



CHORUS

This is the accepted manuscript made available via CHORUS. The article has been published as:

Fast quantum computation at arbitrarily low energy

Stephen P. Jordan

Phys. Rev. A **95**, 032305 — Published 6 March 2017

DOI: [10.1103/PhysRevA.95.032305](https://doi.org/10.1103/PhysRevA.95.032305)

Fast quantum computation at arbitrarily low energy

Stephen P. Jordan^{1,2}

¹*National Institute of Standards and Technology, Gaithersburg, MD*

²*Joint Center for Quantum Information and Computer Science (QuICS),
University of Maryland, College Park, MD*

One version of the energy-time uncertainty principle states that the minimum time T_{\perp} for a quantum system to evolve from a given state to any orthogonal state is $h/(4\Delta E)$ where ΔE is the energy uncertainty. A related bound called the Margolus-Levitin theorem states that $T_{\perp} \geq h/(2\langle E \rangle)$ where $\langle E \rangle$ is the expectation value of energy and the ground energy is taken to be zero. Many subsequent works have interpreted T_{\perp} as defining a minimal time for an elementary computational operation and correspondingly a fundamental limit on clock speed determined by a system's energy. Here we present local time-independent Hamiltonians in which computational clock speed becomes arbitrarily large relative to $\langle E \rangle$ and ΔE as the number of computational steps goes to infinity. We argue that energy considerations alone are not sufficient to obtain an upper bound on computational speed, and that additional physical assumptions such as limits to information density and information transmission speed are necessary to obtain such a bound.

I. INTRODUCTION

It is frequently argued that energy places a fundamental limit on the speed of computation, even if the computation is performed reversibly. Planck's constant provides a conversion factor between energy and frequency, and on dimensional grounds one might expect that a quantum system at energy scale E has maximum computational clock speed of order E/h . More quantitatively, for a quantum system with energy uncertainty ΔE , the time T_{\perp} to evolve to an orthogonal quantum state obeys the bound [1–5]

$$T_{\perp} \geq \frac{h}{4\Delta E}. \quad (1)$$

Several related energetic lower bounds on T_{\perp} have been proven [6–11]. In particular, the Margolus-Levitin theorem [7] shows that for a state with energy expectation $\langle E \rangle$, evolving according to a time-independent Hamiltonian with zero ground energy,

$$T_{\perp} \geq \frac{h}{2\langle E \rangle}. \quad (2)$$

See [12] for a recent review of such quantum speed limits and their applications. Many estimates of computational capacity of physical systems have their starting point in the assumption that T_{\perp}^{-1}

can be interpreted as a maximum computational clock speed [11, 13–25]. Related, but distinct, arguments for limits on computational speed related to energy are given in [26–29]. Another relationship between the energy-time uncertainty principle and computational complexity, which is not based on associating T_{\perp} with the time necessary for a logical operation, was recently given in [30].

For several reasons, equating T_{\perp}^{-1} with a maximum computational clock speed seems quite plausible. In classical computers, each logic gate brings the system to an orthogonal state, and in quantum circuits, each logic gate typically [44] brings the system to a near-orthogonal state. T_{\perp} is the minimum time to flip a qubit from $|0\rangle$ to $|1\rangle$ and it would seem surprising to achieve G logical gates in time less than GT_{\perp} . Furthermore, energy-time uncertainty principles suggest that the uncertainty of the timing of a logic operation scales as $1/\Delta E$, and if the time between elementary logic operations is shorter than this then one would expect their ordering to be uncertain, thereby ruining the computation. (For a contrary conjecture, perhaps presaging the present work, see [31].)

However, we here show that these obstacles can be evaded. We give explicit constructions of quantum time-evolutions using time-independent Hamiltonians, which simulate the operation of a quantum circuit of G elementary gates, while traversing only a constant number of orthogonal states, independent of G . The Hamiltonians achieving this involve only 4-qubit interactions and can be made spatially local in two dimensions. The ratio of the computational clock speed to ΔE is unbounded, specifically growing linearly with G . The same holds for $\langle E \rangle$. In other words, rather counterintuitively, the total time needed to simulate G gates is on the same order as the minimum required time to flip a single bit. We conclude that energy alone does not present a fundamental limit to computational speed; to obtain such a limit one must invoke additional assumptions such as a limit to the spatial density at which qubits can be packed. The argument based on energy-time uncertainty principles is evaded because, although the uncertainty of the timing of each individual logic operation is very large in our construction, these timings are all correlated, leaving no ambiguity as to the ordering of the operations.

II. BASIC CONSTRUCTION

Here we show how to simulate an arbitrary quantum circuit of G gates by evolving for some time T according to a time-independent 4-local Hamiltonian. The computational clock speed achieved $f_{\text{clock}} = G/T$ is such that the ratios $f_{\text{clock}}/\langle E \rangle$ and $f_{\text{clock}}/\Delta E$ both diverge linearly with G . This construction thus suffices to serve as a counterexample for the conjectures that

computational clock speed is limited to some maximum rate proportional to either the energy uncertainty or the energy above the ground state. To those familiar with Feynman-Kitaev clock Hamiltonians, the essential idea of the construction can be concisely expressed: it is to initialize a Feynman-Kitaev-type Hamiltonian with a wavepacket whose breadth is comparable to the number of gates G . Such broad wavepackets have small expected energy and low energy uncertainty. Furthermore, they maintain large overlap with previous states as they propagate, and thus the number of orthogonal states traversed during the time evolution is $O(1)$ even as the number of computational steps G is increased. In the remainder of this section, the construction is given in detail without assuming prior familiarity with the Feynman-Kitaev construction. Note that here we are using the Feynman-Kitaev Hamiltonian to execute quantum computation “ballistically” as in Feynman’s original construction [32] not adiabatically as in [33].

