review of Hamiltonian simulation: given Hermitian matrix H... but how do we get H?

Hamiltonian is a sum of terms that act on only a few qubits (usually only a few, not all of them, leaves most unchanged)

That means we can write H, it's not too big and gnarly

exp(-iH;) is also only going to act on a few qubits - so we can turn it into executable instructions

Can use Suzuki-Trotter formula to add exponents (bc otherwise you can't) v. small error

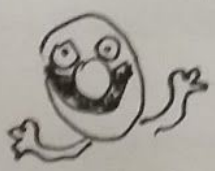

review of Grover search

Hamiltonians are related to energy

energy should be conserved ... H's eigenvectors only change by a phase and are thus stable ... both Hamiltonians and energy don't rlly change, which is the intuition for them being connected

Expect to see a state in its lowest energy state

Read off eigenvalues from using QPE ... not obvious this works

tells you what the state's energy level is

What if the matrix doesn't have eigenvectors? can write it as a sum of them - will give you a dist. over values

The closer you are to a state, the higher its probability will be

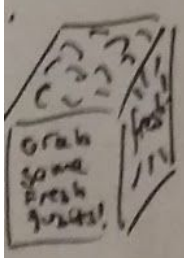

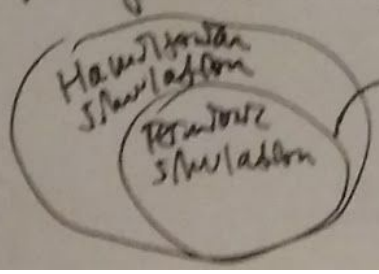

very important!

describes electrons

Physicists use a different basis + notation

Fermions sim. for electrons, Ham sim. for qubits
can't quite interchange since electrons are indistinguishable from each other

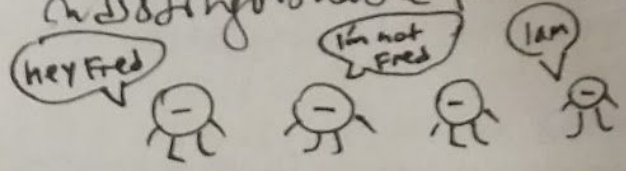

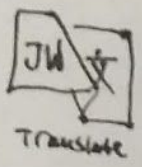

Jordan Wigner transform lets us go from electrons to qubits

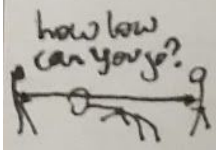

Fermionic simulation - a guess and check strategy to find the lowest energy state
We don't worry about position - leave that to scientists
All we assume about the Hamiltonian is that it's symmetric and a sum of terms - for more specialized structures you get more properties
Introducing the Essential Quantum Algorithm HHL

solves $Ax=b$ for exponentially large A and x, in $\log(N)$ instead of N
$N \times N$
big improvement

caveat: needs A to be well conditioned :(
sus: algo claims to be faster than it takes to write the input!
Ok, b is not a list, it's a quantum state. Want to get x by doing A^{-1}
need to renormalize $|x\rangle$ to return
more caveats: b must be able to be prepared efficiently :(
in every row of A there are mostly 0s :(
The algorithm is actually really complicated so we won't talk about it in detail

Step 1: assume A is Hermitian. If it's not, you can make it Hermitian in constant time

A is not necessarily unitary, but e^{-iA} is

Step 2: use Hamiltonian simulation to apply e^{iA} to $|b\rangle$

Step 3: use phase est. to find value of v (an eigenvalue of A)

Step 4: scale the state by $1/v$

do these for all eigenvectors simultaneously to get A^{-1}

* This is a loaded step. Full details are complicated

Some actions have a chance of failing but that's okay because can fail up to 1/3 of the time

HHL can't be solved classically – called BQP complete

Anything using Hamiltonian is BQP complete

In principle, HHL is useful, but unclear how it's practically useful, would need to show a classical algorithm can't come up with the same distribution of outputs