



Universiteit
Leiden

The Netherlands

The power of one qubit in quantum simulation algorithms

Polla, S.

Citation

Polla, S. (2024, February 22). *The power of one qubit in quantum simulation algorithms*. *Casimir PhD Series*. Retrieved from <https://hdl.handle.net/1887/3719849>

Version: Publisher's Version

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/3719849>

Note: To cite this publication please use the final published version (if applicable).

2.1 Introduction

Ground state preparation is an essential algorithm in the quantum computing toolbox. Any polynomial-time quantum algorithm can be mapped to the problem of estimating the ground state energy of an artificial Hamiltonian given an approximation to its ground state [81], and without such additional input this problem is known to be QMA-hard for even 2-local Hamiltonians [30]. Digital quantum simulation of problems in materials science and chemistry, one of the ‘killer apps’ of a quantum computer, is most often concerned with properties of ground states of the simulated systems [15, 21], and many problems in optimization may be mapped to ground state finding problems [49, 55]. This has led to a wide range of schemes for digital ground state approximation, via adiabatic evolution [49], variational methods [19, 55, 82], phase estimation [83], amplitude amplification [84–86] and approximate imaginary time evolution and other Hamiltonian function techniques [28, 87–89]. However, these algorithms suffer from large computational costs or approximation errors, making designing better schemes an active area of interest.

In nature, ground states are achieved by coupling to a large cold reservoir, which takes energy from the system in keeping with the second law of thermodynamics. Simulating an entire bath would require an impractically

large quantum register, however it has long been suggested that this may be mimicked by coupling to a single qubit which may be reset to its ground state with sufficient frequency [21]. This idea has been since studied in digital quantum computing for the initialization of quantum devices [90, 91] and as an inspiration of an algorithm based on resonant transitions and postselection [92]. This idea was also explored in analog simulation settings, for the preparation of physical [93] and artificial [94, 95] ground states. However, cooling an artificial system in the digital quantum setting provides a set of unique challenges — the system being studied may be completely different from the physical quantum hardware, and the digitized Hamiltonian may be only an approximation to the target of interest. Furthermore, the periodic non-unitary reset may break the unitary evolution in short time-scale chunks which do not conserve energy, implying that one may artificially reheat the system without clever protocol design. This is of critical importance in near-term devices, where limited coherence times compete against the desire for slower cooling cycles.

In this chapter, we detail how one may prepare ground states of an artificial Hamiltonian on a digital quantum computer via quantum digital cooling (QDC). We first present an analytic study of the cooling of a two-level system, from which two different approaches may be outlined to de-excite to the ground state whilst preventing reheating. We investigate the behaviour of both methods in the digitized setting, and find they continue to be robust. The protocols deriving from these two principles are tested in the one-qubit black-box Hamiltonian setting, where the energy gap and matrix elements are unknown. We extend these protocols to N -qubit systems, and investigate their ability to cool small-scale simulations of the transverse-field Ising model numerically. Our LogSweep protocol, based on the weak-coupling approach, is demonstrated to converge to the ground state of all three phases of the transverse-field Ising model. It further shows a relative energy error constant in the system size at a fixed protocol length for the weakly-coupled and critical phases of this model, which corresponds to a requirement to simulate time evolution for $O(N^2)$ and $O(N^3)$ Trotter steps respectively. By contrast, the strong-coupling BangBang protocol shows the ability to prepare low-cost ground-state approximations of the same model in the paramagnetic and ferromagnetic regime, but seems to perform much worse close to the critical point, where the system spectrum shows a less-ordered structure. The small number of calls to the system evolution operator needed to realize this protocol makes it attractive for near term application.

2.2 Cooling a system with a single fridge qubit

In nature, gapped physical systems cool down to a state with high overlap to the ground state when interacting with a bath that is cold and large, under the condition of ergodicity. By *cold*, we mean that temperature T_B of the bath is small compared to the ground state gap Δ_S of the system to be cooled: $k_B T_B \ll \Delta_S$ (with k_B Boltzmann's constant). By *large*, we mean that the bath has a sufficiently large Hilbert space that the above condition is unchanged by the addition of the energy from the system. By *ergodic*, we mean the system must not have symmetries that prevent excitations to be transferred from the system to the bath, or that reduce the effective size of the accessible bath Hilbert space. Given a system with Hamiltonian H_S and eigenstates $H_S|E_j\rangle = E_j|E_j\rangle$, energy conservation implies that the bath must have states at energies $E_j - E_0$ to allow de-excitation of the eigenstates E_j . This is typically achieved by considering a bath with a continuous or near-continuous low-energy spectrum [Fig. 2.1(a)]. However, the bath need not cool all states immediately to the ground state. Instead, a bath typically absorbs single quanta of energy $\epsilon = E_i - E_f$ that correspond to local excitations of the system $|E_i\rangle \rightarrow |E_f\rangle$, at a rate given by Fermi's golden rule:

$$\frac{dP_{i \rightarrow f}}{dt} = \frac{2}{\hbar} \int_0^\infty d\epsilon |\langle E_f, \epsilon | H_C | E_i, 0 \rangle|^2 \rho_B(\epsilon) \times \lim_{t \rightarrow \infty} \frac{\sin[(E_i - E_f - \epsilon)t]}{E_i - E_f - \epsilon} \quad (2.1)$$

$$= \frac{2\pi}{\hbar} |\langle E_f, \epsilon | H_C | E_i, 0 \rangle|^2 \rho_B(E_i - E_f), \quad (2.2)$$

where H_C is the coupling between the system and the bath, and ρ_B is the density of states of the bath¹. This requires the bath to be large enough to prevent re-excitation of these states as the system continues cooling. In other words, the bath must have a large Hilbert space compared to the one of the system. This ensures that, at equilibrium, most of the entropy is distributed in the bath.

To represent such a large bath with an auxiliary register on a quantum device in order to cool a system register would be impractically costly. We propose approximating the presence of a much larger bath B with a single

¹In the rest of this chapter we drop \hbar factors, assuming a choice of units for which $\hbar = 1$.

auxiliary qubit F [Fig. 2.1(b)], with bath Hamiltonian $H_B \rightarrow H_F = \epsilon Z_F/2$. This can be considered a simplified form of a collisional model [96] that does not allow for non-Markovian effects (that would be, in our case, unwanted). The coupling between the bath and the system takes the form $H_C = \gamma X_F \otimes V_S/2$, where γ is the coupling strength, and V_S a *coupling term* that acts on the system alone. A key advantage of the digital approach is that we are free to choose V_S as desired to optimize the cooling protocols. The Hamiltonian of the entire system and bath then takes the form

$$H = H_S + H_F + H_C. \quad (2.3)$$

This has the immediate problem that the bath can only absorb a single quantum of energy ϵ ; however, we may circumvent this by periodically resetting the auxiliary qubit to $|0\rangle$. The non-unitary reset effectively extracts energy and entropy from the auxiliary qubit to a much larger external bath (i.e. the experimenter's environment). For this reason, we call this auxiliary qubit a *fridge qubit* (hence F). The non-unitarity introduced in the process is necessary to dissipate entropy, allowing to prepare the ground state from an arbitrary starting state. As the time between resets is finite, the $t \rightarrow \infty$ limit in Eq. (2.1) is no longer justified and energy is no longer conserved. This is both a blessing and a curse — we need not precisely guess the energy gap $\Delta = E_i - E_f$ of the transition that we need to de-excite, but we run the risk of reheating the system at each cooling round. Minimizing re-heating while maximizing the range of targeted de-excitations is key to the successful design of QDC protocols. In a realistic experiment, qubit re-heating would be effectively increased by reset infidelity on the auxiliary qubit, making the protocol less effective.

2.3 De-exciting a single transition: the 1+1 model

In order to design some basic protocols for QDC, we turn to a toy ‘1+1’ model. We take a single-qubit system with Hamiltonian $H_S = \Delta Z_S/2$, and couple it to a single fridge qubit with coupling term $V_S = X_S$. Although this model is simple, it can represent a pair of levels being targeted for cooling in a larger quantum system. We will make use of this interpretation when extending these cooling protocols in section 2.4.




Figure 2.1: The de-excitation of the system transition $|E_1\rangle_S \rightarrow |E_0\rangle_S$ mediated by: (a) a continuous-spectrum natural bath B, where an excitation $|\epsilon\rangle_B$ at energy ϵ is produced, and (b) a single-qubit digital fridge F, which can be excited if $\epsilon = \Delta$.