Consider a quantum circuit U consisting of a sequence of 2-qubit gates $U = U_G U_{G-1} \dots U_1$ acting on n qubits. For a given initial state $|h_0\rangle$ the circuit proceeds through the states $|h_0\rangle \rightarrow |h_1\rangle \rightarrow \dots \rightarrow |h_G\rangle$ where

$$\begin{aligned} |h_1\rangle &= U_1 |h_0\rangle \\ |h_2\rangle &= U_2 U_1 |h_0\rangle \\ &\vdots \\ |h_G\rangle &= U_G \dots U_1 |h_0\rangle. \end{aligned} \tag{3}$$

The standard Feynman-Kitaev clock Hamiltonian [32, 34] acts on a register of n computational qubits, together with a clock register, as

$$H^{\text{FK}} = \sum_{x=1}^G \left(-U_x \otimes |x\rangle\langle x-1| - U_x^\dagger \otimes |x-1\rangle\langle x| + \mathbb{1} \otimes |x\rangle\langle x| + \mathbb{1} \otimes |x-1\rangle\langle x-1| \right), \tag{4}$$

where $\mathbb{1}$ denotes the identity matrix, the first tensor factor represents the computational qubits, and the second tensor factor represents the clock register. The subspace

$$\mathcal{C} = \text{span}\{|h_0\rangle|0\rangle, |h_1\rangle|1\rangle, \dots, |h_G\rangle|G\rangle\} \tag{5}$$

is preserved by H^{FK} and the block of H^{FK} acting on this subspace looks like

$$H^{\text{FK}}|_{\mathcal{C}} = \begin{bmatrix} 1 & -1 & & & & \\ -1 & 2 & -1 & & & \\ & -1 & 2 & -1 & & \\ & & & \ddots & & \\ & & & & -1 & 1 \end{bmatrix} \tag{6}$$

with all other entries zero. In other words, $H^{\text{FK}}|_{\mathcal{C}}$ is the discretized second derivative on a one-dimensional lattice. Formally, if we think of $H^{\text{FK}}|_{\mathcal{C}}$ as acting on a discretization of the unit interval $[0, 1]$ one has

$$\lim_{G \rightarrow \infty} G^2 H^{\text{FK}}|_{\mathcal{C}} = -\frac{d^2}{dx^2}. \quad (7)$$

Consequently, the dynamics induced by H^{FK} on the computational subspace \mathcal{C} can be intuitively thought of as the dynamics of a free particle on a line. A state $\sum_x \psi(x)|h_x\rangle|x\rangle$ will evolve according to [45]

$$i\frac{d}{dt}\psi(x, t) = H^{\text{FK}}\psi(x, t) \simeq -\frac{1}{G^2}\frac{d^2}{dx^2}\psi(x, t). \quad (8)$$

To perform computation, one can prepare an initial wavepacket at small x with momentum in the direction of increasing x . Under Schrödinger time evolution according to (8) this wavepacket will propagate to larger x and broaden. Eventually, the wavepacket will reach the end of the line, and consist of a superposition with nonzero amplitude on $|\psi_G\rangle$, which is the output of the original quantum circuit.

One might be concerned about two apparent problems with the Feynman-Kitaev construction. Firstly, the initial wavepacket must have some width large compared to the lattice spacing in order for the continuum approximation to be valid. Thus it has support not only on the initial state $|h_0\rangle \otimes |0\rangle$ but also on states $|h_1\rangle \otimes |1\rangle, |h_2\rangle \otimes |2\rangle, \dots$ in which some gates have already been applied. This seems like cheating; we are seeking to simulate the computation performed by the quantum circuit U , but we start with an initial state in which some of this computation has already been done. This problem is easily solved by padding the circuit U with sufficiently many initial identity gates so that the initial wavepacket only has support on states in which no nontrivial gates have been applied. Secondly, upon measuring the final state, the probability of obtaining the desired outcome $|h_G\rangle \otimes |G\rangle$ will be much smaller than one because the final wavepacket has nonzero amplitude on states $|h_{G-1}\rangle \otimes |G-1\rangle, |h_{G-2}\rangle \otimes |G-2\rangle, \dots$ in which not all of the gates have yet been applied. This second problem can be similarly solved by padding the circuit with sufficiently many identity gates at the end.

So far, we have not specified the physical implementation of the clock register. We have simply labeled an orthonormal basis $|0\rangle, |1\rangle, \dots, |G\rangle$ for its Hilbert space. We can implement this using

$G + 1$ qubits in the following encoding.

$$\begin{aligned}
 |0\rangle &\mapsto |100\dots 0\rangle \\
 |1\rangle &\mapsto |010\dots 0\rangle \\
 &\vdots \\
 |G\rangle &\mapsto |000\dots 1\rangle
 \end{aligned} \tag{9}$$

This encoding is inefficient in the sense that only $\log_2(G + 1)$ qubits are actually needed to store numbers in the range $\{0, 2, \dots, G\}$. However, the advantage of this encoding is that the operators such as $|x - 1\rangle\langle x|$ appearing H^{FK} only act on two qubits, namely qubits x and $x - 1$. Thus, if the original circuit U is constructed from a universal gate set of one-qubit and two-qubit quantum gates, then H^{FK} is a 4-local Hamiltonian.




FIG. 1: The circuit $U = U_G \dots U_1$ is padded by G identity gates at the beginning and end. The initial wavepacket has width of order G and support on the initial identity gates. The wavepacket then propagates rightward, ending with its support on the final segment of G identity gates.

Let's first treat broadening of the wavepacket as negligible and take the width of the wavepacket to be G , as illustrated in figure 1. (We defer analysis of broadening to §III.) The original circuit $U = U_G \dots U_1$ should be padded with G identity gates at the beginning and G identity gates at the end, yielding a circuit $U' = U'_{3G} U'_{3G-1} \dots U'_1$ with a total of $3G$ gates and

$$U'_x = \begin{cases} \mathbf{1} & 1 \leq x \leq G \\ U_{x-G} & G + 1 \leq x \leq 2G \\ \mathbf{1} & 2G + 1 \leq x \leq 3G. \end{cases} \tag{10}$$

The initial state can be chosen as a wavepacket with width on the order of G and support only on the initial G identity gates. Let $|h'_x\rangle = U'_x U'_{x-1} \dots U'_1 |h_0\rangle$ while retaining the notation $|h_x\rangle$ defined in (3) for the x^{th} state obtained by the original unpadded circuit. Thus the initial wavepacket is:

$$\sum_{x=0}^G \psi(x) |h'_x\rangle |x\rangle = |h_0\rangle \left(\sum_{x=0}^G \psi(x) |x\rangle \right). \tag{11}$$

That is, the computational register is initialized to the same state that the original circuit starts with, which could be taken as $|0\dots 0\rangle$ without loss of generality[46]. The wavepacket $\psi(x)$ should

be chosen with rightward momentum. Then it will propagate down the line, yielding a final state of

$$\sum_{x=0}^G \psi(x) |h'_{x+2G}\rangle |x+2G\rangle = |h_G\rangle \left(\sum_{x=0}^G \psi(x) |x+2G\rangle \right). \quad (12)$$

Thus, in the final state, the computational register contains the output of the original circuit.

