Three counterexamples refuting Kieu's plan for "quantum adiabatic hypercomputation"; and some uncomputable quantum mechanical tasks

Warren D. Smith
WDSmith@fastmail.fm

May 28, 2005

Abstract — **Tien D. Kieu, in 10 papers posted to the quant-ph section of the xxx.lanl.gov preprint archive [some of which were also published in printed journals such as Proc. Royal Soc. A 460 (2004) 1535] had claimed to have a scheme showing how, in principle, physical "quantum adiabatic systems" could be used to solve the prototypical computationally undecidable problem, Turing's "halting problem," in finite time, with success probability $> 2/3$ (where this $2/3$ is independent of the input halting problem).**

There were several errors in those papers, most which ultimately could be corrected. More seriously, we here exhibit counterexamples to a crucial step in Kieu's argument. The counterexamples are small quantum adiabatic systems in which "decoy" nonground states arise with high probability ($> 99.999\%$). Kieu had wrongly claimed no decoy state could ever acquire occupation probability greater than 50%. These counterexamples destroy Kieu's entire plan and there seems no way to correct the plan to escape them.

Nevertheless, there are some important consequences salvageable from Kieu's idea: we can prove that the "halflife" of Kieu's quantum systems is uncomputably large and no fully general form of the quantum adiabatic theorem can exist that yields computable upper bounds on adiabatic convergence times (both unless Church's thesis is false so that finite time adiabatic quantum system evolution is unsimulable); and we can prove that no algorithm exists to find the ground state energy of Kieu's class of quantum Hamiltonians and hence their long-term thermal behavior is uncomputable.

1 Sketch of Kieu's plan

Kieu wanted to set up a quantum physical system with two Hamiltonians, called H_I and H_P. He assumed one could gradually be homotoped into the other via the time-dependent Hamiltonian

$$\mathfrak{H}(t) = (1 - \frac{t}{T})H_I + \frac{t}{T}H_P \tag{1}$$

over the time interval $0 \le t \le T$. Kieu proposed setting this all up in a Cartesian product of a finite number K of countably-infinite-dimensional Fock spaces. Physically, this space may be regarded as K kinds of bosons, each of which is present in some natural number $(0, 1, 2, 3, \dots)$ of copies. The states of Kieu's system then are precisely described by K natural numbers specifying, for each boson-type, how many of that kind of boson are present. Then H_P describes a weird kind of interaction between those bosons which the experimenter can "ramp up." Other physical interpretations also are possible.

Kieu then proposed one particular universal H_I, namely the Hamiltonian of a "positionally shifted simple Harmonic oscillator." More importantly, he proposed a way to define H_P so that, for any particular "Diophantine function" $D(x_1, x_2, \dots, x_K)$ of K natural number arguments $x_1, x_2, \dots x_K$, there was a simply-constructed corresponding H_P. Kieu's H_P had the property that its energy levels were precisely the same as the attainable integer values of D. Specifically, $H_P = D(a_1^\dagger a_1,\ a_2^\dagger a_2, \dots, a_K^\dagger a_K)$ where $a_1^\dagger,\ a_2^\dagger, \dots,\ a_K^\dagger$ and $a_1,\ a_2, \dots,\ a_K$ are K kinds of raising and lowering operators. By making D be a sum of squares, these attainable values are nonnegative so that Kieu's H_P is nonnegative definite. The diophantine equation $D = 0$ is then soluble exactly if H_P's ground state has energy$= 0$, and insoluble exactly if H_P's ground state has energy ≥ 1. Both of Kieu's H_I and H_P are self-adjoint nonnegative definite operators with discrete (indeed, natural number valued) spectra.

Kieu here made an error about Diophantine equations. He seemed to have the idea that we only need to worry about Diophantine equations $D = 0$ with *unique* solutions, leading to H_P with unique ("non-degenerate") ground states. In fact, it is commonplace for Diophantine equations to have an *infinite* number of solutions, and indeed the only polynomial Diophantine equations presently known to achieve Turing-completeness always do have either an infinite number, or no, solutions (it being Turing-undecidable which) [3]. Thus Kieu was maximally wrong about this. Kieu also had the alternative idea that by adding certain perturbations to his Hamiltonians, the "degeneracy could be broken" but in situations with infinite-fold degeneracy this claim was probably impossible to justify without destroying his soon-to-be-described algorithm by allowing infinitely-close energy levels, or by requiring non-algorithmic steps such as "taking limits"or "guessing" the right perturbation to use.

