Fundamental question of computational complexity:
Which problems can be solved with polynomial resources.

Quantum Church Turing Thesis says: BQP
(the set of problems solvable by quantum circuits of polynomially many gates)

If so, quantum circuits can simulate all other physical systems using gates and qubits scaling polynomially in Energy, time, precision, etc.

Is this really true?

In lecture 1, we solved one example:

$$H_{NR} = -\frac{1}{2m}\frac{d^2}{dx^2} + V(x)$$

$e^{-iH_R t}$ can be approximated within ϵ using $\text{poly}(t, 1/\epsilon)$ quantum gates.

(Method was phase kickback + Trotter-Suzuki Formulas.)

That's a good start, but what about:

- State preparation
- measurement
- more/other degrees of freedom (multiple dimensions, multiple particles, spin, etc.)
- Fermionic statistics
- QFT

Hilroy

Today we'll address these issues one by one, until we reach the research frontier and conclude with some open problems. Due to limited time I'll probably skip:

- fermionic statistics
- more/different degrees of freedom

These can be handled by some extensions of the methods described in lecture 1. Those interested are directed to chapter 4 of Nielsen & Chuang.

Measurement [relegate to end of lecture, skip to pg. 5]

- Recall that in lecture 1, we performed simulation in the X-basis:

$$\sum_{x \in \{0,1\}^n} \psi(x) |x\rangle$$

↑ represents particle position $x \in \mathbb{R}$ to n bits of precision

- If we measure in the "computational basis", i.e. measure each qubit in $\{|0\rangle, |1\rangle\}$ basis, then we obtain a bitstring representing $x \in \mathbb{R}$ to some precision. We thus approximately sample from

$$P(x) = |\psi(x)|^2$$

- Sometimes this may be what we want to know, but often not.

Hilroy

- What if we want to measure energy?
- we can do this using Kitaev's method of phase estimation:

  - suppose $|\psi\rangle$ is an eigenstate of H:

    $$H|\psi\rangle = E|\psi\rangle$$

  - we want to learn E

- We can do it like this:

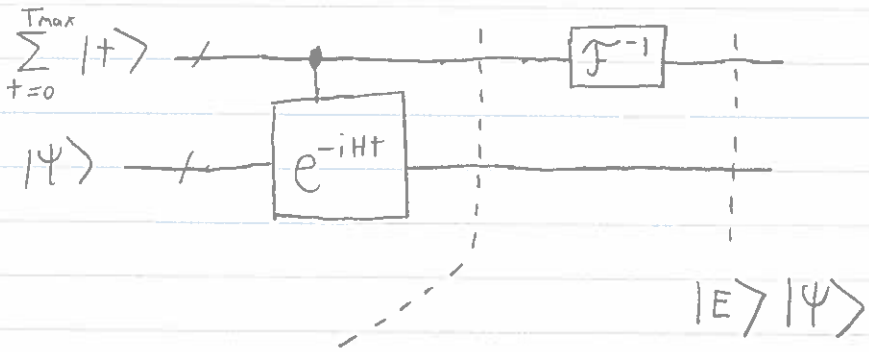

$$\sum_t |t\rangle e^{-iEt} |\psi\rangle$$

$$= \left(\sum_t e^{-iEt}|t\rangle\right) \otimes |\psi\rangle$$

- This lets us measure in eigenbasis of H, with energy resolution $1/T_{max}$

- Some trick works for other observables:
  $\vec{P}$, $\vec{J}$, $\sigma_z^{(i)}\sigma_z^{(j)}$, ...

State Preparation

- state preparation is generically hard. It is attacked large collection of specialized tools.

Hilroy

I'll introduce one of the more general and more interesting methods: adiabatic state preparation. To learn more about state preparation attend Adam Bouland's focus lecture and try problem 4 from my problem set.

Adiabatic Quantum Computing & State Preparation

- Adiabatic Quantum Computing (AQC) is Hamiltonian-based model of computation.

- Let $H(s)$ be a parameterized family of Hamiltonians, e.g.

$$H(s) = (1-s)\,H_{init} + s\,H_{final} \qquad s \in [0,1]$$

- Goal: prepare the ground state of H_{final}.