From figure 1 one sees that, as the wavepacket propagates, it passes through essentially three orthogonal states: the initial state supported on 0 through G , the state supported on G through $2G$, and the final state supported on $2G$ through $3G$. The energy uncertainty in a state of spatial width G for the Hamiltonian H^{FK} scales as $\Delta E \sim 1/G^2$. The ground energy of the Feynman-Kitaev Hamiltonian is zero. One can verify this by noting that the state

$$\frac{1}{\sqrt{G+1}} \sum_{x=0}^G |h_x\rangle |x\rangle \quad (13)$$

is an energy-zero eigenvector of H^{FK} and that, by (4) and (9), H^{FK} is positive semidefinite since it is a sum of G positive-semidefinite terms.

Let \mathcal{G} denote the number of gates in the padded quantum circuit. (If the wavepacket did not spread then we could take $\mathcal{G} = 3G$ while still ensuring that the final superposition is only over states in which the computation is finished. In the present construction we will take \mathcal{G} slightly larger than $3G$, but still only linear in G , as discussed in §III.) Examining (8) with $G \mapsto \mathcal{G}$, one sees that, in the limit of large \mathcal{G} , the wavepacket propagates like a nonrelativistic particle of mass $m = \mathcal{G}^2/2$ on the unit interval. The velocity of such a particle (*i.e.* the group velocity of the wavepacket) is $\sim p/m \sim p/\mathcal{G}^2$, where p is the momentum of the wavepacket. The time needed to propagate to the end of the line is therefore $T \sim \mathcal{G}^2/p$. During this time, G gates are simulated, so the clock speed is $f_{\text{clock}} \sim G/T \sim Gp/\mathcal{G}^2$. As we show in §III, it suffices to choose \mathcal{G} to be a constant multiple of G . Thus, $f_{\text{clock}} \sim p/G$. The expectation value of the energy in the wavepacket state is $\langle E \rangle = p^2/2m \sim p^2/G^2$. Hence, by choosing $p = O(1)$ we can achieve a ratio of clock speed to energy of $f_{\text{clock}}/\langle E \rangle \sim G$. As mentioned above, with wavepacket of width of order G , the energy uncertainty is of order $1/G^2$. Hence, the ratio of clock speed to energy uncertainty is also of order G .

III. DISPERSION

Suppose the initial state is a Gaussian wavepacket

$$\psi(x, 0) = \eta \exp \left[-\frac{x^2}{2\sigma_x^2} + ip_0x \right] \quad (14)$$

where η is a normalization constant. Under Schrödinger's equation, $i\frac{d\psi}{dt} = -\frac{1}{2m}\frac{d^2}{dx^2}$, this evolves to[47]

$$\psi(x, t) = \eta(t) \exp \left[-\frac{(x - x_0(t))^2}{2\sigma_x(t)^2} + i\phi(x, t) \right] \quad (15)$$

where $\phi(x, t)$ is a (real) phase, $\eta(t)$ is a normalization factor, and

$$x_0(t) = \frac{p_0 t}{m} \quad (16)$$

$$\sigma_x(t) = \sqrt{\sigma_x^2 + \left(\frac{t}{m\sigma_x}\right)^2}. \quad (17)$$

In the units being used, where the length of the line is 1 and the lattice spacing is $1/\mathcal{G}$, we have, in the continuum limit (*i.e.* $\mathcal{G} \rightarrow \infty$), an effective particle mass of

$$m = \frac{\mathcal{G}^2}{2} \quad (18)$$

and a propagation velocity

$$v = \frac{p_0}{m} = \frac{2p_0}{\mathcal{G}^2} \quad (19)$$

Thus, the time to propagate down the line is $T \sim 1/v$ and $\sigma_x(T)$ is $O(1)$. In other words, the final superposition has width only a constant factor larger than the initial superposition. To obtain a wavepacket with finite support, one can truncate the Gaussian wavepacket at some multiple of σ_x away from the mean. That is,

$$\psi_{\text{trunc}}(x, 0) = \begin{cases} \eta \exp \left[-\frac{x^2}{2\sigma_x^2} + ip_0 x \right] & |x - x_0| \leq c\sigma_x \\ 0 & \text{otherwise.} \end{cases} \quad (20)$$

The inner product between this wavepacket and the untruncated Gaussian becomes exponentially close to one as the multiple c is increased. By unitarity, the inner product between the ideal and truncated final states will be equal to the inner product between the ideal and truncated initial states. Because initial and final σ_x are $O(1)$, the initial truncated wavefunction has support only on $O(\mathcal{G})$ clock values and can be completely accommodated by padding the circuit with $O(\mathcal{G})$ initial identity gates. The final state will, to an exponentially good approximation for large c , have all its amplitude within a range of $O(\mathcal{G})$ clock values and can also be accommodated by padding with $O(\mathcal{G})$ final identity gates.

As a concrete example, suppose we start with a circuit of G gates, and we pad it with $2G$ initial identity gates and $2G$ final identity gates. Thus, $\mathcal{G} = 5G$. We set the initial state to be a Gaussian superposition of width $\frac{\mathcal{G}}{4}$ and mean G . Thus, we can take $c = 4$ and have zero amplitude outside

the initial pad, whose width is $2G$. This ensures that the initial state has zero amplitude for any of the computational gates to be already computed. An initial superposition of width $\frac{G}{4}$ means, in units where the length of the line is one, $\sigma_x = \frac{1}{20}$. We want the wavepacket to propagate from the middle of the initial pad to the middle of the final pad. Thus the total distance to propagate is $\frac{3}{5}$. By (19) this will take time $T = \frac{3G^2}{10p_0}$. Choosing $p_0 = 240$ (with corresponding energy $\frac{p_0^2}{G^2}$) yields, by (17), a final wavepacket width of $\sigma_x(T) = \frac{\sqrt{2}}{20}$. Measurement of the clock register in the final wavepacket will yield a Gaussian probability distribution with mean $4G$ and standard deviation [48] $2 \times 5G\sigma_x(T) = \frac{1}{\sqrt{2}}G$. Almost all the probability thus lies in the final pad of identity gates covering clock values $3G$ to $5G$. Outcomes in which the clock register has value less than $3G$ correspond to events at least $\sqrt{2}$ standard deviations below the mean, which have total probability about 8%. Hence, including the errors from the truncation of the initial wavepacket, one finds that, upon measuring the final state, the computational register will with probability at least 88% be in a state where all G gates have been computed.

