
Cambridge Quantum

A Generic Compilation
Strategy for the Unitary
Coupled Cluster Ansatz

by Alexander Cowtan, Will Simmons, Ross Duncan

Cambridge Quantum
ALEXANDER COWTAN
alexander.cowtan@cambridgequantum.com
WILL SIMMONS
will.simmons@cambridgequantum.com

Cambridge Quantum
Department of Computer and Information Sciences
University of Strathclyde, United Kingdom
Department of Physics and Astronomy
University College London, United Kingdom
ROSS DUNCAN
ross.duncan@cambridgequantum.com

Cambridge Quantum
Terrington House
13-15 Hills Road
Cambridge CB2 1NL
United Kingdom

Published by Cambridge Quantum

20 July 2020

A Generic Compilation Strategy for the Unitary Coupled Cluster Ansatz

Alexander Cowtan^{1,*} Will Simmons¹ Ross Duncan^{1,2,3}

¹ *Cambridge Quantum Computing Ltd*

9a Bridge Street, Cambridge, United Kingdom

² *Department of Computer and Information Sciences
University of Strathclyde*

26 Richmond Street, Glasgow, United Kingdom

³ *Department of Physics and Astronomy
University College London*

Gower Street, London WC1E 6BT, United Kingdom

We describe a compilation strategy for Variational Quantum Eigensolver (VQE) algorithms which use the Unitary Coupled Cluster (UCC) ansatz, designed to reduce circuit depth and gate count. This is achieved by partitioning Pauli exponential terms into mutually commuting sets. These sets are then diagonalised using Clifford circuits and synthesised using the phase polynomial formalism. This strategy reduces CX depth by 75.4% on average, and by up to 89.9%, compared to naive synthesis for a variety of molecules, qubit encodings and basis sets.

1 Introduction

Many computational problems in quantum chemistry are classically intractable for systems which are large and strongly correlated [50]. Instead, quantum algorithms have been proposed [30] to simulate and calculate chemical properties of such systems. These algorithms leverage useful features of quantum mechanics to perform calculations which would either take too long or yield results too inaccurate using the best known classical algorithms.

However, the resources required by such algorithms tend to be too large for current quantum computers [45], which are limited in the number of qubits and the available circuit depth before decoherence and gate errors overwhelm the system and extracting a correct result from the noise becomes infeasible. These machines are known as Noisy Intermediate Scale Quantum (NISQ) devices [46].

A standard approach to reduce the resource requirements enough to run algorithms successfully on NISQ devices is to only run the quantum circuit as a subroutine in a larger, classical algorithm [41]. In this model, the quantum circuit prepares a parameterised state and measures the expectation value of a relevant operator. The classical routine then performs an optimisation algorithm, using the expectation value as an objective function, and attempts to minimise this value with respect to the circuit's parameters.

The Variational Quantum Eigensolver (VQE) is an archetypal hybrid quantum-classical algorithm, designed for the estimation of ground state energies of quantum systems on NISQ devices [51]. The expectation value of a molecular Hamiltonian is the objective function, and

*alexander.cowtan@cambridgequantum.com

VQE employs the variational principle to approximate the ground state of this Hamiltonian using the parameterised quantum circuit as an ansatz.

In this paper, we focus on the Unitary Coupled Cluster (UCC) ansatz [47], which is motivated by the orbital transitions allowed by the simulated system. We present a compilation strategy for reducing the major source of error affecting the algorithm: noise of the quantum device. This compilation strategy increases circuit fidelity by reducing circuit depth, hence minimising the number of noisy gates and the qubits’ exposure to decoherence.

For NISQ devices, two-qubit gates typically have error rates around an order of magnitude higher than one-qubit gates, as well as taking 2-5x as long [52, 8]. Defining the two-qubit gate depth as the number of two-qubit parallel layers required to complete the circuit, we aim to minimise this metric specifically with our compilation strategy, along with two-qubit gate count. We approximate the hardware-native two-qubit gate metrics with the corresponding CX metrics, noting that in certain scenarios this overstates the number of required gates. Two-qubit gates which are not maximally entangling, particularly tunable ones, can reduce the number of gates required for certain algorithms compared to using CX gates [9, 42].

We begin by partitioning the terms in the UCC ansatz into mutually commuting sets. We describe a well-known equivalence between this sequencing problem and graph colouring. We then show that approximate solutions to this problem enable large-scale synthesis of Pauli exponentials into one- and two-qubit gates, and propose heuristics for performing this synthesis to generate low depth circuits.

Our compilation strategy is valid for any ansatz which is generated by Trotterization of an operator made up of a sum of Pauli tensor products: this means that it is valid for k -UpCCGSD and other variations on the UCCSD ansatz, such as UCCGSD [35] and the parity-disregarding particle-exchange UCCSD [54]. It is not intended for the hardware efficient ansatz [29]. The strategy is generic, and requires no prior knowledge of the qubit encoding, target Hamiltonian or basis set.

We implemented the strategy in `t|ket` [48], our retargetable compiler, and present benchmarks for a variety of UCC circuits for realistic molecules to demonstrate empirically that the strategy significantly reduces the CX gate count and depth compared to previous strategies.

Related work: A similar strategy for optimizing Hamiltonian simulation circuits was recently presented by van den Berg & Temme [11]. This strategy uses different diagonalisation methods. In addition, the strategy is intended for fault-tolerant circuits, and the two-qubit gate reduction is obtained by targeting an ancilla qubit with every CX from each diagonal set, as previously described in Hastings et al. [27], which is impractical for some NISQ devices. However, a thorough comparison of strategies for Pauli partitioning and diagonalisation is presented, which can be applied in the NISQ setting.

Notation: In order to reason about and represent the synthesis of Pauli exponentials, we use notation inspired by the ZX-calculus [14], although our strategy can be followed without requiring any knowledge of the inference rules of the calculus. A brief introduction to the ZX-calculus is found in Fagan & Duncan [20]; for a complete treatment see Coecke & Kissinger [15].

Terminology: We refer to an n -qubit operator of the form $\{I, X, Y, Z\}^{\otimes n}$ as a *Pauli string*, composed of *letters* from the alphabet $\{I, X, Y, Z\}$. The *weight* of a Pauli string is the number of

non- I letters.

2 The Unitary Coupled Cluster Ansatz

The UCC ansatz is defined by the excitation of some reference state by an operator parameterised with coupled cluster amplitudes \vec{t} :

$$|\Psi(\vec{t})\rangle = U(\vec{t})|\Phi_0\rangle = e^{T(\vec{t})-T^\dagger(\vec{t})}|\Phi_0\rangle \quad (1)$$

where operator T is a linear combination of fermionic excitation operators $\vec{\tau}$ such that the parameterised operator can be rewritten:

$$U(\vec{t}) = e^{\sum_j t_j(\tau_j - \tau_j^\dagger)} \quad (2)$$

This parameterised operator cannot be directly implemented on a gate-based quantum computer. It must be mapped to qubits and decomposed into native gates.

In order to generate a quantum circuit, we employ Trotterization, justified by Lloyd [38]. Here we show the first order Lie-Trotter expansion:

$$U(\vec{t}) \approx U_{Trott}(\vec{t}) = \left(\prod_j e^{\frac{t_j}{\rho}(\tau_j - \tau_j^\dagger)}\right)^\rho \quad (3)$$

where ρ is the number of Trotter steps. Since our focus is on the NISQ setting, we will assume that only one Trotter step is taken. It is straightforward to extend the presented techniques to arbitrary step size.

