

PAPER • OPEN ACCESS

## Fault-tolerant, high-level quantum circuits: form, compilation and description

To cite this article: Alexandru Paler *et al* 2017 *Quantum Sci. Technol.* **2** 025003

View the [article online](#) for updates and enhancements.

### Related content

- [Lattice surgery translation for quantum computation](#)  
Daniel Herr, Franco Nori and Simon J Devitt
- [Lattice surgery on the Raussendorf lattice](#)  
Daniel Herr, Alexandru Paler, Simon J Devitt *et al.*
- [A software methodology for compiling quantum programs](#)  
Thomas Häner, Damian S Steiger, Krysta Svore *et al.*

### Recent citations

- [Optimal Synthesis of the Joint Unitary Evolutions](#)  
Hai-Rui Wei *et al*
- [Lattice surgery on the Raussendorf lattice](#)  
Daniel Herr *et al*
- [Optimization of lattice surgery is NP-hard](#)  
Daniel Herr *et al*



**IOP | ebooks™**

Bringing you innovative digital publishing with leading voices to create your essential collection of books in STEM research.

Start exploring the collection - download the first chapter of every title for free.

# Quantum Science and Technology



## PAPER

# Fault-tolerant, high-level quantum circuits: form, compilation and description

### OPEN ACCESS

#### RECEIVED

21 December 2016

#### REVISED

1 March 2017

#### ACCEPTED FOR PUBLICATION

15 March 2017

#### PUBLISHED

14 April 2017

Original content from this work may be used under the terms of the [Creative Commons Attribution 3.0 licence](#).

Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.



Alexandru Paler<sup>1</sup>, Iliia Polian<sup>2</sup>, Kae Nemoto<sup>3</sup> and Simon J Devitt<sup>4,5</sup>

<sup>1</sup> Johannes Kepler University Linz, Altenberger Straße, 69, A-4070 Linz, Austria

<sup>2</sup> University of Passau, Innstraße 43, D-94032, Passau, Germany

<sup>3</sup> National Institute of Informatics, 2-1-2 Hitotsubashi, Chiyoda-ku, Tokyo, Japan

<sup>4</sup> Center for Emergent Matter Science, RIKEN, Wakoshi, Saitama 315-0198, Japan

<sup>5</sup> Author to whom any correspondence should be addressed.

E-mail: [devitt@h-bar.com.au](mailto:devitt@h-bar.com.au)

**Keywords:** quantum compilation, quantum software, topological quantum computing, quantum error correction, quantum computing

## Abstract

Fault-tolerant quantum error correction is a necessity for any quantum architecture destined to tackle interesting, large-scale problems. Its theoretical formalism has been well founded for nearly two decades. However, we still do not have an appropriate compiler to produce a fault-tolerant, error-corrected description from a higher-level quantum circuit for state-of-the-art hardware models. There are many technical hurdles, including dynamic circuit constructions that occur when constructing fault-tolerant circuits with commonly used error correcting codes. We introduce a package that converts high-level quantum circuits consisting of commonly used gates into a form employing all decompositions and ancillary protocols needed for fault-tolerant error correction. We call this form the (I)initialisation, (C)NOT, (M)measurement form (ICM) and consists of an initialisation layer of qubits into one of four distinct states, a massive, deterministic array of CNOT operations and a series of time-ordered X- or Z-basis measurements. The form allows a more flexible approach towards circuit optimisation. At the same time, the package outputs a standard circuit or a canonical geometric description which is a necessity for operating current state-of-the-art hardware architectures using topological quantum codes.

The construction and compilation of quantum algorithms has long been an active area of research [1–10]. Since the first work on designing circuits from abstract quantum algorithms [11–18], research has been centred on many different aspects of compilation and optimisation. These techniques have focused on aspects such as minimising the total number of qubits [19–22], the depth of quantum circuits [23–26], constraints related to hardware [21, 27–32] and the minimisation of certain quantum gates that are resource intensive when implemented in error-corrected hardware [33–38]. However, despite the significant quantity of research in this area, much of the early work omitted details of fault-tolerant error correction protocols. These protocols are not only instrumental in implementing any quantum algorithm, beyond a handful of qubits, but place constraints on the cost metrics used when assessing optimal circuit constructions for a higher-level algorithm [39].

In recent years, theoretical research has begun to take into account the considerations of error-corrected computation. Both high-level quantum programming languages and efforts in optimisation have attempted to re-engineer circuits under more appropriate constraints. Most notably is the increased number of papers recognising that in fault-tolerant quantum computation a Clifford +  $T$  universal gate set is generally used and the resource penalty associated with the single qubit  $T = |0\rangle\langle 0| + e^{i\pi/4} |1\rangle\langle 1|$  gate requires the minimisation of both the number of such gates and the corresponding  $T$ -depth (the number of sequential applications of  $T$ -gates in a circuit) [33–38]. However, much of this work still abstracts out concrete choices of quantum error correction (QEC) coding and/or specific hardware implementations when compiling and optimising quantum circuits.

From a hardware perspective, topological codes have emerged as the preferred method for achieving a scalable design. The majority of scalable architectures proposed either use the two-dimensional (2D) surface code [40–42], or the 3-dimensional (3D) Raussendorf model [43, 44] as the underlying error correction method.

These models are preferred as they can be defined on a 2D or 3D nearest neighbour array of qubits, facilitating experimental fabrication, and have one of the highest fault-tolerant thresholds of any error correction scheme (approaching 1% under certain circuit assumptions) [45, 46]. Additionally, realisation of error-corrected algorithms is a function of how the array of qubits is measured, rather than how it is constructed (i.e. the actual choice of algorithm only dictates the physical size of the lattice needed at the hardware level, nothing more) [39, 47, 48], with error correction decoding and correction occurring concurrently with algorithmic implementation [49, 50]. In recent years, numerous results have been published demonstrating hardware accuracy sufficient for this coding model in both ion traps and superconducting systems [51–55].

These topological codes also have the property, along with more traditional concatenated CSS codes, that a universal, error-corrected gate set is built using the Clifford +  $T$  library. The inherent fault-tolerant operations allowed in these models are initialisations/measurements in the  $Z$ - and  $X$ -bases and a two-qubit CNOT gate. Hadamard ( $H = (X + Z)/\sqrt{2}$ ) and Phase ( $P = \sqrt{Z}$ ) gates can be realised in multiple ways [56–59], to complete the Clifford set, while the  $T$ -gate is realised through state injection, purification protocols and teleportations [56]. State distillation protocols, necessary for the  $T$  gate and some realisations of the  $H$ - and  $P$ -gates, dominate the resources required for a large-scale computation [39, 42].

## 1. Results

In this work, we present a decomposition and compilation package that can be the *starting point* for global optimisation of a fault-tolerant, error-corrected circuit in either the topological models or any other model of error correction where a Clifford +  $T$  library is used. This is enabled by bringing high-level quantum circuits consisting of commonly used gates into the ICM form, which employs all decompositions and ancillary protocols needed for fault-tolerant error correction.

The compiler outputs a circuit description necessary for the execution of the high-level circuit on a large-scale quantum computer using topological QECs. The source code for this compiler can be found at <https://github.com/alexandrupaler/icmconvert>.

In the [appendix](#), we provide detailed explanations about a formalism for time-optimal computation that allows dynamic quantum circuit corrections to be converted into a deterministic circuit [57], each stage of the compilation process, including standard circuit identities used [12], gate approximations based on the results from Ross and Selinger [35], and fault-tolerant ancillary protocols such as state injection and distillation [56].

## 2. Fault-tolerant quantum circuit form

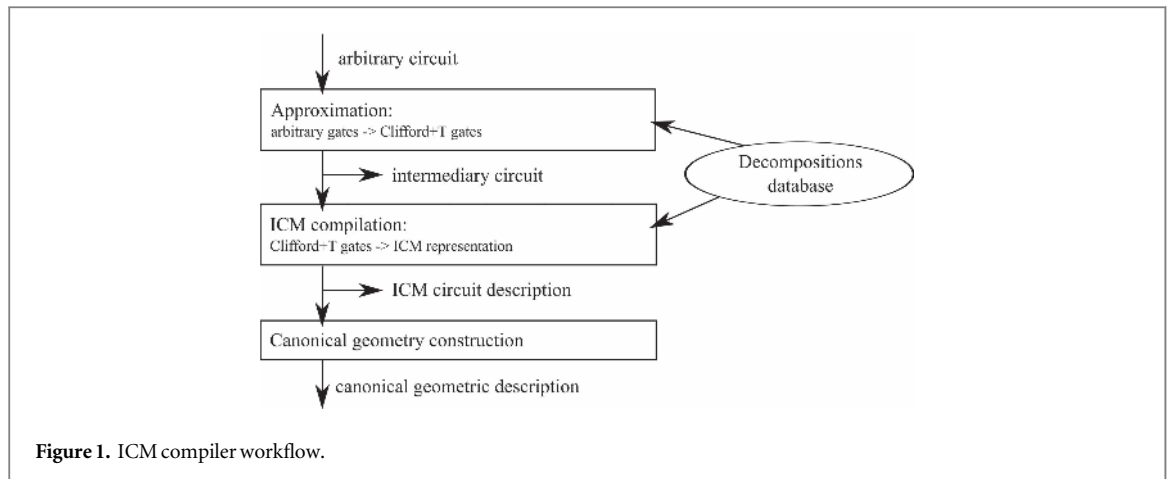
The form of fault-tolerant quantum circuits, which we call the (I)ntialisation, (C)not, (M)easurement form [61], is the result of converting a high-level quantum circuit. A high-level circuit description is the equivalent of an abstract high-level quantum algorithm converted into a standard quantum circuit representation utilising a well-defined universal gate set. The quantum algorithm, described using Quipper [8] or Liquid [9], is output into a circuit using those tools and languages, and compiled into an ICM high-level fault-tolerant quantum circuit using the herein presented framework. The ICM form is presented by employing the formalism of time-optimal quantum computation [60] and recent advances in Clifford +  $T$  gate decompositions [35].

This structure of the form consists of three layers. The first is a series of qubit initialisations in either the  $|0\rangle$  state, the  $|+\rangle$  state or the ancillary states  $|A\rangle = |0\rangle + e^{i\pi/4}|1\rangle$ ,  $|Y\rangle = |0\rangle + i|1\rangle$ . The  $|A\rangle$  and  $|Y\rangle$  states are used to implement teleported  $P = \sqrt{Z}$  and  $T = \sqrt{P}$  gates. The second layer is a massive array of purely CNOT gates, which implement both the high-level decomposed algorithm and *all* ancillary protocols for fault-tolerant computation. The final layer is a *time-ordered* series of  $X$ - and  $Z$ -basis measurements.

Unlike standard measurement based computation, the ICM representation only chooses between  $X$ - and  $Z$ -basis measurements. This is an important distinction, because non-Clifford basis measurements require an extensive ancillary overhead due to fault-tolerant error correction requirements.