IV. DISTANCE TRAVERSED THROUGH HILBERT SPACE

The question whether the number of orthogonal states traversed in a computation is a resource bearing on computational capacity is perhaps interesting independent of connection to energy. It could be asked in the context of computational models where energy may not play a manifest role, such as quantum cellular automata, quantum Turing machines, and the quantum circuit model. To ask this question in a precise way, we first note that the number of orthogonal states traversed during a computation is not a well-behaved metric. Instead, for a discrete-time quantum computation, such as the quantum circuit model, we can formalize the intuitive notion of distance traversed through Hilbert space as

$$\mathcal{L} = \sum_{x=1}^G \left\| |h_x\rangle - |h_{x-1}\rangle \right\| \quad (21)$$

where $|h_x\rangle$ is the quantum state obtained after the first x gates have been applied, as in (3). The continuum analogue of this distance is

$$\mathcal{L}(|\psi(t)\rangle; t_2, t_1) = \int_{t_1}^{t_2} dt \left\| \frac{d|\psi(t)\rangle}{dt} \right\|, \quad (22)$$

which can be applied to Hamiltonian-based models of quantum computation. Note that the continuum version of \mathcal{L} remains invariant if we multiply the speed at which we traverse the path through Hilbert space by some factor, as is fitting for a metric of distance.

The construction of §II demonstrates that it is possible to simulate G gates while keeping $\mathcal{L} = O(1)$. One way to see this is by the following calculation.

$$\mathcal{L} = \int_0^T dt \left\| \frac{d}{dt} |\psi\rangle \right\| \quad (23)$$

$$= T \|H|\psi\rangle\| \quad (24)$$

$$= T \sqrt{\langle \psi | H^2 | \psi \rangle} \quad (25)$$

$$= T \sqrt{(\Delta E)^2 + \langle E \rangle^2} \quad (26)$$

$$= O(1) \quad (27)$$

The last line follows from the results of §II stating that $T = O(G^2)$, $\Delta E = O(1/G^2)$, and $\langle E \rangle = O(1/G^2)$.

V. DISPERSIONLESS DISCRETIZATION

In the construction of §II, the dispersion relation is quadratic and therefore the wavepacket spreads as it propagates. Here we give an alternative construction which avoids this complication by modifying the Feynman-Kitaev Hamiltonian to yield a linear dispersion. This modified construction achieves constant clock speed with energy uncertainty scaling as $\Delta E \sim 1/G$. This is thus a faster method of computation than the basic construction, which achieves $O(1/G)$ clock speed at $O(1/G^2)$ energy uncertainty[49]. Furthermore, in the linear-dispersion construction $\lim_{G \rightarrow \infty} \mathcal{L}$ is easy to calculate in complete quantitative detail. However, it uses wavepackets whose energy is not close to the ground energy. Depending on context this may or may not be relevant, which is why both constructions are presented in this manuscript. Roughly speaking, linear dispersion is achieved by discretizing a one-dimensional analogue of Dirac's equation rather than discretizing the one-dimensional Schrödinger equation. Similar ideas have been used previously in [27, 35].

Consider the Hamiltonian

$$H = v \begin{bmatrix} 0 & -\frac{d}{dx} \\ \frac{d}{dx} & 0 \end{bmatrix} \quad (28)$$

where v is a “velocity.” This is Hermitian because $\frac{d}{dx}$ is antihermitian. Schrödinger's equation then reads

$$\frac{d}{dt} \begin{bmatrix} \Phi \\ \Psi \end{bmatrix} = -iv \begin{bmatrix} 0 & -\frac{d}{dx} \\ \frac{d}{dx} & 0 \end{bmatrix} \begin{bmatrix} \Phi \\ \Psi \end{bmatrix}. \quad (29)$$

Consequently,

$$\frac{d^2}{dt^2} \begin{bmatrix} \Phi \\ \Psi \end{bmatrix} = v^2 \begin{bmatrix} \frac{d^2}{dx^2} & 0 \\ 0 & \frac{d^2}{dx^2} \end{bmatrix} \begin{bmatrix} \Phi \\ \Psi \end{bmatrix}. \quad (30)$$

That is, Φ and Ψ each obey the one-dimensional wave equation. A solution to the Schrödinger equation is therefore

$$\begin{bmatrix} \Phi \\ \Psi \end{bmatrix} = \begin{bmatrix} w(x - vt) \\ iw(x - vt) \end{bmatrix} \quad (31)$$

for any function w , as can easily be verified. w describes the shape of a wavepacket that rigidly propagates in the positive- x direction without distortion.

We can now discretize $\frac{d}{dx}$ and $-\frac{d}{dx}$ using finite differences while maintaining Hermiticity. To achieve this, we use a forward difference to discretize $\frac{d}{dx}$ and a backward difference to discretize $-\frac{d}{dx}$ as illustrated by the following example on a lattice of four sites.

$$H = \begin{bmatrix} 0 & 0 & 0 & 0 & 1 & -1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & -1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & -1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ -1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & -1 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & -1 & 1 & 0 & 0 & 0 & 0 \end{bmatrix} \quad (32)$$

Using $|\Phi_0\rangle, \dots, |\Phi_{G/2}\rangle$ as a basis for the Φ subspace and $|\Psi_0\rangle, \dots, |\Psi_{G/2}\rangle$ for the Ψ subspace, one has

$$\begin{aligned} \langle \Phi_x | H | \Psi_x \rangle &= 1 \\ \langle \Phi_x | H | \Psi_{x-1} \rangle &= -1 \end{aligned} \quad (33)$$

with all other matrix elements zero other than the Hermitian conjugates of the above.