However, this error is repairable. The present author (who was serving as the referee on one of Kieu's papers) was able to modify the proof of Jones & Matiyasevic [6] concerning "singlefold 2-exponential Diophantine equations." By so doing I was able to construct Turing-complete 2-exponential Diophantine functions D which always have a *unique global minimum*. The value of D at this minimum is a nonnegative integer and it is Turing-undecidable whether it is zero. (I call these "singlemin" Diophantines.)

I was then able to show how to modify Kieu's construction to be based on these instead of on polynomial Diophantine equations.[1] So this error was not fatal.

Kieu then proposed starting his quantum system in the (known) ground state of H_I at $t = 0$, and time-evolving it under $\mathfrak{H}(t)$. By the "quantum adiabatic theorem" [2] this would, if T were made large enough, result in a final state when $t = T$ which, with high probability (approaching 1 as $T \to \infty$), would be the ground state of H_P. By measuring this state (even measuring it inexactly would suffice, since everything is integer valued) the solution to the singlemin Diophantine problem could be deduced, then confirmed with an old-fashioned non-quantum computer. Therefore an arbitrary Turing halting problem could be solved, revolutionizing computer science.

To make this work, though, it is crucial to know *how large* T must be before a high ground-state probability was assured. Kieu's plan to evade that objection was to repeatedly do the physical experiment with larger and larger T. (For example, T could be doubled between rounds of experiments.) He claimed to have proved[2] that no matter what the value of $T > 0$, no nonground "decoy" state could ever achieve an occupation probability greater than 1/2. Call this the **half-claim.**

If the half-claim were true, then Kieu could simply keep increasing T until the ground-state occupation probability p became greater than, say, 3/4. He could by performing enough physical experiments at each T, gain high confidence in the conclusion that the most-occupied state at that T was occupied with probability $\le 1/2$ or $\ge 3/4$ (if one of these were true). Kieu would simply keep redoubling T and keep raising his confidence criterion in such a way that eventually, Kieu was guaranteed to acquire high statistical confidence he really had the ground state of H_P.

There was here a second error, or perhaps a better word is "omission," by Kieu: he did not actually perform this probabilistic analysis rigorously. However, this error is not fatal, because I was able to do so. To be precise, for any fixed $\epsilon > 0$, and for any computable value $V(D)$ with $0 < V < 1$ replacing "1/2" in the half-claim, I was able to construct a computable increasing sequence of T's, numbers of experiments to be performed at each T, and confidence criteria that were to be applied after examining the results of those experiments, such that (1) the probability was 1 that the procedure would eventually terminate, and (2) the probability was $> 1 - \epsilon$ that, when it terminated, the correct global minimum of the Diophantine, i.e. the correct ground state of H_P, would be whatever was deduced from the final round of experiments. The analysis utilized Hoeffding's bound [5].

The ultimate source of the super-Turing computational power of Kieu's physical systems was the combination of (1) their infinite dimensionality and (2) the power of the adiabatic homotopy to find the ground state. This somehow seemed to provide "infinite computational parallelism." Kieu noted that quantum mechanics is inherently infinite dimensional because the position and momentum operator relation $[x, p] = \mathbf{1}i$ is not achievable by finite dimensional matrices.[3]

The *fatal* error is that the half-claim is false. We shall show this by constructing counterexamples.

This falsity is somewhat surprising because the half-claim *is* true in various limiting cases:

