

Efficient Hamiltonian programming in qubit arrays with nearest-neighbour couplings

Takahiro Tsunoda, Gaurav Bhole, Stephen A. Jones, Jonathan A. Jones,* and Peter J. Leek
Clarendon Laboratory, University of Oxford, Parks Road, Oxford OX1 3PU, United Kingdom

(Dated: August 14, 2020)

We consider the problem of selectively controlling couplings in a practical quantum processor with always-on interactions that are diagonal in the computational basis, using sequences of local NOT gates. This methodology is well-known in NMR implementations, but previous approaches do not scale efficiently for the general fully-connected Hamiltonian, where the complexity of finding time-optimal solutions makes them only practical up to a few tens of qubits. Given the rapid growth in the number of qubits in cutting-edge quantum processors, it is of interest to investigate the applicability of this control scheme to much larger scale systems with realistic restrictions on connectivity. Here we present an efficient scheme to find near time-optimal solutions that can be applied to engineered qubit arrays with local connectivity for any number of qubits, indicating the potential for practical quantum computing in such systems.

I. INTRODUCTION

In gate-based quantum computation [1], entangling interactions are generated by two-qubit logic gates, which are activated on demand. In order to turn the required two-qubit interactions on and off such gates are usually introduced by controllable external fields [2]. However, such interactions can also be generated directly, without extra control hardware, when there are suitable coupling terms naturally present in the internal Hamiltonian [3–5].

To control such “always-on” couplings, decoupling and selective recoupling techniques were developed in the Nuclear Magnetic Resonance (NMR) community. Just as a spin echo refocuses the coherent dynamics of a spin-1/2 nucleus [6], local NOT gates on selected qubits in a multi-qubit system reverse the evolution of corresponding terms in the internal Hamiltonian [7, 8]. By scheduling the timing and flip patterns properly, it is possible to rescale one-body and two-body terms to simulate any desired effective Hamiltonian [9, 10].

Superconducting qubits commonly use gates generated on demand but can also be tuned to generate always-on two-qubit interactions [11]. In fact always-on interactions are difficult to avoid in superconducting circuits due to their distributed microwave fields, and much effort is invested in mitigating these [12, 13]. Such always-on couplings are present in capacitively-coupled transmon qubits induced by their relatively small anharmonicity [14]. We consider such a background Hamiltonian that contains single qubit interactions (offsets) Ω_j and two qubit interactions (couplings) ω_{jk} which are both diagonal in the computational basis. The most general form

$$\mathcal{H}/\hbar = \sum_j \Omega_j \sigma_j^z/2 + \sum_{j < k} \omega_{jk} \sigma_j^z \sigma_k^z/4. \quad (1)$$

includes couplings between all pairs of qubits [15], but we will soon specialise to the case of *local* couplings.

In principle superconducting qubits could be operated based on this NMR-like architecture [11, 14], provided an efficient decoupling and recoupling scheme can be found to program a desired computation, and a wide range of other technologies have also been considered [16]. We have previously described methods for finding time-optimal rescaling sequences in systems of up to around 20 qubits, and near-optimal sequences for around 100 qubits [15]. However, these methods do not scale to larger systems, as the computational time required to find such solutions grows rapidly with the number of qubits.

The most general form of the Hamiltonian above will have all offsets and all possible couplings present, but couplings are generally only significant between nearby qubits. For example, superconducting qubits are often engineered in a square-lattice [17] and can have couplings limited to only nearest or next-nearest neighbour qubits if circuits are well microwave engineered [18]. Here we describe a method based on graph colouring for rapidly finding near-optimal sequences in these highly practical locally coupled systems. Remarkably we find that in this partially-coupled scenario, which is very realistic for large-scale superconducting circuits, the refocusing and rescaling of always-on couplings can be *efficiently* programmed, requiring only linear time to design control sequences and linear number of control pulses.

