Hadamard gates and amplitudes of computational basis states

A. Montanaro* and D. J. Shepherd†
University of Bristol, Department of Computer Science

September 21, 2006

Abstract

If we are given an adversarially chosen n-qubit state, to which we are allowed to apply any number of single-qubit Hadamard gates, can we always produce a state with all 2^n computational basis states having non-zero amplitudes? In this short note we show that the answer is "yes".

1 Introduction

With respect to some (orthonormal) computational basis, write H for the single-qubit Hadamard map

$$\begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} / \sqrt{2},$$

and write $|\psi\rangle_n = \sum_x \alpha_x |x\rangle$ for some given n-qubit state, in terms of the computational basis $\{ |x\rangle : x \in \mathbb{F}_2^n \}$. We have that $\sum_x |\alpha_x|^2 = 1$ by normalisation. Our goal shall be to analyse the quantity

$$\mu(|\psi\rangle_n) := \max_{p \in \mathbb{F}_2^n} \min_{x \in \mathbb{F}_2^n} \left| \langle x | (H^{p_1} \otimes \ldots \otimes H^{p_n}) |\psi\rangle_n \right|. \quad (1)$$

For the purpose of playing some adversarial games, define also

$$\mu_n := \min_{|\psi\rangle_n} \mu(|\psi\rangle_n). \quad (2)$$

*montanar@compsci.bristol.ac.uk
†shepherd@compsci.bristol.ac.uk, dan.shepherd@cesg.gsi.gov.uk

We will show that μ_n is never zero. This solves Scott Aaronson's 10^{th} most annoying question in quantum computation [1], which is stated as : Given an n-qubit pure state, is there always a way to apply Hadamard gates to some subset of the qubits, so as to make all 2^n computational basis components have nonzero amplitudes?

2 μ_n is never zero

We start this section with a simple lemma, then go on to discuss how this implies a lower bound for μ_n.

For any n-dimensional hypercube whose vertices are decorated each with a single bit (element of $\mathbb{F}_2$), call by the name *reduction* the following process : identifying one of the n directions associated to the hypercube, merging each vertex of the hypercube with the one that neighbours it in the identified direction, decorating the merger with the sum (in $\mathbb{F}_2$,) of the decorations of the original two contributing vertices, thereby obtaining an $n-1$-dimensional hypercube with vertices again decorated each with a single bit.

Lemma 1 *For any hypercube decorated as above, if at least one vertex is decorated with a 1, then there exists a (possibly empty) sequence of* reductions *resulting in a hypercube all of whose vertices are decorated with a 1.*

For a contradiction, consider a smallest hypercube *not* having this property. Because no sequence of reductions on such a hypercube may result in an "all 1" hypercube, no reduction of the hypercube can have a reduction sequence that results in an "all 1" hypercube either. Therefore each of its reductions must be either a *smaller* counterexample to the lemma or else be an "all 0" hypercube. The first of these possibilities we must reject by hypothesis (since our original counterexample was claimed smallest,) and so *every* reduction of the original hypercube must be "all 0". This means that every pair of neighbouring vertices in it must have decorations summing to 0, which means that the decorations must all be the same. But by hypothesis, they can neither be all 0 (since at least one is promised to be 1 in the lemma,) nor can they be all 1 (in this case, we could perform the empty sequence of reductions, resulting in an "all 1" hypercube). Thus there *is* no smallest hypercube not having the property of the lemma, and so the lemma stands by induction. ■

This sets the scene for a proof that $\mu_n \neq 0$. Suppose to the contrary that $\mu_n = 0$ for some n. Then there exists some $|\,\psi\,\rangle_n = \sum_y \alpha_y |\,y\,\rangle$ for which, for all $p \in \mathbb{F}_2^n$, there exists an $x \in \mathbb{F}_2^n$ such that the xth component (in the computational basis) of the vector $T(p)|\,\psi\,\rangle_n$ is zero, where $T(p) = 2^{|p|/2} \cdot H^{p_1} \otimes \ldots \otimes H^{p_n}$. Now let M be a subgroup of $\mathbb{C}$ (complex field under addition) which is the lowest-rank $\mathbb{Z}$-module that happens to contain all of the α_y values, (*i.e.* M is the additive closure of the set $\{\alpha_y : y \in \mathbb{F}_2^n\}$,) and let π be any non-trivial homomorphism from M to $\mathbb{Z}$. Since the entries of $T(p)$ are all integers, it follows that

$$\begin{aligned} \forall p\, \exists x \quad 0 \;=\; \pi(0) \;&=\; \pi\Big(\, \langle\, x \,|\, T(p) \sum_y \alpha_y |\, y\,\rangle \,\Big) \\ &=\; \pi\Big(\, \sum_y \, \{T(p)\}_{xy}\; \alpha_y \,\Big) \\ &=\; \sum_y \, \{T(p)\}_{xy}\; \pi(\alpha_y) \\ &=\; \langle\, x \,|\, T(p) \sum_y \pi(\alpha_y) |\, y\,\rangle, \qquad (3) \end{aligned}$$

and hence there must exist some (unnormalised) non-zero vector $\sum_y \pi(\alpha_y)|\,y\,\rangle$, whose coefficients are integers, that equally well provides a counterexample. Consider such a counterexample after first factoring out any powers of 2 common to all the integer coefficients, so that at least one of the coefficients is odd. Regard the coefficients as being in one-to-one correspondence with the vertices of an n-dimensional hypercube graph. Let the vertices of such a hypercube ($\cong \mathbb{F}_2^n$) be decorated each with a single bit according to the mod 2 value of the corresponding (integer) coefficient. (Since at least one of the coefficients is odd, at least one of the decorations will be a 1.) Then by reading equation (3) "modulo 2", it may be seen that our counterexample directly contradicts lemma 1. (This is because the matrix $T(p)$ (mod 2) emulates the 'reduction sequence' employed in the lemma, and equation (3) says that all reduction sequences terminate with a hypercube that has at least one 0 decoration.) Hence we must conclude that $\mu_n > 0$ after all. ■

3 Upper bounds

Consider the state

$$|\,\psi\,\rangle_n \;=\; \Big(\; \cos(\pi/8)|\,0\,\rangle \;+\; \sin(\pi/8)|\,1\,\rangle \;\Big)^{\otimes n}. \qquad (4)$$

It is easy to see that Hadamard gates on this state do not affect the magnitudes of the amplitudes, and so $\mu(|\,\psi\,\rangle_n) = \sin^n(\pi/8)$. This is an upper bound for μ_n, and is quite probably tight. (The reader may easily check that it is tight in the case $n = 1$.)

4 Further work

It would be nice to have a polynomial-time construction for $\mu(|\,\psi\,\rangle_n)$ whenever only polynomially many coefficients of $|\,\psi\,\rangle_n$ are non-zero. We would also like to have a proof of the tightness of the upper bound quoted above, or else some lower bounds for μ_n.

Acknowledgements

AM would like to thank Simone Severini for introducing him to this problem, and for interesting discussions.

References

[1] Scott Aaronson (2006). The ten most annoying questions in quantum computing.
http://www.scottaaronson.com/blog/2006/08/ten-most-annoying-questions-in-quantum.html