1. It is true in the "adiabatic limit" $T \to \infty$.
2. It is true in the "sudden limit" $T \to 0+$ provided that each inner product between the ground eigenstate g of H_I and any nonground eigenstate ϕ of H_P obeys $|\langle g|\phi\rangle|^2 \le 1/2$ (and a stronger claim than this proviso happens to be valid for the particular H_I and H_P used in Kieu's scheme).
3. It appears to be true if H_I and H_P are $n \times n$ Hermitian matrices (physically, "n-state quantum systems") with $n = 2$, such that H_I has a unique ground state, and H_P has a unique excited state, and the inner product $\langle g|\phi\rangle$ of those two states is $< \sqrt{1/2}$. This has been heavily confirmed by computer experiments by me involving over a million pairs of random 2×2 real symmetric matrices H_I and H_P. Also Kieu (§IV.B of [7]) has given a *proof* of the $n = 2$ case. Although I do not fully follow that proof, it is plausibly correct.
4. The half-claim also appears, in less-extensive computer experiments only involving a few thousand real symmetric 3×3 matrices, to be true for 3-state systems. That is, if H_I and H_P are $n \times n$ Hermitian matrices with $n = 3$, such that H_I has a unique ground state and H_P has two distinct unique excited states, and the two inner products $\langle g|\phi_j\rangle$ of the two H_P excited states ϕ_1 and ϕ_2 with the H_I ground state g both are $< \sqrt{1/2}$, then, at least in all computer experiments so far, neither H_P excited state ever gets occupation probability exceeding 1/2.
5. For any *finite* n, the *modified* version of the half-claim with some computable number $1 - \epsilon$ replacing the "1/2" is true. That is, given a number ϵ with $0 < \epsilon < 1$, an integer $n \ge 2$, and two $n \times n$ Hermitian matrices H_P and H_I with discrete spectra and unique ground states, and such that each inner product between the ground eigenstate g of H_I and any nonground eigenstate ϕ of H_P obeys $|\langle g|\phi\rangle|^2 \le 1/2$, there is an algorithm to compute a positive number T_0 such that Schrödinger time-evolution under $\mathfrak{H}$, starting from the ground state of H_I, from $t = 0$ to $t = T$, for any $T > T_0$, followed

[1] This comes at the cost of making the physical interpretation less attractive and less realistic-sounding.
[2] E.g, on page 11 of quant-ph/0504101 and also on page 14 of quant-ph/0310052v2.
[3] Were x and p finite matrices, the trace of $[x, p]$ would vanish since $\mathrm{tr}(x\,p) = \mathrm{tr}(p\,x)$ while the trace of the identity matrix does not vanish.

by a measurement of H_P, will yield a nonground energy with probability $< \epsilon$. This may be proven by computing the spectra and norms of the matrices and applying [1].

2 A counterexample

However, the half-claim is false for 4, 5, 6, and 7-state systems,[4] and very high occupation probabilities for nonground decoy states can be attained ($> 97\%$). It seems very probable that decoy probabilities *quickly approaching* 1 can be attained by making the matrix dimension large enough, in which case Kieu's plan (involving *infinite* dimensionality) is dead.

Here is a 4-state counterexample found by random trial. (Among pairs of symmetric matrices filled with random normal deviates and which meet the eigenvector correlation< 0.6 condition, each 25 attempts provides on average about 1 counterexample.) We employ natural physical units in which $\hbar = 1$.

$$H_I = \begin{pmatrix} -1.0820 & -0.3509 & 1.8919 & 0.1436 \\ -0.3509 & 0.2864 & -0.5115 & -1.3299 \\ 1.8919 & -0.5115 & 0.4458 & -0.0429 \\ 0.1436 & -1.3299 & -0.0429 & -1.4297 \end{pmatrix}, \quad H_P = \begin{pmatrix} -3.0262 & 0.2740 & -0.4055 & -0.5971 \\ 0.2740 & -5.4990 & 0.6990 & -1.3352 \\ -0.4055 & 0.6990 & -0.3468 & 0.3762 \\ -0.5971 & -1.3352 & 0.3762 & 0.8015 \end{pmatrix}, \quad T = 10.0620 \tag{2}$$

The eigendecomposition of H_P is $Q_P H_P = \Lambda_P Q_P$ where

$$\Lambda_P = \begin{pmatrix} -5.8893 & 0 & 0 & 0 \\ 0 & -3.1501 & 0 & 0 \\ 0 & 0 & -0.2572 & 0 \\ 0 & 0 & 0 & 1.2261 \end{pmatrix}, \quad Q_P = \begin{pmatrix} -0.0719 & 0.9802 & 0.0863 & -0.1632 \\ 0.9681 & 0.0567 & -0.1696 & -0.1757 \\ -0.1406 & 0.1066 & -0.9655 & 0.1917 \\ 0.1947 & 0.1571 & 0.1779 & 0.9517 \end{pmatrix} \tag{3}$$

and the columns of Q_P are the (unit normalized) eigenvectors of H_P. The corresponding eigendecomposition of H_I is

$$\Lambda_I = \begin{pmatrix} -2.4202 & 0 & 0 & 0 \\ 0 & -2.1117 & 0 & 0 \\ 0 & 0 & 0.7050 & 0 \\ 0 & 0 & 0 & 2.0474 \end{pmatrix}, \quad Q_I = \begin{pmatrix} 0.7295 & 0.4011 & 0.2481 & -0.4953 \\ -0.1979 & 0.4473 & 0.7495 & 0.4462 \\ -0.5228 & -0.1942 & 0.4082 & -0.7227 \\ -0.3941 & 0.7754 & -0.4584 & -0.1822 \end{pmatrix} \tag{4}$$