We start by applying the refocusing scheme to the smallest possible case of a nearest-neighbour square lattice, and then generalise to arbitrary lattice sizes. Initially we will assume the square lattice to be engineered with identical couplings between nearest neighbours and no long range couplings at all, but both of these restrictions will subsequently be partly relaxed.

We assume throughout this paper that it is possible to apply ideal single qubit NOT gates to any combination of qubits, despite the presence of the background Hamiltonian, finding the simplest set of control pulses rather than considering how these pulses might themselves be implemented. Experimental implementations will need to use either general approaches such as gradient-ascent pulse engineering (GRAPE) pulses [19], or specific methods tailored to this precise situation [20–22].

*Electronic address: jonathan.jones@physics.ox.ac.uk




FIG. 1: Networks to (a) refocus all offsets and couplings in a square of four qubits with only nearest-neighbour couplings, and (b) keep only the coupling between qubits 3 and 4. Here X indicates a NOT gate while U indicates evolution under the background Hamiltonian. Labels on the left indicate which qubits the NOT gates are applied to.

II. SQUARE LATTICES

We begin with the smallest possible square lattice, containing four qubits with identical nearest neighbour couplings $\omega_{j,j+1}$ and ω_{14} only. Non nearest-neighbour pairs have no coupling, so $\omega_{13} = \omega_{24} = 0$. We can refocus all the couplings in the system, as well as any offsets present, using the circuit of spin-echoes shown in Fig. 1(a). Here, the unitary U describes the evolution of the system under the background Hamiltonian \mathcal{H} given by the propagator $U = \exp(-i\mathcal{H}t/\hbar)$, while the NOT gates are represented by X . As there is no coupling to be refocused between qubits 1 and 3 it is possible to apply the same pattern of NOT gates to them both, and similarly for qubits 2 and 4, so only two distinct patterns of NOT gates are required.

This network can also be extended to retain a single coupling, as shown in Fig. 1(b). In this network the same pattern of NOT gates is applied to qubits 3 and 4, and the corresponding coupling is retained, while the remaining couplings are refocused [7]. These networks can be extended to retain *any* pattern of couplings in a square lattice system.

This extension works by colouring the square lattice, as described in [7]. The system can be described as a non-complete graph, with vertices corresponding to qubits and edges to couplings, with only some of the possible edges present. The graph can be coloured by assigning a colour to each vertex, and is said to be properly coloured, corresponding to complete refocusing, if no two connected vertices are the same colour. Thus in the fully-decoupled square of four qubits we can colour qubits 1 and 3 black and qubits 2 and 4 white, while to retain a single coupling the pair of qubits involved must be assigned the same third colour, say red, as shown in Fig. 2(a). To implement a colouring, we use patterns of NOT gates corresponding to distinct Walsh functions [7, 15] for each colour.

This colouring pattern can be tessellated across a lattice, as shown in Fig. 2(b), by colouring surrounding




FIG. 2: (a) Retaining a nearest neighbour coupling in a square requires three colours: decoupled qubits are coloured black and white, while the coupled qubits are coloured red. Grey dashed lines show couplings which have been refocused, while the coloured solid line indicates a retained coupling. (b) The same result can be achieved in a larger array by colouring surrounding qubits alternately black and white.

qubits alternately black and white. Here we show a four-by-four patch containing sixteen qubits, which can be embedded in a larger lattice, retaining a single coupling while refocusing all the other interactions. The required pulse sequence can be obtained by assigning black qubits B to the first Walsh pattern, white qubits W to the second, and the two red qubits R to the third, to obtain the sequence

$$U X_{\text{WR}} U X_{\text{BR}} U X_{\text{WR}} U X_{\text{BR}}, \quad (2)$$

where X_{WR} indicates that NOT gates are applied to the white and red qubits, and so on. Just like the sequence for the four-qubit system, this sequence requires only 4 time periods, but now requires $2q + 4$ pulses for a system of q qubits. The total time required to implement the network for a π evolution, corresponding to a controlled-NOT gate [15], is $T = 1/2J$, where the nearest-neighbour couplings are of size $\omega = 2\pi J$. This is the same time as is needed for an isolated coupling, as the retained coupling evolves at full strength.

