ELECTRICAL & COMPUTER ENGINEERING

From Adiabatic Quantum Computing and Quantum Adiabatic Optimization to Quantum Annealing: A Lightening Overview

Tameem Albash @ University of New Mexico

LANL Quantum Summer School

21 June 2022

CENTER FOR QUANTUM INFORMATION & CONTROL

Outline

Adiabatic quantum computing (AQC) is a paradigm of quantum computing based on continuous-time evolutions, in contrast to the discrete-time step evolutions of circuit-based algorithms.

- Adiabaticity and the adiabatic condition
- Adiabatic quantum computation and its equivalence to the circuit model
- Unstructured search in the adiabatic paradigm
- Quantum adiabatic optimization
- Beyond adiabatic: Glued-trees problem
- Open quantum system effects
- Quantum annealing, and current experimental realizations
- Final comments

My presentation will try to give enough of the background for these important topics, but given the time constraints I will not dive too deeply into any given topic. The references I provide will hopefully give you a place to explore each topic more deeply. I apologize if I have neglected to cite particular work that might be relevant. Please email me at talbash@unm.edu to let me know and for any other corrections.

Presentation material can be found here:
https://www.unm.edu/~talbash/presentations.html

Terminology

Continuous-time evolution of a quantum state is given by the Schrödinger equation

$$\frac{d}{dt}|\psi(t)\rangle = -\frac{i}{\hbar}H(t)|\psi(t)\rangle$$ ← Will be setting $\hbar = 1$ from now on

Generator of the evolution is a Hermitian operator $H(t)$ called the *Hamiltonian* of the quantum system.

- Eigenvalues are called the *energy levels* of the quantum system. Smallest eigenvalue $E_0(t)$ is called the (instantaneous) *ground state energy* of the quantum system. Higher eigenvalues are called *excited state energy levels*.
- Eigenstates are called *energy (eigen)states*. Eigenstate associated with $E_0(t)$ is called the (instantaneous) *ground state* $|E_0(t)\rangle$.
- Energy difference between ground state and the first excited state energy levels is called the ground state energy gap, $\Delta_{10}(t) = E_1(t) - E_0(t)$.

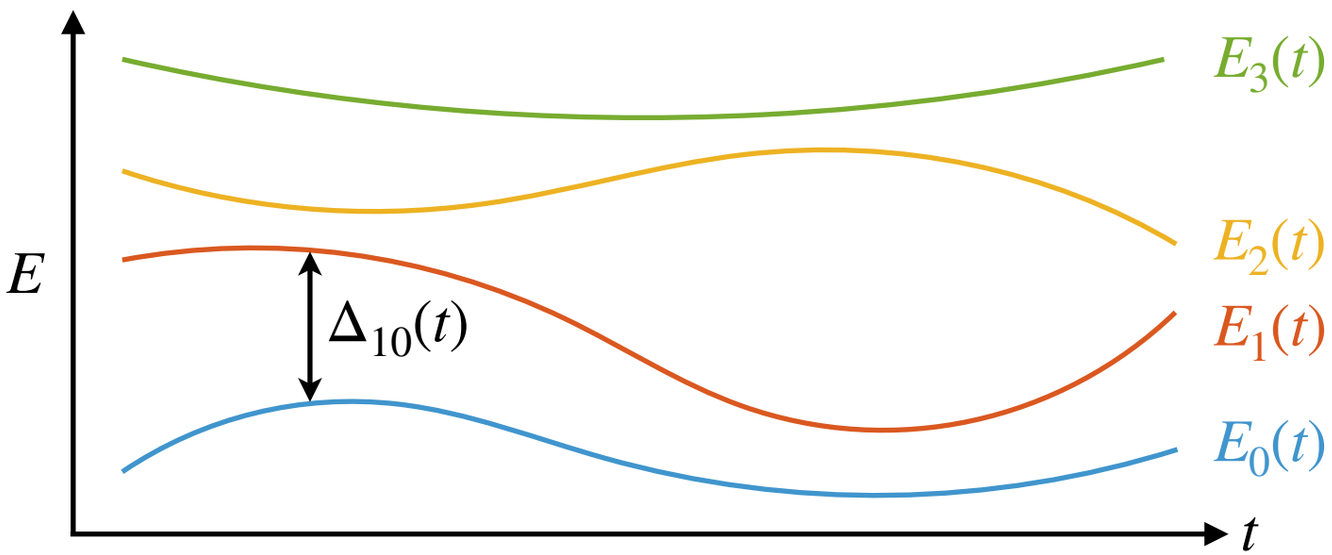

Cartoon depiction of low-lying energy levels of some Hamiltonian

Adiabatic Condition

We will assume that the Hamiltonian only depends on a dimensionless parameter $s = t/t_f \in [0,1]$, where t_f is the total evolution time. $H(t) \to H(s), E_i(t) \to E_i(s)$

Given an initial quantum state prepared in the i-th eigenstate of $H(0)$, an evolution that satisfies the *adiabatic condition* will be in the corresponding i-th eigenstate at time t_f with high probability, assuming no energy level crossings. (The adiabatic condition is a sufficient but not necessary condition.)

$$|\langle E_i(s)|\psi(t=0)\rangle|^2 = 1 \to |\langle E_i(s=1)|\psi(t)\rangle|^2 \approx 1 \ , \ t \in [0,t_f]$$

Take the case of $i = 0$; the adiabatic condition then takes the form:

Physically motivated condition [Ami2009]: $t_f \gg \max_{s\in[0,1]} \frac{|\langle E_1(s)|\frac{d}{ds}H(s)|E_0(s)\rangle|}{\Delta_{10}^2(s)}$ Minimizes the transition rate out of the ground state

Rigorous condition [Kat1950,Jan2007]: $t_f \gg \max\left\{ \max_{s\in[0,1]} \frac{\|\partial_s^2 H(s)\|}{\Delta_{10}^2(s)}, \max_{s\in[0,1]} \frac{\|\partial_s H(s)\|^2}{\Delta_{10}^3(s)}, \max_{s\in[0,1]} \frac{\|\partial_s H(s)\|}{\Delta_{10}^2(s)} \right\}$

Scaling of t_f with system size to satisfy the adiabatic condition *primarily* depends on how the minimum gap $\Delta_{\min} = \min_{s\in[0,1]} \Delta_{10}(s)$ scales with system size.

Examples of evolutions: fast

We will write our quantum state evolutions in the instantaneous energy eigenbasis

$$|\psi(t)\rangle = \sum_k c_k(t)\,|E_k(t)\rangle$$

Let's assume we are initially completely in the ground state
We will have qualitatively three kinds of behavior

(1) Diabatic transition to excited state
Evolution very fast relative to gap

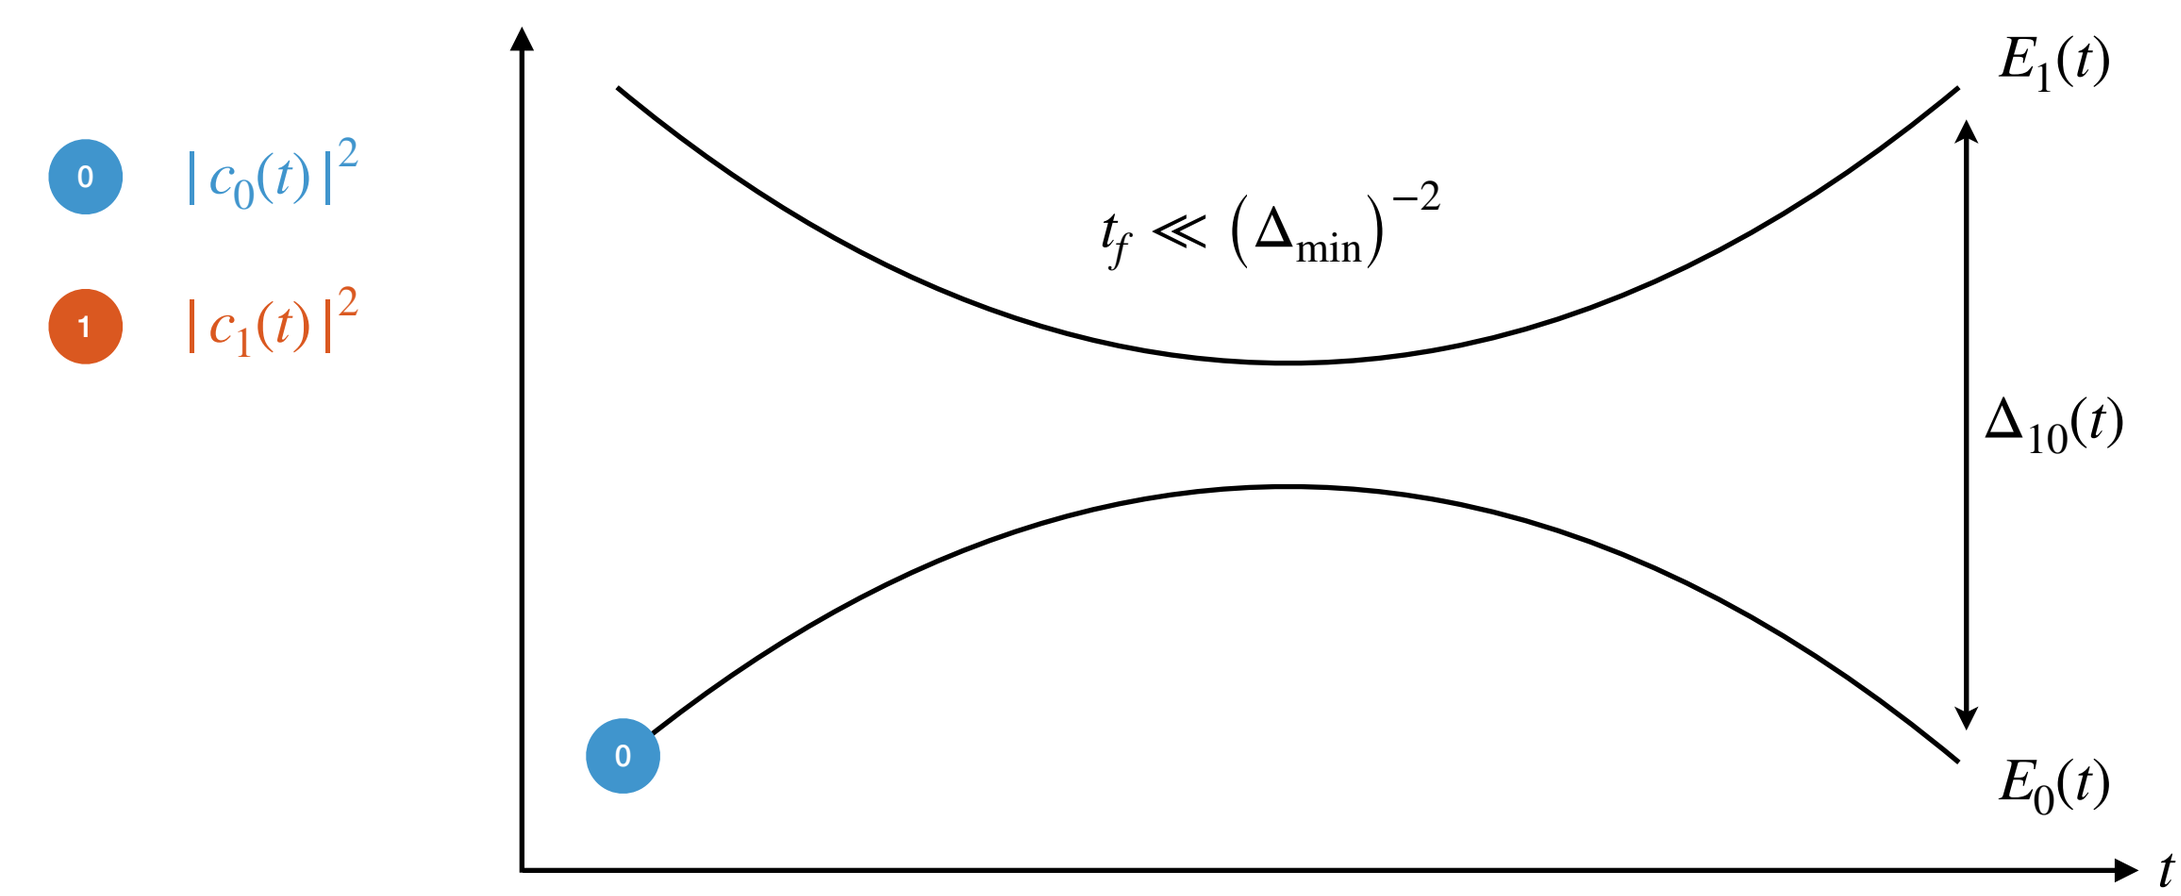

Examples of evolutions: fast

We will write our quantum state evolutions in the instantaneous energy eigenbasis

$$|\psi(t)\rangle = \sum_k c_k(t)\,|E_k(t)\rangle$$

Let's assume we are initially completely in the ground state
We will have qualitatively three kinds of behavior