To implement the Trotterized expression shown in Equation 3 on a quantum computer, we map the τ_j in the product to operations acting on qubits. This can be performed using a variety of qubit encodings, such as Bravyi-Kitaev (BK), Jordan-Wigner (JW) and parity (P) [49]. These encodings have different resource requirements and the qubits represent different physical properties, but regardless of our choice we obtain:

$$(\tau_j - \tau_j^\dagger) = ia_j \sum_k P_{jk} \quad (4)$$

where $a_j \in \mathbb{R}$ and $P_{jk} \in \{I, X, Y, Z\}^{\otimes n}$.

It can be shown that the Pauli strings P_{jk} from a given excitation operator τ_j always commute under multiplication [47]. This gives a simpler expression for the Trotterized operator,

$$U_{Trott}(\vec{t}) = \prod_j \prod_k e^{it_j a_j P_{jk}} \quad (5)$$

where $e^{it_j a_j P_{jk}}$ terms are parameterised with some angle t_j which will be adjusted by the variational algorithm. We refer to these terms as *Pauli exponentials*, and relabel our coefficients $t'_j = t_j a_j$.

Pauli exponentials can be implemented on a quantum computer by decomposition into one- and two-qubit native gates, discussed in Section 4. These gates are appended to a short, constant depth circuit generating the reference state.




Figure 1: Graph colouring to partition Pauli terms into sets of mutually commuting strings. While the parameters are not shown, they must be tracked for synthesis later.

3 Term Sequencing by Graph Colouring

Looking again at Equation 2, note that we can expand the fermionic excitation operators at this stage into Pauli strings using Equation 4, i.e. our chosen qubit encoding:

$$U(\vec{t}) = e^{i \sum_j \sum_k t_j P_{jk}} \quad (6)$$

Since addition is commutative we can freely choose the ordering of P_{jk} terms in this expression. After Trotterization the Pauli exponentials do not commute, so it is sensible at this stage to sequence the Pauli strings in a beneficial order, such that our Trotterization incurs minimal Trotter error and our circuit has low resource requirements. The implications of the ordering of terms for chemical accuracy have been studied by H. R. Grimsley et al. [24]. We justify in Appendix A that reducing Trotter error should be a secondary concern for near-term VQE, and focus on minimising CX gate count and depth.

Our strategy to reduce CX gate count and depth relies on partitioning the set of Pauli exponentials into a small number of subsets, such that within a given subset every Pauli exponential commutes.¹ This partitioning problem can be represented as the well-known graph colouring problem.

We represent each Pauli string as a vertex in an undirected graph. An edge is given between any two vertices that correspond to Pauli strings which anti-commute. Figure 1 shows an example of this graph representation.

Finding the minimum number of mutually commuting sets which cover all vertices in this graph is then equivalent to the *colouring problem*, a well known NP-hard problem [22]. In this instance, the colour assigned to the vertex corresponds to the subset the corresponding Pauli exponential is placed in, and since no two adjacent vertices can have the same colour, all Pauli exponentials within a subset will mutually commute.

¹This problem is common in the literature on measurement reduction [28, 18, 55].




Figure 2: Common circuit gates and their representations in the scalar-free ZX-calculus. The S gate corresponds to $R_Z(\frac{\pi}{2})$, and the V gate to $R_X(\frac{\pi}{2})$.

We use a simple greedy colouring algorithm to partition the Pauli strings. The complexity of this algorithm is $\mathcal{O}(m)$, with m the number of Pauli strings, although building the anti-commutation Pauli graph in the first place scales as $\mathcal{O}(m^2n)$, with n the number of qubits.


Once the vertices have been assigned colours, the UCC reference state $|\Phi_0\rangle$ is prepared and the corresponding Pauli exponential terms are appended, colour by colour, in lexicographical order. For example, given the graph colouring solution from Figure 1, a valid ordering of strings is: IIXY, IYX, XYII, YXII, XXXY, XXYX, XYXX, XYYY, YXXX, YXYY, YXXY, YYYX. Neither the order of the sets nor the order of terms within each set is considered important for optimisation; lexicographical order was an arbitrary choice.

4 Pauli Exponentials

A translation of relevant gates between the quantum circuit and ZX-calculus formalisms is given in Figure 2.

Recall the notation of *phase gadgets* $\Phi_n(\alpha)$, equivalent to the operator $e^{-i\frac{\alpha}{2}Z^{\otimes n}}$. These gadgets were described in Kissinger & van de Wetering [32], and have a natural representation in the ZX-calculus.

Definition 4.1. In ZX-calculus notation we have:




The algebra for phase gadgets and alternate decompositions into one- and two-qubit gates are given in Appendix B. Note that $\Phi_1(\alpha) = R_Z(\alpha)$.

The correspondence between phase gadgets and Pauli- Z exponentials generalises to any Pauli exponential $e^{-i\frac{\alpha}{2}P}$, by conjugating the phase gadget with single-qubit Clifford gates. We recall the *Pauli gadget* diagrammatic notation for the Pauli exponential from Cowtan et al. [17].

	Time complexity	CX complexity
Graph Colouring	$\mathcal{O}(m^2n)$	-
Diagonalisation	$\mathcal{O}(mn^3)$	$\mathcal{O}(n^2)$
GraySynth [5, 4]	$\mathcal{O}(mn^3)$	$\mathcal{O}(mn)$

Figure 3: Summary of subroutine complexities, where m is the total number of Pauli exponentials and n is the number of qubits. Time complexity refers to the compilation time, while CX complexity is defined as the maximum number of CX gates required for circuit synthesis. Graph colouring does not perform circuit synthesis so has no CX complexity.

Definition 4.2. Pauli exponentials are represented succinctly as:



The red, mixed-colour, and green boxes respectively represent the Pauli gadget acting on a qubit in the X , Y , and Z bases. These are formed by a phase gadget on the qubits (generating Z -only interactions), then conjugating the qubits with appropriate single-qubit Cliffords.

Clifford gates may be commuted through Pauli gadgets, but may incur a phase flip or basis change. The exhaustive set of diagrammatic rules required to perform this procedure for relevant Clifford gates are shown in Appendix C, although they are simple to calculate using linear algebra.

The definitions above imply a naive method of circuit synthesis for Pauli gadgets. For a set of m Pauli gadgets over n qubits, this naive synthesis requires $\mathcal{O}(nm)$ CX gates. More, precisely, we require at most $2m(n-1)$, if all the Pauli strings are maximal weight. This gives the baseline performance against which we will compare the method introduced in the next section.

5 Set-based Synthesis

The effect of the transformations in Section 3 is to partition our ansatz into large commuting sets of Pauli gadgets. The next step is to synthesise circuits for these sets, while minimising the CX overhead.

The approach we propose here has two steps:

1. Diagonalisation: every Pauli gadget in a given commuting set is simultaneously converted into a phase gadget by conjugating with an appropriate Clifford circuit.
2. Phase gadget synthesis: the resulting phase gadgets are converted into CX and R_Z gates using the well-studied *phase polynomial* formalism [7].

While diagonalisation incurs a gate overhead, in practice we find that the gate reduction from synthesising using this technique more than makes up for this overhead. Figure 3 summarises the relevant complexities of the different subroutines in our strategy.

5.1 Diagonalisation

Phase gadgets – that is, Pauli gadgets whose Pauli strings contain only the letters Z and I – define unitary maps which are diagonal in the computational basis. For this reason, we’ll call a *set* of Pauli gadgets diagonal when it contains only phase gadgets. Abusing terminology slightly, given a set of Pauli gadgets, we call a *qubit* diagonal over that set when the Pauli letter for that qubit is either Z or I for every Pauli string in the set. Evidently, if every qubit is diagonal then the set as a whole is too.