The matrix of inner products between the eigenvectors of H_I and H_P is

$$Q_I^\dagger Q_P = \begin{pmatrix} -0.2473 & 0.5862 & 0.5311 & -0.5596 \\ 0.5824 & 0.5197 & 0.2843 & 0.5567 \\ 0.5611 & 0.2572 & -0.5813 & -0.5302 \\ 0.5337 & -0.5659 & 0.5470 & -0.3094 \end{pmatrix} \tag{5}$$

making it plain that every inner product obeys $|\langle g|\phi\rangle| < 0.59$.

Evolution of $\vec{\psi}$ from $t = 0$ to $t = T$ under the time-dependent Schrödinger equation

$$i\frac{\partial}{\partial t}\vec{\psi} = \mathfrak{H}(t)\vec{\psi} \tag{6}$$

starting from the unique ground state $\vec{\psi}(0) = (0.7295,\ -0.1979,\ -0.5228,\ -0.3941)$ of H_I (with energy -2.4202), yields a final state $\vec{\psi}$ whose inner products with the columns of Q_P are respectively

$$\vec{\psi}(T)^\dagger Q_P = (0.2952,\ 0.9550,\ 0.0265,\ 0.0061) \tag{7}$$

corresponding to an occupation probability of $0.9550^2 \approx 0.9120$ in the "decoy" first excited state, with energy -3.1501, as compared to the occupation probability of only $0.2952^2 \approx 0.08714$ for the ground state, with energy -5.8893. The expected energy $\vec{\psi}^\dagger H_P \vec{\psi}$ of this final state is -3.3867.

This final state was deduced numerically via

$$\vec{\psi}(T) \ \approx \ \prod_{k=0}^{1023} \exp\left[\frac{iT}{1024}\mathfrak{H}\left(\frac{k+1/2}{1024}T\right)\right] \ \vec{\psi}(0) \tag{8}$$

[4] And, obviously, for n-state systems for each $n \geq 4$, because we can simply add new states to the 4-state counterexample, which are never occupied.

where the product is of 1024 exact matrix exponentials. The program that did it was written in `matlab`. An independently written evolution code by Kieu that employs `mathematica`'s `NDSolve` command to solve the Schrödinger equation, was able to confirm one of my 4×4 counterexample's final $\vec{\psi}$ to about 1% accuracy, so there seems no doubt of its correctness.

I have similarly found counterexamples with $n = 5$, 6, and 7 by random trial, and doing so appears to get easier as n increases, for example with $n = 7$, fewer random trials found a better counterexample, achieving $> 97\%$ probability for a decoy state.

My counterexamples *cannot be escaped* by demanding that the H_I and H_P have integer eigenvalues. (Because: multiply my H_I and H_P by 10^6 and then perturb them slightly to make all the eigenvalues be integer.) They also cannot be escaped by demanding that H_I and H_P be nonnegative- or positive-definite. (Because: add a large multiple of the identity matrix to each matrix, then quantum evolution will be the same except for a phase angle factor.) Finally, they also cannot be escaped by demanding that H_I be diagonal, since we can change bases to make it diagonal. Even combining all of these demands does not allow escape.

3 A stronger counterexample

But conceivably the above counterexample could be escaped by demanding that H_I be some particular and especially nice matrix such as $H_I = \mathrm{diag}(0, 1, 2, 3)$ – which happens to be exactly equivalent to the H_I Kieu was planning to use, expressed in its own eigenbasis and truncated to 4 dimensions. Also, it might be wondered what is the maximum achieveable decoy occupation probability. If instead of Kieu's 1/2 it was bounded by some other absolute constant below 1, such as 16/17, then his whole plan could still be modified to make it work. But, at least for this particular H_I choice, escape is ruled out by the following counterexample, and it achieves decoy probability very close to 1, namely $> 99.999\%$. (This example was found by adding a crude numerical optimizer on top of the random trials.) Note that 99.999% effectively *is* 1 to within the limited accuracy of the numerical time-evolver, even with 2048 timesteps which is what is now being used. I do not know whether exactly 1 is achievable[5] in any finite dimension n, but this makes it plausible that either (a) it is, or (b) probability 1 is approachable very quickly as $n \to \infty$.

