

# Robust quantum compilation and circuit optimisation via energy minimisation

Tyson Jones and Simon C. Benjamin

Department of Materials, University of Oxford, Parks Road, Oxford OX1 3PH, UK

We explore a method for automatically recompiling a quantum circuit  $\mathcal{A}$  into a target circuit  $\mathcal{B}$ , with the goal that both circuits have the same action on a specific input i.e.  $\mathcal{B}|\text{in}\rangle = \mathcal{A}|\text{in}\rangle$ . This is of particular relevance to hybrid, NISQ-era algorithms for dynamical simulation or eigensolving. The user initially specifies  $\mathcal{B}$  as a blank template: a layout of parameterised unitary gates configured to the identity. The compilation then proceeds using quantum hardware to perform an isomorphic energy-minimisation task, and an optional gate elimination phase to compress the circuit. If  $\mathcal{B}$  is insufficient for perfect recompilation then the method will result in an approximate solution. We optimise using imaginary time evolution, and a recent extension of quantum natural gradient for noisy settings. We successfully recompile a 7-qubit circuit involving 186 gates of multiple types into an alternative form with a different topology, far fewer two-qubit gates, and a smaller family of gate types. Moreover we verify that the process is *robust*, finding that per-gate noise of up to 1% can still yield near-perfect recompilation. We test the scaling of our algorithm on up to 20 qubits, recompiling into circuits with up to 400 parameterized gates, and incorporate a custom adaptive timestep technique. We note that a classical simulation of the process can be useful to optimise circuits for today’s prototypes, and more generally the method may enable ‘blind’ compilation i.e. harnessing a device whose response to control parameters is deterministic but unknown.

The code and resources used to generate our results are openly available online [1, 2]. A simple Mathematica demonstration of our algorithm can be found at [questlink.quechtheory.org](http://questlink.quechtheory.org).

Tyson Jones: [tyson.jones@materials.ox.ac.uk](mailto:tyson.jones@materials.ox.ac.uk)

Simon C. Benjamin: [simon.benjamin@materials.ox.ac.uk](mailto:simon.benjamin@materials.ox.ac.uk)

## Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>Variational Simulation</b>	<b>3</b>
2.1	Li’s real time simulation . . . . .	3
2.2	Imaginary time evolution . . . . .	3
2.3	Noisy natural gradient . . . . .	3
<b>3</b>	<b>Overview of compilation</b>	<b>4</b>
3.1	Compilation by energy minimisation . .	4
3.2	The <i>Lure</i> method . . . . .	5
3.3	Adaptive timestep . . . . .	5
3.4	Further gate elimination . . . . .	6
<b>4</b>	<b>Numerical demonstration</b>	<b>6</b>
4.1	Selecting an interesting example circuit $\mathcal{A}$	6
4.2	Performance of the recompilation . . . .	7
<b>5</b>	<b>Robustness to noise</b>	<b>10</b>
<b>6</b>	<b>Numerical scaling</b>	<b>12</b>
6.1	Scaling performance . . . . .	13
<b>7</b>	<b>Possible applications</b>	<b>14</b>
<b>8</b>	<b>Conclusion</b>	<b>15</b>
<b>A</b>	<b>Demonstration of the Lure method</b>	<b>16</b>
<b>B</b>	<b>Li’s Algorithm and Trotter comparison</b>	<b>17</b>
<b>C</b>	<b>The imaginary time eigensolver</b>	<b>19</b>
<b>D</b>	<b>Details of numerical simulation</b>	<b>19</b>
	<b>References</b>	<b>20</b>

## 1 Introduction

In conventional computing, compilers are essential to translate programs into efficient low-level instructions for execution at the hardware level. Quantum computers will also benefit greatly from efficient compilation, but the nature of the compilation goal depends on the nature of the machine. In the immediate future we are entering the so-called NISQ era of noisy, intermediate scale quantum devices [3]. In the further future we expect to see the emergence of fully fault-tolerant, code-protected quantum processors which can be regarded by programmers as error-free.

For fault-tolerant devices, typically only a limited non-universal family of operations can be performed directly on encoded data, so that other (usu-

ally non-Clifford) operations must be performed with the use of additional resources such as magic states [4–6]. Therefore the priority for computation will be to minimise the number of these expensive resource-consuming gates. Substantial efforts have been made to understand how to minimise the number of non-Clifford gates, such as  $T$ -gates, that are required to perform a given task [7–11].

For the era of Noisy Intermediate Scale Quantum (NISQ) devices, the priorities for compilation will be different. Here we may expect that information is stored without the full protection of error-correcting codes and that therefore the difficulty that codes cannot permit universal operations does not arise. Rather, the costly gates are those with the greatest error burden. Typically these are two-qubit gates (and higher degree gates) in today’s prototypes, while single-qubit gates are higher fidelity [12–15]. Moreover a given physical device will have certain operations that are native to it, so that e.g. it may be that a control-NOT is impossible to implement directly but is instead realised through a parity-dependent phase shift together with additional single-qubit gates. Furthermore the device will have a native connectivity: certain qubits will be able to directly link to others, for example in a two-dimensional nearest-neighbour topology or a more flexible networked architecture [16]. Thus one would wish to compile directly to the device’s native gate set and connectivity.

To take further the remark about a device’s native operations: it may be that issues such as cross-talk mean that an operation targeted to a specific qubit inevitably leads to an unwanted, but deterministic, effect on proximal qubits. One can regard that operation as a kind of native gate itself (see e.g. Ref. [17]), albeit one that may be difficult to work with in analytic treatment. A compiler that is capable of targeting an arbitrary family of gates could recompile from a standard ‘white board’ description of a circuit into a truly native format where the gate operations are bespoke for a specific device. As we presently discuss, taking this idea further one could compile into a device whose gates are an unknown function of the control parameters.

This paper describes a general method of translating one quantum circuit into another, i.e. recompiling it. The approach can in principle allow one to do the following:

- Target an arbitrary (user-specified) circuit layout,
- Target an arbitrary (user-specified) set of gates, including bespoke gates not used in analytic treatments,
- Support approximate recompilation, so that if the specified target template is too shallow for perfect recompilation then an approximate circuit will be found,

- Minimise the impact of noise.

We note that if one wished, the user-specification of the target layout and gates could in future be relaxed following VQA studies using ansätze that auto-evolve [18–20].

The present scheme is limited in important ways:

- Compilation of circuits beyond the classical simulation limit will require quantum hardware, and will consume considerable time on that hardware. We note that all-classical software to recompile circuits involving parameterised gates does exist [11] and can make significant savings. However no classical compiler can be expected to approach optimality for general circuits since even the task of verifying that two circuits are near-identical is QMA-complete [21].
- Compilation from the original circuit  $\mathcal{A}$  is not to an equivalent unitary circuit, but rather to a target circuit  $\mathcal{B}$  that (ideally) has the same effect on just one specific input state  $|\text{in}\rangle$ , so that  $\mathcal{B}|\text{in}\rangle = \mathcal{A}|\text{in}\rangle$ . This is a profoundly more permissive goal, but is in fact the *right* goal for many quantum algorithms including so-called hybrid quantum-classical approaches [22].
- While the specific input state  $|\text{in}\rangle$  can have any form, it is necessary that we ‘understand’ it well enough to be able to write down a (fictitious) Hamiltonian for which it is the ground state.
- The approach we describe here has unproven scaling as the circuits grow beyond 20 qubits. We remark on this point presently.
- We restrict our attention to circuits formed of unitary gates, so that our complete circuits  $\mathcal{A}$  and  $\mathcal{B}$  are themselves unitary. Generalising to non-unitary circuits would appear possible however.
- A more comprehensive compiler might automatically propose and test different templates for  $\mathcal{B}$ , rather than requiring the user to specify one. This would be a higher-level process operating above the compilation we describe; prior work on all-classical optimisation could be employed here [11].

We note other recent efforts to perform approximate recompilation of quantum circuits. For example, a scheme presented by Khatri *et al* describes a method to recompile a (possibly unknown)  $n$ -qubit unitary  $U$  into a parametrised circuit  $V$  [23], by maximising their *overlap*. Their cost function is  $|\text{Tr}(V^\dagger U)|^2$ , evaluated via the Hilbert-Schmidt test, requiring  $2n$  qubits and  $2n$  additional controlled gates. They demonstrate their technique with noise-free and noisy simulations

of up to 9 qubits. They thus focus on the less permissive task of recompiling the entire unitary as opposed to its action on a fixed input state, while noting the possibility of fixed-input schemes. Compared to the present work, that study [23] presents a broader range of techniques and includes 1-qubit demonstrations on cloud-based quantum hardware, but investigates considerably more shallow circuits as the input (e.g. 4 rather than 60 layers) and smaller registers (simulating a maximum of 9 qubits rather than the 20 considered here). To perform our larger and more complex compilation tasks, we introduce certain enabling techniques described presently, including adaptive timesteps, elimination of redundant gates, and a ‘lure’ method to improve convergence.

Meanwhile an experimentally-focused effort [24] performed a similar scheme on a photonic processor of 4 qubits, with numerical testing on up to 5 qubits. There, they recompile an unknown unitary  $U$  acting on a fixed, restrictedly separable state  $|\text{in}\rangle$ , and use an ansatz divided into multiple layers  $V_l$ , each targeting one fewer qubits. The parameters in each layer are successively trained by driving the layer’s nominated qubit toward its initial state in the input separable state, i.e. toward parameters which satisfy  $(\prod_l V_l)^\dagger U |\text{in}\rangle \approx |\text{in}\rangle$ .

In contrast to these works, our fixed-input method requires only  $n$  qubits, with a cheap cost function evaluated via measuring a fictitious Hamiltonian. We make no conditions on the input state. When our input unitary happens to be parameterized, our scheme is compatible with an additional novel *luring* technique, which can significantly improve convergence on systems larger than tested by the previous works. We test our scheme on up to 20 qubits, and perform a demonstration of a possible application; extending variational real-time simulation. We also make a comprehensive, density matrix based analysis of the impact of gate noise for a scenario involving over 200 gates.

## 2 Variational Simulation

Our recompilation method is an instance of a variational quantum algorithm; VQA’s have been explored widely for their compatibility with first generation quantum computers. Two recent reviews of VQAs are available which cover the extensive and growing literature [25, 26]. As a subroutine, our algorithm performs energy minimisation, for which there are many quantum algorithms [22, 27], though in this work we opt to use deterministic imaginary-time evolution [28] and a noise-compatible form [29] of quantum natural gradient [30], which we introduce below. To demonstrate an interesting application of our recompilation strategy (Section 4), we will employ a variational real-time simulation algorithm to produce input states. We also summarise this algorithm, which we refer to

as Li’s method [31].

### 2.1 Li’s real time simulation

Li’s method [31] is a variational algorithm for real-time quantum simulation. It prescribes a succession of parameters  $\vec{\theta}_t$  to produce ansatz states

$$|\psi\rangle_t = \hat{U}(\vec{\theta}_t) |\text{in}\rangle, \quad (1)$$

which approximate the states produced from unitary time evolution of input state  $|\text{in}\rangle$  under some Hamiltonian  $\hat{H}$ ;

$$|\psi\rangle_t \approx \exp(-i\hat{H}t) |\text{in}\rangle. \quad (2)$$

In this manuscript, we merely use Li’s method to generate interesting input states to our recompilation strategy. We include a more thorough description of Li’s method in Appendix B, including a short comparison with Trotterisation. See Reference [27] for a detailed summary of variational quantum simulation techniques.

### 2.2 Imaginary time evolution

Imaginary time evolution is a variational algorithm for performing energy minimisation [28], or VQE [32]. It requires the evaluation of a metric object that reconciles between parameter space and Hilbert space, recently shown to be an instance of quantum natural gradient [30]. It is incidentally a Wick rotation of the equations deriving Li’s method, and drives the ansatz state toward the ground-state of the simulated Hamiltonian;

$$|\psi\rangle_t \approx N(t) \exp(-\hat{H}t) |\text{in}\rangle, \quad (3)$$

where  $N(t) = 1/\sqrt{\langle \text{in} | \exp(-2\hat{H}t) | \text{in} \rangle}$  is a normalisation factor. In this work, we use imaginary time evolution in our noise-free simulations as a subroutine to drive an ansatz state into the ground-state of a constructed Hamiltonian. In principle, any generic variational minimisation technique can be employed. However, imaginary time has been found in earlier studies to perform favourably for parameter optimisation [27, 28, 32], and exhibits good resilience to simple decoherence and shot noise [28]. Additionally, it was recently shown to require only marginally more sampling than simple gradient descent [33]. An elaboration of the role of imaginary time evolution is given in Appendix C.

### 2.3 Noisy natural gradient

For simulations of our recompilation algorithm on noise-burdened quantum computers, we replace imaginary time minimisation with an improved method for noisy settings. Recently, the metric evaluated in

quantum natural gradient [30] was adapted for non-unitary ansatzes, using the quantum Fisher information matrix [29]. Moreover, in that work it was established that for the case of noisy circuits producing output  $\rho$ , the elements  $F_{k,l}$  of the metric can be well-approximated by

$$F_{k,l} = \text{Tr}[(\partial_k \rho)(\partial_l \rho)], \quad (4)$$

a quantity that can be evaluated experimentally via a swap-test circuit [34]. We will employ this generalisation of imaginary time evolution to the non-unitary case later, in the context of compilation with a noise-burdened quantum computer.

### 3 Overview of compilation

Consider a quantum circuit  $\mathcal{A}$  acting on an  $n$ -qubit register and involving some arbitrary set of  $m$  unitary gates  $G_i^A$ , each acting on one or more of the qubits. When the input state to the circuit is  $|\text{in}\rangle$ , the output is

$$|\text{out}\rangle = G_m^A G_{m-1}^A \dots G_1^A |\text{in}\rangle. \quad (5)$$

Now suppose that we wish to find an alternative circuit  $\mathcal{B}$  which has the same, or nearly the same, action on the input state. This new circuit acts upon the same  $n$ -qubit register but generally may contain a different number  $m'$  of gate operations  $G_i^B$ . These new gates may be a very different set from those in  $\mathcal{A}$ , even (if we wish) a set that is too restrictive to express the former set exactly.