The linear-dispersion Feynman-Kitaev computational Hamiltonian corresponding to this discretization is

$$H_{\text{lin}} = \sum_{x=1}^{G/2} \left[U_{2x}^\dagger \otimes |\Phi_x\rangle\langle\Psi_x| + U_{2x} \otimes |\Psi_x\rangle\langle\Phi_x| - U_{2x-1} \otimes |\Phi_x\rangle\langle\Psi_{x-1}| - U_{2x-1}^\dagger \otimes |\Psi_{x-1}\rangle\langle\Phi_x| \right]. \quad (34)$$

We can use $\mathcal{G} + 1$ qubits to encode the clock state analogously to (9). That is,

$$\begin{aligned} |\Phi_0\rangle &= |10000 \dots 0\rangle \\ |\Psi_0\rangle &= |01000 \dots 0\rangle \\ |\Phi_1\rangle &= |00100 \dots 0\rangle \\ |\Psi_1\rangle &= |00010 \dots 0\rangle \\ &\vdots \end{aligned}$$

Then, the clock transitions are again 2-local operators. Thus, if the original circuit is built from gates that each act on at most two qubits, H_{lin} is a 4-local Hamiltonian.

Using this encoding, H_{lin} is an operator acting on a $2^{n+\mathcal{G}}$ -dimensional Hilbert space. However, as with the Feynman-Kitaev Hamiltonian, the subspace \mathcal{C} is preserved by this Hamiltonian, and within \mathcal{C} , H_{lin} acts as illustrated in (32). Discretizing $\frac{d}{dx}$ by a forward difference on a lattice of spacing a corresponds to

$$\frac{1}{a} \begin{bmatrix} -1 & 1 & & & \\ & -1 & 1 & & \\ & & -1 & 1 & \\ & & & -1 & 1 \\ & & & & \ddots \end{bmatrix}.$$

So, if we think of the clock register as discretizing the unit interval, the corresponding wave propagation speed v is $1/\mathcal{G}$, up to higher order corrections in $1/\mathcal{G}$.

With H_{lin} we can achieve arbitrary length computations with constant-length paths through Hilbert space just as in the Feynman-Kitaev example, but now the evolution of the wavepacket is simpler and cleaner to analyze. We pad the circuit with G initial identity gates and G final identity gates. Then, we prepare the initial wavepacket state

$$|0\rangle^{\otimes n} \otimes \frac{1}{\sqrt{2}} \sum_{x=0}^{G/2} \left[w\left(\frac{x}{G}\right) |\Phi_x\rangle + iw\left(\frac{x}{G}\right) |\Psi_x\rangle \right] \quad (35)$$

where w is a normalized wavepacket and, without loss of generality, we take $|0\rangle^{\otimes n}$ to be the initial state of the computational register.

We can now compute \mathcal{L} for this construction in the limit $G \rightarrow \infty$. In this continuum limit we have

$$|\psi(t)\rangle = \frac{1}{\sqrt{2}} \int_0^1 dx w(x - vt) |x_\Phi\rangle + \frac{i}{\sqrt{2}} \int_0^1 dx w(x - vt) |x_\Psi\rangle. \quad (36)$$

where $|x_\Phi\rangle$ and $|x_\Psi\rangle$ are the continuum analogues of $|h'_{2x}\rangle|\Phi_x\rangle$ and $|h'_{2x+1}\rangle|\Psi_x\rangle$, respectively. Here we have chosen the normalization so that

$$\int_0^1 w(x)^2 = 1. \quad (37)$$

Thus, one finds

$$\left\| \frac{d}{dt} |\psi\rangle \right\| = \sqrt{2}v \int_0^1 dx |w'(x)|^2, \quad (38)$$

where w' denotes the derivative of w . In our construction padded with identity gates the propagation velocity is $v \simeq \frac{1}{3G}$ and the total duration is $T \simeq 3G$. So

$$\mathcal{L} = \int_0^{3G} dt \frac{\sqrt{2}}{3G} \int_0^1 dx |w'(x)|^2 = \sqrt{2} \int_0^1 dx |w'(x)|^2. \quad (39)$$

As a concrete example of a smooth normalized wavepacket $w(x)$ with support only on $0 \leq x \leq \frac{1}{3}$, one could choose

$$w(x) = \begin{cases} \sqrt{2}(1 - \cos(6\pi x)) & 0 \leq x \leq \frac{1}{3} \\ 0 & \text{otherwise.} \end{cases} \quad (40)$$

In this case, one finds by straightforward calculation that

$$\mathcal{L} = 12\sqrt{2}\pi^2. \quad (41)$$

Next we can compute ΔE . By (36)

$$\langle \psi | H | \psi \rangle = i \langle \psi | \frac{d}{dt} | \psi \rangle \quad (42)$$

$$= iv \int_0^1 dx w(x-vt)w'(x-vt) \quad (43)$$

$$= 0. \quad (44)$$

Thus,

$$\Delta E = \sqrt{\langle \psi | H^2 | \psi \rangle} \quad (45)$$

$$= \|H|\psi\rangle\| \quad (46)$$

$$= \left\| \frac{d}{dt} |\psi\rangle \right\| \quad (47)$$

$$= \frac{\mathcal{L}}{3G} \quad (48)$$

$$= \frac{4\sqrt{2}\pi^2}{G} \quad (49)$$

by (41).

VI. SPATIALLY LOCAL CONSTRUCTION

The constructions of §II and §V are perhaps slightly unphysical in that, although the Hamiltonians are local in the sense of involving only 4-qubit interactions, they are not spatially local. In this section we describe how to use an idea from [36] to modify the constructions so that they become spatially local in two dimensions. A detailed illustrative example is given in appendix A.

We arrange the gates of the original circuit into “layers” such that the gates within each layer act on distinct subsets of the qubits. The number of layers D is called the circuit depth. Correspondingly, we have a sequence $|l_0\rangle, |l_1\rangle, \dots, |l_D\rangle$ where $|l_j\rangle$ is the state of the qubits after j of the layers have been applied.

