

A path towards distributed quantum annealing

Raúl Santos,^{1,2,3} Lorenzo Buffoni,³ and Yasser Omar^{1,2,3}

¹*Instituto Superior Técnico, Universidade de Lisboa, Portugal*

²*Physics of Information and Quantum Technologies Group,*

Centro de Física e Engenharia de Materiais Avançados (CeFEMA), Portugal

³*PQI – Portuguese Quantum Institute, Portugal*

Quantum Annealing has proven to be a powerful tool to tackle several optimization problems. However, its performances are severely limited by the number of qubits we can fit on a single chip and their local connectivity. In order to address these problems, in this work, we propose a protocol to perform distributed quantum annealing. Our approach relies on Trotterization to slice the adiabatic evolution into local and non-local steps, the latter which are distributed using entanglement-assisted local operations and classical communications (eLOCC). Theoretical bounds on the Trotter step size and successful distribution probability of the process have been established, even in the presence of noise. These bounds have been validated by simulating numerically the evolution of the system, for a range of annealing problems of increasing complexity.

INTRODUCTION

Solving optimization problems is a very common problem in science with many practical applications, making this a very active field of research. [1–3]. For these problems physics has contributed with lots of concepts and methods to the field of optimization, starting from the idea of thermal simulated annealing [4], to the applications of replica and cavity methods [3]. Quantum Annealing was initially formulated in the 90’s [5] as a quantum alternative to classical simulated annealing in which quantum tunneling replaces thermal hopping in order for the system to avoid being trapped in local minima and reach the ground state (i.e. the solution of the optimization problem). This idea was later developed to the point at which, today, we have working quantum annealers with thousands of qubits available.

However, Quantum Annealers suffer from some issues that are limiting the size and complexity of the problems we are currently able to solve using this technology. The major problems that we will be addressing here are embedding and connectivity, which are closely related. The connectivity problem is simply an engineering bounds to the number of qubits that one is effectively able to couple on the chip layout (in the latest generation quantum annealers the number of couplings per qubit is ~ 15). This limits the problems that one can solve on a quantum annealer to ones that are sparsely connected. A way to circumvent this problem is to use minor-embedding techniques in which a chain of physical qubits is treated as one logical qubit allowing for more connections. This technique, however, is pretty intense in the number of qubit used thus limiting the size of the problems which we can effectively solve.

Recently, advances in the distribution of quantum operations [6–8] opened the possibility of connecting distant qubits and perform operations at a distance. This has led to the emergence of new paradigms such as distributed quantum computation [9–12], by which multiple quantum computers can be effectively entangled to exploit

their joint computational power. Distribution is an interesting prospect for Quantum Annealers, as it could improve their scalability and connectivity. However, so far, no protocol has been proposed for Quantum Annealing, neither it has been investigated the possibility of expanding these techniques to this particular model of computation.

After a brief formal introduction of Quantum Annealing, in this work we will propose a protocol for the distribution of quantum annealing problems across multiple machines. We will also investigate the robustness of the protocol in the face of imperfect distribution and perform numerical simulations on several problems.

BACKGROUND

Quantum annealers are a special type of analog quantum computers, making their inherent process continuous in time. It is governed by a time-dependent Hamiltonian $H(t)$.

$$H(t) = \left(1 - \frac{t}{t_F}\right) H_0 + \frac{t}{t_F} H_F. \quad (1)$$

At first, the system Hamiltonian is H_0 , which acts on the initial state $|\psi_0\rangle$, the ground state of H_0 . While the system Hamiltonian slowly evolves to H_F , its state evolves to $|\psi_F\rangle$. According to the adiabatic theorem (see [13] for a thorough review on adiabatic quantum computation), the higher the annealing time t_F is, the closer the final state $|\psi_F\rangle$ will be to the ground state of the final Hamiltonian H_F .

To model the system of qubits we will consider the Transverse field Ising Model, and so the starting and final Hamiltonians H_0 and H_F take the form in eqs. (2) and (3), respectively. The qubits are labeled by the ordered set $\Omega = (q_0, q_1, \dots, q_{N-1})$. The h_i terms correspond to local fields applied to the qubit i , which are represented by nodes in a graph fig. 1. With 2-body Hamiltonians, we have the couplings J_{ij} which quantify the strength of the $\sigma^z \sigma^z$ interactions between the

qubits i, j . The qubits in the set Ω then are connected by the edges e_0 that indicate which qubits are coupled, i.e. $e_0 = \{(i, j) : J_{ij} \neq 0\}$.

$$H_0 = \sum_{i \in \Omega} \sigma_i^x \quad (2)$$

$$H_F = \sum_{i \in \Omega} h_i \sigma_i^z + \sum_{(i,j) \in e_0} J_{ij} \sigma_i^z \sigma_j^z \quad (3)$$

In the context of computation, this process is useful as an optimization meta-heuristic, where the final Hamiltonian H_F encodes the objective function. Furthermore, having H_F expressed in the Ising Model, quantum annealing becomes an ideal process for solving Quadratic Unconstrained Binary Optimization (QUBO) problems [14, 15]. These problems, other than being generally hard to solve, are often defined on graphs with a large number of connections. This limits the use of current generation quantum annealers as they suffer from connectivity problems as discussed above.

Distributed quantum computation allows for the use of smaller quantum computers, with more manageable noise and therefore could obtain better qubit coherence. Arbitrary quantum operations between two distant parties can be achieved by having a classical and a quantum channel between both parties [16]. This channel can be made, for example, by entangled photons in fibers [17] or by having a superconducting waveguide cryogenic "cable" connecting them [6]. Another of the main challenges of distributed quantum computation is to have an efficient and low noise interface between the entanglement carriers and the computing qubits. For example, optomechanical transducers are actively being researched for being a promising technology that could allow long distance communication for superconducting technology [7, 18–26]. The last step, is to put all of this together to implement a distributed operation on two qubits. Experimentally, the state-of-the-art distributed CNOT operation has achieved a fidelity of $\sim 85\%$ for trapped ions [8].

It is then clear to imagine how distributed quantum computation techniques applied to quantum annealing can allow for virtually unlimited connectivity between qubits allowed to interact non-locally (we will call them the *interface* qubits). Moreover, having the qubit connectivity not restricted to a fixed hardware topology will help solve the connectivity problem present in current quantum annealing devices. Another advantage of distribution is that it allows to take advantage of smaller quantum devices, which are easier to manufacture and manage, therefore improving the scalability. In the next section we describe the theoretical details of the proposed distribution approach before moving to simulations.