$$H_I = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 3 \end{pmatrix}, \quad H_P = \begin{pmatrix} -0.5289 & 0.2834 & -0.7409 & -0.6673 \\ 0.2834 & -0.9560 & -1.2876 & 1.1387 \\ -0.7409 & -1.2876 & -1.7859 & 0.6167 \\ -0.6673 & 1.1387 & 0.6167 & -1.8111 \end{pmatrix}, \quad T = 7.8102 \tag{9}$$

The eigendecomposition of H_P is $Q_P H_P = \Lambda_P Q_P$ where

$$\Lambda_P = \begin{pmatrix} -3.5542 & 0 & 0 & 0 \\ 0 & -1.9058 & 0 & 0 \\ 0 & 0 & -0.0811 & 0 \\ 0 & 0 & 0 & 0.4591 \end{pmatrix}, \quad Q_P = \begin{pmatrix} 0.0359 & 0.5894 & 0.5465 & 0.5939 \\ -0.5535 & -0.0225 & -0.5841 & 0.5932 \\ -0.5919 & 0.5955 & -0.0134 & -0.5430 \\ 0.5847 & 0.5454 & -0.6000 & -0.0245 \end{pmatrix} \tag{10}$$

and the columns of Q_P are the (unit normalized) eigenvectors of H_P. The corresponding eigendecomposition of H_I involves $\Lambda_I = H_I$ and Q_I being just the 4×4 identity matrix. The matrix of inner products between the eigenvectors of H_I and H_P then is just Q_P, making it plain that every inner product obeys $|\langle g|\phi\rangle| \le 0.60$. Notice also that this matrix seems to contain visible structure.

Evolution of $\vec{\psi}$ from $t = 0$ to $t = T$ under the time-dependent Schrödinger equation starting from H_I's unique ground state $\vec{\psi}(0)$ with energy 0 yields a final state $\vec{\psi}$ whose inner products with the columns of Q_P are respectively

$$\vec{\psi}(T)^\dagger Q_P = (0.0009,\ 1.0000,\ 0.0011,\ 0.0017) \tag{11}$$

corresponding to an occupation probability of $> 99.9995\%$ in the "decoy" first excited state, with energy -1.8995, as compared to the occupation probability of only $0.0009^2 \approx 0.00000081$ for the ground state, with energy -3.5593. The expected energy $\vec{\psi}^\dagger H_P \vec{\psi}$ of this final state is -1.9056.

4 Still another counterexample

Kieu in email still objected that the above counterexamples might not logically suffice to kill his hypercomputer scheme, because his H_I and H_P have special structure which the more random matrices in our counterexamples, do not possess.

Therefore we now devise a 5-state system H_I and H_P which both *exactly* agree with those arising from Kieu's construction (for a certain 1-variable Diophantine problem) and exactly in his "$|n\rangle$ basis." The only thing different is that we truncate all the matrices down to 5 dimensions, i.e. only use the 5-state basis $\{|0\rangle, |1\rangle, |2\rangle, |3\rangle, |4\rangle\}$.

[5] And it might be very hard to settle that question.

Here is Kieu's H_I (exactly his own in the case when we set all his $\alpha_k = 1$) and our H_P:

$$H_I = \begin{pmatrix} 1 & -1 & 0 & 0 & 0 \\ -1 & 2 & -\sqrt{2} & 0 & 0 \\ 0 & -\sqrt{2} & 3 & -\sqrt{3} & 0 \\ 0 & 0 & -\sqrt{3} & 4 & -\sqrt{4} \\ 0 & 0 & 0 & -\sqrt{4} & 5 \end{pmatrix}, \quad H_P = \begin{pmatrix} 2 & 0 & 0 & 0 & 0 \\ 0 & 4 & 0 & 0 & 0 \\ 0 & 0 & 5 & 0 & 0 \\ 0 & 0 & 0 & 3 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix} \tag{12}$$