(1) Diabatic transition to excited state
Evolution very fast relative to gap

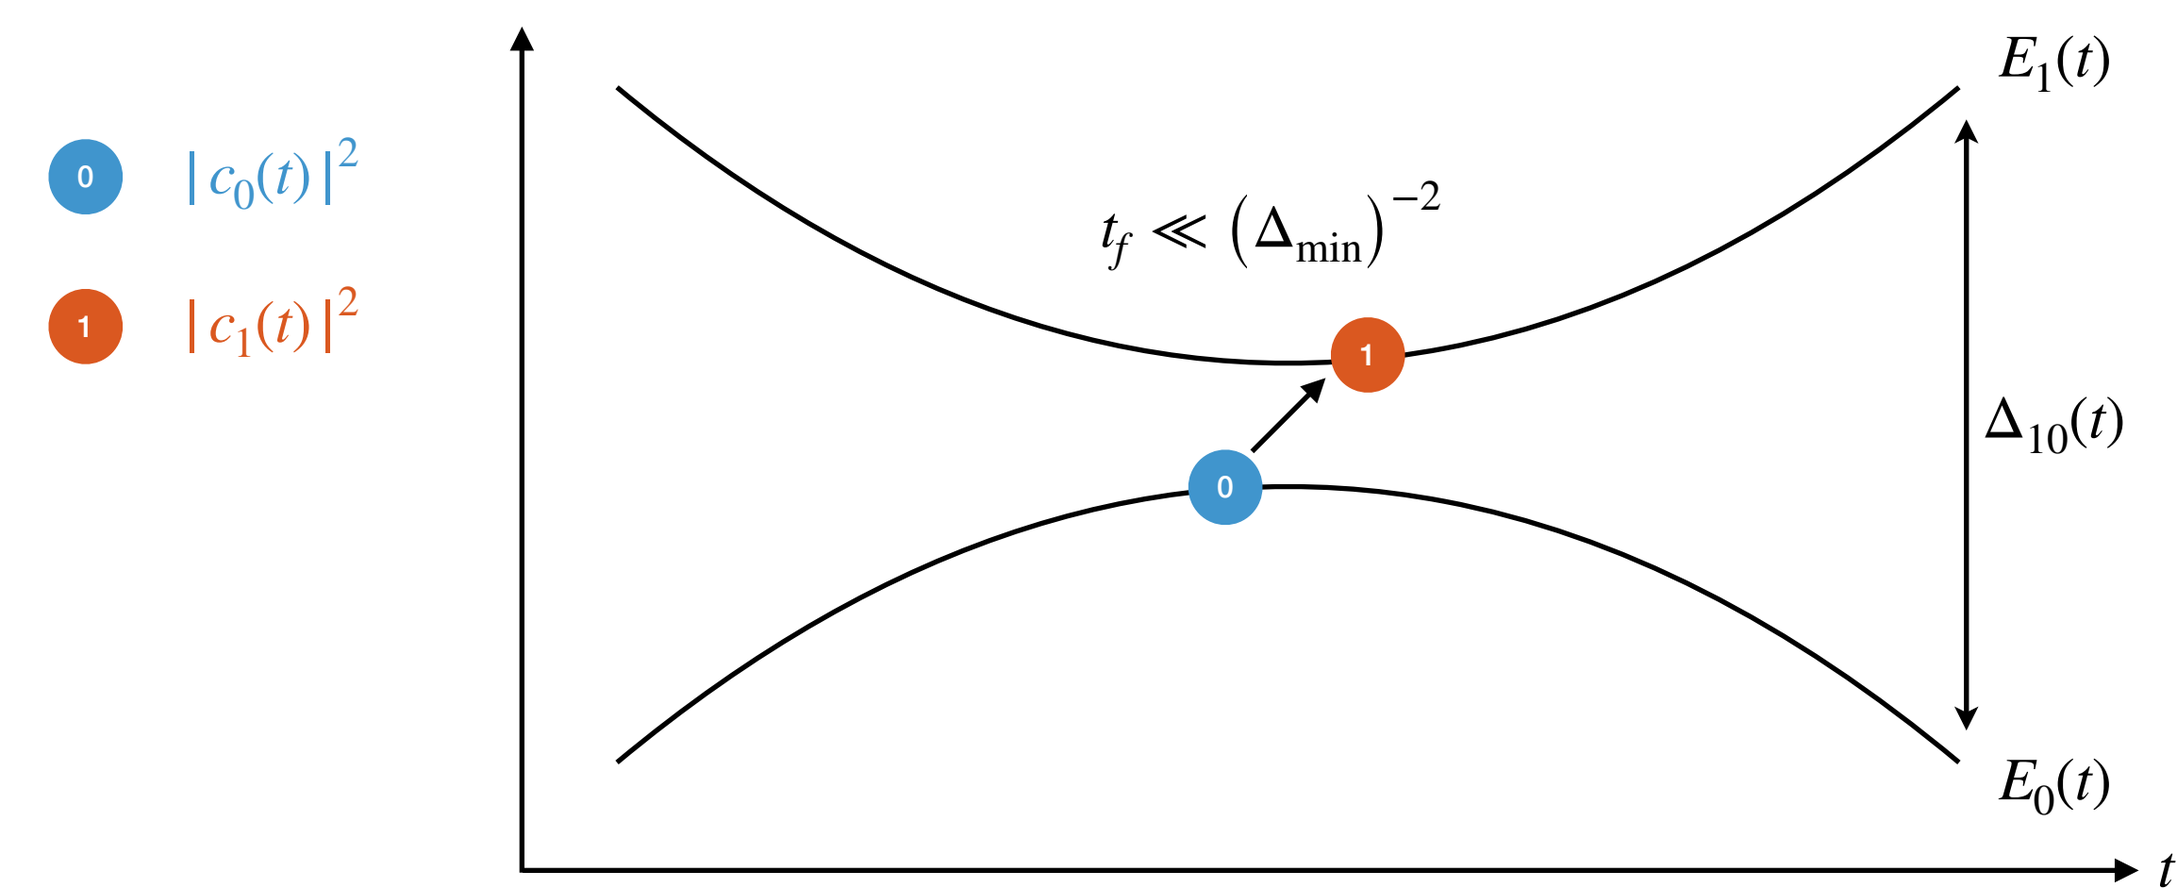

Examples of evolutions: fast

We will write our quantum state evolutions in the instantaneous energy eigenbasis

$$|\psi(t)\rangle = \sum_k c_k(t)\,|E_k(t)\rangle$$

Let's assume we are initially completely in the ground state
We will have qualitatively three kinds of behavior

(1) Diabatic transition to excited state
Evolution very fast relative to gap

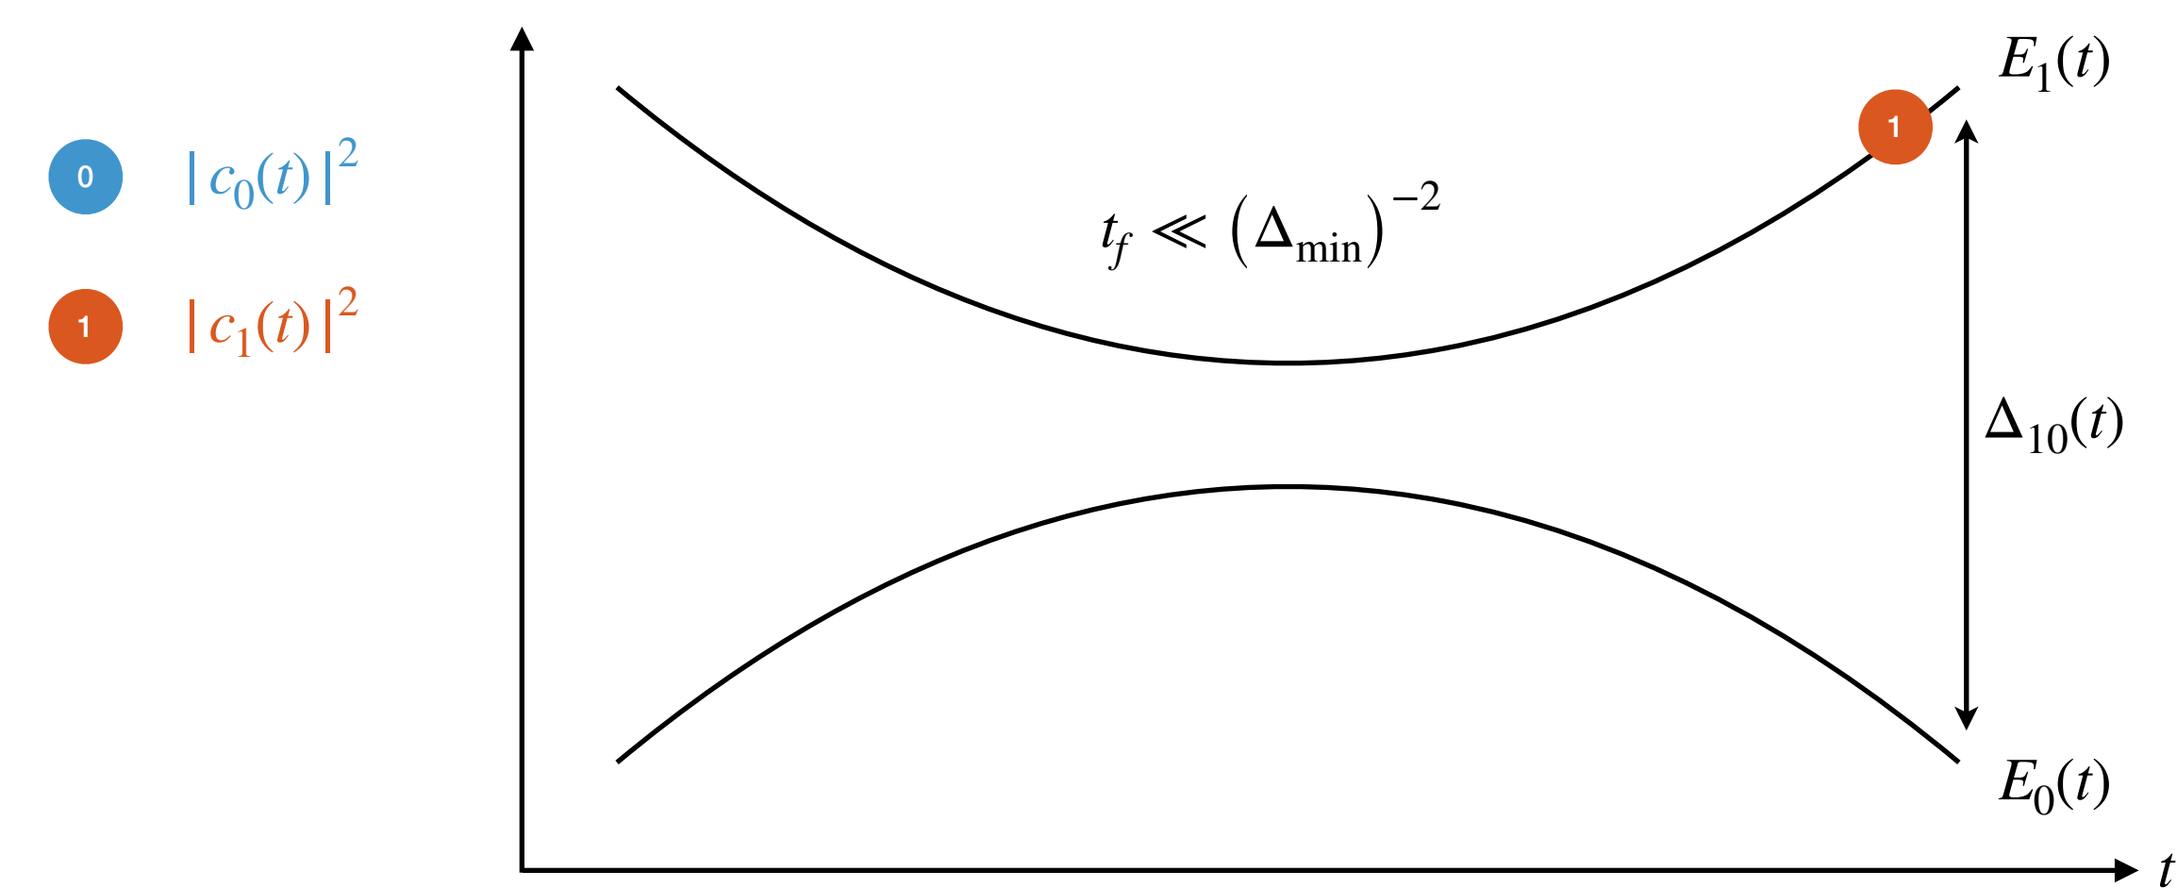

Examples of evolutions: medium

We will visualize wavefunction evolutions in the instantaneous energy eigenbasis

$$|\psi(t)\rangle = \sum_k c_k(t)\,|E_k(t)\rangle$$

Let's assume we are initially completely in the ground state
We will have qualitatively three kinds of behavior

(2) Partial transition to excited state
Evolution still fast relative to gap

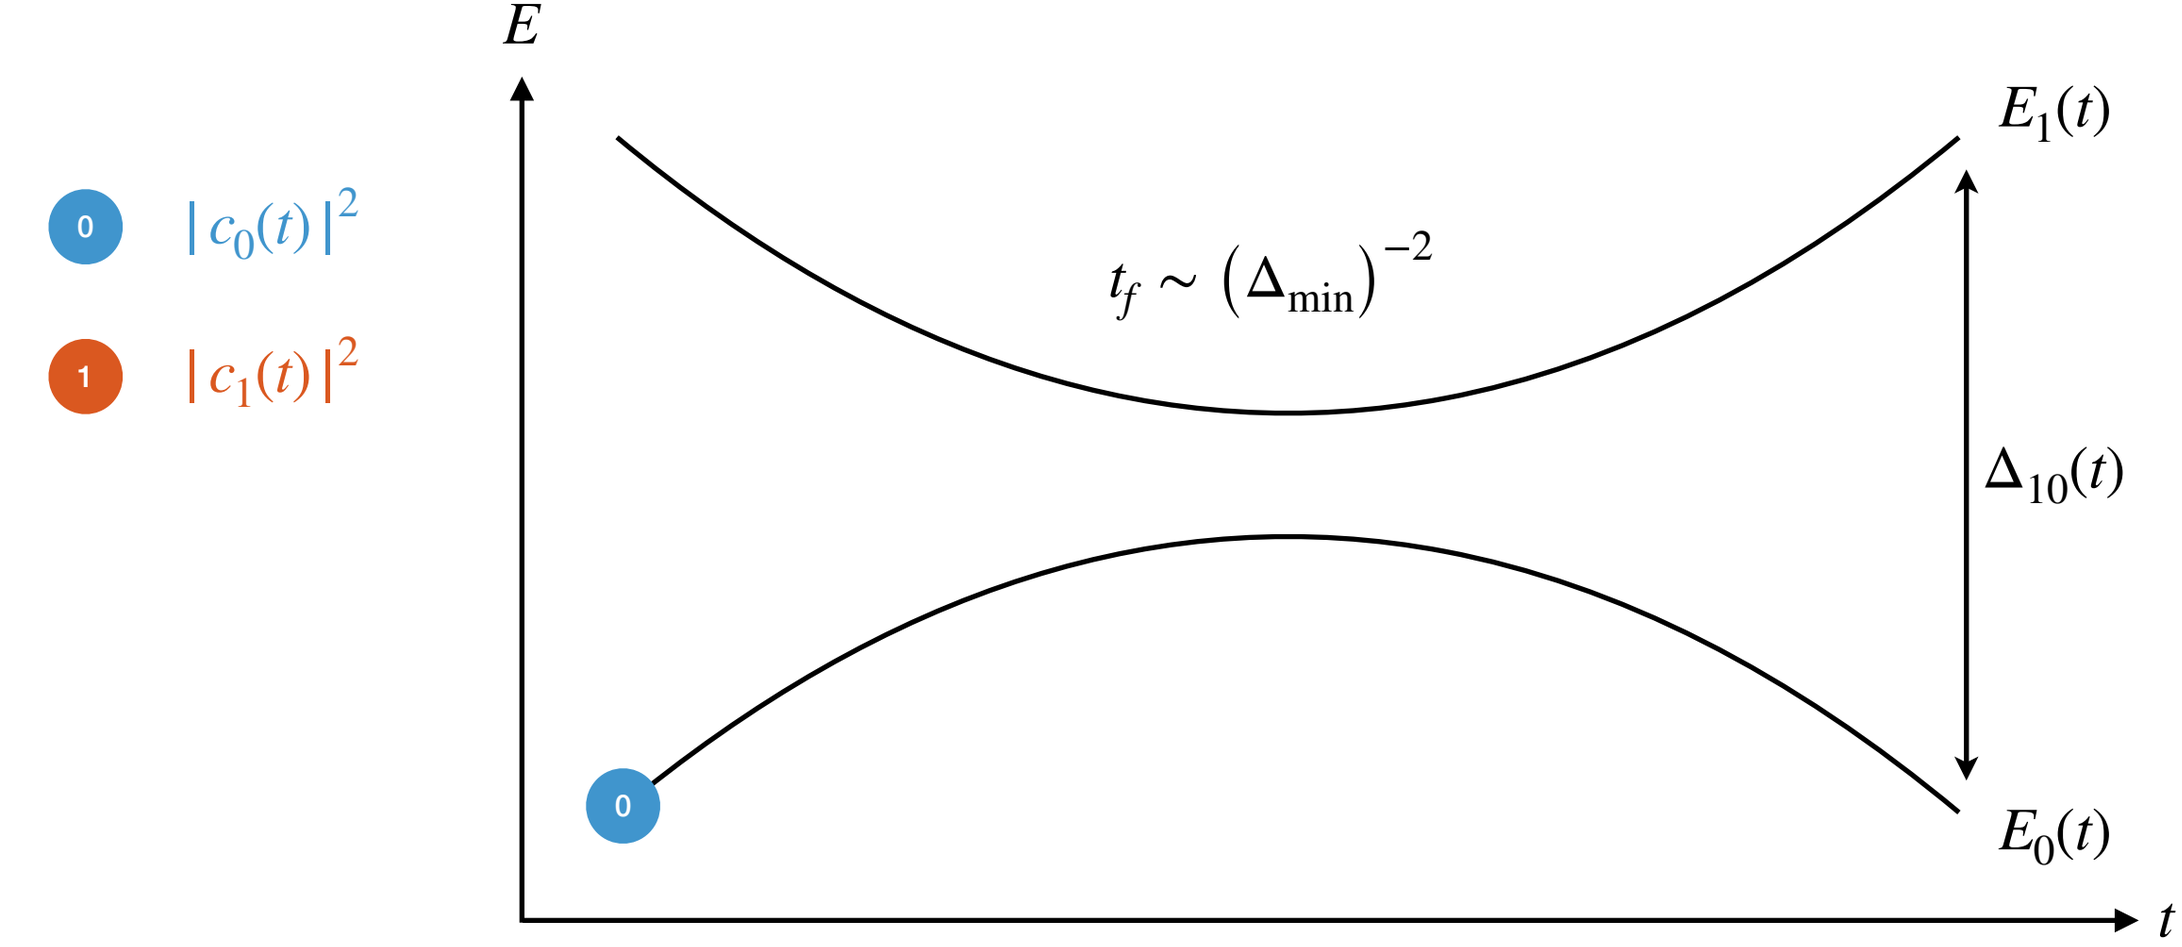

Examples of evolutions: medium

We will visualize wavefunction evolutions in the instantaneous energy eigenbasis

$$|\psi(t)\rangle = \sum_k c_k(t)\,|E_k(t)\rangle$$

Let's assume we are initially completely in the ground state
We will have qualitatively three kinds of behavior