2.3.1 Elementary approaches to digital cooling: strong and weak coupling

Let us first assume Δ is known, in which case resonant cooling ($\epsilon = \Delta$) is the most effective choice of ϵ . With this fixed, the transition probabilities after time t may be calculated exactly to be

$$P_{1 \rightarrow 0} = \sin^2\left(\frac{\gamma}{2}t\right), \quad P_{0 \rightarrow 1} = \frac{\gamma^2 \sin^2(t\Omega)}{4\Omega^2}, \quad (2.4)$$

where $\Omega = \sqrt{\gamma^2/4 + \epsilon^2}$. We wish to maximise the cooling probability $P_{1 \rightarrow 0}$ while minimizing the reheating probability $P_{0 \rightarrow 1}$ by optimizing the remaining free parameters: the coupling strength γ and the cooling time t . To maximize the cooling rate $P_{1 \rightarrow 0} = 1$, we must set

$$t = \pi\gamma^{-1}. \quad (2.5)$$

We assume this constraint throughout this chapter. Note that this choice is beyond the perturbative regime $\gamma t \ll 1$, for which the Fermi golden rule Eq. (2.1) is formulated.

To minimize reheating, we can then choose either of two very different approaches, based on strong or weak coupling. The weak-coupling approach is based on the observation that the off-resonant transition probability $P_{0 \rightarrow 1}$ is bounded by $\gamma^2/4\Omega^2$. As such, we may suppress reheating to

2 Quantum digital cooling

an arbitrary level by choosing sufficiently small γ . The time-cost for Hamiltonian simulation of e^{iHt} scales at best linearly in t [97], so this implies one may obtain the ground state with failure probability p in time $O(p^{-1})$, regardless of the initial state of the qubit. The strong-coupling approach consists of tuning γ so that $\Omega t = \pi$, which is achieved when

$$\gamma = \frac{2}{\sqrt{3}}\epsilon. \quad (2.6)$$

This fixes the reheating exactly to 0, guaranteeing the qubit to be in the ground state in the shortest possible time, but at the cost of requiring fine-tuning.

Unlike in analog quantum simulation, digital devices cannot exactly implement the dynamics of the Hamiltonian in Eq. (2.3), and must approximate it digitally instead. A common approach to such digitization is that of the Suzuki-Trotter expansion [22, 23], which we now explore for the two cooling paradigms. We apply the (second-order) expansion of the coupled system-bath evolution with Trotter number M ,

$$e^{-i(H_S+H_F+H_C)t} \sim \left[e^{-iH_C \frac{t}{2M}} e^{-i(H_S+H_F) \frac{t}{M}} e^{-iH_C \frac{t}{2M}} \right]^M. \quad (2.7)$$

Note that, when we later approach larger systems, we will practically realize $e^{-iH_S t/M}$ as a single second-order Trotter step, effectively implementing a second-order Trotter simulation of the coupled Hamiltonian $H_S + H_F + H_C$ with trotter number M . If we restrict to the subspace containing the states involved in the cooling transition $|10\rangle_{\text{SF}} \rightarrow |01\rangle_{\text{SF}}$, at resonant cooling we have $H_S + H_F \propto \mathbb{1}$ (specifically, in this model $|01\rangle$ and $|10\rangle$ generate a zero-eigenvalue subspace of $H_S + H_F$). Thus, the Trotterized evolution behaves exactly like the continuous one with regards to the cooling transition. We study reheating probabilities as a function of t for different values of M in the weak-coupling regime. We observe (Fig. 2.2) that the digitized evolution approximates well the behavior of the continuum limit whenever $t\Omega/\pi \lesssim M$ (i.e. for the first M Rabi oscillations with pulse Ω). For longer times $t\Omega\pi \gtrsim M$, the second-order Trotter approximation fails, leading to reheating rates far larger than in the continuum limit. This allows us to define a practical choice of M to avoid reheating due to the Trotter approximation – we require

$$M \gg \sqrt{1 + \epsilon^2/\gamma^2}, \quad (2.8)$$

which sets the working point $t = \pi\gamma^{-1}$ before the $M/2$ Rabi oscillation.




Figure 2.2: Effects of Trotterization on cooling and reheating probabilities as a function of the coupling time t , for different numbers of Trotter steps M per cooling cycle. Vertical dotted lines indicate the M -th reheating oscillation, at which point the Trotter approximation fails.

In the strong-coupling case $t\Omega/\pi = \sqrt{3}$, which implies that a single step is sufficient to satisfy the condition above. Indeed, digitized cooling with probability 1 and no reheating can be realized by a *bang-bang* approach (inspired by similar approach in variational methods [98, 99]). This consists in defining the evolution as in Eq. (2.7) with $M = 1$, while adjusting the coupling strength to $\gamma = 2\epsilon$. With this choice, the digitized evolution implements resonant Ramsey interference on the cooling transition $|10\rangle_{\text{SF}} \rightarrow |01\rangle_{\text{SF}}$ and anti-resonant Ramsey interference on the reheating transition $|00\rangle_{\text{SF}} \rightarrow |11\rangle_{\text{SF}}$.

A key difference between the two approaches to digital cooling is in their behavior off-resonance, i.e. when the energy gap is mis-targeted or not precisely known. For the bang-bang approach, detuning reduces the cooling efficiency while symmetrically boosting reheating [Fig. 2.3(a)]. The wide resonance peak around zero detuning makes this approach ideal to quickly cool transitions which energy is known up to a small error. In the weak-coupling approach the resonance peak becomes sharper and the reheating gets more suppressed as the coupling is made smaller [Fig. 2.3(b)], approaching the energy conservation limit. Detuning makes cooling inefficient, but thanks to the low reheating probability this weak-coupling cooling can be iterated with different values of ϵ to try to match the transition energy, without destroying the cooling effect.




Figure 2.3: Effect of fridge-system detuning $\delta = \Delta - \epsilon$ on the cooling (dashed lines) and reheating (solid lines) probabilities for (a) the bang-bang cooling approach, and (b) the weak-coupling cooling approach, where colors indicate different coupling strengths.

2.3.2 Common symmetries and the coupling alternation method

In large systems of interest, we do not know the Hamiltonian's eigenstates, making it more challenging to couple between them. This is required for cooling, as can be seen by the direct dependence of the cooling rate $\frac{dP_{i \rightarrow f}}{dt}$ on the overlap $|\langle E_f, \epsilon | H_C | E_i, 0 \rangle|^2$ in Eq. 2.2. This overlap dependence will prohibit cooling if the system Hamiltonian H_S and the coupling operator V_S share a common symmetry S (i.e. $[S, H_S] = [S, V_S] = 0$). When this is the case, the Hamiltonian may be simultaneously diagonalized with H_S , and a state with some overlap to any eigenspace of S that does not contain the ground state cannot be cooled to the ground state by coupling with V_S . Note that this is a strictly stronger condition than just requiring $[H_S, V_S] \neq 0$. A simple solution is to alternate over a set of couplings $\{V_S^i\}$ as we cool. Then, a symmetry S of H would need to commute with each V_S^i in order to guarantee that a state starting with overlap in a high-energy eigenspace will remain there. Sets of coupling terms $\{V_S^i\}$ on N qubits need only be size $O(N)$ to break commutation with all non-trivial operators (for example, the set of all single-qubit Pauli operators), so although symmetries need to be taken into account, they will not destroy the feasibility of QDC protocols.

This issue may be demonstrated on the prototype 1 + 1 qubit model by considering the system Hamiltonian $H_S = h \vec{n} \cdot \vec{\sigma}$, where \vec{n} is a 3-dimensional unit vector (so H_S points in a random direction on the Bloch

2.3 De-exciting a single transition: the 1+1 model

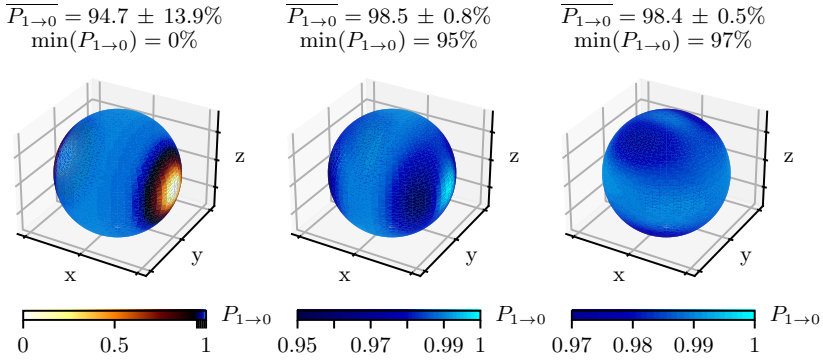


Figure 2.4: Probabilities $P_{1 \rightarrow 0}$ of transitioning from $|1\rangle$ to $|0\rangle$ after three iterations of the weak-coupling ($t\epsilon = 10$) cooling procedure, with coupling potentials $V_S^i = X, X, X$ (left), $V_S^i = X, Y, X$ (center), $V_S^i = X, Y, Z$ (right), on a system qubit with Hamiltonian $H_S = h \vec{n} \cdot \vec{\sigma}$ and known energy splitting h . The orientation of the unit vector \vec{n} is represented on spherical surfaces. The average, standard deviation and minimum of $P_{1 \rightarrow 0}$ are shown above each panel.