DISTRIBUTED UNITARIES FOR QUANTUM ANNEALING

To formulate the proposed distributed quantum annealing protocol, we start by making a distinction between local and non-local Hamiltonians. Local Hamiltonians act on a single annealing processor, while non-local ones are only responsible for coupling qubits which belong in spatially separate processors (see fig. 1). In general, any evolution with time-dependent Hamiltonians can be separated into local $H_L(t)$ and non-local $H_N(t)$ Hamiltonians.




FIG. 1. Example of a configuration for a distributed quantum annealing network. The qubits native to the quantum annealers and their couplings are colored black. The blue areas denote separate annealing systems, which rely on the local, analog evolution of the separate quantum annealers. In this case, $H_L = H_{L_A} + H_{L_B}$. The red dashed lines represent the non-local couplings, belonging to the non-local Hamiltonian H_N . The red points inside the qubit nodes indicate that those qubits are interface nodes, therefore being able to communicate non-locally. The interface qubits can communicate to any other interface qubit reachable by an entanglement distribution network, which immensely increases their connectivity relative to local qubits.

$$H(t) = H_L(t) + H_N(t) \quad (4)$$

The evolution process which happens in quantum annealing can be expressed as a unitary operator, with the use of a time-ordered exponential (5). The process can then be decomposed into a sequence of M unitary operators (6).

$$U(0, t_F) = \mathcal{T}_t \exp \left\{ -i \int_0^{t_F} dt H(t) \right\} \quad (5)$$

$$= \mathcal{T}_{t_k} \prod_{k=0}^{M-1} U(t_k, t_{k+1}) \quad (6)$$

where \mathcal{T}_t is the time-ordering operator, $t_k = k\Delta t_k$ and $\Delta t_k = t_F/M$ is the Trotter step size (chosen to be constant for simplicity). Then, each $U(t_k, t_{k+1})$ is Trotterized into the local and non-local parts

$$U(t_k, t_{k+1}) \rightarrow U_N(t_k, t_{k+1}) U_L(t_k, t_{k+1}). \quad (7)$$

So that we get the whole Trotterized evolution

$$\tilde{U}(0, t_F) = \mathcal{T}_{t_k} \prod_{k=0}^{M-1} U_N(t_k, t_{k+1}) U_L(t_k, t_{k+1}). \quad (8)$$

The local unitary evolution operators $U_L(t_k, t_{k+1})$ represent a partial evolution of the quantum state inside the quantum annealing processor. These alternate with the non-local unitary operators $U_N(t_k, t_{k+1})$, which simulate and distribute the non-local couplings between qubits in separate quantum processors. These are implemented via a unitary distribution protocol, i.e. telegates, such as the one proposed in [16]. Therefore, the distributed annealing process corresponds to the fast switching between local H_L and non-local H_N Hamiltonians, the latter which are implemented via telegates. A similar process which Trotterizes between analog evolution and digital gates on a single quantum processor has been first proposed in [27] and realized experimentally in [28].

Graph Split

As pointed out before, annealing Hamiltonians can be represented with graphs, with vertices labeled $\Omega = (q_0, q_1, \dots, q_{N-1})$ and edges e_0 . The annealing graph then has to be split into multiple quantum annealers, making the identification of the local and non-local couplings. In this work we propose two ways to do so: edge splitting and vertex splitting. Consider two quantum annealers A and B . An original annealing problem represented by \mathcal{G}_0 then turns into local graphs $\mathcal{G}_A, \mathcal{G}_B$ (where $\mathcal{G}_A + \mathcal{G}_B$ has the terms in the local Hamiltonian H_L) and non-local edges e_{AB} . In edge splitting, some of the couplings in e_0 turn into non-local couplings e_{AB} , such that \mathcal{G}_A and \mathcal{G}_B are disjoint. Vertex splitting is a slightly different. Instead of turning a (local) coupling into a non-local coupling, a vertex (qubit) is n -plicated across multiple quantum annealers. For instance, a qubit labeled q in the original problem can be duplicated to annealers A and B , into qubits q_A and q_B . This sounds suspiciously like quantum cloning [29], but it is not the case here. It relies on the initial Hamiltonian H_0 to be such that the state of the system evolves in a way that keeps the n -plicated qubits aligned. The mechanism is described in appendix A. For demonstration purposes, consider the following Hamiltonians representing an annealing problem:

$$\begin{aligned} H_0 &= \sigma_a^x + \sigma_b^x + \sigma_c^x, \\ H_F &= J_{ab}\sigma_a^z\sigma_b^z + J_{ac}\sigma_a^z\sigma_c^z + J_{bc}\sigma_b^z\sigma_c^z. \end{aligned}$$




FIG. 2. Splitting of an annealing graph into multiple ones, with the purpose of distribution to multiple quantum annealers. The red dashed line represents a coupling turned non-local via edge splitting, while the red wavy line represents a non-local coupling required by the vertex splitting approach. In this example, the coupling J_{ab} turns non-local, and the qubit c is duplicated into c_A and c_B . The subgraphs \mathcal{G}_A and \mathcal{G}_B represent the local Hamiltonian that act in each quantum annealer A and B , respectively. The problem of choosing the best possible graph split is analogous to the compilation problem of distributed quantum computation in the circuit model [30].

One way to split these into two annealers is to have

$$\begin{aligned} H_L(t) &= \left(1 - \frac{t}{t_F}\right) ([\sigma_a^x]_A + [\sigma_b^x]_B) \\ &\quad + \frac{t}{t_F} (J_{ac} [\sigma_a^z \sigma_{c_A}^z]_A + J_{bc} [\sigma_b^z \sigma_{c_B}^z]_B) \\ H_N(t) &= \left(1 - \frac{t}{t_F}\right) \sigma_{c_A}^x \sigma_{c_B}^x + \frac{t}{t_F} J_{ab} \sigma_a^z \sigma_b^z \end{aligned}$$

where $[\cdot]_A$ indicates a local term that acts on annealer A (same for B). This transformation is shown in fig. 2. The initial state is the ground state of the new initial Hamiltonian H'_0 , and so the new initial state $|\phi'_0\rangle$ is such that the duplicated qubits are correlated.

$$H'_0 = \sigma_a^x + \sigma_b^x + \sigma_{c_A}^x \sigma_{c_B}^x$$

$$|\phi'_0\rangle = |-\rangle_a |-\rangle_b |\Phi^-\rangle_{c_A c_B}$$

with the Bell state

$$|\Phi^-\rangle = \frac{1}{\sqrt{2}} (|00\rangle - |11\rangle).$$

Convergence of Trotterization

The Trotterized evolution introduces errors due to H_L and H_N not commuting with each other. At each time interval $t \in [t_k, t_{k+1}]$, there may be an equivalent Hamiltonian $H_{eq}(t_k)$ that governs the process as if it was a continuous process. Assuming a uniform time-step, its existence is guaranteed when (see appendix B for proof)

$$t_{k+1} - t_k = \Delta t < \Delta t_M \quad (9)$$




FIG. 3. Distributed CNOT protocol, proposed in [16]. Operations in blue are performed locally, and in red belong to distribution protocols. We consider a noisy entanglement distribution with a shared quantum state denoted by ρ_Φ .

where

$$\Delta t_M = \min_k \left(\|H_L(t_k)\|_L^{-1}, \|H_N(t_k)\|_L^{-1} \right) \quad (10)$$

Furthermore, when $\Delta t/\Delta t_M \ll 1$ we get the bound

$$\|H_{eq}(t_k) - H(t_k)\|_L \leq \frac{\Delta t}{(\Delta t_M)^2} \quad \forall 0 \leq k \leq M. \quad (11)$$

This result indicates that the equivalent Hamiltonian approximates the adiabatic Hamiltonian for small enough Δt . We conclude that, in ideal scenarios (no noise and no limitations for Δt or M), the Trotterized process is equivalent to quantum annealing.