- Computation:
  1) Prepare the ground state of $H(0)$
  2) evolve according to $H(t/T)$ from time 0 to T

- By the adiabatic theorem, choosing T sufficiently large ensures that the system quasi-statically tracks the instantaneous ground state

$$|\Psi(t)\rangle \simeq \underbrace{e^{i\phi}}_{\text{some dynamical phase}} \underbrace{|\Psi_0(s(t))\rangle}_{\text{ground state of } H(s(t))}$$

$|\Psi(t)\rangle$: solution to Schrödinger's equation!

$$\frac{d|\Psi\rangle}{dt} = -i\,H(s(t))\,|\Psi\rangle$$

Hilroy

- Clearly, it's crucial that $H(s)$ have an easily prepared ground state. Otherwise we can't do step 1! Good commonly-used examples of easy-to-prepare states are tensor-product states such as:

$$|00\ldots0\rangle = |0\rangle^{\otimes n}$$

and

$$\frac{1}{\sqrt{2^n}} \sum_{x \in \{0,1\}^n} |x\rangle = \left(\frac{|0\rangle + |1\rangle}{\sqrt{2}} \right)^{\otimes n}$$

- To judge complexity, the key question is: how large does T need to be? Quantitative versions of the adiabatic theorem give sufficient conditions on T. For example:

Adiabatic Theorem: [Elgart & Hagedorn 2012]

To track the ground state it suffices to choose:

$$T = \tilde{O}\left(\frac{\left\| \frac{dH}{ds} \right\|}{\gamma^2} \right)$$

where

$$\gamma = \min_{0 \le s \le 1} E_1(s) - E_0(s) \quad \text{"minimal gap"}$$

- $\tilde{O}$ is a computer science notation which indicates the asymptotic scaling neglecting both constant subleading terms, constant factors, and log factors.

- We restrict to physically realistic Hamiltonians, i.e. with k-body interactions up to some constant k.
- For any $k \geq 2$ this model is polynomially equivalent to the quantum circuit model.
- How is this shown?

Shown by applying the quantum simulation techniques from lecture 1 (Suzuki-Trotter).

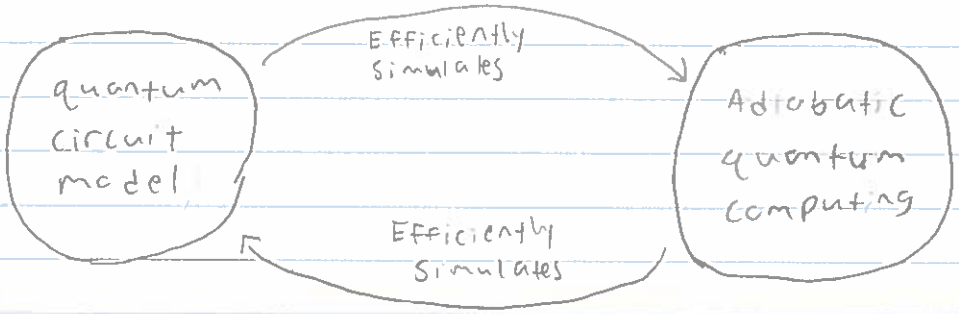

Shown in

[Aharonov, Van Dam, Kempe, Landau, Lloyd, Regev, 2007] (preprint 2004) using "Feynman-Kitaev" clock construction. This is a clever and interesting construction which will be discussed later by Adam Bouland and/or Dorit Aharonov.

Let's now go through some consequences of this equivalence. Some of these will be used in quantum algorithms for simulating QFT. Furthermore, they clarify the importance of polynomial equivalence, which is a central theme of my lectures.

Hilroy

Some implications of AQC ↔ Quantum Circuits

1) Interoperability

We have some powerful tools for quantum simulation.