sphere), $2h$ is a fixed energy splitting, and σ is the vector of Pauli-matrices. For a fixed coupling operator V_S , there is a risk that $[H_S, V_S] \approx 0$. When this is the case, the off-diagonal elements of V_S in the system eigenbasis are zero, preventing cooling. We may prevent this by alternating between different coupling terms during the cooling protocol, such that no non-trivial Hamiltonian can commute with all such coupling terms. This may be achieved for the 1 + 1 model by iterating over $V_S^i \in \{X_S, Y_S, Z_S\}$, or alternatively over $V_S^i \in \{X_S, Z_S\}$. The effectiveness of this scheme is studied in Fig. 2.4 for resonant coupling. We see the probability $P_{1 \rightarrow 0}$ of successful cooling of the weak coupling approach ($t\epsilon = 10$) increased to $\min(P_{1 \rightarrow 0}) = 97\%$ for all choices of \vec{n} when iterating over $V_S^i = X_S, Y_S, Z_S$, and above 95% when iterating $V_S^i = X_S, Z_S, X_S$, compared to the possibility for complete cooling failure [$\min(P_{1 \rightarrow 0}) = 0$] when V_S^i is held constant. Similar results are observed for off-resonant coupling.

2.4 Scalable QDC protocols

We now look to use the insight obtained for cooling in the 1+1 toy model to develop QDC schemes for larger systems. The sub-additivity of entropy places a rough lower bound on the number of cooling steps required to cool an N -qubit system. This limits the entropy ΔS_S that the system can transfer to the fridge qubit before the non-unitary reset to $\Delta S_S \geq -\Delta S_B \geq -1$ bit. If we demand the ability to cool an arbitrary state, a QDC protocol must also be able to cool the maximally-mixed state, which has entropy $S_S = N$. We then require N repetitions of an optimal coupling-and-reset step to reach the pure ground state (which has entropy $S_S = 0$). This can be obtained in the simple example of cooling N non-interacting qubits with known energies, by simply repeating the protocols of the 1 + 1 model. However, this cannot be generalised to arbitrary strongly-correlated systems, as cooling is complicated by irregular and unknown energy gaps and coupling terms between eigenstates. This is to be expected, as preparing ground states of arbitrary Hamiltonians is a known QMA-hard problem [30]. However, as cooling is a physically-motivated process, we hope QDC to be able to achieve polynomial scalings for systems of physical interest, i.e. models of systems that are found in low-temperature equilibrium states in nature. We focus for the rest of this chapter on exploring this thesis.

In the rest of this text, we introduce two scalable QDC protocols for N -qubit systems: the strong-coupling-based BangBang protocol and the weak-coupling-based LogSweep protocol. These extend and generalize the two approaches we established for the 1+1 toy model of section 2.3.1. Each protocol iterates over a sequence of *cooling steps*, each of which consists of coupling the fridge qubit to part of the system for a short time evolution, and then resetting the fridge qubit to its ground state. The protocols differ in the choice of coupling strengths γ_i , coupling terms V_S^i and fridge energies ϵ_i at each i -th cooling step. [The coupling time for each cooling step is fixed by Eq. (2.5)].

2.4.1 The BangBang protocol

We now develop a protocol to extend the strong-coupling approach (Eq. 2.6) to a larger system. This motivation is in line with recently proposed bang-bang approaches to adiabatic state preparation [98, 99], which are known to outperform initial theoretical expectations stemming from a naive Trotter error estimate. We are thus optimistic that this 'BangBang' protocol may provide a low-cost, near-term method for QDC. However,

such a protocol needs to associate each fridge-system coupling with a single fridge energy, that should match the transitions that this coupling triggers. An appropriate choice of fridge energy can be estimated via a perturbation theory approximation. To derive this approximation, we note that the rate of a transition between eigenstates $|E_i\rangle \rightarrow |E_j\rangle$ depends on the matrix element of the coupling V_S :

$$V_{(ij)} := \langle E_i | V_S | E_j \rangle = \frac{\langle E_i | [H_S, V_S] | E_j \rangle}{E_i - E_j}. \quad (2.9)$$

If V_S is local and bounded, $[H_S, V_S]$ is as well, so the matrix element $V_{(ij)}$ will be bounded proportionally to $(E_i - E_j)^{-1}$. The matrix element is additionally bounded by $\|V\|$; this second bound becomes dominant when $E_i - E_j/\|V\|$ falls below the maximum off-diagonal element of $[H, V]$ in any basis, which we define with the notation $\|[H, V]\|_{\perp}$:

$$\|O\|_{\perp} = \max_{\langle \phi | \psi \rangle = 0} |\langle \phi | O | \psi \rangle| = \max_{|\Phi\rangle, |\Psi\rangle} \frac{\langle \Phi | O | \Phi \rangle - \langle \Psi | O | \Psi \rangle}{2}, \quad (2.10)$$

where O is Hermitian and the maxima are taken over all possible states $|\psi\rangle, |\phi\rangle$ and $|\Psi\rangle, |\Phi\rangle$. A simple proof is given in Appendix 2.A. We use this energy scale to set the fridge energy:

$$\epsilon_i = \|[V_S^i, H_S]\|_{\perp} \quad (2.11)$$

for any coupling potential V_S^i . This way, the maximum-energy transitions accessible by V_S are on resonance, while smaller energy ones (which are less important for cooling) still have a higher probability of cooling than of reheating [see Fig. 2.3(a)]. This defines the BangBang protocol: we iterate over coupling to each qubit, with ϵ_i fixed by Eq. (2.11). As this protocol does not attempt to suppress reheating, we choose a single coupling $V_S = Y_n$ for the n -th qubit, instead of iterating over $V_S = X_n, Y_n, Z_n$ (as was suggested in Sec. 2.3.2). In general, the best choice of V_S will depend on the system that we want to cool, and the couplings should be picked to enable as many transitions as possible. We repeat the coupling to each qubit R times, resulting in a total of RN cooling steps. Each cooling step contains two first-order Trotter steps simulating $e^{-iH_C t/2}$ (of depth $d_{H_C}^{(1)}$), a single second-order Trotter step for $e^{-iH_S t}$ (of depth $d_{H_S}^{(2)}$), and a reset gate, resulting in a total circuit depth

$$d_{\text{BangBang}} = RN(2d_{H_C}^{(1)} + d_{H_S}^{(2)} + 1). \quad (2.12)$$




Figure 2.5: Change in energy expectation value for the application of a single cooling step to the maximally mixed state of a $N = 8$ qubit transverse field Ising chain Eq. (2.13), depending on the fridge energy ϵ . The coupling potential is $V_S = Y_3$, the Pauli Y on the third qubit. The relation $B^2 + J^2 = 1$ fixes the energy scale.

To test the BangBang protocol, we study the cooling of a N -qubit transverse-field Ising chain

$$H_S = \sum_{i=0}^N BX_i + \sum_{i=0}^{N-1} JZ_iZ_{i+1}, \quad (2.13)$$

where B represents the transverse magnetic field Zeeman splitting and J is the Ising coupling strength. The relative coupling strength J/B dictates whether the system is in the paramagnetic ($J/B \ll 1$), ferromagnetic ($J/B \gg 1$), or critical ($J/B \sim 1$) phases. This ability to simply tune between three phases of matter with significantly different physical properties make the TFIM a good benchmark model to investigate the ability of different QDC schemes in various scenarios.

We first demonstrate that our choice for the fridge energy Eq. (2.11) is appropriate. In Fig. 2.5, we plot the effect of a single cooling step on the maximally-mixed state. We observe that cooling is maximized for fridge energies around the point defined by Eq. (2.11), for all phases of the TFIM. We find this behaviour to hold for all other (local) choices of coupling potential V_S tested in our work, as predicted.