(2) Partial transition to excited state
Evolution still fast relative to gap

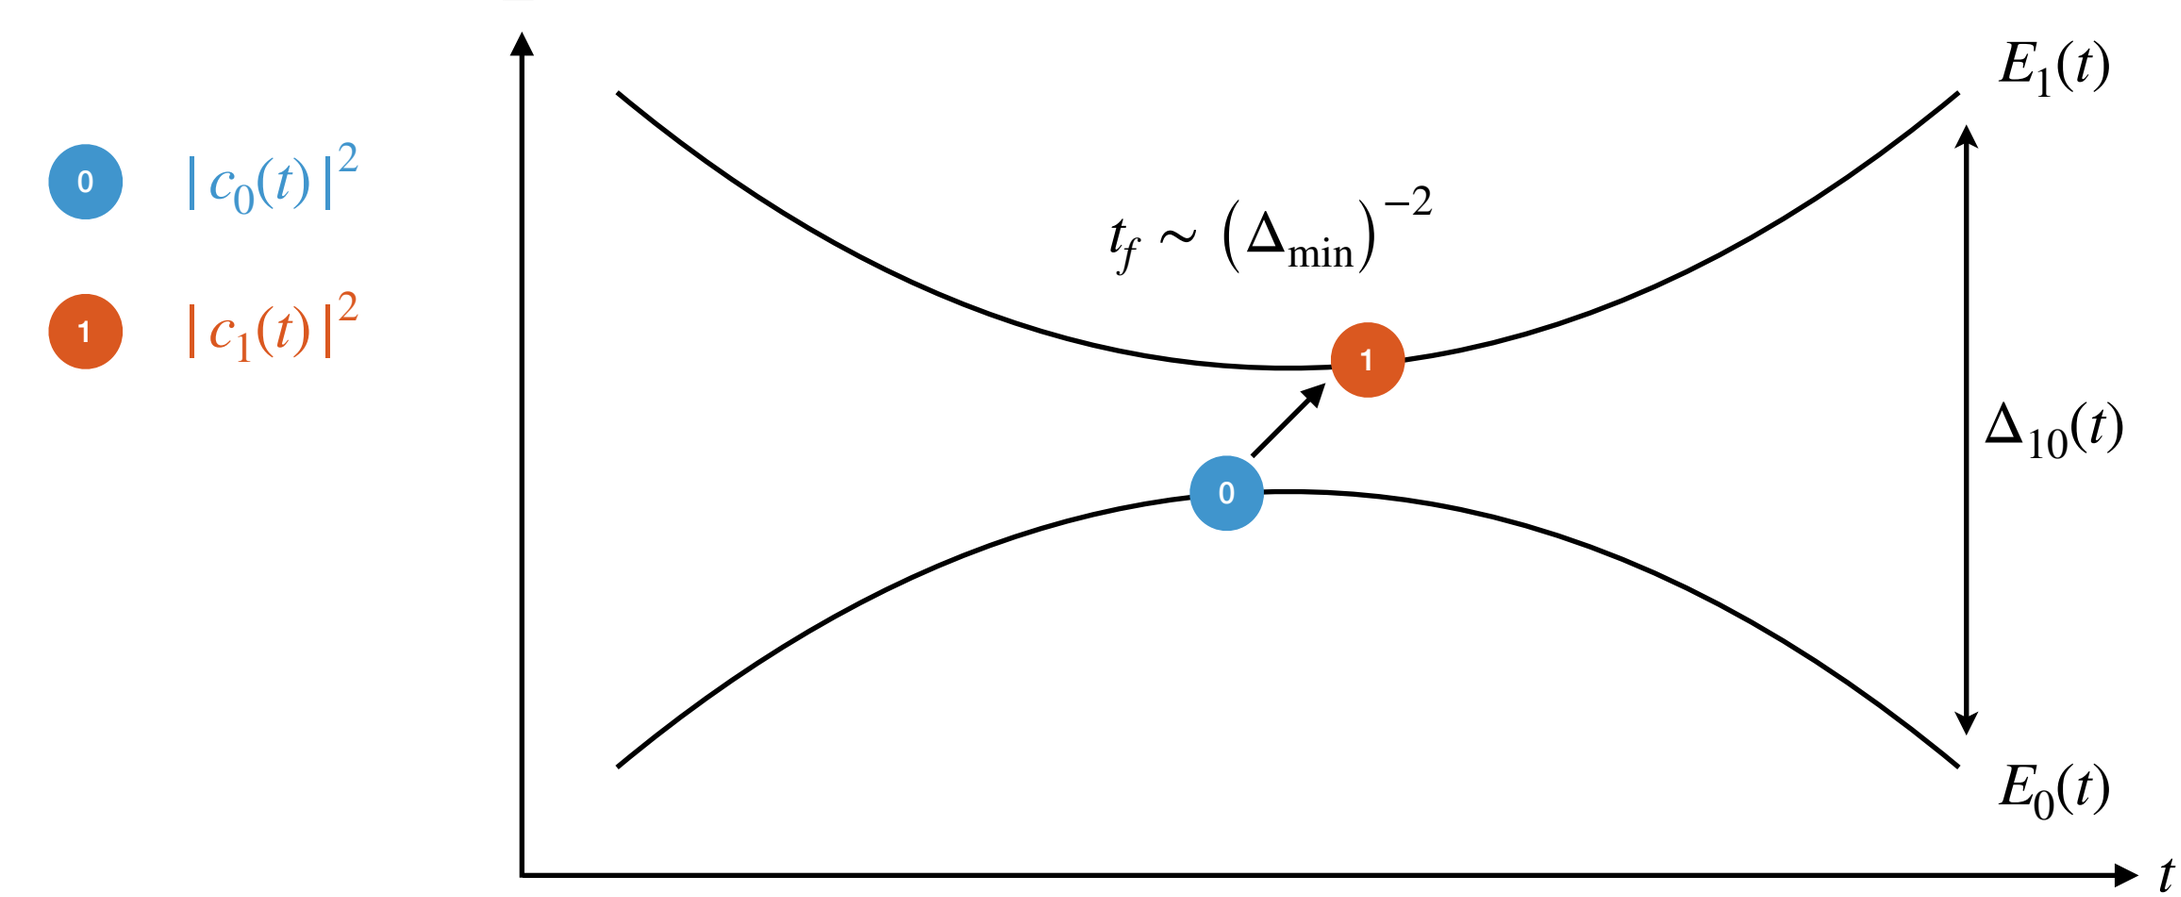

Examples of evolutions: medium

We will visualize wavefunction evolutions in the instantaneous energy eigenbasis

$$|\psi(t)\rangle = \sum_k c_k(t)\,|E_k(t)\rangle$$

Let's assume we are initially completely in the ground state
We will have qualitatively three kinds of behavior

(2) Partial transition to excited state
Evolution still fast relative to gap

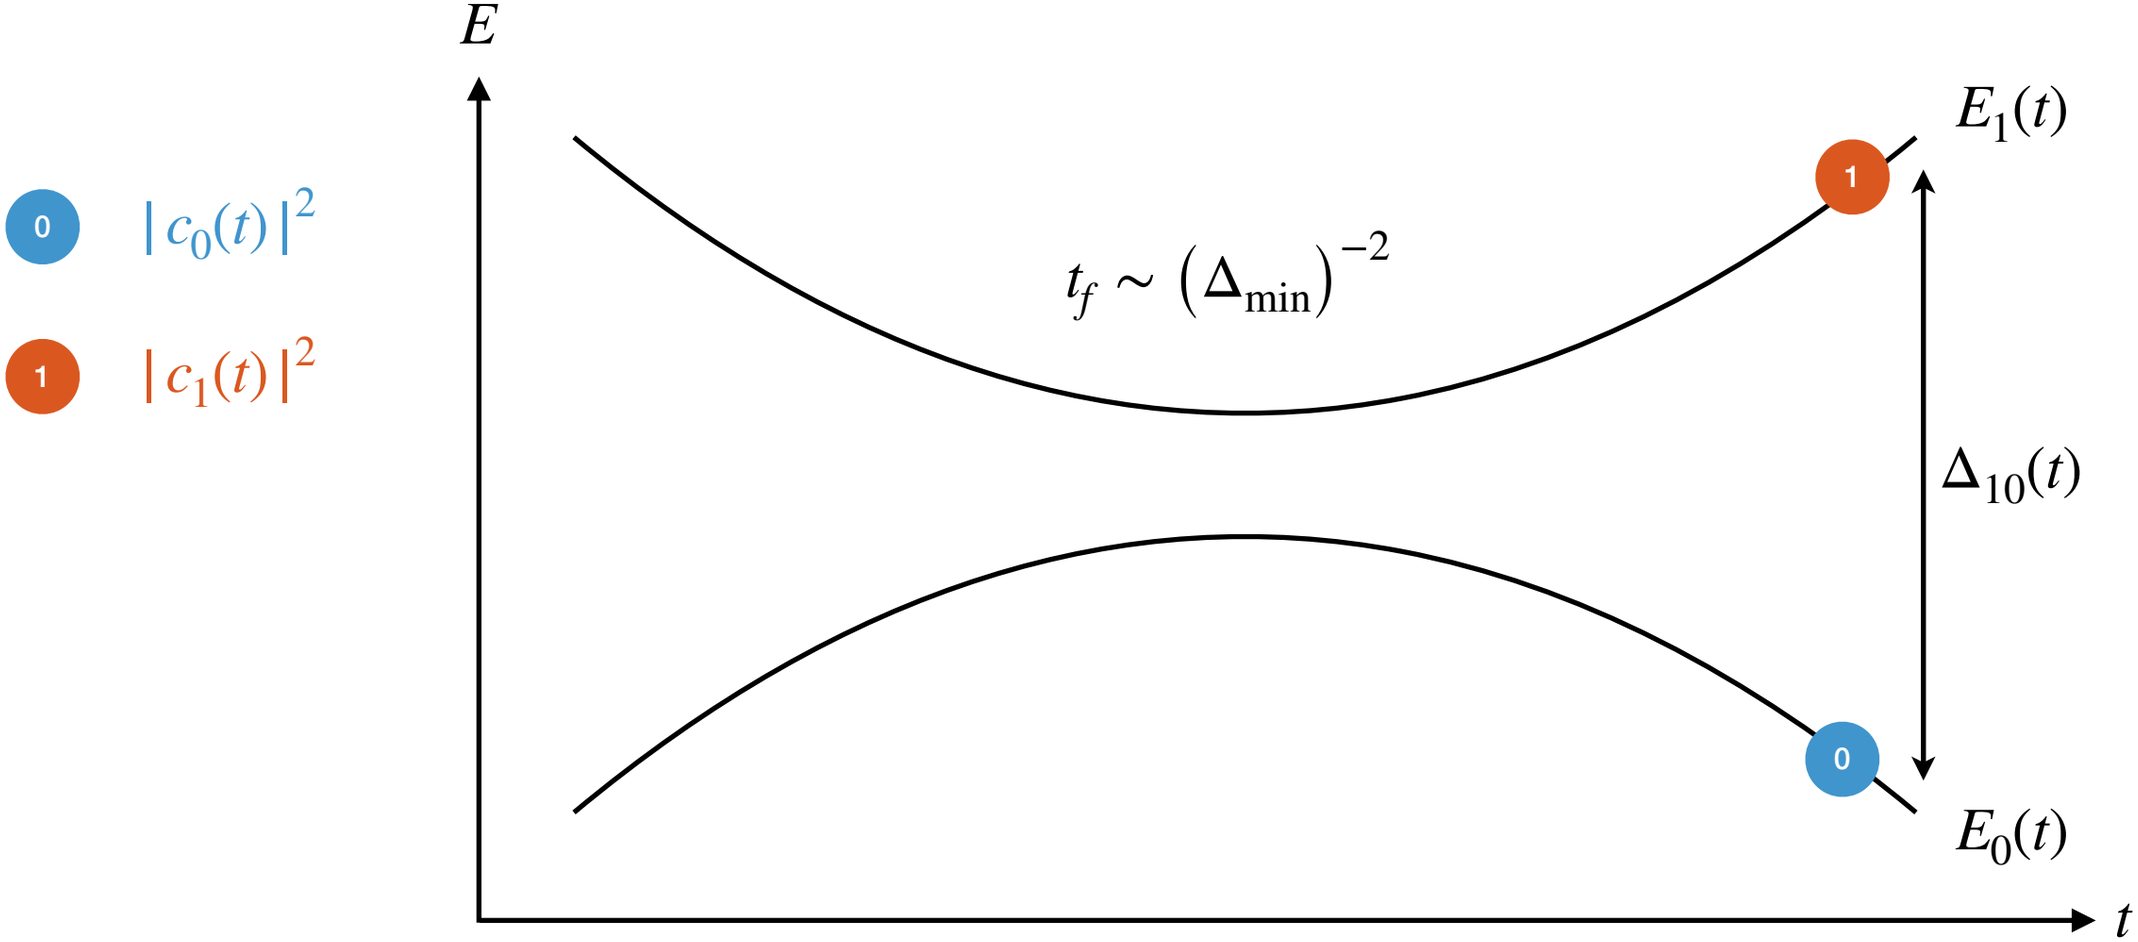

Examples of evolutions: slow

We will visualize wavefunction evolutions in the instantaneous energy eigenbasis

$$|\psi(t)\rangle = \sum_k c_k(t)\,|E_k(t)\rangle$$

Let's assume we are initially completely in the ground state
We will have qualitatively three kinds of behavior