The approach above will refocus all single-qubit interactions, but it is simple to modify the X gates in the network to implement single-qubit rotations directly. This relies on the identity [23]

$$\pi_{\phi_2} \pi_{\phi_1} = 2(\phi_2 - \phi_1)_z, \quad (3)$$

so that applying two π rotations around axes in the xy -plane separated by an angle δ is equivalent to performing a z -rotation through an angle 2δ . These rotations can be performed by modifying any pair of X gates in any refocusing network, and as each qubit is controlled separately different rotations can be applied to different qubits at no cost in time or pulse count.

III. PARALLEL GATES

If it is desired to retain several different coupling interactions then this can be achieved most simply by apply-




FIG. 3: Retaining a more complex pattern of couplings in a square lattice: this pattern requires four colours.

ing such patterns back to back, changing the colouring at each stage, but it is more efficient to as far as possible perform evolutions in parallel. For simple cases this can be achieved as shown in Fig. 3. Qubits which are part of the same coupling island, that is qubits which are connected either directly or indirectly by retained couplings, have been coloured the same colour, with the retained couplings indicated by lines. Note that one single qubit forms an island on its own, and that the choice of colours is completely arbitrary as long as enough distinct shades are used.

This simple approach will only be successful when, as here, all the couplings which could appear in an island are retained. In such cases this simple colouring strategy will retain all the desired couplings while refocusing all the couplings between islands, reproducing the desired pattern in one go. As before, the total time required is just $T = 1/2J$, the same as for an isolated coupling.

It might appear that this approach would require the number of colours to equal the number of islands, but in fact it is never necessary to use more than four, since islands which are completely disconnected (that is, islands which are not connected by couplings in the underlying Hamiltonian) can be safely coloured the same. By the four-colour-map theorem this will never require more than four colours [24]. These four colours can be assigned to four Walsh functions [15], requiring T to be divided up into eight equal time periods. Assigning the most common colours to W_1 and W_2 , which both require two pulses, and the two rarer colours to W_3 and W_4 , which require four pulses, means that the total number of pulses required in a system of q qubits lies between $2q$ and $3q$.

Although this method will work for some simple target patterns, it will not work in general, since many target patterns have one or more missing couplings. Consider, for example, the pattern shown in Fig. 4, where the black couplings must be retained, and the dashed red couplings must not be retained although they connect qubits within the main island. Any colouring which implements all the black couplings that must be retained will also implement the unwanted red couplings, and so this target pattern cannot be implemented with any single colouring pattern.




FIG. 4: A target pattern of couplings which cannot be implemented with a single colouring pattern. Although the grey qubits are easily decoupled from each other and from the black island, it is impossible to retain all the desired black couplings without also retaining the undesired couplings shown as dashed red lines.




FIG. 5: A solution implementing the target couplings in Fig. 4 using two sequential colouring patterns, (a) and (b).

IV. MULTIPLE COLOURINGS

There is, however, a simple method for achieving *any* target pattern with two sequential colourings, each using four colours. As an example a pair of colourings which implements the target couplings in Fig. 4 is shown in Fig. 5. Pattern (a) assigns two colours to the odd-numbered rows and two more to the even numbered rows, thus ensuring horizontal couplings can be controlled while all vertical couplings will be refocused. Along a row, the colour of the first qubit is arbitrary, but the following qubit must be the same colour if the corresponding coupling is to be retained (indicated by a joining line), and the other colour if it is to be refocused (no line). Pattern (b) assigns colours to columns instead of rows to control vertical couplings in an analogous way.

As any subset of horizontal couplings can be selected by the first colouring, and any subset of vertical couplings can be selected by the second, any target pattern at all can be implemented in this way, in a total time $2T$. This total time will be divided into 16 equal time periods, separated by no more than $6q$ individual pulses.

