ARTICLE OPEN

Nested quantum annealing correction

Walter Vinci[1,2,3], Tameem Albash[2,3,4] and Daniel A Lidar[1,2,3,5]

We present a general error-correcting scheme for quantum annealing that allows for the encoding of a logical qubit into an arbitrarily large number of physical qubits. Given any Ising model optimization problem, the encoding replaces each logical qubit by a complete graph of degree C, representing the distance of the error-correcting code. A subsequent minor-embedding step then implements the encoding on the underlying hardware graph of the quantum annealer. We demonstrate experimentally that the performance of a D-Wave Two quantum annealing device improves as C grows. We show that the performance improvement can be interpreted as arising from an effective increase in the energy scale of the problem Hamiltonian or, equivalently, an effective reduction in the temperature at which the device operates. The number C thus allows us to control the amount of protection against thermal and control errors, and, in particular, to trade qubits for a lower effective temperature that scales as $C^{-\eta}$, with $\eta \leqslant 2$. This effective temperature reduction is an important step towards scalable quantum annealing.

INTRODUCTION

Quantum annealing (QA) attempts to exploit quantum fluctuations to solve computational problems faster than it is possible with classical computers.[1–7] As an approach designed to solve optimisation problems, QA is a special case of adiabatic quantum computation (AQC),[8] a universal model of quantum computing.[9–12] In AQC, a system is designed to follow the instantaneous ground state of a time-dependent Hamiltonian whose final ground state encodes the solution to the problem of interest. This results in a certain amount of stability, as the system can thermally relax to the ground state after an error, as well as resilience to errors, as the presence of a finite energy gap suppresses thermal and dynamical excitations.[13–18]

Despite this inherent robustness to certain forms of noise, AQC requires error correction to ensure scalability, just like any other form of quantum information processing.[19] Various error-correction proposals for AQC and QA have been made,[20–33] but an accuracy-threshold theorem for AQC is not yet known, unlike in the circuit model (e.g., ref. 34). A direct AQC simulation of a fault-tolerant quantum circuit leads to many-body (high weight) operators that are difficult to implement[23,24] or myriad other problems.[12] Nevertheless, a scalable method to reduce the effective temperature would go a long way towards approaching the ideal of closed-system AQC, where quantum speedups are known to be possible.[9,35–37]

Motivated by the availability of commercial QA devices featuring hundreds of qubits,[38–41] we focus on error correction for QA. There is a consensus that these devices are significantly and adversely affected by decoherence, noise and control errors,[42–49] which makes them particularly interesting for the study of tailored, practical error-correction techniques. Such techniques, known as quantum annealing correction (QAC) schemes, have already been experimentally shown to significantly improve the performance of quantum annealers,[26,30–32] and they are theoretically analysed using a mean-field approach.[33] However, these QAC schemes are not easily generalisable to arbitrary optimisation problems, as they induce an encoded graph

that is typically of a lower degree than the qubit-connectivity graph of the physical device. Moreover, they typically impose a fixed code distance, which limits their efficacy.

To overcome these limitations, here we present a family of error-correcting codes for QA, based on a 'nesting' scheme, that has the following properties: (1) it can handle arbitrary Ising model optimisation problem; (2) it can be implemented on present-day QA hardware; and (3) it is capable of an effective temperature reduction controlled by the code distance. Our 'nested quantum annealing correction' (NQAC) scheme thus provides a very general and practical tool for error correction in quantum optimisation.

We test NQAC by studying antiferromagnetic complete graphs numerically, as well as on a D-Wave Two (DW2) processor featuring 504 flux qubits connected by 1,427 tunable composite qubits acting as Ising-interaction couplings, arranged in a non-planar Chimera-graph lattice[50] (complete graphs were also studied for a spin glass model in ref. 51). We demonstrate that our encoding schemes yield a steady improvement for the probability of reaching the ground state as a function of the nesting degree, even after minor-embedding the complete graph onto the physical graph of the quantum annealer. We also demonstrate that NQAC outperforms classical repetition code schemes that use the same number of physical qubits.

RESULTS

Nested quantum annealing correction

In QA, the system undergoes an evolution governed by the following time-dependent, transverse-field Ising Hamiltonian:

$$H(t) = A(t)H_X + B(t)H_P, \quad t \in [0, t_f], \tag{1}$$

with, respectively, monotonically decreasing and increasing 'annealing schedules' $A(t)$ and $B(t)$. The 'driver Hamiltonian' $H_X = -\sum_i \sigma_i^x$ is a transverse field whose amplitude controls the tunnelling rate. The solution to an optimisation problem of interest is encoded in the ground state of the Ising-problem Hamiltonian H_P, with

[1]Department of Electrical Engineering, University of Southern California, Los Angeles, CA, USA; [2]Department of Physics and Astronomy, University of Southern California, Los Angeles, CA, USA; [3]Center for Quantum Information Science & Technology, University of Southern California, Los Angeles, CA, USA; [4]Information Sciences Institute, University of Southern California, Marina del Rey, CA, USA and [5]Department of Chemistry, University of Southern California, Los Angeles, CA, USA.
Correspondence: W Vinci (wltvinci@gmail.com)

$$H_\mathrm{P} = \sum_{i \in \mathcal{V}} h_i \sigma_i^z + \sum_{(i,j) \in \varepsilon} J_{ij} \sigma_i^z \sigma_j^z, \qquad (2)$$

where the sums run over the weighted vertices $\mathcal{V}$ and edges ε of a graph $G = (\mathcal{V}, \varepsilon)$, and $\sigma_i^{x,z}$ denotes the Pauli operators acting on qubit i. The D-Wave devices use an array of superconducting flux qubits to physically realise the system described in Equations (1) and (2) on a fixed 'Chimera' graph (Figure 1) with programmable local fields $\{h_i\}$, couplings $\{J_{ij}\}$ and annealing time t_f.[39–41]