Noisy Distribution

Ideal distribution is equivalent to quantum annealing, but it is not realistic. Therefore we will investigate the behaviour of our protocol in the case where, instead of distributing with pure Bell states, we consider a noisy quantum channel with some entanglement fidelity $F_\Phi < 1$. This analysis will also help us understand if this protocol is robust enough to noise to be implemented in a real-world scenario. The shared quantum state is

$$\rho_\Phi = x |\Phi\rangle \langle \Phi| + \frac{1-x}{4} \mathbb{I}_{2 \times 2} \quad (12)$$

with fidelity

$$F_\Phi = x + \frac{1-x}{4}.$$

The CNOT distribution protocol (DCNOT) consumes a Bell state in the process

$$|\Phi\rangle = \frac{1}{\sqrt{2}} (|00\rangle + |11\rangle)$$

The fidelity of the state between the ideal distribution and the noisy distribution has the bound

$$p_N(\delta) > \left(1 - \frac{4}{3}\delta\right)^{M_D} \quad (13)$$

where M_D is the number of consumed Bell pairs (same as the number of DCNOT gates, and is proportional to the number of Trotter steps M), and $\delta = 1 - F_\Phi$ is the infidelity of the shared entangled pair. This expression is deduced in appendix C. We have found this to be true for any distributed process that employs noisy telestates.

NUMERICAL SIMULATIONS

In this section, we look at how the method performs with different annealing problems, starting with the spin chain, and later with other more densely connected networks. There are two important quantities that we will use to compare the Trotterized process (T , ideal or noisy distribution) with the annealing evolution (A) or the ground state space of H_F , corresponding to the limit when the annealing time $t_F \rightarrow \infty$. These are the fidelity p_{*0} between the evolved state ($* \in \{A, T\}$) and the ground state space of H_F and the relative energy error $\varepsilon_{*\bullet}$ between $* \in \{A, T\}$ and $\bullet \in \{0, A\}$, in eqs. (14) and (15).

$$p_{*0} = \sum_{|\phi_i\rangle \in gs(H_F)} \langle \phi_i | \rho_* | \phi_i \rangle, \quad * \in \{A, T\} \quad (14)$$

where $gs(H_F)$ is the ground state space of the final Hamiltonian H_F .

$$\varepsilon_{*\bullet} = \frac{E_* - E_\bullet}{E_{\text{mix}} - E_\bullet} \quad * \in \{A, T\}, \quad \bullet \in \{0, A\} \quad (15)$$

$E_* = \langle \psi_* | H_F | \psi_* \rangle$ is the energy of the state after evolution, measured in the final Hamiltonian H_F , with E_A or E_T for adiabatic and Trotterized evolutions, respectively. $E_0 = \langle \psi_0 | H_F | \psi_0 \rangle$ is the ground state energy of H_F . $E_{\text{mix}} = \langle \phi_0 | H_F | \phi_0 \rangle$ is the energy of the initial state $|\phi_0\rangle$ measured in H_F . This corresponds to the average energy of a state after a process that randomly guesses the minimum of H_F , while respecting the evolution dynamics of H_0 . It is, in principle, the worst case scenario.

With this definition of the renormalized energy error, we have that $\varepsilon_{*0} \sim 0$ when the evolution successfully finds the ground state (or one of the ground states, if H_F is degenerate), and $\varepsilon_{*0} \sim 1$ when no useful information about the ground state can be obtained from $|\psi_*\rangle$. These qualities are also true for the probability $1 - p_{*0}$. There is also the case when the adiabatic evolution doesn't find the ground state $\varepsilon_{A0} \sim 1$. In this case, the Trotterized process might still be able to approximate the adiabatic evolution in case $\varepsilon_{TA} \sim 0$.

Toy model

The first set of simulations will be about a simple graph network - the four qubit spin chain. The graph splitting for the spin chain is shown in fig. 5 (a). There are 3

tunable parameters related to the Trotterization process:

- The annealing time t_F , which controls the speed of the adiabatic evolution;
- The number of Trotter steps M , or equivalently the Trotter step size $\Delta t = t_F/M$, which control the accuracy of the Trotterization;
- The noise of the distribution process, tuned via the entanglement fidelity F_Φ .

The overall effect of each of the parameters can be seen in fig. 4 (a), indicated with different colored arrows. Firstly, a gradient is seen as the annealing time t_F increases and the relative energy error ε_{T0} decreases. This effect is due to the adiabatic nature of the process: the slower the process is, the closer the final state is to the ground state of H_F , and therefore $\varepsilon_{A0} \rightarrow 0$. Secondly, there is a clear transition region at constant Δt , which was expected from the theoretical analysis. This transition is visible at roughly $\Delta t \sim 1.0$, while the theoretical prediction is

$$\Delta t_M \equiv \min_k \left(\| |H_L(t_k)| \|_L^{-1}, \| |H_N(t_k)| \|_L^{-1} \right) = 0.5,$$

being less than the one observed in the simulated Trotterization. The region where the algorithm appears to converge for $\Delta t > \Delta t_M$ is likely related to the conditional convergence of the equivalent Hamiltonian H_{eq} , where our theoretical analysis cannot establish any bound. Thirdly, a gradient is visible as the number of Trotter steps M increases. Since $F_\Phi < 1$ there is distribution noise. With bigger M there are more distribution steps, and naturally the errors due to distribution start being more noticeable.

