

TKET — A RETARGETABLE COMPILER FOR NISQ DEVICES

*by Seyon Sivarajah, Silas Dilkes, Alexander Cowtan,
Will Simmons, Alec Edgington, Ross Duncan*

Cambridge Quantum
SEYON SIVARAJAH
seyon.sivarajah@cambridgequantum.com
SILAS DILKES
silas.dilkes@cambridgequantum.com
ALEXANDER COWTAN
alexander.cowtan@cambridgequantum.com
WILL SIMMONS
will.simmons@cambridgequantum.com
ALEC EDGINGTON
alec.edgington@cambridgequantum.com

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

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

$t|ket\rangle$: A Retargetable Compiler for NISQ Devices

Seyon Sivarajah^{1*†}, Silas Dilkes^{1*}, Alexander Cowtan^{1*},
Will Simmons^{1*}, Alec Edgington¹, and Ross Duncan^{1,2}

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

* These authors contributed equally to this work.

Abstract. We present $t|ket\rangle$, a quantum software development platform produced by Cambridge Quantum Computing Ltd. The heart of $t|ket\rangle$ is a language-agnostic optimising compiler designed to generate code for a variety of NISQ devices, which has several features designed to minimise the influence of device error. The compiler has been extensively benchmarked and outperforms most competitors in terms of circuit optimisation and qubit routing.

1. Introduction

Quantum computing devices promise significant advantages for a wide variety of information processing tasks [1, 2, 3]. For some tasks, notably the simulation of condensed-matter physics, the abstract structure of the problem may be sufficiently similar to the physical structure of the device that translation from one to the other is natural and (relatively) straightforward [4]. However, for most problems, and most quantum computers, this is not the case. Quantum algorithms are often described in terms that facilitate proving correctness or deriving asymptotic complexity estimates, without reference to a specific computing device on which to execute them. The translation from a high-level description of the algorithm to a machine-specific sequence of physical operations is called *compilation*, and is essential to realising the supposed computational advantage of quantum algorithms.

In computer science, the term *compiler* was introduced by Grace Hopper in the early 1950s [5], and originally referred to a routine which “compiled” a desired program from pre-existing pieces. Today the term denotes a program that translates a human-readable programming language into the binary language of the machine that will execute it. Early compilers produced code that was grossly inefficient, compared to what an average human programmer could write; however, today’s sophisticated *optimising* compilers reliably generate code that runs more quickly and uses less memory than even the best human programmer could manage [6]. The steady improvement of compiler technology

† seyon.sivarajah@cambridgequantum.com

has, in turn, enabled programming languages to increase in power and sophistication, increasing the conceptual distance between the programmer and the executable machine language.

By comparison, almost all programming systems available for quantum computing are conceptually primitive, remaining extremely close to the basic quantum circuit model [7]. Although higher-level application-oriented toolkits are becoming available [8, 9], the programmer must usually describe the algorithm to be run in terms of basic unitary gates. On the other hand, quantum computing hardware displays great diversity. Superconducting and ion-trap-based quantum processors are now available from multiple commercial companies [10, 11, 12, 13, 14], while other technologies such as photonics are not far behind [15, 16]. Different underlying technologies have very different performance parameters and trade-offs, and even broadly similar devices may differ in what basic operations are available. Even in the context of the simple, circuit-centric, programming model, the requirement to translate an abstract circuit into something suitable for the chosen device creates the need for a compiler. Naive approaches to this translation can significantly increase the size of the circuit; therefore the other major task for a quantum compiler is *circuit optimisation*, to minimise the resources required by the program.

Circuit optimisation is especially pertinent on so-called *noisy intermediate-scale quantum* (NISQ) devices. Preskill [17] defines a NISQ device as having a memory size of 50–100 qubits, and sufficient gate fidelity to carry out around 1000 two-qubit operations with tolerable error rates. We will adopt a wider definition: a NISQ device is any quantum computer for which general-purpose quantum error correction [18] is not feasible and hardware errors are expected. Because of these ineradicable errors, mere qubit count is poor measure of the capability of NISQ devices. The longer the computation runs, the more noise builds up.