We next turn to the ability of the BangBang protocol to prepare an approximation ρ of the ground state, starting from a maximally-mixed (i.e. infinite temperature) initial state. We benchmark by the final state

with two figures of merit: the ground state fidelity

$$F = \text{Tr} [|E_0\rangle\langle E_0| \rho], \quad (2.14)$$

and the energy relative to the ground state energy $\text{Tr}[H_S \rho]/|E_{\text{GS}}|$. This last property is local in local system, and represents an energy density in TFIM. To verify convergence, we compare cooling results to a *reheating limit*, obtained by running the protocol with the ground state as initial state. We observe that all TFIM phases converge after $R \approx N$ repetitions (with the weakly-coupled phase system converging already at the first repetition). In Fig. 2.6 we plot the energy density of the cooled state, as well as the reheating limit, as a function of the number of sites in the system. This shows that convergence is indeed achieved for $R = N$ independently on the system phase and size, and that the final energy density stays approximately constant, without showing any other trend. The BangBang protocol achieves a final energy density close to 90% and 95% of $\|H_S\|_{\perp}$ for the ferromagnetic and paramagnetic regime respectively, while performing significantly worse in the critical regime. This is to be expected, as in this regime the spectrum is no longer banded, and excitation energies are not as uniform as in the paramagnetic or ferromagnetic regimes. Following Eq. (2.12), the protocol's circuit depth is $7NR$ for a gate-set containing single- and double-qubit rotations (and the reset gate). Given the low cost of the protocol, we suggest that this is of particular interest for near-term experiments, and may be further refined by other cooling protocols, or methods such as quantum phase estimation, in the long term.

2.4.2 The LogSweep protocol

Refrigeration at weak-coupling suppresses reheating, but only allows for the cooling of transitions within a narrow energy band [as shown in Fig. 2.3(a)]. We may take advantage of this in a larger system, where a wide range of energy gaps are present, by sweeping the fridge energy ϵ_k from high to low as we iterate over cooling steps. (As low-energy transitions are more susceptible to re-heating than high-energy transitions, this will in general be more efficient than sweeping from low to high.)

To construct a full protocol, we further need to fix the set of fridge energies ϵ_k and linewidths $\delta_k = t_k^{-1} = \pi \gamma_k$ that we plan to use for each cooling step. We will be guided by two principles. First, the target band of fridge energies $(E_{\text{min}}, E_{\text{max}})$ should be tightly covered by the cooling lines $\epsilon_k \pm \delta_k$. Second, the reheating should be suppressed to a certain degree throughout the protocol. As by Eq. (2.4) the reheating suppression




Figure 2.6: Performance of the BangBang protocol as a function of the system size N for the three different phases of the transverse-field Ising model (detailed in legend). The coupling potentials are $V_S^i = Y_i$. Dots correspond to result when the protocol is applied to the maximally-mixed state, shaded regions corresponds to result when protocol is applied to the true ground state (which gives a bound on protocol re-heating). Data generated by Trotterized wave-function simulations of the protocol, and random sampling of the initial mixed state and of nonunitary operations (details in App. 2.F). All points are run with 200 samples, and average results are plotted with the sampling error.

depends on γ_k/ϵ_k , we fix this value to a small constant throughout the protocol (i.e. we choose $\gamma_k \propto \epsilon_k$). Thus we define the LogSweep protocol, where the fridge energy ϵ_k is swept over (E_{\min}, E_{\max}) in a logarithmic gradation. Specifically, given the *gradation number* K , we index each cooling step $k = 1, \dots, K$, and we define

$$\epsilon_k = E_{\min}^{\frac{k-1}{K-1}} E_{\max}^{1-\frac{k-1}{K-1}}, \quad (2.15)$$

and choose δ_k to fix $\epsilon_{k+1} + \delta_{k+1}/\zeta = \epsilon_k - \delta_k/\zeta$, with ζ a constant (potentially dependent on K). In App. 2.B, we prove that such a scheme will cool a single transition in the range (E_{\min}, E_{\max}) with probability 1 as $K \rightarrow \infty$, and in App. 2.C we demonstrate that the logarithmic gradation is optimal for such a scheme for a choice of $\zeta(K) \sim \log(K)$. To make sure all system excitations have a chance of being dissipated, we further iterate

the couplings V_S over a set of local couplings $\{V_S^i\}$ throughout the system: for the considered spin systems we choose $\{V_S^i\} \equiv \{X_n, Y_n, Z_n\}$ for each qubit n (see Sec. 2.3.2), for a total of $3NK$ cooling steps. The number of Trotter steps M_k for each cooling step k is chosen to prevent re-heating. This follows Eq. 2.8, but as transition energies between system eigenstates may be as large as the Hamiltonian spread $2\|H_S\|_{\perp}$, we set

$$M_k = 2\sqrt{1 + \frac{2\|H_S\|_{\perp}^2}{\gamma_k^2}}. \quad (2.16)$$

The choice of the fridge energy range $[E_{\min}, E_{\max}]$ will generally depend on heuristics on the system. E_{\max} should be greater or equal than the largest energy of the transitions that we are able to de-excite with the chosen couplings V_S (for local Hamiltonians we can estimate this with the techniques described in 2.4.1). For ground state cooling, E_{\min} should be close to the system ground state gap Δ_{GS} , as no transition with an energy lower than Δ_{GS} can lead from an excited state to the ground state.

We first test the LogSweep protocol as applied to the 1+1 model defined in Sec. 2.3.1, with the system gap Δ now taking an unknown value between E_{\min} and E_{\max} (Fig. 2.7). At each step $k = 1, \dots, K$ we want to maximise cooling of transitions $\Delta \sim \epsilon_k$, while minimizing reheating of previously-cooled transitions $\Delta \sim \epsilon_{k'}, k' < k$. As demonstrated by the black dashed curve in Fig. 2.7, when $E_{\max}/E_{\min} = 5$ this can be achieved well with only $K \approx E_{\max}/E_{\min}$ steps. Note that, to maintain a constant relative linewidth (and thus constant maximum reheating per step), we should scale $K \sim E_{\max}/E_{\min}$. This implies $K \rightarrow \infty$ as $E_{\min} \rightarrow 0$, in line with the third law of thermodynamics.

In a larger system, the situation is more complex than in the model above. Instead of a single transition from the excited state $|E_1\rangle \rightarrow |E_0\rangle$ which occurs with unit probability when $\epsilon = \Delta = E_1 - E_0$, our system may transition to one of many eigenstates $|E_j\rangle$, to each with a transition probability $A_{i,j}$ (assuming a initial state $|E_i\rangle$). As there are many possible target states, the maximum transition probability might be very small ($\max_j A_{i,j} \ll 1$). If we restrict to a single transition $|E_i\rangle \rightarrow |E_j\rangle$ with the above reduced cooling rate, one may show that the LogSweep protocol still cools that transition with unit probability as $K \rightarrow \infty$, albeit at a rate that scales exponentially in $A_{i,j}$. Luckily, we do not need to ensure any specific transition occurs, instead we may cool sequentially

$$|E_i\rangle \rightarrow |E_{j_0}\rangle \rightarrow |E_{j_1}\rangle \rightarrow \dots \rightarrow |E_0\rangle, \quad (2.17)$$

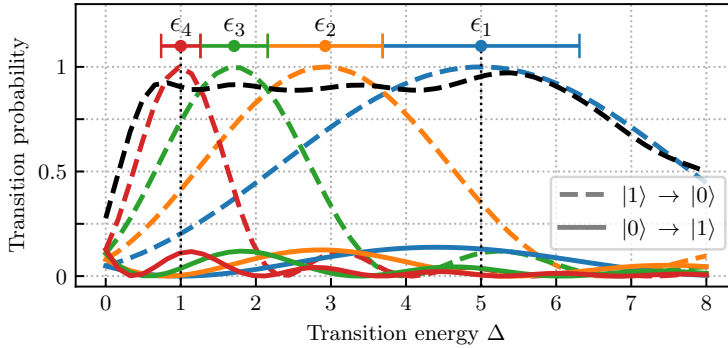


Figure 2.7: Choices of energies ϵ_k and linewidths δ_k (bars at the top of the graph showing $\epsilon_k \pm \delta_k$) for a $K = 4$ LogSweep protocol applied to the model introduced in section 2.3.1 with an unknown $\Delta \in (E_{\min} = 1, E_{\max} = 5)$. Colored lines show cooling (dashed) and reheating (solid lines) probabilities for each j -th step alone, the dashed black line shows the cooling probability after sequential application of the 4 steps.

