


Rescaling Interactions for Quantum Control

Gaurav Bhole¹, Takahiro Tsunoda¹, Peter J. Leek¹, and Jonathan A. Jones^{1*}
Clarendon Laboratory, University of Oxford, Parks Road, Oxford OX1 3PU, United Kingdom

 (Received 13 November 2019; revised manuscript received 13 January 2020; accepted 3 February 2020; published 2 March 2020)

A powerful control method in experimental quantum computing is the use of spin echoes, employed to select a desired term in the system's internal Hamiltonian, while refocusing others. Here, we address a more general problem, describing a method to not only turn on and off particular interactions but also to rescale their strengths so that we can generate any desired effective internal Hamiltonian. We propose an algorithm based on linear programming for achieving time-optimal rescaling solutions in fully coupled systems of tens of qubits, which can be modified to obtain near-time-optimal solutions for rescaling systems with hundreds of qubits.

DOI: [10.1103/PhysRevApplied.13.034002](https://doi.org/10.1103/PhysRevApplied.13.034002)

I. INTRODUCTION

In the circuit model of quantum computation, quantum algorithms are implemented by applying a network of quantum logic gates to a set of qubits. Typically, these qubits are defined by the internal Hamiltonian of the underlying system, while logic gates are applied using external control fields. In some cases, such as trapped-ion implementations [1], the free evolution under the internal Hamiltonian corresponds to single-qubit z rotations and two-qubit gates are generated as and when required; for example, by coupling internal and external degrees of freedom [2]. In other cases, exemplified by nuclear magnetic resonance (NMR), the internal Hamiltonian contains two-qubit zz interactions and two-qubit gates are generated by free evolution [3]. Many other proposed implementations [4–13] fit into one of these two paradigms or a mixture [14]: for example, N-V centers are NMR-like when controlling the spins around a single center [15] but generate long-range two-qubit gates on demand [16]. Superconducting qubits commonly use gates generated on demand [17] but can also be tuned to generate always-on two-qubit interactions [18].

With always-on two-qubit interactions, free evolution provides a universal quantum logic gate when combined with single-qubit rotations but this gate is not a convenient one, as it corresponds to a complex pattern of evolutions. For this reason, methods such as spin echoes [19] are used to replace the background Hamiltonian with a more convenient average Hamiltonian, in which desired interactions are isolated while undesirable interactions are refocused. As a concrete example system, we consider NMR spin

systems, but equivalent Hamiltonians can be found in other quantum-information-processing (QIP) implementations.

Previous studies in NMR [20–23] have largely concentrated on methods for refocusing all the interactions, or for isolating one single interaction while refocusing everything else. However, a more general problem is to *rescale* the size of interactions in the Hamiltonian, to produce a desired effective Hamiltonian. Here, we describe a method for finding a rescaling sequence with the shortest possible total time and with a fairly small number of echo pulses. In its simplest form, this is practical for systems of up to about 20 qubits, but for larger systems we develop a pragmatic method for finding short, although not perfectly optimal, rescaling sequences, which works with more than 100 qubits, meaning that the method can be of use for practically useful quantum computation in scalable platforms. Related ideas have been explored in systems with other similar Hamiltonians [14,24,25].

The initial sections below introduce the terminology of spin echoes and reprise previous results for refocusing sequences. This is followed by a description of our rescaling method with some sample applications in systems with a small number of qubits. Finally, we show how random sampling can be used to extend these methods to much larger systems.

II. SPIN ECHOES

The background Hamiltonian of an NMR spin system consists of one-spin interactions (resonance offsets) and two-spin interactions (J couplings). Consider a system of q spins, where the i th spin has a resonance offset Ω_i and the pair of spins i and j have a J -coupling frequency ω_{ij} . The Hamiltonian for this system in the weak coupling

*jonathan.jones@physics.ox.ac.uk

limit is

$$\mathcal{H} = \sum_i \Omega_i I_z^i + \sum_{i < j} \omega_{ij} I_z^i I_z^j, \quad (1)$$

where, following NMR notation, factors of \hbar have been dropped and $I_z^i = \sigma_z^i/2$ is the Pauli spin-1/2 operator acting on the i th spin. In practice, some of these interactions could either be zero or set to zero, allowing them to be dropped. For example, working in a suitable rotating frame allows some of the resonance offsets to be set to zero, while many J couplings can be negligible in partially coupled spin systems. However, for generality and completeness, we shall initially consider fully coupled systems with q nonzero resonance-offsets and $p = q(q-1)/2$ nonzero J couplings between the p pairs of spins.

During a period τ of free evolution, each spin evolves under all the $q + p$ interactions, given by the propagator $U = \exp(-i\mathcal{H}\tau)$. Since this Hamiltonian is diagonal in the chosen z basis, all terms in the Hamiltonian commute and the one-spin and two-spin evolutions can be summarized by the acquired phases

$$\Phi_i = \Omega_i \tau, \quad \phi_{ij} = \omega_{ij} \tau. \quad (2)$$

To sculpt the effective Hamiltonian into the desired form, it is necessary to isolate the spin interactions that we want while suppressing the unwanted interactions. This essentially requires controlling the evolution of the spins such that the unwanted interactions finally acquire a phase of zero while letting the required interactions evolve to reach the desired values.

A spin echo is a period of free evolution interrupted by π rotations in the middle of the evolution period. The notation π^i is used to denote a pulse that causes a 180 degree rotation on spin i , about the x axis unless otherwise stated. The effect of a pair of π^i rotations is to negate the effective frequency Ω_i of the spin for the time period between the two pulses. Thus, the sequence $\tau\pi^i\tau\pi^i$, where time periods indicate free evolution under the internal Hamiltonian, will refocus the offset Ω_i as the phase Φ_i acquired during the first period of evolution gets nullified by the phase $-\Phi_i$ acquired in the second half.

Clearly, the one-spin interaction of any given spin i is only affected by π rotations applied to spin i and so individual interactions can be controlled independently. However, for two-spin interactions, a π rotation applied to either spin i or j reverses the frequency ω_{ij} while a simultaneous π rotation applied on both spins i and j leaves ω_{ij} unchanged. A sequence $\tau\pi^i\tau\pi^i$ will thus refocus Ω_i and ω_{ij} , whereas a sequence $\tau\pi^{ij}\tau\pi^{ij}$ will refocus both Ω_i and Ω_j but not ω_{ij} . In this manner, the two-spin interactions can be controlled but this control cannot be achieved independently from one-spin interactions.

A general spin-echo sequence comprises a series of free-evolution time periods τ_m , sometimes called delays,

separated by π pulses applied to one or more spins. As long as the total number of π pulses applied to a given spin is even, the overall evolution can still be summarized by a set of phases, but now

$$\Phi_i = \Omega_i \sum_m S_m^i \tau_m, \quad \phi_{ij} = \omega_{ij} \sum_m S_m^i S_m^j \tau_m, \quad (3)$$

where S is a sign matrix, containing only the elements ± 1 , with a sign change whenever a π pulse is applied to spin i . For convenience, we will also refer to the two-qubit sign matrix $S_m^{ij} = S_m^i S_m^j$, although this is obviously not independent from the one-qubit matrix. The complete sign matrix can be obtained by combining the one- and two-spin matrices.

