|} \sum_{h \in H} \chi_\sigma(h)^* \quad (11)$$

$$= \frac{d_\sigma |H|}{|G|} (\chi_\sigma, \chi_1)_H, \quad (12)$$

or in other words, $d_\sigma |H|/|G|$ times the number of times the trivial representation appears in $\operatorname{Res}_H^G \sigma$, the restriction of σ to H . We may now ask whether polynomially many samples from this distribution are sufficient to determine H , and if so, whether H can be reconstructed from this information efficiently.

Normal subgroups

If G is abelian, then all of its representations are one-dimensional, so weak Fourier sampling reveals all of the available information about ρ_H . (In this case there is no difference between weak Fourier sampling and strong Fourier sampling, which we will discuss later.) Indeed, for an abelian group, we saw that the information provided by Fourier sampling can be used to efficiently determine H .

Weak Fourier sampling succeeds for a similar reason whenever H is a *normal subgroup* of G (denoted $H \trianglelefteq G$), i.e., whenever $gHg^{-1} = H$ for all $g \in G$. In this case, the hidden subgroup state within the irrep $\sigma \in \hat{G}$ is proportional to

$$\sigma(H)^* = \frac{1}{|G|} \sum_{g \in G, h \in H} \sigma(ghg^{-1})^*. \quad (13)$$

This commutes with $\sigma(g)^*$ for all $g \in G$, so by Schur's Lemma, it is a multiple of the identity. Thus $\hat{\rho}_H$ is proportional to the identity within each block, and again weak Fourier sampling reveals all available information about H .

Furthermore, when $H \trianglelefteq G$, the distribution under weak Fourier sampling is a particularly simple generalization of the abelian case: we have

$$\Pr(\sigma) = \begin{cases} d_\sigma^2 |H|/|G| & H \leq \ker \sigma \\ 0 & \text{otherwise,} \end{cases} \quad (14)$$

where $\ker \sigma := \{g \in G : \sigma(g) = I_{d_\sigma}\}$ is the kernel of the representation σ (a normal subgroup of G). To see this, note that if $H \not\leq \ker \sigma$, then there is some $h' \in H$ with $\sigma(h') \neq 1$; but then $\sigma(h')\sigma(H) = \sum_{h \in H} \sigma(h'h) = \sigma(H)$, and since $\sigma(h')$ is unitary and $\sigma(H)$ is a scalar multiple of the identity, this can only be satisfied if in fact $\sigma(H) = 0$. On the other hand, if $H \leq \ker \sigma$, then $\chi_\sigma(h) = d_\sigma$ for all $h \in H$, and the result is immediate.

To find H , we can simply proceed as in the abelian case: perform weak Fourier sampling $O(\log |G|)$ times and compute the intersection of the kernels of the resulting irreps (assuming this can be done efficiently). Again, it is clear that the resulting subgroup contains H , and we claim that it is equal to H with high probability. For suppose that at some stage during this process, the

intersection of the kernels is $K \trianglelefteq G$ with $K \neq H$; then the probability of obtaining an irrep σ for which $K \leq \ker \sigma$ is

$$\frac{|H|}{|G|} \sum_{\sigma: K \leq \ker \sigma} d_\sigma^2 = \frac{|H|}{|K|} \leq \frac{1}{2} \quad (15)$$

where we have used the fact that the distribution (14) remains normalized if H is replaced by any normal subgroup of G . Since each repetition of weak Fourier sampling has a probability of at least $1/2$ of cutting the intersection of the kernels at least in half, $O(\log |G|)$ repetitions suffice to converge to H with substantial probability. In fact, applying the same approach when H is not necessarily normal in G gives an algorithm to find the *normal core* of H , the largest subgroup of H that is normal in G .

Strong Fourier sampling

Despite the examples we have just discussed, weak Fourier sampling does *not* provide sufficient information to recover the hidden subgroup for the majority of hidden subgroup problems. For example, weak Fourier sampling fails to solve the HSP in the symmetric group and the dihedral group.

To obtain more information about the hidden subgroup, we can perform a measurement on the d_σ^2 -dimensional state that results when weak Fourier sampling returns the outcome σ . Such an approach is referred to as *strong Fourier sampling*.