A set S of commuting Pauli gadgets can be simultaneously diagonalised using only Clifford operations. Our goal is to find a Clifford circuit C and a diagonal set S' such that

$$S = CS'C^\dagger \tag{7}$$

where C is as small as possible. Several methods have been proposed [28, 18, 1, 40] to compute a suitable polynomially-sized circuit C .

Note that, since $[A, B] = 0 \iff [UAU^\dagger, UBU^\dagger] = 0$ for unitaries A , B and U , conjugating the gadgets preserves commutation, so the required C can be constructed by conjugation with Cliffords. Below, we use this approach on *compatible pairs* of qubits, where one qubit can be used to diagonalise the other. In the worst case C has $\mathcal{O}(n^2)$ CX gates; however in practice, on realistic examples of UCC ansatz circuits, we find our method typically produces Clifford diagonalisers much smaller than the asymptotic worst case.

Remark 5.1. Jena et al. [28] presented an algorithm guaranteed to give such a C for qudits of any dimension, of size quadratic in the number of qudits. For m Pauli gadgets, Crawford et al. [18] recently presented two efficient constructions of C with a bound of $mn - m(m+1)/2$ and $\mathcal{O}(mn/\log m)$ CX gates respectively, when $m < n$. When $m \geq n$, the construction provided by Aaronson & Gottesman requires $\mathcal{O}(n^2/\log n)$ CX gates [1, 40].

5.1.1 Diagonalising a compatible pair

In the following, let S be a set of m commuting Pauli gadgets acting on n qubits, and let σ_{kl} denote the Pauli letter on qubit k from gadget l .


Definition 5.2. Let $i, j \in \{1, \dots, n\}$ with $i \neq j$. Qubits i and j are called *compatible* in S if the following relation holds:

$$\exists A, B \in \{X, Y, Z\} \text{ s.t. } \forall l \in \{1, \dots, m\}, \sigma_{il} \in \{I, A\} \iff \sigma_{jl} \in \{I, B\} ; \tag{8}$$


In this case i and j are called a *compatible pair*.

Theorem 5.3. *If qubits i and j form a compatible pair in S then one of them can be diagonalised by conjugating S with a single CX and at most two single qubit Cliffords acting on qubits i and j .*


Proof. Without loss of generality, assume $i = 1$ and $j = 2$, and let P_2 be a 2-qubit Pauli string. Applying the transformation $P_2 \mapsto \text{CX} \circ P_2 \circ \text{CX}$ (with the second qubit the target) will diagonalise the second qubit when $P_2 \in \{II, IZ, XX, XY, YX, YY, ZI, ZZ\}$. This set satisfies the property $\sigma_1 \in \{Z, I\} \iff \sigma_2 \in \{Z, I\}$, and all other sets of 2-qubit Paulis which satisfy this property are subsets of this one. (Note that the control qubit will be diagonal after conjugation iff it was diagonal before.) Conjugating the first and/or second qubit by an additional single qubit Clifford allows this relation to be generalised to any pair of Paulis as in (8), giving the result. \square



(a) A compatible pair.



(b) Conjugation with appropriate single-qubit Cliffords.



(c) Conjugation with CX gates to diagonalise the second qubit.

Figure 4: Application of Theorem 5.3 to diagonalise a qubit.

If i and j are a compatible pair, then the specific values of A and B in relation (8) determine which single-qubit Cliffords are required before conjugation by CX gates will diagonalise a qubit. An example is shown in Figure 4a. The first two qubits are compatible, with $A = Y$ and $B = Y$, which implies that V and V^\dagger gates are required to prepare the second qubit for diagonalisation as shown in Figure 4b. The diagonalisation is completed by CX conjugation as shown in Figure 4c.

Applying Theorem 5.3 to compatible pairs of qubits is the key subroutine in our diagonalisation algorithm, described in the next section.

Corollary 5.4. *For any commuting set of m gadgets over n qubits, if $m < 4$ a Clifford circuit C exists which diagonalises this set of gadgets using at most $n - 1$ CX gates.*

Proof. See Appendix D. □

Corollary 5.5. *For any commuting set of gadgets over n qubits, if $n < 5$ a Clifford circuit C exists which diagonalises this set of gadgets using at most $n - 1$ CX gates.*

Proof. See Appendix E. □

5.1.2 Diagonalising a commuting set

This section describes our method for diagonalising a set of commuting Pauli gadgets. The basic approach is to repeatedly apply three methods which diagonalise a single qubit :

1. Diagonalise the trivially diagonalisable qubits
2. Diagonalise qubits in compatible pairs
3. Synthesise a single gadget to diagonalise one of its qubits.

Detailed pseudo-code for the algorithm² is presented in Figure 5. The overall time complexity for this algorithm is $\mathcal{O}(mn^3)$, with m the number of Pauli gadgets in the commuting set and n the number of qubits.

To make the algorithm clearer, we'll go through a worked example in Figures 6 and 7. Initially we have the mutually commuting gadgets shown in Figure 6a, corresponding to the Pauli strings $IXZIZ$, $IYIZY$, $XXIYI$, $YYXII$, $ZIYXX$, $ZXIZZ$, and $ZYZIY$. We proceed as follows.

1. First, check whether there is a trivially diagonalisable qubit: that is, a qubit i for which $\exists P \in \{X, Y, Z\}$ s.t. $\forall l, \sigma_{il} \in \{I, P\}$. Any such qubits may be diagonalised with only single-qubit Clifford gates, and ignored from now on. This check takes time $\mathcal{O}(mn)$. Figure 6a contains no such qubits.
2. Now search for a compatible pair of qubits (Defn. 5.2) Theorem 5.3 for any choice of Paulis A and B , and apply the conjugation of Theorem 5.3 to diagonalise a qubit and remove it from consideration. The choice of qubit within the pair is arbitrary. This search can be performed in $\mathcal{O}(mn^2)$. The example of Figure 6a does not contain a compatible pair.
3. If no compatible pair is found, we adopt a greedy approach as a backup strategy. In $\mathcal{O}(m)$ time, find the Pauli string with the lowest weight; if there are multiple pick arbitrarily. Conjugate the corresponding Pauli gadget with single-qubit Clifford and CX gates to convert the Pauli string to $II \dots IZ$, demonstrated in Figure 6b. Then, commute the Clifford gates through the rest of the gadgets, as shown in Figure 6c, until all Clifford gates are outside