III. REFOCUSING

Methods for removing all the interactions (sometimes called decoupling), or for isolating one single interaction while refocusing everything else, have been widely studied. The most effective methods to achieve this rely on choosing sign matrices the rows of which are taken from Walsh–Hadamard matrices [21–23], so that each row is a Walsh function [26]. These matrices differ from other Hadamard matrices in that they are only defined for dimensions equal to a power of 2, the rows are not normalized, and the ordering of the rows is different.

A Walsh function W_n is defined by a vector with length equal to a power of 2 and with all the entries equal to ± 1 . For W_0 , all the entries are $+1$, while for every other W_n , half the entries are $+1$ and half are -1 , with the entries arranged such that there are n regularly spaced sign changes along the vector. For example, the 4×4 Walsh–Hadamard matrix contains the four rows

$$\begin{pmatrix} W_0 \\ W_1 \\ W_2 \\ W_3 \end{pmatrix} = \begin{pmatrix} +1 & +1 & +1 & +1 \\ +1 & +1 & -1 & -1 \\ +1 & -1 & -1 & +1 \\ +1 & -1 & +1 & -1 \end{pmatrix}. \quad (4)$$

Strictly, the name of the Walsh function must specify the number of columns as well as the number of sign changes but this is left implicit here: the number of columns is equal to the smallest power of 2 larger than the highest Walsh number considered.

In a system of three spins, it is possible to remove all three one-spin and all three two-spin interactions by using four equal time periods τ and a sign matrix obtained by choosing $S^i = W_i$, avoiding W_0 . This relies on two key properties of Walsh functions. First, all Walsh functions except W_0 contain an equal number of ± 1 values and so all one-qubit interactions will be refocused when equal-length time periods are used. Second, the product of two

Walsh functions is itself a Walsh function [26], defined by

$$W_p \circ W_q = W_{p \oplus q}, \quad (5)$$

where “ \circ ” indicates element-wise multiplication, sometimes called the Schur product [27], and “ \oplus ” indicates bitwise addition modulo 2. Thus all two-qubit interactions will also be refocused.

A decoupling network is easily modified [21,22] to retain a single interaction: to retain a one-spin interaction Ω_i , use $S_i = W_0$ for this spin, while to retain a coupling ω_{ij} , set $S_i = S_j$ so that $S_{ij} = W_0$. To take a concrete example, the coupling ω_{12} can be isolated in a three-spin system by choosing $S_1 = S_2 = W_1$ and $S_3 = W_2$. The π pulses required can be deduced by applying a pulse to a spin whenever the corresponding row of S changes sign, including a final π pulse if the S row ends in -1 , giving the sequence

$$\tau \pi^3 \tau \pi^{1,2} \tau \pi^3 \tau \pi^{1,2}. \quad (6)$$

Note that ω_{12} evolves with sign $+1$ at every stage and so is retained at full strength. The total evolution time required is given by $4\tau = \phi_{12}/\omega_{12}$. If this expression gives a negative time, then this can be resolved by applying additional π pulses to one spin at the beginning and end of the sequence to negate the evolution frequency.

Now consider how this approach scales to a system of q spins. Retaining a single interaction can be done efficiently: the number of time periods required is given by the smallest power of 2 larger than q , which is upper bounded by $2q$, while the number of individual π pulses required is clearly less than $2q^2$, which corresponds to applying a pulse to every spin after every time period. A more careful analysis (see Appendix A) shows that only around $q^2/2$ pulses are required, which is still $O(q^2)$ but with a smaller prefactor. As the single interaction is retained at full strength, this is also a minimum-time solution.

IV. RESCALING

So far, we have only considered retaining a single one-spin or two-spin interaction while refocusing the remainder. However, a more general problem is to *rescale* the size of interactions in the Hamiltonian. In other words, we desire to achieve a certain set of nonzero target phases for *all* the spin interactions. The obvious approach is just to place spin-echo sequences that isolate the individual interactions back to back. As there is a total of $r = q + p = q(q+1)/2$ single-spin and two-spin interactions to be considered, it is clear that the number of time periods is $O(q^3)$ and the number of pulses is $O(q^4)$. The total time required is given by the sum of the times required to evolve under

each individual interaction:

$$T = \sum_i \left| \frac{\Phi_i}{\Omega_i} \right| + \sum_{i < j} \left| \frac{\phi_{ij}}{\omega_{ij}} \right|. \quad (7)$$

This naive approach is expensive, both in time and the number of π pulses. We thus want to find a more efficient sequence, by carrying out as many evolutions in parallel as possible. Although this might sound challenging, we propose here a straightforward way to achieve this using linear programming. This approach also greatly reduces the number of pulses and time periods required.

A. Setting up the problem

We consider a system of q coupled spins described by $r = q(q+1)/2$ one- and two-spin interactions. Our aim is to rescale all r interactions simultaneously such that they reach the desired target phases, which, for generality, we assume to be all different. We begin by constructing an overcomplete Walsh basis by building a one-spin sign matrix with the rows given by Walsh functions numbered 2^j , where $j = 0, 1, \dots, (q-1)$. Next, we use this matrix to construct the two-spin sign matrix by taking products of corresponding rows in the one-spin matrix.

Combining these by stacking the two matrices together gives the complete sign matrix S of r rows and $s = 2^q$ columns. The single-spin functions correspond to binary numbers with precisely 1 bit set, while the two-spin functions correspond to binary numbers with precisely 2 bits set. As these numbers are all distinct, it is guaranteed that the complete sign matrix has enough flexibility to permit every interaction to be controlled separately. This is quite different from refocusing sequences, where many functions are repeated.

B. Linear programming

This overcomplete Walsh basis guarantees that solutions to Eq. (3) can be found for any target values of $\{\Phi_i\}$ and $\{\phi_{ij}\}$ by choosing 2^q appropriate values of $\{\tau_m\}$, but it is not obvious how these can be found. As the basis is overcomplete, multiple solutions will exist but these can be distinguished by requiring that all the times $\{\tau_m\}$ must be non-negative and by preferring the solutions with the shortest value of total time $T = \sum_m \tau_m$. These criteria for desirable solutions suggest a powerful method, namely linear programming [28].

The general linear programming problem varies some inputs (here, the times $\{\tau_m\}$), seeking to minimize some linear function of these inputs (here, the total time T) subject to a number of equality constraints [here Eq. (3)] and inequality constraints (here, that each $\tau_m \geq 0$). We adopt a simple approach, using the inbuilt MATLAB function `linprog`.

It is important to consider the computational complexity of linear programming, as this determines how the time required to find a solution scales with the number of qubits q . The precise computational complexity of linear programming is known to be poorly defined, depending on both the algorithm used and the details of the problem and with the worst-case behavior being very different from the typical case [28]. The MATLAB function `linprog` has a choice of two algorithms: the original simplex algorithm developed by Dantzig [29] and a more modern interior-point algorithm [30]—these are briefly discussed in Appendix B. Both algorithms typically have computational complexity between $O(n^2)$ and $O(n^3)$, where n , the dimension of the problem, can be taken as the sum of the number of rows and columns in the constraint matrix, so that here $n = r + s \approx 2^q$. We investigate this question experimentally by simply timing the program. All results are for the simplex algorithm unless otherwise stated.