(3) Negligible transitions (adiabatic)
Evolution slow enough relative to gap

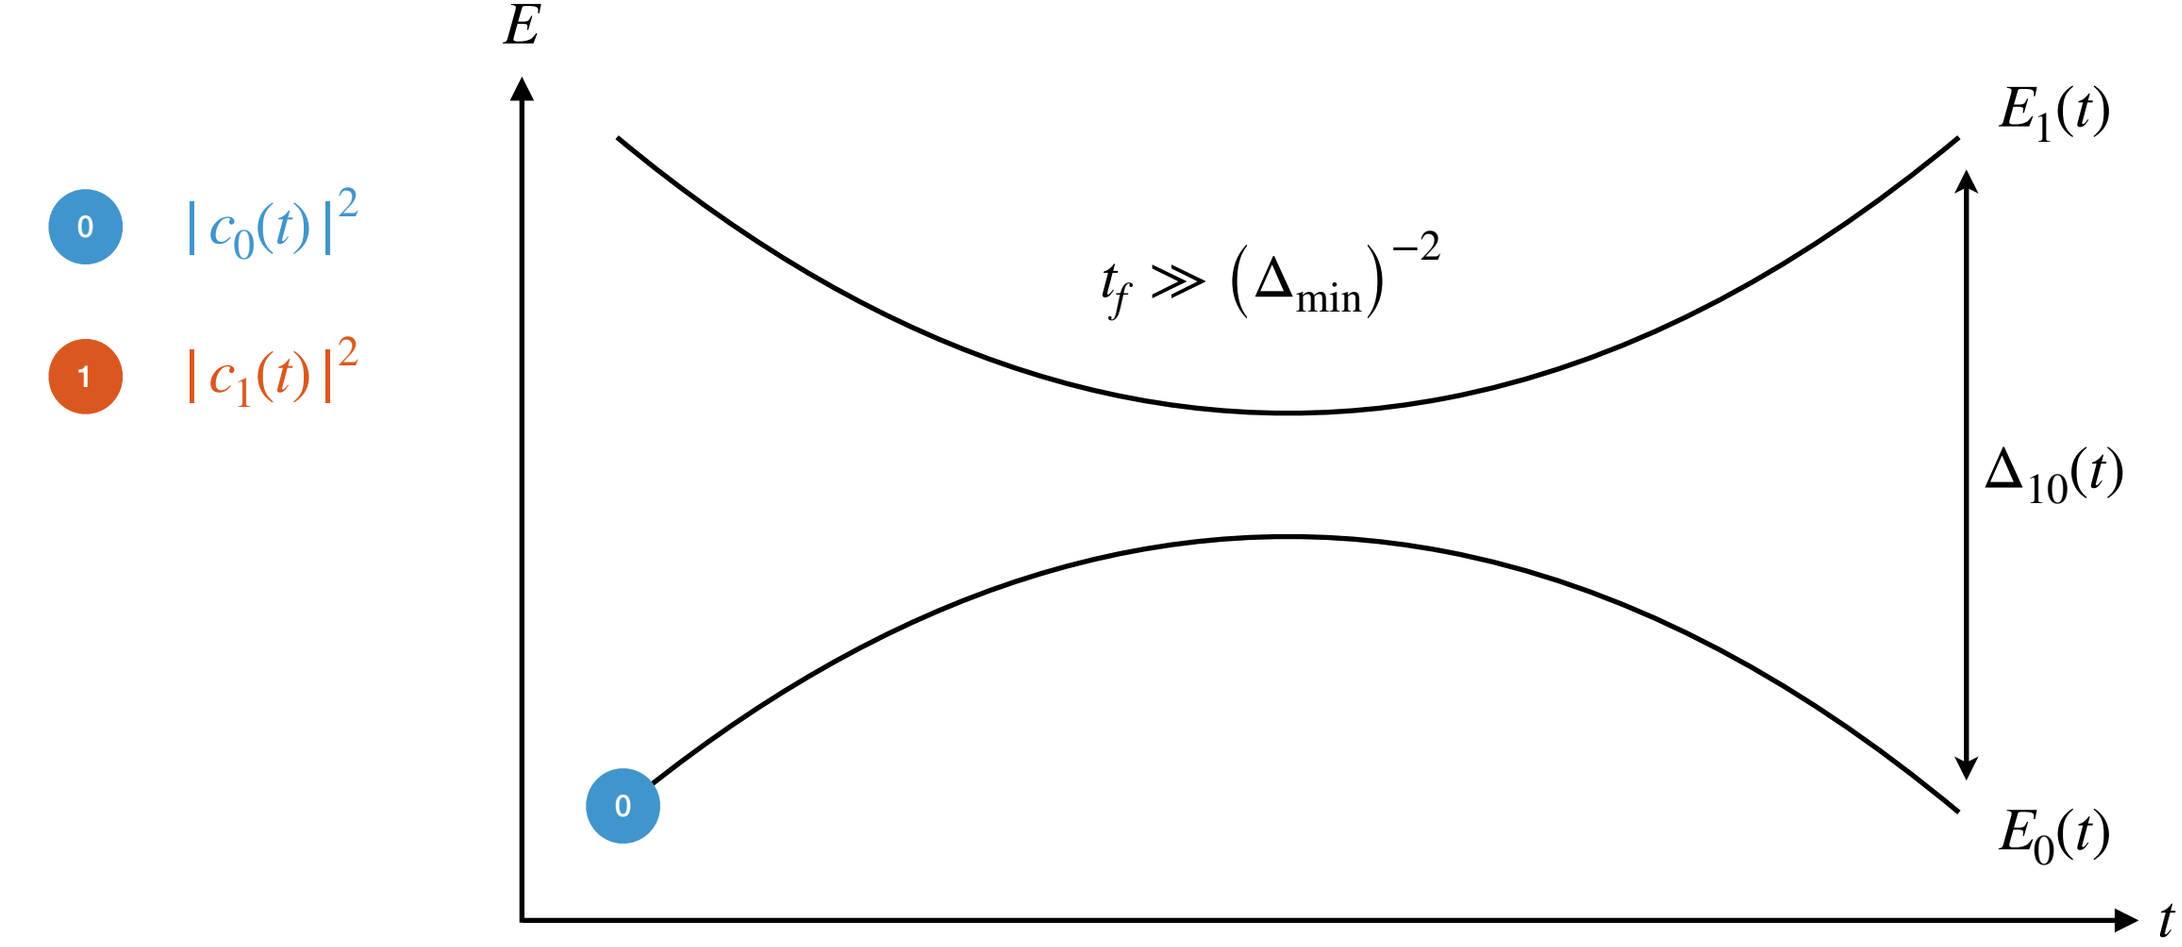

Examples of evolutions: slow

We will visualize wavefunction evolutions in the instantaneous energy eigenbasis

$$|\psi(t)\rangle = \sum_k c_k(t)\,|E_k(t)\rangle$$

Let's assume we are initially completely in the ground state
We will have qualitatively three kinds of behavior

(3) Negligible transitions (adiabatic)
Evolution slow enough relative to gap

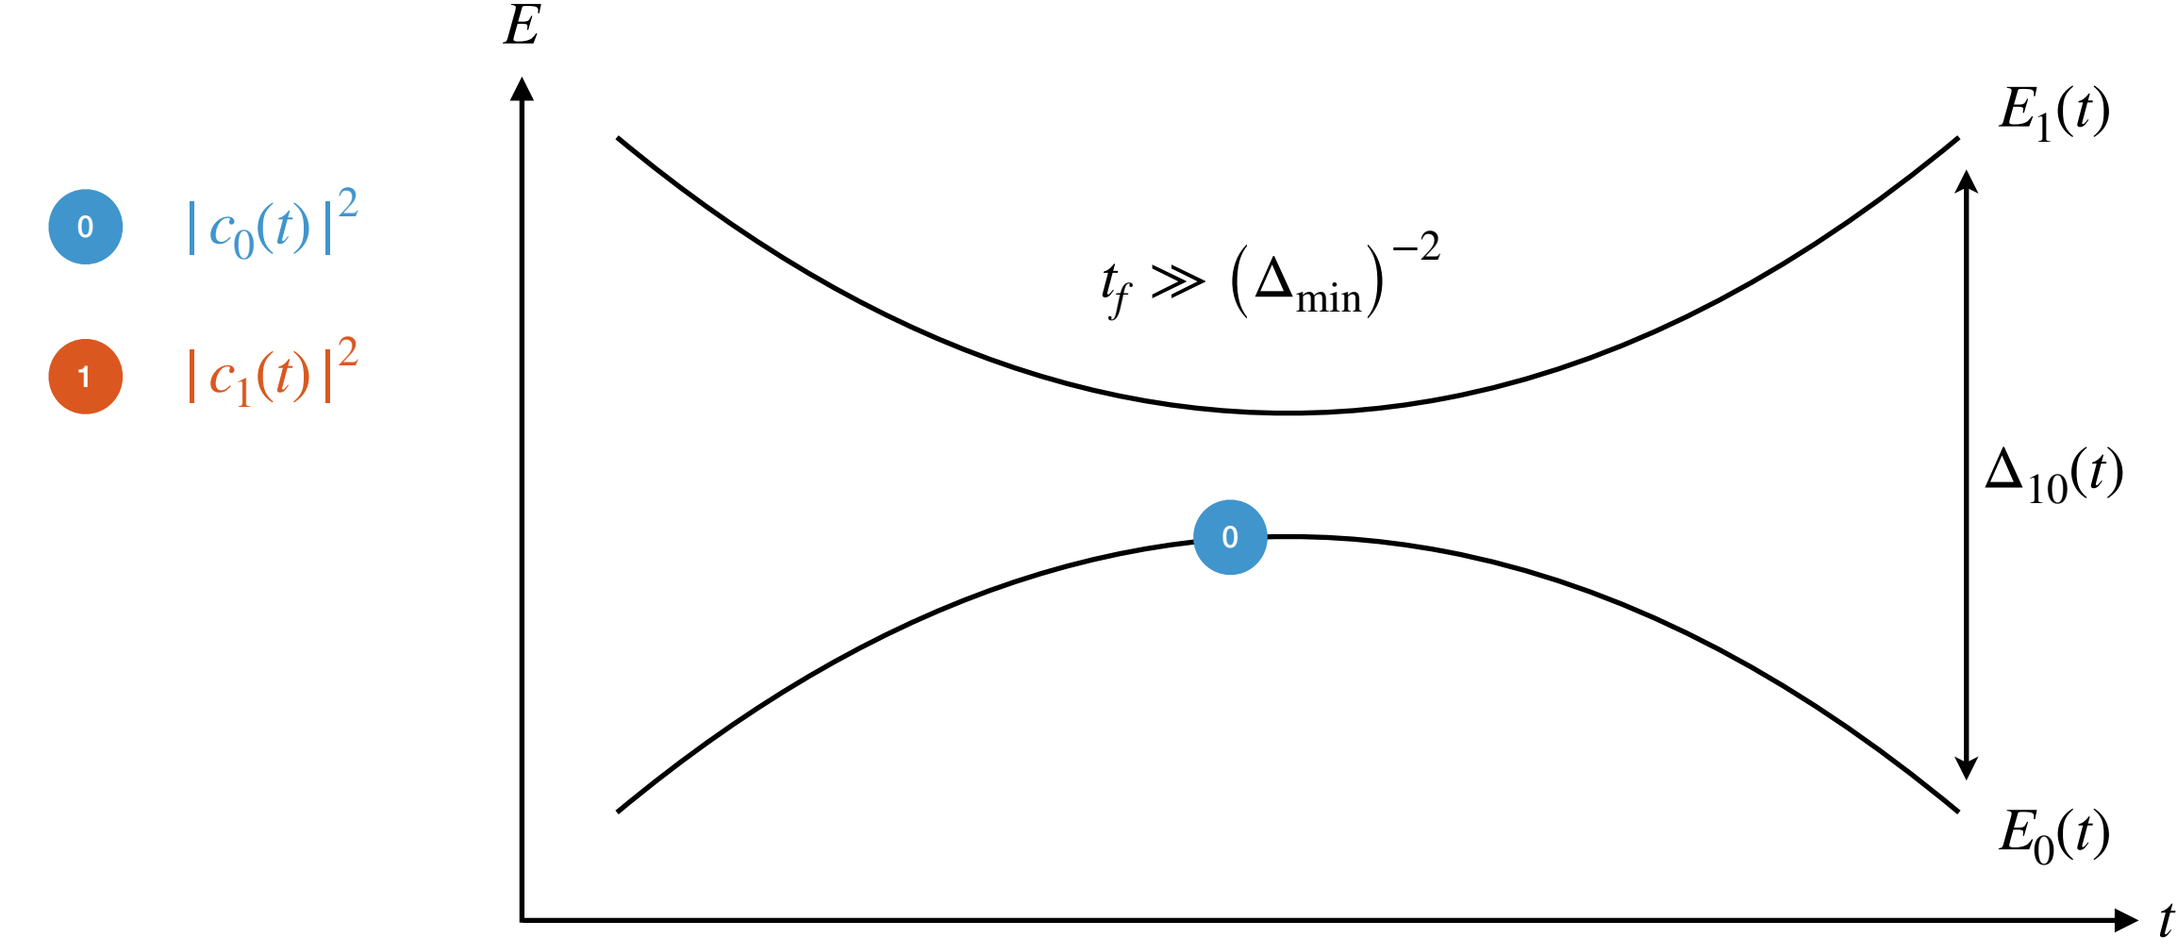

Examples of evolutions: slow

We will visualize wavefunction evolutions in the instantaneous energy eigenbasis

$$|\psi(t)\rangle = \sum_k c_k(t)\,|E_k(t)\rangle$$

Let's assume we are initially completely in the ground state
We will have qualitatively three kinds of behavior