Next, we construct a new equivalent quantum circuit on $n \times D$ qubits as follows. We lay out an $n \times D$ square lattice of qubits in two dimensions. Each column (of n qubits) is to be initialized to $|l_0\rangle$ which, without loss of generality, is the all zeros state. The first stage in the new circuit is to obtain the state $|l_1\rangle$ in the second column. This is done by applying the gates in the first layer of the original circuit in order from top to bottom on the qubits of the first column, and following each gate with a SWAP operation that brings the qubits it acted on into the second column.

Any qubit that was not acted on in the first layer of the circuit can be thought of as acted on by an identity gate. That is, it is swapped into the second column. Next, the same thing is done to implement the second layer of the circuit and swap the qubits into the third column, except the gates are implemented in order from bottom to top. Gates are implemented on successive layers alternating between bottom-to-top order, and top-to-bottom order, until all D layers are complete. After the last layer there is no need for swap operations. This procedure ensures that, at the end of the computation, the last column contains state $|l_D\rangle$, which is the output from the original circuit. (See figures 2-5 in appendix A.)

The modified circuit can then be simulated using a clock Hamiltonian with bounded-range interactions in two-dimensions. Recall the clock encoding (9). Here, we use the same encoding, except the number of clock qubits will equal the total number of gates of the modified circuit including the swap gates. These clock qubits can then be “snaked” between the layers so that each clock bit geometrically neighbors the computational qubits acted on by the gate that the clock bit corresponds to, and the hopping terms that move the single 1 among the clock qubits are also spatially local. The construction is thus local in two spatial dimensions and involves 4-body interactions (and fewer). The wavepacket propagates down the snaking path of the clock qubits, which is of length at most nD , which is upper bounded by nG and often much smaller. (See figure

6 in appendix A.)

VII. COMPRESSED CLOCK

In the above clock-based constructions, the total number of qubits used is on the order of $n + G$, whereas the original quantum circuit acted only on n qubits. Such an increase in qubit requirement is undesirable, especially in the context of limited qubit density and signal propagation speed, as discussed in §VIII. In this section we show how to modify the encoding of the clock so that only $G^{1/r}$ clock qubits are used, at the cost of requiring $(2r + 2)$ -local interactions. Achieving such clock compression together with spatial locality remains an open problem.

The most efficient encoding for the clock, in terms of qubit count, would be to store the number $x \in \{0, 1, \dots, \mathcal{G}\}$ as a binary number in $b = \lceil \log_2(\mathcal{G} + 1) \rceil$ qubits. However, a Hamiltonian based on this clock would involve $(b + 2)$ -qubit interactions, and is therefore, in most circumstances, not physically realistic. To achieve computational universality using a k -local Hamiltonian of the sort described in section §II one requires that the clock bit string encoding x can be incremented to $x + 1$ by flipping at most $k - 2$ bits, and deciding whether a given clock bitstring encodes a given number $x \in \{0, 1, \dots, \mathcal{G}\}$ requires examining at most $k - 2$ qubits. The “pulse” encoding (9) is optimal in this respect: incrementing the clock requires flipping two qubits, and deciding whether a string encodes x requires only examining the x^{th} qubit. However, as noted above, the pulse encoding is highly suboptimal in terms of number of qubits required, namely \mathcal{G} .

We can interpolate between these extremes as follows. We take all $\binom{b}{r}$ strings of b bits in which the number of ones (Hamming weight) is r . We number these from 0 to $\binom{b}{r} - 1$. In such an encoding, the operators $|x + 1\rangle\langle x|$ and $|x\rangle\langle x + 1|$ appearing in the clock Hamiltonian act nontrivially on at most $2r$ bits, and therefore the clock Hamiltonian is $(2r + 2)$ -local. The number of bits needed for the clock register is the minimum b such that $\binom{b}{r} \geq \mathcal{G} + 1$, which scales as $b = O(\mathcal{G}^{1/r})$. We leave open the problem of finding the optimal tradeoff between locality and number of clock qubits.

To ensure that, given a quantum circuit, computing a description of the corresponding clock Hamiltonian is efficient, one must choose the numbering of the $\binom{b}{r}$ bit strings so that encoding of numbers into bitstrings and decoding of bitstrings into numbers are both efficient. That is, the map should admit polynomial-time classical encoding and decoding. The lexicographical numbering of the Hamming weight r bitstrings admits simple methods for polynomial-time encoding and decoding for any r as described for example in [37].

VIII. CONCLUSION AND OPEN PROBLEMS

The examples given above show that energy limitations alone do not impose an upper bound on computational clock speed, even if we restrict our attention to Hamiltonians that are spatially local and involve only 4-body interactions. Should one conclude that our universe admits unlimited computation speed, in principle? This seems unlikely. By introducing more detailed assumptions about physics, beyond just a limit on energy, one may recover limits on computational speed.

Two goals one could consider are achieving constant clock speed with asymptotically shrinking energy (as is done in §V) or achieving asymptotically growing clock speed with constant energy (ΔE and/or $\langle E \rangle$). This latter goal could be achieved by the constructions in this manuscript but with the Hamiltonian rescaled by an appropriate factor, namely a factor of order G^2 for H^{FK} , and a factor of order G for H_{lin} . However, these rescaled Hamiltonians would then be sums of 4-body interactions each with norm scaling as G or G^2 . Whether a construction is possible achieving unbounded clock speed while keeping $\langle E \rangle$, ΔE , and the strength of the local interactions bounded remains an open question. In other words, limiting the interaction strength to $O(1)$ might be a sufficient additional physical assumption to recover a bound on computational speed.

If we keep the normalization used throughout this paper, where the interaction strengths are $O(1)$, we no longer obtain unbounded clock speed, but we still obtain unbounded ratio of clock speed to the energy scales $\langle E \rangle$ and ΔE . Is this fully physically realistic? It seems the most fundamental aspect of this question is whether universal quantum computation by states of vanishing energy and energy uncertainty evolving according to time-independent Hamiltonians can be made fault tolerant against the influences of imperfect implementation and environmental noise. To our knowledge this is an open question.