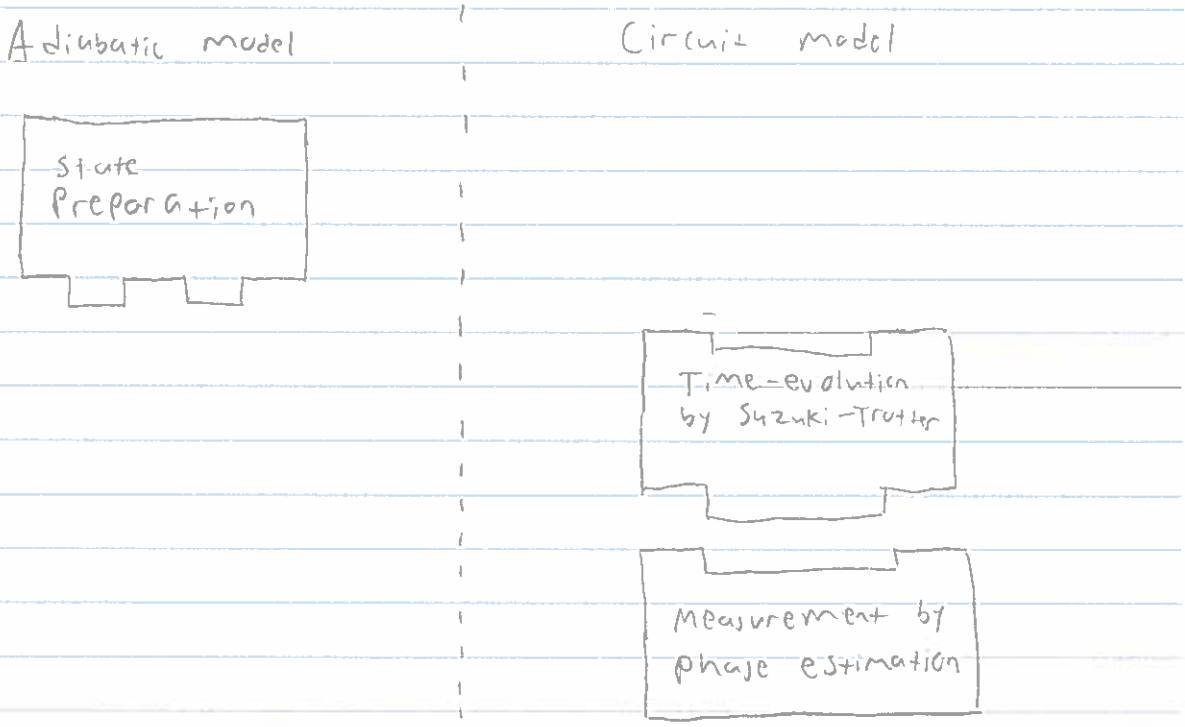

These are like legos that don't fit. How do we join them together to make a full simulation algorithm?

State prep → time evolution → measurement ?

Answer: We don't. Instead just convert the adiabatic state preparation method into a circuit-model state-preparation method by simulating it.