² We have omitted the greedy diagonalisation method from the pseudo-code, as it is straightforward.

```

function GADGETDIAG( $S$ )
   $Q \leftarrow \text{Qubits}(S)$ 
   $C \leftarrow \text{EmptyCircuit}(Q)$ 
  while  $Q$  non-empty do
    ( $S, Q, C$ )  $\leftarrow$  UpdateSingleQubits( $S, Q, C$ )
    if  $Q$  empty then
      break
    end if
     $Q' \leftarrow Q$ 
    ( $S, Q, C$ )  $\leftarrow$  UpdatePairQubits( $S, Q, C$ )
    if  $Q = Q'$  then
      ( $S, Q, C$ )  $\leftarrow$  GreedyDiagonalisation( $S, Q, C$ )
    end if
  end while
  return ( $S, C$ )
end function

function UPDATESINGLEQUBITS( $S, Q, C$ )
  for  $q \in Q$  do
     $p \leftarrow \text{FindCommonPauli}(S, q)$   $\triangleright p : \text{Maybe Pauli}$ 
    if  $p \neq \text{None}$  then
       $S \leftarrow \text{UpdateGadgetsSingleQubit}(S, p, q)$ 
       $Q \leftarrow Q \setminus \{q\}$ 
       $C \leftarrow \text{AddCliffordsSingleQubit}(C, p, q)$ 
    end if
  end for
  return ( $S, Q, C$ )
end function

function UPDATEPAIRQUBITS( $S, Q, C$ )
  for  $q_a \in Q$  do
    for  $q_b \in Q \setminus \{q_a\}$  do
      ( $p_a, p_b$ )  $\leftarrow$  FindValidPaulis( $S, q_a, q_b$ )  $\triangleright (p_a, p_b) : \text{Maybe Pair Pauli}$ 
      if ( $p_a, p_b$ )  $\neq$  None then
         $S \leftarrow \text{UpdateGadgetsPairQubit}(S, p_a, p_b, q_a, q_b)$ 
         $Q \leftarrow Q \setminus \{q_b\}$ 
         $C \leftarrow \text{AddCliffordsPairQubit}(C, p_a, p_b, q_a, q_b)$ 
        return ( $S, Q, C$ )
      end if
    end for
  end for
  return ( $S, Q, C$ )
end function

```

Figure 5: Diagonalisation algorithm

the adjacent Pauli gadgets. Every gadget must still commute with the $II\dots IZ$ string, and therefore the last qubit must be diagonal. This is a similar method to Jena et al. [28].

These steps are repeated until all qubits are diagonal over the set of Pauli gadgets. Following our example, we find that Figure 6c has the same two-qubit chain on the first and second qubits as our example from Figure 4a, and can therefore be diagonalised in the same way, resulting in the circuit in Figure 7a. The backup strategy is not required for the remaining qubits. See Figure 7b for the circuit after full diagonalisation.

Since each iteration will diagonalise at least one qubit, $\mathcal{O}(n)$ repetitions are required, so Algorithm 5 has time complexity $\mathcal{O}(mn^3)$. In the worst case, the greedy approach is required repeatedly, so C will require at most $\frac{1}{2}n(n-1)$ CX gates. If the greedy approach is not required at all, C will require at most $n-1$ CX gates. For our small, 5-qubit example circuit, the greedy approach was required at one iteration, and C used 5 CX gates.

5.2 Phase Polynomials

After diagonalisation, we have a circuit $CS'C^\dagger$ where the interior section S' is comprised entirely of phase gadgets, with each gadget acting on a different set of qubits. This phase gadget circuit can be expressed as a *phase polynomial*.

Proposition 5.6 (Nam et al. [43]). *Let D be a quantum circuit containing only CX gates and the gates $R_Z(\theta_1), R_Z(\theta_2), \dots, R_Z(\theta_m)$. The action of D on a basis state $|x_1, x_2 \dots x_n\rangle$ has the form:*

$$D|x_1, x_2 \dots x_n\rangle = e^{ip(x_1, x_2, \dots, x_n)} |h(x_1, x_2, \dots, x_n)\rangle \quad (9)$$


where $h(x_1, x_2, \dots, x_n)$ is a linear reversible function and

$$p(x_1, x_2, \dots, x_n) = \sum_{i=1}^m \theta_i f_i(x_1, x_2, \dots, x_n) \quad (10)$$

is a linear combination of linear Boolean functions $f_i: \{0, 1\}^n \rightarrow \{0, 1\}$.

Definition 5.7. Given D as above, the *phase polynomial* of circuit D is $p(x_1, x_2, \dots, x_n)$, and each f_i is called a *parity*.


Example 5.8. The circuit shown below has the required form.




We can read that the corresponding phase polynomial is $p(q_1, q_2) = \alpha(q_1 \oplus q_2)$, defined on the single parity $q_1 \oplus q_2$.

Phase polynomials have been studied for the purposes of circuit optimisation³ [43, 40, 7, 3]. Phase polynomial *synthesis* refers to the task of generating a circuit over a chosen gate set, usually CX and $R_Z(\theta)$, which implements a given phase polynomial with minimal resources. Optimal synthesis is NP-complete in specific cases [5], but the time complexity of the general


³The representation of phase gadget circuits as phase polynomials has also inspired circuit optimisation techniques [10].




(a) Example set of adjacent commuting Pauli gadgets.



(b) Leftmost Pauli gadget before and after decomposition.




(c) Pauli gadget set after commuting Cliffords through.




(d) Theorem 5.3 is satisfied by $A = Y$ and $B = Y$ for the first two qubits. Single-qubit rotations are applied to convert to the Z -basis.


Figure 6: Set diagonalisation example.



(a) Diagonalise the second qubit with a CX.



(b) Repeat the procedure for the remaining qubits to fully diagonalise the set.




(c) Convert phase gadgets to CX and R_Z gates using GraySynth [5].

Figure 7: Set diagonalisation example (cont'd).

case remains open. In practice, heuristic methods such as the GraySynth algorithm of Amy et al. [5] can achieve significant reductions in CX count.

The circuit of Example 5.8 can be equivalently written as a phase gadget over two qubits. In fact, every n -qubit phase gadget is equivalent to a phase polynomial with a single term in the summation of Eq. (10), so each phase gadget corresponds to a parity f_i and a rotation θ_i .



More generally, a circuit comprising only phase gadgets can be represented by a phase polynomial where the linear reversible function $h(x_1, x_2, \dots, x_n)$ of Eq. (9) is the identity. This allows us to use techniques based on phase polynomials to synthesise a circuit for S' .

While any synthesis method could be used, for the results described here we chose the heuristic GraySynth method [5] because it produces an efficient circuit⁴ at reasonable computational cost. If a specific qubit architecture was required, then an architecture-aware synthesis method would be more appropriate [44, 23]. GraySynth runs in time $\mathcal{O}(mn^3)$, and requires a maximum of $\mathcal{O}(mn)$ CX gates when the linear reversible function h is identity [4]. For reasons of space we omit the algorithm.

Returning to the running example, the synthesised circuit generated from the interior phase gadgets is shown in Figure 7c. Using a naive decomposition, as described in Definitions 4.1 and 4.2, the initial set from Figure 6a would have required 34 CX gates, and 34 CX depth. Our strategy has reduced the CX count to 22, and the CX depth to 18.

6 Results and Discussion

We implemented our strategy in `t|ket` [48], our retargetable compiler. We benchmarked this implementation on a suite of ansatz circuits for electronic structure UCCSD (Unitary Coupled Cluster Singles and Doubles) VQE problems. We included the molecules H_2 , H_4 , H_8 , LiH , BeH_2 , NH , H_2O , CH_2 , NH_3 , HNO , HCl , N_2 , C_2 , H_2CO , CO_2 and C_2H_4 in the ‘sto-3g’ basis set. For the smaller molecules, we also used the ‘631g’ basis. We tested using the Bravyi-Kitaev (BK), Jordan-Wigner (JW) and parity (P) encodings.

The comparisons made are:

1. Naive decomposition: circuits generated from Equation 5 by decomposing Pauli gadgets naively into CX and single-qubit gates, as described in Section 4.
2. Pairwise synthesis: circuits generated by graph colouring and then synthesising Pauli gadgets within a set in a pairwise manner with CX balanced trees using the methods from Cowtan et al. [17].
3. Set synthesis: our full compilation strategy. Graph colouring, diagonalisation and phase polynomial synthesis.