One might also recover computational speed limits by bringing in spacetime considerations. In particular, assume a maximum speed v for signal propagation, and a maximum density ρ at which qubits can be packed. Then, for a computer of n qubits in three-dimensional space, the distance between nearest-neighbor qubits is $\sim \rho^{-1/3}$ and the average distance between qubits is $(n/\rho)^{1/3}$. Consequently, a two qubit gate must take time at least

$$t_{\text{min}} \sim \rho^{-1/3} v^{-1} \tag{50}$$

to act on neighboring qubits and on the order of $\sim n^{1/3} \rho^{-1/3} v^{-1}$ to act on generic pairs of qubits.

The relevant limits to qubit density may appear at different scales depending on context. In a present-day practical context the limit on qubit density may be set by the atomic scale. In consid-

ering the computational complexity implications of relativistic quantum field theory, one can use consider a length scale $1/E$ as a cutoff[50], where E is the available energy. More concretely, the quantum simulation algorithms of [38–40] demonstrate (in simple examples of quantum field theories) that scattering processes involving particles with total energy E can be faithfully simulated by discretizing space onto a lattice of spacing $O(1/E)$ and associating a register of logarithmically many qubits with each lattice site. This suggests that, effectively, the density of qubits accessible by experiments at energy scale E is limited to $\tilde{O}(E^d)$ in d spatial dimensions. Perhaps the most fundamental limits to qubit density come from quantum gravity considerations such as the Bekenstein bound [41–43].

The practical limit on speed of information transmission often coincides, at least approximately, with the fundamental limit set by the speed of light. At the fundamental scale, it is thought that the maximum number of bits of entropy supportable within a region of spacetime faces a limit proportional to the surface area of the region, with the constant of proportionality approximately 10^{69} bits per square meter [41–43]. It may be tempting therefore, for a given n to compute from the Bekenstein bound a corresponding ρ and consequently a maximum clock speed via (50) with the speed of light taking the place of v . However, it is not clear that this is a valid argument because in the regime where quantum gravity effects are significant the question of spatial locality may become subtle.

Acknowledgments: I thank Scott Aaronson, Ning Bao, Michael Jarret, Manny Knill, Carl Miller, and Aaron Ostrander for useful discussions. I thank Ilya Bogdanov for suggesting the key idea behind §VII (via mathoverflow). This paper is a contribution of NIST, an agency of the US government, and is not subject to US copyright.

Appendix A: Example of Spatially Local Construction

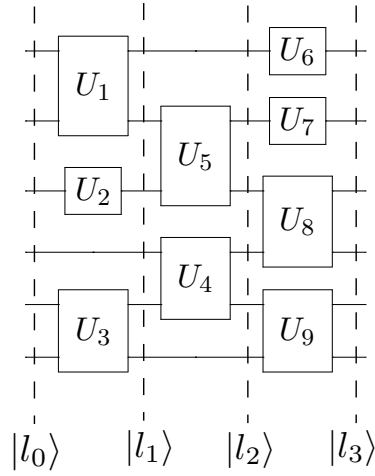


FIG. 2: Here is an example of a quantum circuit, referred to henceforth as “the original circuit”. $|l_j\rangle$ labels the state after the first j layers of the circuit have been applied. We have assumed that all gates in the original circuit act on nearest-neighbors in one dimension. This can always be achieved through the use of SWAP gates.




FIG. 3: The original circuit has depth 3 and 6 qubits. Correspondingly we have a grid of 3 columns with 6 qubits each. The above sequence of steps implements the first layer of gates. Initially each column is in the state $|l_0\rangle$, the input to the original circuit, which can be taken to be the all zeros state. After these steps are complete, the middle column contains state $|l_1\rangle$. The paired \times symbols indicate swap gates.




FIG. 4: These steps implement the second layer of gates. After these steps are complete, the right column contains state $|l_2\rangle$.




FIG. 5: These steps implement the third layer of gates. After these steps are complete, the right column contains state $|l_3\rangle$, which is the output of the original circuit.




FIG. 6: The numbered circles represent clock qubits. The valid states of the clock are Hamming weight one strings. If the x^{th} clock qubit is 1 then this corresponds to the x^{th} step in figures 3-5. The clock qubits are “snaked” among the computational qubits to ensure that the implementation of the x^{th} step triggered by the x^{th} clock qubit is spatially local, as is the hopping of the 1 from the x^{th} clock qubit to the $(x + 1)^{\text{th}}$ clock qubit.

-
- [1] L. Mandelstam and I. Tamm. The uncertainty relation between energy and time in nonrelativistic quantum mechanics. *Journal of Physics (USSR)*, 9:249–254, 1945.
- [2] G. N. Fleming. A unitarity bound on the evolution of nonstationary states. *Nuovo Cimento A*, 16(2):232–240, 1973.
- [3] J. Anandan and Y. Aharonov. Geometry of quantum evolution. *Physical Review Letters*, 65:1697–1700, 1990.
- [4] Lev Vaidman. Minimum time for the evolution to an orthogonal quantum state. *American Journal of Physics*, 60:182–183, 1992.
- [5] J. Uffink. The rate of evolution of a quantum state. *American Journal of Physics*, 61:935–936, 1993.
- [6] Asher Peres. *Quantum Theory: Concepts and Methods*. Kluwer, 1995.
- [7] Norman Margolus and Lev B. Levitin. The maximum speed of dynamical evolution. *Physica D: Nonlinear Phenomena*, 120:188–195, 1998. arXiv:quant-ph/9710043.
- [8] P. Pfeifer. How fast can a quantum state change with time? *Physical Review Letters*, 70:3365–3368, 1993.
- [9] Samuel Braunstein, Carl Caves, and Gerard Milburn. Generalized uncertainty relations: theory, examples, and Lorentz invariance. *Annals of Physics*, 247:135–173, 1996.
- [10] Vittorio Giovannetti, Seth Lloyd, and Lorenzo Maccone. Quantum limits to dynamical evolution. *Physical Review A*, 67:052109, 2003.
- [11] Lev B. Levitin and Tommaso Toffoli. Fundamental limit on the rate of quantum dynamics: the unified bound is tight. *Physical Review Letters*, 103:160502, 2009.
- [12] Michael R. Frey. Quantum speed limits—primer, perspectives, and potential future directions. *Quantum Information Processing*, 15(10):3919–3950, 2016.
- [13] Seth Lloyd. Ultimate physical limits to computation. *Nature*, 406:1047–1054, 2000. arXiv:quant-ph/9908043.
- [14] Igor L. Markov. Limits on fundamental limits to computation. *Nature*, 512:147–154, 2014.
- [15] Seth Lloyd and Jack Ng. Black hole computers. *Scientific American*, 291(5):52–61, November 2004.
- [16] Seth Lloyd. Computational capacity of the universe. *Physical Review Letters*, 88:237901, 2002.
- [17] Michael P. Frank. The physical limits of computing. *Computing in Science and Engineering*, 4(3):16–26, 2002.
- [18] Norman Margolus. Looking at nature as a computer. *International Journal of Theoretical Physics*, 42(2):309–327, 2003.
- [19] Adam R. Brown, Daniel A. Roberts, Leonard Susskind, Brian Swingle, and Ying Zhao. Holographic complexity equals bulk action? *Physical Review Letters*, 116:191301, 2016. arXiv:1509.07876.
- [20] Adam R. Brown, Daniel A. Roberts, Leonard Susskind, Brian Swingle, and Ying Zhao. Complexity, action, and black holes. *Physical Review D*, 93:086006, 2016. arXiv:1512.04993.