(3) Negligible transitions (adiabatic)
Evolution slow enough relative to gap

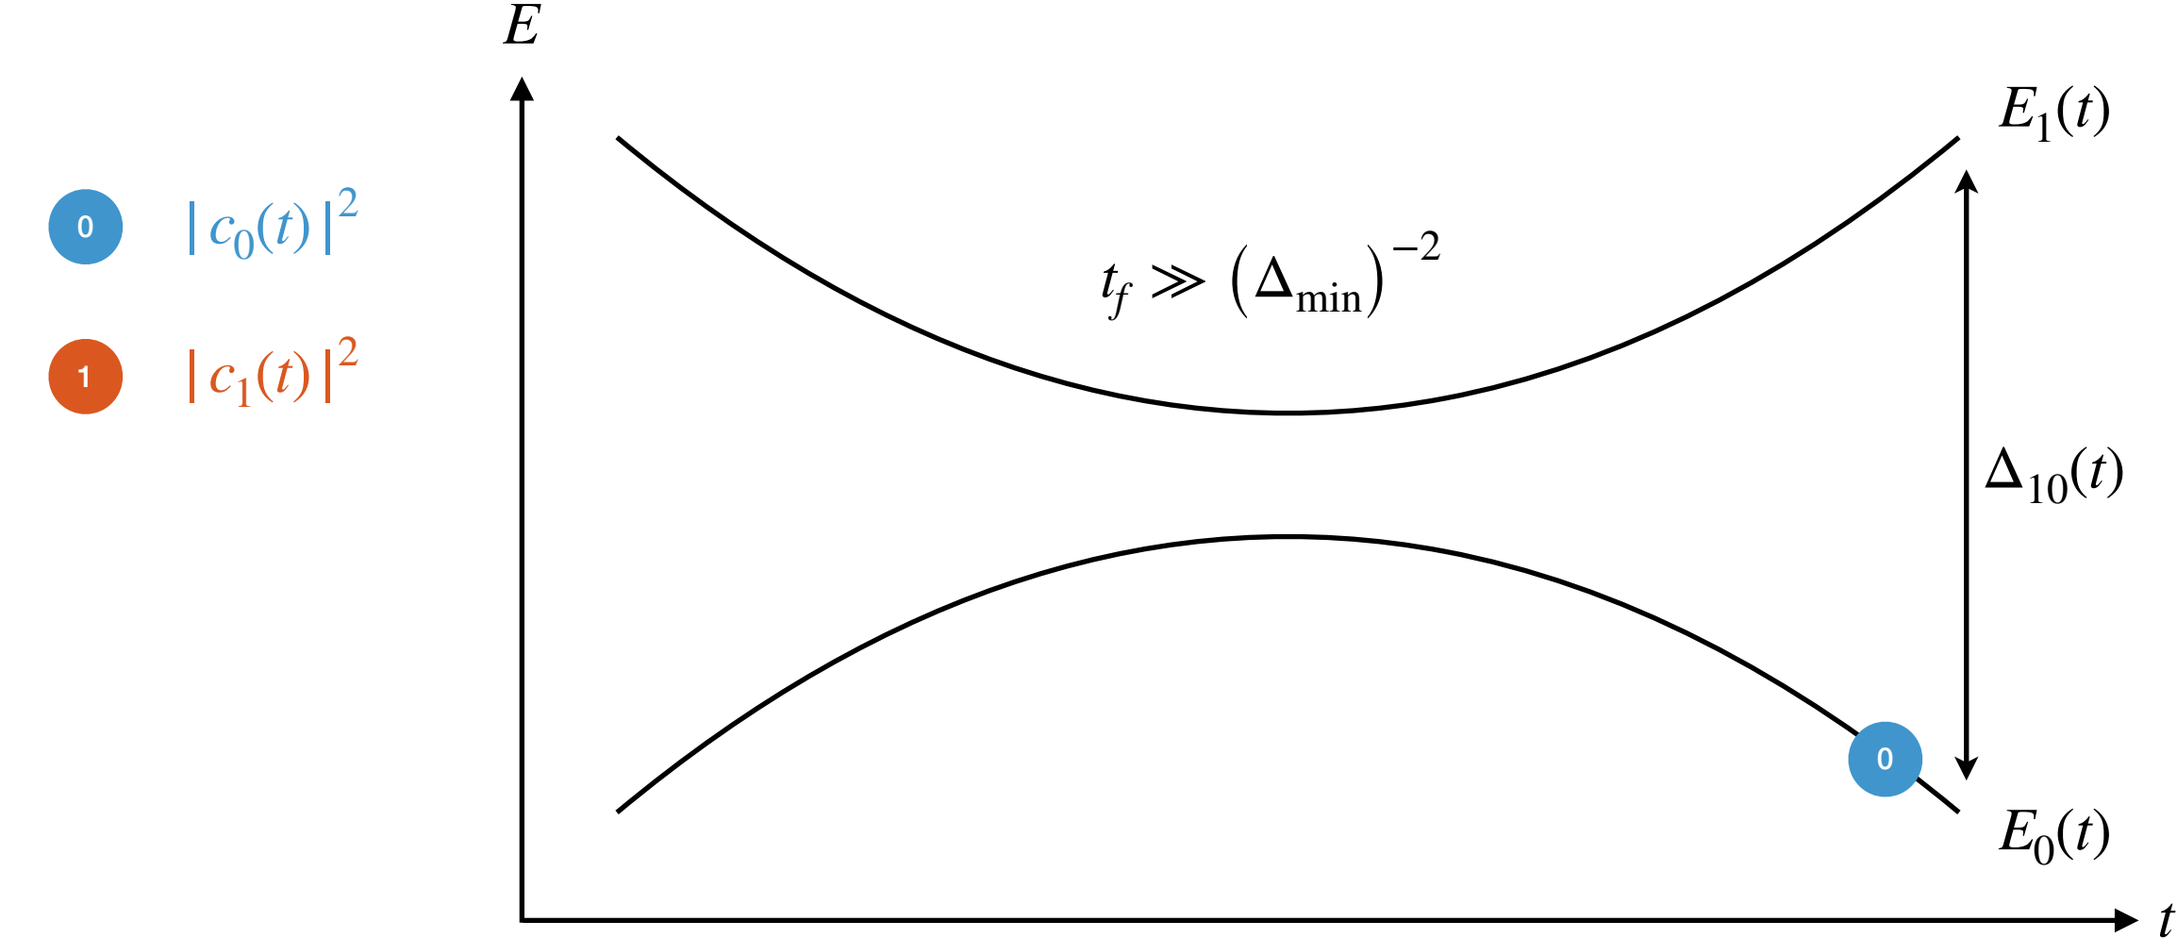

Adiabatic Evolution as an Algorithm

Adiabatic algorithm [Apo1988,Apo1989,Som1991,Ama1993,Fin1994,Kad1998,Far2000] is defined by

(1) a Hamiltonian H_0 with an easily (efficiently) prepared ground state.

ex. Transverse field Hamiltonian: $H_0 = -\sum_{i=1}^{n} \sigma_i^x \rightarrow |E_0\rangle = |+\rangle^{\otimes n}$

(2) a Hamiltonian H_1 whose ground state encodes the solution to a computational problem. The objective is to find this ground state. (Examples to come)

(3) a continuous interpolating Hamiltonian $H(s)$ that interpolates between $H(s=0) = H_0$ and $H(s=1) = H_1$.

ex. Linear interpolation: $H(s) = (1-s)H_0 + sH_1 \quad (s = t/t_f)$

Starting from the ground state of H_0, an evolution satisfying the adiabatic condition gives a guarantee of finding the ground state of H_1 with high probability.

Gives a very generic algorithm for preparing ground states of Hamiltonians that can be physically implemented, ex. Molecular Hamiltonians

Efficiency of Adiabatic Algorithm

Efficiency of the adiabatic algorithm is given by how t_f must scale with system size n to ensure adiabatic condition is satisfied.

Numerator of the adiabatic condition scales at most polynomially with system size, so typically the dominant dependence on system size comes from the minimum gap along the interpolation $\Delta_{\min}$. Reporting the scaling of the minimum gap is then equivalent to reporting the efficiency of the algorithm.

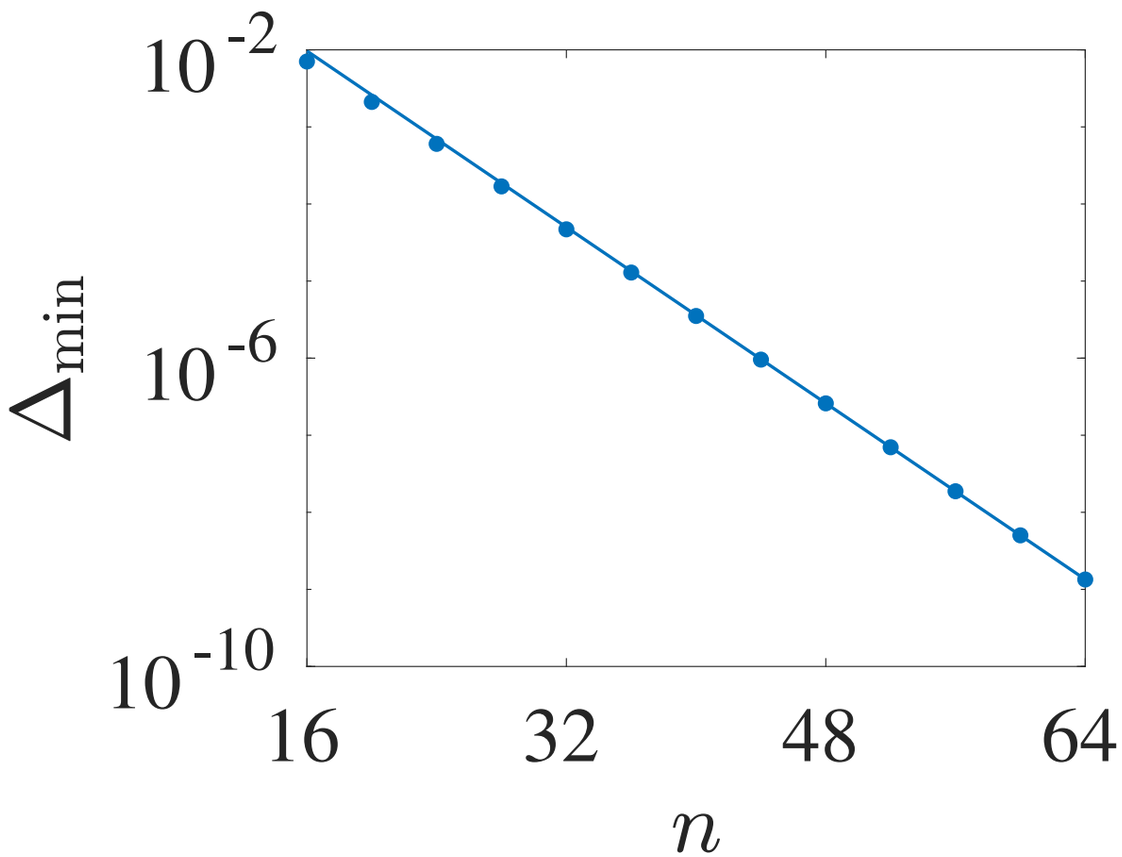

Exponentially closing gap → exponential time algorithm

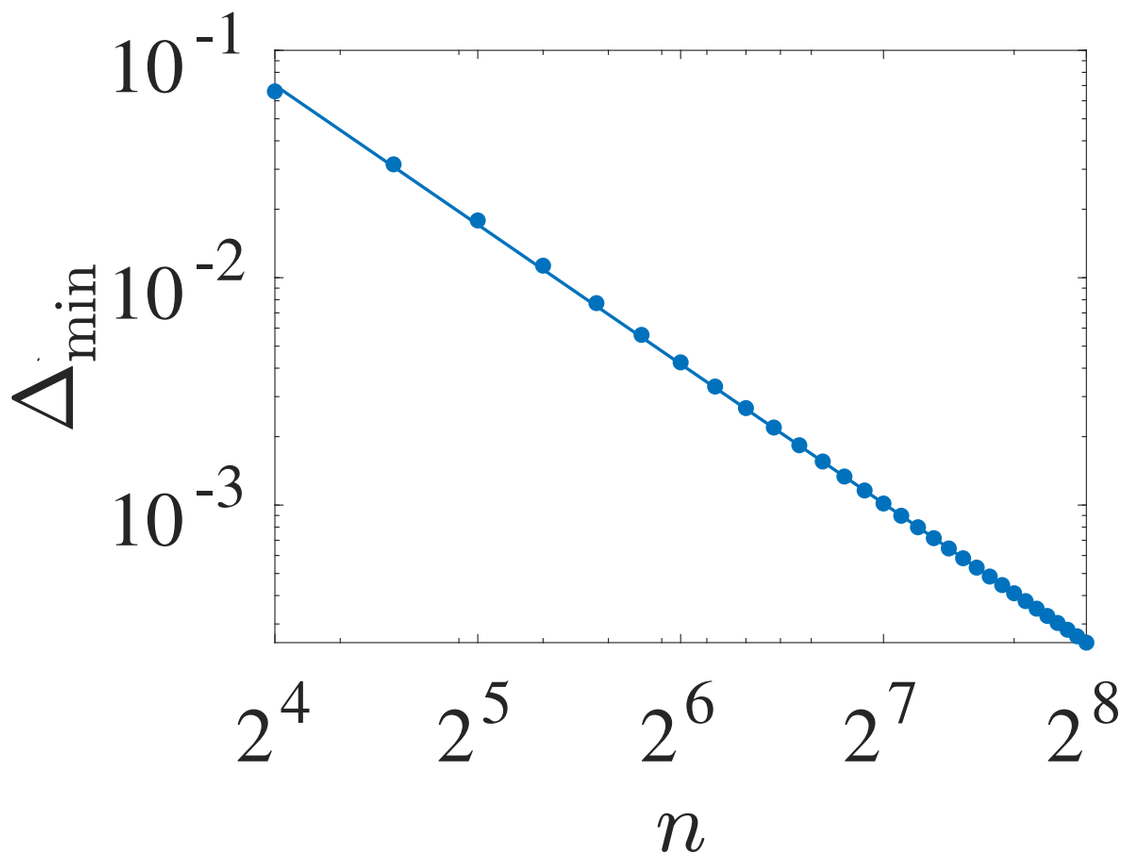

Polynomially closing gap → polynomial time algorithm

Equivalence of Adiabatic and Circuit Model

We can encode the output of *any* quantum circuit of depth L (to simplify the construction, we assume all single and two qubit gates are applied sequentially) into the ground state of a Hamiltonian $H(1)$ [Kit2002] such that there exists an interpolation $H(s)$ with a minimum gap $\Delta_{\min}$ that scales polynomially with L [Aha2007].

Better estimates of the gap were presented in [Dei2007,Doo2020]; improvements on ground state overlap with the state at the end of the circuit [Cah2018,Bau2018]; we can expect limits to how much the gap can be improved in these constructions [Gan2013,Gon2018].

This Circuit-to-Hamiltonian mapping shows that adiabatic quantum computation is equivalent to circuit quantum computation up to a polynomial overhead, $t_f \sim \mathrm{poly}(L)$.

- ➡ Any exponential speedups in the circuit model can be realized in principle in the adiabatic model.
- ➡ Overhead of the mapping may mean that polynomial speedups are not preserved.

The above Circuit-to-Hamiltonian mapping can be thought of as a “proof-of-principle” construction and not necessarily the most efficient way to realize adiabatic algorithms. More native constructions can reproduce circuit speedups (see [Hen2014] for an example).

Universal Adiabatic Quantum Computer

A minimal set of programmable Hamiltonian terms with real coefficients to implement universal adiabatic quantum computing are [Bia2008]:

$$\left\{ \sigma_i^x, \sigma_i^z, \sigma_i^x \sigma_j^x, \sigma_i^z \sigma_j^z \right\}$$

$$\left\{ \sigma_i^x, \sigma_i^z, \sigma_i^x \sigma_j^z, \sigma_i^z \sigma_j^x \right\}$$

ex. $$H(s) = \sum_i (h_i^x(s)\sigma_i^x + h_i^z(s)\sigma_i^z) + \sum_{i<j} (J_{ij}^x(s)\sigma_i^x\sigma_j^x + J_{ij}^z(s)\sigma_i^z\sigma_j^z)$$

The local Hamiltonian problem associated with this Hamiltonian (at a fixed interpolation point s) is also QMA-complete (for a review, see [Aha2002,Gha2013,Gha2015])

Even if you had a universal adiabatic quantum computer, there is no guarantee that you could find the ground state of all the programmable Hamiltonians efficiently (QMA is not expected to be in BQP)

Example: Unstructured Search (Grover search)

(1) Computational task is to find a 'marked' state m among $N = 2^n$ possible states [Gro1996]. Encode problem in the n-qubit Hamiltonian:

$$H_1 = 1 - |m\rangle\langle m|$$

Ground state is $|m\rangle$ with energy 0, and all other states have energy 1.

Energy landscape has no structure and gives no indication of direction to go in. "Searching for a needle in a haystack". Only classical approach is to randomly search all states, which takes on average $\mathcal{O}(N)$ time (the time it takes to find the marked state scales linearly with the number of states to search)

(2) Pick initial Hamiltonian with easily prepared ground state [Rol2002]:

$$H_0 = 1 - |\phi\rangle\langle\phi|$$

where $|\phi\rangle = |+\rangle^{\otimes n} = \frac{1}{\sqrt{N}} \sum_x |x\rangle$. Ground state is $|E_0\rangle = |\phi\rangle$ with energy 0.

(3) Consider the linear interpolation $H(s) = (1 - s)H_0 + sH_1$

Example: Unstructured Search (Grover problem)

Hamiltonian spectrum is very simple for this problem

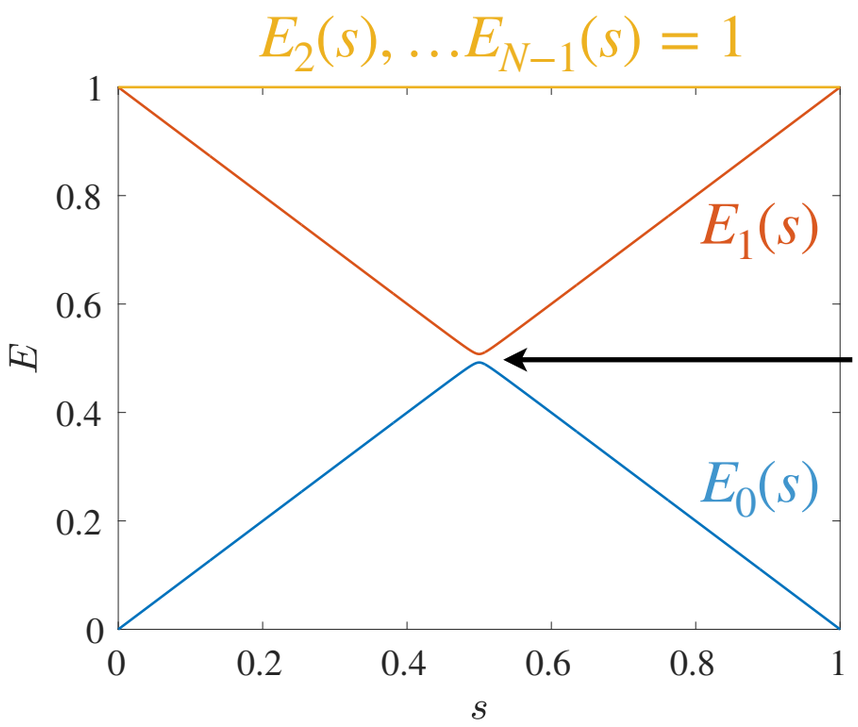

$$\Delta_{10}(s = 1/2) = \frac{1}{\sqrt{N}}$$

Adiabatic condition predicts a runtime scaling of $\mathcal{O}(N)$. This does not achieve the quadratic speedup of Grover's algorithm [Gro1996] in the circuit model.