⁴ Across a suite of Clifford+ T benchmark circuits, the implementation of Amy et al. reduced the CX gate count by 23% with a maximum reduction of 43% [5].

4. Templated lexicographical operator sequence (TLOS): for ansatz circuits generated using the JW encoding we compare against a mock implementation of the best known previous strategy for JW circuit synthesis: operator sequencing methods from Hastings et al. [27] allowing for CX cancellation between excitations, with templated excitation operators from Nam et al. [42] for low CX count excitations⁵. We are not aware of similar strategies for the BK or P encoding.
5. van den Berg & Temme (vdBT): a strategy for optimizing Hamiltonian simulation circuits [11]. While similar in aims to our strategy, the methods for diagonalisation and two-qubit gate reduction are quite different. We compare against the best performing method, the "cz-rnd" method, which executes the procedure stochastically several times and picks the best result. Data is available only for the CX count with the JW encoding on some molecules.

Circuits in our test set were chosen to have a CX count and depth of less than 10^6 when decomposed naively. All results were obtained using `pytket v0.5.5`, on a machine with a 2.3 GHz Intel Core i5 processor and 8 GB of 2133 MHz LPDDR3 memory, running MacOS Mojave v10.14.

A benchmark script for reproducing the results, along with the input operators, can be found at https://github.com/CQCL/tket_benchmarking/tree/master/compilation_strategy. The methodology for generating and serialising these operators is described in Appendix F.

A comparison of CX metrics for different compilation strategies, active spin orbitals and qubit encoding methods is shown in Figure 8.


The set-based synthesis strategy outperforms pairwise and naive strategies on all encodings, but is on average outperformed by the TLOS method for the JW encoding, particularly for larger systems with more active spin orbitals.

Set-based synthesis gives greater fractional reductions for larger circuits than for smaller ones. For the largest circuits, up to 89.9% CX depth reduction can be achieved, compared to the mean CX depth reduction of 75.4% shown in Figure 9. As the compilation strategy is composed of several heuristics in sequence, we do not at this stage argue that the asymptotic complexity of the UCC ansatz can be reduced - in order to do this, we would need to prove sufficient bounds on the size of sets found by graph colouring, the CX complexities of Clifford circuits required for diagonalisation and the number of CX gates produced by phase polynomial synthesis.


Remark: This compilation strategy assumes that the qubits have all-to-all connectivity, so CX gates are allowed between any two qubits. When connectivity is constrained, *routing* is required to ensure the circuit conforms to the constraints. The typical approach to this problem is SWAP network insertion [12, 56, 34, 16].

Remark: While VQE using the UCC ansatz is a candidate for quantum advantage, there are no complexity-theoretic guarantees of success. Should the advantage be sufficiently small, the low-degree polynomial compilation time required for this strategy could be too slow. In this case, we emphasise that *co-design* of a compilation strategy with the qubit encoding can give large reductions, shown by the TLOS method, while reducing compilation time.


⁵We do not allow the use of ancilla qubits for this method, which Hastings et al. showed can reduce CX overhead significantly. Additionally, Nam et al. used a bosonic excitation technique relating molecular and spin orbitals, which we do not include here.



(a) Bravyi-Kitaev qubit encoding.



(b) Jordan-Wigner qubit encoding.



(c) Parity qubit encoding.

Figure 8: Comparison of compilation strategies for molecules with varying active spin orbital counts using different qubit encoding methods. A 4th-degree polynomial least-squares fit has been added to suggest scaling.

	Mean CX count reduction (%)	Mean CX depth reduction (%)
Pairwise Synthesis	40.0	56.9
Set-based Synthesis	63.6	71.9

(a) Bravyi-Kitaev qubit encoding.

	Mean CX count reduction (%)	Mean CX depth reduction (%)
Pairwise Synthesis	49.9	67.9
Set-based Synthesis	78.0	82.1
TLOS Synthesis	82.6	84.5
vdBT Synthesis	64.0	—

(b) Jordan-Wigner qubit encoding.

	Mean CX count reduction (%)	Mean CX depth reduction (%)
Pairwise Synthesis	38.1	55.7
Set-based Synthesis	65.3	72.1

(c) Parity qubit encoding.

	Mean CX count reduction (%)	Mean CX depth reduction (%)
Pairwise Synthesis	42.7	60.2
Set-based Synthesis	69.0	75.4

(d) All encodings.

Figure 9: Mean CX metric reductions. All reductions are measured against the naive decomposition method.

7 Conclusions and Future Work

The primary contribution of our paper is an empirically successful method to efficiently synthesise the UCC ansatz to one- and two-qubit gates. We have shown large average reductions in CX metrics for the Bravyi-Kitaev, Jordan-Wigner, and parity qubit encodings; although alternative methods are competitive with ours for the JW encoding, we emphasise that our strategy is valid for any other qubit encodings which generate similar Trotterized excitation operators. We note that the reductions for the JW encoding are the greatest, with respect to both metrics and both the pairwise and set-based synthesis methods. This may suggest that this encoding has more exploitable redundancy than the BK or P encodings.

We briefly discuss four future directions to explore.

7.1 Applications to Measurement Reduction

Measurement reduction for VQE is a method to simultaneously measure terms in a Hamiltonian which commute, and therefore reduce the number of circuits required to run [28, 18, 55]. For

realistic devices, assuming that the only available native measurements are single-qubit Z -basis measurements, generating a Clifford circuit to diagonalise this set is required. Minimising this Clifford circuit using applications of Theorem 5.3 can reduce the CX overhead required for measurement reduction.

7.2 Architecture-Aware Synthesis

Instead of introducing a SWAP network to enforce connectivity constraints on NISQ devices, recent work has explored the possibility of resynthesising the circuit in a topologically aware manner, for limited gate sets [31, 44, 53]. This constrained synthesis has been found to typically produce lower CX counts than SWAP networks, and phase polynomials are a viable class of circuit for constrained synthesis [6, 23]. If topologically constrained phase polynomials can be composed with Clifford regions in a manner that respects architecture, this would appear to be a suitable strategy for those devices with limited connectivity.

7.3 Applications to QAOA

The Quantum Approximate Optimisation Algorithm (QAOA) [21] for combinatorial optimisation problems consists of repeated blocks of ‘mixing’ and ‘driver’ exponentiated Hamiltonians. The driver Hamiltonians are already diagonal, as they encode a classical system, and typically the mixing Hamiltonians correspond to single-qubit gates only. However, recent work on a so-called Quantum Alternating Operator Ansatz [26] introduces more complicated mixing Hamiltonians. These mixing Hamiltonians could be amenable to our compilation strategy.

7.4 Applications to Fault Tolerant Computation

While this strategy was designed specifically for VQE, it can be directly ported over to non-variational quantum algorithms for Hamiltonian dynamics which require gate counts and qubit numbers too high for NISQ computers. Gui et al. [25] and van den Berg & Temme [11] have performed term sequencing similar to our work in Section 3 for digital quantum simulation, in which a quantum evolution defined by a time-dependent Hamiltonian is mapped to a quantum circuit. Reducing Trotter error is more important for fault-tolerant algorithms than for VQE, as it is the only significant non-correctable source of error, and Gui et al. argue that our term sequencing method would also minimise Trotter error.

The efficacy of our proposed compilation strategy is greatly dependent on the model of fault-tolerant computation. For example, in the model presented by Litinski [36], all non-Clifford Pauli exponentials are performed natively by performing ancilla measurements. In this model, the exponentials do not need to be converted to one- and two-qubit gates at all.