C. Extracting solutions

As the linear programming algorithm is fundamentally trying to minimize T , subject to the positivity constraint and the target phases, the algorithm prefers solutions where many of the $\{\tau_m\}$ are zero. (This is not specific to this problem, but is a general feature of linear programming solutions [31].) The linear programming solution has at most only as many nonzero times as the number of constraints r in the problem. If the problem involves extensive refocusing rather than rescaling, then solutions with an even smaller number of nonzero times can be found.

It is obviously not necessary to explicitly implement the time periods of length 0 and so the overcomplete sign matrix, S , can be replaced by a reduced matrix, R , by selecting only r or fewer columns from S that correspond to nonzero evolution times.

D. Optimizing the solutions

One subtlety is that the order of columns in the R matrix does not affect the phases produced but different orderings of these columns can lead to pulse sequences with different numbers of pulses. As minimization of the number of pulses is desirable, it is useful to explore different permutations of the R matrix, seeking for the arrangement that gives the smallest number of sign changes.

If the matrix is not too large, then exhaustive permutation can be practical, but in larger cases it is more sensible to select a number of random permutations and keep the best one found. Experience so far suggests that different permutations can require numbers of pulses that differ by a factor of around 2. We also find that the pulse pattern corresponding to the original R matrix is typically relatively good, although rarely the absolute best.

E. Building the pulse sequence

From this optimal reduced matrix R , a pattern of pulses can be generated by applying a π pulse to every spin the sign of which changes. It is important to remember that all interactions start at $+1$ and must end at $+1$, which can be modeled by adding initial and final columns to R containing entirely $+1$. These additional columns have evolution times set to zero and so are not actually implemented but the resulting sign changes make it necessary to apply pulses to some spins at the start and end of the sequence. This also ensures that the number of π pulses applied to each spin is even, which is required to create true spin echoes.

As the reduced matrix has $r \approx q^2/2$ times, the final pulse sequence will have $O(q^2)$ time periods and $O(q^3)$ individual π pulses, which is a very significant improvement on naive methods.

V. EXAMPLE CALCULATIONS

We illustrate our method of rescaling with the help of two specific examples before returning to the general case, when we will consider the computational time required to find these optimal solutions.

A. Homonuclear three-spin fully coupled system

First, consider the example of iodotrifluoroethene [32], C_2F_3I , with the three ^{19}F nuclei forming our spin system, in a magnetic field such that the Larmor frequency of 1H nuclei is 600 MHz, with the excitation frequency set in the middle of the spectral range. The molecular structure and the Hamiltonian of the system are shown in Fig. 1(a).

This system has $r = 6$ interactions altogether. Our aim here is to rescale all three of the two-spin interactions such that they have equal effective strength, while refocusing the three one-spin interactions. This is equivalent to achieving the target phases $\{\Phi_1, \Phi_2, \Phi_3\} = \{0, 0, 0\}$ and $\{\phi_{12}, \phi_{23}, \phi_{13}\} = \{\pi, \pi, \pi\}$. The quantum circuit shown in Fig. 1(a) uses the coupling gate [33,34], which is equivalent to a controlled- Z gate, with the notation indicating the complete symmetry of a controlled gate implemented through zz couplings, with no distinction between control and target spins. Actual controlled- Z gates would require additional single-spin Z rotations, so that $\{\Phi_1, \Phi_2, \Phi_3\} = \{\pi, \pi, \pi\}$. These can be either be implemented directly in the pulse sequence [35–37] or tracked and corrected at the end [33,34]. Both approaches have been used in the literature, with tracking being particularly popular for fixed phases that are integer multiples of $\pi/2$, as seen here, as this is relatively simple in many experimental implementations.

Here, we will consider both cases, beginning with tracking, so that the single-qubit target phases are all 0. Note that even though the target phases are all 0 or π , the

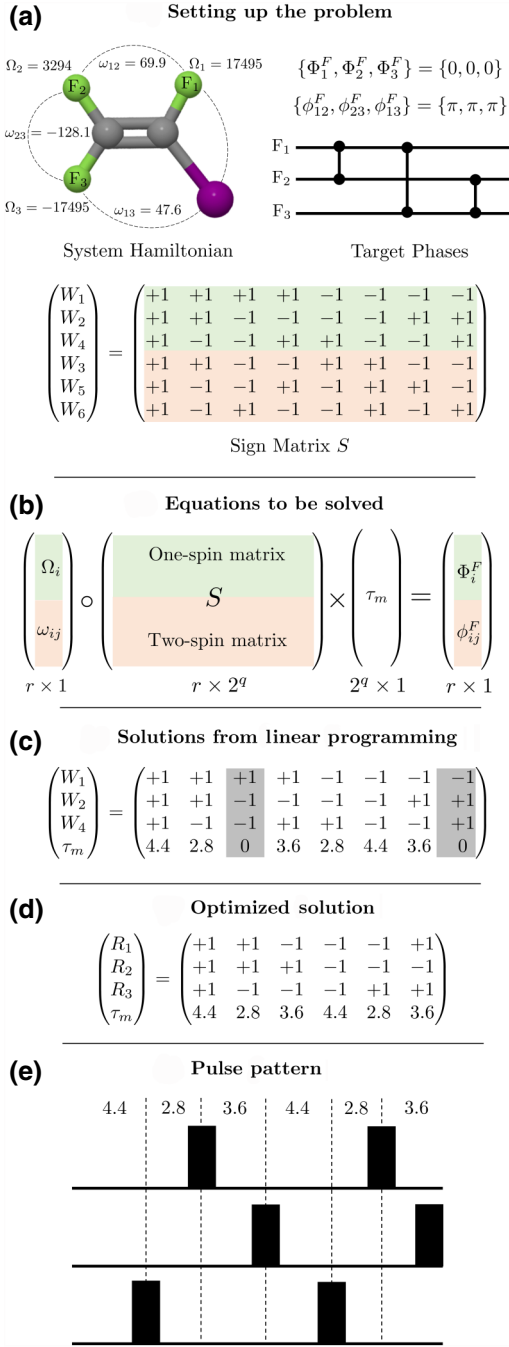


FIG. 1. The implementation of our algorithm using the three ^{19}F nuclei in iodotrifluoroethene, forming a three-spin full coupled homonuclear NMR system: (a) setting up the problem; (b) equations to be solved; (c) solutions from linear programming; (d) the optimized solution; (e) the pulse pattern. All interactions are written in hertz and times in milliseconds. The stages of the implementation are described in the main text.

corresponding evolution times will be different, as the one- and two-spin interactions have different values. Thus, we expect to require up to six time periods, consistent with $r = 6$ interactions.

We begin by constructing an overcomplete basis from a one-spin sign matrix with three rows and eight columns, corresponding to the Walsh functions $\{W_1, W_2, W_4\}$, with the explicit form shown in Fig. 1(a). The two-spin interaction matrix can be constructed by taking the three unique products of rows in the one-spin interaction matrix. The resulting set of equations is shown schematically in Fig. 1(b). Here, \circ indicates an element-wise (Schur) product as before, while \times indicates a conventional matrix product.