The approach described here involves first compiling to a circuit  $\mathcal{B}^{-1}$ , which we define as the gate-by-gate inverse of  $\mathcal{B}$ . A comparable technique has been used for the related problem of learning an unknown state [35], where one seeks a mapping to a known target in order that the reverse process can define the original state. We write the individual gates as  $g_i^B = (G_i^B)^{-1}$  so that our circuits are,

$$\begin{aligned} \mathcal{B} &\rightarrow G_{m'}^B \dots G_2^B G_1^B \\ \mathcal{B}^{-1} &\rightarrow g_1^B g_2^B \dots g_{m'}^B. \end{aligned}$$

Importantly, the gates  $G_i^B$  are parameterised: each gate takes a single parameter  $\phi_i$ , such that  $G_i^B(\phi_i = 0)$  is the identity. Thus each  $G_i^B$  is really a continuous family of unitaries. An example would be  $\exp(i\phi \sigma_x^a \sigma_x^b)$  acting on qubits  $a$  and  $b$ . It follows that each inverse gate  $g_i^B$  is also a function of  $\phi_i$ . We are implicitly assuming that the gates are sufficiently simple (e.g. low enough qubit counts) that the mapping between  $G_i^B$  and  $g_i^B$  is tractable. Moreover, for the case that the recompilation process is using a quantum computer (as opposed to an all-classical implementation) we are implicitly assuming that  $G_i^B(\phi_i)$  and  $g_i^B(\phi_i)$  are physically implementable for all values of  $\phi_i$ . (One could tolerate certain restrictions, but we do not explore that here).

To recompile  $\mathcal{A}$ , the circuit is applied to the input  $|\text{in}\rangle$  and then circuit  $\mathcal{B}^{-1}$  is applied, ultimately producing state  $|\text{fin}\rangle$ .

$$|\text{fin}\rangle = (g_1^B g_2^B \dots g_{m'}^B) (G_m^A G_{m-1}^A \dots G_1^A) |\text{in}\rangle. \quad (6)$$

The compiler seeks to find a set of parameter values  $\vec{\phi}$  such that (as nearly as possible),

$$|\text{in}\rangle = \mathcal{B}^{-1}(\vec{\phi}) \mathcal{A} |\text{in}\rangle \Rightarrow \mathcal{B}(\vec{\phi}) |\text{in}\rangle = \mathcal{A} |\text{in}\rangle, \quad (7)$$

where it is understood that these equations are up to a meaningless global phase.

Initially  $\phi_i = 0, \forall i$  so that all gates  $g_i^B$  are simply the identity. As the ‘user’ we specify a *circuit template* for  $\mathcal{B}$ , since we fix the gate types and the sequence, but this template is ‘blank’ i.e. free of parameter information.

Generally finding satisfying  $\vec{\phi}$  is a hard search problem since there may be thousands of parameters even for NISQ-era machines. One must therefore select the strategy carefully, giving consideration to potential problems such as becoming ‘stuck’ in a local minimum as we evolve the parameter set. There are also issues relating to the device size: ideally recompilation will be achieved without the need for additional qubits.

The approach we take here coopts recent ideas relating to finding the ground state energy of a Hamiltonian. The device size remains  $n$  qubits, and moreover although the scaling and performance of such approaches are not fully understood there is a developing literature on these topics [22, 27, 36].

#### 3.1 Compilation by energy minimisation

As noted in our list of restrictions, we are assuming that we understand the specific input state  $|\text{in}\rangle$  sufficiently to create a Hamiltonian for which that state would be the unique ground state. We stress that this Hamiltonian does not correspond to any real physical system of interest, it is a fictitious construct purely to enable the recompilation process. We denote it  $H_{\text{rec}}$  where the subscript stands for ‘recompilation’. Finding  $H_{\text{rec}}$  is of course trivial for any input that has a product form: for example if  $|\text{in}\rangle = |00\dots 0\rangle$  then the obvious choice would be  $H_{\text{rec}} = \sum_j \sigma_j^z$ .

It may be desirable to ensure that there is a well-defined gap to the first excited state that is itself only  $n$ -fold degenerate (or to break such degeneracy if we wish). For the example just given, the gap is unity regardless of  $n$ . These properties are helpful in terms of the efficiency of the ground state finding protocol.

Given that we have selected a suitable  $H_{\text{rec}}$ , then the recompilation process has become an eigensolving task:

*Given a (fixed) input  $\mathcal{A}|\text{in}\rangle$  to ‘ansatz’ circuit  $\mathcal{B}^{-1}(\vec{\phi})$ , find the parameter values  $\vec{\phi}$  for which the output has the lowest possible energy with respect to  $H_{\text{rec}}$ .*

We can adopt any one of several quantum minimisation techniques to solve this problem [22]. As motivated in Section 2, we employ imaginary time evolution [28] for our noise-free simulations, and a non-unitary adaptation [29] for our noisy simulations. We outline these methods in Appendices B and C.

Any variational eigensolving technique may become slow to evolve in specific cases. Anticipating this problem, we have explored a solution involving a series of proximal targets which in effect ‘lure’ the process toward the eventual target. We present the technique in a following section, and offer an intuitive example in Appendix A. Note the specific demonstration of compilation in Section 4 does not in fact require such a lure (we do not observe any slowdown issues).

In the case that the compilation process is being performed with a quantum computer (rather than a classical emulation of the process), one may wonder whether the user would be able to determine how successful the compilation has been. A reasonable measure is the fidelity between the input state  $|\text{in}\rangle$  and its attempted reconstruction  $\mathcal{B}^{-1}\mathcal{A}|\text{in}\rangle$ , equivalent to that between  $\mathcal{A}|\text{in}\rangle$  and  $\mathcal{B}|\text{in}\rangle$ . We report this measure from our simulations, however in the lab a user may not be able to evaluate this directly. Fortunately one can lower-bound the fidelity using the expected energy  $\langle H_{\text{rec}} \rangle$  which is measurable (indeed the imaginary-time variational approach involves repeatedly estimating quantities of this kind). Since  $|\text{in}\rangle$  is the ground state of  $H_{\text{rec}}$  with energy  $E_0$ , and presuming that we ‘understand’ our fictitious Hamiltonian  $H_{\text{rec}}$  sufficiently to know its first excited state energy  $E_1$ , then

$$\min\{\langle H_{\text{rec}} \rangle\} - E_0 = (1 - F)(E_1 - E_0)$$

where the min denotes ‘minimum observable value given that the fidelity is  $F$ ’. It follows that

$$F \geq \frac{E_1 - \langle H_{\text{rec}} \rangle}{E_1 - E_0}. \quad (8)$$

Note that if, as in the example above, our fictitious Hamiltonian  $H_{\text{rec}}$  has a gap  $E_1 - E_0$  of unity then the accuracy with which we can bound  $F$  simply depends on the shot noise in our estimate  $\langle H_{\text{rec}} \rangle$ .

### 3.2 The Lure method

The present task of mapping the state  $\mathcal{B}^{-1}(\vec{\phi})\mathcal{A}|\text{in}\rangle$  ‘back’ to state  $|\text{in}\rangle$  can be made more robust by seeking a series of intermediate mappings. In effect, we ‘lure’  $\vec{\phi}$  towards the eventual goal incrementally. To achieve this, we generalise each gate  $G_i^A$  in our original circuit  $\mathcal{A}$  by introducing a global parameter  $\alpha$  such that  $G_i^A(\alpha = 0)$  corresponds to the identity, and  $\alpha = 1$  produces the original gate in its full effect.

As a single global variable common to all gates within  $\mathcal{A}$ ,  $\alpha$  then interpolates between the identity circuit  $\mathcal{A}_{\alpha=0} = \mathbb{1}$ , and the full circuit  $\mathcal{A}_{\alpha=1}$ . We

step  $\alpha$  in small increments, at each step running the imaginary-time protocol to find a parameter set  $\vec{\phi}$  that produces a sufficiently low energy state  $\mathcal{B}^{-1}(\vec{\phi})\mathcal{A}_{\alpha}|\text{in}\rangle$  according to our constructed Hamiltonian  $H_{\text{rec}}$ . We increment  $\alpha$  by  $\delta\alpha$  whenever we come within hyper-parameter  $\delta E$  of the ground state energy  $E_0$ . The final step sees  $\alpha = 1$  and the parameters  $\vec{\phi}$  converge to their final values, which should ‘undo’ the action of  $\mathcal{A}$  on  $|\text{in}\rangle$ . Other variant algorithms where  $\alpha$  develops by a continuous change should have a similar performance.

An example implementation of this lure method is provided in Appendix A. Despite its simplicity, luring can markedly improve the convergence of the minimisation subroutine of recompilation. The method was *not* necessary for the simulations performed in Sec. 4, but *was* in an anomalous simulation performed in Sec. 6. We note that the lure technique has some similarity to the incrementally evolving Hamiltonian employed in the adiabatic VQA demonstration of Ref. [37], however there the gradual adaption is to the cost function rather than the input state as here.

### 3.3 Adaptive timestep

In variational minimisation techniques like imaginary time evolution, the timestep  $\Delta t$  is often a fixed hyper-parameter. Indeed in the preceding Section 4,  $\Delta t$  will be chosen and fixed using a simple stability heuristic (as in [28]). However, an adaptive timestep algorithm can scale this hyper-parameter during minimisation, to significantly speed-up convergence in drastically changing landscapes. This becomes important in increasingly large Hilbert spaces, where motivating a fixed parameter becomes difficult, like the during the scaling tests in Section 6. We here present a custom adaptive timestep algorithm based on line search [38] which sees the timestep automatically change through orders of magnitude during optimisation.

Assume  $\Delta t$  was the timestep used in a particular iteration of parameter optimisation (e.g. imaginary time evolution). The next iteration involves computing the energy gradient  $\nabla \langle E(\vec{\theta}) \rangle$ , solving a linear equation for  $\frac{d\vec{\theta}}{dt}$ , and updating the parameters according to  $\vec{\theta} \rightarrow \vec{\theta} + \Delta t \frac{d\vec{\theta}}{dt}$ . In general, experimentally measuring the energy gradient  $\nabla \langle E(\vec{\theta}) \rangle$  is a task that scales linearly with the number of parameters, unlike measuring the energy at a point,  $\langle E(\vec{\theta}) \rangle$ . Hence, after computing  $\nabla \langle E(\vec{\theta}) \rangle$ , we can afford to make several additional energy measurements at little relative total cost. It is worthwhile to use the current gradient to produce several choices of  $\Delta t$ , measure their associated energies, and pick the best. This is advantageous because the energy landscape in parameter space can itself vary drastically, so a fixed  $\Delta t$  can become unsuitable. This is the basis of our algorithm, which

we present formally in Algorithm 1 .

---

**Algorithm 1:** Adaptively updating  $\Delta t$  each iteration of parameter minimisation.

---

**Require:**  $\Delta t$  is the previous iteration’s timestep  
**Require:**  $E_0$  is the previous iteration’s energy  
**Require:**  $\vec{\theta}$  are the current parameters and  $\frac{d\vec{\theta}}{dt}$  their pre-computed derivatives

```
// energy of nominated timestep under current
// gradient
define EN( $\Delta t$ )
  return  $\langle E(\vec{\theta} + \Delta t \frac{d\vec{\theta}}{dt}) \rangle$ 

// returns adapted timestep
define TIMESTEP( $\Delta t, \epsilon_{\Delta t/2}, \epsilon_{\Delta t}, \epsilon_{2\Delta t}$ )
   $\epsilon_{\min} = \min(\epsilon_{\Delta t/2}, \epsilon_{\Delta t}, \epsilon_{2\Delta t})$ 

  // halt when converged
  if  $|\epsilon_{\min}| \leq 10^{-8}$ 
    return  $\Delta t$ 

  // shrink rapidly when  $\Delta t$  is unstably large
  if  $\epsilon_{\min} > E_0$ 
    return TIMESTEP( $\Delta t/8, EN(\Delta t/16), EN(\Delta t/8), EN(\Delta t/4)$ )

  // adjust  $\Delta t$  until optimal, re-using
  // previous energies
  if  $\epsilon_{\min} = \epsilon_{\Delta t}$ 
    return  $\Delta t$ 
  if  $\epsilon_{\min} = \epsilon_{\Delta t/2}$ 
    return TIMESTEP( $\Delta t/2, EN(\Delta t/4), \epsilon_{\Delta t/2}, \epsilon_{\Delta t}$ )
  if  $\epsilon_{\min} = \epsilon_{2\Delta t}$ 
    return TIMESTEP( $2\Delta t, \epsilon_{\Delta t}, \epsilon_{2\Delta t}, EN(4\Delta t)$ )

// adapt timestep, starting from previous
 $\Delta t = \text{TIMESTEP}(\Delta t, EN(\Delta t/2), EN(\Delta t), EN(2\Delta t))$ 
```

---

Via Algorithm 1,  $\Delta t$  can be enlarged exponentially quickly if the current  $\vec{\theta}$  lies on a persistent slope, and shrunk exponentially quickly if  $\Delta t$  has become unstably large. It includes the circumstance whereby no choice of  $\Delta t/2$ ,  $\Delta t$  or  $2\Delta t$  decrease the energy, implying that  $\Delta t$  is unstably large, since it is known that a sufficiently small timestep of variational imaginary time evolution should exhibit a monotonic decrease in energy [28].

### 3.4 Further gate elimination

After recompiling  $\mathcal{A} |in\rangle$  into  $\mathcal{B}(\vec{\phi}) |in\rangle$  using the above methods, we can attempt to further shrink the circuit

by eliminating gates with small parameters (thus deviating from the user-specified template). We choose the parameter  $\phi_j$  whose current value  $\phi_j = \delta$  is closest to 0 (or more strictly, since we may be dealing with periodic functions, we identify  $j$  for which  $g_j^B(\phi_j)$  is closest to the identity). We then continue our imaginary-time evolution under modified variational equations, where we additionally constrain that within each iteration,

$$\Delta\phi_j = -\frac{\delta}{N}. \quad (9)$$

Here  $N$  limits the change in  $\phi_j$  in a single iteration. This simultaneously drives  $\phi_j$  toward zero while retaining the pressure toward the ground state. Generally we will reach  $\phi_j = 0$  having suffered a small penalty in energy (and thus fidelity of the new circuit). Once zero is reached, we remove gate  $g_j^B = \mathbb{1}$  and then repeat the process. This continues until the energy has unacceptably risen. What is ‘unacceptable’ will depend on the application, but for the examples here we set the threshold to be twice the gap between true ground and the original energy of  $\mathcal{B}^{-1}\mathcal{A}|in\rangle$  under  $\hat{H}_{\text{rec}}$ . In other words, we permit the energy defect to double in return for reducing the number of gates in the output circuit. We found this threshold enabled excellent circuit compression despite only a modest reduction in fidelity. We emphasise that this entire phase is an optional post-process after the main recompilation. We will denote the resulting circuit of this additional gate elimination process as  $\mathcal{B}_{\text{elim}}$ .