Due to the adiabaticity condition (slow evolution, large t_F) the Trotterized adiabatic evolution is similar to the evolution of Floquet systems for intermediate time intervals $t \in [\tau, \tau + \Delta\tau]$, $\tau \in [0, t_F - \Delta\tau]$ with $\Delta t \ll \Delta\tau \ll t_F$. The phase transition we have observed is also a common feature of Floquet systems [31–34]. The region that remains at low energy, i.e. where $\varepsilon_{T0} \sim 0$, is the localized phase, and is associated with low ergodicity. In contrast, the region where $\varepsilon_{T0} \sim 1$ has energy close to the average of the eigenenergies of the Hamiltonian, and so is analogous to a system with infinite temperature, i.e. high ergodicity. Thus it is commonly referred to as the thermalization phase. In the literature this phase transition is known as the many-body localization (MBL) transition.

Fig. 4 (b) shows the plot of the relative energy error between the Trotterized and the adiabatic evolution ε_{TA} , versus Δt . We can identify 3 regions. An initial one with $\Delta t < \Delta t_M$, where we are assured of the convergence of the equivalent Hamiltonian H_{eq} , and where we verify that $\varepsilon_{TA} < \varepsilon_{A0}$, meaning that the energy of the state after Trotterization E_T is closer to the energy of the state after annealing E_A , than E_A is to the ground




FIG. 4. (a) Relative energy error of the Trotter process, varying the annealing time t_F , the number of Trotter steps M , and considering noisy distribution with $F_\Phi = 0.999$. The three arrows disclose the effect of the tunable parameters in the Trotterization process. The blue arrow points towards a gradient of the error due to improved adiabaticity conditions, as t_F grows. The green arrow points across the many-body localization transition that is natural to this process. The gradient directed by the red arrow is a result of noisy distribution, since as the number of distribution steps increases, also do the errors related to distribution. The theoretical convergence Δt_M is drawn with a dashed green line and an estimation for the transition is drawn with a light gray dashed line. (b) Relative energy error between the Trotterized process and the ideal annealing evolution ε_{TA} , with varying Trotter step size Δt . The annealing times have been fixed, and are indicated in different colors. The horizontal dash-dotted lines mark the annealing error ε_{A0} , for each annealing time t_F . The annealing error ε_{A0} at the absolute convergence transition point is larger than the Trotter error ε_{TA} , meaning that the absolute convergence condition appears to be enough to select an adequate Δt .

state energy of H_F . It follows that in this region the energy error of the Trotterized evolution ε_{T0} is almost all due to the annealing error ε_{A0} , and corresponds to the localized phase. Then, a region where $\Delta t > \Delta t_M$ and the error is still bounded, is where the phase transition occurs. We speculate that it is characterized by the conditional convergence of H_{eq} , which we have not been able to determine analytically. Finally, a third region where $\varepsilon_{TA} \sim 1$ is linked to high ergodicity and infinite temperature, and corresponds to the divergence of H_{eq} . Here, the Trotterized evolution cannot be replaced by a continuous process governed by the Schrödinger equation, and as such the solution reached by distributed quantum annealing will not resemble quantum annealing.




FIG. 5. Example of models with Spin Chain and Sparse Network graph topology. The coupling graphs have been parametrized in different ways to keep control on the number of qubits, non-local couplings and local systems. **(a) Spin Chain:** Parametrized by the number of qubits N and the number of sublocal systems s . The shown example has $(N, s) = (6, 3)$. The local fields are $h_i = 0$, the local couplings are $J_L = 1$ and the non-local couplings $J_N = -2$. The number of non-local connections is tuned by changing s . **(b) Sparse Network:** Parametrized by the number of qubits N , number of qubits on system 0 l , with $l < N$. The example graph has $N = 7$ total qubits, $l = 3$ qubits on sublocal system 0 and $N - l = 4$ qubits on sublocal system 1 . The non-local connections are tuned by changing the number of interface qubits on either side. These are i_0, i_1 for interface qubits in sublocal system $0, 1$ respectively. Therefore, the example graph is parametrized by $(N, l, i_0, i_1) = (7, 3, 3, 2)$. The local fields are $h_i = 0$ and all couplings shown in the graph have $J = \pm 1$, with the sign chosen at random, given a generation seed.

Spin Chain and Sparse Networks

In this section we consider two different models that have more complex coupling networks. These are represented in fig. 5. Firstly, for all of the different network topology, we fixed the annealing energy error $\varepsilon_{A0} \sim 10^{-4}$ by tuning the annealing time t_F . Then for the ideal Trotterized evolution we have verified that with $\Delta t = \Delta t_M$, the relative energy error $\varepsilon_{TA} \lesssim \varepsilon_{A0}$ with ideal distribution ($\delta = 0$). From these simulations we know that the ideal Trotterization approximates well the evolution by annealing with $\Delta t = \Delta t_M$, indicating also that we are on the localized phase.

When we introduce noise, we have observed that the distribution fidelity eq. (13) is a crude lower bound. We introduce a free parameter β , and we observe that it tightens the bound considerably (see fig. 6 (a)). The distribution fidelity then becomes

$$p_N(\delta; \beta) = \left(1 - \frac{4}{3} \delta (1 - \beta)\right)^{M_D}. \quad (16)$$

The motivation for the β parameter is given in appendix C. The estimation in eq. (13) is too coarse as




FIG. 6. **(a)** Log negativity of the probability of measuring the Trotterized state at the ground state of H_F . It is plotted as a function of the entangled pair infidelity δ , for some spin chain models labeled by N, s (explanation of the model in fig. 5 (a)). The simulated data is shown in contiguous lines. The predicted curves where no distribution fault is allowed correspond to the dashed lines. The adjusted theory, with the β parameter fitted to the data, is shown in dotted lines. These overlap for with the data for the majority of the values of δ . This is not the case for larger δ , where the observed log negativity plateaus. This is because p_{T0} evaluated for the maximally mixed state is $1/2^{n(\Omega)} > 0$. This plot also shows that for small enough distribution error (small δ), the Trotterization/annealing errors dominate, and thus the distribution error can be ignored. **(b)** The fit parameter β for a few of the tested models, plotted against the total number of qubits in the graph N . The color of the points indicates the degeneracy of the ground state of the final Hamiltonians H_F ($n(gs)$). For all of the models we have $\beta > 0$, which indicates that eq. (13) is a safe bound. The darker colors correspond to higher degeneracy and also larger β . This correlation can be understood as follows: the larger the ground state space of H_F , the more is a state affected by noise to be measured within it. Interestingly, the sparse networks, which have more non-local connections appear to be less sensitive to noise in each connection, when compared to the spin chain models with not many non-local couplings.

it neglects contributions to the probability of distributions with at least one distribution fault. Having a $\beta > 0$ indicates that the process can still have a few distribution faults, without hindering the fidelity evolved state with the ground state space of H_F . An important point to note is that this is not a property inherent to the distributed quantum annealing protocol, as β it is problem dependent, as shown in fig. 6 (b). For instance, an annealing Hamiltonian with high ground state degeneracy will appear to be more tolerant to distribution faults (greater β) than an Hamiltonian with low ground state degeneracy. Unfortunately, we have not been able to determine β *a priori*, and therefore it is best to consider $\beta = 0$ when making estimations of the distribution fidelity.