NISQ devices therefore impose strict limitations both on the number of qubits available to algorithms and on the maximum circuit depth that can be achieved. Aside from the obvious requirement to use this limited hardware budget in the most efficient manner possible, the noisiness of NISQ machines introduces further complications. Since many common textbook algorithms such as quantum phase estimation ‡ are not feasible in the available circuit depth, hybrid algorithms such as the Variational Quantum Eigensolver (VQE) [20] and the Quantum Approximate Optimisation Algorithm (QAOA) [21] have been proposed instead. While the circuit depths required by these algorithms are more favourable to NISQ devices, they are based on repeatedly executing circuits inside a classical optimisation loop, where both the rate of convergence and the accuracy of the final result can be adversely affected by device noise. In consequence, any compiler for NISQ devices should aim to maximise the overall fidelity of the computation. Minimising the number of operations helps, but other techniques may be employed [22, 23, 24, 25].

This paper describes t|ket>, a compiler system for NISQ devices that aims to achieve

‡ But see O’Brien *et al* [19].

```

 $C(\cdot) \leftarrow \text{GenerateParameterisedCircuit}()$ 
 $M \leftarrow \text{GenerateListOfPauliMeasurements}()$ 
 $\bar{\theta} \leftarrow \bar{\theta}_0$ 
 $E \leftarrow 0$ 
while  $E$  has not converged do
  for  $i$  from 1 to  $N\_ITERS$  do
    for  $m \in M$  do
       $r_{m,i} \leftarrow \text{Execute}(C(\bar{\theta}); m)$ 
    end for
  end for
   $E' \leftarrow \text{EstimatorFunction}(\bar{r})$ 
   $\bar{\theta} \leftarrow \text{ClassicalOptimiser}(\bar{\theta}, E, E')$ 
   $E \leftarrow E'$ 
end while
return  $E$ 

```

Figure 1: The typical structure of a variational quantum algorithm

these objectives. The core of t|ket> is a flexible optimising compiler which supports multiple programming frameworks, and multiple quantum devices. It is specifically designed for NISQ devices, and includes features that minimise the influence of device errors on computation. As we demonstrate in Section 9, t|ket>’s optimisation and qubit mapping routines reliably outperform other compilers. The system also includes runtime management features to facilitate the variational algorithms typical of NISQ devices.

1.1. NISQ devices and their software

Before addressing the t|ket> system, we consider a schematic variational algorithm in the context of the system architecture of an idealised NISQ device. For purposes of illustration, a toy VQE algorithm is shown in Figure 1.

The first point to note in this example is that the central “Execute($C(\bar{\theta}); m$)” subroutine is the only part that runs on the quantum device: the other subroutines and the main loop are classical. The subroutines “EstimatorFunction” and “ClassicalOptimiser” are used repeatedly inside the main loop – the characteristic of a hybrid algorithm – and their outputs are used in the next quantum execution. The first two subroutines, “GenerateParameterisedCircuit” and “GenerateListOfPauliMeasurements”, are tasks that are usually considered part of the compiler, but observe that inside the main loop a fresh quantum circuit must be built using the parameterised circuit $C(\cdot)$, the measurement m , and the new parameters $\bar{\theta}$. How does an algorithm like this map onto a realistic quantum computer system?

While it is common to talk about a “quantum computer” as a unified device, in practice it consists of multiple subsystems, each of which is a computer in its own right.

Running a quantum algorithm therefore involves a large number of software components in a mixture of runtime environments, with very different performance demands. Figure 2 displays a realistic architecture for a NISQ computer. The lowest level comprises the programmable *devices* which drive the evolution of the qubits and read out their states. An example of this kind of device is an arbitrary waveform generator, as found in many superconducting architectures. The microwave pulse sequences output by these devices are generated by simple low-level programs optimised for speed of execution. These devices, and the *real-time controller* which synchronises them, operate in a hard real-time environment where the computation takes place on the time-scale of the coherence time of the qubits. These components combine to execute a single instance of a quantum circuit, possibly with some classical control. By analogy with GPU computing, we refer to this layer as a *kernel*.