As the sign matrix has more columns than rows, there is no unique solution but linear programming can be used to find a solution that minimizes the total time while keeping all time periods positive. The program finds an optimal solution that uses only six columns as expected, by assigning zero time periods to two of the columns, as shown in Fig. 1(c). The total time T of the evolution is 21.6 ms, which in this case is identical to the time required for a naive implementation with sequential gates. This is a general rule for rescaling multiple couplings in three-spin systems, as it is impossible to control any coupling without affecting all the others.


However, even in this case, the resulting pulse sequence, requiring eight pulses, is much simpler than naive sequential designs, which require 18 pulses. We can further improve this by permuting the columns of the reduced sign matrix R , as shown in Fig. 1(d), to find a solution with a smaller number of sign changes. From the reduced sign matrix, we can now calculate the pattern of π pulses, shown in Fig. 1(e), with a π pulse applied wherever we encounter a sign change. Since there is one sign change in row R_2 , it might seem that we need only one π pulse. However, since the row ends in -1 , we must apply another π pulse at the end, leading to a total of six π pulses. Similarly, reduced sign matrix rows might in some cases begin with -1 and in this case it is necessary to apply a π pulse before the first time period.

Next, we consider the case in which tracking is not used, so that the single-qubit target phases of π need to be implemented directly. In this case, a naive sequential approach would require a small additional evolution time divided up by 12 additional pulses. By contrast, our efficient method can achieve this with no additional evolution time or pulses by slightly unbalancing the identical pairs of times seen in the case above. It is always possible to do this as long as all the single-qubit offset frequencies are large in comparison with the two-qubit coupling strengths, which is usually the case and can always be achieved by changing the central excitation frequency.

We have assumed here that the π pulses are implemented perfectly—that is there is no evolution under the drift Hamiltonian while the π pulse is applied. This will be approximately true in a heteronuclear spin system, while in a homonuclear system equivalent pulses can be designed using techniques such as gradient-ascent pulse engineering (GRAPE) [38].

B. Homonuclear four-spin linear chain

Next, consider an NMR system provided by four ^{13}C nuclei in labeled crotonic acid [39]. The four spins form a rough linear chain, with large one-spin interactions, moderate nearest-neighbor couplings, and significantly smaller long-range couplings, as shown in Fig. 2. Our typical target for this system is to refocus all the one-spin interactions and the three small couplings but to retain the three nearest-neighbor couplings. We also want to rescale the retained couplings such that they have the same target phases. We have a total of ten interactions, from which we need to refocus seven while rescaling three. This can also be represented, up to single-spin phases, by the following quantum circuit:



In other words, we want the target phases to be

$$\{\Phi_1, \Phi_2, \Phi_3, \Phi_4\} = \{0, 0, 0, 0\} \quad (9)$$


and

$$\{\phi_{12}, \phi_{13}, \phi_{14}, \phi_{23}, \phi_{24}, \phi_{34}\} = \{\pi, 0, 0, \pi, 0, \pi\}. \quad (10)$$

We begin by constructing the overcomplete basis by building a sign matrix S with ten rows and 16 columns that correspond to the Walsh functions $\{W_1, W_2, W_4, W_8\}$ and their six distinct products, which are $\{W_3, W_5, W_6, W_9, W_{10}, W_{12}\}$. Running linear prediction and following the procedure as in the previous example, we now find a solution with nine time periods,

$$\begin{pmatrix} R_1 \\ R_2 \\ R_3 \\ R_4 \\ \tau_m \end{pmatrix} = \begin{pmatrix} +1 & +1 & -1 & -1 & -1 & -1 & +1 & +1 & +1 \\ +1 & +1 & +1 & -1 & -1 & -1 & -1 & +1 & +1 \\ +1 & +1 & +1 & +1 & -1 & -1 & -1 & -1 & -1 \\ +1 & -1 & -1 & +1 & +1 & -1 & -1 & -1 & +1 \\ 3.5 & 1.3 & 1.8 & 3.0 & 1.8 & 3.0 & 1.8 & 1.7 & 1.3 \end{pmatrix}, \quad (11)$$

which can be implemented with ten pulses. Note that since our target consists mostly of zeros, the number of time



C_1	C_2	C_3	C_4	(Hz)
-11 962.2	41.6	1.5	7.1	C_1
	7306.0	69.6	1.2	C_2
		3972.1	72.3	C_3
			10 626.1	C_4

FIG. 2. The molecular structure of fully labeled crotonic acid and the corresponding Hamiltonian for four ^{13}C nuclei forming a four-spin chainlike homonuclear NMR system.

periods required here is less than r . In this case, the total time required is only 19.2 ms, significantly shorter than the naive sequential time of 26.1 ms.

As a second example, we now consider the case in which we only seek to implement the first and third controlled- Z gates, so that the coupling phase are now

$$\{\phi_{12}, \phi_{13}, \phi_{14}, \phi_{23}, \phi_{24}, \phi_{34}\} = \{\pi, 0, 0, 0, 0, \pi\}. \quad (12)$$

In this case, only six time periods are required, although once again ten pulses are needed to implement the desired sign changes. The total time required is now 12.0 ms, which is identical to the time required to implement just the first gate by naive methods. Thus, in this case, our algorithm has found a solution that implements two gates in parallel, at no additional cost in total time. Clearly, it is not possible to implement the gates any more quickly than this, as the time required to implement the slowest gate on its own provides a firm lower bound. An upper bound is provided by the time required to implement each of the r evolutions in sequence [Eq. (7)], which is simply the sum of the times required for each gate.

C. General solutions

We repeat calculations of this kind in a large variety of fictitious spin systems with increasing numbers of spins, up to $q = 18$, and solutions are always found. For larger values of q , these solutions are always more time-efficient than the sequential approach and are usually far quicker. The greatest savings are found in cases in which a moderate number of gates need to be implemented in parallel and particularly when unused long-range couplings are significantly weaker than the couplings being controlled.

Until now, we have described the problem as if there was a unique optimal solution that the linear programming locates. In fact, there are multiple equivalent solutions, from among which the linear programming chooses one. These alternative solutions can be easily generated by permuting the columns of the S matrix before running the algorithm but as they all have the same number of individual time periods and take the same total time, there is no good reason to do this.

We note in passing that these optimal solutions are only optimal for implementations containing only delays and π pulses. If it is desired to implement an evolution corresponding to a weak long-range coupling, then it may be quicker to use SWAP gates and related methods to implement long-range interactions through a chain of stronger short-range interactions [40,41]. Even in such cases, however, the ideas described here can be used to assist in the design of such indirect gates.