## 4 Numerical demonstration

Having thus described the compilation and optimisation process in general terms, we now illustrate it with a specific example: recompilation of a 7-qubit, 186-gate circuit into a quite different template.

### 4.1 Selecting an interesting example circuit $\mathcal{A}$

This section describes how we select an interesting and complex  $\mathcal{A}$  circuit as the object that we will attempt to recompile into a new form. In essence: we choose a 7-qubit circuit relevant to a certain simulation task and specify it in Fig. 1; we also choose initial state  $|in\rangle = |1\rangle|+\rangle^{\otimes 6}$ . *Readers who are concerned only with the recompilation process may care to skip the rest of this section.*

Rather than randomly generating the circuit  $\mathcal{A}$ , we focus on the likely application areas for our recompilation technique: hybrid algorithms that aim at dynamical simulation or eigensolving. Given that the recompilation technique itself involves a kind of eigensolver, for clarity we opt instead to make the circuit  $\mathcal{A}$  relevant to a dynamical simulation task. Specifically, we assume that we wish to model the evolution of a

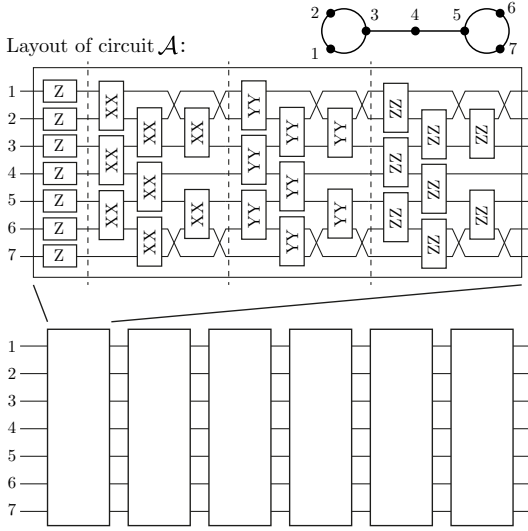


Figure 1: The circuit  $\mathcal{A}$  which we opt to use as the input to our compilation process. The upper right figure summarises the two-qubit gate connections. The circuit is related to a quantum dynamics problem as described in Section 4.1, but for recompilation purposes one can regard it as an arbitrary pattern of 186 unique non-Clifford gates (including 144 two-qubit gates). Here  $Z(\theta) = \exp(-i\frac{\theta}{2}\sigma^Z)$ ,  $ZZ(\theta) = \exp(-i\frac{\theta}{2}\sigma^Z \otimes \sigma^Z)$ , and similarly for the  $Y$  and  $X$  gates. The angle  $\theta$  is unique for each gate, and for completeness we specify the values in a table in the appendix.

certain 7–spin network, with the topology of spin-spin interactions shown in the upper right of Fig. 1. We take it that the Hamiltonian of this system is

$$H_{\text{sys}} = \sum_i B_i \sigma_i^z + \sum_{i,j} \sum_{S \in \{x,y,z\}} J_{i,j}^S \sigma_i^S \sigma_j^S \quad (10)$$

where  $\sigma$  are the standard Pauli operators and the constants  $B_i < 0$  and  $J_{i,j}^S > 0$  as listed in Appendix D. This is therefore a rather general spin network with irregular antiferromagnetic interactions and local fields. In order to create an interesting evolution we select the initial state  $\Psi(0) = |1\rangle |+\rangle^{\otimes 6}$  i.e. a product state where one qubit is orientated such that it has maximum energy with respect to its local field and the others have zero expected energy in their local fields. We choose a simple recompilation Hamiltonian for which  $\Psi(0)$  is the ground state, namely

$$\hat{H}_{\text{rec}} = \sigma_1^z - \sum_{j=2}^7 \sigma_j^x.$$

As a relevant test of our recompilation technique, we stipulate that the purpose of original circuit  $\mathcal{A}$  is to model the evolution of this system, i.e. to create (a good approximation to) the state

$$|\Psi(t)\rangle = \exp(iH_{\text{sys}}t) |\Psi(0)\rangle$$

for some time  $t$  which we presently specify. A naive approach might be to use a number of Trotter ‘cycles’ [39, 40] i.e. to use a number  $q$  of identical circuit

blocks each of which contains one gate for each Pauli term in  $H_{\text{sys}}$ . Each gate  $\mathcal{T}_j$  within the first Trotter cycle would therefore be of the form

$$\mathcal{T}_j = \exp\left(-i\frac{\theta_j}{2}K_j\right) = \left(\cos\frac{\theta_j}{2}\right)\mathcal{I} - i\left(\sin\frac{\theta_j}{2}\right)K_j$$

where  $j$  runs  $1 \dots 31$  in our case. Here  $K_j$  is the Pauli operator from the  $j^{\text{th}}$  term of  $H_{\text{sys}}$ , i.e. either a single or double  $\sigma$  operator. Meanwhile  $\theta_j/2 = C_j t/q$  where  $C_j$  is the constant in the  $j^{\text{th}}$  term of  $H_{\text{sys}}$ , i.e. a  $B$  or  $J$  value. Thus for  $\mathcal{T}_1$  we would use  $K_1 = \sigma_z$  and  $\theta_1 = 2B_1 t/q$ . Gates deeper into the circuit each replicate a gate in the first cycle, i.e.  $\mathcal{T}_{j+31} = \mathcal{T}_j$ . For sufficiently small  $t$  this approach is guaranteed to provide an accurate simulation. Fig. 1 shows a circuit of this kind.

Following this basic Trotter rule for selecting the  $\theta_j$  values leads to rather too simple a structure to test our recompilation protocol fully because of the gate recurrence  $\mathcal{T}_{j+31} = \mathcal{T}_j$ . (As an aside we remark that we have verified this: Setting the circuit’s  $\theta$  values to correspond to  $t = 0.75$  where its simulation fidelity is 0.9983, we find we can recompile to a circuit with only about half the number of two-qubit gates and yet still retain simulation fidelity above 0.998). The ease of recompiling this standard Trotter circuit is related to the fact that it does not make optimal use of the gates available. Given the same gate layout, as shown in Fig. 1, one can instead use the variational algorithm described in Ref. [31] to adjust the ‘strength’  $\theta_j$  of each gate independently of all others in an optimal fashion. We indeed apply this algorithm, which we introduced earlier as Li’s algorithm and we describe more completely in Appendix B, to create our circuit  $\mathcal{A}$ . We choose the time  $t = 1.75$  as this is toward the outer limit of the range for which the circuit structure in Fig. 1 can produce an accurate simulation using Li’s algorithm. The resulting parameters  $\theta_j$  are specified in Table 2 in Appendix D; configured this way, the circuit in Fig. 1 successfully replicates the state of the simulated spin system at time  $t = 1.75$  with a fidelity of 0.995.

Although the purpose here was simply to create a complex but meaningful circuit for recompilation, in doing so we did make a number of interesting observations about the power of Li’s algorithm as compared to Trotter approaches. For the interested reader these are described in Appendix B.

## 4.2 Performance of the recompilation

To make the recompilation task interesting, we specify a template for  $\mathcal{B}$  which differs from  $\mathcal{A}$  in several ways: Firstly, we select a different set of gates. The original circuit  $\mathcal{A}$  involves 6 types of gate (single qubit rotations about the  $X$ ,  $Y$  and  $Z$  axes, and two qubit rotation gates around  $XX$ ,  $YY$  or  $ZZ$ ). For the new template  $\mathcal{B}$  we opt to use the same set of single-qubit

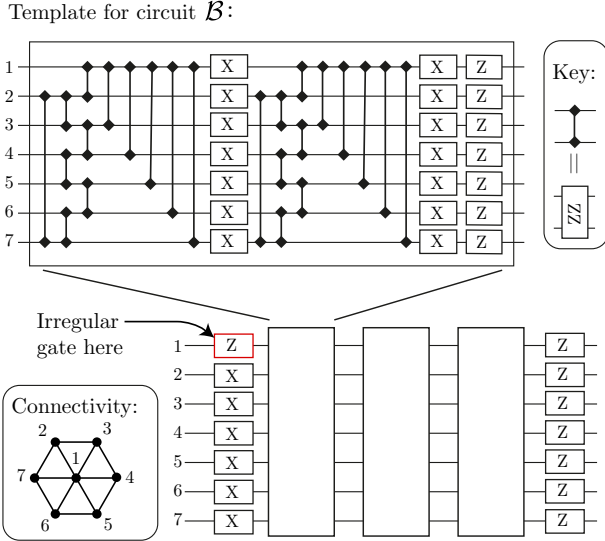


Figure 2: The template for the recompiled circuit  $\mathcal{B}$ . The template is user-specified, and the recompilation process will determine the  $\phi$  value for each gate. The template's structure is quite different to the original circuit: The template has a smaller family of gate types ( $YY$  and  $XX$  type gates are omitted), it has half as many two-qubit gates in total (72 rather than 144), but a larger number of single-qubit gates (77 versus 42). Moreover the topology of the two-qubit gates is different: it is a triangular lattice forming a hexagon as shown in the inset.




Figure 3: The process of recompiling circuit  $\mathcal{A}$  into the template  $\mathcal{B}$ . We use an imaginary time variational algorithm [28] as described in Appendices B and C, according to which the  $\phi_i$  parameters evolve through an iterative process until no further reduction in  $\langle H_{\text{rec}} \rangle$  occurs. The light blue curve shows the fidelity with which  $\mathcal{B}|\text{in}\rangle$  matches the target  $\mathcal{A}|\text{in}\rangle$ . Since the fidelity is not directly experimentally measurable, we also plot the user measurable quantity  $\langle H_{\text{rec}} \rangle$  in red, which bounds the fidelity as stated in Eqn. (8). (Also shown in dark purple is the fidelity with respect to the true state of the simulated spin system, see Section 4.1. The circuit  $\mathcal{A}$  itself has a finite simulation fidelity of 0.995 and thus the recompiled circuit is somewhat lower.)

gates but we restrict ourselves to only the  $ZZ$ -based two-qubit gate. (A restriction of this kind is relevant to real devices which typically have a native type of two-qubit gate that is the least onerous to perform, and therefore it is natural to attempt to recompile into a template featuring only one kind of two-qubit gate). We also vary the connectivity: instead of mimicking the Trotter terms of the Hamiltonian  $H_{\text{sys}}$ , as in circuit  $\mathcal{A}$ , we now adopt a centred hexagon as shown in Fig. 2.

The recompilation process then proceeds as described earlier. The task of the classical computer, i.e. the solution of the simultaneous equations which yields the appropriate parameter updates at each step, is performed via truncated singular value decomposition as described in Appendix D. Figure 3 shows how the parameters develop. The upper part of this figure indicates the performance of the recompilation process; reaching the target energy of  $\langle H_{\text{rec}} \rangle = -7$  would indicate perfect recompilation. Also shown is the fidelity of  $\mathcal{B}|\text{in}\rangle$  versus  $\mathcal{A}|\text{in}\rangle$ , which the user would not have access to. The final fidelity of the process is 0.998. Given the restricted nature of the template, and particularly the fact that it has only half as many two-qubit gates, the recompile is remarkably effective.

We remark that, in other examples that we have studied we have used a template with a higher gate count and recompiled to (essentially) perfect fidelity: infidelity of order  $10^{-5}$  has been observed, and it seems probable that this is non-zero only because of imperfections in the numerical solution methods. We anticipate that these methods can be further refined.

As a post-recompile stage, we apply the gate elimination process described in Section 3.4. It is necessary to select a tolerance for the process, i.e. a level of reduction in the overall quality of the recompiled circuit which we are prepared to tolerate in return for further ‘compressing’ the circuit. For the present example, we assume that we will tolerate a doubling in the energy defect with respect to the ideal value of  $-7$ . The performance of this process is indicated in Fig. 4. Remarkably we find that we can eliminate a further 30 gates (11 single-qubit, 19 two-qubit) in this fashion, and that the fidelity of the resulting circuit is still high at 0.995. The final circuit includes only 119 gates (whereas the original  $\mathcal{A}$  has 186) and moreover the number of two-qubit gates has been reduced almost to a third (from 144 to 53). Recall that all remaining two-qubit gates are all of a single type,  $ZZ = \exp(-i\frac{\theta}{2}\sigma^Z \otimes \sigma^Z)$ .

We emphasise that these recompilation fidelities have been achieved with a non-trivial circuit  $\mathcal{A}$  that is *already* optimised, with respect to a simple Trotter circuit, through the use of Li’s algorithm. As noted earlier, applying our recompilation to a simple Trotter circuit will produce more dramatic results, e.g. compression to  $\sim 50\%$  depth with only a very small loss






Figure 4: The gate elimination post-process, following the recompilation stage. We require certain of the  $\phi_j$  parameters to diminish to zero during additional iterations of imaginary time evolution, so that the corresponding gate can be removed. This continues until the energy  $\langle H_{\text{rec}} \rangle$ , shown in red in the top panel, deviates from the known minimum (here,  $-7$ ) by some chosen threshold (here, double the original defect). The lower panel shows how the corresponding parameters are eliminated; in total thirty are set to zero, reducing the circuit size from 149 to 119. Vertical lines indicate each time a parameter is eliminated.