One level higher, the *scheduler* is responsible for dispatching circuits to be run and packaging the results for the higher layers. It is also likely to be heavily involved in the device calibration process. (Calibration data are an important input to the compiler.) This layer and those below may be thought of as the low-level system software of the quantum computer, and must normally be physically close to the device. In the layer above we find service-oriented middleware, principally the *task manager*, which may distribute jobs to different quantum devices or simulators, GPUs, and perhaps conventional HPC resources, to perform the various subroutines of the quantum algorithm. This layer may also allocate access to the quantum system among multiple users. Finally, at the highest level is the *user runtime*, which defines the overall algorithm and integrates the results of the subcomputations to produce the final answer.

With this picture in mind, we see that the path from a high-level program describing a quantum algorithm to its final result involves many stages of decomposition and compilation in order to run in this heterogeneous environment. In practice, some of these stages may be amalgamated or absent. In this paper we will focus on the generation of the kernel, since this is the indispensable part of the process, and can be (to some extent) decoupled both from the high-level architecture and the low-level system-specific parts.

The picture is further complicated when considering quantum computers capable of error correction. The “logical” kernel must be translated to an encoded equivalent; subroutines to perform gate synthesis must be added; and error detection correction stages must be interleaved in the main algorithm. However, every part of the NISQ process is also required in the error-corrected case, so we focus on compilation in the NISQ context.

1.2. Related work

The last few years have seen an explosion of interest in quantum programming languages, and the problems of quantum compilation have been explored at various levels of abstraction [26], from high-level algorithm design to pulse control at the machine level.

Several languages have been developed for quantum programming. Quipper [27] is a

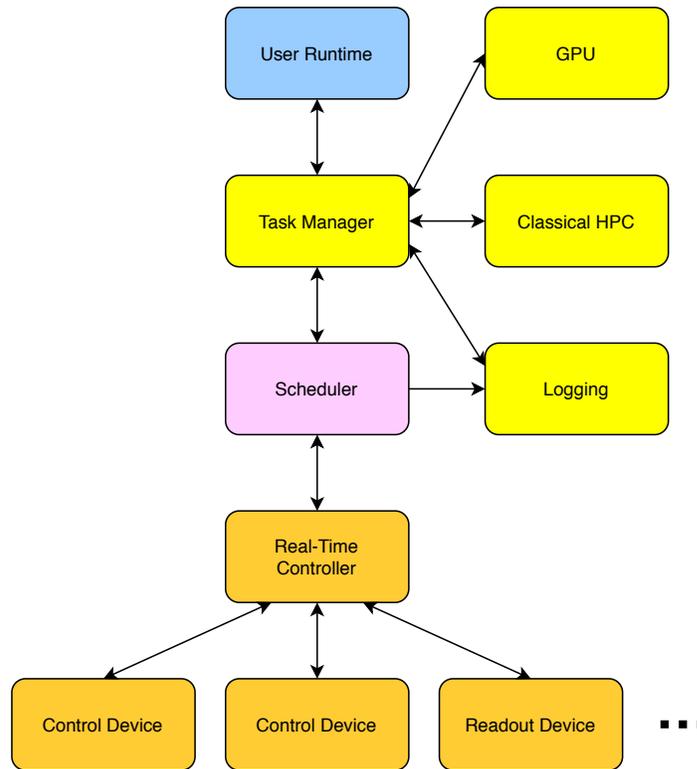


Figure 2: Idealised system architecture for a NISQ Computer

functional language for quantum circuits embedded in Haskell. The ScaffCC compiler [28], based on the LLVM framework, compiles an extension of C, and can be configured for routing to specific architectures. Q# [29] is a hybrid classical–quantum language designed to facilitate the development of programs that can be run on a simulator (and eventually on actual hardware). Strawberry Fields [30] is a Python-based quantum programming framework that is based on the “continuous variable” (CV) model of computation.