Recall that the state $\hat{\rho}_H$ from (8) is maximally mixed over the row register, as a consequence of the fact that the left and right regular representations commute. Thus we may discard this register without loss of information, so that strong Fourier sampling is effectively faced with the d_σ -dimensional state

$$\hat{\rho}_{H,\sigma} := \frac{\sigma(H)^*}{\sum_{h \in H} \chi_\sigma(h)^*}. \quad (16)$$

In fact, this state is proportional to a projector whose rank is simply the number of times the trivial representation appears in $\text{Res}_H^G \sigma^*$. This follows because

$$\sigma(H)^2 = \sum_{h,h' \in H} \sigma(hh') = |H| \sigma(H), \quad (17)$$

which gives

$$\hat{\rho}_{H,\sigma}^2 = \frac{|H|}{\sum_{h \in H} \chi_\sigma(h)^*} \hat{\rho}_{H,\sigma}, \quad (18)$$

so that $\hat{\rho}_{H,\sigma}$ is proportional to a projector with $\text{rank}(\hat{\rho}_{H,\sigma}) = \sum_{h \in H} \chi_\sigma(h)^* / |H|$.

It is not immediately clear how to choose a good basis for strong Fourier sampling, so a natural first approach is to consider the effect of measuring in a random basis (i.e., a basis chosen uniformly with respect to the Haar measure over \mathbb{C}^{d_σ}). There are a few cases in which such *random strong Fourier sampling* produces sufficient information to identify the hidden subgroup—in particular, Sen showed that it succeeds whenever $\text{rank}(\hat{\rho}_{H,\sigma}) = \text{poly}(\log |G|)$ for all $\sigma \in \hat{G}$.

However, in many cases random strong Fourier sampling is unhelpful. For example, Grigni et al. showed that if H is sufficiently small and G is sufficiently non-Abelian (in a certain precise sense), then random strong Fourier sampling is not very informative. In particular, they showed this for the problem of finding hidden involutions in the symmetric group. Another example was

provided by Moore et al., who showed that random strong Fourier sampling fails in the metacyclic groups $\mathbb{Z}_p \rtimes \mathbb{Z}_q$ (subgroups of the affine group $\mathbb{Z}_p \rtimes \mathbb{Z}_p^\times$) when $q < p^{1-\epsilon}$ for some $\epsilon > 0$.

Even when measuring in a random basis is information-theoretically sufficient, it does not give an efficient quantum algorithm, since it is not possible to efficiently measure in a random basis. It would be interesting to find informative pseudo-random bases that can be implemented efficiently. However, in the absence of such techniques, we can instead hope to find explicit bases in which strong Fourier sampling can be performed efficiently, and for which the results give a solution of the HSP. The first such algorithm was provided by Moore et al., for the aforementioned metacyclic groups, but with $q = p/\text{poly}(\log p)$. Note that for these values of p, q , unlike the case $q < p^{1-\epsilon}$ mentioned above, measurement in a random basis is information-theoretically sufficient. Indeed, we do not know of *any* example of an HSP for which strong Fourier sampling succeeds, yet random strong Fourier sampling fails; it would be interesting to find any such example (or to prove that none exists).

Note that simply finding an informative basis is not sufficient; it is also important that the measurement results can be efficiently post-processed. This issue arises not only in the context of measurement in a pseudo-random basis, but also in the context of certain explicit bases. For example, Ettinger and Høyer gave a basis for the dihedral HSP in which a measurement gives sufficient classical information to infer the hidden subgroup, but no efficient means of post-processing this information is known.

For some groups, it turns out that strong Fourier sampling simply fails. Moore, Russell, and Schulman showed that, regardless of what basis is chosen, strong Fourier sampling provides insufficient information to solve the HSP in the symmetric group. Specifically, they showed that for any measurement basis (indeed, for any POVM applied to a hidden subgroup state), the distribution of outcomes in the cases where the hidden subgroup is trivial and where the hidden subgroup is an involution are exponentially close. Thus, in general one has to consider *entangled measurements* on multiple copies of the hidden subgroup states. (Indeed, entangled measurements on $\Omega(\log |G|)$ copies may be necessary, as Hallgren et al. showed for the symmetric group.) In the next two lectures, we will see some examples of quantum algorithms for the HSP that make use of entangled measurements.