ICM allows for a deterministic array of qubit initialisations and CNOT gates to describe *any* high-level quantum algorithm, incorporating all necessary fault-tolerant protocols. The standard implementation of Clifford +  $T$  circuits involves active  $P$ -gate corrections that must be adaptively patched into the quantum circuit as it is executed. A correction is required 50% of the time for each  $T$  gate.

Optimising circuits beyond individual  $T$  gates is problematic as we either need to reserve resources for *possible* corrections or attempt to optimise the exponential number of possibilities for circuits containing many  $T$ -gates. The time-optimal approach to Clifford +  $T$  circuits folds this indeterminism into whether certain qubits are measured in the  $X$ - or  $Z$ -basis. The approach provides a deterministic initialisation and CNOT structure for an arbitrary large circuit.



Hence from this formalism, it is now possible to optimise circuits starting from any point, from individual  $T$ -gates with required ancillary protocols up to the entire high-level circuit.

### 3. Compilation

The ICM form of arbitrary high-level quantum circuits is obtained by using an ICM compiler. This takes an arbitrary quantum circuit as input. It outputs a standard circuit description and a canonical geometric structure, also called geometric description, that represents its implementation for topological QEC codes. Figure 1 illustrates the compiler workflow. The following will detail the steps utilised in the compiler to allow for a fault-tolerant compatible description (quantum circuit or geometric) of the higher-level circuit.

It should be noted that this compiler is acting at the *encoded* level, i.e. each qubit in the circuit description is encoded with an appropriate QEC. The physical details of how individual qubits are manipulated are assumed to be compatible with the QEC implementation (i.e. physical gate can be switched on or off and have error rates low enough to satisfy the fault-tolerant threshold of the QEC code).

The *arbitrary circuit* input to this workflow is a quantum algorithm that is already decomposed into a standard circuit gate set consisting of gates such as Toffoli, arbitrary single qubit rotations and CNOT gates. This circuit may be optimised for qubits and/or computational time before further decompositions. This work does not address this level of circuit-level optimisation. Once a circuit has been specified (under whatever constraints), it is input into this workflow which will apply further decompositions to ultimately reach a fault-tolerant, error-corrected compatible representation.

The input circuit is specified as a list of non-ICM gates (the arbitrary circuit in figure 1) and the ICM form is algorithmically constructed by replacing all non-ICM gates with their corresponding fault-tolerant ICM implementations. Arbitrary non-ICM circuits include gates from a previously chosen universal gate set. There are two possible situations: either the universal gate set is finite (e.g. Clifford +  $T$  decompositions of a Toffoli gate) or infinite (Arbitrary single qubit rotations are *approximated* up to a residual error rate).

As a consequence, compiling an ICM circuit defined over an infinite gate set is a two-step process. First, our software uses in the *approximation step* (the first layer in figure 1) as detailed in [35] and outputs an intermediate circuit composed of gates from the finite (approximate) universal gate set  $Q_1 = \{H, P, T\}$ . Second, during the *compilation step*, the intermediate circuit is further decomposed into the ICM representation (second layer in figure 1). For example, after decomposing an arbitrary  $Z$ -rotational gate into a gate sequence from  $Q_1$ , the  $H, P, T \in Q_1$  gates are compiled into ICM.

For the compilation step, we chose to use the  $\{H, T, T^\dagger, P, P^\dagger, CNOT\}$  finite gate set (Clifford +  $T$ , where  $T^\dagger$  and  $P^\dagger$  are realised by re-interpreting measurement results of an ICM circuit), but treat the decomposition of the Toffoli, CPHASE, Controlled- $V$ , Controlled- $V^\dagger$  gates as belonging to the approximation step. Although these gates have exact decompositions, their optimality in terms of  $T$  count is still researched (e.g. the Toffoli gate [62, 63]). The controlled- $V$  gate, where  $V = \sqrt{X}$  is a finite controlled rotation as it does not require approximating arbitrary rotations and is heavily utilised in reversible circuit construction. Arbitrary controlled rotations can also be converted into ICM form by using the identities shown in the appendix and then performing an approximation to arbitrary single qubit rotations using the algorithm of Ross and Selinger [35].

Figure 1 mentions a decomposition database which includes known gate decompositions. The database stores the corresponding circuit representing the matching gate to be decomposed. For example, the Toffoli gate decomposition is loaded from the database during the approximation step. The  $T$  gate ICM decomposition is

loaded during the compilation step. The current version of our compiler uses a text based human readable database file format.

The resulting ICM circuit description includes a qubit initialisation list, a CNOT gate list and a measurement list (see output of second layer in figure 1). In general, a circuit's qubit initialisations and measurements are configurable because practical circuits are designed to support multiple input-output transformations. Exceptions exist for qubits used in specific protocols (e.g. distillation, teleportation) or ancillary workbench qubits (e.g. in adders). For example, as shown in the code segment below, in the circuit description of a  $T$  gate (see figure 8(f) for the teleported  $T$  gate in the appendix), the first line specifies the initialisation basis of the second qubit, the second line specifies a CNOT controlled by the second qubit and targeting the first qubit. The third line indicates that the first qubit is measured in the  $Z$  basis. The initialisation of the first qubit is not specified, because it is an input and thus configurable. At the same time, the measurement of the second qubit is not specified. Note: for a circuit that represents an entire algorithm, there are no input or output qubits. Every qubit is initialised and eventually measured. Certain techniques in topological optimisation, specifically bridging [64], has significant power when we do not have to consider inputs/outputs which restrict which qubits can be bridged and which cannot.

---

```
init 2 A
cnot 2 1
measure 1 Z
```

---

For input circuits using finite sets, the complexity of the ICM transformation algorithm is linear in the number of gates in the input circuit. However, because gates from an infinite gate set have to first be approximated, the complexity of this procedure is a function of the chosen approximation method. In the specific case of this software, the number of gates needed scales logarithmically with the approximation error [35]. The overhead associated with the ICM representation depends on the ICM compilation step, because the optimality of approximating arbitrary quantum gates is a function of the resulting  $T$  gate count, which in turn is algorithm specific. For example, Ross and Selinger's algorithm achieves a  $T$  count bounded by  $4 \log_2(1/\varepsilon) + \mathcal{O}(\log(\log(1/\varepsilon)))$  for an approximation precision  $\varepsilon$  [35].

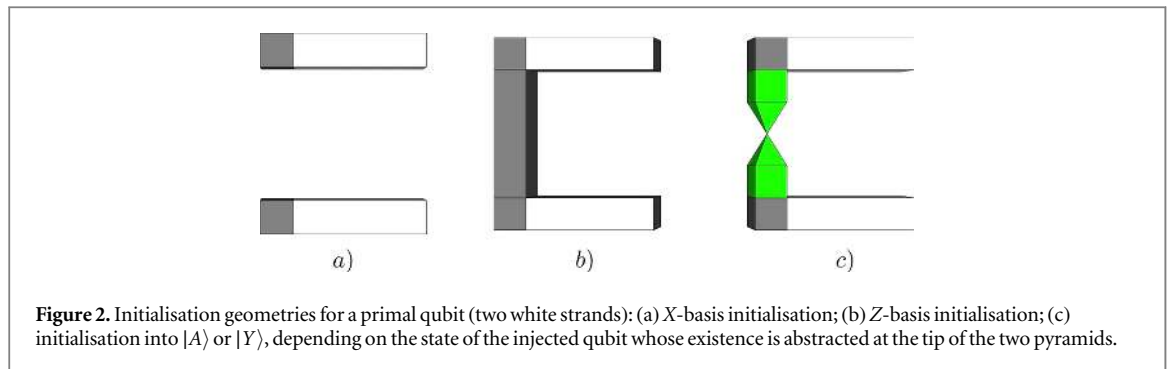
For a  $q$ -qubit non-ICM circuit consisting of  $n$  gates chosen from the finite universal gate set, the resulting ICM circuit description (not including distillation protocols) will require  $\mathcal{O}(q)$  qubits and  $\mathcal{O}(n)$  gates, because the fault-tolerant teleportation-based gate constructions introduce a constant number of ancillary qubits (maximum 5 for the  $T$  gate) and of gates (maximum 6 for the  $T$  gate). Including distillation protocols into the ICM description introduces a polynomially bounded overhead of additional resources (qubits and gates).

#### 4. Geometric description


Topologically error-corrected computations have a visual representation. The canonical geometric description is a 3D bijective mapping of the circuit description. Bijectivity is the property of each circuit description element to have a single corresponding structure in the geometric description, and vice versa. The geometric description can be optimised [48, 59, 64, 65], but its canonicity and bijectivity would be difficult to regain. Similarly to the circuit description, the geometric description consists of three geometric regions: initialisation, CNOT and measurement. The bijective relation is recognised also in terms of the temporal axis: because the geometry occupies a cuboid region of the 3D space, the inputs are on the left of the cuboid, the outputs on the right, and CNOTs are ordered inside from left to right. This is exactly the quantum circuit formalism convention.

From a geometrical point of view, qubits are defined as pairs of strands (marked white in figure 2). There are two possible types of strands: primal and dual. These two types are a function of how the surface and Raussendorf codes work and are necessary to achieve topological braided logic [42, 43]. Consequently, two types of qubits can be defined: primal and dual. For example, a primal qubit is determined by a pair of primal strands. Due to the manner how strands are constructed in the surface code, they are also commonly called defects [42]. The difference between primal and dual defects corresponds to regions of the physical quantum hardware that is 'switched off' in order to create them. Restricting the discussion to the surface code [42], a primal defect is created when plaquette stabilisers are switched off and dual defects when vertex stabilisers are switched off.

The CNOT gate is implemented, both in the surface code (2D) and Raussendorf code (3D), by braiding a primal strand (corresponding to the target qubit) around a dual strand (corresponding to the control qubit). This represents a primal-dual CNOT gate implementation, and can be extended to a primal-primal CNOT by using the circuit identity from figure 3(a). The qubits labelled with  $|c_i\rangle$ ,  $|c_o\rangle$ ,  $|t_i\rangle$  and  $|t_o\rangle$  will be mapped to primals, and the ancillary qubit (initialised into  $|+\rangle$  and measured in the  $Z$  basis) is going to be dual in the geometric description. The ancillary is the control for the three CNOTs, while the other qubits are targets. The



**Figure 2.** Initialisation geometries for a primal qubit (two white strands): (a)  $X$ -basis initialisation; (b)  $Z$ -basis initialisation; (c) initialisation into  $|A\rangle$  or  $|Y\rangle$ , depending on the state of the injected qubit whose existence is abstracted at the tip of the two pyramids.



**Figure 3.** A primal-primal CNOT: (a) The circuit identity; (b) Primal strands are red and the duals are blue. The control qubit is represented in the upper region of the geometry, and the target qubit in the lower. By comparing the geometry with figure 3(a), the dual strand is identified as the ancillary initialised into  $|+\rangle$  and measured in the  $X$  basis. Each of the three braids between the blue dual and the red primal strands has a corresponding CNOT.

resulting geometric representation of the dual qubit (figure 3(b)) hints at the vertical representation of the CNOT gate in the circuit formalism.