Several other compilation systems have been developed as Python modules targeting specific hardware. These include the Forest SDK/pyQuil [31] (for Rigetti backends), Qiskit [32] and ProjectQ [33] (for IBM backends), and Cirq [34] (for Google backends). Other projects have adopted a backend-agnostic approach. XACC [35] is a quantum programming framework that can target several different backends as plug-ins. TriQ [36] uses ScaffCC to compile quantum software for several different architectures in order to study their performance characteristics. Even for the full-stack systems, the compiler element (e.g. Quilc [37] from the Forest SDK or the transpiler passes in Qiskit Terra) can be invoked to compile for arbitrary devices.

A range of gate-level circuit-optimisation techniques have been explored, including the use of phase polynomials [38] and constraint programming [39]. There are also promising results for using information on the noise characteristics and fidelities of the target device to assist compilation [40, 41, 42]. Meanwhile at the level of machine control there have been efforts to optimise the implementation of variational algorithms using

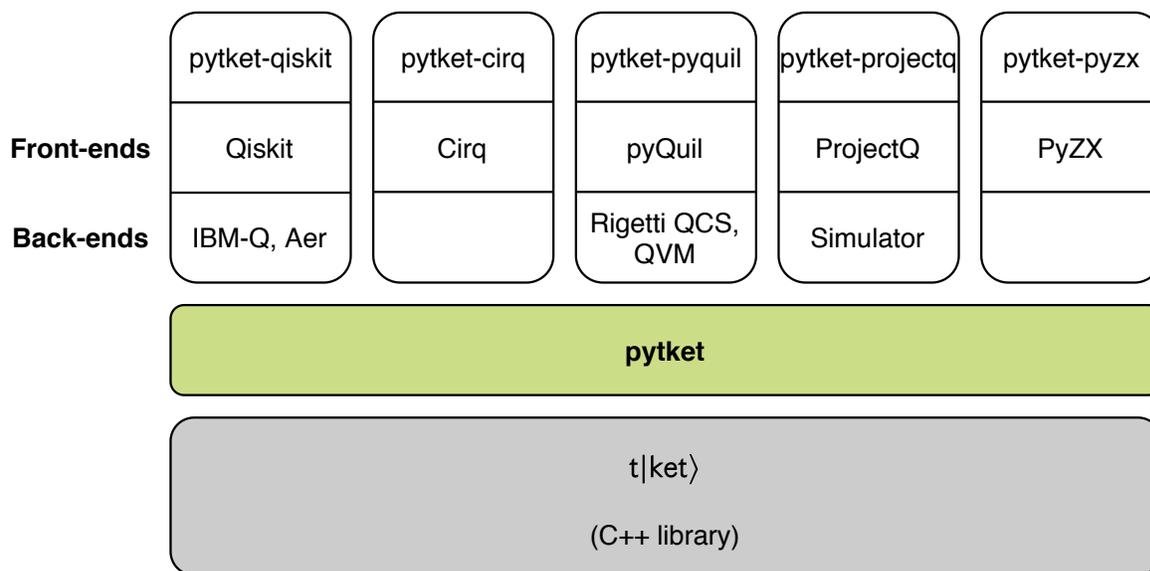


Figure 3: Components of t|ket)

automatic differentiation and interleaving compilation with execution [43, 44].

1.3. Synopsis

Here we give an overview of the t|ket) system. Subsequent sections give detailed descriptions of the modular front- and back-ends, the intermediate representation used, the transform system, some of the optimisation methods, and the system of qubit placement and routing. Section 9 provides comparative benchmark results for the performance of the optimiser and qubit allocation engine. The benchmark set of circuits is available for download.

1.4. How to get t|ket)

While the core of t|ket) is a highly optimised C++ library, the system is available as the Python module `pytket`, which provides the programming interface and interoperability with other systems. It can be installed on Linux and MacOS using the command:

```
pip install pytket
```

and the documentation is available online at: <https://cqcl.github.io/pytket/>. Figure 3 shows the components that are installed by this command.