For closed systems, so-called eigenstate steering[52] or shortcuts to adiabaticity[53] can keep a system in its ground state, but they typically require either knowledge of the final ground state or the addition of auxiliary counterdiabatic fields whose determination may be overly complicated. The adiabatic theorem for closed systems[54,55] guarantees that if the system is initialised in the ground state of $H(0) = A(0)H_X$, a sufficiently slow evolution relative to the inverse minimum gap of $H(t)$ will take the system with high probability to the ground state of the final Hamiltonian $H(t_f) = B(t_f)H_\mathrm{P}$. Dynamical errors then arise because of diabatic transitions, but they can be made arbitrarily small via boundary cancellation methods that control the smoothness of $A(t)$ and $B(t)$, as long as the adiabatic condition is satisfied.[56–58] Specifically,

it was shown in ref. 58 that for annealing functions with infinite smoothness (belonging to the so-called Gevrey class) the error can be made exponentially small in the total runtime t_f. This means that in particular the probability of Landau–Zener transitions is exponentially suppressed, although of course t_f is still controlled by an inverse (cubic) power of the minimum gap. We shall assume that the problem of Landau–Zener transitions is addressed by such boundary cancellation methods (although the experiments we describe below do not include such methods, as we had no control over the smoothness of $A(t)$ and $B(t)$ and focus here on addressing the errors that occur in open systems. For the latter, specifically a system that is weakly coupled to a thermal environment, the final state is a mixed-state $\rho(t_f)$ that is close to the Gibbs state associated with $H(t_f)$ if equilibrium is reached throughout the annealing process.[18,59,60] In the adiabatic limit, the open-system QA process is thus better viewed as a Gibbs distribution sampler. The main goal of QAC is to suppress the associated thermal errors and to restore the ability of QA to act as a ground-state solver. In addition, QAC should suppress errors due to noise-driven deviations in the specification of H_P.[25]

Error correction is achieved in QAC by mapping the logical Hamiltonian $H(t)$ to an appropriately chosen encoded
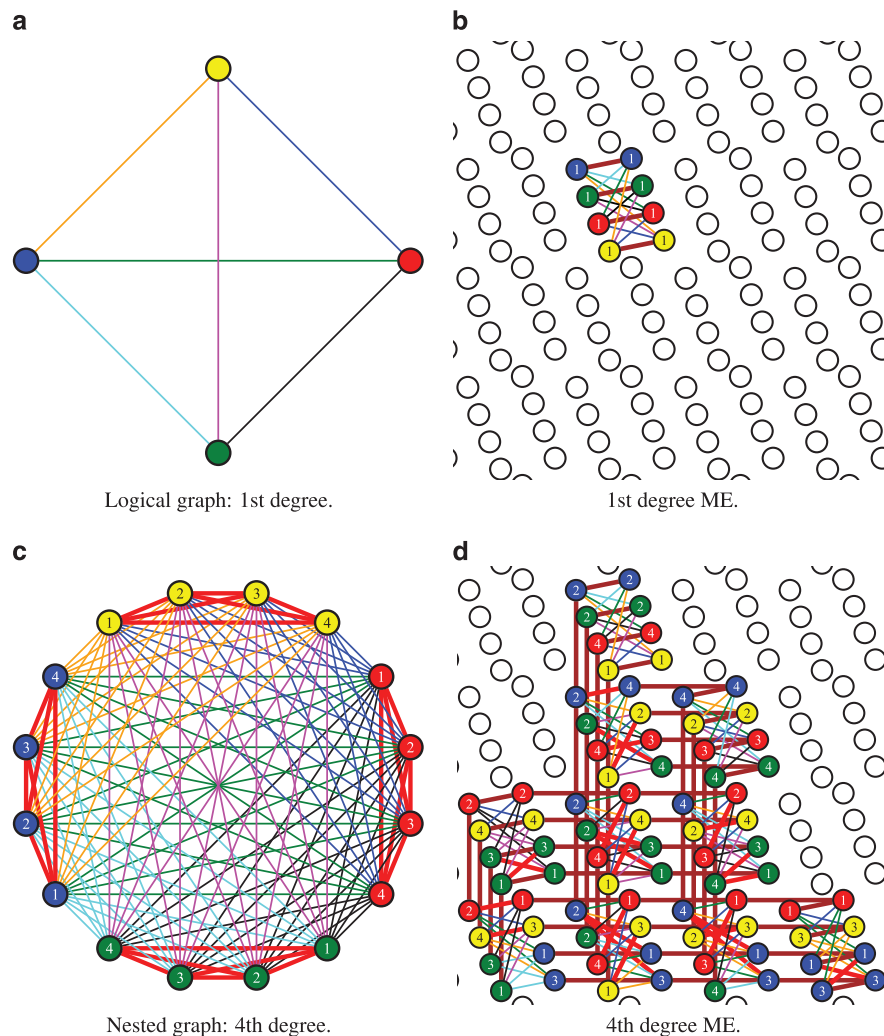
Figure 1. Illustration of the nesting scheme. In the left column, a C-degree nested graph is constructed by embedding a K_N into a $K_{C \times N}$, with $N = 4$ and $C = 1$ (**a** and **b**) and $C = 4$ (**c** and **d**). Red, thick couplings are energy penalties defined on the nested graph between the (i, c) nested copies of each logical qubit i. The right column shows the nested graphs after ME on the DW2 Chimera graph. Brown, thick couplings correspond to the ferromagnetic chains introduced in the process.