Strands are composed of multiple segments parallel to either the  $x$ ,  $y$  or  $z$  axis [43, 44]. Three-dimensional coordinates are sufficient to describe computations protected using the 2D or 3D codes. The properties of the coordinates are dictated by the structure of the surface codes: primal segment end-point coordinates consist entirely of odd integer values, and dual segment end points of even integer values.

The generated circuit geometric description references qubit initialisations, measurements and the structure of the strands (list of segments). The braids do not have to be explicitly described because, using geometric algorithms, it is possible to efficiently infer the braids (CNOTs) existing in a geometric description consisting of segments specified using their end-point coordinates [66].

Geometrically described qubit initialisations are illustrated in figure 2. The measurement geometries are similar, but mirrored along the time axis (the  $x$ -axis for each figure). The interpretation of the geometries is interchanged between qubit types. For example, initialising a primal qubit into the  $X$  basis uses the same geometry as dual qubit  $Z$  initialisation. The primal  $X$  initialisation (figure 2(a)) requires, from a geometrical perspective, to keep the strands separated, whereas a  $Z$  initialisation requires connecting the strands into an U-shaped structure (figure 2(b)). Measuring a primal qubit into the previous two bases uses the same geometric structures, but vertically flipped due to the time axis assumed in the canonical representation. ICM circuit qubits can also be initialised into the  $|A\rangle$  or  $|Y\rangle$  states. The initialisation geometries are similar to primal  $Z$ -basis initialisation. The major difference is that instead of a single connecting strand, there is a point (injection point) from which two disjoint strands are defined (figure 2(c)). Between a qubit's strand end points there could be either no connecting strand, a connecting strand, or two disjoint strands intersecting at a point. The previous three options are valid for qubit initialisation and for measurements only the first options are valid in the context of this work.

Similar to the quantum circuit description, the geometric description is configurable at circuit inputs and outputs, too. We achieved this by introducing a configuration point at the geometric middle of an imaginary segment connecting the qubit's strands end points (figure 4(a)). The point corresponds to a possible injection point (compare figures 4(a) and 2(c)). Configuring an ICM circuit to perform an  $X$  measurement on a specific primal qubit is equivalent to deleting the corresponding point and the incident segments (figure 4(b)), while  $Z$  measurement requires joining the incident segments and deleting the point (figure 4(c)). The same scenarios can be applied for  $X$  or  $Z$  initialisation, but the point and its incident segments are kept into the geometry for ancillary state initialisation (figure 4(d)).

The generated geometric description includes a list of 3D end-point coordinates and a list of segments existing between the end points. Furthermore, the initialisation and measurement points are marked accordingly:  $|A\rangle$  initialisation points,  $|Y\rangle$  initialisation points and ICM circuit inputs and outputs (configurable



**Figure 4.** The qubits represented by a canonical geometric description can be configured, similar to a circuit description, to be initialised or measured into multiple bases. The procedure requires a configuration point (red) and two segments (green) that connect to the end points (grey) of the qubit segments (black).

**Table 1.** The geometric description of the primal-primal CNOT from figure 3. Due to its length, the description is split into five columns.

|               |        |              |              |       |
|---------------|--------|--------------|--------------|-------|
| 4             | 11, 13 | 1, 0, 0, 0   | 14, 2, 12, 0 | 3, i  |
| 23            | 12, 13 | 2, 0, 0, 2   | 15, 2, 12, 2 | 10, i |
| 23            | 14, 16 | 3, 0, 0, 1   | 16, 2, 12, 1 | 13, o |
| 3, 10, 13, 16 | 15, 16 | 4, 0, 6, 0   | 17, -1, 9, 1 | 16, o |
| 1, 3          | 11, 14 | 5, 0, 6, 2   | 18, -1, 5, 1 |       |
| 2, 3          | 12, 15 | 6, 0, 8, 0   | 19, 1, 5, 1  |       |
| 4, 5          | 17, 18 | 7, 0, 8, 2   | 20, 1, 5, -1 |       |
| 1, 4          | 18, 19 | 8, 0, 12, 0  | 21, 3, 5, -1 |       |
| 2, 5          | 19, 2  | 9, 0, 12, 2  | 22, 3, 5, 1  |       |
| 6, 7          | 20, 21 | 10, 0, 12, 1 | 23, 3, 9, 1  |       |
| 8, 1          | 21, 22 | 11, 2, 0, 0  |              |       |
| 9, 1          | 22, 23 | 12, 2, 0, 2  |              |       |
| 6, 8          | 17, 23 | 13, 2, 0, 1  |              |       |
| 7, 9          |        |              |              |       |

geometry). Table 1 contains the description, split into five columns, of the geometry from figure 3. The first value (4) indicates the number of configuration points. The following two lines indicate the number of 3D coordinates (23) and segments (23). The coordinates, enumerated in the second and third columns, are of the form  $id, x, y, z$ , where  $id$  is a running index. In the first column, the line 3, 10, 13, 16 lists the coordinate indices representing the four initialisation and measurement points.

The remaining entries in the first column and the entire second column are index pairs indicating end points of geometric segments. The coordinates of primal configuration points contain an odd integer value, due to their definition at the middle of primal segments. For example, point 3 is the middle of segment 1, 2, the sum of segments 1, 3 and 2, 3.

Finally, the fourth column specifies the configuration point types. The third (control) and tenth (target) points are inputs, while the thirteenth (control) and sixteenth (target) are corresponding to outputs (compare with figures 3(a) and (b)).

## 5. Example ICM circuits

Diagrammatically, even comparatively small quantum circuits (such as a single Toffoli gate) are extremely large, especially when state distillation protocols for  $|A\rangle$  and  $|Y\rangle$  states are introduced. The previously presented compiler can operate on arbitrary sized circuits and outputs to multiple formats (e.g. the geometric description).

Due to the size restrictions, we illustrate two specific examples: a  $T$  and a controlled- $V$  gate. Figure 5 illustrates the canonical geometric form for a topological circuit derived from the deterministic  $T$  gate described in appendix, which is shown in ICM form). In this circuit, input 3 and output 60 represent the actual input/ outputs for the circuit while the remaining numbered nodes correspond to ancilla initialisations ( $|0\rangle, |+\rangle, |A\rangle$  and  $|Y\rangle$ ) and possible  $X$ - and  $Z$ -basis measurements to enact the teleported  $T$ -gate and possible  $P$ -gate correction.

We also illustrate the series of decomposition steps to convert a controlled- $V$  gate over two qubits into ICM form *without including* distillation protocols. At the input level, a controlled- $V$  gate is a single two-qubit operation as illustrated in figure 6(a). This gate can be expressed exactly in the Clifford +  $T$  form (which may or may not have already been performed before being compiled into ICM form) (figure 6(b)). Compilation converts this circuit into the one shown in figure 6(c), where  $T$ -gates are compiled into teleportation circuits




Figure 5. Example: the canonical geometry of the  $T$  gate circuit (see the appendix).

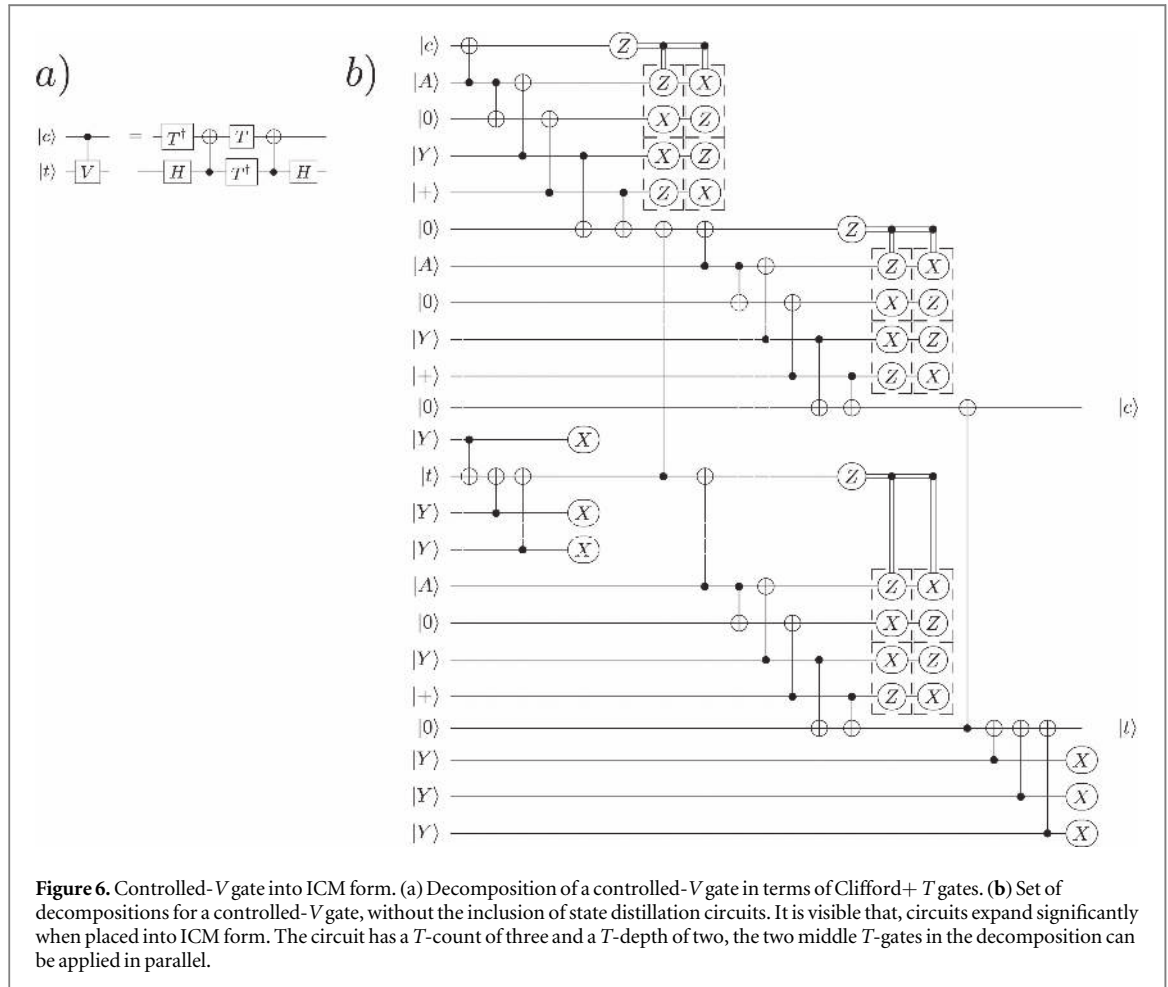


Figure 6. Controlled- $V$  gate into ICM form. (a) Decomposition of a controlled- $V$  gate in terms of Clifford+  $T$  gates. (b) Set of decompositions for a controlled- $V$  gate, without the inclusion of state distillation circuits. It is visible that, circuits expand significantly when placed into ICM form. The circuit has a  $T$ -count of three and a  $T$ -depth of two, the two middle  $T$ -gates in the decomposition can be applied in parallel.