It is useful to compare this result with timings found by our previous algorithm [15] based on linear programming with an exhaustive basis set, which is guaranteed to find a true minimum-time solution. The computational complexity of this algorithm renders it impractical for large arrays, but for four-by-four arrays with nearest-neighbour couplings it is perfectly practical. We have analysed a very large number of randomly chosen targets in four-by-four arrays, and in every case the optimal solution required a total time of either T (in cases with no missing couplings, so a single colouring is possible) or $1.5T$ (in cases where this is not possible).

The implementation time required for our colouring based networks, $2T$, is slightly longer than the absolute minimum required, but these colourings are far easier to design, with a computation time scaling only *linearly* with the total number of qubits, and so can be applied to systems any number of qubits. The number of NOT gates required is also greatly reduced, from $O(q^2)$ for linear programming based solutions to $O(q)$.

V. NEXT-NEAREST NEIGHBOURS

Until now we have assumed that only nearest-neighbour interactions are important, with all others being too small to matter. In practice, real physical systems are likely to also have non-zero couplings at longer range. It is reasonable to expect such couplings to drop with distance in well-engineered devices, making the next-nearest neighbour interactions, across the diagonals of the square array, the most significant. In the original two-colour refocusing scheme qubits connected by next-nearest neighbour interactions will be the same colour, both black or both white, and so these interactions will be retained rather than refocused, leading to a significant error in the final gate implementation.

This is not a problem with the general four-colour scheme, as the use of different colours for alternate rows and columns guarantees that diagonally connected qubits will be of different colours. Thus these networks automatically suppress any unwanted diagonal couplings, which can therefore simply be ignored. Suppressing even longer range couplings is more complicated, but can be achieved using larger numbers of colouring patterns with more colours used in each pattern, as described in the Appendix.

VI. DIFFERENT EVOLUTION TIMES

The parallel gates approach, however, hides a further important assumption: it is not sufficient to retain two different couplings if they are required to evolve for different times, either because different couplings must evolve through different angles, or because apparently equivalent couplings will have different strengths, and so will need different evolution times to achieve the same angle.

It might seem necessary to apply such gates using different echo sequences, but in fact they can be partly combined. Consider two couplings in the same group, where one requires evolution for a total time T_A and the second for a time T_B , with $T_B > T_A$. The naive approach is to use two different colouring patterns, one implementing the first coupling for time T_A , and another implementing the second coupling for T_B . In fact these periods can be carried out partly in parallel: during the first period, which lasts for time T_A , *both* couplings are retained, while for the second period, which lasts for time $T_B - T_A$, only the second coupling is retained. Thus both couplings can be carried out in a total time T_B , and the method generalises for any number of distinct couplings strengths.

It follows that *any* pattern of couplings can be achieved in an evolution time equal to the sum of the longest evolution times required for horizontal and vertical couplings, which itself is no more than $2T_{\max}$, where $T_{\max} = 1/2J_{\min}$ is the evolution time required for the slowest gate. However the resulting sequences will contain $O(q)$ evolution times and $O(q^2)$ pulses, while designing such sequences requires sorting the evolution times into ascending order, with computational time complexity $O(q \log q)$. They will also be impractical to implement experimentally, as the differences between very similar times may be smaller than the clock resolution.

Rather than implementing a very large number of distinct evolution times precisely, it makes more sense to use a much smaller number of evolutions to approximate all the desired times. This is most easily achieved by dividing the range of evolution times by successive powers of two, in effect encoding each evolution time as a binary number. Using k different evolution times results in k -bit precision, with an exponential increase in precision with a linear increase in the number of evolution times used. By accepting a degree of approximation one can reduce the time complexity from $O(q \log q)$ to $O(kq)$, and reduce the pulse count from $O(q^2)$ to $O(kq)$, where the constant k depends on the accuracy required. For example, using 20 distinct delays will allow angles to be approximated to a precision better than 10^{-6} .