Even for models which perform individual gates explicitly, our proposed compilation strategy is optimised to reduce two-qubit gate count and depth, which are not considered as important on planned fault tolerant devices as non-Clifford gates, as the CX gate can be performed without magic state distillation. However, on surface code-based computers performing lattice surgery, a two-qubit gate between distant logical qubits can be as costly in the worst case as magic state distillation [37]; in general, two-qubit gates increase the overhead from routing on surface codes. Therefore, two-qubit gate reduction may still be a valuable optimisation.

Moreover, the circuits produced by the strategy are structured such that all non-Clifford rotations reside in partitioned phase polynomials, and will be approximated with T gates and

one-qubit Cliffords. T -count and T -depth optimisation has been successfully performed using the phase polynomial formalism via matroid partitioning [7]. The T -count of phase polynomials generated via diagonalisation cannot be optimised using phase-folding, as the parities are guaranteed to be unique, but T -depth reduction could still be enabled by our strategy.

Acknowledgements

The authors would like to thank John van de Wetering, Arianne Meijer, David Zsolt-Manrique and Irfan Khan for helpful discussions, and Matthew Amy for correspondence on the GraySynth algorithm.

References

- [1] S. Aaronson & D Gottesman (2004): *Improved Simulation of Stabilizer Circuits*. *Phys. Rev. A* 70(052328), doi:10.1103/PhysRevA.70.052328. arXiv:quant-ph/0406196v5.
- [2] Gadi Aleksandrowicz, et al. (2019): *Qiskit: An Open-source Framework for Quantum Computing*, doi:10.5281/zenodo.2562110.
- [3] M. Amy, D. Maslov, M. Mosca & M. Roetteler (2013): *A Meet-in-the-Middle Algorithm for Fast Synthesis of Depth-Optimal Quantum Circuits*. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 32(6), pp. 818–830, doi:10.1109/TCAD.2013.2244643. arXiv:1206.0758.
- [4] Matthew Amy (2020): *Personal correspondence*.
- [5] Matthew Amy, Parsiad Azimzadeh & Michele Mosca (2018): *On the controlled-NOT complexity of controlled-NOT-phase circuits*. *Quantum Science and Technology* 4(1), p. 015002, doi:10.1088/2058-9565/aad8ca. arXiv:1712.01859.
- [6] Matthew Amy & Vlad Gheorghiu (2020): *staq – A full-stack quantum processing toolkit*. *Quantum Science and Technology* 5(3), p. 034016, doi:10.1088/2058-9565/ab9359. arxiv:1912.06070.
- [7] Matthew Amy, Dmitri Maslov & Michele Mosca (2014): *Polynomial-Time T -Depth Optimization of Clifford+ T Circuits Via Matroid Partitioning*. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 33(10), pp. 1476–1489, doi:10.1109/TCAD.2014.2341953.
- [8] Frank Arute, et al. (2019): *Quantum supremacy using a programmable superconducting processor*. *Nature* 574(7779), pp. 505–510, doi:10.1038/s41586-019-1666-5.
- [9] Frank Arute, et al. (2020): *Quantum Approximate Optimization of Non-Planar Graph Problems on a Planar Superconducting Processor*. arXiv.org. arXiv:2004.04197.
- [10] Niel de Beaudrap, Xiaoning Bian & Quanlong Wang (2019): *Techniques to reduce $\pi/4$ -parity phase circuits, motivated by the ZX calculus*. arXiv.org. arXiv:1911.09039.
- [11] Ewout van den Berg & Kristan Temme (2020): *Circuit optimization of Hamiltonian simulation by simultaneous diagonalization of Pauli clusters*. arXiv. arXiv:2003.13599.
- [12] Andrew M. Childs, Eddie Schoute & Cem M. Unsal (2019): *Circuit Transformations for Quantum Architectures*. In Wim van Dam & Laura Mancinska, editors: *14th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2019)*, *Leibniz International Proceedings in Informatics (LIPIcs)* 135, pp. 3:1–3:24, doi:10.4230/LIPIcs.TQC.2019.3. Available at <http://drops.dagstuhl.de/opus/volltexte/2019/10395>. arXiv:1902.09102.
- [13] Andrew M. Childs, Yuan Su, Minh C. Tran, Nathan Wiebe & Shuchen Zhu (2019): *A Theory of Trotter Error*. arXiv.org. arXiv:1912.08854.

- [14] Bob Coecke & Ross Duncan (2011): *Interacting Quantum Observables: Categorical Algebra and Diagrammatics*. *New J. Phys* 13(043016), doi:10.1088/1367-2630/13/4/043016. Available at <http://iopscience.iop.org/1367-2630/13/4/043016/>. arXiv:0906.4725.
- [15] Bob Coecke & Aleks Kissinger (2017): *Picturing Quantum Processes: A First Course in Quantum Theory and Diagrammatic Reasoning*. Cambridge University Press, doi:10.1017/9781316219317.
- [16] Alexander Cowtan, Silas Dilkes, Ross Duncan, Alexandre Krajenbrink, Will Simmons & Seyon Sivarajah (2019): *On the qubit routing problem*. In Wim van Dam & Laura Mancinska, editors: *14th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2019), Leibniz International Proceedings in Informatics (LIPIcs)* 135, pp. 5:1–5:32, doi:10.4230/LIPIcs.TQC.2019.5. Available at <http://drops.dagstuhl.de/opus/volltexte/2019/10397>.
- [17] Alexander Cowtan, Silas Dilkes, Ross Duncan, Will Simmons & Seyon Sivarajah (2020): *Phase Gadget Synthesis for Shallow Circuits*. In Bob Coecke & Matthew Leifer, editors: *Proceedings 16th International Conference on Quantum Physics and Logic (QPL 2019), Chapman University, Orange, CA, USA., 10-14 June 2019*, 318, pp. 213–228, doi:10.4204/EPTCS.318.13. arXiv:1906.01734.
- [18] Ophelia Crawford, Barnaby van Straaten, Daochen Wang, Thomas Parks, Earl Campbell & Stephen Brierley (2019): *Efficient quantum measurement of Pauli operators*. arXiv.org. arXiv:1908.06942.
- [19] Andrew W. Cross, Lev S. Bishop, John A. Smolin & Jay M. Gambetta (2017): *Open Quantum Assembly Language*. arXiv.org. Available at <https://github.com/IBM/qiskit-openqasm/blob/master/spec/qasm2.pdf>. arXiv:1707.03429.
- [20] Andrew Fagan & Ross Duncan (2019): *Optimising Clifford Circuits with Quantomatic*. In Peter Selinger & Giulio Chiribella, editors: *Proceedings of the 15th International Conference on Quantum Physics and Logic*, Halifax, Canada, 3-7th June 2018, *Electronic Proceedings in Theoretical Computer Science* 287, Open Publishing Association, pp. 85–105, doi:10.4204/EPTCS.287.5. arXiv:1901.10114.
- [21] Edward Farhi, Jeffrey Goldstone & Sam Gutmann (2014): *A Quantum Approximate Optimization Algorithm*. arXiv.org. arXiv:1411.4028.
- [22] M. R. Garey, D. S. Johnson & L. Stockmeyer (1974): *Some Simplified NP-Complete Problems*. In: *Proceedings of the Sixth Annual ACM Symposium on Theory of Computing*, STOC '74, Association for Computing Machinery, New York, NY, USA, pp. 47–63, doi:10.1145/800119.803884.
- [23] Arianne Meijer van de Griend & Ross Duncan (2020): *Architecture-aware synthesis of phase polynomials for NISQ devices*. In: *Proceedings of QPL2020 (to appear)*. arXiv:2004.06052.
- [24] Harper R. Grimsley, Daniel Claudino, Sophia E. Economou, Edwin Barnes & Nicholas J. Mayhall (2019): *Is the Trotterized UCCSD Ansatz Chemically Well-Defined?* *Journal of Chemical Theory and Computation* 16(1), pp. 1–6, doi:10.1021/acs.jctc.9b01083.
- [25] Kaiwen Gui, Teague Tomesh, Pranav Gokhale, Yunong Shi, Frederic T. Chong, Margaret Martonosi & Martin Suchara (2020): *Term Grouping and Travelling Salesperson for Digital Quantum Simulation*. arXiv.org. arXiv:2001.05983.
- [26] Stuart Hadfield, Zihui Wang, Bryan O’Gorman, Eleanor Rieffel, Davide Venturelli & Rupak Biswas (2019): *From the Quantum Approximate Optimization Algorithm to a Quantum Alternating Operator Ansatz*. *Algorithms* 12(2), p. 34, doi:10.3390/a12020034. arXiv:1709.03489.
- [27] Matthew B. Hastings, Dave Wecker, Bela Bauer & Matthias Troyer (2015): *Improving Quantum Algorithms for Quantum Chemistry*. *Quantum Information and Computation* 15. arXiv:1403.1539.
- [28] Andrew Jena, Scott Genin & Michele Mosca (2019): *Pauli Partitioning with Respect to Gate Sets*. arXiv.org. arXiv:1907.07859.
- [29] Abhinav Kandala, Antonio Mezzacapo, Kristan Temme, Maika Takita, Markus Brink, Jerry M. Chow & Jay M. Gambetta (2017): *Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets*. *Nature* 549, doi:10.1038/nature23879.