The eigenvectors of H_I are the columns of

$$Q_I = \begin{pmatrix} -0.6198 & -0.6541 & -0.4122 & 0.1336 & 0.0162 \\ -0.6127 & 0.0855 & 0.6350 & -0.4527 & -0.0959 \\ -0.4232 & 0.5151 & 0.0487 & 0.6701 & 0.3226 \\ -0.2300 & 0.4861 & -0.5055 & -0.1675 & -0.6536 \\ -0.0922 & 0.2513 & -0.4111 & -0.5478 & 0.6777 \end{pmatrix} \tag{13}$$

with corresponding eigenvalues (energy levels)

$$0.0114,\ 1.1307,\ 2.5406,\ 4.3884,\ 6.9288 \tag{14}$$

After time evolution, starting from the ground state of H_I at $t = 0$, to $t = T = 13.3444$ we get this final state $\vec{\psi}$ (absolute values of the $\vec{\psi}$ entries are shown, ordered in the same order as the energies $1, 2, 3, 4, 5$ of the H_P eigenstates)

$$(0.0139,\ 0.9997,\ 0.0062,\ 0.0210,\ 0.0015). \tag{15}$$

which note has probability $> 99.9\%$ of being measured as the first *excited* state of H_I, with energy 2, instead of the ground state with energy 1. The expected final energy is 2.0007.

Incidentally, in this example, throughout the evolution from H_I to H_P, the Hamiltonian $\mathfrak{H}$ has 5 *distinct* energy levels which never cross. Kieu has confirmed the 0.999 decoy probability in this example with `mathematica`, and made a plot of the 5 eigenenergy levels versus t for $0 \le t \le T \approx 13.34$, which confirmed my claim that they never cross throughout the homotopy. However, Kieu found that at $t \approx 10.37$ there is something that looks very much *like* a crossing of the ground and first excited levels! When we examine the picture with high magnification (see figures at end) we discover that this is *not* a crossing! Kieu believes that it is during this near-crossing that a "massive transfer of occupation probability from the ground state to the 1st excited state" occurs.

These counterexamples are interesting not only because they refute Kieu's hypercomputer – they also cast a great deal of light on the (bad) behavior of quantum adiabatics.

5 Some additional remarks

Another misconception in Kieu's papers is the notion that Turing's proof of the nonexistence of an algorithm to solve the halting problem, still might allow a *probabilistic* algorithm for that purpose. In fact

Theorem 1 (Halting remains undecidable even for probabilistic decision procedures). *Let a "probabilistic decision procedure" (PDP) be a Turing-machine program that accepts both (1) a finite input, and (2) an unbounded string of random bits as a second kind of input, and terminates in finite expected time no matter what input1 is with output "yes" or "no" with that yes/no being correct for whatever input1 it is, with probability* $> 2/3$. *Then: there is no probabilistic decision procedure for the halting problem.*

Proof. We shall prove that any PDP may be converted into an ordinary deterministic decision procedure; and Turing proved the impossibility of that. Simply run the PDP on every possible n-bit "random" string (all 2^n of them) as input2 and count the number of "yes," "no," and "not done yet – need more random bits" results that occur. Keep increasing n and redoing this until either "yes" or "no" happens $> 99\%$ of the time ("not done" happens $< 1\%$). This must happen for some n, due to the finite expected run time of the PDP. Now if $> 50\%$ of the answers are yes, then output "yes," otherwise output "no." Due to the $> 2/3$ correctness probability of the PDP this will yield 100% accuracy. □

The correction of this misconception actually would have made Kieu's (false) claim that his physical system yields such a PDP *more* impressive. Even if Kieu's claim had been correct – i.e. if he really had a physical system to solve Diophantine equations – there would be another problem:

Theorem 2 (Infinite manufacturing precision). *It is not possible to solve general Diophantine systems, even polynomial ones, with any Kieu-like approach unless the Hamiltonian is manufactured infinitely precisely.*

Proof.[6] Consider the two-equation Diophantine system

$$x^2 - 2y^2 = 0, \qquad x = a^2 + b^2 + c^2 + d^2 + 1. \tag{16}$$

This system has *no* solutions because $\sqrt{2}$ is irrational. (We have implicitly used the well known Lagrange theorem that any natural number is a sum of 4 squares.)

However if the coefficient "2" were replaced by any squared rational number, then the resulting system would have an *infinite* number of solutions (a, b, c, d, x, y). There are an infinite number of squared rational numbers arbitrarily near 2. Therefore, in order to solve this Diophantine system by physical methods, we would need to implement the Hamiltonian (i.e. all the beam splitters and Kerr nonlinear media, or whatever physical objects are used to produce this Hamiltonian) with *infinite* manufacturing precision − presumably impossible. □

For this reason, had Kieu succeeded I would have preferred to call that success an "unsimulable quantum adiabatic system" rather than a "quantum adiabatic hypercomputer."