FIG. 7. Annealing graphs for the different models considered in the comparison of annealing methods. The models labeled A_i correspond to the annealing models, where A_0 is the original annealing graph, a fully connected 4-vertex graph, and A_1 is its embedded version onto the square grid hardware topology. The graphs labeled D_i correspond to the distributed models, with D_0 corresponding to vertex splitting, and D_1 corresponds to edge splitting of the original graph A_0 .

Comparison with standard annealing methods

So far, we have been concerned with the characterization of the process, increasing the complexity of the simulated models while doing so. In this section, we will directly compare the embedding tricks used in quantum annealing with the distributed alternative. We restrict the allowed topology of the annealing (local) hardware to a preset topology. We denominate the original annealing problem by A_0 , whose coupling graph is the fully connected 4 vertex graph. The hardware embedded equivalent is A_1 , and has an additional parameter J_M that corresponds to the strength of the negative couplings that are needed to keep the qubits in the same chain correlated. Then, there are two ways to distribute A_0 : either by vertices (D_0), or by edges (D_1). Both require 4 non-local connections, but in D_0 they do not depend on the annealing problem, while in D_1 the strength of the non-local couplings is controlled by the annealing problem (encoded into H_F). These graphs are displayed in fig. 7.

While simulating the evolution of model A_1 we have observed that the system is harder to anneal (larger ε_{A_0} for a given t_F) as $-J_M$ increases. This is likely due to a reduction of the minimum energy gap $\Delta\epsilon$, increasing the annealing time t_F . With a reasonable energy error $\varepsilon_{A_0} \sim 10^{-2}$, the required annealing time for the models A_0, D_0 and D_1 is $t_F \sim 4$, while for A_1 we need $t_F > 10$ (see fig. 8). We assess that for our tested models, the annealing curves (ε_{A_0} vs t_F) of the distributed models D_i are equivalent to the annealing curve of the original model A_0 . The reason for this is that, for the edge splitting model D_1 , its annealing Hamiltonian ($H = H_L + H_N$) is precisely the same as the original




FIG. 8. Annealing curve comparing the annealing energy error ε_{A_0} for a wide range of annealing times t_F , for the original model A_0 , embedded model A_1 and vertex distributed model D_0 . It is also tried different values of J_M for the model A_1 . The blue dotted line corresponds to the model D_0 . The lines for A_0 and D_0 overlap, meaning that these models are equivalent when evolving via quantum annealing. For the model A_1 , as $-J_M$ is increased, the annealing problem becomes harder and harder, visible from the color gradient.

problem A_0 . This is not the case for the vertex duplication model D_0 , although we see an exact overlap of the annealing curves. Because of this we are confident that, as long as no qubit chains are employed, the distributed annealing problems have the same annealing curves as the original problem. From here, we then conclude that distributed quantum annealing without noise has the same time scaling as annealing with no embedding. This means that the distributed protocol offers no speedup when compared to a quantum annealer that can already natively embed the graph in its hardware.

When the hardware either does not have enough qubits, or requires long chains to embed an annealing problem, distributed quantum annealing can be a viable option. After finding an adequate graph split, the distribution fidelity $p_N(\delta)$ gives a lower bound for the probability of measuring a state that was evolved by annealing. Then, the average number of runs until the anneal state is measured is $[p_N(\delta)]^{-1}$. This means that the average time complexity (time that it takes to measure a state evolved by annealing) of the noisy protocol is $t_F [p_N(\delta)]^{-1}$. In addition, the average quantum communication complexity (how much of the entanglement resource is required by the protocol) is $M_D [p_N(\delta)]^{-1}$.

CONCLUSIONS

To summarize, we have introduced a protocol for distributed quantum annealing based on a combination of trotter decomposition and unitary distribution. We motivated the need for such a protocol by the increasing connectivity and size requirements that Quantum Annealing has to face in order to solve bigger and more complex problems. We showed that, using our protocol it is indeed possible to connect multiple annealers to

represent a bigger problem that wouldn't fit on a single machine and successfully recover the global ground state with high probability. We have also investigated the robustness of our protocol in the presence of non-ideal distribution of entanglement and proved that it could, indeed, work under realistic assumptions. We find that the number of non-local operations on each Trotter step and the number of Trotter steps M scale differently depending on how the graph is split, in the future it would be interesting to investigate if a recipe for optimal splitting can be achieved. The technology required to perform such distributions either in a local area network or at a larger scale are still in a proof of principle stage, but showing promise for the future. We thus conclude that distributed quantum annealing can be a viable way to augment the capabilities of a quantum annealer, when it cannot natively embed an annealing problem due to hardware limitations, either by its connectivity or by the number of required qubits. Having a cluster of quantum annealers connected with quantum communication channels could thus improve significantly the capability of solving larger and more complex optimization problems, as well as the possibility to explore optimization in a secure multi-party setting.

ACKNOWLEDGMENTS

The authors thank the support from FCT – Fundação para a Ciência e a Tecnologia (Portugal), namely through project UIDB/04540/2020, as well as from projects QuantHEP and HQCC supported by the EU QuantERA ERA-NET Cofund in Quantum Technologies and by FCT (QuantERA/0001/2019 and QuantERA/004/2021, respectively), from the EU Quantum Flagship projects QIA (820445) and QMiCS (820505). This work is also funded by the Advanced Computing/EuroCC MSc Fellows Programme, which is funded by EuroHPC under grant agreement No 951732.

Appendix A: Graph split: n -plication mechanism

The idea behind qubit n -plication in quantum annealing is to restrict the evolution of the n -plicated qubits in a way that they are kept aligned when measured. For instance, duplicating a qubit q means to transform the label $q \rightarrow q_1, q_2$ and the state

$$|\psi\rangle_q = \alpha |0\rangle_q + \beta |1\rangle_q \rightarrow |\psi\rangle_{q_1 q_2} = \alpha |00\rangle_{q_1 q_2} + \beta |11\rangle_{q_1 q_2} \quad (\text{A1})$$

This means that the annealing Hamiltonian $H(t)$ also has to be transformed. The transformed final Hamiltonian H_F has to be such that when the labels q_i are replaced by q , the original final Hamiltonian is obtained. For instance, the local field terms can be transformed as such

$$\sigma_q^z \rightarrow \frac{1}{n} \sum_{i=1}^n \sigma_{q_i}^z \quad (\text{A2})$$

For the starting Hamiltonian H_0 the transformation is different. The ground state of the starting Hamiltonian corresponds to the superposition of all possible states that can be a solution of the optimization posed by H_F . For the n -plicated qubits we want to restrict the solution to $\{|0\rangle^{\otimes n}, |1\rangle^{\otimes n}\}$, which means the initial state has to be

$$|GHZ_n^-\rangle = \frac{1}{\sqrt{2}} (|0\rangle^{\otimes n} - |1\rangle^{\otimes n}). \quad (\text{A3})$$

Then, for a system with the n -plicated qubit q , the starting Hamiltonian has to have the term

$$H_0 \sim \sigma_{q_1}^x \sigma_{q_2}^x \cdots \sigma_{q_n}^x = \prod_{i=1}^n \sigma_{q_i}^x. \quad (\text{A4})$$

This term has $|GHZ_n^-\rangle$ as one of the ground states. Because of this and because the action of $\sigma^z \sigma^z$ terms only change a relative phase in the superposition, the evolution by annealing will keep the n -plicated qubit to the space $\{|0\rangle^{\otimes n}, |1\rangle^{\otimes n}\}$.