But we are free to choose other interpolation schedules. We can choose an interpolation the slows down around the minimum gap and speeds up elsewhere. Specifically, we can choose an interpolation that is locally adiabatic [Rol2002]

$$H(s) = \big(1 - \Phi(s)\big)\, H_0 + \Phi(s) H_1$$

$$\Phi(s) = \frac{1}{2} + \frac{1}{2\sqrt{N-1}} \tan\left[(2s-1)\tan^{-1}\sqrt{N-1}\right]$$

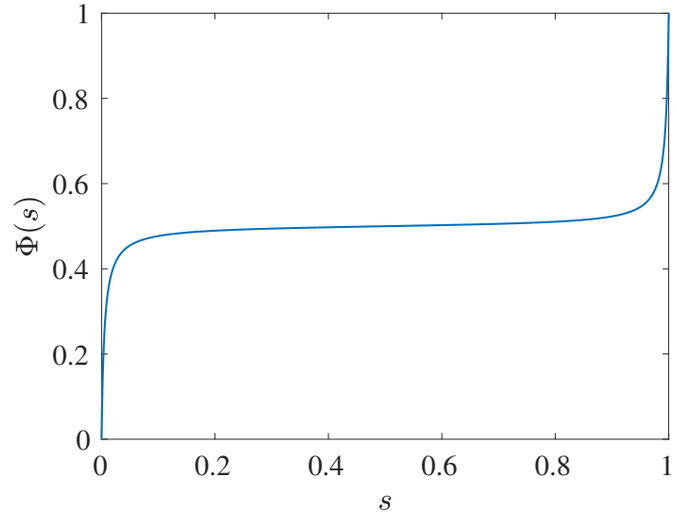

New runtime scaling is $\mathcal{O}(\sqrt{N})$ [Rol2002], recovering quadratic speedup of Grover's algorithm [Gro1996].

Example: Unstructured Search (Grover problem)

The energy landscape of H_1 is different from that of $H(s)$

Using the spin-coherent path integral formalism, we can associate with our interpolating Hamiltonian a semi-classical energy potential $V(s)$ [Kla1976]:

$$V(s) = \langle \Omega | H(s) | \Omega \rangle$$

$$|\Omega\rangle = \otimes_{j=1}^{n} \left(\cos\frac{\theta_j}{2} |0_j\rangle + e^{i\varphi_j} \sin\frac{\theta_j}{2} |1_j\rangle \right)$$

Semiclassical energy landscape of $2n$ variables

The global minimum of this landscape is the spin-coherent approximation of the quantum ground state, and an adiabatic evolution attempts to follow this global minimum. The bottleneck is when the global minimum of this landscape changes discontinuously.

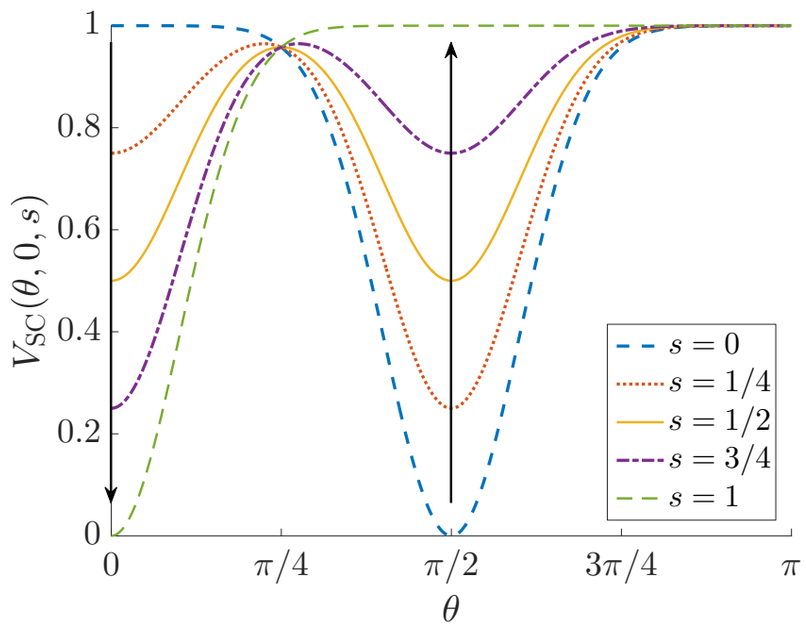

If marked state is all-zero state, Hamiltonian has qubit-permutation symmetry, and we can consider the case where $\theta_j = \theta, \varphi_j = 0, \forall j$.

Scaling properties of the barrier separating the two minima coincides with scaling of the minimum gap [Kon2015,Mut2016].

Example: Unstructured Search (Grover problem)

Important caveats:

Expressing the problem Hamiltonian, $H_1 = 1 - |m\rangle\langle m|$, in terms of Pauli operators requires up to n-body Pauli operators $\rightarrow$ a highly non-local operator.

Example of an oracular speedup: given access to such a Hamiltonian, a quadratic speedup over classical approaches is possible.

Optimized interpolation schedule works because the spectrum and the location of the minimum gap are fixed and independent of the marked state.

➡ Interpolation schedules can improve performance but requires knowledge of the position and magnitude of the minimum gap to be effective. Remains an active area of research on how to do this [Jar2018].

Example: Quantum Adiabatic Optimization

General optimization cost functions $C(x)$ over bit strings can be straightforwardly encoded as a Hamiltonian. Special case: Quadratic Unconstrained Binary Optimization problems (QUBOs) are equivalent to Ising Hamiltonians:

$$H_1 = \sum_x C(x)\,|x\rangle\langle x| \;\rightarrow \text{ex. QUBO: } H_1 = \sum_i h_i \sigma_i^z + \sum_{\langle i,j \rangle} J_{ij} \sigma_i^z \sigma_j^z$$

Finding the ground state of the Ising problem (on a non-planar graph in the absence of local fields, or on a planar graph in the presence of local fields) is NP-hard [Bar1982], so all problems in NP can be mapped to it (with some overhead)

NP Problem	Application
Traveling salesman	Logistics, vehicle routing
Minimum Steiner tree	Circuit layout, network design
Graph coloring	Scheduling, register allocation
MAX-CLIQUE	Social networks, bioinformatics
QUBO	Machine learning, software V&V
0-1 Integer Linear Programming	Natural language processing
Sub-graph isomorphism	Chem-informatics, drug discovery
Job shop scheduling	Manufacturing
MAX-2SAT	Artificial intelligence

Example: Quantum Adiabatic Optimization

Initial Hamiltonian, often called the driver Hamiltonian, is taken to be the uniform transverse field Hamiltonian

$$H_0 = -\sum_i \sigma_i^x \qquad\qquad H(s) = A(s)H_0 + B(s)H_1\ , \quad A(1) = 0,\ B(0) = 0$$

Evolving adiabatically to find the ground state of the “diagonal-in-the-computational basis” Hamiltonian H_1 is called *quantum adiabatic optimization*.

For special class of problems minimum gap can be bounded [Dam2001,Rei2004], so efficiency of adiabatic algorithm can be determined. These problems are again oracular in nature and are not physically meaningful problems.

In the case of local Ising problems, no known example of a speedup exists. This is both due to the fact that we cannot determine the gap scaling at large system sizes to determine the efficiency of the algorithm or has been ruled out by lots of hard work [Far2001,You2008,You2010,Far2012]. At the same time, we do not expect to be able to solve NP-hard problems in polynomial time on a quantum computer. But the hope is that there may still be restricted classes of problems that we can solve more efficiently than we can classically (ex. maybe not an exponential speedup but a quadratic speedup) .

Stoquastic vs Non-Stoquastic

There is folklore that Hamiltonians of the form:

$$H(s) = -A(s)\sum_i \sigma_i^x + B(s)H_{\text{Ising}}$$

are not amenable to a quantum advantage. This belief stems from the fact that these Hamiltonians are 'stoquastic'. A Hamiltonian is stoquastic if its off-diagonal elements (in the computational basis) are real and non-positive, up to local unitary transformations [Bra2008,Bra2009].

- Adiabatic ground state evolution using stoquastic Hamiltonians is not expected to be universal [Bra2008].
- The ground state and thermal equilibrium properties of the ferromagnetic transverse-field Ising Model can be efficiently computed classically [Bra2015,Bra2017].
- Quantum Monte Carlo (QMC) methods provide an efficient mapping (in terms of space) to a classical system to study thermal states of stoquastic Hamiltonians. QMC is also efficient in time for 1d Stoquastic Hamiltonians [Cro2021a].

- QMC is known to face obstructions that the adiabatic algorithm does not [Has2013,Jar2016].
- Recent result shows that adiabatic evolutions with stoquastic Hamiltonians can exhibit super-polynomial quantum speedups over all classical algorithms [Has2020,Gil2020]. Whether this translates to speedups for real-world problems remains an open question.

Non-Stoquastic Quantum Adiabatic Optimization

So why not use a non-stoquastic Hamiltonian, ex.

$$H(s) = -(1-s)\sum_i \sigma_i^x + s(1-s)\sum_{\langle i,j\rangle} \sigma_i^x \sigma_j^x + sH_{\text{Ising}}$$

- In the generic case, the non-stoquastic Hamiltonian has a smaller gap that its stoquastic counterpart [Cro2014,Hor2017,Cro2020].

- Only a handful of problems are known to exhibit a gap enhancement for the non-stoquastic case [Sek2012,Dur2019,Alb2019]. Majority of examples are based on highly symmetric problems, such as the *p*-spin model:

$$H_{\text{Ising}} \to -n\left(\frac{1}{n}\sum_i \sigma_i^z\right)^p$$

which exhibit a change from exponentially closing to polynomially closing minimum gap, but are classically trivial problems to solve.

It remains an open question under what conditions a non-stoquastic Hamiltonian can help improve the efficiency of quantum adiabatic optimization.

Beyond Adiabatic

Recall that the adiabatic condition is only a sufficient condition on reaching the ground state with high probability

We can still use an interpolating Hamiltonian but not restrict the evolution to be adiabatic.

Why should we think this could be advantageous?

- Adiabatic evolution in excited states of stoquastic Hamiltonians is universal [Jor2010]
- A very convincing example: polynomial time algorithm for solving the glued trees problem [Som2012]. An exponential speedup over classical algorithms.

An overview of the promise of this approach, dubbed diabatic quantum annealing (DQA), can be found in [Cro2021b].

Example: Glued Trees Problem

Glued Trees Problem [Chi2003]: Two binary trees “glued” randomly. All vertices except for the “ENTRANCE” and “EXIT” vertices have degree three. Each vertex is labeled by a random $2n$-bit string. A total of $N = 2^{n+2} - 2$ vertices.

Computational task: Starting from the “ENTRANCE” vertex, find the “EXIT” vertex. You can only query the adjacency matrix of the graph A; for a given vertex, the query tells you the three/two connected vertices.

$j = 0$ $j = 1$ $\cdots$ $\cdots$ $j = n$ $j = n+1 \cdots$ $\cdots$ $j = 2n+1$

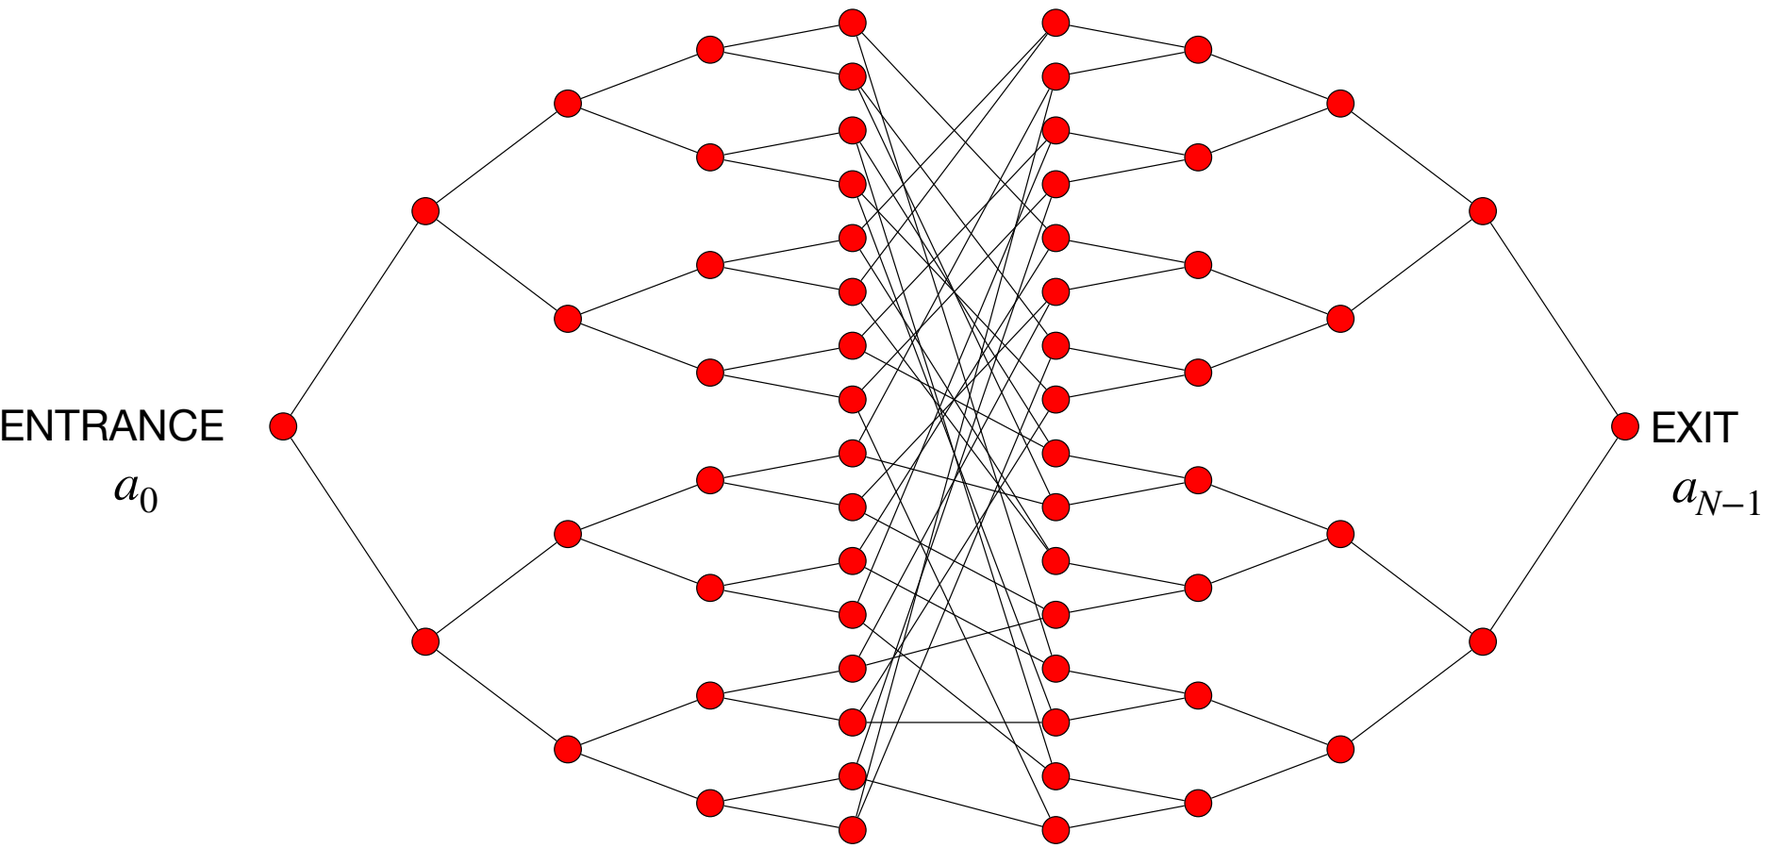

Random labelling of vertices means classical algorithms “get lost” and take exponential time in the size n to find the EXIT vertex.