D. Computation time complexity

The principal downside of this approach is that the computational time needed to run linear programming




FIG. 3. The computation time required to find a linear programming solution on a desktop computer as a function of q . For small values of q this time is almost constant but it rises exponentially for larger q , rendering this method impractical beyond around 20 qubits. Timings are shown for the interior-point algorithm, but the results are very similar for the simplex algorithm. The error bars (estimated by repetition) are comparable to the size of the symbols and the line simply joins the individual points.

increases with the size of the basis set and this grows as $s = 2^q$ for the current method. As discussed above, a computation time scaling proportional to $(s + r)^2$ or $(s + r)^3$, where r is the number of equality constraints, is likely, which for large q is dominated by the exponential growth in s . This is confirmed by a plot of time required on a desktop computer (Intel Core i7-9700, 3.0–4.7 GHz, with a 12-MB cache and 40-GB RAM), shown in Fig. 3. The linear behavior at large q on this semilog plot indicates exponential computational time complexity and the gradient is consistent with the time required, scaling as about 4^q .

Although attempts have been made to parallelize linear programming algorithms, progress so far has been limited [42]. Thus, this method seems to be practical only up to a small number of spins, perhaps $q = 20$. Indeed, much above 20 qubits it becomes difficult even to hold the S matrix in memory on a desktop computer, although this could be sidestepped with a customized algorithm. While it is true that 20 is quite a large number of spins in the context of conventional NMR or even NMR QIP, we do not want to restrict ourselves to NMR spin systems but to extend to more general quantum systems that have the potential for a scalable quantum computing architecture. Fortunately, a more sophisticated approach is available that takes time only polynomial in q , albeit with a high power.

VI. STABILIZING SOLUTIONS

Having located an optimal solution, it is important to consider whether it can be implemented reliably. This will obviously require the ability to implement large numbers of π pulses (single-qubit NOT gates) reliably but we do not consider that problem here. Instead, we consider the more




FIG. 4. The infidelity of three-qubit circuits as a function of the precision of the underlying clock. Both the infidelity $\mathcal{I} = 1 - \mathcal{F}$ and the time resolution are plotted on logarithmic scales, with the time resolution running from 1 ms to 1 ns. The lower line, plotted in blue, shows the infidelity for the first three-qubit circuit, in which the single-qubit target phases are all zero. The upper line, plotted in red, corresponds to the second three-qubit circuit, where the single-qubit target phases are all π .

fundamental question of the precision that is necessary for the free-evolution times that lie between the pulses.

So far, we have assumed that these times can be implemented with effectively perfect precision but any real implementation will be built around a clock, so that any delay period must be rounded to the nearest multiple of some underlying cycle time. We show in Fig. 4 the effects of such rounding on the two three-qubit circuits considered in the previous section. This figure shows the infidelity $\mathcal{I} = 1 - \mathcal{F}$, where \mathcal{F} , the fidelity between the desired propagator, U , and the propagator actually implemented, V , is the propagator fidelity [43]

$$\mathcal{F} = \left| \frac{\text{tr}(U^\dagger V)}{\text{tr}(U^\dagger U)} \right|^2. \quad (13)$$

This is based on the Hilbert–Schmidt inner product between two unitary operators [44], with the square-modulus value taken to remove global phase differences and the denominator acting to normalize the fidelity for the dimension of the underlying Hilbert space. Many other fidelity definitions are in use [45–50], some of which are more suitable for the experimental measurement of fidelities, but for comparing two unitary operations, these are all closely related.

The lower line, plotted in blue, shows the infidelity for the first three-qubit circuit implemented on iodotrifluoroethene, in which the single-qubit target phases are all zero. Clearly, this circuit works well, achieving an acceptable fidelity with a time resolution of $1 \mu\text{s}$, which is easily achievable in NMR experiments. The upper line, plotted in red, corresponds to the second three-qubit circuit, where the single-qubit target phases are all π . This

circuit is unacceptably sensitive, requiring an unrealistic time resolution of about 1 ns to reach an acceptable fidelity.

The sensitivity of the second circuit arises because it uses evolution at the one-spin frequencies to achieve one-qubit phases. In NMR, one-spin frequencies (typically several kilohertz) are much larger than the coupling frequencies that generate two-qubit phases (typically tens of hertz) and so very small errors in evolution time will give large phase errors. (This will be less of a problem in other technologies where the qubits are individually addressed, as the offset frequency can in principle be chosen at will and so be set near zero.) The first circuit overcomes this by using pairs of *identical* time periods to cancel the single-spin evolution entirely: when these identical time periods are rounded, they remain identical and so perfectly generate the desired single-qubit phase equal to 0. If a nonzero single-qubit phase is required, it can be generated in other ways, most simply by altering the relative phase of two π pulses [36,37].

However, this restriction is not sufficient to guarantee a stable solution. The four-qubit network described in Eq. (11) has all single-qubit phases equal to zero and yet contains time periods without corresponding balancing pairs. The resulting infidelity is shown as the upper red curve in Fig. 5 and shows a very erratic dependence on time resolution. Fortunately, there is an easy method to stabilize such networks. If a reduced sign matrix R is replaced by its negative $-R$ with the evolution times unchanged, then this negated network will clearly generate identical two-spin phases to the original network (since the two-spin signs are negated twice and so left unchanged) but the opposite one-spin phases. Thus if a matrix R is combined with $-R$ and all the evolution times are halved, the combined network will generate the original two-spin phases but will have all one-spin phases equal to zero, even in the presence of rounding errors in time periods. The infidelity of the stabilized four-qubit network is shown as the lower blue curve in Fig. 5 and shows a greatly improved dependence on time resolution.

By its construction, the stabilized network will require the same total evolution time as the original network and the same computational time to find it, and so the cost of stabilization is confined to a doubling in the number of time periods and π pulses required for the implementation. This doubling can be reduced by subsequent optimization, locating equivalent evolutions in R and $-R$. Of the nine sign combinations found in Eq. (11), six can also be found in $-R$ and so the corresponding evolution times can be combined. Thus the symmetrized circuit can be reduced from 18 periods down to only 12 distinct evolution periods.

As symmetrization is effectively essential to generate a stable network in the presence of large single-qubit interactions, it is simpler just to generate symmetric networks by design. Since such networks are guaranteed to remove single-qubit phases, there are only $p = q(q-1)/2$




FIG. 5. The infidelity of a four-qubit circuit (upper curve in red) and its stabilized equivalent (lower curve in blue) as a function of the precision of the underlying clock.

two-qubit phases that can be controlled. These phases can be controlled using only p underlying times, each of which must occur twice, once in R and once in $-R$. Thus it is unsurprising that a symmetrized four-qubit network contains 12 time periods. Such networks can be found by setting no constraints on the single-qubit phases, since these will be cancelled by the second half of the network. It is sufficient to use only columns from the first half of S in the linear programming search, as columns from the second half will be used in $-R$. Thus the search for symmetrized networks is actually faster than for direct networks.

VII. RESCALING IN LARGER SPIN SYSTEMS

The fundamental problem that slows down the rescaling algorithm is the exponentially large size of the overcomplete basis. This basis contains $s = 2^q$ columns, from which the linear programming selects at most $r = q(q+1)/2$, equal to the number of interactions, and in many cases fewer. For moderate values of q , these numbers r and s are quite similar, but the difference grows rapidly with q , leading one to wonder whether there is some way to cut down the size of S before starting the linear programming step. Is it really necessary to include a very large number of columns, the great majority of which will eventually be discarded? In this regard, we note that once one has identified the appropriate reduced sign matrix, then the times required can be found by direct inversion of this $r \times r$ square matrix. However, it is clear that most of the hard work is done in locating the appropriate columns used to construct the reduced sign matrix.