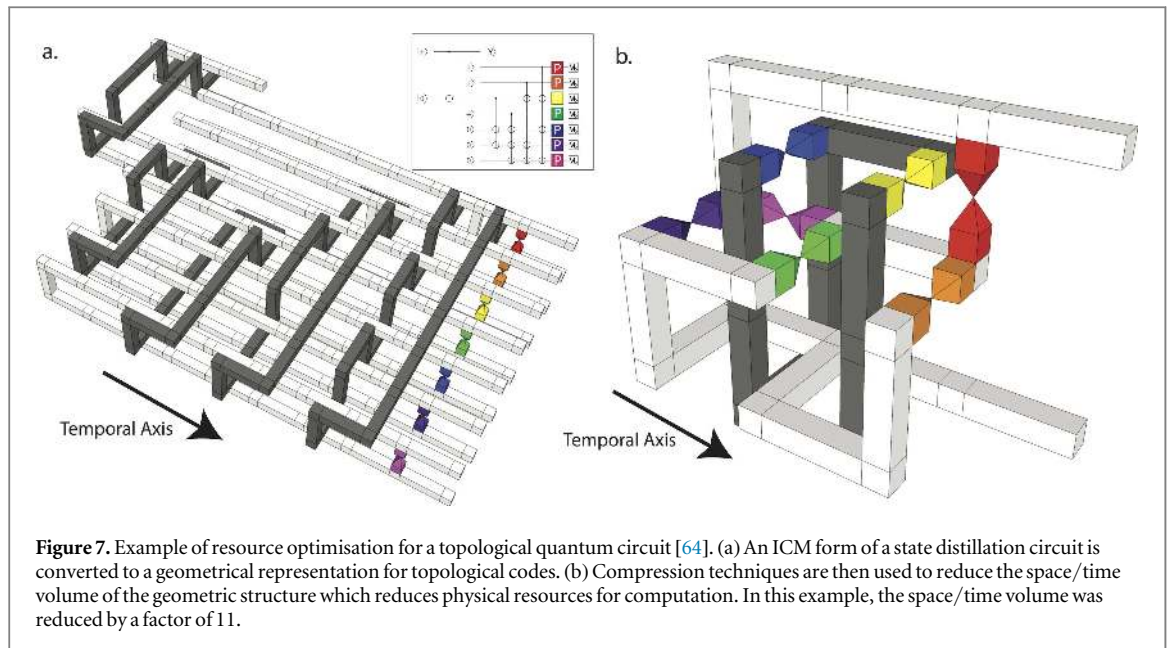
with both  $|A\rangle$  and  $|Y\rangle$  ancillary states and each Hadamard gate is decomposed into three teleportations, each requiring a  $|Y\rangle$  ancilla.

The pair of dotted measurement boxes represents bases choices for qubit measurement that are predicated on the logical measurement results from previous teleportations (the classical control dependence is indicated). Temporal ordering is indicated by staggering each  $T$ -gate subcircuits measurements. This staggering is required in case classical  $X$ -corrections are required which can convert subsequent  $T$ -gates to  $T^\dagger$ -gates and vice versa. This ensures that the  $T$ -depth of the ICM circuit is identical to the higher-level Clifford +  $T$  description.

### 6. Discussion

The ICM form allows for an arbitrary high-level circuit, obtained from high-level quantum languages like e.g. Quipper, to be decomposed into a time-optimal, deterministic form. Deterministic form implies that the array of CNOTs are fixed and the only probabilistic elements in the circuit are the  $X/Z$  basis choices for measurements that occur through the decomposition of  $T$ -gates. This form is, at least at the circuit description level,





deliberately wasteful. We intentionally introduce a large number of ancillary qubits in order to have a deterministic array of qubit initialisations and CNOTs which can then be converted to a geometric structure (for topological codes) or to another output format if different error correction codes are used. This provides us with a starting point for further optimisation at any level. Individual high-level gates (Toffoli gates or arbitrary logical rotations) could be optimised or the entire high-level quantum circuit, if desired.

## 7. Optimal geometrical descriptions

Restricting the discussion to topological quantum codes, namely the surface code and the Raussendorf code, it has been shown that various topological techniques can be used to significantly reduce the space/time volume of a quantum circuit [48, 59, 64, 65, 72]. Figure 7 illustrates what has been derived previously for a distillation circuit for  $|Y\rangle$  states. In this example, the space/time volume of a topological circuit was reduced by a factor of 11 from the original, canonical, ICM circuit. This space/time optimisation relates directly to the number of physical qubits and computational time required to implement a circuit and is a purely classical problem [39]. Compression of topological circuits result in structures that bare very little resemblance to the original circuit specification and consists of topological structures that no longer directly relate to individual qubits and CNOT gates in the original circuit [48]. Therefore, the fan-out of ancillary qubits in the ICM representation may not be detrimental to a compressed and optimised topological structure. A formal solution to the compression of topological circuits has still not been developed and therefore the degree in which resources from the ICM conversion can be recovered is still unclear. However, the ICM representation allows us to have an appropriate circuit to *start* the optimisation procedure. Without the ICM form, the initial topological circuit would be undefined (as every  $T$ -gate requires a quantum correction 50% of the time, and one cannot predict which corrections will be necessary before computation begins). Therefore circuit compression would be restricted to individual  $T$ -gates or be deliberately wasteful (by including space/time regions in the computation for corrections regardless of whether or not they are needed) or attempt to optimise as many of the  $2^{\#T\text{-gates}}$  combinations of circuits as possible. Neither of these options is desirable.

## 8. Time-ordered ICM operations

A second significant issue that needs to be considered when determining at what level quantum circuits are converted into ICM form and optimised is related to the classical information required in the circuit and when that information becomes available. As described in the previous sections, various measurement results in the ICM representation dictate which basis subsequent measurements are performed in. These basis measurements are always in either the  $X$  or  $Z$ -basis and are therefore related through a Hadamard operation that can be implemented using code deformation [57] rather than additional teleportations, without destroying the ICM structure. However, when optimising a circuit globally, extraordinary care must be taken to ensure that classical

information from these measurements is obtained as soon as possible and without disrupting the time-ordering required from the ICM representation.

Rules for topological compression do not take into account the aspects of time ordering. The example shown in figure 7 does not have time ordering (as the entire circuit is based on Clifford elements), so no further constraints need to be imposed when performing optimisation. However, for general circuits, this will not be the case and further constraints need to be placed on any optimisation technique.

As we are still developing compression techniques and optimisation software for topological quantum circuits, it remains unclear how these constraints will effect the level to which circuits can be optimised. Manual techniques were used in benchmarking Shor's algorithm [39], where circuits were optimised at the level of individual state distillation circuits and  $T$ -gates [48, 64]. The ICM form allows a circuit or topological structure to be defined for an arbitrarily large circuit, but it may be inefficient or undesirable to compress and optimise extremely large circuits. Instead it may be more preferable to optimise individual compound gates, or small subroutines that are routinely used. Until we have a more concrete formalism for optimisation, particularly for the case of topological circuits, the level to which the ICM form is used for the high-level algorithm remains an open question.

## 9. Conclusion

In this paper we introduced software to convert a high-level quantum algorithm into a fault-tolerant representations that consists of a set of qubit initialisations, a large CNOT array and measurements in the  $X$ - and  $Z$ -basis. This allows us to design a circuit compatible with all fault-tolerant protocols for a large range of error correcting codes whose only non-deterministic element is the basis choices for measurement. The ICM representation allows us to generate a canonical geometric description that fully describes the implementation of the algorithm using topological codes that forms the basis of all current hardware models. The ICM representation for quantum circuits is directly related to algorithmically specific graph states that can be produced by mapping a quantum circuit to a standard 2D cluster state and performing all Clifford basis measurements in the algorithmic specification prior to computation. Our work extends this to fully error-corrected circuits, their necessary decomposition and ancillary protocols; it also can be used to connect work in higher-level circuit synthesis and optimisation to the reality of implementing and optimising these algorithms on actual physical hardware models.

## Acknowledgements

SJD acknowledges support from the JSPS Grant-in-aid for Challenging Exploratory Research. IP acknowledges support from BAYFOR grants BayIntAn\_Uni\_Passau\_2012\_21 and BayIntAn\_Uni\_Passau\_2014\_52.