To interface with other software packages, and to use back-ends that depend on external software, the user must also install plug-in packages. At the time of writing, the available plug-ins are: `pytket_qiskit`, `pytket_cirq`, `pytket_pyquil`, `pytket_projectq`, and `pytket_pyzx`. All of these can be installed using `pip` in the same manner as the core package. The `pytket` module is free for non-commercial use. We encourage the reader to try it out for themselves!

2. System Overview

The t|ket> system consists of two main components: a powerful optimising compiler written in C++, and a lightweight user interface and runtime system written in Python. This Python layer allows the user to define circuits and invoke compiler functions, while the runtime environment marshalls and dispatches kernels for execution, and provides convenience methods for defining variational loops, updating parameters, and collating statistics across circuit evaluations. Optional Python extensions provide interfaces to third-party quantum software systems. The overall structure is illustrated in Figure 3.

In the classical setting, a compiler translates a human-readable programming language into machine-executable object code. This process can be divided into three stages: a *front-end*, which handles lexing, parsing, semantic analysis, and other tasks which depend on the source language; a *back-end*, which allocates registers and generates suitable instruction sequences in the target machine language; and an intermediate stage, which performs data and control-flow analysis on an *intermediate representation* (IR) of the program, which is independent of both the source and the target languages. Modern compiler systems, such as LLVM [45], use a standard IR to decouple these three stages, making it relatively simple to add support for a new programming language or machine architecture to an existing compiler framework.

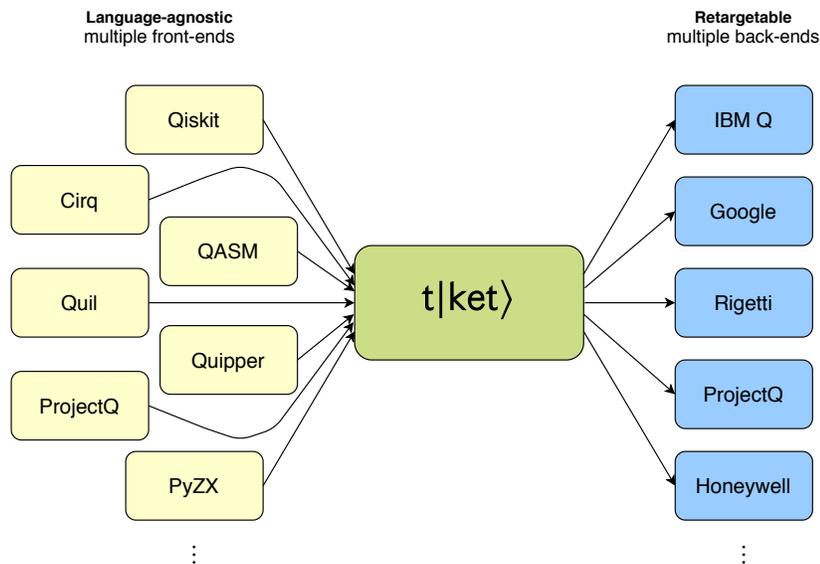


Figure 4: Modular front-ends and back-ends for t|ket>

t|ket> was designed from the ground up to be *retargetable*, meaning that it can generate code for many different quantum devices, and *language agnostic*, meaning that it accepts input from most of the major quantum software platforms. For this reason, its overall structure, shown in Figure 4, follows the same basic pattern as the LLVM. A variety of lightweight front-end units translate the desired input language

into the $t|ket\rangle$ IR. This internal representation is a generalisation of the usual language of quantum circuits based on hierarchical non-planar maps; this is described in more detail in Section 4. Standard quantum circuits are easily embedded into the $t|ket\rangle$ IR – a fact which eases the task of adding new front-ends – but many node types that are not unitary gates are also available.