with a growing number of possible cooling paths as the system grows and the transition probabilities spread over more eigenstates. A good choice of the fridge energy interval $[E_{\min}, E_{\max}]$ and of the coupling potentials $\{V_S^i\}$ allows all eigenstates to be connected to the ground state by sequences of transitions $|E_{j_l}\rangle \rightarrow |E_{j_{l+1}}\rangle$ that have unit probability of being de-excited for $K \rightarrow \infty$. However, a single transition probability approaches 1 only over the entire LogSweep protocol. In particular, if the transition $|E_{j_l}\rangle \rightarrow |E_{j_{l+1}}\rangle$ during step k^* of the protocol corresponds to an energy loss $E_{j_l} - E_{j_{l+1}} \gg \epsilon_{k^*}$, this transition will be off-resonance for the entire remaining duration of the protocol (as $\epsilon_k < \epsilon_{k^*}$ for $k > k^*$), making it unlikely to occur. This can cause convergence issues especially when cooling systems with banded spectra. For such systems, as we set $E_{\min} \approx \Delta_{\text{GS}}$ as detailed above, there may be a point k^* in the protocol after which ϵ_k will become smaller than the average interband gap, but never as small as the spread of a single band. After this point, states at the bottom of a band might transition to states in the lower band, but states at the top of each band never have any resonant transitions to lower energy states, thus becoming absorbing states. This effect is clearly shown in Fig. 2.8, representing the LogSweep-cooled states of the transverse-field Ising model in different regimes. We start with the maximally-mixed state, and plot

the resultant distribution over the eigenstate energies. In the banded regimes (side panels), we observe that the distribution of energies in any given band is tilted towards the higher-energy states in that band (i.e. the aforementioned absorbing states), by some orders of magnitude. This dead-ends ultimately hinder sequential cooling, and prevent the LogSweep cooling from converging to the same state independently on the initial state. The effect worsens as K is increased, as transition linewidths δ_k become smaller making off-resonant transitions less and less probable. This issue can be fixed in practice by using an initial state with fewer high-energy excitation (e.g. a classical approximation of a low-energy state). We solve the issue in principle, by constructing an *iterative LogSweep* protocols, where the LogSweep cooling is repeated with growing K . The early, lower-cost iterations cool the highest energy excitations, while the larger K iterations grant vanishing reheating, and probabilities approaching unity for the cooling transitions allowed by symmetries. Thus, adding iterations with larger and larger K , will make the whole protocol converge to the system ground state (unless symmetries forbid all paths from some states to the ground state). Note that this adjustment is not required for systems with a continuous spectra (i.e. critical systems), as in such a system there will be on-resonance transitions for any state with an energy E_{\min} or more above the ground state.

We now investigate the performance of the (iterative) LogSweep protocol on different phases of the transverse-field Ising model. In Fig. 2.9, we plot the ground state infidelity of the prepared state ρ [$1 - F$ with F as in Eq. (2.14)], as a function of K . The protocol consists in $K - 1$ sweeps of a LogSweep QDC protocol, each sweep having gradation number $g_l = 2, \dots, K$. The Hamiltonian simulation is performed by second-order Trotter approximation. We investigate the protocol effect on two initial states ρ_0 : the maximally-mixed state $\rho_0 = 1/2^n$ to check for cooling capabilities (dots), and the ground state $\rho_0 = |E_0\rangle\langle E_0|$ (crosses) to show the lower bound originated by reheating. We observe polynomial convergence to the ground state in all three phases of the model, attaining an infidelity of $\varepsilon = 1 - F$ in approximately $K \sim O(\varepsilon^{-1/\beta})$ energy gradation steps for $\beta \approx 0.4$ -0.8. Additionally, we verify that the protocol converges to the reheating limit for the critical and strongly-coupled regimes. In the weakly-coupled regime instead, although the cooling is far more efficient because of the local nature of the system excitations, the reheating bound is not saturated. We attribute this to the very small linewidths $\{\delta_k\}$, consequence of the well-defined transition energies, together with the strong banding of the system spectrum.

The number of Trotter steps for a single iteration of the LogSweep pro-




Figure 2.8: Effect of banding on single LogSweep iterations. A maximally mixed state in the three different phases of the 7-qubit TFIM spin chain is evolved by the LogSweep protocol for three different values of K . We plot the distribution of the result here over the system’s eigenstates (indexed by energy), at three different values of K . We see that while the critical system demonstrates an approximate thermal or exponential distribution, the weak and strongly-coupled systems demonstrate an inversion in the population of the system within each band, which increases with K . Data generated by continuous-evolution density-matrix simulation (details in App. 2.F).

tolcol with gradation number g_l on a system of N spins with Hamiltonian H_S scales as $O\left(\|H_S\|_{\perp} \Delta_{\text{GS}}^{-1} N g_l^2 \log(g_l)^{-1}\right)$. Thus, the iterative implementation required to deal with the banded cases needs a total number of Trotter steps

$$M_{\text{tot}} \sim O\left(\|H_S\|_{\perp} \Delta_{\text{GS}}^{-1} N K^3 \log(K)^{-1}\right) \quad (2.18)$$

The gate complexity required to attain an error (infidelity) ε for the models studied scales thus as $O(\varepsilon^{-3}) - O(\varepsilon^{-8})$.

We next investigate the scaling of the LogSweep protocol as a function of the system size. In Fig. 2.10 we plot the relative error in the ground state energy as a function of the system size for a single (not iterated) LogSweep with gradation number $K = 5$. We see a constant error in the ground state energy as a function of the system size for the weakly-coupled and critical systems. Thus, here we expect no need to scale K with N for the protocol to be accurate. Let us also note that the gap

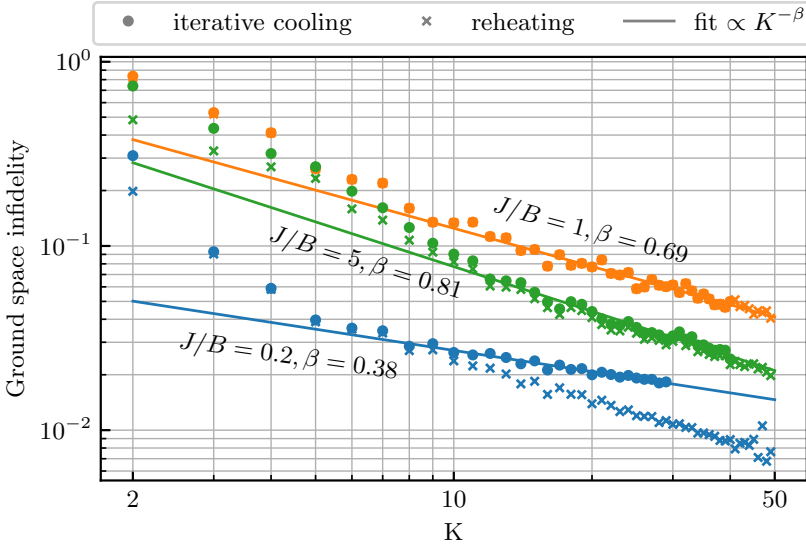


Figure 2.9: Convergence of the LogSweep protocol to the ground state as a function of the gradation number K , starting from the maximally mixed state (dots) and the ground state (crosses), for three phases of the transverse-field Ising model (detailed in legend). Data was generated by deterministic density-matrix simulations of the iterative LogSweep protocol, with second-order Trotter Hamiltonian simulation (details in App. 2.F).

in these two cases shrinks as $\Delta_{\text{GS}}/\|H\| \sim N^{-1}$ and $\Delta_{\text{GS}}/\|H\| \sim N^{-2}$ respectively. Using the above arguments and the estimate (2.18), one can find how the circuit length (in terms of time evolution steps), required to obtain a constant energy error, scales with N . We obtain $O(N^2)$ for the weakly-coupled and $O(N^3)$ for the critical case. From this analysis, we expect that the QDC protocol may be asymptotically competitive with methods such as adiabatic state preparation, whose runtime naively scales as $O(1/\Delta_{\text{GS}}^2)$ [15, 49]. In the strongly-correlated phase, we do not see such positive results; the energy error increases with the system size, though the relative error remains beneath 10% for up to 14 spins. This may be explained by the relative growth of the extension of excitations within the strongly correlated phase, while cooling is performed with local couplings. Due to the error in the simulation, we are unable to reliably extract an




Figure 2.10: Performance of the LogSweep protocol as a function of the system size for the three different phases of the transverse-field Ising model (detailed in legend), with fixed $K = 5$. Dots correspond to result when protocol is applied to the maximally-mixed state, shaded region corresponds to result when protocol is applied to the true ground state (which gives a bound on protocol re-heating). Data generated by Trotterized wave-function simulations of the protocol, and random sampling of the initial mixed state and of nonunitary operations (details in App. 2.F). All points are run with 100 samples, and average results are plotted with the sampling error.

estimate of the computational cost in the same way as for the critical and weakly-coupled systems. Future work may explore whether this error may be improved on by adjusting the form of the coupling terms $\{V_S^i\}$ based on heuristics on the considered system.