Hamiltonian $\overline{H}(t)$:

$$\overline{H}(t) = A(t)H_X + B(t)\overline{H}_P, \quad t \in [0, t_f],$$ (3)

defined over a set of physical qubits $\overline{N}$ larger than the number of logical qubits $N = |\mathcal{V}|$. Note that $\overline{H}_P$ also includes penalty terms, as explained below. The logical ground state of H_P is extracted from the encoded system's state $\overline{\rho}(t_f)$ through an appropriate decoding procedure. A successful error-correction scheme should recover the logical ground state with a higher probability than a direct implementation of H_P or with a higher probablilty than a classical repetition code using the same number of physical qubits $\overline{N}$. Because of practical limitations of current QA devices that prevent the encoding of H_X, only H_P is encoded in QAC. In the future, it may be possible to circumvent this limitation using coupling to ancilla qubits.[61] At present, it results in a tradeoff, as it requires us to optimise the penalty strength, and may also result in a need to optimise the nesting degree, as without encoding H_X the minimum gap may shrink relative to the unencoded problem.

To allow for the most general N-variable Ising optimisation problem, we now define an encoding procedure for problem Hamiltonians H_P supported on a complete graph K_N. The first step of our construction involves a 'nested' Hamiltonian $\tilde{H}_P$ that is defined by embedding the logical K_N into a larger $K_{C \times N}$. The integer C is the 'nesting degree', and it controls the amount of hardware resources (qubits, couplers and local fields) used to represent the logical problem. $\tilde{H}_P$ is constructed as follows. Each logical qubit i ($i = 1, \ldots, N$) is represented by a C-tuple of encoded qubits (i, c), with $c = 1, \ldots, C$. The 'nested' couplings $\tilde{J}_{(i,c),(j,c')}$ and local fields $\tilde{h}_{(i,c)}$ are then defined as follows:

$$\tilde{J}_{(i,c),(j,c')} = J_{ij}, \quad \forall c, c', i \neq j,$$ (4a)

$$\tilde{h}_{(i,c)} = Ch_i, \quad \forall c, i,$$ (4b)

$$\tilde{J}_{(i,c),(i,c')} = -\gamma, \quad \forall c \neq c'.$$ (4c)

This construction is illustrated in the left column of Figure 1. Each logical coupling J_{ij} has C^2 copies $\tilde{J}_{(i,c),(j,c')}$, thus boosting the energy scale at the encoded level by a factor of C^2. Each local field h_i has C copies $\tilde{h}_{(i,c)}$; the factor C in Equation (4b) ensures that the energy boost is equalised with the couplings. For each logical qubit i, there are $C(C-1)/2$ ferromagnetic couplings $\tilde{J}_{(i,c),(i,c')}$ of strength $\gamma > 0$ (to be optimised), representing energy penalties that promote agreement among the C encoded qubits, i.e., that bind the C-tuple as a single logical qubit i.

The second step of our construction is to implement the fully connected problem $\tilde{H}_P$ on the given QA hardware, with a lower-degree qubit-connectivity graph. This requires a minor embedding (ME).[62–66] The procedure involves replacing each qubit in $\tilde{H}_P$ by a ferromagnetically coupled chain of qubits, such that all couplings in $\tilde{H}_P$ are represented by inter-chain couplings. The intra-chain coupling represents another energy penalty that forces the chain qubits to behave as a single logical qubit. The physical Hamiltonian obtained after this ME step is the final encoded Hamiltonian $\overline{H}_P$. We can minor-embed a $K_{C \times N}$ nested graph representing each qubit (i, c) as a physical chain of length $L = \lceil CN/4 \rceil + 1$ on the Chimera graph.[62] This is illustrated in the right column of Figure 1. The number of physical qubits necessary for a ME of a $K_{C \times N}$ is $N_C^{phys} = CNL \sim C^2N^2/4$.

At the end of a QA run implementing the encoded Hamiltonian $\overline{H}_P$ and a measurement of the physical qubits, a decoding procedure must be used to recover the logical state. For the sake of simplicity, we only consider majority vote decoding over both the length-L chain of each encoded qubit (i, c) and the C-encoded qubits comprising each logical qubit i (decoding over the length-L chain first, and then over the C-encoded qubits, does not affect performance; Supplementary Information S1). The encoded and logical qubits can thus be viewed as forming repetition codes

with, respectively, distance L and C. Other decoding strategies are possible wherein the encoded or logical qubits do not have this simple interpretation: e.g., energy minimisation decoding, which tends to outperform majority voting.[31] In the unlikely event of a tie, we assign a random value of $+1$ or -1 to the logical qubit.