Figure 5: Realtime simulation of the spin system specified in Eqn. (10) from  $t = 0$  to  $t = 2.5$  using Li’s algorithm [31] with various possible circuits. In blue, the original structure  $\mathcal{A}$ , which we previously ‘froze’ at  $t = 1.75$ . We note if we allow its parameters  $\theta$  to continue to evolve using Li’s algorithm, past  $t = 1.75$  the fidelity of the simulation drops precipitously. In red, the recompiled circuit  $\mathcal{B}$  is seen to fare somewhat worse than when its parameters  $\phi$  are similarly evolved by Li’s algorithm. However, the green line corresponds to the performance of the recompiled and augmented circuit (as described in the main text) and is superior: it can sustain high fidelity simulation for a further 0.4 time units. We highlight the resource cost of the featured circuits in Fig. 6.

of fidelity ( $\Delta F = 0.0002$  in that case).

This concludes our description of our recompilation example. However, given the success of the process, we were interested to consider the following questions which relate to the nature of circuit  $\mathcal{A}$  as a simulation task. As explained earlier in Section 4.1, the meaning of  $\mathcal{A}$  is that it reproduces, with fidelity 0.995, the state of a certain physical spin system (Eqn. (10)) at time  $t = 1.75$ .

(1) If we take the template form of  $\mathcal{B}$  with all  $\phi_j = 0$  and use it with Li’s algorithm ‘in the first place’ without ever considering the Trotter-inspired layout of  $\mathcal{A}$ , will we obtain a high performance simulation out to  $t = 1.75$ ? If this were the case it would rather obviate the need for recompilation, at least for this type of application. However the answer is ‘no’. As shown in Fig. 16 by the red-dashed line, the template  $\mathcal{B}$  performs very poorly as a basis for the simulation. Its fidelity drops to  $\sim 0.98$  almost immediately. This is because a small increment in the parameters in  $\mathcal{B}$  does not correspond to a small shift in time for the simulated system (circuit  $\mathcal{B}$  has the ‘wrong’ topology, as shown by the inset in Fig. 2 versus the inset in Fig. 1). To create our compact representation of the circuit that simulates the state of the spin system at  $t = 1.75$ , we did indeed need to use a circuit whose structure reflects the Hamiltonian (Eqn. (10)) to reach  $t = 1.75$  and only then recompile it.

(2) What happens if we now augment the (recompiled, compressed) circuit by ‘pasting on’ a number of ‘blank’ Trotter cycles so as to recover roughly the

two-qubit gate count of the original circuit  $\mathcal{A}$ , and then proceed with the simulation which motivated that circuit? Substituting the new recompiled and augmented circuit, will the simulation of the spin system’s dynamics using Li’s algorithm proceed forward from  $t = 1.75$  with good fidelity? The results are shown in in Fig. 5. We see that this does indeed produce a superior performance versus simply continuing past the  $t = 1.75$  point with a circuit of the form of  $\mathcal{A}$ . Fig. 6 visualises the resource costs of the tested simulation strategies, colour coded to match their performances in Fig. 5. It highlights that the significant improvement of circuit  $\mathcal{T}^3\mathcal{B}$  over  $\mathcal{A}$  in the future phase of real-time simulation was achieved without increasing the number of two qubit gates.




Figure 6: Resource count of the considered circuits in Fig. 5, and those involved in recompilation. Recall that  $\mathcal{A}$  is the input circuit of 6 Trotter cycles though with parameters freely evolving under Li’s method.  $\mathcal{A}$  is recompiled into  $\mathcal{B}$  via the use of ansatz  $\mathcal{B}^{-1}\mathcal{A}$  in an imaginary-time extension of Li’s method. At a further fidelity penalty, one can eliminate the weakest gates in  $\mathcal{B}$  to obtain  $\mathcal{B}_{\text{elim}}$ . After recompiling, real-time simulation can be continued by appending additional Trotter cycles  $\mathcal{T}$  to  $\mathcal{B}$  or  $\mathcal{B}_{\text{elim}}$  (as indicated in Fig. 5).

## 5 Robustness to noise

The methods described in this paper are relevant to two scenarios: compiling a quantum circuit using a classical machine, as in the simulations here, or doing so with the use of an actual quantum computer to perform the optimisation. The latter is necessary for large circuits given that classical machines cannot efficiently emulate general circuits of more than about 50 qubits. In scenarios where a quantum computer is available for the recompilation, it might be a fault-tolerant machine capable of practically perfect execution or it might be a NISQ-era machine with inherent noise in gates, measurements, etc. Indeed the latter is a likely scenario given that NISQ-era computers will have a critical need of optimised circuits. However the possibility of compilation with a noise-burdened system raises the important question of whether the output of the compilation, i.e. the classical param-




Figure 7: Recompilation using a quantum computer that suffers significant noise. The task is that explored in the previous section, namely recompilation of  $\mathcal{A}|\text{in}\rangle$  (Fig. 1) into  $\mathcal{B}(\vec{\phi})|\text{in}\rangle$  (Fig. 2), though using the cost function given in Eq. 12. The horizontal axis  $\xi$  is the severity of the two-qubit depolarising noise following each two-qubit gate (single qubit gates receive severity  $\xi/10$ ). The green lines show the infidelity between the target state  $\mathcal{A}|\text{in}\rangle$  and the pure state generated by  $\mathcal{B}$  with parameters  $\phi_{\xi}^n$ , as learned from  $n$  iterations of noisy quantum natural gradient. Specifically, this is  $1 - \langle \text{in} | \mathcal{A}^\dagger \mathcal{B}(\phi_{\xi}^n) | \text{in} \rangle$ , and is the error introduced by the noisy compilation process itself. The red line (uppermost) is included for context, and shows the infidelity between the ideal target state  $\mathcal{A}|\text{in}\rangle$  and its noisy counterpart  $\mathcal{D}_{\mathcal{A}}(|\text{in}\rangle\langle\text{in}|)$ , where  $\mathcal{D}_{\mathcal{A}}$  is the channel prescribed by inserting noise into  $\mathcal{A}$ . Concretely, the red line shows  $1 - \langle \text{in} | \mathcal{A}^\dagger \mathcal{D}_{\mathcal{A}}(|\text{in}\rangle\langle\text{in}|) \mathcal{A} | \text{in} \rangle$ . One observes that the blue line is orders of magnitude lower than the red, confirming that recompilation is remarkably noise-robust. Generally 9 independent runs were collated for each plotted value of  $\xi$ , except  $\xi = 0.01$  for which there were 14 runs.

eters specifying the gates in the recompiled circuit, would be severely degraded by the noise in the compilation process. This is a question we address through a series of further simulations.

An encouraging study [41] by Sharma *et al* notes that VQE-type optimisation of a cost function can be remarkably robust against noise in the system. Furthermore, the methods mentioned earlier in subsection 2.3 provide a mechanism to formally incorporate noise (or indeed other non-unitary processes) into the imaginary descent mechanism which we have adopted throughout this paper. Specifically, in Ref. [29] the metric required for imaginary time (or ‘quantum natural gradient’) methods is shown to relate to quantum Fisher Information in the general case, and an experimentally-practical method to estimate that quantity is described. It is therefore possible, if computationally demanding, to investigate the generalisation of the cases considered earlier in this paper to scenarios involving circuit noise.

We revisit the previous problem presented of recompilation from  $\mathcal{A}|\text{in}\rangle$  (shown in Fig. 1) into the format of circuit  $\mathcal{B}|\text{in}\rangle$  (shown in Fig. 2). As before, we seek the parameter set  $\vec{\phi}$  that allows  $\mathcal{B}$  to most

closely match the action of  $\mathcal{A}$  on a given input state (see Eqn. 7). The performance given error-free gates was described in subsection 4.2 and Fig. 3. We now introduce errors afflicting *all* quantum gates, including those in both circuits (since both are required in the recompilation task). Our goal is still to find the optimal parameter set for the noise-free case, therefore we can expect that the noise in both circuits is deleterious for this task.

We therefore modify our previous simulation by introducing per-gate noise, relevant error mitigation, and a modified cost function, as described below:

1. We apply depolarising noise subsequent to all initialisation and gate operations. We apply either single-qubit or two-qubit correlated noise depending on the gate that has just occurred. Two-qubit depolarising is of severity  $\xi$  and acts on qubits  $i, j$  with the mapping

$$\rho \rightarrow (1 - \xi)\rho + \frac{\xi}{15} \sum_k N^k \rho N^{k\dagger}, \quad (11)$$

where  $N^i$  are the fifteen operators  $\mathcal{I}_i \sigma_j^X, \mathcal{I}_i \sigma_j^Y, \dots$  though to  $\sigma_i^Z \sigma_j^Z$ . Meanwhile the single-qubit depolarising channel has the analogous form with just three operators  $N$  and lower severity  $\xi/10$ . This reflects the fact that in experiments single qubit gates typically have higher fidelity.

2. In order to generalise the imaginary time (a.k.a. natural gradient) method to these non-unitary, noise-burdened gates, the metric tensor key to that approach must be generalised [29]. In the lab, the evaluation of the metric requires a swap-test involving two copies of the circuit output together with an ancilla, implying  $2 \times 7 + 1 = 15$  qubits for the present case. Note that in our present density matrix simulations, both circuit outputs are represented identically going into the swap-test step, and can therefore employ a single copy and apply a generous degree of noise in modelling the swap-test process itself. Further details are in Appendix D.
3. We presume that an experimental effort would employ error mitigation to negate noise as far as possible, and so we incorporate such a method in our model. We adopt the extrapolation technique which is a well-established [31, 42] and experimentally validated [43] method with a simple protocol: Each observable required for the optimisation process is estimated by measuring at noise level  $\xi$ , and again at noise level  $2\xi$  (in our simulations, this is done without the addition of further noise due to finite sampling, i.e. shot noise, so that values will correspond to the high-sampling limit). The recorded value is then the observable extrapolated to  $\xi \rightarrow 0$  using the presumption of exponential decay [44].

4. We select a fictitious Hamiltonian with a cost function which we find efficient for this more computationally-demanding scenario:

$$H_{\text{rec}} = \mathcal{I} - |\Psi(0)\rangle\langle\Psi(0)|. \quad (12)$$

As before,  $|\Psi(0)\rangle = |1\rangle|+\rangle^{\otimes 6}$ . This cost function is straightforward to evaluate experimentally by measuring each qubit in the relevant basis. For problems beyond the 7-qubit case considered here, we would likely adopt a localised version of this global function, in line with recent work arguing for superior scaling via local functions [45].

We performed simulations for various values of the noise severity  $\xi$  ranging between 0.1% and 1%, presented in Fig. 7. We emphasise  $\xi$  is the noise for each two-qubit gate (with additional, weaker, noise on single-qubit gates). One might expect the compilation process to fail completely for  $\xi$  approaching 1% given that the total recompilation circuit has about 250 gates (as it involves both circuits  $\mathcal{A}$  and  $\mathcal{B}$ , shown in Figs. 1 and 2 respectively). We note that even the original circuit  $\mathcal{A}$  acting alone will produce a very noisy output with this level of gate imperfection; in Fig. 7 the solid red line shows the infidelity in the output of  $\mathcal{A}$  versus the noise-free case, and we observe severe degradation (order 10% infidelity) even for noise level  $\xi = 0.1\%$ , rising above 70% infidelity at the  $\xi = 1\%$  level.

Despite these high levels of noise, recompilation in fact succeeds remarkably well. Recall that the purpose of the recompilation is to determine the best set of classical parameters  $\vec{\phi}$  with which to configure variational circuit  $\mathcal{B}$  so when it acts on input  $|1\rangle|+\rangle^{\otimes 6}$ , its output most closely matches that of circuit  $\mathcal{A}$  given the same input. Therefore to measure the effectiveness of the recompile we plot the infidelity

$$G = 1 - F\left(\mathcal{B}(\vec{\phi}_\xi)|\text{in}\rangle, \mathcal{A}|\text{in}\rangle\right), \quad (13)$$

where circuits are now noiseless but circuit  $\mathcal{B}$  is configured with the classical parameters  $\vec{\phi}_\xi$  obtained from our (noisy, with severity  $\xi$ ) recompilation process.

The solid blue line in Fig. 7 shows this infidelity  $G$ . We observe that even for severe noise at the 1% per gate level, the recompile still produces good classical parameters  $\vec{\phi}_\xi$ . The use of  $\vec{\phi}_{\xi=0.01}$  instead of the ideal  $\vec{\phi}_{\xi=0}$  causes  $G$  to rise from 0.3% to about 1.1%, whereas the noise is now so severe that the output of any variant, even the original circuit  $\mathcal{A}$ , is now highly mixed. One can reasonably state that *in circumstances where the noise is not too high for circuits to be of some value*, the recompilation remains useful and indeed remarkably accurate. In this sense, we conclude recompilation with noisy circuits is entirely practical, and the use of the advanced imaginary time optimisation method does not impede this.

We should note two caveats to this encouraging conclusion: Firstly, that only the depolarising model has

been studied so far, and secondly that there are increasing sample costs for increasing noise severity. For the former, while this is something that could be investigated comprehensively in a followup study, we feel confident that the very high degree of robustness suggests that other models will be at least moderately well-tolerated. We note that VQA approaches are inherently blind to systematic, deterministic errors (albeit one would need to correct the obtained-parameters for relevant shifts) whereas dephasing is an element of the depolarising process already modelled, and even damping errors are first-order correctable by the extrapolation mitigation that we incorporate.

For the latter point regarding sampling costs, we emphasise that our numerical studies as reported in Fig. 7 correspond to the high-sampling limit, i.e. the performance is that which is achieved when sampling is cheap enough that shot-noise is never the dominant imperfection. Thus the present work does not incorporate optimisations for frugal sampling; indeed exploration of VQAs using this constraint is an evolving theme in its own right (see e.g. Refs. [46, 47]). We can introduce an *ad hoc* shot-noise factor to our model, but because our core gradient estimation does not employ shot-frugal techniques any resulting estimate of absolute sample count required would have little value as an indicator of the performance of a mature, full-optimised application. However, we can obtain some broad indications of the *scaling* of sample cost as we introduce gate noise. To do so, we indeed wrote a function to synthesise shot-noise variability and set a high sample count of  $10^8$  for the low-noise (0.1%) limit of the circuit explored in Fig. 7. At such a high sampling rate, there is negligible impact of shot noise versus infinite sampling. We then explore the increase in sampling count needed to maintain near-optimal performance for gate noise scenarios in the range 0.1% to 1%, where by ‘near-optimal’ we mean that the final cost function value should change only modestly (well within a factor of two). We found that a further factor of 3 in the sample count sufficed for per-gate noise of 0.4%, while the factor increased to 15 for the high-noise limit of 1%.