2.5 Conclusion

In this chapter, we investigated how cooling can be simulated on a digital quantum computer, and demonstrated that this can be exploited for the design of scalable algorithms for preparing ground states of N -qubit systems. We identified how one can meet many of the fundamental challenges that the digital approach to cooling raises and use the leverage offered exclusively by digital quantum hardware, namely the freedom of choice in the coupling strength and fridge energy. We laid out a general

approach of simulating a cold bath with a single auxiliary qubit, which is iteratively coupled to various locations in the system and reset periodically to extract entropy and energy. We studied how to digitize the system- fridge coupling simulation without causing additional reheating, and how to avoid symmetries which produce non-ergodic behavior that hinders cooling. By tuning coupling parameters beyond the perturbative regime described by Fermi’s golden rule, efficient cooling of targeted transitions can be realized. Following these principles we proposed two protocols for preparing approximate ground states of N -qubit systems — the BangBang protocol and the LogSweep protocol. We studied numerically how these protocols perform on the three phases of the 1D transverse-field Ising model. We found that the BangBang protocol quickly cools the system near to the ground state in the paramagnetic and in the ferromagnetic regime, but has difficulty in the critical regime. The LogSweep protocol is observed to cool all three phases to the ground state at a polynomial cost in the overlap error. In the weakly-coupled and critical phases, the LogSweep protocol further demonstrates a constant energy error as a function of the system size (for fixed gradation number), making it a competitive state preparation method.

Thanks to the low number of steps required, we believe the BangBang protocol has the potential of finding a near-term application, specially if integrated with projective or variational methods to improve its performance. The requirement of precise Hamiltonian simulation and multiple cooling steps makes the LogSweep algorithm in its current form unsuitable for near-term implementation. Nevertheless, the scaling arguments discussed at the end of section 2.4.2 show that our algorithm can be competitive with other non-NISQ methods such as projective quantum phase estimation (QPE) and adiabatic state preparation. QPE in its standard form requires multiple auxiliary qubits and an initial state with a finite ground state overlap, while as our has no requirement on the initial state. Adiabatic state preparation requires an integrable Hamiltonian which can be adiabatically connected to the required Hamiltonian, and requires time scaling as $O(1/\Delta_{\min}^2)$ [15, 49] where $\Delta_{\min} < \Delta_{\text{GS}}$ is the minimum gap along the adiabatic path.

The introduction of quantum digital cooling opens future research directions related to the characterization of proposed protocols, their optimization, and their extension beyond ground state preparation. A study of the effect of noise on currently proposed QDC protocols, and the optimization of such protocols for noise resilience, are in order to establish their applicability on near-term devices. Applying QDC to more complex physical systems, in areas such as quantum spin liquids, many-body

localization and quantum chemistry, would bring new challenges to the protocol construction. A thorough study of the role in the cooling process played by the symmetries and locality of coupling could lead to the design of more optimized protocols. Furthermore, various extensions to the QDC protocols we proposed can be suggested. In a parallelized version of QDC, the use of multiple fridge qubits coupled to various locations in the system might allow to trade space complexity for time complexity. A variationally-optimized QDC protocol might be devised, that can efficiently prepare a state in the ground state manifold of some Hamiltonian starting from an arbitrary initial state — differently from the variational quantum eigensolver [19] which requires the preparation of a fiducial state at every iteration. The principles of QDC might inspire a new class of efficient non-unitary quantum algorithms, where non-unitary operations are mediated by a single auxiliary qubit, with possible application e.g. in the simulation of open quantum system dynamics.

One application of particular future interest for QDC protocols is in the preparation of Gibbs thermal states, which are useful e.g. for semi-definite programming [100]. This seems especially promising given the near-thermal distribution in Fig. 2.8 of the critical system under the evolution of the LogSweep scheme. However, it is as of yet unclear how to overcome the finite width of the distribution, and how well these protocols behave in the banded case (or for more general systems). Adjustment of the LogSweep protocol to produce robust thermal state preparation schemes is an obvious target for future research.

2.A Proof of Eq. (2.10)

To prove Eq. (2.10) we first show that

$$|\langle \phi | O | \psi \rangle| \leq \max_{|\Phi\rangle, |\Psi\rangle} \frac{\langle \Phi | O | \Phi \rangle - \langle \Psi | O | \Psi \rangle}{2}, \quad (2.19)$$

for all $|\psi\rangle, |\phi\rangle : \langle \psi | \phi \rangle = 0$. We can assume without loss of generality $\langle \phi | O | \psi \rangle$ is real and nonnegative (if it's not, we can multiply one state by an irrelevant global phase), and drop the absolute value. As $\langle \psi | \phi \rangle = 0$ we

can define the states $|\pm\rangle = \frac{|\phi\rangle \pm |\psi\rangle}{\sqrt{2}}$ we can then write

$$\begin{aligned} \langle\phi|O|\psi\rangle &= \frac{1}{2}(\langle\phi|O|\psi\rangle + \langle\psi|O|\phi\rangle) \\ &= \frac{\langle+|O|+\rangle - \langle-|O|-\rangle}{2} \end{aligned}$$

immediately proving Eq. (2.19). The opposite inequality is proven by noticing that the $|\Psi\rangle$ and $|\Phi\rangle$ that maximize the right of Eq. (2.10) have to be eigenvalues (by the variational principle). With these, we can redefine the states $|\pm\rangle = \frac{|\Phi\rangle \pm |\Psi\rangle}{\sqrt{2}}$ which are also granted to be orthogonal, thus

$$\begin{aligned} \frac{\langle\Phi|O|\Phi\rangle - \langle\Psi|O|\Psi\rangle}{2} &= \text{Re}[\langle+|O|-\rangle] \\ &\leq |\langle+|O|-\rangle| \\ &\leq \max_{\langle\phi|\psi\rangle=0} |\langle\phi|O|\psi\rangle| \end{aligned}$$

which combined with Eq. (2.19) proves Eq. (2.10).

2.B Asymptotic reheating and cooling probabilities for QDC protocols

Let us consider a two-state subsystem of a larger Hilbert space with a gap energy E , evolving under a QDC protocol on the k th step via a coupling term that does not mix the $\{|01\rangle, |10\rangle\}$ and $\{|00\rangle, |11\rangle\}$ subspaces (where the second index denotes the fridge). Under this assumption, the evolution of the system within this space is a Markov process. Following the main text, let the fridge energy on the k th step be ϵ_k , the coupling strength be γ_k , and the time evolved for in the cooling protocol t_k . Additionally, let the spacing of the fridge energies to be

$$(\epsilon_k - \epsilon_{k+1}) = \zeta(\delta_k + \delta_{k+1}) = \frac{\alpha}{2}(\gamma_k + \gamma_{k+1}),$$

for some K -dependent $\alpha = \alpha(K) = \frac{2}{\pi\zeta(K)}$. We may calculate the transition matrix for the Markov process, $p^{(k)}(E)$ (defined by $p_{a,b}^{(k)}(E) = P(|a\rangle \rightarrow |b\rangle)$)

2 Quantum digital cooling

in a single cooling step) as

$$p^{(k)}(E) = \begin{pmatrix} 1 - \sin^2(\Omega_k t_k/2) \frac{\gamma_k^2}{\Omega_k^2} & \sin^2(\omega_k t_k/2) \frac{\gamma_k^2}{\omega_k^2} \\ \sin^2(\Omega_k t_k/2) \frac{\gamma_k^2}{\Omega_k^2} & 1 - \sin^2(\omega_k t_k/2) \frac{\gamma_k^2}{\omega_k^2} \end{pmatrix}, \quad (2.20)$$

where

$$\omega_k = \sqrt{(E - \epsilon_k)^2 + \gamma_k^2} \quad (2.21)$$

$$\Omega_k = \sqrt{(E + \epsilon_k)^2 + \gamma_k^2}. \quad (2.22)$$

Assuming no additional cooling or heating to the rest of the system during the protocol, the transition matrix for the $k_0 \rightarrow k_1$ block takes the form

$$P_{k_0, k_1}(E) = \prod_{k=k_0}^{k_1} p^{(k)}(E), \quad (2.23)$$

and the transition matrix for the entire process may be written $P(E) = P_{1, K}(E)$.