Starting from the other extreme, one could just select r random columns from the full sign matrix and try to invert this. However, this process can fail in a number of ways. First, the reduced matrix might not be of full rank (although this is unlikely when $s \gg r$) and so may not have an inverse. Even for a reduced matrix with full rank, the set of times obtained from the inversion process is very

likely to include some negative times that are not physically implementable. Lastly, even if all the times are non-negative, the total time will not normally be the desired minimum and so the sequence will not be time optimal.

Between these extremes there is a middle way: using linear programming on a reduced, but still overcomplete, basis set. The current linear programming approach starts from the largest conceivable basis set, containing all of the $s = 2^q$ possible sign patterns, which guarantees the finding of an optimal solution but also makes the process slow. One might imagine choosing some subset of columns at random and attempting linear programming on this subset. For large values of q , the gap between the full size $s = 2^q$ and the minimum size r becomes very significant. It is thus worth exploring how many columns need to be picked so that linear programming generally finds a solution. There is no guarantee that such solutions will be time optimal, but as long as the random choice contains all the components of at least one optimal solution, then linear programming will find this solution. Given the very large number of equivalent solutions identified for moderate values of q , it seems plausible that this could be achieved with quite a small subset.

A. The RROS method

In the random-reduced-overcomplete-set (RROS) method, instead of using all s columns of the sign matrix S , as in the exhaustive approach, we choose just kr columns from S , for some $k > 1$, and run linear programming. Of course, one does not have to explicitly construct the entire S matrix and then choose the kr columns, as these kr columns correspond to the binary representations of kr distinct decimal numbers chosen at random from 0 to $s - 1$. Our experience so far suggests that the probability of finding a possible solution that achieves the desired phases using only positive evolution times increases as k increases, with a transition point around $k = 2$, at which the probability of a random set giving a solution reaches 50%.

This observation is substantiated by the empirical evidence in Fig. 6. RROS is run 500 times for values of q ranging from 10 to 22 with k varied between 1 and 3, and the fraction of occasions f when linear programming finds a suitable solution is calculated. The error bars on these estimates are calculated using Bayesian credible intervals [51], corresponding to the region of the probability density function within $\pm 34\%$ of the median, equivalent to one standard deviation for a normal distribution. It is clear that the probability of success rises sharply as k passes some critical value, with this transition becoming sharper as q is increased. To help locate this transition point, a sigmoidal logistic function [51]

$$f(k) = \frac{1}{1 + \exp[-b(x - c)]} \quad (14)$$




FIG. 6. The effectiveness of the RROS method selecting kr columns at random, showing how the probability of the algorithm finding a solution depends on q and k . Calculations are performed for q ranging from 10 to 22 but only four selected values are shown. The error bars on the $q = 12$ points are estimated using Bayesian credible intervals; error bars for other curves are very similar. Fitted sigmoidal curves are used to estimate the value of k at which this success probability reaches 50%. The inset shows the location of this transition point as a function of q . The error bars are estimated using error propagation from the sigmoidal-fit results; the smooth curve has no significance and is plotted simply to guide the eye.

is fitted for each value of q , with c being the transition point at which the success probability passes 50% and b indicating the sharpness of the transition. Although this function is chosen for convenience, it clearly fits fairly well. A plot of the value of c as a function of q , shown as an inset to Fig. 6, suggests that the transition lies just above $k = 2$. For large values of q , this transition becomes sharp, so that for $k \geq 2.5$ it is almost certain that a solution will be found.

If a solution is found, then this solution will be time optimal for the subset of columns chosen, but there is no guarantee that this will be the overall optimum, taking the shortest possible time. Unsurprisingly, the probability of finding a solution reaching the shortest possible time increases as k increases, but investigating this question in detail is challenging, as for large values of q the overall time-optimal solution cannot be located in a reasonable time. Nevertheless, our preliminary studies suggest that, for large q , the quality of solutions plateaus around $k \approx 4$ and so there is little point going beyond this in practice. For small values of q , it seems to be necessary to use a slightly larger value of k , but in these cases it is more sensible just to use direct solution of the full S matrix anyway.

B. Results for large numbers of qubits

The use of a smaller basis set permits RROS to be extended to much larger values of q . This is investigated by running the algorithm for q between 10 and 60, as shown in Fig. 7. For RROS, the time required to perform linear programming depends not only on the randomly chosen

Hamiltonian and target phases but also on the precise choice of columns, and so timings are repeated ten times using different choices. As expected, for $k = 4$ a solution is located in every case.

The simplex algorithm is found to be slightly faster than the interior-point algorithm for $q \leq 31$, but the interior-point algorithm is faster for $q \geq 32$ and becomes much faster at high q . The discontinuity in the interior-point timings between $q = 31$ and $q = 32$ may indicate a change in the precise algorithm used by MATLAB.

For RROS, the dimension of the problem is $n = r + kr$ and so we expect a computational complexity between $O(r^2)$ and $O(r^3)$. The smooth curve in the figure is fitted to the timings from the interior-point algorithm for the range $32 \leq q \leq 60$ using a power function, and takes the form Ar^3 , consistent with our expectations (including the power as an additional fitting parameter gives a value indistinguishable from 3). The timings for the simplex algorithm appear to lie between $O(r^4)$ and $O(r^5)$, which is much slower than the interior-point algorithm.

With the simplex algorithm, the number of time periods required is always approximately equal to r (in some cases, a small number of time periods are negligibly short and so can be dropped). The number of pulses required is determined for the R matrix as originally found, without further optimization, and is typically around $qr/2 \approx q^3/4$, confirming our previous expectations. With the interior-point algorithm, the solutions always use all kr time periods, with a large number of very short times rather than a clear division into zero and nonzero times. These initial




FIG. 7. The computation time required to perform linear prediction using RROS with $k = 4$. The simplex method is found to be fastest for $q \leq 31$, with the interior-point method faster for $q \geq 32$. The error bars for $q \leq 60$ indicate the mean and SD over ten repetitions. The smooth curve with the form Ar^3 is fitted to interior-point timings in the range $32 \leq q \leq 60$. Data points for $q > 60$ are single repetitions and are not included in the fit. The fact that these points lie very close to the extrapolated fit suggests that our fitted curve is a good model. The red points on the left show timings for exhaustive calculations, demonstrating the huge time gains possible with RROS.

solutions can then be simplified by choosing a smaller subset of $k'r$ columns, with $k' \approx 1.2$, corresponding to the columns with the largest time values in the original solution. Surprisingly, the computational time for these recalculations appears to be linear in k , while a quadratic or cubic dependence might have been expected.

The calculation is then extended to higher values of q , using only a single repetition, and these data points are found to lie very close to the extrapolated curve. As $r \approx q^2/2$, this gives a pragmatic computational time scaling of $O(q^6)$, which is polynomial in the number of qubits. Extrapolating this curve still further suggests that calculations with up to about 150 qubits could be practical, but the MATLAB implementation `linprog` runs out of memory above $q = 125$. This limit should be solvable by a custom implementation of linear programming.