Once the input has been translated to the IR, the central circuit transformation engine can begin its work. The transformation engine performs a user-configurable sequence of rewrites of the IR; some examples are described in Section 6. Typically this proceeds in two phases: an architecture-independent optimisation phase, which aims to reduce the size and complexity of the circuit; and an architecture-dependent phase, which prepares the circuit for execution on the target machine. This phase itself decomposes into a *rebase*, which maps the gates present in the circuit into those supported by the device, and a *qubit mapping* phase. The mapping phase is necessary to ensure that all qubits that are required to interact during the program are physically able to do so; this typically increases the size of the circuit, since most devices do have restricted interactions between their qubits. This is described in detail in Section 7.

The end product of this process is a *kernel*: a circuit that can be executed on the chosen target device. The kernel may then be scheduled for execution by the runtime environment, or simply saved for later.

In keeping with its focus on NISQ devices, the design of $t|ket\rangle$ is minimalistic compared to the schema proposed by Häner *et al* [26]. There is no error correction, and $t|ket\rangle$ does not include a linker, preferring to rely on application programming frameworks such as CQC’s Eumen or IBM’s Qiskit Aqua to provide libraries of common routines. The lowest layer of compilation – translation of kernels to control signals for the lasers, microwave generators, and so on – is left to the hardware implementor. However, since $t|ket\rangle$ takes into account the device’s architectural constraints during the compilation phase, this last stage of translation can be minimal.

Thanks to its retargetability, $t|ket\rangle$ can be used as a cross-compiler: source programs produced from any supported front-end can be compiled to run on hardware produced by any vendor.

3. Front-ends and Back-ends

The `pytket` interface can be used directly to build quantum circuits from individual gates in the standard way. While this may be acceptable for small experiments, more powerful high-level tools are preferable for larger or more complex tasks. For this reason $t|ket\rangle$ sports a range of lightweight front-end modules for different quantum programming systems. The industry-standard OpenQASM [46] and the functional language Quipper [27] are supported via direct source-file input. Python converters provide support for IBM’s Qiskit [32], Google’s Cirq [34] and Rigetti’s pyQuil [31], as well as the independent open-source projects ProjectQ [33] and PyZX [47]. These Python libraries in turn support higher-level application programming frameworks such

```
1 circuit = circuit_from_qasm('input.qasm')
2 q0, q1 = circuit.qubits[:2]
3 circuit = circuit.H(q0).CX(q0, q1).measure_all()
4
5 backend = IBMQBackend('ibmq_ourense')
6 if not backend.valid_circuit(circuit):
7     circuit = backend.compile_circuit(circuit)
8     assert backend.valid_circuit(circuit)
9
10 backend.process_circuits([circuit], n_shots=100)
11 results = backend.get_shots(circuit)
```

Figure 5: Code example showing front-end and back-end use. A circuit is read in from a QASM file; operations are appended to it using the `pytket` interface; the circuit is compiled to satisfy the constraints of a back-end, and then executed. The `IBMQBackend` class is included in the `pytket_qiskit` extension package.

as OpenFermion [48] or Rigetti Grove. Support for Q# [29] is planned for the next release.

Similarly, t|ket) offers multiple back-ends, each supporting a different quantum hardware platform or classical simulator. Supporting a given platform implies, firstly, generating a circuit that respects the constraints of the hardware or simulator (generally, connectivity and primitive gate limitations); secondly, the back-end must dispatch the kernel for execution and collate results. The first of these tasks is handled by the system described in Section 5. Each back-end class provides a default compiler pass, which guarantees that a compiled circuit will respect the relevant constraints.

t|ket) attempts to provide a uniform interface across the various back-end platforms, so that a user can easily change back-ends for an experiment without changing anything else in their code. At the time of writing, t|ket) supports all IBM Q and Rigetti devices via their online access services, and experimental devices produced by Honeywell Quantum Systems, Oxford Quantum Circuits and the University of Maryland. In addition, t|ket) can use the ProjectQ, IBM Aer and Rigetti QVM simulators. Various other machines are supported indirectly using either QASM output or t|ket)'s integration with Qiskit and Cirq. Figure 5 shows an example of front-end input followed by circuit compilation, execution and result retrieval via a back-end. For back-ends that support it, circuit submission and job retrieval can be performed separately, allowing asynchronous execution of the quantum circuit.