Exact analytic evaluation of this expression in the large K limit is quite difficult. Instead, we aim for a conservative estimate, bounding the final cooling probability $p_c = [P(E)]_{01}$ from below. For this, given the energy E , we first lower bound the ‘initial’ cooling around the resonant step k_c , i.e. such k_c that $|\epsilon_{k_c} - E|$ is minimal. Then we give an upper bound on reheating during the following protocol steps $k = k_c, \dots, K$. Given the estimated cooling probability $p_c^{(k_c)}$ and reheating probability $p_{rh}^{(k_c; K)}$, we can obtain a lower bound for p_c :

$$p_c > (1 - p_{rh}^{(k_c; K)}) p_c^{(k_c)} \quad (2.24)$$

The value of $p_c^{(k_c)}$ can be conservatively estimated from the formula:

$$1 - p_c^{(k_c)} < \prod_{k=1}^K (1 - \sin^2(\omega_k t_k/2) \frac{\gamma_k^2}{\omega_k^2}) \quad (2.25)$$

$$< \prod_{k, \frac{|E - \epsilon_k|}{\gamma_k} < 1} ((E - \epsilon_k)^2 / \gamma_k^2), \quad (2.26)$$

where the second line follows from the inequality $\sin\left(\frac{\pi\sqrt{1+x^2}}{2}\right)/(1+x^2) \geq$

$\min(0, 1 - x^2)$ applied to each term in the product. In the perfect resonance scenario, $|E - \epsilon_{k_c}| = 0$ and the cooling probability is exactly 1. The worst case scenario is when E is right between the two neighbouring ϵ_k 's, thus $|E - \epsilon_{k_c}| = \frac{\alpha}{2}\gamma_k$. In this case, we can calculate the logarithm of (2.26) in the leading order of K^{-1} , α :

$$2 \sum_{k=k_c^{(-)}}^{k_c^{(+)}} \log \left| \frac{\epsilon_k - E}{\gamma_k} \right| = 2 \int_{\epsilon^{(-)}}^{\epsilon^{(+)}} \log \left| \frac{\epsilon - E}{\gamma(\epsilon)} \right| d\epsilon \frac{dk}{d\epsilon} \quad (2.27)$$

$$= \frac{2}{\alpha} \int_{\epsilon^{(-)}}^{\epsilon^{(+)}} \log \left| \frac{\epsilon - E}{\gamma(\epsilon)} \right| \frac{d\epsilon}{\gamma(\epsilon)}. \quad (2.28)$$

Here, we used the fact that $\gamma\alpha$ defines energy spacing (and so $\frac{d\epsilon}{dk} = \alpha\gamma(\epsilon)$), and introduced summation limits $k_c^{(\pm)}$, $\epsilon^{(\pm)}$ as the points where $\frac{\epsilon - E}{\gamma} = \pm 1$. As this implies scaling $\epsilon^{(\pm)} = E + O(\gamma)$, (2.28) should scale as $O(1/\alpha)$. The calculation can be completed for the LogSweep gradation ϵ_k , γ_k , which implies $\epsilon'_k \propto \gamma(\epsilon) \propto \epsilon$. In particular, if $x = \frac{\epsilon - E}{\gamma}$ then $dx = \frac{E d\epsilon}{\epsilon\gamma} = \frac{d\epsilon}{\gamma} (1 + O(1/K))$, and we have:

$$\frac{2}{\alpha} \int_{\epsilon^{(-)}}^{\epsilon^{(+)}} \log \left| \frac{\epsilon - E}{\gamma(\epsilon)} \right| \frac{d\epsilon}{\gamma(\epsilon)} = \frac{4}{\alpha} \int_0^1 \log x \, dx = -\frac{4}{\alpha}. \quad (2.29)$$

Substituting into Eq. 2.26, we find the initial cooling probability bounded by

$$p_c^{(k_c)} \gtrsim 1 - \exp(-4/\alpha(K)). \quad (2.30)$$

The reheating accumulated between steps k_c and K , $p_{rh}^{(k_c;K)}$, can be upper bounded as:

$$p_{rh}^{(k_c;K)} \leq 1 - \prod_{k=k_c}^K \left(1 - \sin^2(\Omega_k t_k/2) \frac{\gamma_k^2}{\Omega_k^2} \right) \quad (2.31)$$

The product in Eq. (2.31) can be further estimated as:

$$\prod_{k=k_c}^K \left(1 - \sin^2(\Omega_k t_k/2) \frac{\gamma_k^2}{\Omega_k^2} \right) \geq \prod_{k=k_c}^K \left(1 - \frac{\gamma_k^2}{\Omega_k^2} \right) \quad (2.32)$$

$$\geq \prod_{k=k_c}^K \left(1 - \frac{\gamma_k^2}{(E + \epsilon_k)^2} \right) \quad (2.33)$$

$$\simeq \exp \left(- \sum_{k=k_c}^K \frac{\gamma_k^2}{(E + \epsilon_k)^2} \right), \quad (2.34)$$

where in the last line we assumed that $\gamma_k \ll E + \epsilon_k$ for all k . As we are most concerned about the large K asymptotics of the total cooling probability, let us now analyze how the expression (2.34) behaves in this limit. Since γ_k^2 scales as $O(1/K^2)$ and we have K terms in the sum, we generally expect $O(1/K)$ scaling for the sum. Such scaling would imply a rapidly vanishing reheating for a large- K protocol. In the specific case of the LogSweep protocol, to the leading order in $1/K$ one indeed obtains:

$$p_{rh}^{(k_c;K)} \lesssim \sum_{k=k_c}^K \frac{\gamma_k^2}{(E + \epsilon_k)^2} \approx \frac{1}{\alpha(K)} \int_{E_{\min}}^E \frac{\gamma(\epsilon)}{(E + \epsilon)^2} d\epsilon \quad (2.35)$$

$$\approx \frac{\log \frac{E_{\max}}{E_{\min}}}{\alpha^2(K)K} \left(\frac{1}{2} - \frac{E}{E + E_{\min}} + \log \left(\frac{2E}{E + E_{\min}} \right) \right) \quad (2.36)$$

$$\equiv \frac{\mathbf{R}(E_{\min}, E_{\max}, E)}{\alpha^2(K)K}. \quad (2.37)$$

Here, we used Eq. (2.15) and the fact that $\alpha(K)\gamma_k$ defines energy spacing $|\epsilon_{k+1} - \epsilon_k|$. Finally, combining Eqq. (2.24) - (2.37), we obtain an asymptotic lower bound to the final cooling probability:

$$p_c = \left(1 - \exp \left(- \frac{4}{\alpha(K)} \right) \right) \cdot \left(1 - \frac{\mathbf{R}(E_{\min}, E_{\max}, E)}{\alpha^2(K)K} \right). \quad (2.38)$$

This estimate implies $p_c \rightarrow 1$ for large K , provided that both $e^{-4\alpha^{-1}(K)} \rightarrow 0$ and $\frac{1}{K\alpha^2(K)} \rightarrow 0$.

To ensure that the infidelity is minimized and thus $\alpha(K)$ is optimal, we solve the extremum condition $\partial_\alpha (e^{-4\alpha^{-1}} + \frac{R}{\alpha^2 K}) = 0$ for α . The solution can be expressed in terms of the product logarithm function W , $\alpha(K) = 4 W^{-1}(8K/R)$. For large K , at the leading order we obtain simply:

$\alpha(K) = 4 \log^{-1}(8K/R)$. The infidelity then scales down almost linearly with K : $1 - p_c = \frac{\log^2(8K/R)}{16K}$. This asymptotically optimal $\alpha(K)$ yields the choice $\zeta(K) = \frac{1}{2\pi} \log(8K/R)$, which we use in all our simulations.

2.C Optimizing energy spacing in LogSweep protocol

In Sec. 2.4.2, we argued that the energy spacing of the LogSweep protocol is optimal for the protocol precision for a K -step protocol. This was based on the reheating estimate taken from the cooling step k_c only. One may ask, if this persists when one includes the total reheating into account. In the large K limit, we can use the estimate (2.34) for this check. Fixing the constraint $\gamma_k = \frac{|\epsilon_{k+1} - \epsilon_k|}{\alpha}$, we proceed by means of variational calculus:

$$\frac{\delta}{\delta \epsilon_k} \sum_{k=k_c}^K \frac{\gamma_k^2}{(E + \epsilon_k)^2} = 0 \quad (2.39)$$

$$\Rightarrow \frac{\delta}{\delta \epsilon(k)} \int_{k_c}^K \frac{(\epsilon'(k))^2}{(E + \epsilon(k))^2} dk = 0 \quad (2.40)$$

$$\Rightarrow \epsilon''(k) \cdot (E + \epsilon(k)) = (\epsilon'(k))^2. \quad (2.41)$$

The solution to Eq. (2.41) that satisfies boundary conditions $\epsilon(k_c) = E$, $\epsilon(K) = E_{\min}$, is as follows:

$$\epsilon_k = (2E)^{\frac{K-k}{K-k_c}} (E + E_{\min})^{\frac{k-K}{K-k_c} + 1} - E. \quad (2.42)$$

This shows that the logarithmic character of the optimal spacing persists when we consider total reheating (cf. Eq. (2.15)). However, we cannot directly use the embellished result (2.42) for our cooling protocol. That is because this formula uses the targeted energy E as a reference, whereas we are targetting a continuum of energies. Therefore, we keep using the simpler and more practical formula Eq. (2.15) for the LogSweep protocol.