C. Stability of solutions

With large numbers of qubits, it is not easy to directly evaluate the fidelity of a solution, as even writing down the diagonal elements of the underlying propagator requires 2^q terms. It is, however, possible to estimate the infidelity directly using

$$\mathcal{I} \approx \sum_{i < j} \frac{(\delta\phi_{ij})^2}{16}, \quad (15)$$

where $\delta\phi_{ij}$ is the error in the coupling angle for the pair of qubits i and j and the approximation is only valid when the infidelity is small. This assumes that the solution has been symmetrized so that the errors in the single-spin angles are all zero. Networks with large numbers of qubits have higher infidelities than those shown in Sec. VI, reflecting the larger number of phases that need to be controlled. This is not a specific limitation of our methods but is inherent in any attempt to control a large number of qubits simultaneously.

VIII. CONCLUSIONS

We present a practical algorithm that is guaranteed to find the minimum time solution to rescaling z and zz terms in the internal Hamiltonian of a quantum computer with up to about 20 qubits. Above 20 qubits, the direct approach becomes intractable but random sampling will extend this (although the solution might not be quite time optimal) to more than 100 qubits. Beyond about 150 qubits, the q^6 scaling of the computational time for a fully coupled system renders any known approach impractical. It is, however, very unlikely that any system of that size will still be fully coupled, as couplings are usually only substantial between nearby spins. In partially coupled systems, r is linear in q rather than quadratic and allowing for this should permit the process to be extended to hundreds or even thousands of qubits.

Although described in the language of NMR, the techniques used are applicable to any equivalent system, where an always-on two-qubit interaction commutes with the single-qubit background terms. They may be of particular value in solid-state platforms, such as superconducting circuits, in which two-dimensional lattices of qubits are developed with a sparse coupling network (generally nearest-neighbor). In such very large systems, it is likely that the underlying structure in the pattern of interactions will allow symmetries to be exploited, potentially simplifying the problem greatly.

ACKNOWLEDGMENTS

We thank Paul Goldberg (Department of Computer Science, University of Oxford) for useful discussions. G.B. is supported by a Felix Scholarship. T.T. is supported by the Masason Foundation and the Nakajima Foundation.

APPENDIX A: PULSES FOR REFOCUSING

Consider a system of q fully coupled qubits, where we seek to refocus all the one-qubit interactions and retain one single two-qubit interaction. This is most efficiently achieved by assigning the two coupled qubits to the Walsh function W_1 , with the remaining qubits assigned from W_2 up to W_{q-1} in sequence. As W_n contains n sign changes, it requires $n\pi$ pulses when n is even and $n+1$ pulses when n is odd, with the additional pulse needed to restore the Hamiltonian to its original sign.

The overall number of pulses required is then obtained by summing these pulse counts from $n=1$ to $q-1$, remembering to include W_1 twice, giving a total of

$$\frac{q(q-1)}{2} + \left\lceil \frac{q-1}{2} \right\rceil + 2 \leq \frac{q^2}{2} + 2.$$

Thus the total number of pulses required to retain a single two-qubit interaction is approximately $q^2/2$.

APPENDIX B: LINEAR PROGRAMMING

The general linear programming problem seeks to maximize some linear function of its inputs subject to a set of linear equality and inequality constraints. In our application, we seek to minimize the total evolution time: this is easily achieved by maximizing the negative of the total time and our only inequality constraints are that all inputs (the individual evolution times for each period) must be positive. The equality constraints then describe a convex polytope (a multidimensional polyhedron) called a simplex.

It can be shown that the solution to the linear programming problem always lies on one of the vertices and so it is only necessary to compare the values of the function at all its vertices. However, for a high-dimensional problem,

the number of vertices becomes very large and direct search is simply impractical. Dantzig's original simplex algorithm provides a pragmatic search strategy that usually locates the optimum in a reasonable time. Typically, this time is no worse than $O(n^3)$, where n , the dimension of the problem, is in our case the sum of the number of inputs and the number of equality constraints. However, it is known that particular problems can have worse scaling and, indeed, some problems require full exhaustive comparison of every vertex.

A second common approach, the interior-point algorithm, does not confine itself to vertices but instead travels across the inside of the polytope in an attempt to maximize the desired function while seeking not to cross the boundary faces of the polytope. This algorithm can be either faster or slower than Dantzig's algorithm, depending on details of the problem.

An important distinction between the two algorithms is that simplex will locate the true optimal vertex, and for our problem this has the property that some (typically, the great majority) of the input times are strictly equal to zero, allowing these times to be dropped in any experimental implementation. By contrast, the interior-point algorithm only locates an approximate optimum, near one of the vertices, and so the solution contains a large number of input times that are small but not quite zero. These erroneous terms cannot simply be dropped without thought, as they act to cancel other error terms elsewhere in the solution. An important aspect of the stabilization method discussed in Sec. VI is that it enables these error terms to be dropped with comparative safety.

-
- [1] K. R. Brown, J. Kim, and C. Monroe, Co-designing a scalable quantum computer with trapped atomic ions, [npj Quantum Inf.](#) **2**, 16034 (2016).
 - [2] V. M. Schäfer, C. J. Ballance, K. Thirumalai, L. J. Stephenson, T. G. Ballance, A. M. Steane, and D. M. Lucas, Fast quantum logic gates with trapped-ion qubits, [Nature](#) **555**, 75 (2018).
 - [3] J. A. Jones, Quantum computing with NMR, [Prog. NMR Spectrosc.](#) **59**, 91 (2011).
 - [4] A. Barenco, C. H. Bennett, R. Cleve, D. P. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J. A. Smolin, and H. Weinfurter, Elementary gates for quantum computation, [Phys. Rev. A](#) **52**, 3457 (1995).
 - [5] B. E. Kane, A silicon-based nuclear spin quantum computer, [Nature](#) **393**, 133 (1998).
 - [6] D. Loss and D. P. DiVincenzo, Quantum computation with quantum dots, [Phys. Rev. A](#) **57**, 120 (1998).
 - [7] G. K. Brennen, C. M. Caves, P. S. Jessen, and I. H. Deutsch, Quantum Logic Gates in Optical Lattices, [Phys. Rev. Lett.](#) **82**, 1060 (1999).
 - [8] D. Jaksch, H.-J. Briegel, J. I. Cirac, C. W. Gardiner, and P. Zoller, Entanglement of Atoms via Cold Controlled Collisions, [Phys. Rev. Lett.](#) **82**, 1975 (1999).