- [30] Ivan Kassal, James D. Whitfield, Alejandro Perdomo-Ortiz, Man-Hong Yung & Alán Aspuru-Guzik (2011): *Simulating Chemistry Using Quantum Computers*. *Annual Review of Physical Chemistry* 62(1), pp. 185–207, doi:10.1146/annurev-physchem-032210-103512. arXiv:1007.2648.
- [31] Aleks Kissinger & Arianne Meijer van de Griend (2020): *CNOT circuit extraction for topologically-constrained quantum memories*. *Quantum Information and Computation* 20(7&8), pp. 581–596, doi:10.26421/QIC20.7-8. arXiv:1904.00633.
- [32] Aleks Kissinger & John van de Wetering (2019): *Reducing T-count with the ZX-calculus*. arXiv.org. arXiv:1903.10477.
- [33] Vadym Kliuchnikov & Dmitri Maslov (2013): *Optimization of Clifford circuits*. *Phys. Rev. A* 88, p. 052307, doi:10.1103/PhysRevA.88.052307.
- [34] Lingling Lao, Daniel M. Manzano, Hans van Someren, Imran Ashraf & Carmen G. Almudever (2019): *Mapping of quantum circuits onto NISQ superconducting processors*. arXiv.org. arXiv:1908.04226.
- [35] Joonho Lee, William J. Huggins, Martin Head-Gordon & K. Birgitta Whaley (2019): *Generalized Unitary Coupled Cluster Wave functions for Quantum Computation*. *Journal of Chemical Theory and Computation* 15(1), pp. 311–324, doi:10.1021/acs.jctc.8b01004. arXiv:1810.02327.
- [36] Daniel Litinski (2019): *A Game of Surface Codes: Large-Scale Quantum Computing with Lattice Surgery*. *Quantum* 3, p. 128, doi:10.22331/q-2019-03-05-128. Available at <https://doi.org/10.22331/q-2019-03-05-128>.
- [37] Daniel Litinski (2019): *Magic State Distillation: Not as Costly as You Think*. *Quantum* 3, p. 205, doi:10.22331/q-2019-12-02-205.
- [38] Seth Lloyd (1996): *Universal Quantum Simulators*. *Science* 273(5278), pp. 1073–1078, doi:10.1126/science.273.5278.1073.
- [39] Guang Hao Low, Vadym Kliuchnikov & Nathan Wiebe (2019): *Well-conditioned multiproduct Hamiltonian simulation*. arXiv.org. 1907.11679.
- [40] Dmitri Maslov & Martin Roetteler (2017): *Shorter stabilizer circuits via Bruhat decomposition and quantum circuit transformations*. *IEEE Transactions on Information Theory* (7), pp. 4729–4738, doi:10.1109/TIT.2018.2825602. arXiv:1705.09176.
- [41] Jarrod R McClean, Jonathan Romero, Ryan Babbush & Alán Aspuru-Guzik (2016): *The theory of variational hybrid quantum-classical algorithms*. *New Journal of Physics* 18(2), p. 023023, doi:10.1088/1367-2630/18/2/023023. arXiv:1509.04279.
- [42] Yunseong Nam, Jwo-Sy Chen, Neal C. Pienti, Kenneth Wright, Conor Delaney, Dmitri Maslov, Kenneth R. Brown, Stewart Allen, Jason M. Amini, Joel Apisdorf, Kristin M. Beck, Aleksey Blinov, Vandiver Chaplin, Mika Chmielewski, Coleman Collins, Shantanu Debnath, Andrew M. Ducore, Kai M. Hudek, Matthew Keesan, Sarah M. Kreikemeier, Jonathan Mizrahi, Phil Solomon, Mike Williams, Jaime David Wong-Campos, Christopher Monroe & Jungsang Kim (2020): *Ground-state energy estimation of the water molecule on a trapped ion quantum computer*. *npj Quantum Information* 6(33), doi:10.1038/s41534-020-0259-3. arXiv:1902.10171.
- [43] Yunseong Nam, Neil J. Ross, Yuan Su, Andrew M. Childs & Dmitri Maslov (2018): *Automated optimization of large quantum circuits with continuous parameters*. *npj Quantum Information* 4(1), p. 23, doi:10.1038/s41534-018-0072-4.
- [44] Beatrice Nash, Vlad Gheorghiu & Michele Mosca (2020): *Quantum circuit optimizations for NISQ architectures*. *Quantum Science and Technology* 5(025010), doi:10.1088/2058-9565/ab79b1. arXiv:1904.01972.
- [45] David Poulin, Matthew B. Hastings, Dave Wecker, Nathan Wiebe, Andrew C. Doherty & Matthias Troyer (2015): *The Trotter Step Size Required for Accurate Quantum Simulation of Quantum Chemistry*. *Quantum Information and Computation* 15(5&6), pp. 361–384. Available at <http://www.rintonpress.com/xxqic15/qic-15-56/0361-0384.pdf>. arXiv:1406.4920.