Example: Glued Trees Problem

Initial Hamiltonian has “ENTRANCE” vertex as ground state; Final Hamiltonian has “EXIT” vertex as ground state

$$H_0 = -|a_0\rangle\langle a_0| \qquad H_1 = -|a_{N-1}\rangle\langle a_{N-1}|$$

Interpolation Hamiltonian uses adjacency matrix A of the graph to *drive* transitions [Som2012]

$$H(s) = \alpha(1-s)H_0 - s(1-s)A + \alpha s H_1$$

(α is a constant that does not depend on n)

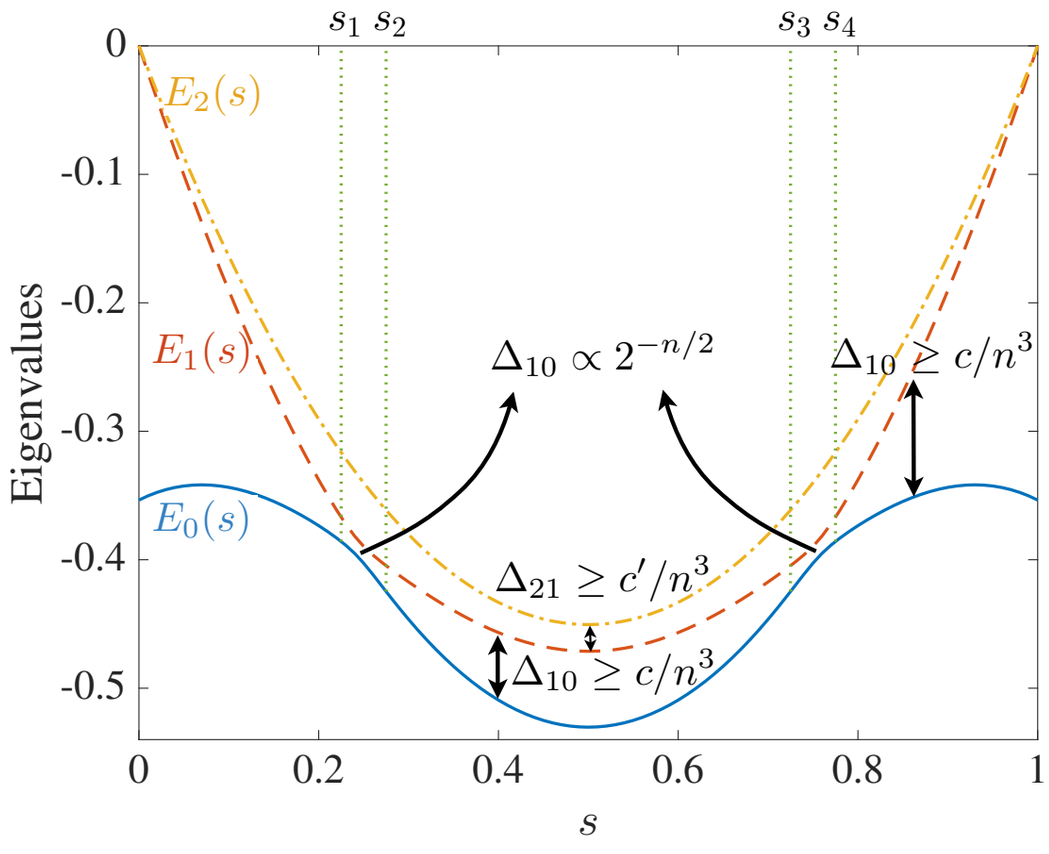

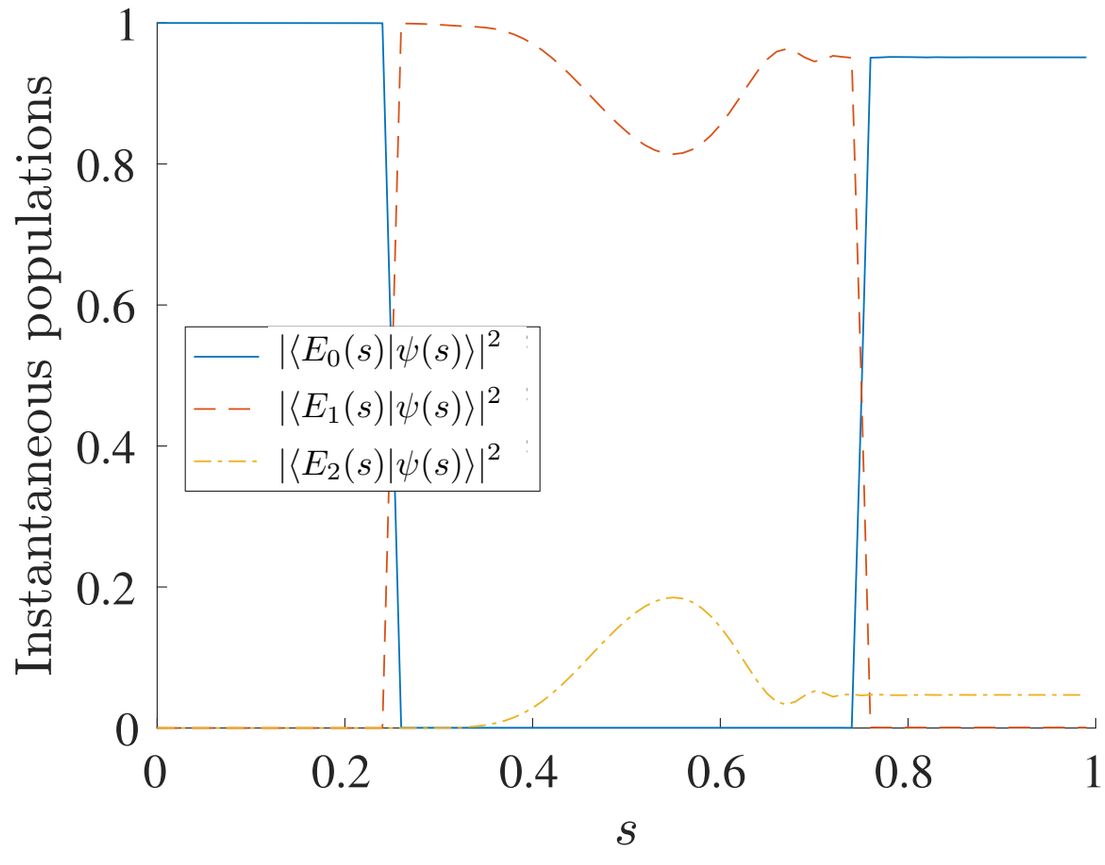

$t_f \sim \mathcal{O}(\text{poly}(n))$ finds EXIT vertex with high probability. Example of exponential speedup over classical algorithms.

Open Quantum System

Our discussion so far has assumed the evolution of the quantum system is described by Schrödinger dynamics → closed system dynamics

In the presence of an environment (ex. another quantum system) with which the quantum system interacts, the effective dynamics of our quantum system is no longer given by Schrödinger dynamics → open system dynamics [Bre2002,Lid2019]

The exact dynamics depends on the type of environment and interaction. Here we will consider a model that is perhaps the most innocuous for adiabatic evolutions: the weak-coupling Markovian adiabatic master equation (AME) [Dav1978,Alb2012]

Express density matrix in instantaneous energy eigenbasis:

$$\rho(t) = \sum_{i,j} \rho_{ij}(t) \, | E_i(t) \rangle \langle E_j(t) |$$

Key properties of AME for continuous-time evolutions [Alb2015b]:

- Dephasing in instantaneous energy eigenbasis: off-diagonal elements $\rho_{ij}(t)$ decay exponentially
- Dynamics pushes system density matrix towards instantaneous Gibbs state

$$\lim_{t_f \to \infty} \rho(t) \to \frac{e^{-\beta H(t)}}{Z}$$

Open Quantum Effects

Dephasing in instantaneous energy eigenbasis: off-diagonal elements $\rho_{ij}(t)$ decay

In the adiabatic setting, system follows instantaneous energy eigenstate with some overall phase. Phase relationship between instantaneous energy eigenstates not relevant, so dephasing is innocuous.

In the non-adiabatic setting, coherent transitions between instantaneous energy eigenstates requires coherent phase relationship, so dephasing is harmful.

Dynamics pushes system density matrix towards instantaneous Gibbs state
(we will call this thermalization)

If thermalization time scale is fast (faster than the algorithm time t_f), population in the ground state depends on Hamiltonian spectrum and temperature.

Open Quantum Effects

Initial state is the ground state at $H(0)$. Assume the time scale t_f would result in an adiabatic evolution in the closed-system case.

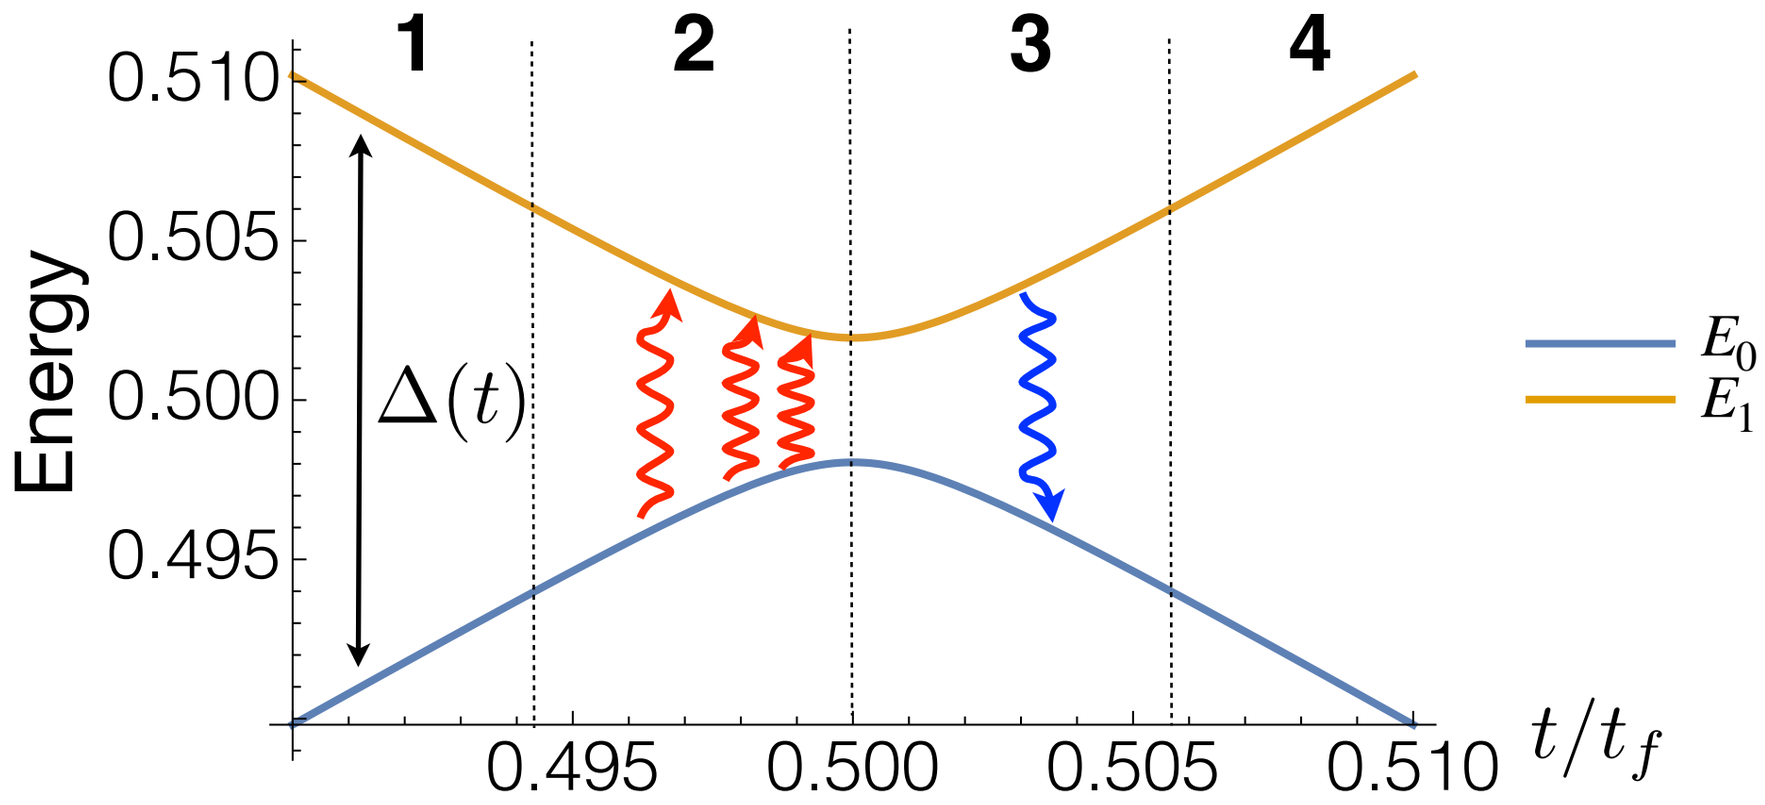

1

$\Delta(t)/k_BT > 1$

No significant change in the state

2

$\Delta(t)/k_BT \lesssim 1$

Population loss to the first excited state via thermal excitations

3

$\Delta(t)/k_BT \lesssim 1$

Population recovery from the first excited state via thermal relaxation

4

$\Delta(t)/k_BT > 1$

Dynamics slowed down (depends on the interaction)

Quantum Annealing

Currently used as a generic term that encompasses continuous-time evolution algorithms with an interpolating Hamiltonian for finding the ground state of some target classical Hamiltonian, either adiabatically or not, in the closed-system setting or not.

How do we quantify the time cost of such algorithms [Røn2014]?

Each anneal ends with a computational basis measurement and returns a bit-string x with probability $p_x(t_f) = |\langle x|\Psi(t_f)\rangle|^2$. Let $p_S(t_f)$ denote the probability of success (success can be defined as finding the ground state or finding a solution below a given target energy) for an anneal with total time t_f.

Time-to-solution (TTS) metric: Number of independent anneals/repetitions $R(t_f)$ of the algorithm required to succeed at least once with 99% confidence, times the time of a single anneal t_f

$$\mathrm{TTS}(t_f) = t_f R(t_f) = t_f \frac{\log(1-0.99)}{\log(1-p_S(t_f))}$$

Gives a way to compare computational cost for a wide range of algorithms

Scaling Analysis

$$\mathrm{TTS}(t_f) = t_f R(t_f) = t_f \frac{\log(1-0.99)}{\log(1-p_S(t_f))}$$

For a given instance from a given problem class, different t_f's give a different TTS.

How to then choose t_f ?

For each t_f , calculate quantiles of the TTS for the sample of instances from the problem class at a given size L. For each percentile, pick the t_f value that minimizes the TTS for that problem size.