Hilroy

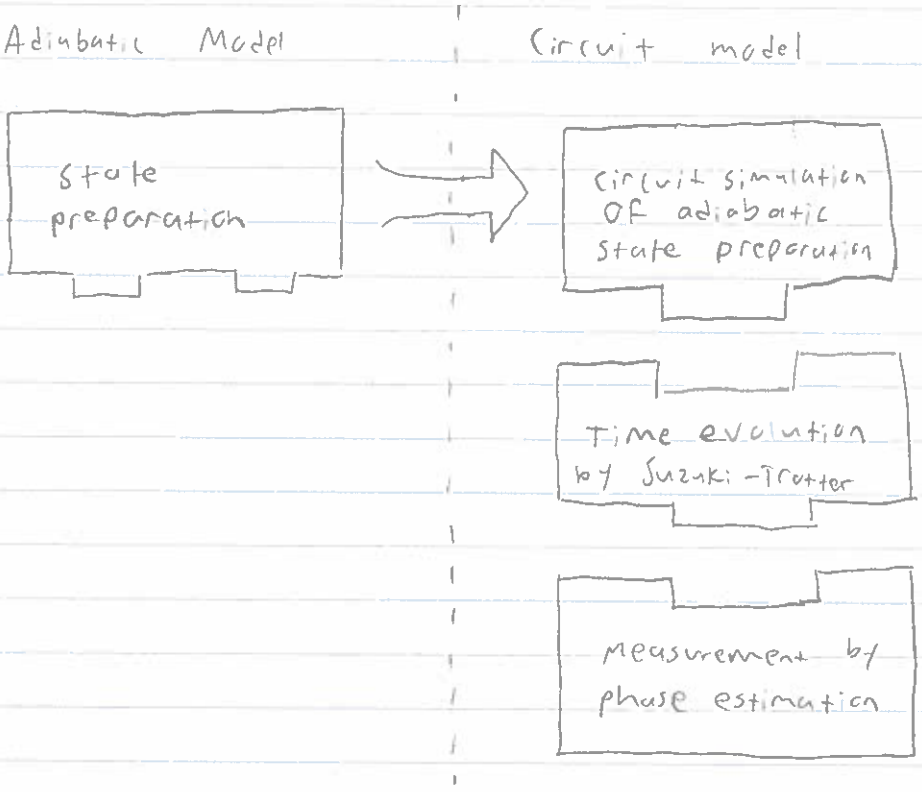

2) Fault Tolerance

- No threshold theorem is known for adiabatic quantum computation.
- Nevertheless, adiabatic algorithms providing exponential speedups over classical computing represent genuine quantum advantage. They can be simulated with polynomial overhead by fault tolerant quantum circuits.
- Some quantum algorithms might be easier to discover in one model than another.

3) Different models provide different architectures for building quantum computers in the lab, each with their own advantages and challenges.

Hilroy

Let's now apply these 3 tools:

1) state prep (simulated adiabatic)
2) time evolution (Suzuki-Trotter)
3) measurement (phase estimation)

to simulate a quantum field theory.
We'll follow [Jordan, Lee, Preskill, 2012]

Simulating QFT on a Quantum Computer

Motivation

efficient quantum algorithms for simulation:

- chemistry ✓
- condensed matter ✓
- QFT partial progress
- quantum gravity ???

→ The remaining challenges to the quantum Church-Turing thesis come from high energy physics.

Q. Why might simulating QFT be different from simulating low energy physics?

- field has infinitely many degrees of freedom
- particle # not conserved
- relativistic

→ So, efficient quantum algorithms to simulate QFT do not follow directly from quantum simulation algorithms scaling polynomially with # degrees-of-freedom or # particles.

Hilroy

Another motivation is that in certain regimes we have no polynomial-time classical simulation algorithms:

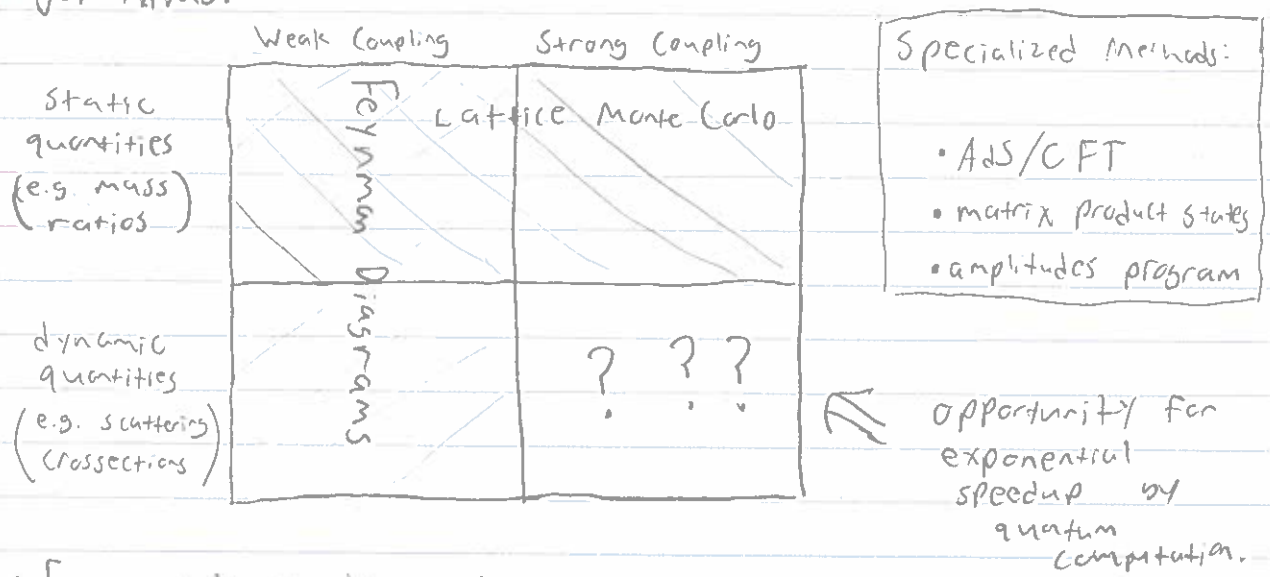

- For weak coupling we can use perturbative expansions such as Feynman diagrams

- For static quantities like binding energies we can do imaginary-time lattice gauge theory by Monte Carlo.