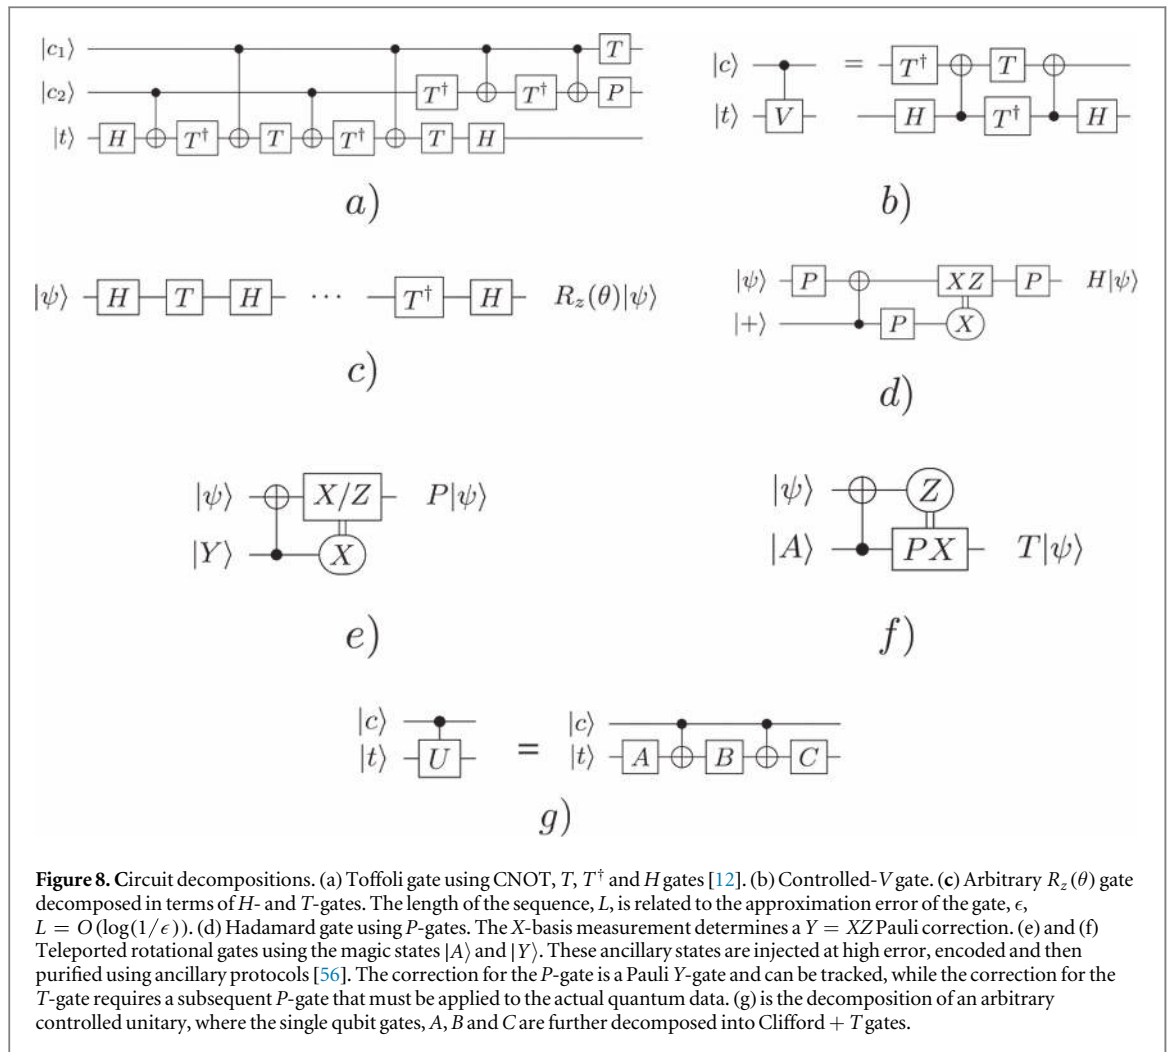
## Appendix

### A.1. Gate primitives and decompositions

We begin with a set of circuit decompositions from higher-level circuit primitives into an appropriate Clifford +  $T$  set of gates, compatible with fault-tolerant error correction. Our compiler uses the following high-level circuit primitives: CNOT, Toffoli, Controlled- $V^{(\dagger)}$  ( $V = \sqrt{X}$ ), Hadamard and other arbitrary single qubit gates. The Toffoli and Controlled- $V^{(\dagger)}$  gates are replaced with standard decompositions in terms of the Clifford +  $T$  library as shown in figures 8(a) and (b). For the Toffoli gate, we use the original decomposition containing seven  $T$ -gates [12], and the software will be periodically updated with other decompositions with reduced  $T$ -counts and  $T$ -depths [62, 63].

Hadamard gates are decomposed in terms of teleportation circuits using the ancilla state  $|Y\rangle$ , as shown in figure 8(d). For topological implementations there are methods of applying the Hadamard using code deformation [57], which allow for the ancilla state  $|Y\rangle$  to be repeatedly used. Otherwise it needs to be distilled each time such a state is required [59]. We choose not to utilise this technique because the code deformation protocol acts somewhat like a black box, and optimisation has to occur *around* these objects. While utilising a teleported Hadamard requires more resources (as  $|Y\rangle$  states can no longer be recycled), this maintains the ICM form for the entire circuit. Adding in black box elements (such as code-deformed Hadamards or other circuit substructures) may prove to be beneficial in future optimisation algorithms. However, at this stage a global ICM representation is a good starting point.

For arbitrary single qubit rotations, we utilise the algorithm of Ross and Selinger [35]. This algorithm comes close to achieving the theoretical lower bound for approximating arbitrary  $Z$ -rotations in terms of Clifford +  $T$  gates



**Figure 8.** Circuit decompositions. (a) Toffoli gate using CNOT,  $T$ ,  $T^\dagger$  and  $H$  gates [12]. (b) Controlled- $V$  gate. (c) Arbitrary  $R_z(\theta)$  gate decomposed in terms of  $H$ - and  $T$ -gates. The length of the sequence,  $L$ , is related to the approximation error of the gate,  $\epsilon$ ,  $L = O(\log(1/\epsilon))$ . (d) Hadamard gate using  $P$ -gates. The  $X$ -basis measurement determines a  $Y = XZ$  Pauli correction. (e) and (f) Teleported rotational gates using the magic states  $|A\rangle$  and  $|Y\rangle$ . These ancillary states are injected at high error, encoded and then purified using ancillary protocols [56]. The correction for the  $P$ -gate is a Pauli  $Y$ -gate and can be tracked, while the correction for the  $T$ -gate requires a subsequent  $P$ -gate that must be applied to the actual quantum data. (g) is the decomposition of an arbitrary controlled unitary, where the single qubit gates,  $A$ ,  $B$  and  $C$  are further decomposed into Clifford +  $T$  gates.

up to a specified error,  $\epsilon$  [67]. Recent techniques in probabilistic circuits have demonstrated a scaling for approximating  $Z$ -axis rotations *better* than this lower bound, but they reintroduce probabilistic corrections that this work is attempting to eliminate [68, 69]. Arbitrary axis rotations are then achieved via a standard Euler angle decomposition, with  $X$ -axis rotations approximated via  $Z$ -axis rotations and Hadamard gates,  $R_{\vec{n}}(\omega) = R_z(\alpha)HR_z(\beta)HR_z(\gamma)$ .

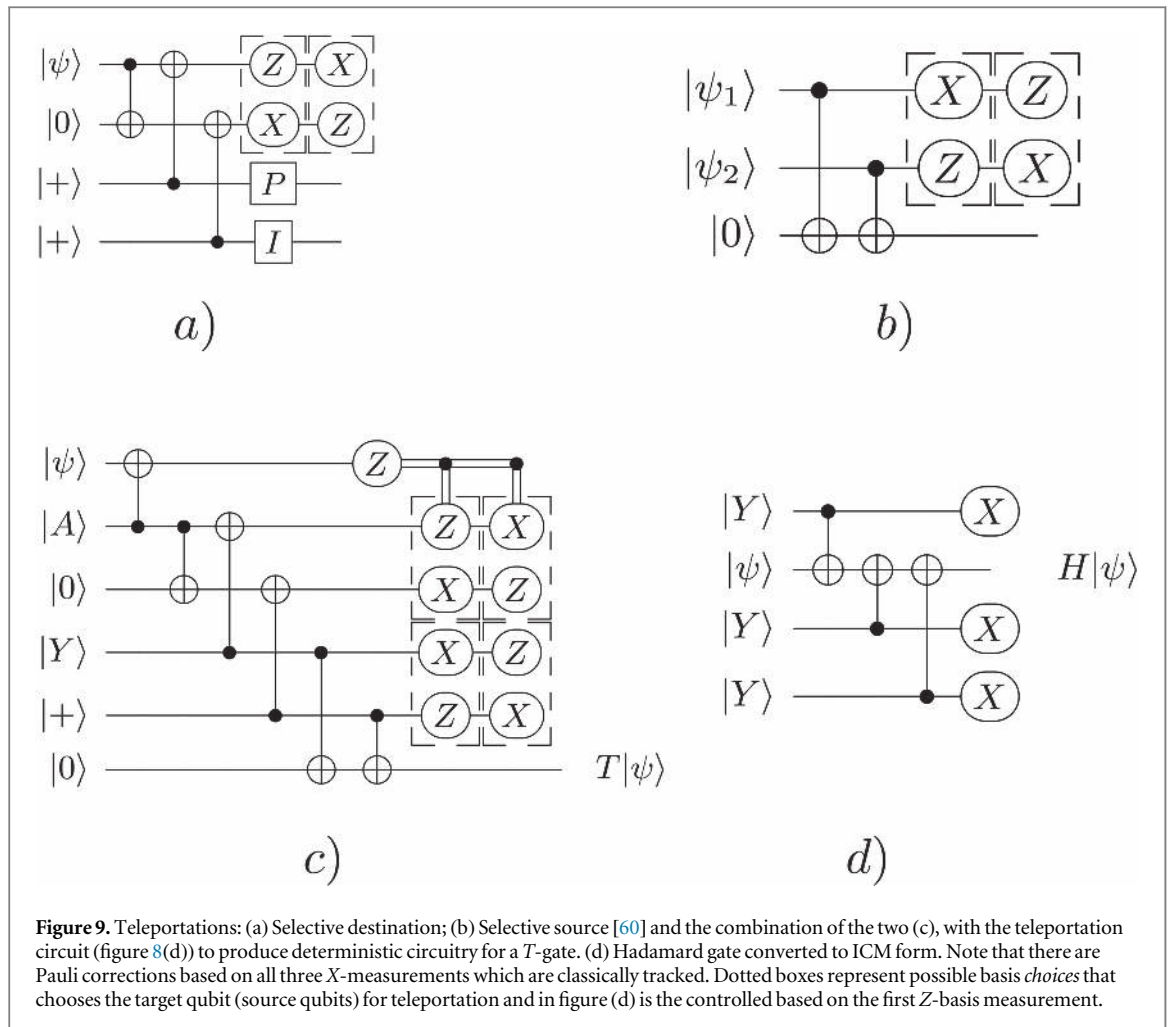
### A.2. Selective source and destination teleportation: eliminating dynamic corrections

In many fault-tolerant, error-corrected models,  $P$ - and  $T$ -gates are achieved via the teleportation circuits shown in figures 8(e) and (f). These gates are intrinsically probabilistic, and with 50% probability may apply the gates  $P^\dagger$  or  $T^\dagger$  instead. For the  $P$ -gate, this indeterminism is not a problem, as the  $Z$  gate is used as a correction ( $ZP^\dagger = P$ ); this can be tracked by updating the Pauli frame [70]. However, the correction for the  $T$ -gate is a  $P$ -gate ( $PT^\dagger = T$ ), which cannot be classically tracked and therefore must be applied using active quantum circuitry. As this correction occurs with a probability of 50%, the initial quantum circuit (consisting of many  $T$ -gates when fully decomposed) cannot be assumed and optimised prior to computation.

In a result from Fowler [60], the following trick can be used to construct deterministic circuitry regardless of this probabilistic teleportation and transfer the indeterminism of the global circuit to basis choices for qubit measurements after the deterministic circuit is executed. This is similar to measurement based quantum computation [71], where a universal, algorithmically independent, resource state is constructed and measurement outcomes inform basis choices for subsequent measurements.

The two relevant circuits are known as selective source and selective destination teleportation, illustrated in figures 9(a) and (b).

The circuit in figure 9(a) allows us to selectively teleport a state to a qubit line where a  $P$ -gate is waiting to perform a correction or not, depending on the pattern ( $Z/X$ ) or ( $X/Z$ ) of the chosen measurements. Once such a choice is made, the circuit in figure 9(b) allows us to selectively choose one of two source qubits to teleport to the output. These two circuits, combined, and commuting the  $P$ -gate correction through the control line in figure 9(a) allows us to perform a  $T$ -gate with deterministic circuitry (figure 9(c)). The original  $Z$ -measurement



for the  $T$ -gate teleportation is now used to classically control a set of  $X$ - and  $Z$ -basis measurements on the other qubits (the dotted boxes represent the two basis choices). The  $P$ -gate correction can be commuted through the control of a CNOT gate and rotates one of the  $|+\rangle$  qubits to the  $|Y\rangle = P|+\rangle$  state.