This simple exercise supports the intuition that it may ultimately be profligate sampling requirements, rather than any inherent limitations of the scheme, that will bound the use of noisy circuits in practical compilation scenarios. (As a footnote we remark that for all these simulations the sampling count for the quantities in the metric tensor was set ten times higher, as this object is crucial to correct parameter evolution but, in an optimised setting, is relatively cheap to estimate as it does not need to be updated frequently [33]).




Figure 8: An example of circuit  $\mathcal{A}$  (Eqn. 14) which prepares a 4-qubit squeezed state [49] when applied to  $|0\rangle$ , as used as input to test our recompilation scheme’s scaling performance.

## 6 Numerical scaling

We now explore the scaling of our recompilation technique as the number of qubits increases. For computational tractability we take the gate-noise-free case and perform only the base recompilation strategy via imaginary time evolution (without additional gate elimination nor luring). However we retain the use of our adaptive timestep scheme as described earlier.

As input to the recompilation, we choose a scalable family of circuits which generate squeezed states [48, 49], chiefly

$$\mathcal{A}(\phi_1, \phi_2, \phi_3) = \prod_t Y_t(\pi/2) \prod_{c,t>c} C_c Z_t(\phi_1) \times \prod_{t,c>t} C_c Z_t(\phi_1) \prod_t X_t(\phi_2) Z_t(\phi_3), \quad (14)$$

where  $U_t$  notates gate  $U$  acting on qubit  $t$ , and  $C_c U_t$  notates the same gate with additional control qubit  $c$ . As before,  $X, Y, Z$  denote rotations. For each test of a given circuit size, the parameters  $\phi_1, \phi_2, \phi_3$  are chosen uniformly randomly in  $[0, 2\pi)$ . An example circuit which prepares a 4-qubit squeezed state is shown in Figure 8, and the qubit connectivity illustrated in Figure 9.

With a family of input circuits chosen, we then choose a family of target template circuits with a distinct gate topology and gate set from the input circuit. We nominate the template

$$\mathcal{B}(\vec{\theta}) = V \text{reverse}(V), \quad (15)$$

$$\text{where } V = \left( \prod_t Z_t X_t Y_t Z_t \prod_t U_{t,t+1} \right)^2, \quad (16)$$

$$\text{and } U = \begin{pmatrix} 1 & & & \\ & e^{i\theta} & & \\ & & e^{i\theta} & \\ & & & 1 \end{pmatrix}, \quad (17)$$

where every gate features a unique parameter  $\theta$ , and  $U$  is the parameterised SWAP gate. An example 4-qubit  $\mathcal{B}(\vec{\theta})$  circuit is shown in Figure 10, and the qubit connectivity in Figure 9.

Note that  $\mathcal{A}(\phi_1, \phi_2, \phi_3)$  and  $\mathcal{B}(\vec{\theta})$  differ significantly, and there is no exact assignment  $\vec{\theta}$  for which  $\mathcal{A}(\phi_1, \phi_2, \phi_3) |0\rangle \equiv \mathcal{B}(\vec{\theta}) |0\rangle$  besides when both are the identity ( $\phi_i = \theta_j = 0$ ). Furthermore, the *optimum*




Figure 9: Connectivity of the qubits in the input circuit  $\mathcal{A}$  (producing squeezed states), and the target ansatz  $\mathcal{B}$ , shown here for 6 qubits.




Figure 10: A 4-qubit example of  $V$  (the first half of  $\mathcal{B}(\vec{\theta})$  in Equ 17), used a template to recompile into in our scaling tests. The second half of  $\mathcal{B}$  is merely the reverse of this circuit, with independent parameters.

value of  $\vec{\theta}$  for which  $|\langle 0 | \mathcal{B}^\dagger(\vec{\theta}) \mathcal{A}(\phi_1, \phi_2, \phi_3) | 0 \rangle|$  is maximum, may admit a progressively less-accurate recompilation as the number of qubits increases. Hence, qualifying the performance of the scheme in a scale-invariant way is difficult.

Recompilation will be driven by the simple  $n$ -qubit Hamiltonian

$$H_{\text{rec}} = \mathbb{1} - |0\rangle\langle 0|^{\otimes n}. \quad (18)$$

We opt to use this since its spectral gap remains 1 as the system size  $n$  grows; the first excited state is  $(2^n - 1)$ -degenerate, and the fidelity of the recompiled state is exactly  $1 - \langle E \rangle$ , where  $\langle E \rangle$  is the expected energy under  $H_{\text{rec}}$ . Furthermore, for the sake of numerical simulation, it offers quick classical evaluation ( $\mathcal{O}(1)$  and  $\mathcal{O}(2^n)$  respectively) of the energy and its partial derivative in parameter space,

$$\langle E \rangle = \langle \psi | H_{\text{rec}} | \psi \rangle = 1 - |\psi_0|^2, \quad (19)$$

$$\begin{aligned} \nabla \langle E \rangle_j &= \text{Re} \left\{ \langle \psi | H_{\text{rec}} \left| \frac{\partial \psi}{\partial \theta_j} \right\rangle \right\} \\ &= \text{Re} \left\{ \langle \psi | \frac{\partial \psi}{\partial \theta_j} \rangle - \psi_0^* \left( \frac{\partial \psi}{\partial \theta_j} \right)_0 \right\}. \end{aligned} \quad (20)$$

For a general Hamiltonian, these quantities could otherwise take time  $\mathcal{O}(2^{2n})$  to evaluate. Here  $\psi_0$  indicates the first complex element of the statevector  $|\psi\rangle$ . These quantities are calculated for every assignment of  $\vec{\theta}$  (i.e. every iteration).

We will pick a number of iterations sufficiently large such that all tests have converged (in our tests, 200 iterations), and measure the final fidelity achieved. Recompilation will be driven by imaginary time evolution, without luring, though *with* the adaptive time-step subroutine introduced in Section 3.3.

In the subsequent simulations, use of adaptive time-step over a heuristically-motivated fixed time-step im-




Figure 11: An example of how the time-step exponentially evolves under the adaptive scheme used in the scaling tests.




Figure 12: Performance of recompilation of squeezed states, driven by 200 iterations of imaginary time evolution, tested 20 times for each system size. All tests employ no luring except the red point; this corresponds to the use of 10 stages of luring achieving a fidelity increase from 59% to 94%.

proves convergence drastically. The cost is very modest: on average about 1% more energy measurements. A demonstration of the time-step evolving adaptively during one of the subsequent scaling tests is presented in Figure 11.

## 6.1 Scaling performance

Using this adaptive time-step scheme, we simulated recompilation of squeezed states of 3 to 20 qubits. For each state size, we perform 20 tests with random squeezed state parameters  $\vec{\phi}$ , and random initial ansatz parameters  $\vec{\theta}$ . At each of the 200 iterations of variational imaginary-time evolution, the linear equations were solved using singular value decomposition, with no truncation. The results are presented in Figure 12. Despite excellent recompilation of the 16-qubit state to 99% fidelity, larger squeezed states performed significantly worse. The maximum fidelity achieved for 19 qubits (highlighted along the red line in Figure 12) is 59%, though the average is 11%. We speculate this particular squeezed state may be anomalously difficult to recompile due to, for example, prescribing strong entanglement between non-neighbouring qubits, for which there may be a




Figure 13: Recompilation of the anomalously difficult 19-qubit squeezed state in Fig 12, using 10 stages of luring. The ultimate fidelity of 94% is denoted in Figure 12 as a red dot.

small space of proximate ansatz states. Converging into such a space from an initial random set of parameters becomes decreasingly likely with increasing number of qubits [50].

As cautioned however, a test of scaling performance may be more difficult than recompiling a particular state. We demonstrate this by recompiling the anomalously difficult 19 qubit circuit, using the same total number of iterations (200), but introducing 10 stages of luring, and parameter truncation. The ansatz parameters now begin at zero, and all parameters in  $\mathcal{A}|0\rangle$  (including the  $Y(\pi/2)$  terms) are initially scaled to  $\times 0.1$ , and then incrementally restored to their full values over the course of recompilation. We furthermore introduce truncation of singular values below  $10^{-5}$  when solving the linear equations involved in variational imaginary-time evolution. At no extra algorithmic cost, the same 19 qubit case is now recompiled to a final 94% fidelity, and is denoted by the red point in Figure 12. The changing fidelity achieved by recompilation against the lured state is shown in Figure 13.

## 7 Possible applications

In Section 4, we demonstrated one possible application area: Using recompilation periodically during some task (modelling dynamics, or eigensolving, etc) so as to ‘compress’ the current circuit and thus ‘make room’ for additional gates. This may be relevant when considerations such as noise accumulation make deeper circuits undesirable. However as a caveat we should stress that in the present demonstration it is actually the recompilation process itself that requires the deepest circuit, since it involves the concatenation  $\mathcal{B}^{-1}\mathcal{A}$ . A full demonstration of this possibility would require one to show that repeated recompilation can reduce the maximum circuit depth, and this is a topic for further work.

Generally when considering applications it is important to remember that the present technique needs quantum hardware in order to perform the recompilation (unless we are doing so for small circuits rele-

vant to today’s small prototypes). Moreover, the recompile process presented here involves a variational eigensolver and therefore will require that circuit  $\mathcal{A}$  be executed a large number of times in order to complete the recompile. One might ask, would the user not have done better simply to use the (non-optimal) circuit  $\mathcal{A}$  in their application, rather than bothering to recompile it? The answer depends on how many state preparations  $\mathcal{A}|in\rangle$  are required for that application. We now mention two important cases where one might expect that this number is large relative to the recompilation cost, thus making it prudent to indeed recompile and perform subsequent state preparations with  $\mathcal{B}|in\rangle$ .

One such application is that of preparing Gibbs states [51, 52]; states of this kind are vital in understanding the equilibrium properties of physical systems and are also relevant to application areas such as machine learning via Boltzmann machines [53]. A full understanding of recompilation in the context of Metropolis algorithms would likely require extending the present techniques to include projective measurement, but this does not appear to present an in-principle difficulty. Once one has a recompiled circuit for generation of the Gibbs state, downstream applications involving sampling from that state (typically, a very great number of times) will be correspondingly accelerated. A important context where repeated sampling of a Gibbs state is essential, is the emerging field of quantum semi-definite programming, see e.g. Refs. [54, 55].

Variational algorithms for the estimation of molecular energies are currently the subject of much interest [22]. Typically some circuit  $\mathcal{A}$  of depth  $d$  generates an  $n$ -qubit state that is a good approximation to the ground state of the (translated, qubit-based) chemical Hamiltonian  $H_{\text{chem}}$  [56]. The estimation of  $\langle H_{\text{chem}} \rangle$  with respect to this state is costly since  $H_{\text{chem}}$  can consist of up to  $O(n^4)$  terms [36] and the result must be known to at least 3 decimal places to achieve chemical accuracy. Thus the process may take time of order  $T \propto dn^4/\epsilon^2$  where  $\epsilon$  the tolerable shot noise of order  $10^{-3}$  or less. In principle it is then efficient to recompile the circuit prior to the energy estimate, in order to reduce  $d$  to  $d'$  the depth of the recompiled circuit, provided that (a) the recompile process is rapid compared to  $T$  and (b) the circuit recompilation introduces error small compared to  $\epsilon$ . These criteria can potentially be met since the Hamiltonian within the recompile process itself is very simple, consisting of  $n$  terms.

A different type of application would be to use the techniques described to recompile a circuit, initially expressed in some standard language (CNOTs, single-qubit rotations, etc) into the real gates occurring in some specific device. Such gates will be the device’s direct response to control signals (laser or microwave pulses, electrode potential shifts, etc) and may have

non-trivial effects not only on targeted qubits but also collateral effects on others, i.e. crosstalk may be significant. To the extent that these effects are deterministic and can be characterised, they can certainly be the building blocks with which we construct our template  $\mathcal{B}$ .

This raises an interesting further prospect of what one might call ‘blind compilation’ on a quantum computer: Provided that the state  $\mathcal{A}|\text{in}\rangle$  can be prepared, then the recompilation into template  $\mathcal{B}$  can proceed by varying the parameters  $\phi_j$  *without* understanding the effect of those parameters. We note that the imaginary time eigensolver [28] which we have applied in the present paper may require modification to be used in this way, but more simple approaches such as gradient descent would be immediately applicable. However the caveat that  $\mathcal{A}|\text{in}\rangle$  must be prepared is significant and would seem to imply that some mode of operation of the device does allow the realisation of standard gates (unless  $\mathcal{A}|\text{in}\rangle$  has been prepared by another system).

Note that all these potential applications must carefully consider, and may be ultimately limited by, the total samplings costs. Since a practical implementation must make use of strategies to budget and optimise sampling, the final sampling costs cannot be quickly estimated here. We do at least know that the cost of the metric tensors involved in imaginary time evolution and noisy quantum natural gradient are marginal [29].

## 8 Conclusion

We have described a method for recompiling a quantum circuit  $\mathcal{A}$  into a new, user-specified ‘template’ with the goal that the resulting circuit  $\mathcal{B}$  has the same effect on a *specific* input state  $|\text{in}\rangle$ , i.e. we aim to achieve  $\mathcal{A}|\text{in}\rangle = \mathcal{B}|\text{in}\rangle$ . This is a quite different (and more permissive) requirement than seeking a new circuit that represents the same overall unitary. However our criterion is exactly the relevant one for many quantum algorithms currently under study. An important class is that of variational hybrid algorithms, whether intended for eigensolving, or modelling the dynamics of quantum systems, or materials science applications.

The methods we describe include the recompilation process itself, which is achieved through an isomorphism to a ground-state finding protocol (i.e. an eigensolver) and is therefore a hybrid algorithm itself. We identify an approach which can improve the performance of the recompilation by using a ‘lure’, i.e. providing a proximal goal and then moving the goal further off once it has been approached (this is discussed fully in Appendix A). Furthermore we have described a post-processing gate elimination routine, which can be applied following the recompilation in order to remove ‘weak’ gates at relatively little cost

in the fidelity of the circuit.