Free energy

Using a mean-field analysis that reduces the model to an equivalent classical one by using the Suzuki-Trotter formula (see refs 67,68 for an early similar analysis of the Sherrington–Kirkpatrick model in a transverse field), and similar to the approach pursued in ref. 33 we can compute the partition function associated with the nested Hamiltonian $A(t)H_X + B(t)\tilde{H}_P$ for the case with uniform antiferromagnetic couplings. This leads to the following free-energy density in the low temperature and thermodynamic limits (Supplementary Information S2):

$$\beta F = C^2 \beta \left(\sqrt{[A(t)/C]^2 + [2\gamma B(t)m]^2} - \gamma B(t)m^2 \right)$$ (5)

where m is the mean-field magnetisation. There are two key noteworthy aspects of this result. First, the driver term is rescaled as $A(t) \mapsto C^{-1}A(t)$. This shifts the crossing between the A and B annealing schedules to an earlier point in the evolution and is related to the fact that QAC encodes only the problem Hamiltonian term proportional to $B(t)$. Consequently, the quantum critical point is moved to earlier in the evolution, which benefits QAC, as the effective energy scale at this new point is higher.[33] Second, the inverse temperature is rescaled as $\beta \mapsto C^2\beta$. This corresponds to an effective temperature reduction by C^2, a manifestly beneficial effect. The same conclusion, of a lower effective temperature, is reached by studying the numerically computed success probability associated with thermal distributions (Supplementary Information S3). We shall demonstrate that this prediction is born out by our experimental results, although it is masked to some extent by complications arising from the ME and noise.

NQAC results

The hardness of an Ising optimisation problem, using a QA device, is controlled by its size N, as well as by an overall energy scale α.[45] The smaller this energy scale, the higher the effective temperature and the more susceptible QA becomes to (dynamical and thermal) excitations out of the ground state and misspecification noise on the problem Hamiltonian. This provides us with an opportunity to test NQAC. As in our experiments we were limited by the largest complete graph that can be embedded on the DW2 device, a K_{32} (see Supplementary Information S4 for details), we tuned the hardness of a problem by studying the performance of NQAC as a function of α via $H_P \mapsto \alpha H_P$, with $0 < \alpha \leq 1$. Note that we did not rescale γ; instead, γ was optimised for optimal post-decoding performance (see Supplementary Information S5). It is known that, for the DW2, intrinsic coupler control noise can be taken to be Gaussian with s.d. $\sigma \sim 0.05$ of the maximum value for the couplings.[48] Thus, we may expect that, without error correction, Ising problems with $\alpha \lesssim 0.05$ are dominated by control noise.

We applied NQAC to completely antiferromagnetic ($h_i = 0 \; \forall i$) Ising problems over K_4 ($J_{ij} = 1 \; \forall i, j$), and K_8 (random $J_{ij} \in [0.1, 1]$ with steps of 0.1) with nesting up to $C = 8$ and $C = 4$, respectively. We denote by $P_C(\alpha)$ the probability to obtain the logical ground state at energy scale α for the C-degree nested implementation (see Supplementary Information S1 for the data collection methods). The experimental QA data in Figure 2 (left) shows a monotonic increase of $P_C(\alpha)$ as a function of the nesting degree C over a wide range of energy scales α. As expected, $P_C(\alpha)$ drops from $P_C(1) = 1$ (solution always found) to $P_C(0) = 6/16$ (random
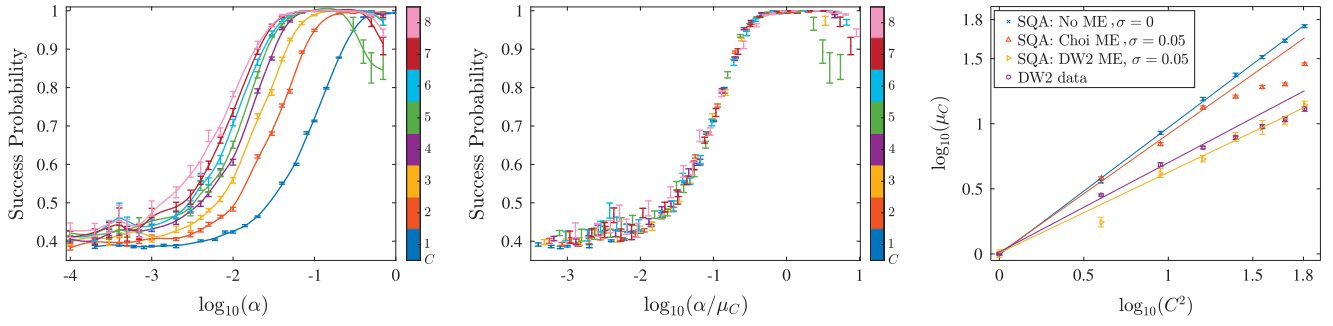
Figure 2. Experimental and numerical results for the antiferromagnetic K_4, after encoding, followed by ME and decoding. Left: DW2 success probabilities $P_C(\alpha)$ for eight nesting degrees C. Increasing C generally increases $P_C(\alpha)$ at fixed α. Middle: rescaled $P_C(\alpha\mu_C)$ data, exhibiting data collapse. Right: scaling of the energy boost μ_C versus the maximal energy boost μ_C^{max}, for both the DW2 and SQA. Purple circles: DW2 results. Blue stars: SQA for the case of no ME (i.e., for the problem defined directly over $K_{C\times N}$ and no coupler noise). Red up-triangles: SQA for the Choi ME[62] (for a full Chimera graph), with $\sigma = 0.05$ Gaussian noise on the couplings. Yellow right-triangles: SQA for the DW2 heuristic ME[63,64] (applied to a Chimera graph with eight missing qubits) with $\sigma = 0.05$ Gaussian noise on the couplings. The flattening of μ_C suggests that the energy boost becomes less effective at large C. However, this can be remedied by increasing the number of SQA sweeps (Supplementary Information S3), fixed here at 10^4. Thus, the lines represent best fits to only the first four data points, with slopes 0.98, 0.91, 0.62 and 0.69, respectively. In all panels, $N_{phys} \in [8,288]$.
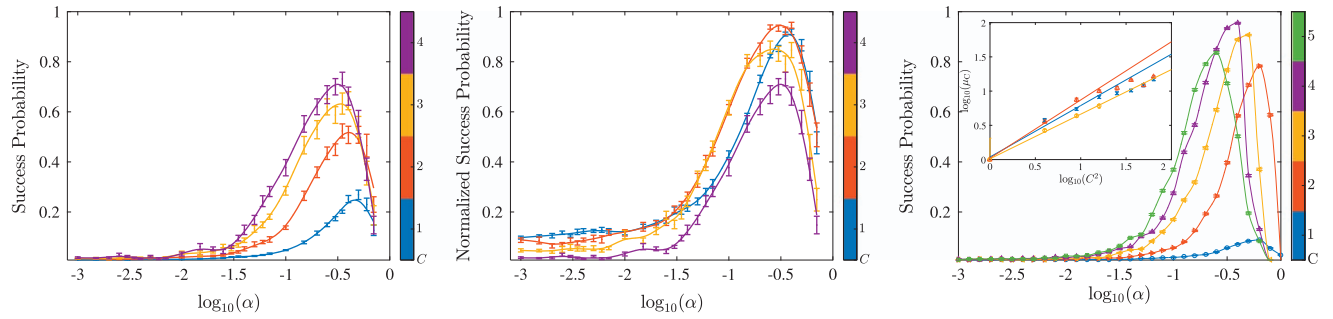
Figure 3. Random antiferromagnetic K_8: experimental and numerical results. Left: success probabilities $P_C(\alpha)$ for four nesting degrees. Middle: success probabilities $P'_C(\alpha)$ adjusted for classical repetition. Right: numerical results for SQA simulations with 20,000 sweeps, $\sigma = 0.05$ Gaussian noise on the couplings and with the Choi embedding, showing five nesting degrees. Inset: scaling of the energy boost μ_C versus the maximal energy boost μ_C^{max}, for both the DW2 and SQA. Yellow circles: DW2 results. Blue crosses and red up-triangles: SQA for the Choi ME with 10,000 (crosses) and 20,000 (up-triangles) sweeps, and with $\sigma = 0.05$ Gaussian noise on the couplings. The flattening of μ_C for $C > 4$ suggests that the energy boost becomes less effective at large C, but increasing the number of sweeps recovers the effectiveness. The lines represent best fits to only the first four data points, with respective slopes $\eta/2 = 0.65$, 0.75 and 0.85.