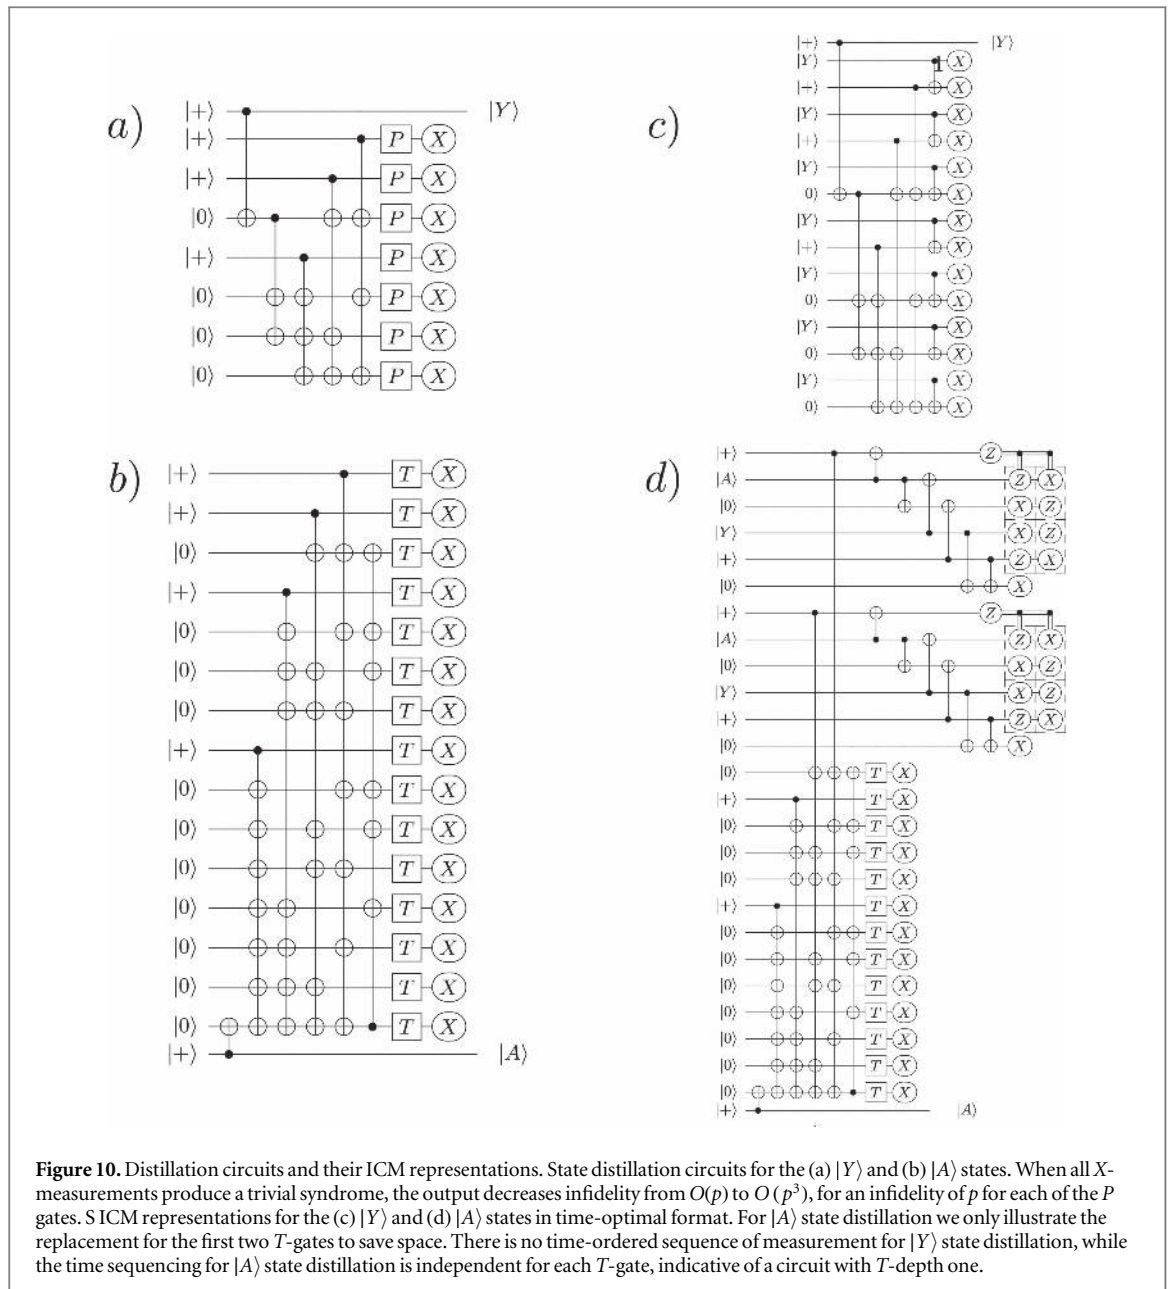
This basic element represents the building block of the ICM representation for an arbitrary, high-level, quantum circuit. The structure is such that we have a vertical line of qubit initialisations (I), followed by an array of only CNOT gates (C) and then by a time-ordered series of either  $X$ - or  $Z$ -measurements (M).

The Hadamard gate, using  $|Y\rangle$  state resources, as shown in figure 8(b), can also be expressed in terms of teleportation circuits and hence be put into ICM form [12]. In figure 9(d), there are classical corrections based on the three  $X$ -basis measurements. However, these corrections are all Pauli corrections and can be classically tracked.

### A.3. Adding in fault-tolerant Protocols

Based on the circuit decompositions shown so far, we take a high-level quantum circuit, that may already be optimised with respect to a certain set of metrics. Additionally, we assume a universal gate set that is not automatically compatible with fault-tolerant error correction and we first convert it into a Clifford +  $T$  representation, useful for both standard topological encoding and CSS concatenated encoding. Once this decomposition is done, we introduce selective source and destination circuits to place the entire circuit into an ICM form. This form of the higher-level circuit is time optimal, given the details of the decomposition to the Clifford +  $T$  level (i.e. the  $T$ -depth before introducing selective source and destination circuits determines the run-time of the computation, reducing  $T$ -depth at this level will reduce computation time [60]). We do not consider optimisation at this higher-level, as we do not focus on determining the best structure of the high-level circuit in terms of reduced  $T$ -depths.

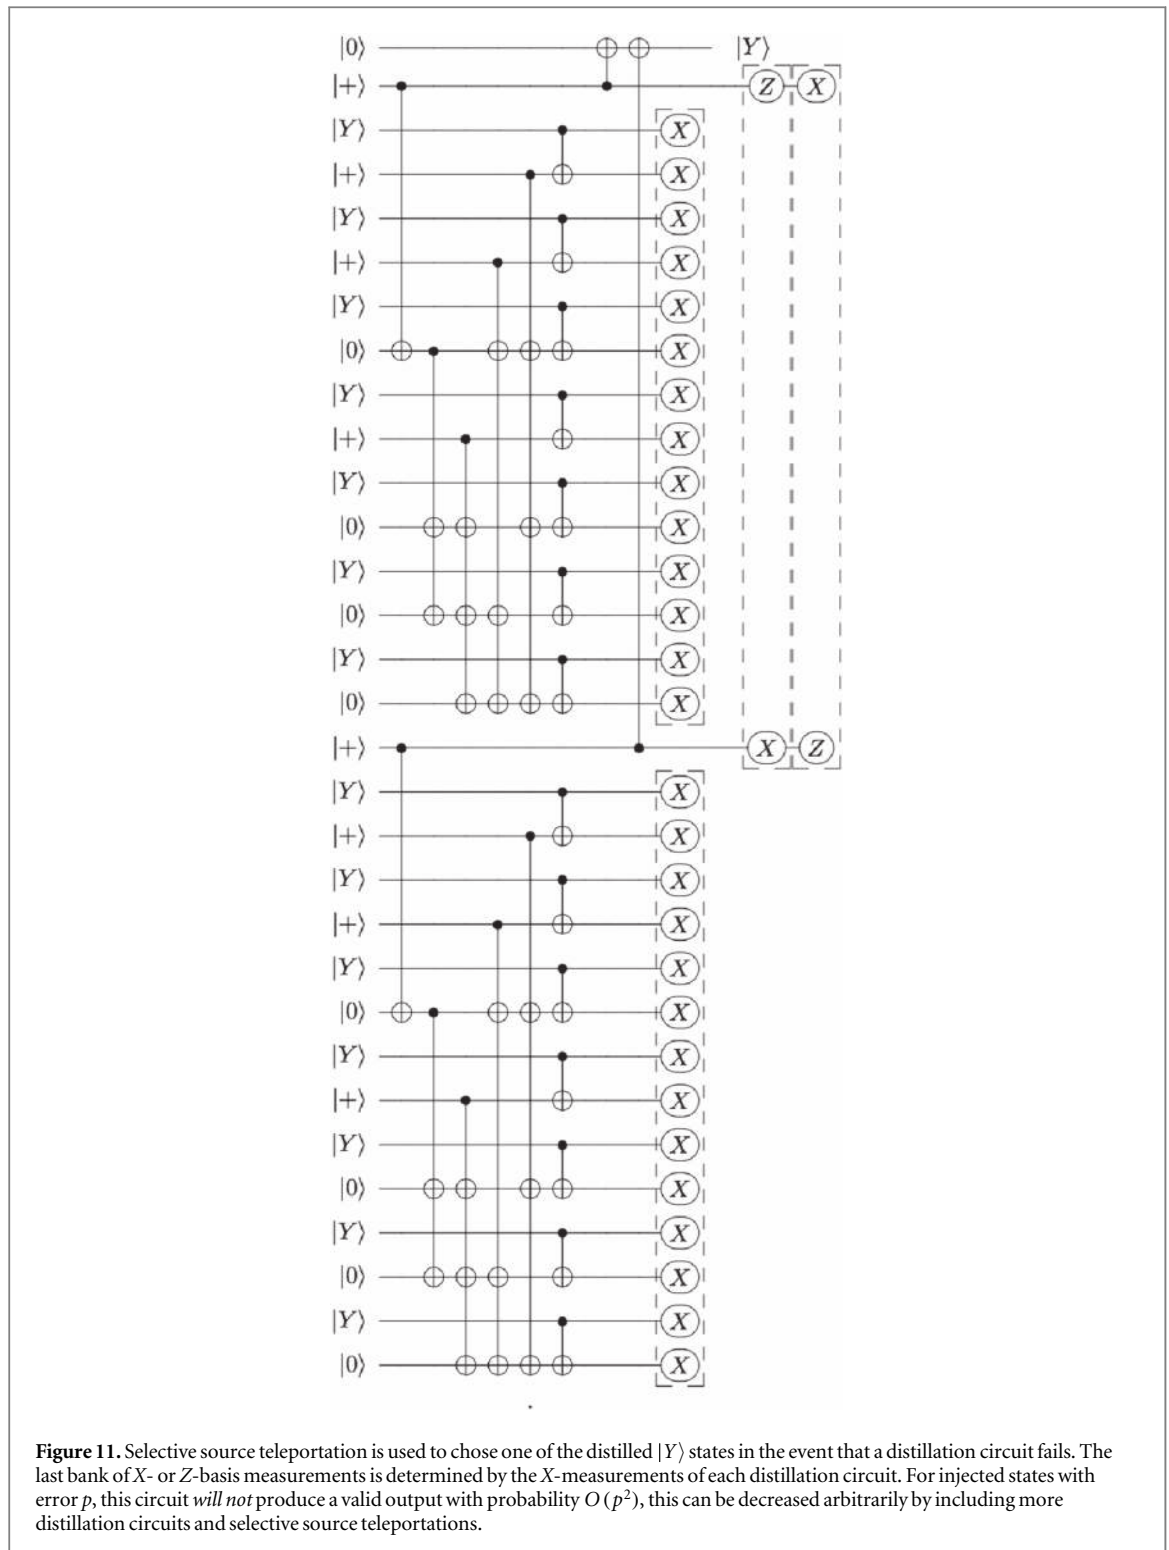
At the initialisation layer, we have a selection of four basis states,  $\{|0\rangle, |+\rangle, |Y\rangle, |A\rangle\}$ . The  $|0\rangle$  and  $|+\rangle$  states can be initialised fault-tolerantly, intrinsically, with all relevant QEC codes for large-scale hardware architectures. The states  $|A\rangle$  and  $|Y\rangle$ , in general, cannot be. This is particularly true for QEC models currently considered for large-scale hardware architectures, the topological surface and Raussendorf codes. These states