- To calculate dynamical quantities at strong coupling we don't have any efficient classical algorithm except in special cases such as supersymmetric models.

→ Furthermore, general purpose polynomial-time classical algorithms are probably impossible. If one existed it could simulate quantum computers implying $BQP = P$.

→ Even QFTs much simpler than the Standard Model support universal quantum computation. [JLP, to appear] So, even these should be impossible to efficiently simulate classically.

Hilroy

In the rest of this lecture we'll treat a simple QFT in detail:

Quantum Algorithm to Simulate ϕ^4-theory

- What problem do we solve?

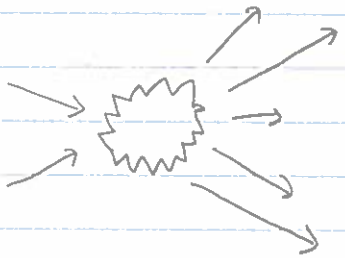

Input: A list of momenta of incoming particles.

Output: A list of momenta of outgoing particles.

ϕ^4-theory:

A high energy theorist would describe it by the Lagrangian-density

$$\mathcal{L} = \frac{1}{2}\partial^\mu \phi \partial_\mu - \frac{1}{2}m^2\phi^2 - \frac{\lambda}{4!}\phi^4$$

Spacetime dim.

$(D = d+1)$

↑ Spatial dim.

(This goes into a path integral like this: $e^{-i\int D\phi\, \mathcal{L}}$)

As usual, we can convert to an equivalent Hamiltonian formulation:

$$H = \int d^d x \left[\frac{1}{2}\Pi^2 - \frac{1}{2}(\nabla\phi)^2 + \frac{1}{2}m^2\phi^2 + \frac{\lambda}{4!}\phi^4\right]$$

Now $\Pi(x)$ and $\phi(x)$ are canonically-conjugate operators.

$$[\phi(x), \Pi(y)] = i\,\delta^{(d)}(x-y)\,\mathbb{1}$$

Hilroy

Next, we discretize space. To keep things simple, I'll restrict to one spatial dimension for the rest of this lecture.

lattice: Ω

Then:

$$H = \sum_{x \in \Omega} a \left[\frac{1}{2} \Pi(x)^2 + \frac{1}{2} \left(\frac{\phi(x+a) - \phi(x)}{a} \right)^2 + \frac{1}{2} m_0^2 \phi(x)^2 + \frac{\lambda}{4!} \phi(x)^4 \right]$$

$$[\phi(x), \Pi(y)] = \frac{i}{a} \delta_{x,y} \mathbb{1}$$

- Now we just have a conventional quantum many-body system.
- $\{ \phi(x) \mid x \in \Omega \}$ form a complete set of commuting observables, which define a basis.
- The wavefunction can be expressed as:

$$|\Psi(t)\rangle = \int d\phi_1 \int d\phi_2 \cdots \int d\phi_N \; \Psi(\phi_1, \phi_2, \ldots, \phi_N, t) |\phi_1, \phi_2, \ldots, \phi_N\rangle$$

(In the limit $a \to 0$ $(N \to \infty)$ this becomes a wavefunctional.)

If we store each of $\phi_1, \ldots, \phi_N$ with m bits of precision, then we can store an approximation to this state using Nm qubits.