6 What can be salvaged?

First of all, finding the energies of ground states is one of the most fundamental computational tasks in quantum mechanics.

Theorem 3 (Ground energies are uncomputable). *No algorithm exists which, given a description of the Hamiltonian H of a quantum system in a Cartesian product of K Fock spaces (for some finite K; and the "description" could be a finite-length formula in terms of the K kinds of raising and lowering operators a and $a^\dagger$), finds its ground energy accurate to within $\pm 1/3$. This remains true even if H is known to have a unique ground state (or if H is known to have an infinitely nonunique ground state – either demand works), and to have only natural numbers as energy levels; and then even just the question of whether the ground energy is zero or ≥ 1, is undecidable.*

Proof sketch. If such an algorithm were available, then via Kieu's Diophantine↔Hamiltonian correspondence, it could solve singlemin 2-exponential (respectively polynomial) Diophantine equations – a task which Smith (respectively Matiyasevic) proved to be undecidable. □

One way to find the ground state of a system is to create the system, let it "cool down" by allowing it to interact weakly with a very cold heat bath (whose temperature is well below the lowest excited energy level of the system), or let it "decay" by allowing to interact with a quantum radiation field into which it can "emit photons," and then measure its energy.

The amount of time this takes is the "halflife" of the system. (For the "heat bath" one could simply use another Fock space.)

It is a common observation that halflives can be very long. For example, the ground state of carbon is believed to be graphite. Yet tiny diamonds (2nm diameter) are common in meteorites and evidently have lasted billions of years. We have:

Theorem 4 (Halflives are uncomputably large). *Either: No algorithm exists which, given a description of the Hamiltonian H of a quantum system in a Cartesian product of K Fock spaces (for some finite K; and the "description" could be a finite-length formula in terms of the K kinds of raising and lowering operators a and $a^\dagger$), and a description of its initial density matrix, finds an upper bound on its (radiative, or thermal) halflife.*

Or: Church's thesis (that quantum systems with specified Hamiltonians are simulable for finite times to arbitrary specified accuracy) is false for our systems.

Another way to find ground states is via the "quantum adiabatic theorem" (QAT).

The QAT dates to P.Ehrenfest in 1916 and is of basic importance, especially in experimental physics. According to [2], all versions of the QAT proved before 1981 involved *bounded* Hamiltonian operators. If physics is regarded as having an *un*bounded Hamiltonian (and, e.g., the usual flat-space quantum field theories, and the usual treatments of the "particle in a box," the "simple harmonic oscillator" and the "hydrogen atom" found in undergraduate textbooks all involve unbounded infinite dimensional Hamiltonians) we can conclude that every use of the QAT before 1981 was, technically speaking, unjustified either mathematically or physically. Avron et al [2] published a proof of a QAT version allowing unbounded Hamiltonians in 1987. However, neither this, not any previous proof, gave any explicit or rigorous bound on the convergence time. That complaint was addressed by Ambainis & Regev [1] who in 2004 gave an elementary proof of a QAT version which did give a rigorous bound relating ϵ to the time T required to force ground state probability $> 1 - \epsilon$. However, this bound is useless (since in general it merely says "$T \leq \infty$") if the Hamiltonians are unbounded, as in textbook physics.[7]

So then a fundamental question is: can [2] and [1] somehow be combined to yield an a priori bound on T, applicable even for arbitrary unbounded Hamiltonians? The answer appears to be *no*.

[6] A similar theorem and proof was invented independently and earlier by Andrew Hodges, best known as the author of Turing's biography [4].

[7] In email to Ambainis, I suggested that apparently his bound on T could be re-expressed not in terms of the *operator* norms of the Hamiltonians H, but instead in terms of the maximum *vector* norm of $H\psi$ over all states ψ arising during the adiabatic evolution. This would be a simple improvement that would allow unbounded Hamiltonians to be used, *but* only in an *a posteriori* fashion, because it only can be used if the evolution of $\psi(t)$ is already known.