- [9] E. Knill, R. Laflamme, R. Martinez, and C. Negrevergne, Benchmarking Quantum Computers: The Five-Qubit Error Correcting Code, *Phys. Rev. Lett.* **86**, 5811 (2001).
- [10] M. N. Leuenberger and D. Loss, Quantum computing in molecular magnets, *Nature* **410**, 789 (2001).
- [11] W. Harneit, Fullerene-based electron-spin quantum computer, *Phys. Rev. A* **65**, 032322 (2002).
- [12] D. Suter and K. Lim, Scalable architecture for spin-based quantum computers with a single type of gate, *Phys. Rev. A* **65**, 052309 (2002).
- [13] Z. Chen, Y. Zhou, and J.-T. Shen, in *Conference on Lasers and Electro-Optics* (Optical Society of America, San Jose, 2019), p. FTh3A.4.
- [14] A. Parra-Rodriguez, P. Lougovski, L. Lamata, E. Solano, and M. Sanz, Digital-analog quantum computation, arXiv:1812.03637 (2018).
- [15] Y. Wu, Y. Wang, X. Qin, X. Rong, and J. Du, A programmable two-qubit solid-state quantum processor under ambient conditions, *npj Quantum Inf.* **5**, 9 (2019).
- [16] D. L. Moehring, P. Maunz, S. Olmschenk, K. C. Younge, D. N. Matsukevich, L.-M. Duan, and C. Monroe, Entanglement of single-atom quantum bits at a distance, *Nature* **449**, 68 (2007).
- [17] F. Arute *et al.*, Quantum supremacy using a programmable superconducting processor, *Nature* **574**, 505 (2019).
- [18] J. M. Gambetta, J. M. Chow, and M. Steffen, Building logical qubits in a superconducting quantum computing system, *npj Quantum Inf.* **3**, 2 (2017).
- [19] E. L. Hahn, Spin echoes, *Phys. Rev.* **80**, 580 (1950).
- [20] R. Freeman, *Spin Choreography* (Spektrum, Oxford, 1997).
- [21] J. A. Jones and E. Knill, Efficient refocusing of one-spin and two-spin interactions for NMR quantum computation, *J. Magn. Reson.* **141**, 322 (1999).
- [22] D. W. Leung, I. L. Chuang, F. Yamaguchi, and Y. Yamamoto, Efficient implementation of coupled logic gates for quantum computation, *Phys. Rev. A* **61**, 42310 (2000).
- [23] D. Leung, Simulation and reversal of n -qubit Hamiltonians using Hadamard matrices, *J. Mod. Opt.* **49**, 1199 (2002).
- [24] D. Hayes, S. T. Flammia, and M. J. Biercuk, Programmable quantum simulation by dynamic Hamiltonian engineering, *New J. Phys.* **16**, 083027 (2014).
- [25] J. Welch, D. Greenbaum, S. Mostame, and A. Aspuru-Guzik, Efficient quantum circuits for diagonal unitaries without ancillas, *New J. Phys.* **16**, 033040 (2014).
- [26] K. G. Beauchamp, *Applications of Walsh and Related Functions* (Academic Press, London, 1984).
- [27] M. S. Lynn, On the Schur product of H -matrices and non-negative matrices, and related inequalities, *Math. Proc. Camb. Philos. Soc.* **60**, 425 (1964).
- [28] R. G. Bland, The allocation of resources by linear programming, *Sci. Am.* **244**, 126 (1981).
- [29] G. B. Dantzig, Reminiscences about the origins of linear programming, *Oper. Res. Lett.* **1**, 43 (1982).
- [30] I. Adler, M. Resende, G. Veiga, and N. Karmarkar, An implementation of Karmarkar's algorithm for linear programming, *Math. Program.* **44**, 297 (1989).
- [31] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes in C* (Cambridge University Press, Cambridge, 1992), 2nd ed.
- [32] J. Du, J. Zhu, M. Shi, X. Peng, and D. Suter, Experimental observation of a topological phase in the maximally entangled state of a pair of qubits, *Phys. Rev. A* **76**, 042121 (2007).
- [33] E. Knill, R. Laflamme, R. Martinez, and C. H. Tseng, An algorithmic benchmark for quantum information processing, *Nature* **404**, 368 (2000).
- [34] M. D. Bowdrey, J. A. Jones, E. Knill, and R. Laflamme, Compiling gate networks on an Ising quantum computer, *Phys. Rev. A* **72**, 032315 (2005).
- [35] R. Freeman, T. A. Frenkiel, and M. H. Levitt, Composite Z pulses, *J. Magn. Reson.* **44**, 409 (1981).
- [36] J. J. L. Morton, A. M. Tyryshkin, A. Ardavan, S. C. Benjamin, K. Porfyakis, S. A. Lyon, and G. A. D. Briggs, Bang-bang control of fullerene qubits using ultrafast phase gates, *Nat. Phys.* **2**, 40 (2006).
- [37] J. A. Jones, Designing short robust NOT gates for quantum computation, *Phys. Rev. A* **87**, 052317 (2013).
- [38] N. Khaneja, T. Reiss, C. Kehlet, T. Schulte-Herbrüggen, and S. J. Glaser, Optimal control of coupled spin dynamics: Design of NMR pulse sequences by gradient ascent algorithms, *J. Magn. Reson.* **172**, 296 (2005).
- [39] N. Boulant, E. M. Fortunato, M. A. Pravia, G. Teklemariam, D. G. Cory, and T. F. Havel, Entanglement transfer experiment in NMR quantum information processing, *Phys. Rev. A* **65**, 024302 (2002).
- [40] D. Collins, K. W. Kim, W. C. Holton, H. Sierzputowska-Gracz, and E. O. Stejskal, NMR quantum computation with indirectly coupled gates, *Phys. Rev. A* **62**, 022304 (2000).
- [41] N. Khaneja, B. Heitmann, A. Spörl, H. Yuan, T. Schulte-Herbrüggen, and S. J. Glaser, Shortest paths for efficient control of indirectly coupled qubits, *Phys. Rev. A* **75**, 012322 (2007).
- [42] J. A. J. Hall, Towards a practical parallelisation of the simplex method, *Comput. Manage. Sci.* **7**, 139 (2010).
- [43] J. A. Jones, Robust Ising gates for practical quantum computation, *Phys. Rev. A* **67**, 012317 (2003).
- [44] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (CUP, Cambridge, 2000).
- [45] I. L. Chuang and M. A. Nielsen, Prescription for experimental determination of the dynamics of a quantum black box, *J. Mod. Opt.* **44**, 2455 (1997).
- [46] M. D. Bowdrey, D. K. L. Oi, A. J. Short, K. Banaszek, and J. A. Jones, Fidelity of single qubit maps, *Phys. Lett. A* **294**, 258 (2002).
- [47] M. A. Nielsen, A simple formula for the average gate fidelity of a quantum dynamical operation, *Phys. Lett. A* **303**, 249 (2002).
- [48] H. F. Hofmann, Complementary Classical Fidelities as an Efficient Criterion for the Evaluation of Experimentally Realized Quantum Operations, *Phys. Rev. Lett.* **94**, 160504 (2005).
- [49] L. H. Pedersen, N. M. Møller, and K. Mølmer, Fidelity of quantum operations, *Phys. Lett. A* **367**, 47 (2007).
- [50] K. Mayer and E. Knill, Quantum process fidelity bounds from sets of input states, *Phys. Rev. A* **98**, 052326 (2018).
- [51] J. K. Kruschke, *Doing Bayesian Data Analysis* (Academic Press, Oxford, 2011).