Superposition over classical field configurations

↓

Superposition over bit strings

= state of qubits

Could skip or defer

- We know how to simulate e^{-iHt} by Suzuki-Trotter:

$$H = H_\phi + H_\pi$$

$$H_\pi = \sum_{x \in \Omega} a \, \frac{1}{2} \pi(x)^2$$

$$H_\phi = \sum_{x \in \Omega} a \left[\frac{1}{2} \left(\frac{\phi(x+a) - \phi(x)}{a} \right)^2 + \frac{1}{2} m_0^2 \phi(x)^2 + \frac{\lambda}{4!} \phi(x)^4 \right]$$

$$e^{-iHt} = \left(e^{-iH_\phi t/n} \, e^{-iH_\pi t/n} \right)^n + \mathcal{O}\left(\frac{1}{n}\right)$$

↑ Diagonal in ϕ-basis

↑ diagonal in Fourier basis

And we can read off the corresponding circuit for e^{-iHt}:

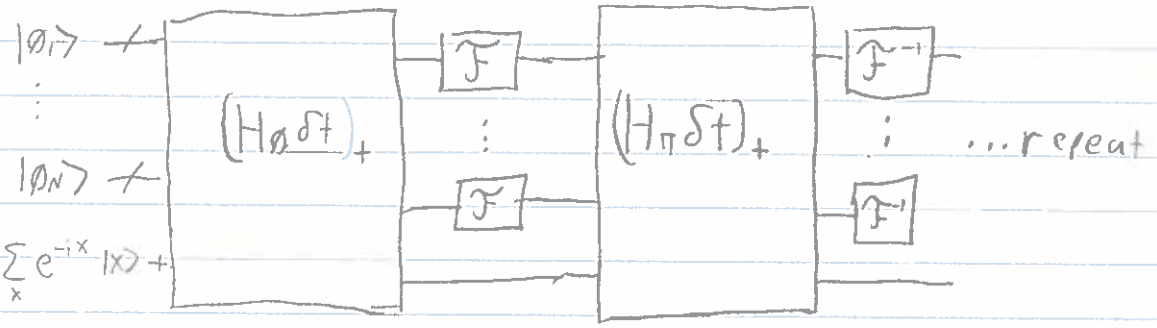

The hard parts:

1) bounding discretization errors
2) State preparation
3) measurement

1) Bounding discretization errors

One way to define a QFT is as a continuum limit of a sequence of theories on successively finer lattices.

continuum

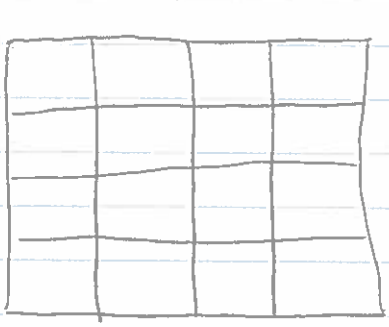

Sometimes it is easier to follow the flow through Hamiltonians down instead of up:

$$H_{a/2} = \sum_x a\left[\frac{1}{2}(\nabla\phi)^2 + \frac{1}{2}m_0\phi^2 + \frac{\lambda_0}{4!}\phi^4\right]$$

coarse grain

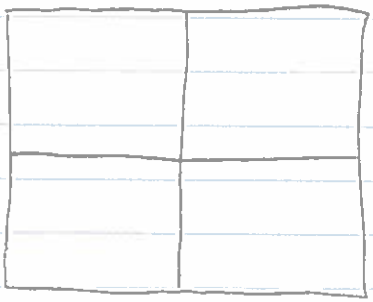

$$H_a = \sum_x a\left[\frac{1}{2}(\nabla\phi)^2 + m_0'\phi^2 + \frac{\lambda_0'}{4!}\phi^4 + \underbrace{g\phi^6 + \dots}_{\text{error terms}}\right]$$

- However, the coefficients λ_0 and m_0 in H are functions of the lattice spacing!

Hilroy

• In the case of ϕ^4-theory in $d=1,2,3$ the error terms converge to zero as a^2.

2) State preparation

• with $\lambda=0$ the model is exactly solvable
- free bosons
- change of variables ("Bogoliubov transformation") shows that the system is equivalent to a collection of decoupled harmonic oscillators
- ground state is multivariate gaussian
- we know how to make Gaussian superpositions efficiently. (Zalka's method of conditional rotations, see Problem set. Better: Kitaev-Webb.)