sampling of 6 ground states, where 4 out of the 6 couplings are satisfied, out of a total of 16 states).

Note that $P_1(\alpha)$ (no nesting) drops by $\sim 50\%$ when $\alpha \sim 0.1$, which is consistent with the aforementioned $\sigma \sim 0.05$ control noise level, whereas $P_8(\alpha)$ exhibits a similar drop only when $\alpha \sim 0.01$. This suggests that NQAC is particularly effective in mitigating the two dominant effects that limit the performance of quantum annealers: thermal excitations and control errors. To investigate this more closely, the middle panel of Figure 2 shows that the data from the left panel can be collapsed via $P_C(\alpha) \mapsto P_C(\alpha/\mu_C)$, where μ_C is an empirical rescaling factor discussed below (see also Supplementary Information S6). This implies that $P_1(\mu_C\alpha) \approx P_C(\alpha)$, and hence that the performance enhancement obtained at nesting degree C can be interpreted as an energy boost $\alpha \mapsto \mu_C\alpha$ with respect to an implementation without nesting.

The existence of this energy boost is a key feature of NQAC, as anticipated above. Recall (Equations (4a)–(4c)) that a nested graph $K_{C\times N}$ contains C^2 equivalent copies of the same logical coupling J_{ij}. Hence, a degree-C nesting before ME can provide a maximal energy boost $\mu_C^{max} = C^{\eta^{max}}$, with $\eta^{max} = 2$. This simple argument agrees with the reduction of the effective temperature by C^2 based on the calculation of the free energy (Equation (5)). The right panel of Figure 2 shows μ_C as a function of μ_C^{max}, yielding $\mu_C \sim C^\eta$ with $\eta \approx 1.37$ (purple circles). To understand why $\eta < \eta^{max}$, we

performed simulated quantum annealing (SQA) simulations (see Supplementary Information S7 for details). We observe in Figure 2 (right) that without ME and control errors the boost scaling matches μ_C^{max} (blue stars). When including ME and control errors, a performance drop results (red triangles). Both factors thus contribute to the sub-optimal energy boost observed experimentally. However, the optimal energy boost is recovered for a fully thermalised state with a sufficiently large penalty (see Supplementary Information S3). To match the experimental DW2 results using SQA, we replace the Choi ME designed for full Chimera graphs[62] by the heuristic ME designed for Chimera graphs with missing qubits,[63,64] and achieve a near match (yellow triangles) (see Supplementary Information S4 for more details on ME).