- [46] John Preskill (2018): *Quantum Computing in the NISQ era and beyond*. *Quantum* 2, p. 79, doi:10.22331/q-2018-08-06-79.
- [47] Jonathan Romero, Ryan Babbush, Jarrod R McClean, Cornelius Hempel, Peter J Love & Alán Aspuru-Guzik (2018): *Strategies for quantum computing molecular energies using the unitary coupled cluster ansatz*. *Quantum Science and Technology* 4(1), p. 014008, doi:10.1088/2058-9565/aad3e4. Available at <https://iopscience.iop.org/article/10.1088/2058-9565/aad3e4>.
- [48] Seyon Sivarajah, Silas Dilkes, Alexander Cowtan, Will Simmons, Alec Edgington & Ross Duncan (2020): *t|ket>: A Retargetable Compiler for NISQ Devices*. *Quantum Science and Technology*, doi:10.1088/2058-9565/ab8e92. arXiv:2003.10611.
- [49] Mark Steudtner & Stephanie Wehner (2018): *Fermion-to-qubit mappings with varying resource requirements for quantum simulation*. *New Journal of Physics* 20(6), p. 063010, doi:10.1088/1367-2630/aac54f. Available at <https://iopscience.iop.org/article/10.1088/1367-2630/aac54f>.
- [50] Péter G Szalay, Thomas Müller, Gergely Gidofalvi, Hans Lischka & Ron Shepard (2012): *Multiconfiguration self-consistent field and multireference configuration interaction methods and applications*. *Chem Rev* 112(1), pp. 108–181, doi:10.1021/cr200137a.
- [51] Dave Wecker, Matthew B. Hastings & Matthias Troyer (2015): *Progress towards practical quantum variational algorithms*. *Phys. Rev. A* 92, p. 042303, doi:10.1103/PhysRevA.92.042303.
- [52] K. Wright, K. M. Beck, S. Debnath, J. M. Amini, Y. Nam, N. Grzesiak, J. S. Chen, N. C. Panti, M. Chmielewski, C. Collins, K. M. Hudek, J. Mizrahi, J. D. Wong-Campos, S. Allen, J. Apisdorf, P. Solomon, M. Williams, A. M. Ducore, A. Blinov, S. M. Kreikemeier, V. Chaplin, M. Keesan, C. Monroe & J. Kim (2019): *Benchmarking an 11-qubit quantum computer*. *Nature Communications* 10(1), p. 5464, doi:10.1038/s41467-019-13534-2.
- [53] Bujiao Wu, Xiaoyu He, Shuai Yang, Lifu Shou, Guojing Tian, Jialin Zhang & Xiaoming Sun (2019): *Optimization of CNOT circuits on topological superconducting processors*. *arXiv.org*. 1910.14478.
- [54] Rongxin Xia & Sabre Kais (2020): *Coupled cluster singles and doubles variational quantum eigensolver ansatz for electronic structure calculations*. *arXiv.org*. 2005.08451.
- [55] Andrew Zhao, Andrew Tranter, William M. Kirby, Shu Fay Ung, Akimasa Miyake & Peter Love (2020): *Measurement reduction in variational quantum algorithms*. *Phys. Rev. A* 101(6), p. 062322, doi:10.1103/PhysRevA.101.062322. arXiv:1908.08067.
- [56] Alwin Zulehner, Alexandru Paler & Robert Wille (2019): *An Efficient Methodology for Mapping Quantum Circuits to the IBM QX Architectures*. In: *DAC '19: Proceedings of the 56th Annual Design Automation Conference 2019*, Association for Computing Machinery, doi:10.1145/3316781.3317859. arXiv:1712.04722.

A Trotter Error

Given the Lie-Trotter formula for our parameterised operator:

$$U_\rho(\vec{t}) = \left(\prod_j e^{\frac{t_j}{\rho}(\tau_j - \tau_j^\dagger)}\right)^\rho \tag{11}$$

The Trotter error bound for this expansion, δ_ρ , is described in Low et al. [39]:


$$\|\delta_\rho\| = \mathcal{O}\left(\left(\frac{1}{\rho} \sum_j \|t_j(\tau_j - \tau_j^\dagger)\|\right)^2\right) \tag{12}$$

Given a well-chosen reference state with a large overlap with the exact wavefunction, low amplitudes will parameterise the ansatz, i.e. $\forall j, t_j \ll 1$ [47]. Therefore, the $\|t_j(\tau_j - \tau_j^\dagger)\|$ term will be small, particularly when compared to current two-qubit gate errors of the order of 0.1% or greater [52, 8].

Conveniently, Gui et al. give evidence to suggest that term grouping by mutual commutation such as we describe in this paper minimises Trotter error compared to other methods [25]. For a rigorous theory of Trotter error, see Childs et al. [13].

B Phase Gadgets

Theorem B.1. *We have the following laws for decomposition, commutation, and fusion of phase gadgets [17].*



Phase gadgets are invariant under qubit permutation. For an n -qubit phase gadget, this gives a choice of $C_{n-1}n!$ different CX arrangements, where C_n is the n -th Catalan number. Figure 10 shows example arrangements.






Figure 10: Comparing worst-case and best-case patterns for constructing phase gadgets with respect to CX depth. The left shows a CX ladder, with a linear CX depth, and the right shows the balanced-tree form, with a logarithmic CX depth.

C Clifford Gates and Pauli Gadgets

Theorem C.1. *We have the following laws for commuting single-qubit Clifford gates through Pauli gadgets [17].*



Theorem C.2. *We have the following laws for commuting CX gates through Pauli gadgets [17].*



D Proof of Corollary 5.4

All commuting sets of m Pauli gadgets over n qubits are diagonalisable using Theorem 5.3 if all sets of m pairs of Paulis are either compatible with Theorem 5.3 or already contain a diagonal qubit. We prove by enumerating over these sets of pairs that this compatibility is satisfied for the case $m = 3$, and therefore for $m < 4$. Compatibility is not satisfied for $m = 4$, and therefore any $m > 3$.

A short script to verify this can be found at https://github.com/CQCL/tket_benchmarking/blob/master/compilation_strategy/corollaries/corollary54.py.

E Proof of Corollary 5.5

Enumerating over all commuting sets of gadgets over 4 qubits and finding at least one pair of compatible qubits for each commuting set is sufficient proof, as each commuting set of gadgets over fewer than 4 qubits is just a special case of a 4-qubit set. However, each unique commuting set is defined by a Clifford circuit. There are more than 4.7×10^{10} Clifford circuits over 4 qubits, ignoring global phase [33]. As an optimisation, we instead search over all the *generators* of each commuting set of gadgets. Each commuting set over 4 qubits can be generated by taking products from a commuting set of 4 Pauli strings. It is therefore sufficient to find at least one

pair of compatible qubits for each commuting set of 4 Pauli strings.

A short script to verify this can be found at https://github.com/CQCL/tket_benchmarking/blob/master/compilation_strategy/corollaries/corollary55.py.

F Operator generation

For the Jordan-Wigner and Bravyi-Kitaev encodings, `t|ket>` `QubitPauliOperator` objects were produced using EUMEN, an under-construction software platform for quantum chemistry on quantum computers. Excitation operators τ_j are calculated from the molecules' spin orbitals, after which they are converted into `QubitPauliOperator` objects, which represent the $U(\vec{t})$ operators from Equation 6. These objects contain a python dictionary from Pauli string to symbolic expression representing t'_j . The coefficients a_j are dependent on the molecular geometry, and are unimportant to the compilation strategy.

The qubit operators for the parity encoding were obtained from Qiskit [2], and converted to `t|ket>` native `QubitPauliOperator` objects.

All `QubitPauliOperator` objects are serialised and stored at https://github.com/CQCL/tket_benchmarking/tree/master/compilation_strategy/operators.

For the templated lexicographical operator sequence (TLOS) method, the circuits are generated from the excitation operators, rather than a dictionary of Pauli strings to expressions, and therefore bypass the $U(\vec{t})$ operators stage of Equation 6 entirely. Rather than serialising the corresponding operators, the relevant TLOS circuits are stored in the OpenQASM format [19] at https://github.com/CQCL/tket_benchmarking/tree/master/compilation_strategy/TLOS_qasm_files.

We do not include the operations required to generate a chosen reference state $|\Phi_0\rangle$, as they are irrelevant to the strategy.