Idea:
1) Prepare free (i.e. $\lambda=0$) vacuum (i.e. ground state) which is a gaussian superposition.
2) Prepare wavepackets of free theory
3) Adiabatically turn on λ (simulate)
4) scatter
5) make measurements

This almost works. However, wavepackets are not energy eigenstates.

$$\sum_j \alpha_j |E_j(s=0)\rangle \longmapsto \sum_j \alpha_j e^{\int dt\, E_j(s(t))\, t} |E_j(s=1)\rangle$$

The different eigenstates accrue different dynamical phases.

In other words, the wavepacket time evolves! It propagates and broadens.

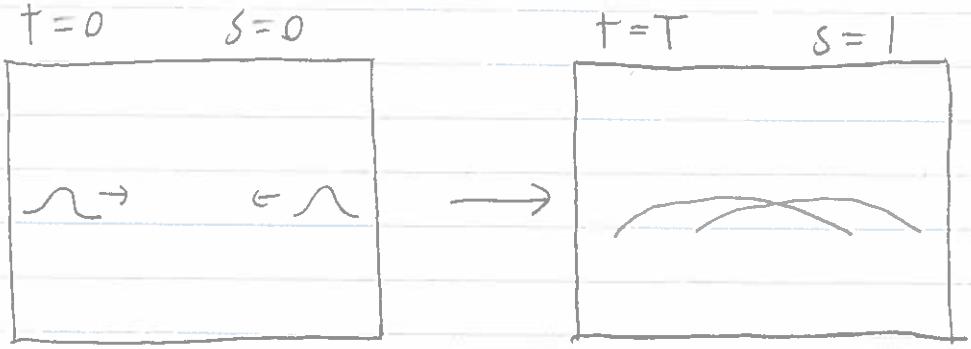

This is bad:

- particles may collide prematurely. Then we are simulating a theory with smaller λ than we wanted to.
- diffuse wavepackets means particles might miss each other. This is a boring experiment.

Solution: Intersperse backward time evolutions with time-independent Hamiltonians.

This winds back dynamical phase without undoing adiabatic basis change.

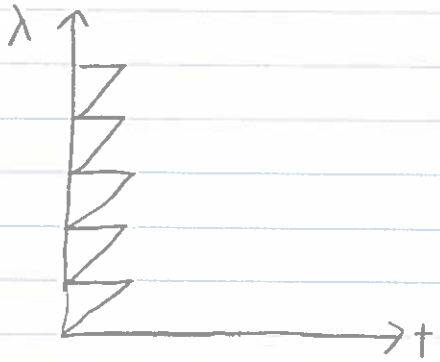

Measurement

- Phase estimation lets us measure observables.
- the question is what observables to measure

Idea Calorimetric observables

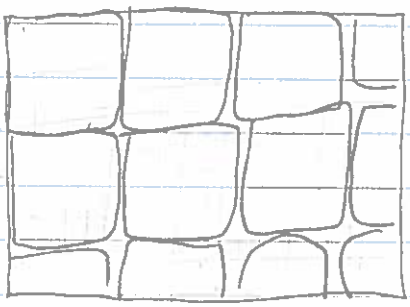

- divide space into pixels
- measure energy in each one
- if energy is higher by m, we've found a particle

- What is an observable for local energy in region $R \subseteq \Omega$?
- Naive idea:

$$H = \sum_{x \in \Omega} H(x)$$

$$H_R = \sum_{x \in R} H(x)$$

- Doesn't work!

$\langle H_R^2 \rangle - \langle H_R \rangle^2$ diverges as $a \to 0$

- This is a signal to noise problem. (Possibly related to vacuum entanglement.
- Instead use: (some smooth envelope function like a gaussian on R)

$$H_R = \sum_{x \in R} f(x) H(x)$$

- works ok but we'd like systematic understanding of this phenomenon.

Hilroy

- works fine at strong coupling
- extends to fermions

Conclusion

- It is thought that quantum computers can efficiently simulate all physical systems.
- If so, the set of feasible computations in our universe is BQP.
- This answers the most fundamental problem in CS: what can we compute and what can't we?
- However, it might not be true!
- can quantum computers simulate all of high energy physics?
  - massless QFTs?
  - Fermion doubling problem?
  - q. gravity?
- Nobody knows. Let's find out!

Hilroy