We applied these ideas to a specific, complicated 7-qubit circuit  $\mathcal{A}$ , which we created by considering a simulation task involving 7 spins. The creation of the example circuit was itself an interesting task, and we make a number of remarks about it in the appendices. The circuit involved 186 unique non-Clifford gates and was therefore quite substantial. Our recompilation process successfully realised an equivalent smaller circuit with half the number of two-qubit gates and a smaller number of gate types; the new circuit replicated the action of the original with a 0.998 fidelity.

Applying the optional post-compile gate elimination, we removed a further thirty gates and reduced the two-qubit gate count to only about a third of the number in original circuit  $\mathcal{A}$ . This came at the cost of a small further reduction in fidelity, to 0.995.

Given that circuit recompilation is likely to be of value in the NISQ era, it is very relevant to ask whether the methods described here can be successful even when the quantum machine employed for the optimisation is imperfect. To explore this question we repeated our initial study but with the introduction of noise throughout the circuits involved. We confirmed a remarkable degree of robustness in the compilation process, concluding that *for any level of noise for which the original circuit can be of practical use*, the imperfect recompilation process would introduce very little further degradation. This observation was made with caveats regarding noise model and sample costs, as explained in the main text. We note that a related but distinct question, is whether compilation (or ‘learning’ of tasks) can discover a more noise-resilient way to approximately perform a task under a given noise model; there have been encouraging results reported on this theme quite recently [57].

Having thus explored the 7-qubit case in some detail, we then studied the scaling of the recompilation protocol over a range of tasks from 3 to 20 qubits. Such a study inevitably faces a number of difficulties in obtaining a clear conclusion. In particular, because we wish to see the scaling of a challenging task (recompilation of a meaningful circuit into a ‘template’ with a less rich gate topology) it follows that the task itself may become harder with increasing number of qubits – then, even an ideal recompilation system would return circuits of diminishing fidelity. We observed that the protocol successfully achieved recompilation at all scales tested, although there was an appreciable drop-off in fidelity for the largest cases. An interesting observation was that the case with worst performance, a certain randomly-generated task involving 19 qubits, could be dramatically improved by altering the protocol to enable the ‘lure’ technique. A reasonable inference is that, as the size of the recompilation task increases, naïve implementations of the variational technique must be replaced with more nuanced protocols.

Ultimately the recompilation method we de-




Figure 14: A failed direct recompilation of the 5-qubit circuit  $\mathcal{C}(\vec{\theta})|0\rangle$  (Eqn. 21) into  $\mathcal{C}(\vec{\phi})|0\rangle$ . The top panel shows the 15 parameters converge to a non-ground state, distinct from their optimal (original) parameters on the right. The bottom plot shows the energy of  $\mathcal{C}(\vec{\phi})^{-1}\mathcal{C}(\vec{\theta})|0\rangle$  and fidelity  $F(\mathcal{C}(\vec{\phi})|0\rangle, \mathcal{C}(\vec{\theta})|0\rangle)$  plateau without reaching their optima.

scribe here, being dependent on variational energy-minimisation, will be subject to essentially the same scaling as all variational methods. The corresponding issues of ‘barren plateaus’ [50] and local versus global cost functions [45], etc, are under very active investigation in the field [27, 29, 30].

We concluded the main part of the paper by indicating some potential application areas for the technique.

**Acknowledgements** We thank Xiao Yuan, Sam

## A Demonstration of the *Lure* method

To illustrate the *Lure* method, we here consider a 5-qubit 15-gate circuit

$$\mathcal{C}(\vec{\theta}) = \prod_{q=0}^4 X_q(\theta_{5q}) Z_q(\theta_{5q+1}) C_q Y_{(q+1)\%5}(\theta_{5q+2}), \quad (21)$$

where  $Z_q(\theta_j) = \exp(-i\theta_j\sigma^z/2)$  indicates a rotation of qubit  $q$  (indexed from 0) around the  $z$ -axis of the Bloch sphere by angle  $\theta_j$  (and similarly for  $Y$  and  $Z$ ),  $C_q$  indicates the preceding gate is controlled on qubit  $q$ , and  $\%$  indicates modulus. We randomly assign  $\vec{\theta}$  and test recompilation of  $\mathcal{C}(\vec{\theta})|0\rangle$  into the same circuit,  $\mathcal{C}(\vec{\phi})|0\rangle$ , which can in theory be done with perfect fidelity via  $\vec{\phi} = \vec{\theta}$ . We use time-step  $\Delta\tau = 0.1$  and employ a simple Hamiltonian  $\hat{H}_{\text{rec}} = \mathbb{1} - |0\rangle\langle 0|$  which under ideal imaginary time evolution, drives  $\mathcal{C}(\vec{\phi})^{-1}\mathcal{C}(\vec{\theta})|0\rangle \rightarrow |0\rangle$ . Despite this, attempts to directly recompile  $\mathcal{C}(\vec{\theta})$  can fail as shown in Fig. 14, where the parameters become trapped in a non-ground state. To combat this, Fig. 15 demonstrates *luring* by successively recompiling an intermediate




Figure 15: A successful recompilation of the 5-qubit ansatz  $\mathcal{C}(\vec{\theta})|0\rangle$  (Eqn. 21) into  $\mathcal{C}(\vec{\phi})|0\rangle$  using *luring*: successively recompiling  $\mathcal{C}(\vec{\theta}_{n/10})|0\rangle$  where  $n$  is incremented each time the measurable energy falls within 0.1 of ground (indicated by the dashed lines). The top panel shows the evolving parameters  $\vec{\phi}$  recover their original counterparts. The bottom panel shows the fidelity and energy reach their optima.

McArdle and Earl Campbell for very helpful comments. This work was supported by EPSRC grant EP/M013243/1. SCB acknowledges support from the EPSRC QCS Hub EP/T001062/1, from U.S. Army Research Office Grant No. W911NF-16-1-0070 (LOGIQ), and from EU H2020-FETFLAG-03-2018 under the grant agreement No 820495 (AQ-TION). The authors would like to acknowledge the use of the University of Oxford Advanced Research Computing (ARC) facility in carrying out this work (dx.doi.org/10.5281/zenodo.22558).






Figure 16: Comparative performance of the Li and Trotter methods for realtime simulation of the spin system specified in Eqn. (10), without any compilation. The top panel shows the fidelity with true evolution. The Li method freely evolves the parameters in an ansatz with the structure of 6 Trotter cycles. Meanwhile, the Trotter method uses fixed parameters and every second cycle is reversed except in the simulation indicated by the orange dashed line; this is discussed further in the main text. The bottom panel shows how the 186 parameters in Li’s method evolve. The subplot compares 16 of the 32 parameters (summed over Trotter cycles) to their fixed Trotter counterparts in the early stages of the simulation. The vertical dashed line in both panels indicates the time  $t = 1.75$  time units at which simulation is stopped and the Li circuits recompiled according to the main text.

state  $\mathcal{C}(\vec{\theta}n/10)$  for  $n = 1, 2, \dots, 10$ , updated whenever the energy falls within 0.1 of the known 0 energy ground state. With luring, the evolved parameters  $\vec{\phi}$  are seen to restore their  $\theta$  counterparts.

## B Li’s Algorithm and Trotter comparison

In the main paper we generated our input circuit  $\mathcal{A}$  by considering a simulation task. Rather than using a direct Trotter method we instead used a gate layout suitable for Trotter but selected the gate parameters using a more efficient method. That method is the one described in Ref. [31] by Li and Benjamin, and we refer to it as Li’s Algorithm.

Two variants are described in Ref. [31], each resulting from a different initial variational equation. In the present paper we use the method that results from McLachlan’s variational principle, which in practice means selecting  $\eta = -i$  in Eqns. (5) and (6) of that paper, which we reproduce below:

$$\sum_q M_{k,q} \dot{\lambda}_q = V_k, \quad (22)$$

where  $\lambda_q$  are the parameters in the evolving circuit, equivalent to the  $\theta_i$  or  $\phi_i$  in our treatments. Meanwhile

$$M_{k,q} = \frac{\partial \langle \Psi |}{\partial \lambda_k} \frac{\partial | \Psi \rangle}{\partial \lambda_q} + \text{H.C.}, \quad (23)$$

$$V_k = -i \frac{\partial \langle \Psi |}{\partial \lambda_k} H | \Psi \rangle + \text{H.C.} \quad (24)$$

In Ref. [31] it is shown that these matrix and vector elements are of the form  $a \text{Re}(e^{i\theta} \langle \bar{0} | U | \bar{0} \rangle)$  where  $U$  is a unitary involving the  $k^{\text{th}}$  and  $q^{\text{th}}$  gates of the parameterised circuit, and  $|\bar{0}\rangle$  is the all-zero state.

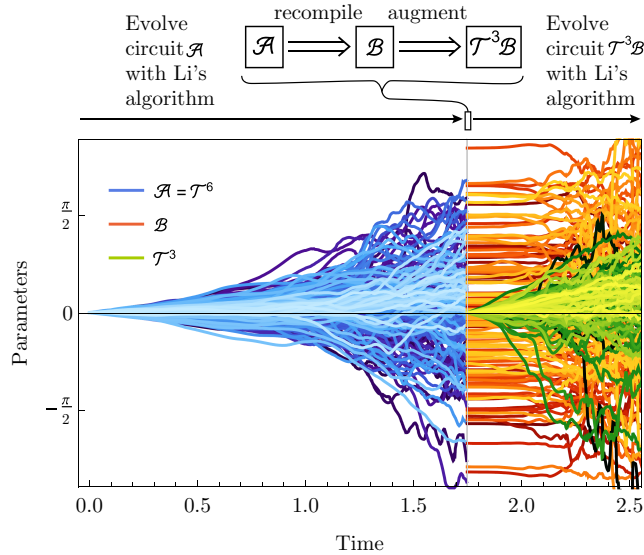


Figure 17: A demonstration of continuing realtime simulation by recompiling ansatz  $\mathcal{A}(\vec{\theta})$  into  $\mathcal{B}(\vec{\phi})$  and appending 3 Trotter cycles with freely evolving parameters  $\vec{\omega}$ . The blue lines indicate  $\vec{\theta}$  as in Fig. 16 until  $t = 1.75$ . Thereafter, the red lines indicate  $\vec{\phi}$  and the green lines,  $\vec{\omega}$ , which are initially zero.

It is further shown that  $\text{Re}(e^{i\theta} \langle 0|U|\bar{0}\rangle) = \text{Tr}(X_{\text{anc}}\rho)$ , i.e. it relates directly to expected value of an ancilla that probes the parameterised circuit in a simple fashion (see Fig. 2 of Ref. [31], which is a variant of a circuit proposed in Ref. [58] in 2002).

Therefore a quantum computer is able to provide the  $M$  and  $V$  values involved in equation (22), which can then be solved classically to yield the  $\lambda_i$  values and thus to update the parameters for the next incremental state of the circuit.

In the present work, a small modification was applied to the algorithm specified in Ref. [31]. We now ensure there is a parameter in the problem which controls the global phase. This can be significant since the McLachlan principle is sensitive to global phase, even though it has no physical meaning. Generally one need not introduce an actual global phase gate, instead such a gate can be virtually present and allowed for in the update equations. Fortunately in the specific case we describe in the main paper, the input state is such that a gate performing a  $z$ -rotation gate acts on a qubit in state  $|1\rangle$ , thus this gate indeed generates a global phase.

In creating our trial circuit  $\mathcal{A}$ , we made some interesting observations concerning the way in which the circuit derived using Li's algorithm compares to a standard Trotter circuit. A comparison of the Li and Trotter methods is shown in Fig. 16, which includes a visualisation of the Li parameters in the bottom panel. By 'standard Trotter' we mean that the choice of parameters (angles  $\theta$  within the gates) is determined simply from the elapsed simulation time and the strength of the corresponding Hamiltonian term.

Figure 16 includes three lines (orange-dashed, orange and green) which correspond to Trotter-type solutions of this kind. The orange-dashed line is a direct realisation of the circuit pattern shown in Fig. 1. Its performance is poor since the simple repetition of identical cycles causes Trotter errors to directly accumulate. The orange line represents the best realisation of the Trotter-type approach that we found: it involves alternately inverting the order of the gates (versus that shown in the middle panel of Fig. 1). This is related to approaches in Ref. [59], but without the randomisation element; in our simulations random re-ordering was not advantageous, but this may relate to the small system size and circuit depth. We acknowledge that there are many approaches including the sampling method of Campbell which has just been presented [60], and we have not comprehensively searched over them. In short, the orange line is a fair indicator of a 'smart' Trotter approach (as opposed to the naive one represented by the orange-dashed line). The green line is the same approach of alternating inversions, but now with a total of 16 cycles.

The figure indicates that Li's algorithm does better than a (smart, if not necessarily optimal) Trotter approach for a given circuit depth. This is to be expected since Li's approach has greater freedom: it does not seek to approximate the time evolution unitary but rather it seeks the circuit with the closest-to-ideal action on the defined input state.

As an aside, we remark that since the ansatz has the structure of 6 Trotter cycles, it is natural to wonder whether Li's algorithm will, at early times, choose the same set of parameters which Trotterisation prescribes to the rotation gates in the time-evolution unitary for those times. Indeed, around half of Li's parameters (when

summed according to their Trotter term) precisely follow their Trotter counterparts, as shown in the subfigure of Fig. 16. Interestingly, the other parameters which differ do improve the fidelity achieved by Li’s algorithm versus Trotterisation by an appreciable factor even in the early evolution (the fidelity ‘gap’ is  $4 \times 10^{-4}$  after 100 iterations of  $\Delta t = 2.5 \times 10^{-3}$ ).

This behaviour is even seen in the parameters of the additional Trotter cycles appended to the recompiled ansatz, as shown in Fig. 17.

## C The imaginary time eigensolver

The recompilation process described in the main paper employs a ground state finding algorithm at its core. We seek the ground state of a fictitious Hamiltonian  $H_{\text{rec}}$  because by doing so we necessarily find a circuit that can invert the original circuit  $\mathcal{A}$ , and thus we obtain an alternative realisation of  $\mathcal{A}$ .