must be injected into the computer at high error and then recursively distilled. For the  $|A\rangle$  state, we utilise 15 low-fidelity states and then run a distillation circuit based on the  $[[15,1,3]]$  Reed-Muller code. This, when successful, will reduce the error from  $O(p)$  to  $O(p^3)$ ; if this is still not sufficient, the process is concatenated, and 15 of these output states are used in a further distillation circuit. For the  $|Y\rangle$  state, a similar procedure is performed, this time utilising seven resource states and the  $[[7,1,3]]$  Steane code. For states with input error rates of  $p$ , each of these circuits fail with probability  $O(p)$  [56].

Examples of these circuits are shown in figures 10(a) and (b), which are well known in the literature. In figures 10(c) and (d) we convert both of them into their own ICM representations.

These structures are in the ICM form, and can be used to replace any  $|A\rangle$  or  $|Y\rangle$  state in the initialisation stage for a higher-level circuit. A single level of state distillation is illustrated, being intended to take the error associated with the output state from an initial  $O(p)$  to  $O(p^3)$ . If this is not sufficient, then each  $|A\rangle$  or  $|Y\rangle$  state in these circuits are themselves replaced with identical structures. Thus the ICM form is maintained and takes the error rate of each  $O(p)$  injected state, to an  $O(p^9)$  distilled output. This recursion can continue as much as required (but often two levels are sufficient for large-scale algorithms) [39, 42]. State distillation only succeeds when no errors are detected by the final  $X$ -basis measurements. Hence, for states injected with an error rate  $p$ , the circuit will fail with probability  $O(p)$ . To maintain a deterministic array of CNOT gates in the ICM representation, we use the same trick with selective source and destination teleportation and duplicate the distillation circuits. This is illustrated in figure 11.



**Figure 11.** Selective source teleportation is used to choose one of the distilled  $|Y\rangle$  states in the event that a distillation circuit fails. The last bank of  $X$ - or  $Z$ -basis measurements is determined by the  $X$ -measurements of each distillation circuit. For injected states with error  $p$ , this circuit *will not* produce a valid output with probability  $O(p^2)$ , this can be decreased arbitrarily by including more distillation circuits and selective source teleportations.

To protect against failures of a distillation circuit, we utilise selective source teleportation on multiple copies. The  $X$ - and  $Z$ -basis measurements on the teleportation circuits are predicated on the  $X$ -measurements returning a trivial syndrome on at least one of the distillation circuits. As a consequence, the probability of *not* producing a distilled  $|Y\rangle$  state is reduced from  $O(p)$  to  $O(p^2)$ . This procedure can be repeated to reduce the probability of failure arbitrarily. To reduce resources, if multiple distillation circuits succeed, outputs can be routed using selective source and destination teleportations to other regions of the computer where distillation has failed.

The initial decomposed ICM form, combined with these distillation structures for  $|A\rangle$  and  $|Y\rangle$  states now allows to represent an arbitrary high-level algorithm into an ICM representation, incorporating all necessary fault-tolerant ancillary protocols.

## Compiler code

The compiler source code is released under the Microsoft Reference Source License (Ms-RSL, <http://referencesource.microsoft.com/license.html>) at <https://github.com/alexandrupaler/icmconvert>. A first application of the compiler exists at <https://github.com/alexandrupaler/tqec>, which is gradually updated with respect to ICM compilation [72], too.

The ICM compiler consists of three components whose interaction will be explained and exemplified in the following. Converting an arbitrary circuit into ICM requires the decomposition of the non-ICM quantum gates, and there are two possibilities: either a gate can be exactly decomposed into ICM (first component: the `convertft` tool), or it needs to be approximated first into gates that are easily ICM representable (second component: the `decompose` and `processraw` tools). Irrespective of the chosen method, the decompositions are stored into and retrieved from a text database accessed using the methods of the `database reader` class (third component).

The database contains three types of decompositions. The `icmdist` decomposition refers to subcircuits used to replace state injections in an ICM circuit. The `nicm` decomposition, used for example for Toffoli gates, represents the decomposition of a (non-)ICM gate into non-ICM elements (gates, measurements). Finally, `icm` stands for ICM decompositions of non-ICM gates.

From an arbitrary quantum circuit towards an ICM one, the database plays a central role. In general, each quantum gate is expressed as a unitary matrix and has a name, such that a quantum circuit is a list of quantum gate names and the qubits operated on. The `nicm` decompositions are obtained by the `decompose` tool, and undecomposed quantum gates are specified in a text file, where the complex matrix entries are written in polar form. For example, the Hadamard gate (although a direct ICM representation is known) is

---

```
1
Hadamard
0.70710678118 0
0.70710678118 0
-0.70710678118 0
```

---

The first line of the file indicates the number of specified gates (e.g. 1). Each gate specification consists of name (e.g. Hadamard) and four lines for the complex numbers of the unitary matrix (e.g. radius 0.70710678118 and angle 0). The `decompose` tool reads such a file and outputs corresponding `nicm` database entries.

After decomposing single qubit unitaries and storing their `nicm` decompositions into the database, the `processraw` tool takes the input circuit and replaces all occurrences of `nicm` gates (e.g. Toffoli) with their corresponding database entries. Thus, if in the input circuit appears the name `toffoli`, its occurrence is replaced by the decomposition stored into the database. For example, the `nicm` Toffoli gate decomposition is

---

```
=toffoli
nicm
0
WIRE WIRE WIRE CTRL WIRE WIRE WIRE CTRL WIRE CTRL WIRE CTRL TGATE
WIRE CTRL WIRE WIRE WIRE CTRL WIRE WIRE TGATE TGT TGATE TGT PGATE
HGATE TGT TGATE TGT TGATE TGT TGATE TGT TGATE HGATE WIRE WIRE WIRE
```

---

The name of the decomposed gate is specified on the first line after the '=' character. The second line (e.g. `nicm`) indicates the decomposition type, and the third line the number of required ancillae (e.g. zero). The convention is that for a controlled gate, as the Toffoli, the first qubits are the controls and the last ones are the targets. The final three lines represent the Toffoli decomposition. For each database entry the decomposition elements (e.g. WIRE, TGATE) are encoded as strings defined in `gatenumbers.h`.

The intermediate circuit obtained after executing `processraw` is finally processed by `convertft`. The tool outputs the ICM circuit description (`.circ` file), a geometrical description (`.geom` file) and a Postscript representation of the ICM circuit (`.ps` file). Furthermore, `convertft` takes as a parameter the number of distillation rounds to be included into the ICM circuit. The functionality of this tool is illustrated using the `icm` decomposition of the  $T$  gate.

---

```
=TGATE
icm
1
EMPTY AA
c 2 1
MZ EMPTY
```

---

The  $T$  gate requires a single ancilla, and because of its ICM form the decomposition consists of the same three regions. The fourth line in the listing indicates the initialisations of the qubits: the initialisation of the first qubit is left unchanged (EMPTY) and the second qubit is initialised into  $|A\rangle$  (AA). The lines starting with the character  $c$  represent CNOTs (e.g.  $c\ 2\ 1$  is a CNOT controlled by the second qubit and targeting the first qubit). The last decomposition line specifies measurement types: the first qubit is measured in  $Z$  (MZ), and the second qubit is left unmeasured (EMPTY).

After replacing in the intermediate circuit each occurrence of TGATE with the icm decomposition, `convertft` can include distillation subcircuits into the ICM output. For example, the  $|A\rangle$  state distillation circuit is specified as (\ indicates a continuation of the same text line)

---

```
=AA
icmdist
15 PLUS PLUS ZERO PLUS ZERO ZERO ZERO PLUS \
  ZERO ZERO ZERO ZERO ZERO ZERO ZERO PLUS EMPTY
c 1 6 15
c 1 3 5 7 9 11 13 15
c 2 3 6 7 10 11 14 15
c 4 5 6 7 12 13 14 15
c 8 9 10 11 12 13 14 15
c 15 3 5 6 9 10 12
MA MA MA MA MA MA MA MA MA MA MA MA MA MA MA MA MA EMPTY
```

---

Distillations are included by replacing each occurrence of the AA initialisation with the above specified circuit. The measurements in the  $|A\rangle$  basis (MA) are not corresponding to the ICM definition offered in the introduction in the main text, and in a second step each MA occurrence is replaced by a  $T$  gate applied before an  $X$ -basis measurement. Finally, the  $T$  gates are ICM decomposed.

---

```
=MA
nicm
0
TGATE MX
```

---

The above  $T$  gate decompositions did not include the selective source and destination teleportations, but these can be easily included by augmenting the TGATE database entry. In this case, the variable ancillae measurement types (effect of probabilistic corrections) are specified using the operations MXZ and MZX also defined in `gatenumbers.h`.

---

```
=TGATE
icm
5
EMPTY AA ZERO YY PLUS ZERO
c 2 1
c 2 3
c 4 2
c 5 3
c 4 6
c 5 6
MZ MZX MXZ MXZ MZX EMPTY
```

---

The circuit description presented in the results section of the main text is a reformulation of the database entry format. The initialisation and measurement lines in the circuit description are compressed to a single line in the database entries, and instead of the `cnot` command, the database uses the `c` command. This fact shows that, once a circuit is ICM transformed, it can be easily stored into the database for future transformations to use it as a sub circuit (e.g. modular adders used in modular multiplications).