- [21] Stephen D. H. Hsu. Physical limits on information processing. *Physics Letters B*, 641(1):99–100, 2006. arXiv:hep-th/0607082.
- [22] Stephen D. H. Hsu. Information, information processing, and gravity. *International Journal of Modern Physics A*, 22(17&17):2895–2907, 2007. arXiv:0704.1154.
- [23] Rodolfo Gambini, Rafael A. Porto, and Jorge Pullin. Fundamental gravitational limitations to quantum computing. In P. Fiziev and M. Todorov, editors, *Gravity astrophysics and strings at the black sea*, 2006. arXiv:quant-ph/0507262.
- [24] Michael P. Frank. On the interpretation of energy as the rate of quantum computation. *Quantum Information Processing*, 4(4):283–334, 2005.
- [25] Y. Jack Ng. Spacetime foam: from entropy and holography to infinite statistics and nonlocality. *Entropy*, 10(4):441–461, 2008. arXiv:0801.2962.
- [26] Hans J. Bremermann. Optimization through evolution and recombination. In M. C. Yovits, G. T. Jacobi, and G. D. Goldstein, editors, *Self-Organizing Systems, 1962*. Spartan Books, 1962.
- [27] Hans J. Bremermann. Minimum energy requirements of information transfer and computing. *International Journal of Theoretical Physics*, 21(3):203–217, 1982.
- [28] J. D. Bekenstein. Energy cost of information transfer. *Physical Review Letters*, 46:623, 1981.
- [29] L. B. Levitin. Physical limitations of rate, depth, and minimum energy in information processing. *International Journal of Theoretical Physics*, 21:299–309, 1982.
- [30] Yosi Atia and Dorit Aharonov. Fast-forwarding of Hamiltonians and exponentially precise measurements. *arXiv:1610.09619*, 2016.
- [31] Charles H. Bennett and Rolf Landauer. The fundamental physical limits of computation. *Scientific American*, 253(1):48–56, 1985.
- [32] Richard P. Feynman. Quantum mechanical computers. *Optics News*, 11(2):11–20, 1985.
- [33] Dorit Aharonov, Wim van Dam, Julia Kempe, Zeph Landau, Seth Lloyd, and Oded Regev. Adiabatic quantum computation is equivalent to standard quantum computation. *SIAM Journal of Computing*, 37(1):166–194, 2007. arXiv:quant-ph/0405098.
- [34] A. Yu. Kitaev, A. H. Shen, and M. N. Vyalyi. *Classical and Quantum Computation*, volume 47 of *Graduate Studies in Mathematics*. American Mathematical Society, 2002.
- [35] Rolando Somma and Sergio Boixo. Spectral gap amplification. *SIAM Journal on Computing*, 42:593–610, 2013. arXiv:1110.2494.
- [36] Daniel Nagaj. Universal two-body-Hamiltonian quantum computing. *Physical Review A*, 85:032330, 2012. arXiv:1002.0420.
- [37] Thomas M. Cover. Enumerative source encoding. *IEEE Transactions on Information Theory*, IT-19(1):73–77, 1971.
- [38] Stephen P. Jordan, Keith S. M. Lee, and John Preskill. Quantum algorithms for quantum field theories. *Science*, 336(6085):1130–1133, 2012. arXiv:1111.3633.
- [39] Stephen P. Jordan, Keith S. M. Lee, and John Preskill. Quantum computation of scattering in

- scalar quantum field theories. *Quantum Information and Computation*, 14(11/12):1014–1080, 2014. arXiv:1112.4833.
- [40] Stephen P. Jordan, Keith S. M. Lee, and John Preskill. Quantum algorithms for fermionic quantum field theories. *arXiv:1404.7115*, 2014.
- [41] Jacob D. Bekenstein. Black holes and entropy. *Physical Review D*, 7(8):2333–2346, 1973.
- [42] Jacob D. Bekenstein. Universal upper bound on the entropy-to-energy ratio for bounded systems. *Physical Review D*, 23:287, 1981.
- [43] Raphael Bousso. The holographic principle. *Reviews of Modern Physics*, 74:825–874, 2002. arXiv:hep-th/0203101.
- [44] As a rough argument, one could model a typical state at an intermediate step of a quantum computation by a Haar random state on n qubits. In this case, the root-mean-square inner product between the states before and after a one-qubit gate has been performed is $1/\sqrt{2^n}$.
- [45] For notational simplicity, we use units where $\hbar = 1$ here and throughout the remainder of this paper.
- [46] Quantum circuits acting on the all zeros string are computationally universal because any classical input string can be hard-coded via initial NOT gates.
- [47] An easy way to obtain this is by going to Fourier space, where the Hamiltonian is diagonal.
- [48] The 2 comes from squaring the amplitudes. The 5 comes from $\mathcal{G} = 5G$.
- [49] These speeds and energy scales are determined by our choice of normalization of the Hamiltonian, but this choice is not arbitrary. Physically, one expects the individual 4-local terms to have $O(1)$ norm.
- [50] $\hbar c/E$ in units where explicit factors of c and \hbar are kept