There are a variety of possible algorithms that can configure a parameterised circuit so that it maps a fixed input state onto (an approximation to) the ground state of a specified Hamiltonian. In the present paper we employed a recent imaginary time method, as described in Ref. [28]. In essence this algorithm is a variant of Li’s algorithm described in Appendix B above; as proved in Ref. [28] we need only modify Eqn. (24) with a factor  $i$  in order to switch from a real time evolution (i.e. an evolution according to the Schrödinger equation) to an evolution under  $\exp(-Ht)$  (with renormalisation) which drives the system to its ground state – or rather, to the lowest eigenstate with finite projection on the initial state.

## D Details of numerical simulation

We simulate the quantum circuits involved in the introduced technique using the Quantum Exact Simulation Toolkit [61], and QuESTlink [62]. All numerics were performed using GPU parallelisation, on a desktop machine with a 24 GiB Quadro P6000 GPU.

Several computational shortcuts are taken in the evaluation of the equations involved in the Li method, which we stipulate in the Supplementary Materials of previous works [28, 32]. Tables 1 and 2 list some numerical values used in the simulated Hamiltonian and recompilation process. The full simulation code for the demonstration in Sec. 4, which depends only on the GNU Scientific Library [63] and standard/included C libraries, is available on GitHub [1]. This includes extensive documentation about employing the recompiler as a simulation tool. The QuESTlink and Mathematica code for the scaling tests in Sec. 6 are also available on GitHub [2].

The Li method, in both its original form [31] and imaginary-time adaptation [28], involves populating a family of linear equations which are then numerically solved by a classical machine. We make some observations about the choice of the linear solving algorithm, which was previously seen to affect the ground state convergence rate of imaginary-time simulations [32]. In this work, we discover that Tikhonov regularisation and least squares minimisation [64] work well to ensure the parameters vary smoothly in realtime Li simulation. However, imaginary time Li simulation benefited from truncated singular value decomposition (TSVD) [64] which imposes no constraint on the parameter smoothness and saw much faster convergence to the ground state than the previously mentioned methods.

We now outline the approach taken for tractable simulation of the fully noise-burdened circuits in Section 5. A lab-based realisation of the method used here, specifically the experimental determination of the crucial metric tensor employed in the imaginary time (or ‘natural gradient’) method, would require two copies of the main register in order to perform swap-test-like process as explained in Ref. [29]. However, in the present case of numerical simulation, we can benefit from the fact that the density matrix corresponding to each register is identical (presuming the same noise model through the device). We can avoid the very considerable increase in simulation costs that would arise from explicitly representing the density matrix of  $1 + 2n$  qubits, instead representing  $1 + n$ . This is exact until the actual swap-test process; the noise afflicting this stage (which is far more shallow than the preceding ansatz-based circuits) can be modelled by estimating how severely the observables comprising the metric tensor would be impacted, and then indeed applying this level of corruption directly to the tensor. In our study we ‘err on the side of caution’ by making these noise events quite severe, in order to be certain that the imperfection of the swap-test phase does not undermine the robustness of the overall compilation process. Specifically: we first implement a proxy for the inter-register swap-test via a control-permute of all qubits in the (sole) register, such that every qubit position is changed (formally, a derangement). We then characterise the degree of corruption in a random sample of 30 single- and two-qubit Pauli observables (themselves computed using noisy gates). The noise-free metric tensor is then corrupted by corresponding variances.

$B_1$	-0.433333	$J_{1,2}^x$	0.730767	$J_{4,5}^y$	0.412000
$B_2$	-1.006667	$J_{2,3}^x$	0.745333	$J_{5,6}^y$	0.641200
$B_3$	-0.941460	$J_{1,3}^x$	0.830400	$J_{6,7}^y$	0.697453
$B_4$	-0.312000	$J_{3,4}^x$	0.521333	$J_{7,5}^y$	0.345333
$B_5$	-0.478667	$J_{4,5}^x$	0.543800	$J_{1,2}^z$	0.362667
$B_6$	-0.347867	$J_{5,6}^x$	0.338771	$J_{2,3}^z$	0.864433
$B_7$	-0.314533	$J_{6,7}^x$	0.700544	$J_{1,3}^z$	0.527000
		$J_{7,5}^x$	0.771333	$J_{3,4}^z$	0.543000
		$J_{1,2}^y$	0.367333	$J_{4,5}^z$	0.547667
		$J_{2,3}^y$	0.913333	$J_{5,6}^z$	1.014333
		$J_{1,3}^y$	0.570667	$J_{6,7}^z$	0.368100
		$J_{3,4}^y$	0.315333	$J_{7,5}^z$	0.701000

Table 1: The coefficients of  $\hat{H}_{\text{sys}}$ , the 7-qubit spin network Hamiltonian. These are available in plaintext in the online repository [1] (direct link).

As a side note: We employed the following heuristic rule, entirely compatible with experimental realisation, which somewhat improved performance in the high-noise limit. Specifically, when taking a step in the optimisation walk, the protocol tries both the ‘new’ metric tensor and the ‘old’ tensor, where the latter was the last successfully computed tensor. If the old one gives rise to a superior step down in cost function, then this is employed instead of the new (which is deemed to be too corrupted to be useful). This approach exploits the fact that the metric tensor should be smoothly changing function of the ansatz parameters, therefore one can ‘fall back’ to an immediately prior tensor without significant impact on efficacy.

Finally, we now list the chosen numerical constants. For realtime and imaginary time simulation, we use step sizes of  $\Delta t = 2.5 \times 10^{-3}$  and  $\Delta \tau = 1 \times 10^{-2}$  respectively. We regularise the change in parameters per-iteration using truncated singular value decomposition (TSVD) and Tikhonov regularisation. We use a TSVD tolerance of  $10^{-5}$  and choose the Tikhonov parameter as the corner of a 3-point L-curve [64]. Wavefunction derivatives are estimated using a fourth-order finite difference method with a step-size of  $\Delta \theta = 10^{-5}$ . When we desire an initially identity unitary, we set all parameters  $\theta_j = 10^{-8}$  to avoid singularities in the matrix inversion during the first iteration. During the gate elimination subroutine, we restrict the to-be-zeroed parameter to vary by at most  $|\Delta \phi_j| \leq 0.1$  radians each iteration.

## References

- [1] [github.com/QTechTheory/DissipativeRecompiler](https://github.com/QTechTheory/DissipativeRecompiler)
- [2] [github.com/QTechTheory/RecompilerSqueezeScaling](https://github.com/QTechTheory/RecompilerSqueezeScaling)
- [3] John Preskill, *Quantum Computing in the NISQ era and beyond*, *Quantum* **2**, 79 (2018).
- [4] S. Bravyi and A. Kitaev, *Universal quantum computation with ideal Clifford gates and noisy ancillas*, *Phys. Rev. A* **71**, 022316 (2005).
- [5] Ying Li, *A magic state’s fidelity can be superior to the operations that created it*, *New J. Phys.* **17**, 023037 (2015).
- [6] E. T. Campbell and M. Howard, *Unified framework for magic state distillation and multiqubit gate synthesis with reduced resource cost*, *Phys. Rev. A* **95**, 022316 (2017).
- [7] David Gosset, Vadym Kliuchnikov, Michele Mosca, and Vincent Russo, *An algorithm for the T-count*, *Quantum Information & Computation* **14**, 1261 (2014).
- [8] Neil J. Ross and Peter Selinger, *Optimal ancilla-free Clifford+T approximation of z-rotations*, *Quantum Information & Computation* **16**, 0901 (2016).
- [9] Matthew Amy and Michele Mosca, *T-count optimization and Reed-Muller codes*, *IEEE Transactions on Information Theory* **65**, 8 (2019)
- [10] Luke E. Heyfron and Earl T. Campbell, *An Efficient Quantum Compiler that reduces T count*, *Quantum Sci. Technol.* **4**, 015004 (2019).
- [11] Yunseong Nam, Neil J. Ross, Yuan Su, Andrew M. Childs and Dmitri Maslov, *Automated optimization of large quantum circuits with continuous parameters*, *npj Quantum Information* **4**, 23 (2018).
- [12] R. Barends, J. Kelly, A. Megrant, A. Veitia, D. Sank, E. Jeffrey, T. C. White, J. Mutus, A. G. Fowler, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, C. Neill, P. O’Malley, P. Roushan, A. Vainsencher, J. Wenner, A. N. Korotkov, A. N. Cleland, and J. M. Martinis, *Superconducting quantum circuits at the surface code threshold for fault tolerance*, *Nature* **508**, 500 (2014).

n	$\hat{O}_n$	$\theta_n$															
1	$Z_0$	-0.00078	32	$Z_0$	0.00612	63	$Z_0$	-1.92631	94	$Z_0$	0.07784	125	$Z_0$	0.26960	156	$Z_0$	-0.80412
2	$Z_1$	-0.76524	33	$Z_1$	-0.72202	64	$Z_1$	-0.72235	95	$Z_1$	-0.28875	126	$Z_1$	-0.14207	157	$Z_1$	-1.95611
3	$Z_2$	-0.07102	34	$Z_2$	-0.89156	65	$Z_2$	-0.98678	96	$Z_2$	0.30971	127	$Z_2$	0.00518	158	$Z_2$	-1.35826
4	$Z_3$	-1.57210	35	$Z_3$	1.92348	66	$Z_3$	0.34984	97	$Z_3$	-1.02248	128	$Z_3$	0.10542	159	$Z_3$	-0.83901
5	$Z_4$	-0.57014	36	$Z_4$	-0.31031	67	$Z_4$	-0.42137	98	$Z_4$	0.08347	129	$Z_4$	-0.61129	160	$Z_4$	-0.21782
6	$Z_5$	-0.03702	37	$Z_5$	-0.54383	68	$Z_5$	0.20952	99	$Z_5$	-0.68914	130	$Z_5$	1.16104	161	$Z_5$	-0.85344
7	$Z_6$	-0.56917	38	$Z_6$	0.42303	69	$Z_6$	-0.64586	100	$Z_6$	-2.11670	131	$Z_6$	-0.49693	162	$Z_6$	-0.36461
8	$X_{0,1}$	-0.23393	39	$X_{0,1}$	-0.23944	70	$X_{0,1}$	-1.19517	101	$X_{0,1}$	0.93936	132	$X_{0,1}$	0.98615	163	$X_{0,1}$	1.56663
9	$X_{1,2}$	-0.77487	40	$X_{1,2}$	1.10291	71	$X_{1,2}$	0.19727	102	$X_{1,2}$	0.99647	133	$X_{1,2}$	0.31287	164	$X_{1,2}$	0.67684
10	$X_{0,2}$	0.87379	41	$X_{0,2}$	0.62165	72	$X_{0,2}$	0.43687	103	$X_{0,2}$	0.02575	134	$X_{0,2}$	1.18330	165	$X_{0,2}$	0.73790
11	$X_{2,3}$	-1.09072	42	$X_{2,3}$	0.22439	73	$X_{2,3}$	0.77102	104	$X_{2,3}$	0.90254	135	$X_{2,3}$	0.24543	166	$X_{2,3}$	0.29939
12	$X_{3,4}$	0.81605	43	$X_{3,4}$	1.58883	74	$X_{3,4}$	-0.08047	105	$X_{3,4}$	-0.33305	136	$X_{3,4}$	0.50658	167	$X_{3,4}$	0.10799
13	$X_{4,5}$	-0.10714	44	$X_{4,5}$	0.69977	75	$X_{4,5}$	-0.11872	106	$X_{4,5}$	0.04207	137	$X_{4,5}$	0.50517	168	$X_{4,5}$	0.31331
14	$X_{5,6}$	-1.61988	45	$X_{5,6}$	-0.43299	76	$X_{5,6}$	1.13964	107	$X_{5,6}$	1.74800	138	$X_{5,6}$	-1.30005	169	$X_{5,6}$	0.48987
15	$X_{6,4}$	-0.58697	46	$X_{6,4}$	0.38593	77	$X_{6,4}$	0.26760	108	$X_{6,4}$	0.19480	139	$X_{6,4}$	0.18067	170	$X_{6,4}$	0.38603
16	$Y_{0,1}$	1.03841	47	$Y_{0,1}$	0.28201	78	$Y_{0,1}$	0.03540	109	$Y_{0,1}$	0.64047	140	$Y_{0,1}$	-0.98891	171	$Y_{0,1}$	0.59770
17	$Y_{1,2}$	-0.20021	48	$Y_{1,2}$	0.77338	79	$Y_{1,2}$	0.40912	110	$Y_{1,2}$	0.75941	141	$Y_{1,2}$	0.78280	172	$Y_{1,2}$	0.89393
18	$Y_{0,2}$	1.25312	49	$Y_{0,2}$	0.17709	80	$Y_{0,2}$	0.25703	111	$Y_{0,2}$	0.04256	142	$Y_{0,2}$	0.62293	173	$Y_{0,2}$	0.22623
19	$Y_{2,3}$	0.19017	50	$Y_{2,3}$	0.37472	81	$Y_{2,3}$	-0.00260	112	$Y_{2,3}$	0.43709	143	$Y_{2,3}$	0.25405	174	$Y_{2,3}$	0.35892
20	$Y_{3,4}$	0.59210	51	$Y_{3,4}$	0.59450	82	$Y_{3,4}$	-0.06945	113	$Y_{3,4}$	0.89857	144	$Y_{3,4}$	0.03360	175	$Y_{3,4}$	0.25520
21	$Y_{4,5}$	-1.71059	52	$Y_{4,5}$	-0.22396	83	$Y_{4,5}$	0.25205	114	$Y_{4,5}$	0.28955	145	$Y_{4,5}$	0.42917	176	$Y_{4,5}$	0.35523
22	$Y_{5,6}$	0.05980	53	$Y_{5,6}$	-0.02429	84	$Y_{5,6}$	-0.44813	115	$Y_{5,6}$	2.37157	146	$Y_{5,6}$	-0.84443	177	$Y_{5,6}$	0.88797
23	$Y_{6,4}$	-0.12607	54	$Y_{6,4}$	-0.23800	85	$Y_{6,4}$	0.38207	116	$Y_{6,4}$	0.29358	147	$Y_{6,4}$	0.43788	178	$Y_{6,4}$	0.31737
24	$Z_{0,1}$	0.18838	55	$Z_{0,1}$	-1.31227	86	$Z_{0,1}$	1.60603	117	$Z_{0,1}$	0.02855	148	$Z_{0,1}$	-0.00031	179	$Z_{0,1}$	0.26033
25	$Z_{1,2}$	-0.20432	56	$Z_{1,2}$	0.17508	87	$Z_{1,2}$	0.06147	118	$Z_{1,2}$	0.74683	149	$Z_{1,2}$	0.24798	180	$Z_{1,2}$	0.99595
26	$Z_{0,2}$	0.49788	57	$Z_{0,2}$	0.26480	88	$Z_{0,2}$	-0.01839	119	$Z_{0,2}$	0.19942	150	$Z_{0,2}$	0.33195	181	$Z_{0,2}$	0.62794
27	$Z_{2,3}$	0.58855	58	$Z_{2,3}$	1.11408	89	$Z_{2,3}$	-0.95135	120	$Z_{2,3}$	0.37599	151	$Z_{2,3}$	0.81622	182	$Z_{2,3}$	-0.03170
28	$Z_{3,4}$	0.09731	59	$Z_{3,4}$	0.13049	90	$Z_{3,4}$	1.39875	121	$Z_{3,4}$	-0.13800	152	$Z_{3,4}$	0.41878	183	$Z_{3,4}$	-0.14826
29	$Z_{4,5}$	-1.04750	60	$Z_{4,5}$	0.31341	91	$Z_{4,5}$	0.10627	122	$Z_{4,5}$	0.26683	153	$Z_{4,5}$	0.71386	184	$Z_{4,5}$	0.11074
30	$Z_{5,6}$	0.34316	61	$Z_{5,6}$	1.52420	92	$Z_{5,6}$	-1.00678	123	$Z_{5,6}$	0.63415	154	$Z_{5,6}$	-0.83722	185	$Z_{5,6}$	0.25884
31	$Z_{6,4}$	-0.12666	62	$Z_{6,4}$	0.39354	93	$Z_{6,4}$	-0.07047	124	$Z_{6,4}$	0.35838	155	$Z_{6,4}$	0.70935	186	$Z_{6,4}$	0.01146