Performance of NQAC versus classical repetition

Recall that $N_C^{phys} = CNL$ is the total number of physical qubits used at nesting degree C; let C_{max} denote the highest nesting degree that can be accommodated on the QA device for a given K_N—i.e., $C_{max}NL \leq N_{tot} < (C_{max}+1)NL$, where N_{tot} is the total number of physical qubits (504 in our experiments). Then, $M_C = \lfloor N_{C_{max}}^{phys}/N_C^{phys} \rfloor$ is the number of copies that can be implemented in parallel. For NQAC at degree C to be useful, it must be

more effective than a classical repetition scheme in which M_C copies of the problem are implemented in parallel. If a single implementation has success probability $P_C(\alpha)$, the probability to succeed at least once with M_C statistically independent implementations is $P'_C(\alpha) = 1 - [1 - P_C(\alpha)]^{M_C}$. It turns out that the antiferromagnetic K_4 problem, for which a random guess succeeds with probability 6/16, is too easy (i.e., $P'_C(\alpha)$ approaches 1 too rapidly), and we therefore consider a harder problem: an antiferromagnetic K_8 instance with couplings randomly generated from the set $J_{ij} \in \{0.1, 0.2, \ldots, 0.9, 1\}$ (see Supplementary Information S5 for more details and the data on this and additional instances). Problems of this type turn out to have a sufficiently low success probability for our purposes, and can still be nested up to $C = 4$ on the DW2 processor.

Results for $P_C(\alpha)$ are shown in Figure 3 (left), and again increase monotonically with C, as in the K_4 case. For each C, $P_C(\alpha)$ peaks at a value of α, for which the maximum allowed strength of the energy penalties $\gamma = 1$ is optimal ($\gamma > 1$ would be optimal for larger α, as shown in Supplementary Information S5; the growth of the optimal penalty with problem size, and hence chain length, is a typical feature of minor-embedded problems[51]). An energy-boost interpretation of the experimental data of Figure 3 is possible for α values to the left of the peak; to the right of the peak, the performance is hindered by the saturation of the energy penalties.

Figure 3 (middle) compares the success probabilities $P'_C(\alpha)$ adjusted for classical repetition, where we have set $C_{\max} = 4$, and it shows that $P'_2(\alpha) > P'_1(\alpha)$—i.e., even after accounting for classical parallelism, $C = 2$ performs better than $C = 1$. However, we also find that $P'_4(\alpha) < P'_3(\alpha) \lesssim P'_2(\alpha)$, and thus no additional gain results from increasing C in our experiments. This can be attributed to the fact that even the K_8 problem still has a relatively large $P_1(\alpha)$. Experimental tests on QA devices with more qubits will thus be important to test the efficacy of higher nesting degrees on harder problems.

To test the effect of increasing C, and also to study the effect of varying the annealing time, we present in Figure 3 (right) the performance of SQA on a random K_8 antiferromagnetic instance with the Choi ME. The results are qualitatively similar to those observed on the DW2 processor with the heuristic ME (Figure 3 (left)). Interestingly, we observe a drop in the peak performance at $C = 5$ relative to the peak observed for $C = 4$. We attribute this to both a saturation of the energy penalties and a sub-optimal number of sweeps. The latter is confirmed in Figure 3 (right, inset), where we observe that the scaling of μ_C with C is better for the case with more sweeps—i.e., again $\mu_C \sim C^\eta$, and η increases with the number of sweeps.

DISCUSSION

Nested QAC offers several significant improvements over previous approaches to the problem of error correction for QA. It is a flexible method that can be used with any optimisation problem, and it allows the construction of a family of codes with arbitrary code distance. We have given experimental and numerical evidence that nesting is effective by performing studies with a D-Wave QA device and numerical simulations. We have demonstrated that the protection from errors provided by NQAC can be interpreted as arising from an increase (with nesting degree C) in the energy scale at which the logical problem is implemented. This represents a very useful tradeoff: the effective temperature drops as we increase the number of qubits allocated to the encoding, and thus that these two resources can be traded. Thus, NQAC can be used to combat thermal excitations, which are the dominant sources of errors in open-system QA, and are the bottleneck for scalable QA implementations, assuming that closed-system Landau–Zener transitions have been suppressed using other methods. We have also demonstrated that an appropriate nesting degree can outperform classical repetition with the same number of qubits, with improvements to be expected when next-generation QA devices with larger numbers of physical qubits become available. We, therefore, believe that our results are of immediate and near-future practical use, and constitute an important step towards scalable QA.

METHODS

For more details on the experimental and numerical methodologies used, as well as for details on the mean-field calculation of the free energy as in Equation 5; see Supplementary Information.

COMPETING INTERESTS

The authors declare no conflict of interest.