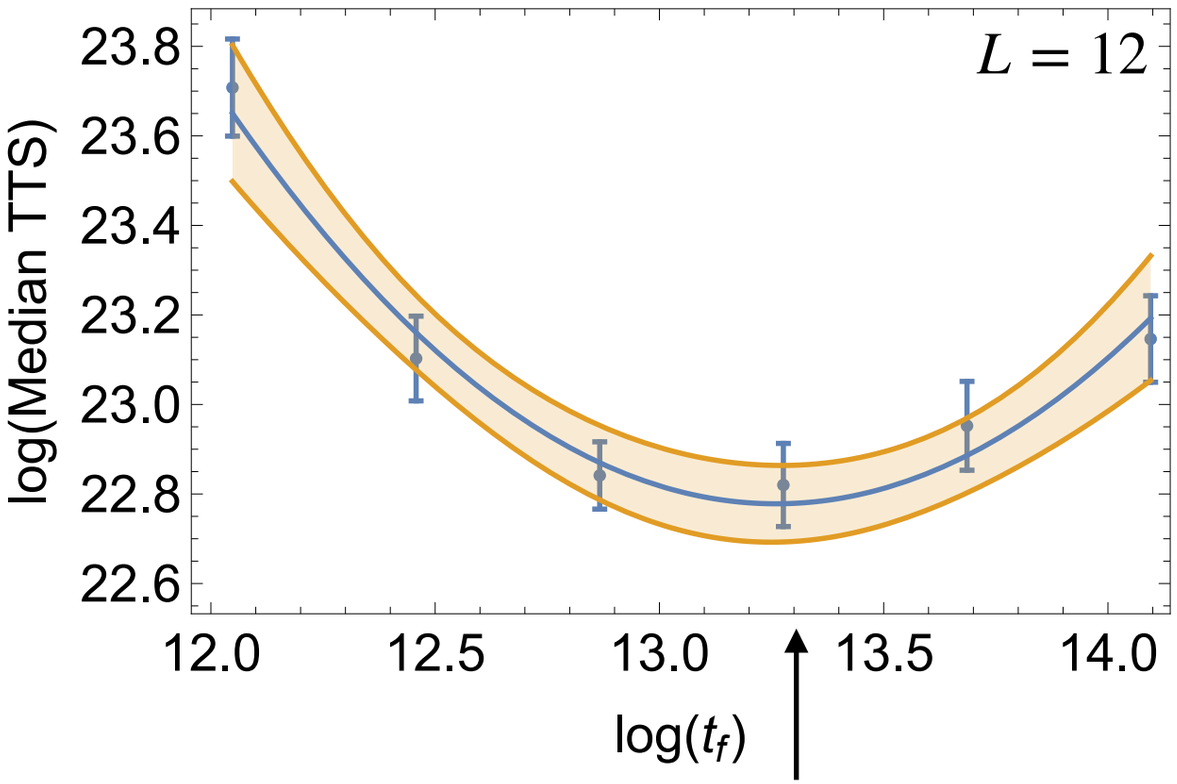

t_f value that minimizes TTS

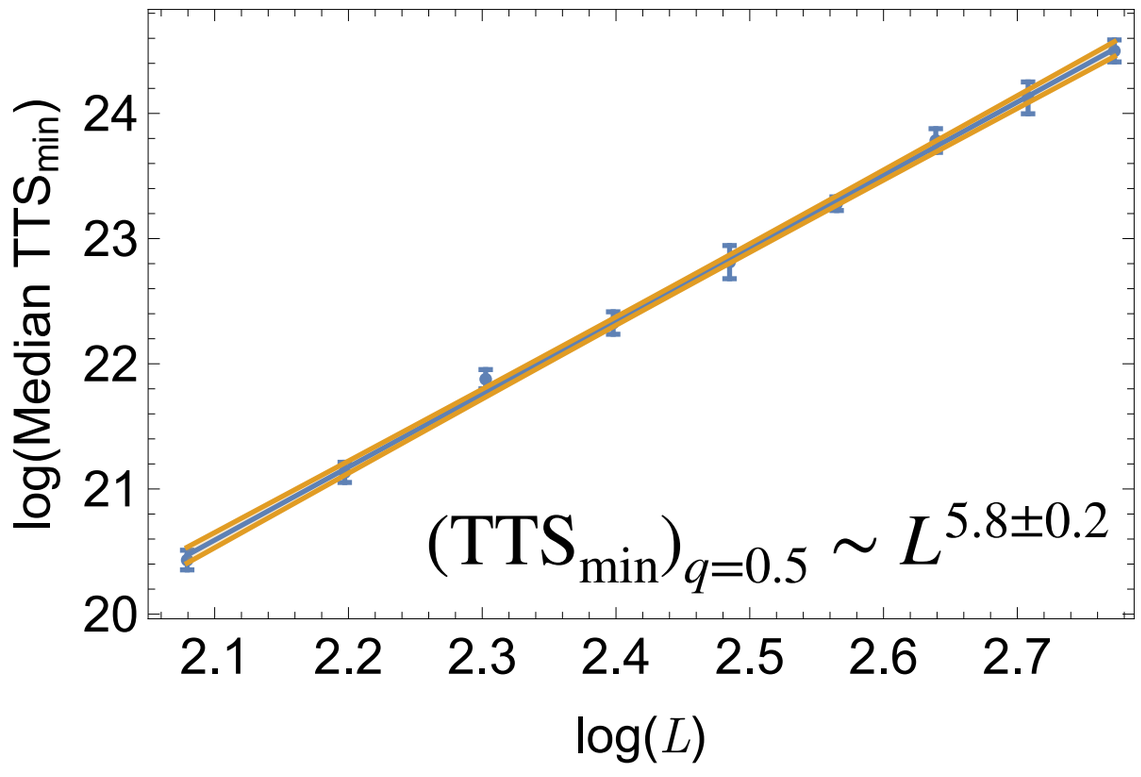

D-Wave Quantum Annealer

The first physical realization of quantum annealing using superconducting circuits on large scales [Joh2011]. Low energy spectrum realizes transverse field Ising model

$$H(s) = -A(s)\sum_i \sigma_i^x + B(s)H_{\text{Ising}}$$

$$H_{\text{Ising}} = \sum_{i\in\mathscr{V}} h_i\sigma_i^z + \sum_{\langle i,j\rangle\in\mathscr{E}} J_{ij}\sigma_i^z\sigma_j^z$$

with the Ising Hamiltonian limited to interactions on a specific connectivity graph $\mathscr{G} = (\mathscr{V}, \mathscr{E})$ with vertices $\mathscr{V}$ and edges $\mathscr{E}$. Until recently, the graph of the D-Wave devices was a Chimera graph [Cho2011], and its current generation device is based on a Pegasus graph [Boo2020] .

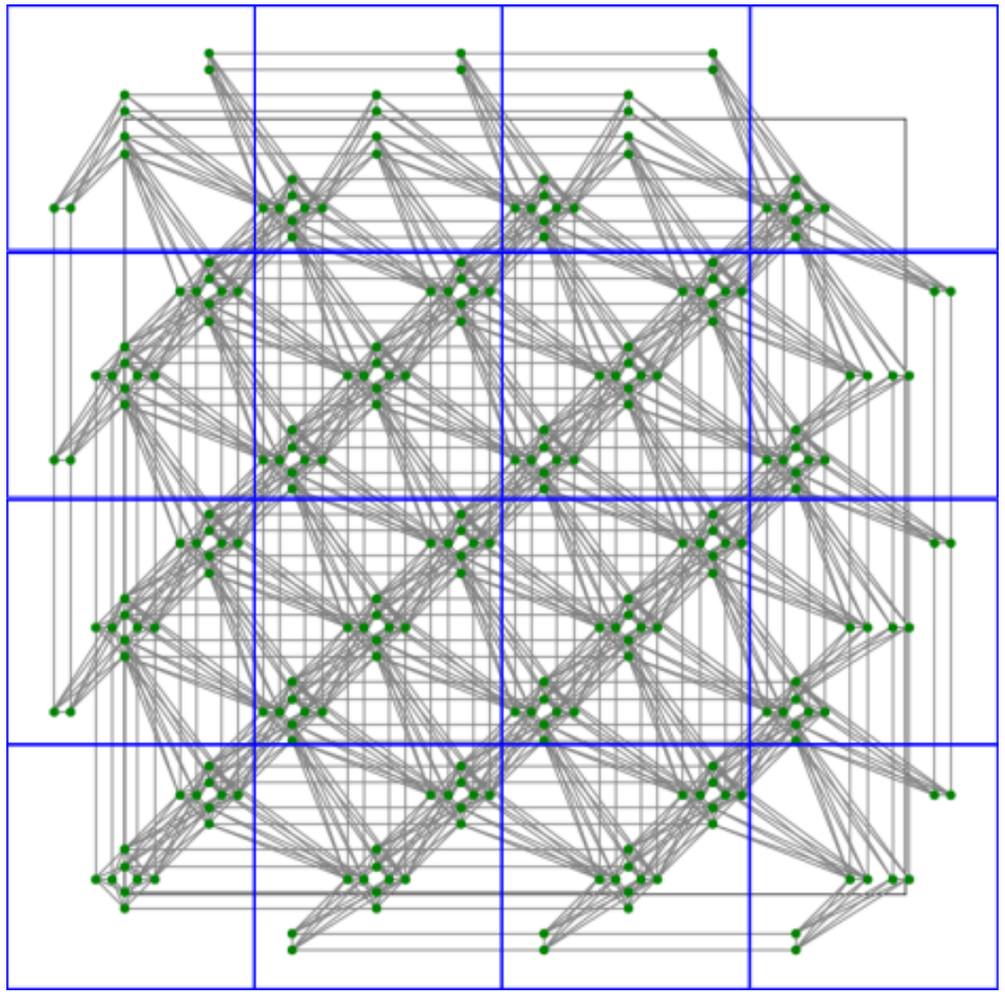

Short D-Wave History

Thermal and noise effects have dominated the output statistics of the device.

- At small system sizes, devices behaves comparably to AME [Alb2015a,Alb2015c] with noisy Ising parameters [Vuf2020,Nel2022], but actually AME description breaks down near small gaps and more accurate models are needed to quantitatively match D-Wave output [Boi2016].
- At larger sizes and long times, output is well captured by “simulated quantum annealing” (SQA) [Boi2014], a simulation technique using a sequence of (path integral) QMC simulation replicating the interpolating Hamiltonian $H(s)$ [San2002,Mar2002].
- At larger sizes, output is also well captured by the semi-classical counterpart of SQA, often referred to as Spin-vector Monte Carlo (SVMC) [Shi2014]. This is primarily because the DW operating temperature is too high, so for a large number of realization of $H(s)$ there is little difference between the behavior of SVMC and SQA. This is particularly problematic because SVMC is a model that does not include important quantum effects such as tunneling and entanglement.
- Not all experiments on the D-Wave processor can be reproduced by SQA and SVMC [Lan2014,Alb2015a].

To zeroth order, one can think of the dynamics as simply thermalizing in the instantaneous Hamiltonian $H(s)$ with the energy scales set by $A(s), B(s)$ and the physical temperature of the device.

Thermalizing Quantum Annealers for Optimization

To date (June 2022), no quantum advantage observed for optimization problems on D-Wave hardware (as far as I am aware).

Cautionary points in the context of optimization (finding minima of classical cost functions):

- A Gibbs/thermal state at a fixed temperature will have exponentially decreasing overlap with the ground state with increasing system size. Algorithm is doomed to fail at growing system sizes unless the temperature is decreased accordingly [Alb2017].
- Quantum annealers are inherently analog devices. Implementation errors in the final Hamiltonian (for example in the Ising parameters $\{h_i \to h_i' = h_i + \delta h_i\}$, $\left\{J_{ij} \to J_{ij}' = J_{ij} + \delta J_{ij}\right\}$) may mean that the implemented Hamiltonian H_1' has a different ground state than the desired Hamiltonian H_1. These errors must also be scaled down as system sizes grow in order to ensure the "right" ground state is reached [You2013,Alb2019b].

Beating state-of-the-art classical algorithms for optimization is not easy! Whether analog hardware can be made useful for classical optimization (real world problems always have a finite size!) in-spite of these limitation is not clear and should be investigated.

A universal lesson though: Identifying and characterizing the noise is critical in order to predict its effect on performance, and we should continue to do it for these devices [Vuf2020,Nel2021,Nel2022] .

QAOA vs QA

The Quantum Approximate Optimization Algorithm (QAOA) [Far2014] is an optimization algorithm for the circuit model.

$$|\psi_p\rangle = \prod_{k=1}^{p} e^{-i\beta_k \sum_j \sigma_j^x} e^{-i\gamma_k H_{\text{Ising}}} |\psi_0\rangle$$

Empirical/Experimental studies of QAOA [Pag2020,Zho2020] suggest that the optimal QAOA angles approximate a smooth curve

→ QAOA at fixed large p approximates a continuous-time (diabatic) evolution.

This connection was made more concrete in Ref. [Bra2021], where it was shown that the true optimal protocol is neither QAOA nor continuous-time evolution but a hybrid of the two (bang-anneal-bang).

This at least suggests that from the point of view of a heuristic algorithms (note that QAOA need not be used as a heuristic algorithm), QA may be a complementary approach to QAOA to pursue. However, learning the optimal annealing schedule may not be as easy as learning the optimal QAOA angles.

QAOA vs QA: Experimental Demonstrations

Neutral atom quantum simulators with up to 289 qubits on a square grid solving Maximum Independent Set (MIS) problems [Eba2022]

$$H_{\mathrm{MIS}} = -\sum_i n_i + \alpha \sum_{\langle i,j \rangle} n_i n_j \ , \quad \alpha > 0$$

$$n_i = |1\rangle_i\langle 1|$$

Their quantum annealing algorithm outperformed their QAOA algorithm

~ 70 μs coherence time, ability to optimize annealing schedule to a certain extent

Platform has limited programmability of individual 2-qubit couplings, so limited to problems like MIS with uniform antiferromagnetic couplings

No comparison to state of the art spin glass algorithms like Parallel Tempering with Isoenergetic cluster moves (PT-ICM) [Zhu2015]

Using Quantum Annealers as Quantum Simulators

Quantum annealers can be used more generally to prepare interesting states

• Thermal phases: On D-Wave hardware, there have been studies of a (quantum) Kosterlitz–Thouless phase transition [Kin2018], (quantum) 3d cubic lattice spin glass phase transition [Har2018], (classical) Shastry-Sutherland Ising Model [Kai2020], (classical) spin ice [Kin2021a].

• Thermalization: Recent results suggest that D-Wave device may be able to thermalize faster than QMC simulations when starting from specific initial states with non-trivial winding [Kin2021b]

•Open-System Kibble-Zurek mechanism: On D-Wave hardware at annealing times between 1 μs and 2ms, consistent results with open-system Kibble-Zurek predictions on the transverse field Ising model with up to 800 qubits [Ban2020]. Results actually consistent with spin-Langevin equation [Sub2022].

• Closed-System Kibble-Zurek mechanism: On D-Wave hardware at very short annealing times ($\lesssim 100ns$), verification of Kibble-Zurek predictions on the transverse field Ising model with up to 2000 qubits [Kin2022].

• Phases of antiferromagnetic transverse field Ising models on square and triangular lattices with up to 256 atoms have been studied on neutral atom platforms [Sch2021,Eba2021].

As I understand it, none of these demonstrations are beyond classical simulation

Final Thoughts

Adiabatic quantum computing, while out of fashion these days, is a paradigm of quantum computing with obvious connections to many areas: Hamiltonian complexity, quantum phase transitions, etc.

Quantum annealing, while originally formulated in the adiabatic setting, has become an umbrella term for more general continuous-time algorithms with an interpolating Hamiltonian. It remains a useful generic approach for preparing ground states.

Open question whether the algorithm can provably demonstrate quantum advantages on real-world problems, but also whether quantum annealers at some fixed size can be used in a setting to solve such problems faster than can be done on classical hardware.

New programs in Europe and Japan are being pursued to develop coherent quantum annealers. The US government (DARPA) expects to have a testbed available for research.

Scalable quantum computing will require fault tolerant quantum error correction. Such protocols are not yet known for the adiabatic paradigm, but error mitigation strategies could improve performance at intermediate sizes.

- Passive error suppression [Jor2006,Jia2017,Mar2017,Mar2019]:
- Active error correction [Ata2021]