## References

- [1] Knill E 1996 Conventions for quantum pseudocode *Technical Report LAUR-96-2724*, Los Alamos National Laboratory
- [2] Omer B 2000 Quantum programming in QCL *Master's Thesis, Institute of Information Systems* (Technical University of Vienna, Austria)
- [3] Bettelli S, Calarco T and Serafini L 2003 Towards an architecture for quantum programming *Eur. Phys. J. D* **25** 181–200
- [4] Selinger P 2004 Towards a quantum programming language *Math. Struct. Comput. Sci.* **14** 527–86



- [5] van Tonder A 2004 A lambda calculus for quantum computation *SIAM J. Comput.* **33** 1109–35
- [6] Altenkirch T and Grattage J 2005 A functional quantum programming language *Logic in Computer Science, 2005. 20th Annual IEEE Symposium on*, 26–29 249–58
- [7] Selinger P and Valiron B 2006 A lambda calculus for quantum computation with classical control *Math. Struct. Comput. Sci.* **16** 527–52
- [8] Green A S, Selinger P, LeFanu Lumsdaine P, Ross N J and Valiron B 2013 An introduction to quantum programming in quipper *Lecture Notes in Computer Science, Springer* **7948** 110–24
- [9] Wecker D and Svore K M 2014 Liquid: a software design architecture and domain-specific language for quantum computing arxiv:1402.4467
- [10] Devitt S J, Munro W J and Nemoto K 2011 High performance quantum computing *Progress in Informatics* **8** 1–7
- [11] Barenco A et al 1995 Elementary gates for quantum computation *Phys. Rev. A.* **52** 3457
- [12] Nielsen M A and Chuang I L 2000 *Quantum Computation and Information* 2nd edn Cambridge University Press)
- [13] Kabanets V and Cai J-Y 2000 Circuit minimization problem *Proceedings of the Thirty-second Annual ACM Symposium on Theory of Computing (ACM)* pp 73–9
- [14] Yao A C 1993 Quantum circuit complexity *Foundations of Computer Science. Proceedings., 34th Annual Symposium on (Palo Alto, CA)* 352–61
- [15] Vedral V, Barenco A and Ekert A K 1996 Quantum networks for elementary arithmetic operations *Phys. Rev. A.* **54** 147
- [16] Patel K N, Markov I L and Hayes J P 2008 Optimal synthesis of linear reversible circuits *Quant. Inf. Comp.* **8** 282–94
- [17] Van Meter R and Itoh K M 2005 Fast quantum modular exponentiation *Phys. Rev. A.* **71** 052320
- [18] Dias da Silva R, Pius R and Kashefi E 2013 Global quantum circuit optimization arxiv:1301.0351
- [19] Beauregard S 2003 Circuit for Shor’s algorithm using  $2n+3$  qubits *Quant. Inf. Comp.* **3** 175
- [20] Shende V V, Markov I L and Bullock S S 2004 Minimal universal two-qubit controlled-NOT based circuits *Phys. Rev. A.* **69** 062321
- [21] Maslov D 2007 Linear depth stabilizer and quantum fourier transformation circuits with no auxiliary qubits in finite-neighbor quantum architectures *Phys. Rev. A.* **76** 052310
- [22] Golubitsky O and Maslov D 2012 A study of optimal 4-bit reversible toffoli circuits and their synthesis *Computers, IEEE Transactions on* **61** 1341–53
- [23] Høyer P and Špalek R 2005 Quantum fan-out is powerful *Theory of Computing* **1** 81–103
- [24] Broadbent A and Kashefi E 2009 Parallelizing quantum circuits *Theor. Comput. Sci.* **410** 2489–510
- [25] Amy M, Maslov D, Mosca M and Roetteler M 2013 A meet-in-the-middle algorithm for fast synthesis of depth-optimal quantum circuits *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on* **32** 818–30
- [26] Abdessaied N, Wille R, Soeken M and Drechsler R 2013 Reducing the depth of quantum circuits using additional circuit lines *Lecture Notes in Computer Science, Springer* **7948** 221–33
- [27] Fowler A G, Devitt S J and Hollenberg L C L 2004 Implementation of Shor’s algorithm on a linear nearest neighbour qubit array *Quant. Inf. Comp.* **4** 237–51
- [28] Draper T G, Kutin S A, Rains E M and Svore K M 2006 A logarithmic-depth quantum carry-lookahead adder *Quant. Inf. Comp.* **6** 351–69
- [29] Stephens A, Fowler A G and Hollenberg L C L 2008 Universal fault-tolerant computation on bilinear nearest neighbor arrays *Quant. Inf. Comp.* **8** 330
- [30] Saeedi M, Wille R and Drechsler R 2011 Synthesis of quantum circuits for linear nearest neighbor architectures *Quant. Inf. Proc.* **10** 355–77
- [31] Saeedi M, Shadaei A and Pedram M 2013 Constant-factor optimization of quantum adders on 2d quantum architectures *Lecture Notes in Computer Science, Springer* **7948** 58–69
- [32] Wiebe N, Kapoor A and Svore K M 2015 Quantum nearest-neighbor algorithms for machine learning *Quant. Inf. Comp.* **15** 0318–58
- [33] Selinger P 2015 Efficient clifford + T approximations of single-qubit operators *Quant. Inf. Comp.* **15** 159–80
- [34] Giles B and Selinger P 2013 Exact synthesis of multiqubit Clifford + T circuits *Phys. Rev. A.* **87** 032332
- [35] Ross N J and Selinger P 2014 Optimal ancilla-free clifford + T approximation of z-rotations arXiv:1403.2975
- [36] Kliuchnikov V, Maslov D and Mosca M 2013 Asymptotically optimal approximation of single qubit unitaries by clifford and T circuits using a constant number of ancillary qubits *Phys. Rev. Lett.* **110** 190502
- [37] Welch J, Bocharov A and Svore K M 2016 Efficient approximation of diagonal unitaries over the clifford + T basis *Quant. Inf. Comp.* **16** 87–104
- [38] Gosset D, Kliuchnikov V, Mosca M and Russo V 2014 An algorithm for the T-count *Quant. Inf. Comp.* **14** 1277–301
- [39] Devitt S J, Stephens A M, Munro W J and Nemoto K 2013 Requirements for fault-tolerant factoring on an atom-optics quantum computer *Nat. Commun.* **4** 2524
- [40] Dennis E, Kitaev A, Landahl A and Preskill J 2002 Topological quantum memory *J. Math. Phys.* **43** 4452
- [41] Fowler A G, Stephens A M and Groszkowski P 2009 High threshold universal quantum computation on the surface code *Phys. Rev. A.* **80** 052312
- [42] Fowler A G, Mariantoni M, Martinis J M and Cleland A N 2012 Surface codes: towards practical large-scale quantum computation *Phys. Rev. A.* **86** 032324
- [43] Raussendorf R, Harrington J and Goyal K 2007 Topological fault-tolerance in cluster state quantum computation *New J. Phys.* **9** 199
- [44] Fowler A G and Goyal K 2009 Topological cluster state quantum computing *Quant. Inf. Comp.* **9** 721
- [45] Wang D S, Fowler A G and Hollenberg L C L 2011 Quantum computing with nearest neighbor interactions and error rates over 1% *Phys. Rev. A.* **83** 020302 (R)
- [46] Stephens A M 2014 Fault-tolerant thresholds for quantum error correction with the surface code *Phys. Rev. A.* **89** 022321
- [47] Devitt S J and Nemoto K 2012 Programming a topological quantum computer *Proceedings of the 2012 IEEE 21st Asian Test Symposium (IEEE Computer Society, Niigata, Japan)* 55–60
- [48] Fowler A G, Devitt S J and Jones C 2010 Surface code implementation of block code state distillation *Int. J. Quant. Inf.* **8** 1–27
- [49] Devitt S J, Fowler A G, Tilma T, Munro W J and Nemoto K 2013 Classical processing requirements for a topological quantum computing system *Int. J. Quantum Inform.* **08** 121
- [50] Fowler A G 2015 Minimum weight perfect matching of fault-tolerant topological quantum error correction in average  $O(1)$  parallel time *Quant. Inf. Comp.* **15** 0145–58
- [51] Devitt S J et al 2009 Architectural design for a topological cluster state quantum computer *New J. Phys.* **11** 083032
- [52] Cody Jones N et al 2012 A layered architecture for quantum computing using quantum dots *Phys. Rev. X.* **2** 031007
- [53] Yao N Y et al 2012 Scalable architecture for a room temperature solid-state quantum information processor *Nat. Commun.* **3** 800
- [54] Nemoto K et al 2014 Photonic architecture for scalable quantum information processing in NV-diamond *Phys. Rev. X.* **4** 031022

- [55] Lekitsch B *et al* 2015 Blueprint for a microwave ion trap quantum computer arxiv:[1508.00420](#)
- [56] Bravyi S and Kitaev A 2005 Universal quantum computation with ideal clifford gates and noisy ancillas *Phys. Rev. A.* **71** 022316
- [57] Fowler A G 2012 Low overhead surface code logical H *Quant. Inf. Comp.* **12** 970
- [58] Aliferis P 2007 Level reduction and the quantum threshold theorem *PhD Thesis* Caltech
- [59] Cody Jones N 2013 Logic synthesis for fault-tolerant quantum computers *PhD Thesis* Stanford University
- [60] Fowler A G 2012 Time optimal quantum computation arxiv:[1210.4626](#)
- [61] Paler A, Polian I, Nemoto K and Devitt S J 2015 A regular representation of quantum and reversible circuits *Lecture Notes in Computer Science, Springer* **9138** 139–54
- [62] Jones C 2013 Novel constructions for the fault-tolerant Toffoli gate *Phys. Rev. A.* **87** 022328
- [63] Jones C 2013 Composite Toffoli gate with two-round error detection *Phys. Rev. A.* **87** 052334
- [64] Fowler A G and Devitt S J 2012 A bridge to lower overhead quantum computation arXiv:[1209.0510](#)
- [65] Paetznick A and Fowler A G 2013 Quantum circuit optimization by topological compaction in the surface code arXiv:[1304.2807](#)
- [66] Paler A, Devitt S J, Nemoto K and Polian I 2014 Mapping of topological quantum circuits to physical hardware *Sci. Rep.* **4** 4657
- [67] Dawson C M and Nielsen M A 2006 The Solovay-Kitaev algorithm *Quant. Inf. Comp.* **6** 81
- [68] Bocharov A, Rotteler M and Svore K M 2015 Efficient synthesis of probabilistic quantum circuits with fallback *Phys. Rev. A.* **91** 052317
- [69] Bocharov A, Roetteler M and Svore K M 2015 Efficient synthesis of universal repeat-until-success circuits *Phys. Rev. Lett.* **114** 080502
- [70] Paler A, Devitt S J, Nemoto K and Polian I 2014 Software-based pauli tracking for topological quantum circuits *Proceedings of the Conference on Design, Au-tomation & Test in Europe* 1–4 European Design and Au-tomation Association
- [71] Raussendorf R and Briegel H-J 2001 A one way quantum computer *Phys. Rev. Lett.* **86** 5188
- [72] Paler A, Devitt S J and Fowler A G 2016 Synthesis of arbitrary quantum circuits to topological assembly *Sci. Rep.* **6** 3600