REFERENCES

1. Kadowaki, T. & Nishimori, H. Quantum annealing in the transverse Ising model. *Phys. Rev. E* **58**, 5355 (1998).
2. Brooke, J., Bitko, D., F., Rosenbaum, T. & Aeppli, G. Quantum Annealing of a Disordered Magnet. *Science* **284**, 779–781 (1999).
3. Brooke, J., Rosenbaum, T. F. & Aeppli, G. Tunable quantum tunnelling of magnetic domain walls. *Nature* **413**, 610–613 (2001).
4. Farhi, E. *et al.* A Quantum Adiabatic Evolution Algorithm Applied to Random Instances of an NP-Complete Problem. *Science* **292**, 472–475 (2001).
5. Morita, S. & Nishimori, H. Mathematical foundation of quantum annealing. *J. Math. Phys* **49**, 125210–125247 (2008).
6. Das, A. & Chakrabarti, B. K. *Colloquium*: Quantum annealing and analog quantum computation. *Rev. Mod. Phys.* **80**, 1061–1081 (2008).
7. Suzuki B. S. & Das A. guest eds. Discussion and Debate - Quantum Annealing: The Fastest Route to Quantum Computation? *Eur. Phys. J. Spec. Top* **224**, 75–88 (2015).
8. Farhi, E., Goldstone, J., Gutmann, S. & Sipser, M. Quantum Computation by Adiabatic Evolution. http://arxiv.org/abs/quant-ph/0001106 (2000)
9. Aharonov, D. *et al.* Adiabatic Quantum Computation is Equivalent to Standard Quantum Computation. *SIAM J. Comput.* **37**, 166–194 (2007).
10. Mizel, A., Lidar, D. A. & Mitchell, M. Simple Proof of Equivalence between Adiabatic Quantum Computation and the Circuit Model. *Phys. Rev. Lett.* **99**, 070502 (2007).
11. Gosset, D., Terhal, B. M. & Vershynina, A. Universal Adiabatic Quantum Computation via the Space-Time Circuit-to-Hamiltonian Construction. *Phys. Rev. Lett.* **114**, 140501 (2015).
12. Lloyd, S. & Terhal, B. Adiabatic and Hamiltonian computing on a 2D lattice with simple 2-qubit interactions. http://arXiv.org/abs/1509.01278 (2015).
13. Childs, A. M., Farhi, E. & Preskill, J. Robustness of adiabatic quantum computation. *Phys. Rev. A* **65**, 012322 (2001).
14. Sarandy, M. S. & Lidar, D. A. Adiabatic Quantum Computation in Open Systems. *Phys. Rev. Lett.* **95**, 250503 (2005).
15. Amin, M. H. S., Love, P. J. & Truncik, C. J. S. Thermally Assisted Adiabatic Quantum Computation. *Phys. Rev. Lett.* **100**, 060503 (2008).
16. Lloyd, S. Robustness of Adiabatic Quantum Computing. http://arXiv.org/abs/0805.2757 (2008).
17. Amin, M. H. S., Averin, D. V. & Nesteroff, J. A. Decoherence in adiabatic quantum computation. *Phys. Rev. A* **79**, 022107 (2009).
18. Albash, T. & Lidar, D. A. Decoherence in adiabatic quantum computation. *Phys. Rev. A* **91**, 062320 (2015).
19. Lidar D. & Brun T. (eds.) *Quantum Error Correction* (Cambridge University Press, 2013).
20. Jordan, S. P., Farhi, E. & Shor, P. W. Error-correcting codes for adiabatic quantum computation. *Phys. Rev. A* **74**, 052322 (2006).
21. Lidar, D. A. Towards Fault Tolerant Adiabatic Quantum Computation. *Phys. Rev. Lett.* **100**, 160506 (2008).

22. Quiroz, G. & Lidar, D. A. High-fidelity adiabatic quantum computation via dynamical decoupling. *Phys. Rev. A* **86**, 042333 (2012).

23. Young, K. C., Sarovar, M. & Blume-Kohout, R. Error Suppression and Error Correction in Adiabatic Quantum Computation: Techniques and Challenges. *Phys. Rev. X* **3**, 041013 (2013).

24. Sarovar, M. & Young, K. C. Error suppression and error correction in adiabatic quantum computation: non-equilibrium dynamics. *New J. of Phys.* **15**, 125032 (2013).

25. Young, K. C., Blume-Kohout, R. & Lidar, D. A. Adiabatic quantum optimization with the wrong Hamiltonian. *Phys. Rev. A* **88**, 062314 (2013).

26. Pudenz, K. L., Albash, T. & Lidar, D. A. Error-corrected quantum annealing with hundreds of qubits. *Nat. Commun.* **5**, 3243 (2014).

27. Ganti, A., Onunkwo, U. & Young, K. Family of [[6k,2k,2]] codes for practical, scalable adiabatic quantum computation. *Phys. Rev. A* **89**, 042313 (2014).

28. Bookatz, A. D., Farhi, E. & Zhou, L. Error suppression in Hamiltonian-based quantum computation using energy penalties. *Physical Review A* **92**, 022317 (2015).

29. Mizel, A. Fault-tolerant, Universal Adiabatic Quantum Computation. http://arXiv.org/abs/1403.7694 (2014).

30. Pudenz, K. L., Albash, T. & Lidar, D. A. Quantum annealing correction for random Ising problems. *Phys. Rev. A* **91**, 042302 (2015).

31. Vinci, W., Albash, T., Paz-Silva, G., Hen, I. & Lidar, D. A. Quantum annealing correction with minor embedding. *Phys. Rev. A* **92**, 042310 (2015).

32. Mishra, A., Albash, T. & Lidar, D. A. Performance of two different quantum annealing correction codes. *Quant. Inf. Proc* **15**, 609–636 (2015).

33. Matsuura, S., Nishimori, H., Albash, T. & Lidar, D. A. Mean Field Analysis of Quantum Annealing Correction. http://arXiv.org/abs/1510.07709 (2015).

34. Aliferis, P., Gottesman, D. & Preskill, J. Quantum accuracy threshold for concatenated distance-3 codes. *Quantum Inf. Comput.* **6**, 97 (2006).

35. Roland, J. & Cerf, N. J. Quantum search by local adiabatic evolution. *Phys. Rev. A* **65**, 042308 (2002).

36. Somma, R. D., Nagaj, D. & Kieferová, M. Quantum Speedup by Quantum Annealing. *Phys. Rev. Lett.* **109**, 050501 (2012).

37. Hen, I. Period Finding with Adiabatic Quantum Computation. *Europhysics Letters* **105**, 50005 (2014).

38. Johnson, M. W. *et al.* Quantum annealing with manufactured spins. *Nature* **473**, 194–198 (2011).

39. Johnson, M. W. *et al.* A scalable control system for a superconducting adiabatic quantum optimization processor. *Superconductor Science and Technology* **23**, 065004 (2010).