Theorem 5 (There is no computable bound on the adiabatic time). *If the Church thesis (that the evolution for finite times of finite-energy quantum adiabatic systems is algorithmically simulable) is correct, then there DOES NOT EXIST any algorithm which, given a description of two self-adjoint nonnegative definite operators H_I and H_P with discrete spectra (indeed, one may demand natural number valued spectra) and unique ground states (with a known ground state of energy* 0 *in the case of H_I), computes an ϵ with $0 < \epsilon < 1$ and a time $T > 0$ such that quantum adiabatic evolution under $\mathfrak{H}(t)$ from $t = 0$ to $t \geq T$ starting from the known ground state of H_I, will yield an occupation probability for the ground state of H_P that exceeds $1 - \epsilon$.*

Proof sketch. If such an algorithm were available, then Kieu's plan, as corrected by Smith as sketched in the present paper, with experiments running for that time T, would suffice to yield a PDP for the halting problem. But we have shown in theorem 1 that no such PDP exists. □

These are the first examples we know of where uncomputably large numbers arise naturally in theoretical physics. This suggests that there is a good reason why the many provers of the QAT during the period 1916-2004 have never been able to produce a useful upper bound on the convergence time; the best they've been able to to is merely to prove convergence occurs.

Fundamental open questions.

1. Exactly what class of Hamiltonians is physically realizable?
2. Exactly what class of Hamiltonians is algorithmically simulable? Note that Smith [8] showed that Hamiltonians for n particles in 3-space with a wide class of interparticle potentials (including Coulombic) were algorithmically simulable, and that the eigenenergies of such systems were computable real numbers. But the Hamiltonians of Kieu's type, constructible as formulas in terms of raising and lowering operators $a^\dagger$ and a in Fock spaces (and when these formulae of polynomials, these Hamltonians arguably are physically realizable) have uncomputable ground energies. The problem is to draw that borderline more precisely.
3. Why do the counterexamples start working for 4-state systems and not 2 or 3 – what is so special about 4?
4. Can such counterexamples exactly achieve decoy probability=1? (These are both fundamental questions about quantum adiabatics.)

Finally, let me remark that all of the claims made here that are insufficiently justified here (in the desire to keep this piece short) *are* justified with full formal mathematical rigor, in a $>$ 20-page manuscript jointly authored by Kieu and Smith. (This includes: the construction of uncomputable singlemin Diophantines, the rigorous demonstration that the QAT applies to Kieu's systems, and the rigorous probabilistic analysis.) That manuscript is not yet publically available for various reasons (such as that Kieu objects to its present form) but probably will eventually surface.

References

[1] Andris Ambainis and Oded Regev: An Elementary Proof of the Quantum Adiabatic Theorem, Nov. 2004, quant-ph/0411152v1.

[2] J. Avron, R. Seiler, L. G. Yaffe: Adiabatic theorems and applications to the quantum Hall effect, Commun. Math. Phys., 110,1 (1987) 33-49.

[3] Martin Davis, Yuri Matijasevic, Julia Robinson: Hilbert's Tenth Problem: Diophantine Equations: Positive Aspects of a Negative Solution, Proceedings of Symposia in Pure Mathematics 28 (1976) 323-378 (Mathematical developments arising from Hilbert Problems); reprinted in The Collected Works of Julia Robinson (S.Feferman, ed.) Amer. Math. Soc. 1996, pp.269-.

[4] Andrew Hodges: Alan Turing: The Enigma, Simon & Schuster 1983.

[5] Wassily Hoeffding: Probability inequalities for sums of bounded random variables, J. Amer. Statist. Assoc. 58,301 (March 1963) 13-30.

[6] J.P. Jones & Y.V. Matijasevic: Register Machine Proof of the Theorem on Exponential Diophantine Representation of Enumerable Sets, Journal of Symbolic Logic 49,3 (1984) 818-829.

[7] Tien D. Kieu: Quantum adiabatic algorithm for Hilbert's tenth problem: I. The algorithm, Oct. 2003, quant-ph/0310052v2. The other quant-ph papers by Kieu on this (or very related) subjects are: 0110136, 0111020, 0111062, 0111063, 0203034, 0205093, 0304114, 0407090, 0504101.

[8] Warren D. Smith: Church's thesis meets quantum mechanics, #49 on `http://math.temple.edu/~wds/homepage/works.html`.

FIGURE CAPTION. Two plots provided by Tien D. Kieu showing the 5 eigenenergies versus t in the example of §4. The top plot exhibits what appears to be a level-crossing near $t \approx 10.37$, but in the bottom plot heavy magnification shows that to be an illusion.