Appendix B: Trotterization convergence details

Since the local and non-local Hamiltonians do not commute, i.e. $[H_L(t), H_N(t)] \neq 0$ for $t \in]0, t_F[$, the Trotterization process will not be ideal. In the end, we want the Trotterized evolution to approximate quantum annealing. At step k the evolution is given by

$$U_N^{(k)} U_L^{(k)} = \exp\{-i H_N(t_k) \Delta t\} \exp\{-i H_L(t_k) \Delta t\}. \quad (\text{B1})$$

The equivalent Hamiltonian $H_{eq}(t)$, when it exists, corresponds to the Hamiltonian that governs the evolution as if it was a continuous process governed by the Schrödinger equation. We want this equivalent Hamiltonian to converge to the annealing Hamiltonian $H(t)$ with increasingly smaller step sizes. The Baker-Campbell-Hausdorff expansion allows us to write an expression for $H_{eq}(t_k)$ at step k [35].

$$\begin{aligned} & -i H_{eq}(t_k) \Delta t = \\ & = \log \left[\exp\{-i H_N(t_k) \Delta t\} \cdot \exp\{-i H_L(t_k) \Delta t\} \right] \\ & = -i H(t_k) \Delta t + \sum_{n \geq 2} n^{-1} \sum_{|\omega|=n} g_\omega [\omega] \end{aligned} \quad (\text{B2})$$

where $\omega = \omega_0 \omega_1 \dots \omega_{n-1}$ is a word of length $n = |\omega|$ and ω_i can be X or Y where

$$\begin{aligned} X &= -i H_L(t_k) \Delta t, \\ Y &= -i H_N(t_k) \Delta t. \end{aligned}$$

$[\omega] = [[\omega_0, \omega_1], \dots, \omega_{n-1}]$ represents nested commutators of the word ω and g_ω is a scalar coefficient dependent on ω . In essence, we want that $H_{eq}(t)$ to approximate the annealing Hamiltonian $H(t)$. It then follows that we want to study the convergence properties of $\|H_{eq}(t) - H(t)\|$. With this aim, we use the relation eq. (B4) obtained by [35]. It makes use of the Lie norm, which satisfies eq. (B3).

$$\|[X, Y]\|_L \leq \|X\|_L \|Y\|_L \quad (\text{B3})$$

$$\left\| \frac{1}{n} \sum_{|\omega|=n} g_\omega [\omega] \right\|_L \leq \frac{2\mathcal{M}^n}{n} \quad (\text{B4})$$

where

$$\begin{aligned} \mathcal{M} &= \max(\|X\|_L, \|Y\|_L) \\ &= \Delta t \cdot \max(\|H_L(t_k)\|_L, \|H_N(t_k)\|_L) \\ &= \frac{\Delta t}{\Delta t_M} \end{aligned} \quad (\text{B5})$$

Consequently, we obtain a bound of the Hamiltonian difference at step k

$$\|H_{eq}(t_k) - H(t_k)\|_L \Delta t \leq \sum_{n \geq 2} \frac{2}{n} \left(\frac{\Delta t}{\Delta t_M} \right)^n \quad (\text{B6})$$

where $H(t) = H_L(t) + H_N(t)$ is the adiabatic Hamiltonian. Two observations can be made about this bound. The first is that for $\Delta t > \Delta t_M$, the absolute series diverges, meaning that the results of the Trotterization process cannot be predicted in this analysis. The series may conditionally converge for $\Delta t > \Delta t_M$. In the case that $H_{eq}(t)$ diverges, then the evolution cannot be described by the Schrödinger equation, and thus is not a continuous process. The second observation is that when $\Delta t / \Delta t_M \ll 1$, we have the bound

$$\|H_{eq}(t_k) - H(t_k)\|_L \leq \frac{\Delta t}{(\Delta t_M)^2} \quad (\text{B7})$$

Appendix C: Noisy distribution details

A generic distributed operation consists of a sequence of local unitaries and CNOT telegates. The noise is considered to come solely from the shared quantum state

$$\rho_\Phi = x |\Phi\rangle\langle\Phi| + \frac{1-x}{4} \mathbb{I}_\Phi,$$

where $|\phi\rangle = (|00\rangle + |11\rangle) / \sqrt{2}$ for the CNOT telegate. The CNOT distribution protocol shown in fig. 3 is equivalent to the circuit fig. 9, and so the CNOT telegate can be written as

$$\begin{aligned} \text{DCNOT}[\rho] &= \text{Tr}_\Phi \left[U_{\text{DCNOT}}(\rho \otimes \rho_\Phi) U_{\text{DCNOT}}^\dagger \right] \\ &= x \text{CNOT}[\rho] + (1-x) \Theta[\rho] \end{aligned} \quad (\text{C1})$$




FIG. 9. CNOT telegate circuit. It is equivalent to the one shown in fig. 3. This representation of the circuit allows us to write the CNOT telegate with eq. (C1).

with

$$\Theta[\rho] = \frac{1}{4} \text{Tr}_\Phi \left[U_{\text{DCNOT}}(\rho \otimes \mathbb{I}_\Phi) U_{\text{DCNOT}}^\dagger \right].$$

The $\text{CNOT}[\rho]$ operation corresponds to the state after ideal distribution, without noise, while $\Theta[\rho]$ is the state after a distribution fault occurs, i.e. $\rho_\Phi = \mathbb{I}_\Phi / 4$.