40. Berkley, A. J. *et al.* A scalable readout system for a superconducting adiabatic quantum optimization system. *Superconductor Science and Technology* **23**, 105014 (2010).

41. Harris, R. *et al.* Experimental investigation of an eight-qubit unit cell in a superconducting optimization processor. *Phys. Rev. B* **82**, 024511 (2010).

42. Boixo, S. *et al.* Evidence for quantum annealing with more than one hundred qubits. *Nat. Phys.* **10**, 218–224 (2014).

43. Shin, S. W., Smith, G., Smolin, J. A. & Vazirani, U. How 'Quantum' is the D-Wave Machine? http://arXiv.org/abs/1401.7087 (2014).

44. Albash, T., Rønnow, T. F., Troyer, M. & Lidar, D. A. Reexamining classical and quantum models for the D-Wave One processor. *Eur. Phys. J. Spec. Top.* **224**, 111–129 (2015).

45. Albash, T., Vinci, W., Mishra, A., Warburton, P. A. & Lidar, D. A. Consistency tests of classical and quantum models for a quantum annealer. *Phys. Rev. A* **91**, 042314 (2015).

46. Crowley, P. J. D., Durić, T., Vinci, W., Warburton, P. A. & Green, A. G. Quantum and classical dynamics in adiabatic computation. *Phys. Rev. A* **90**, 042317 (2014).

47. Martin-Mayor, V. & Hen, I. Unraveling Quantum Annealers using Classical Hardness. http://arXiv.org/abs/1502.02494 (2015).

48. King, A. D., Lanting, T. & Harris, R. Performance of a quantum annealer on range-limited constraint satisfaction problems. http://arXiv.org/abs/1502.02098 (2015).

49. Vinci, W. *et al.* Hearing the Shape of the Ising Model with a Programmable Superconducting-Flux Annealer. *Sci. Rep.* **4**, 5703 (2014).

50. Bunyk, P. I. *et al.* Architectural Considerations in the Design of a Superconducting Quantum Annealing Processor. *IEEE Transactions on Applied Superconductivity* **24**, 1–10 (2014).

51. Venturelli, D. *et al.* Quantum Optimization of Fully Connected Spin Glasses. *Phys. Rev. X* **5**, 031040 (2015).

52. Emmanouilidou, A., Zhao, X. G., Ao, P. & Niu, Q. Steering an Eigenstate to a Destination. *Physical Review Letters* **85**, 1626–1629 (2000).

53. Deffner, S., Jarzynski, C. & del Campo, A. Classical and Quantum Shortcuts to Adiabaticity for Scale-Invariant Driving. *Physical Review X* **4**, 021013 (2014).

54. Kato, T. On the adiabatic theorem of Quantum Mechanics. *J. Phys. Soc. Jap.* **5**, 435 (1950).

55. Jansen, S., Ruskai, M.-B. & Seiler, R. Bounds for the adiabatic approximation with applications to quantum computation. *J. Math. Phys.* **48**, 102111 (2007).

56. Lidar, D. A., Rezakhani, A. T. & Hamma, A. Adiabatic approximation with exponential accuracy for many-body systems and quantum computation. *J. Math. Phys.* **50**, 102106 (2009).

57. Wiebe, N. & Babcock, N. S. Improved error-scaling for adiabatic quantum evolutions. *New J. Phys.* **14**, 013024 (2012).

58. Ge, Y., Molnár, A. & Cirac, J. I. Rapid adiabatic preparation of injective PEPS and Gibbs states. http://arXiv.org/abs/1508.00570 (2015).

59. Avron, J. E., Fraas, M., Graf, G. M. & Grech, P. Adiabatic Theorems for Generators of Contracting Evolutions. *Comm. Math. Phys.* **314**, 163–191 (2012).

60. Venuti, L. C., Albash, T., Lidar, D. A. & Zanardi, P. Adiabaticity in open quantum systems. http://arXiv.org/abs/1508.05558 (2015).

61. Subasi, Y. & Jarzynski, C. Simulating highly nonlocal Hamiltonians with less non-local Hamiltonians. http://arXiv.org/abs/1601.02922 (2016).

62. Choi, V. Minor-embedding in adiabatic quantum computation: II. Minor-universal graph design. *Quant. Inf. Proc.* **10**, 343–353 (2011).

63. Cai, J., Macready, W. G. & Roy, A. A practical heuristic for finding graph minors. http://arXiv.org/abs/1406.2741 (2014).

64. Boothby, T., King, A. D. & Roy, A. Fast clique minor generation in Chimera qubit connectivity graphs. *arXiv:1507.04774* (2015) URL http://arXiv.org/abs/1507.04774.

65. Kaminsky, W. M., Lloyd, S. & Orlando, T. P. *Quantum Computing and Quantum Bits in Mesoscopic Systems* chap. 25, 229–236 (Springer, 2004).

66. Klymko, C., Sullivan, B. D. & Humble, T. S. Adiabatic quantum programming: minor embedding with hard faults. *Quant. Inf. Proc* **13**, 709–729 (2014).

67. Ray, P., Chakrabarti, B. K. & Chakrabarti, A. Sherrington-Kirkpatrick model in a transverse field: Absence of replica symmetry breaking due to quantum fluctuations. *Phys. Rev. B* **39**, 11828–11832 (1989).

68. Thirumalai, D., Li, Q. & Kirkpatrick, T. R. Infinite-range Ising spin glass in a transverse field. *Journal of Physics A: Mathematical and General* **22**, 3339 (1989).

Supplemental Information accompanies the paper on the *npj Quantum Information* website (http://www.nature.com/npjqi)