Table 2: The gates  $\hat{O}$  and parameters  $\theta$  which produce ansatz  $\mathcal{A}|1\rangle|+\rangle^{\otimes 6}$ , which is a result of Li simulation of the 7-qubit spin network for 700 iterations at  $\Delta t = 2.5 \times 10^{-3}$ . Note  $Z_i = \exp(-i\theta\sigma_i^z/2)$ ,  $Z_{i,j} = \exp(-i\theta\sigma_i^z \otimes \sigma_j^z/2)$  and similarly for  $Y$  and  $X$ . These values are available in plaintext in the online repository [1] (direct link).

- [13] T. P. Harty, D. T. C. Allcock, C. J. Ballance, L. Guidoni, H. A. Janacek, N. M. Linke, D. N. Stacey, and D. M. Lucas, *High-fidelity preparation, gates, memory and readout of a trapped-ion quantum bit*, Phys. Rev. Lett. **113**, 220501 (2014).
- [14] C. J. Ballance, T. P. Harty, N. M. Linke, M. A. Sepiol, and D. M. Lucas, *High-fidelity quantum logic gates using trapped-ion hyperfine qubits*, Phys. Rev. Lett. **117**, 060504 (2016).
- [15] J. P. Gaebler, T. R. Tan, Y. Lin, Y. Wan, R. Bowler, A. C. Keith, S. Glancy, K. Coakley, E. Knill, D. Leibfried, and D. J. Wineland, *High-fidelity universal gate set for  $^9\text{Be}^+$  ion qubits*, Phys. Rev. Lett. **117**, 060505 (2016).
- [16] Naomi H. Nickerson, Joseph F. Fitzsimons and Simon C. Benjamin, *Freely Scalable Quantum Technologies Using Cells of 5-to-50 Qubits with Very Lossy and Noisy Photonic Links*, Phys. Rev. X **4**, 041041 (2014).
- [17] Kentaro Heya, Yasunari Suzuki, Yasunobu Nakamura, and Keisuke Fujii, *Variational Quantum Gate Optimization*, arxiv:1810.12745 (2018).
- [18] Harper R. Grimsley, Sophia E. Economou, Edwin Barnes and Nicholas J. Mayhall, *An adaptive variational algorithm for exact molecular simulations on a quantum computer*, Nature Comms. **10** 3007 (2019).
- [19] Arthur G. Rattew, Shaohan Hu, Marco Pistoia, Richard Chen and Steve Wood, *A Domain-agnostic, Noise-resistant, Hardware-efficient Evolutionary Variational Quantum Eigensolver*, arXiv:1910.09694 (2019).
- [20] M. Bilkis, M. Cerezo, Guillaume Verdon, Patrick J. Coles, Lukasz Cincio, *A semi-agnostic ansatz with variable structure for quantum machine learning*, arXiv:2103.06712 (2021).
- [21] P. Wocjan, D. Janzing, and T. Beth, *Two QCMA-complete problems*, Quant. Inf. & Comp., **3**, 6 (2003).
- [22] For a recent review describing hybrid variational techniques in chemistry, see e.g. Sam McArdle, Suguru Endo, Alan Aspuru-Guzik, Simon Benjamin and Xiao Yuan, *Quantum computational chemistry*, Rev. Mod. Phys. (2019).
- [23] Sumeet Khatri, Ryan LaRose, Alexander Poremba1, Lukasz Cincio1, Andrew T. Sornborger, and Patrick J. Coles, *Quantum-assisted quantum compiling*, Quantum **3**, 140 (2019).
- [24] Jacques Carolan, Masoud Mohseni, Jonathan P. Olson, Mihika Prabhu, Changchen Chen, Darius Bunandar, Murphy Yuezhen Niu, Nicholas C. Harris, Franco N. C. Wong, Michael Hochberg, Seth Lloyd and Dirk Englund, *Variational quantum unsampling on a quantum photonic processor*, Nat. Phys. (2020).
- [25] M. Cerezo, Andrew Arrasmith, Ryan Babbush, Simon C. Benjamin, Suguru Endo, Keisuke Fujii, Jarrod

- R. McClean, Kosuke Mitarai, Xiao Yuan, Lukasz Cincio, Patrick J. Coles, *Variational quantum algorithms*, Nat. Rev. Phys. 3(9), pp.625-644 (2021).
- [26] Suguru Endo, Zhenyu Cai, Simon C. Benjamin, Xiao Yuan, *Hybrid quantum-classical algorithms and quantum error mitigation*, J. Phys. Soc. Jpn., **90**, 032001 (2021).
- [27] Xiao Yuan, Suguru Endo, Qi Zhao, Ying Li, and Simon C. Benjamin, *Theory of variational quantum simulation*, Quantum **3**, 191 (2019).
- [28] Sam McArdle, Tyson Jones, Suguru Endo, Ying Li, Simon Benjamin and Xiao Yuan, *Variational ansatz-based quantum simulation of imaginary time evolution*, npj Quantum Information **5**, 75 (2019).
- [29] Bálint Koczor, Simon C. Benjamin, *Quantum natural gradient generalised to non-unitary circuits*, arXiv:1912.08660v2 (2020).
- [30] James Stokes, Josh Izaac, Nathan Killoran, Giuseppe Carleo, *Quantum Natural Gradient*, Quantum **4**, 269 (2020).
- [31] Y. Li and S. C. Benjamin, *Efficient variational quantum simulator incorporating active error minimization*, Phys. Rev. X **7**, 021050 (2017).
- [32] Tyson Jones & Suguru Endo, Sam McArdle, Xiao Yuan and Simon Benjamin, *Variational quantum algorithms for discovering Hamiltonian spectra*, Phys. Rev. A **99**, 062304 (2019).
- [33] Barnaby van Straaten, Bálint Koczor, *Measurement cost of metric-aware variational quantum algorithms*, PRX Quantum **2**, 030324 (2021).
- [34] Lukasz Cincio, Yigit Subasi, Andrew T Sornborger and Patrick J Coles, *Learning the quantum algorithm for state overlap*, New J. Phys. **20** 113022 (2018).
- [35] Sang Min Lee, Jinhyoung Lee and Jeongho Bang, *Learning unknown pure quantum states*, Phys. Rev. A **98**, 052302 (2018).
- [36] Jarrod R McClean, Jonathan Romero, Ryan Babbush and Alán Aspuru-Guzik, *The theory of variational hybrid quantum-classical algorithms*, New J. Phys. **18**, 023023 (2016).
- [37] Ming-Cheng Chen, Ming Gong, Xiaosi Xu, Xiao Yuan, Jian-Wen Wang, Can Wang, Chong Ying, Jin Lin, Yu Xu, Yulin Wu, Shiyu Wang, Hui Deng, Futian Liang, Cheng-Zhi Peng, Simon C. Benjamin, Xiaobo Zhu, Chao-Yang Lu, and Jian-Wei Pan, *Demonstration of Adiabatic Variational Quantum Computing with a Superconducting Quantum Coprocessor*, Phys. Rev. Lett., **125**, 180501 (2020).
- [38] Grippo, Luigi, Francesco Lampariello, and Stephano Lucidi, *A nonmonotone line search technique for Newton's method*. SIAM Journal on Numerical Analysis **23**, 4 (1986).
- [39] Hale F. Trotter, *On the product of semi-groups of operators*, Proceedings of the American Mathematical Society **10**, 4 (1959).
- [40] Seth Lloyd, *Universal quantum simulators*, Science 1073-1078 (1996).
- [41] Kunal Sharma, Sumeet Khatri, M Cerezo, Patrick J Coles, *Noise resilience of variational quantum compiling*, New J. Phys. **22** 043006 (2020).
- [42] Kristan Temme, Sergey Bravyi and Jay M. Gambetta, *Error Mitigation for Short-Depth Quantum Circuits*, Phys. Rev. Lett. **119**, 180509 (2017).
- [43] Abhinav Kandala, Kristan Temme, Antonio D. Córcoles, Antonio Mezzacapo, Jerry M. Chow & Jay M. Gambetta *Error mitigation extends the computational reach of a noisy quantum processor*, Nature **567**, 491 (2019).
- [44] Suguru Endo, Simon Benjamin and Ying Li, *Practical Quantum Error Mitigation for Near-Future Applications*, Phys. Rev. X, **8**,031027 (2018).
- [45] M. Cerezo, Akira Sone, Tyler Volkoff, Lukasz Cincio and Patrick J. Coles, *Cost function dependent barren plateaus in shallow parametrized quantum circuits*, Nature Comms. **12**, 1791 (2021).
- [46] Jonas M. KÅEbler, Andrew Arrasmith, Lukasz Cincio and Patrick J. Coles, *An Adaptive Optimizer for Measurement-Frugal Variational Algorithms*, Quantum **4**, 263 (2020).
- [47] Andrew Arrasmith, Lukasz Cincio, Rolando D. Somma and Patrick J. Coles, *Operator Sampling for Shot-frugal Optimization in Variational Algorithms*, arXiv:2004.06252 (2020).
- [48] Jian Ma, Xiaoguang Wang, C. P. Sun, Franco Nori, *Quantum spin squeezing*, Physics Reports **509**, 2-3 (2011).
- [49] Duger Ulam-Orgikh and Masahiro Kitagawa, *Spin squeezing and decoherence limit in Ramsey spectroscopy*, Phys. Rev. A **64**, 052106 (2001).
- [50] Edward Grant, Leonard Wossnig, Mateusz Ostaszewski, and Marcello Benedetti, *An initialization strategy for addressing barren plateaus in parametrized quantum circuits*, Quantum **3**, 214 (2019).
- [51] K. Temme, T. J. Osborne, K. G. Vollbrecht, D. Poulin and F. Verstraete, *Quantum Metropolis sampling*, Nature **471**, 87 (2011).
- [52] Man-Hong Yung and Alán Aspuru-Guzik, *A quantum-quantum Metropolis algorithm*, PNAS **109**, 754 (2012).

- [53] Mohammad H. Amin, Evgeny Andriyash, Jason Rolfe, Bohdan Kulchytskyy and Roger Melko, *Quantum Boltzmann Machine* Phys. Rev. X **8**, 021050 (2018).
- [54] Fernando G.S.L. Brandao and Krysta Svore, *Quantum Speed-ups for Semidefinite Programming*, Proceedings FOCS 2017 (2017) and arXiv:1609.05537 (2016).
- [55] Joran van Apeldoorn, András Gilyén, *Improvements in Quantum SDP-Solving with Applications*, arXiv:1804.05058 (2018).
- [56] Dave Wecker, Matthew B. Hastings, Nathan Wiebe, Bryan K. Clark, Chetan Nayak and Matthias Troyer, *Solving strongly correlated electron models on a quantum computer*, Phys. Rev. A **92**, 062318 (2015).
- [57] Lukasz Cincio, Kenneth Rudinger, Mohan Sarovar and Patrick J. Coles, *Machine Learning of Noise-Resilient Quantum Circuits*, PRX Quantum **2** 010324 (2021).
- [58] Artur K. Ekert, Carolina Moura Alves, Daniel K. L. Oi, Michal Horodecki, Pawel Horodecki, and L. C. Kwek, *Direct estimations of linear and nonlinear functionals of a quantum state*, Phys. Rev. Lett. **88**, 217901 (2002).
- [59] Andrew M. Childs, Aaron Ostrander and Yuan Su, *Faster quantum simulation by randomization* Quantum **3**, 182 (2019).
- [60] Earl Campbell, *A random compiler for fast Hamiltonian simulation*, Phys. Rev. Lett. **123**, 070503 (2019).
- [61] Tyson Jones, Anna Brown, Ian Bush and Simon Benjamin, *QuEST and High Performance Simulation of Quantum Computers*, Scientific Reports **9**, 10736 (2019).
- [62] Tyson Jones, and Simon C Benjamin, *QuESTlink - Mathematica embiggened by a hardware-optimised quantum emulator*, Quantum Sci. Technol. **5** 034012 (2020).
- [63] Contributors and GSL Project, *GSL - GNU Scientific Library - GNU Project - Free Software*, The GNU Operating System (2010).
- [64] Gene Golub and Charles Van Loan, *Matrix computations*, The Mathematical Gazette **83**, 498 (1999).