Consider $\mathbb{D}_k[\rho]$ the state that results from the distribution of the initial state ρ , with all possible k distribution faults in M_D distribution steps. We have that $\text{Tr} \mathbb{D}_k[\rho] = 1$. For example, a state ρ_2 which has to use 2 noisy DCNOT operations can be rewritten as

$$\rho_2 = x^2 \mathbb{D}_0[\rho_0] + 2x(1-x) \mathbb{D}_1[\rho_0] + (1-x)^2 \mathbb{D}_2[\rho_0]. \quad (\text{C2})$$

$$\mathbb{D}_0[\rho_0] = U_2 \circ \text{CNOT} \circ U_1 \circ \text{CNOT} \circ U_0[\rho_0]$$

$$\begin{aligned} \mathbb{D}_1[\rho_0] &= \frac{1}{2} U_2 \circ \Theta \circ U_1 \circ \text{CNOT} \circ U_0[\rho_0] \\ &\quad + \frac{1}{2} U_2 \circ \text{CNOT} \circ U_1 \circ \Theta \circ U_0[\rho_0] \end{aligned}$$

$$\mathbb{D}_2[\rho_0] = U_2 \circ \Theta \circ U_1 \circ \Theta \circ U_0[\rho_0]$$

with $U_i[\rho] = U_i \rho U_i^\dagger$ being an arbitrary unitary gate operation (noiseless) and $f \circ g[x] = f[g[x]]$ being the composition operation. Given that the probability of a perfect distributed operation (telegate) is given by x , then the state ρ_0 after M_D distribution steps is given by the binomial distribution:

$$\rho_{M_D} = \sum_{k=0}^{M_D} \binom{M_D}{k} x^{M_D-k} (1-x)^k \mathbb{D}_k[\rho_0] \quad (\text{C3})$$

Writing $F_\Phi = 1 - \delta$, we have $x = 1 - \frac{4}{3}\delta$. After M_D DCNOT operations, the fidelity between the states after noisy (ρ_{M_D}) and noiseless ($\mathbb{D}_0[\rho_0]$) distribution is $p_N(\delta)$, and is lower bounded by the first term in the sum

$$p_N(\delta) \geq x^{M_D} = \left(1 - \frac{4}{3}\delta \right)^{M_D}. \quad (\text{C4})$$

With $|\Psi\rangle\langle\Psi| = \mathbb{D}_0[\rho_0]$ being the state after ideal (noiseless) distribution, $p_N(\delta)$ is given by

$$p_N(\delta) = \langle\Psi|\rho_{M_D}|\Psi\rangle = \sum_{k=0}^{M_D} \binom{M_D}{k} x^{M_D-k} (1-x)^k f_k, \quad (\text{C5})$$

with $f_k = \langle\Psi|\mathbb{D}_k[\rho_0]|\Psi\rangle$ being a coefficient which depends on the number of faults in the distribution k . Assuming that successive faults reduce the probability of measuring $|\Psi\rangle$ by β , we get the ansatz

$$f_k = \beta^k, \quad \beta < 1 \quad (\text{C6})$$

where $\beta = \langle\Psi|\mathbb{D}_1[\rho_0]|\Psi\rangle$. With this ansatz, $p_N(\delta)$ can be reduced to a closed-form expression using the binomial theorem, and thereby obtaining eq. (16), reproduced here.

$$p_N(\delta; \beta) = \left(1 - \frac{4}{3}\delta(1-\beta)\right)^{M_D}$$

-
- [1] Bernhard H Korte, Jens Vygen, B Korte, and J Vygen. *Combinatorial optimization*, volume 1. Springer, 2011.
- [2] Vangelis Th Paschos. *Applications of combinatorial optimization*, volume 3. John Wiley & Sons, 2014.
- [3] F. Krzakala et al. *Statistical Physics, Optimization, Inference, and Message-Passing Algorithms*. Lecture Notes of the Les Houches School of Physics. 2013.
- [4] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi. Optimization by simulated annealing. *Science*, 220(4598):671–680, 1983. URL <http://science.sciencemag.org/content/220/4598/671>.
- [5] T. Kadowaki and H. Nishimori. Quantum annealing in the transverse ising model. *Physical Review E*, 58, 1998.
- [6] P. Magnard, S. Storz, P. Kurpiers, J. Schär, F. Marxer, J. Lütolf, T. Walter, J.-C. Besse, M. Gabureac, K. Reuer, A. Akin, B. Royer, A. Blais, and A. Wallraff. Microwave Quantum Link between Superconducting Circuits Housed in Spatially Separated Cryogenic Systems. 125(26):260502. ISSN 0031-9007, 1079-7114. doi: 10.1103/PhysRevLett.125.260502. URL <https://link.aps.org/doi/10.1103/PhysRevLett.125.260502>.
- [7] K. Stannigel, P. Rabl, A. S. Sørensen, P. Zoller, and M. D. Lukin. Opto-mechanical transducers for long-distance quantum communication. 105(22):220501, . ISSN 0031-9007, 1079-7114. doi: 10.1103/PhysRevLett.105.220501. URL <http://arxiv.org/abs/1006.4361>.
- [8] Yong Wan, Daniel Kienzler, Stephen D. Erickson, Karl H. Mayer, Ting Rei Tan, Jenny J. Wu, Hilma M. Vasconcelos, Scott Glancy, Emanuel Knill, David J. Wineland, Andrew C. Wilson, and Dietrich Leibfried. Quantum gate teleportation between separated qubits in a trapped-ion processor. 364(6443):875–878. doi: 10.1126/science.aaw9415. URL <https://www.science.org/doi/10.1126/science.aaw9415>.
- [9] W. Dür, R. Lamprecht, and S. Heusler. Towards a quantum internet. doi:10.1088/1361-6404/aa6df7.
- [10] Stephanie Wehner, David Elkouss, and Ronald Hanson. Quantum internet: A vision for the road ahead. 362(6412):eaam9288. doi:10.1126/science.aam9288. URL <https://www.science.org/doi/10.1126/science.aam9288>.
- [11] Daniele Cuomo, Marcello Caleffi, and Angela Sara Cacciapuoti. Towards a Distributed Quantum Computing Ecosystem. 1(1):3–8. ISSN 2632-8925, 2632-8925. doi:10.1049/iet-qtc.2020.0002. URL <http://arxiv.org/abs/2002.11808>.
- [12] Stephen DiAdamo, Marco Ghibaudi, and James Cruise. Distributed Quantum Computing and Network Control for Accelerated VQE. 2:1–21. ISSN 2689-1808. doi: 10.1109/TQE.2021.3057908.
- [13] Tameem Albash and Daniel A. Lidar. Adiabatic Quantum Computing. 90(1):015002. ISSN 0034-6861, 1539-0756. doi:10.1103/RevModPhys.90.015002. URL <http://arxiv.org/abs/1611.04471>.
- [14] Fred Glover, Gary Kochenberger, and Yu Du. Quantum Bridge Analytics I: A Tutorial on Formulating and Using QUBO Models. page 46.
- [15] Andrew Lucas. Ising formulations of many NP problems. 2. ISSN 2296-424X. doi:10.3389/fphy.2014.00005. URL <http://arxiv.org/abs/1302.5843>.
- [16] J. Eisert, K. Jacobs, P. Papadopoulos, and M. B. Plenio. Optimal local implementation of non-local quantum gates. 62(5):052317. ISSN 1050-2947, 1094-1622. doi: 10.1103/PhysRevA.62.052317. URL <http://arxiv.org/abs/quant-ph/0005101>.
- [17] Alexander S. Solntsev and Andrey A. Sukhorukov. Path-entangled photon sources on nonlinear chips. 2:19–31. ISSN 2405-4283. doi:10.1016/j.revip.2016.11.003. URL <https://www.sciencedirect.com/science/article/pii/S2405428316300168>.
- [18] K. Stannigel, P. Rabl, A. S. Sørensen, M. D. Lukin, and P. Zoller. Optomechanical transducers for quantum information processing. 84(4):042341, . ISSN 1050-2947, 1094-1622. doi:10.1103/PhysRevA.84.042341. URL <http://arxiv.org/abs/1106.5394>.
- [19] V. Singh, S. J. Bosman, B. H. Schneider, Y. M. Blanter, A. Castellanos-Gomez, and G. A. Steele. Optomechanical coupling between a multilayer graphene mechanical resonator and a superconducting microwave cavity. 9(10):820–824. ISSN 1748-3395. doi:10.1038/nnano.2014.168. URL <https://www.nature.com/articles/nnano.2014.168>.
- [20] Dong Yan, Zhi-Hai Wang, Chun-Nian Ren, Hang Gao, Yong Li, and Jin-Hui Wu. Duality and bistability in an optomechanical cavity coupled to a Rydberg superatom. 91(2):023813. doi:10.1103/PhysRevA.91.023813. URL <https://link.aps.org/doi/10.1103/PhysRevA.91.023813>.