2.D Cooling rate for LogSweep protocol in a large system

In a large system, the above analysis is complicated by the presence of multiple transitions from every energy level. We now give a simplified

analysis that focuses on a pair of states $|E_i\rangle, |E_j\rangle$, in a spirit similar to Appendix 2.B. This means we formulate the protocol as a Markov process equivalent to (Eq. 2.20), where the transitions to levels other than i and j are ignored. Note that in the perturbative limit, this is a good approximation of the actual Markov process as restricted onto the subspace $|E_i\rangle, |E_j\rangle$. Specifically, even though we ignore the indirect transitions between i and j via other levels, this is justified at the first order of perturbation theory. Unlike in the $1 + 1$ model however, the transitions here are imperfect. If our total coupling has strength γ (i.e. $\|H_C\| = 2^N \gamma$), following the analysis in Sec. 2.4.1 the coupling between states $|E_i\rangle$ and $|E_j\rangle$ will take the form $\gamma \sqrt{A_{i,j}}$ with $\sqrt{A_{i,j}}$ scaling down as $O((E_i - E_j)^{-2})$. This has the effect of scaling both the cooling and re-heating rates by $A_{i,j}$, recasting the Markov process (Eq. 2.20) as

$$p_{i,j}^{(k)} = \begin{pmatrix} 1 - A_{i,j} \sin^2\left(\frac{\Omega_k t_k}{2}\right) \frac{\gamma_k^2}{\Omega_k^2} & A_{i,j} \sin^2\left(\frac{\omega_k t_k}{2}\right) \frac{\gamma_k^2}{\omega_k^2} \\ A_{i,j} \sin^2\left(\frac{\Omega_k t_k}{2}\right) \frac{\gamma_k^2}{\Omega_k^2} & 1 - A_{i,j} \sin^2\left(\frac{\omega_k t_k}{2}\right) \frac{\gamma_k^2}{\omega_k^2} \end{pmatrix}.$$

As this only reduces both the heating and cooling rates, our claim that reheating in the LogSweep protocol tends to 0 as $K \rightarrow \infty$ still holds. However, we need to repeat the analysis of App. 2.B to bound the cooling rate $p_c^{(k_c)}$ below and check that it continues to tend to 1. For the sake of generality, we drop the i, j indices, and consider a cooling probability restricted by a k -independent factor A .

With this adjustment, we may recast Eq.2.26 when $A \ll 1$ as

$$1 - p_c^{k_c} < \prod_{k, \frac{|E - \epsilon_k|}{\gamma_k} < 1} \left[\left(1 - \frac{A\pi^2}{4}\right) + \frac{A\pi^4}{48} \frac{(E - \epsilon_k)^2}{\gamma_k^2} \right]. \quad (2.43)$$

Then, taking the log and converting again to an integral, we obtain

$$\log(1 - p_c^{k_c}) < \frac{1}{\alpha} \int_{\epsilon^-}^{\epsilon^+} \log \left[B + A' \frac{(E - \epsilon)^2}{\gamma(E)^2} \right] \frac{d\epsilon}{\gamma(\epsilon)}, \quad (2.44)$$

where $A' = \frac{A\pi^4}{48} \sim 2A$, and $B = 1 - \frac{A\pi^2}{4} < 1$. Next, setting $x = \frac{E - \epsilon}{\gamma(\epsilon)}$, and using the fact that for the LogSweep protocol $\gamma(\epsilon) \sim \epsilon$, we find

$$\log(1 - p_c^{k_c}) < \frac{2}{\alpha} \int_{-1}^{+1} \log(B + A'x^2) dx. \quad (2.45)$$

This may be evaluated by integrating by parts, giving

$$\log(1 - p_c^{k_c}) < -\frac{2}{\alpha} \int_{-1}^{+1} \frac{x^2}{BA'^{-1} + x^2} \quad (2.46)$$

$$= \frac{-4}{\alpha} [1 - BA'^{-1} \tan^{-1}(A'B^{-1})] \quad (2.47)$$

$$\sim -\frac{4}{3\alpha} A'^2 B^{-2} + O(A^4). \quad (2.48)$$

Using the optimal scaling $\alpha(K) = 4 \log^{-1}(K)$ we identified in Appendix 2.B, this adjusts our bound in the cooling rate to

$$p_c^{k_c} \gtrsim 1 - K^{-\frac{1}{3}} A'^2 B^{-2}, \quad (2.49)$$

which continues to tend to 1 as $K \rightarrow \infty$, albeit at a rate reduced proportional to A .

This result requires some consideration in a large system — if our coupling Γ from a state $|E_i\rangle$ is spread over transitions to J states $|E_j\rangle$, we have $A_{i,j} \sim J^{-1}$, and the probability of any transition being cooled can be found to be

$$\prod_j (1 - p_{c,j}^{k_c}) \sim e^{-\sum_{j=1}^J \frac{1}{3\alpha} A_j^2 (1 - A_j^2)} \sim e^{-\frac{1}{3\alpha J}}. \quad (2.50)$$

This implies that we require $\alpha \sim J^{-1}$ in order to maintain a constant cooling rate, which in turn may require adjustments to the optimal scaling identified in Appendix 2.B. As such adjustments are highly system-dependent, we do not investigate them further here.

2.E Effect of banding on QDC protocols

In this appendix we demonstrate the effect of banding on single sweeps of the LogSweep protocol. In Fig. 2.11, we plot the infidelity of a single shot of the LogSweep protocol with gradation number K acting on the maximally-mixed state, as a function of K (triangular markers). We see that in the critical case, the system continues to tend to the ground state polynomially in K . However, for the TFIM chain in the weakly- and strongly-coupled phases, we find that the protocol fails to converge as a function of K , due to the banding issue described above. This lack of convergence is rectified in the series marked by dots (same data as in Fig. 2.9) by repeating the LogSweep protocol as a function of K . We note




Figure 2.11: Difference between cooling by applying the a single LogSweep protocol with gradation number K (round markers), and iterating LogSweep for all $g_l = 2, \dots, K$ (solid lines). The iterative and reheating data are the same as in Fig. 2.9, the same context and simulation techniques apply.

that the failure in the strong-coupling case is not of the same degree as in the weak-coupling case, which we ascribe to the fact that the banding is not as strongly pronounced in Fig. 2.8, and so the result has not yet presented itself.

2.F Details on numerical methods

In this appendix we report the methods used to simulate QDC protocols on many-qubit systems. The python code is packaged and available on github, at <https://github.com/aQaLeiden/QuantumDigitalCooling> [101], and makes cirq [102] to build and simulate the required quantum circuits. All the data reported in this chapter, along with more simulations results that

are left out in the interest of space and clarity, can be found in the same repository.

The simulations for the BangBang protocol energy validation Fig. 2.5, as well as the study of LogSweep performance with increasing K (Figures 2.9 and 2.11) were performed using cirq density matrix simulator. This stores the system's state in a density matrix, to which are applied sparse unitaries (representing the circuit's unitary gates) and the eventual quantum channel representing the reset gate. For BangBang, as prescribed by the protocol, the unitary circuit applied before each fridge reset gate is defined by second-order Trotter expansion of the coupled system-fridge Hamiltonian with a single step (Trotter number $M = 1$). This corresponds to Eq. 2.7 where also e^{-iH_S} is substituted by its $M = 1$ second-order Trotter expansion.

To push to a larger number of qubits the results on scaling of both protocols, data for in Figures 2.6 and 2.10 were generated with cirq state vector simulator. The non-unitary reset required by QDC protocols and the initially-mixed state used to benchmark cooling cannot be represented deterministically in state vector simulations. These are instead implemented by random sampling. Each sample is constructed by choosing at random an initial computational basis state (these are enough to sample the maximally mixed state, because of the density matrix equivalence class). For each non-unitary reset gate, the outcome is sampled probabilistically. This process is repeated for 100 samples for each data point. The mean value of the quantity of interest is plotted, together with an interval representing the standard deviation of the mean.

In both the density matrix and the wavefunction simulations performed with cirq, the numerical error causes the final state to be often non-normalized. In the worst cases, the deviations from unit L2 norm (for wavefunctions) and trace norm (for density matrices) are up to few parts per thousand and percent respectively. This is attributed to the large number of short-time Trotter steps required to produce LogSweep data in Fig. 2.9, which translate to sparse unitaries with small entries and a building up of numerical error. As all operations are linear, the first-order error can be dealt with by forcing normalization on the final state. This technique was used for results reported in figures 2.9, 2.10 and 2.11

The numerical error of Trotterized sparse-unitary simulations still is too large to show the final state occupations in Fig. 2.8, which range over more than 18 orders of magnitude. For this reason, these simulations were performed by constructing the continuous evolution operator $e^{-i(H_S+H_F+H_C)t}$ for each unitary evolution step. These results were validated by comparing with the Trotterized approach the results for large occupations and the energy expectation values (which are less sensitive to the numerical error).