VII. CONCLUSIONS

The colouring techniques described here allow the efficient control of interactions in qubit arrays of arbitrary size, provided the couplings are constrained to be local. While the resulting pulse sequences take slightly longer

to implement than the absolute minimum time required, the computational time is vastly reduced, from $O(4^q)$ for exhaustive linear programming, or $O(q^6)$ for RROS [15], right down to $O(q)$, rendering them practical in systems with thousands or even millions of qubits. The number of control pulses is also greatly reduced, from $O(q^2)$ to $O(q)$, thus reducing implementation errors. The method can handle unwanted next-nearest neighbour couplings, and is easily extended in a scalable way to systems with variable coupling strengths or evolution angles.

Although it is not yet common to use always-on couplings as a computational resource in superconducting qubits, such implementations might fit well into the context of Near-term Intermediate Scale Quantum (NISQ) applications [25]. For example, the native interaction can be used for hardware-efficient implementations of the variational quantum eigensolver (VQE) [26] and quantum approximate optimisation algorithms (QAOA) [27]. Since entangling gates are realized without external fields, there is less need to calibrate the control hardware, as the calibration problem can be focused only on local control. Furthermore, quantum error mitigation methods might allow slow native interactions to output useful results by extended time evolution [28–30]. Efficient refocusing and rescaling methods, such as those described here, will be essential to extend experiments to the next generation of devices.

Acknowledgments

TT is supported by the Masason Foundation and the Nakajima Foundation. GB is supported by a Felix Scholarship. PL acknowledges support from the EPSRC under grant EP/T001062/1.

APPENDIX A: LONGER RANGE COUPLINGS

In the main text we describe a method for implementing a desired coupling pattern in an engineered system containing only nearest-neighbour couplings. While

longer range couplings are likely to be significantly weaker than nearest neighbour couplings, they will not in general be zero, and so could lead to significant errors. Fortunately our control scheme also refocusses next-nearest-neighbour couplings, across the diagonals of each square, which are likely to be the biggest error couplings. Here we consider the problem of refocussing longer range error couplings, which can be achieved using larger numbers of colouring patterns each of which contains more colours.

We begin by considering the case of error couplings between qubits separated by two edges horizontally or vertically (qubits separated by one horizontal and one vertical edge give rise to the diagonal error couplings considered previously). As before we will first colour the horizontal edges and then the vertical edges. When colouring horizontal edges we now have to divide the qubits into *three* groups of rows, rather than two, so that members of each group are at least three rows apart. Within each row a further problem arises, as we cannot directly retain couplings along two sequential horizontal edges without also retaining the error coupling between the two end qubits. Instead we have to use two separate colourings, retaining the two edges separately. This is most simply achieved by retaining odd and even numbered edges in separate steps. Within each row it is necessary to use three colours so that a colour is never repeated two qubits apart.

Putting this together we use two colouring patterns for the horizontal edges, and two for the vertical edges, making four patterns overall. Each pattern requires three separate colours for each of three sets of rows or columns, making nine colours in total. As before all diagonal error couplings will be automatically refocussed if the corresponding horizontal and vertical errors are refocussed.

This approach can be generalised to error couplings over any distance: to correctly handle all couplings between qubits n horizontal or vertical edges apart it is sufficient to use $2n$ colouring patterns, each of which has $(n+1)^2$ different colours. For the simplest case of nearest-neighbour couplings we have $n = 1$ and so two different four-colourings suffice, as described in the main text.