- 91.023813.
- [21] Zhang-qi Yin, W. L. Yang, L. Sun, and L. M. Duan. Quantum network of superconducting qubits through an optomechanical interface. 91(1):012333. ISSN 1050-2947, 1094-1622. doi:10.1103/PhysRevA.91.012333. URL <https://link.aps.org/doi/10.1103/PhysRevA.91.012333>.
- [22] Ondřej Černotík and Klemens Hammerer. Measurement-Induced Long-Distance Entanglement of Superconducting Qubits using Optomechanical Transducers. 94(1):012340. ISSN 2469-9926, 2469-9934. doi:10.1103/PhysRevA.94.012340. URL <http://arxiv.org/abs/1512.00768>.
- [23] Mohammad Mirhosseini, Alp Sipahigil, Mahmoud Kalaei, and Oskar Painter. Superconducting qubit to optical photon transduction. 588(7839):599–603. ISSN 1476-4687. doi:10.1038/s41586-020-3038-6. URL <https://www.nature.com/articles/s41586-020-3038-6>.
- [24] Yiwen Chu and Simon Gröblacher. A perspective on hybrid quantum opto- and electromechanical systems. 117(15):150503. ISSN 0003-6951, 1077-3118. doi:10.1063/5.0021088. URL <http://arxiv.org/abs/2007.03360>.
- [25] Chun-Ling Zhang, Xiang Chen, Chang-Geng Liao, and Xiu-Min Lin. Reversible quantum state transfer in a three-mode optomechanical system. 18(6):065206. ISSN 1612-202X. doi:10.1088/1612-202X/abfe0d. URL <https://dx.doi.org/10.1088/1612-202X/abfe0d>.
- [26] Niccolò Fiaschi, Bas Hensen, Andreas Wallucks, Rodrigo Benevides, Jie Li, Thiago P. Mayer Alegre, and Simon Gröblacher. Optomechanical quantum teleportation. 15(11):817–821. ISSN 1749-4885, 1749-4893. doi:10.1038/s41566-021-00866-z. URL <http://arxiv.org/abs/2104.02080>.
- [27] Adrian Parra-Rodriguez, Pavel Lougovski, Lucas Lamata, Enrique Solano, and Mikel Sanz. Digital-Analog Quantum Computation. 101(2):022305. ISSN 2469-9926, 2469-9934. doi:10.1103/PhysRevA.101.022305. URL <http://arxiv.org/abs/1812.03637>.
- [28] D. V. Babukhin, A. A. Zhukov, and W. V. Pogosov. Hybrid digital-analog simulation of many-body dynamics with superconducting qubits. 101(5):052337. ISSN 2469-9926, 2469-9934. doi:10.1103/PhysRevA.101.052337. URL <http://arxiv.org/abs/1909.10732>.
- [29] W. K. Wootters and W. H. Zurek. A single quantum cannot be cloned. 299(5886):802–803. ISSN 1476-4687. doi:10.1038/299802a0. URL <https://www.nature.com/articles/299802a0>.
- [30] Davide Ferrari, Angela Sara Cacciapuoti, Michele Amoretti, and Marcello Caleffi. Compiler Design for Distributed Quantum Computing. 2:1–20. ISSN 2689-1808. doi:10.1109/TQE.2021.3053921. URL <https://ieeexplore.ieee.org/document/9334411/>.
- [31] Luca D’Alessio and Anatoli Polkovnikov. Many-body energy localization transition in periodically driven systems. 333:19–33. ISSN 0003-4916. doi:10.1016/j.aop.2013.02.011. URL <https://www.sciencedirect.com/science/article/pii/S0003491613000389>.
- [32] Tomotaka Kuwahara, Takashi Mori, and Keiji Saito. Floquet-Magnus Theory and Generic Transient Dynamics in Periodically Driven Many-Body Quantum Systems. 367:96–124. ISSN 00034916. doi:10.1016/j.aop.2016.01.012. URL <http://arxiv.org/abs/1508.05797>.
- [33] Markus Heyl, Philipp Hauke, and Peter Zoller. Quantum localization bounds Trotter errors in digital quantum simulation. 5(4):eaau8342. doi:10.1126/sciadv.aau8342. URL <https://www.science.org/doi/10.1126/sciadv.aau8342>.
- [34] Supanut Thanasilp, Jirawat Tangpanitanon, Marc-Antoine Lemonde, Ninnat Dangniam, and Dimitris G. Angelakis. Quantum supremacy and quantum phase transitions. 103(16):165132. ISSN 2469-9950, 2469-9969. doi:10.1103/PhysRevB.103.165132. URL <http://arxiv.org/abs/2012.06459>.
- [35] Robert C. Thompson. Convergence proof for Goldberg’s exponential series. 121:3–7. ISSN 0024-3795. doi:10.1016/0024-3795(89)90688-5. URL <https://www.sciencedirect.com/science/article/pii/0024379589906885>.