-
- [1] A. Barenco, C. H. Bennett, R. Cleve, D. P. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J. A. Smolin, and H. Weinfurter, Phys. Rev. A **52**, 3457 (1995).
 - [2] B. E. Kane, Nature **393**, 133 (1998).
 - [3] D. G. Cory, M. D. Price, and T. F. Havel, Physica D **120**, 82 (1998).
 - [4] X. Zhou, Z.-W. Zhou, G.-C. Guo, and M. J. Feldman, Phys. Rev. Lett. **89**, 197903 (2002).
 - [5] S. C. Benjamin and S. Bose, Phys. Rev. Lett. **90**, 247901 (2003).
 - [6] E. L. Hahn, Phys. Rev. **80**, 580 (1950).
 - [7] J. A. Jones and E. Knill, J. Magn. Reson. **141**, 322 (1999).
 - [8] D. W. Leung, I. L. Chuang, F. Yamaguchi, and Y. Yamamoto, Phys. Rev. A **61**, 42310 (2000).
 - [9] D. Leung, J. Mod. Opt. **49**, 1199 (2002).
 - [10] A. Parra-Rodriguez, P. Lougovski, L. Lamata, E. Solano, and M. Sanz, Phys. Rev. A **101**, 022305 (2020).
 - [11] J. M. Gambetta, J. M. Chow, and M. Steffen, npj Quantum Information **3**, 2 (2017).
 - [12] Y. Chen, C. Neill, P. Roushan, N. Leung, M. Fang, R. Barends, J. Kelly, B. Campbell, Z. Chen, B. Chiaro, et al., Phys. Rev. Lett. **113**, 220502 (2014).
 - [13] P. Mundada, G. Zhang, T. Hazard, and A. Houck, Phys. Rev. Applied **12**, 054023 (2019).
 - [14] L. DiCarlo, J. M. Chow, J. M. Gambetta, L. S. Bishop,

- B. R. Johnson, D. I. Schuster, J. Majer, A. Blais, L. Frunzio, S. M. Girvin, et al., *Nature* **460**, 240 (2009).
- [15] G. Bhole, T. Tsunoda, P. J. Leek, and J. A. Jones, *Phys. Rev. Applied* **13**, 034002 (2020).
- [16] T. Satoh, Y. Matsuzaki, K. Kakuyanagi, W. J. Munro, K. Semba, H. Yamaguchi, and S. Saito, *Phys. Rev. A* **91**, 052329 (2015).
- [17] F. Arute et al., *Nature* **574**, 505 (2019).
- [18] P. A. Spring, S. Sosnina, T. Tsunoda, B. Vlastakis, and P. J. Leek, *Modelling of TM modes in periodically-shortened cavities for circuit QED* (2019), 1909.02104.
- [19] N. Khaneja, T. Reiss, C. Kehlet, T. Schulte-Herbrüggen, and S. J. Glaser, *J. Magn. Reson.* **172**, 296 (2005).
- [20] P. Sengupta and L. P. Pryadko, *Phys. Rev. Lett.* **95**, 037202 (2005).
- [21] A. De and L. P. Pryadko, *Phys. Rev. Lett.* **110**, 070503 (2013).
- [22] A. De and L. P. Pryadko, *Phys. Rev. A* **89**, 032332 (2014).
- [23] J. J. L. Morton, A. M. Tyryshkin, A. Ardavan, S. C. Benjamin, K. Porfyrakis, S. A. Lyon, and G. A. D. Briggs, *Nature Physics* **2**, 40 (2006).
- [24] K. Appel and W. Haken, *Sci. Am.* **237**(4), 108 (1977).
- [25] J. Preskill, *Quantum* **2**, 79 (2018).
- [26] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow, and J. M. Gambetta, *Nature* **549**, 242–246 (2017).
- [27] E. Farhi, J. Goldstone, and S. Gutmann, *A quantum approximate optimization algorithm* (2014), 1411.4028.
- [28] K. Temme, S. Bravyi, and J. M. Gambetta, *Phys. Rev. Lett.* **119**, 180509 (2017).
- [29] Y. Li and S. C. Benjamin, *Phys. Rev. X* **7**, 021050 (2017).
- [30] A. Kandala, K. Temme, A. D. Córcoles, A. Mezzacapo, J. M. Chow, and J. M. Gambetta, *Nature* **567**, 491–495 (2019).