Utility functions are also provided for postprocessing of results, such as calculation of expectation values. Generic mitigation of classical state-preparation-and-measurement (SPAM) [49] errors across back-ends is scheduled for release in 2020.

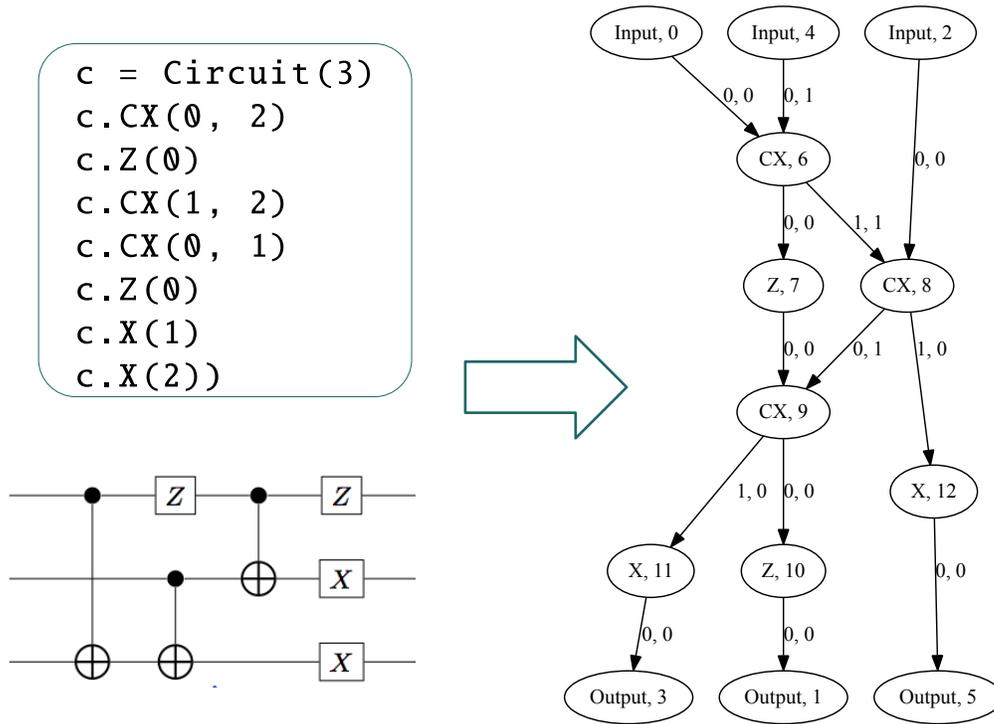


Figure 6: t|ket> circuit internal representation. Pairs of values labeling edges correspond to port numbers at the source and target vertices. Note that the Input and Output label numbers are even and odd respectively, so that qubit 0 corresponds to the path from “Input, 0” to “Output, 1”, qubit 1 is the path from “Input, 2” to “Output, 3”, and so on.

4. Representing Circuits

The standard intermediate representation (IR) in t|ket> is the *circuit*. A circuit is a labelled directed acyclic graph (DAG) with some additional structure. Vertices in the DAG correspond to operations, usually quantum or classical logic gates, but also *boxes*, a kind of opaque container which we will define later, and certain compiler-internal meta-operations. Edges in the DAG track the flow of computational resources from operation to operation. Typically, these resources are qubits, and the operations are unitary gates. Since many operations do not act symmetrically on their inputs, we add port labels for the incoming and outgoing edges at each vertex to distinguish between, for example, the control and target qubits of a CX gate. Each input port of a quantum operation is paired with an output port, allowing the path of a particular resource unit to be traced through the circuit, as shown in Figure 6.

For qubits, this linear resource management is justified by the no-cloning [50] and no-deleting [51] theorems. However, for classical operations, this effectively means that output values must overwrite their previous value. This treatment of classical bits is an

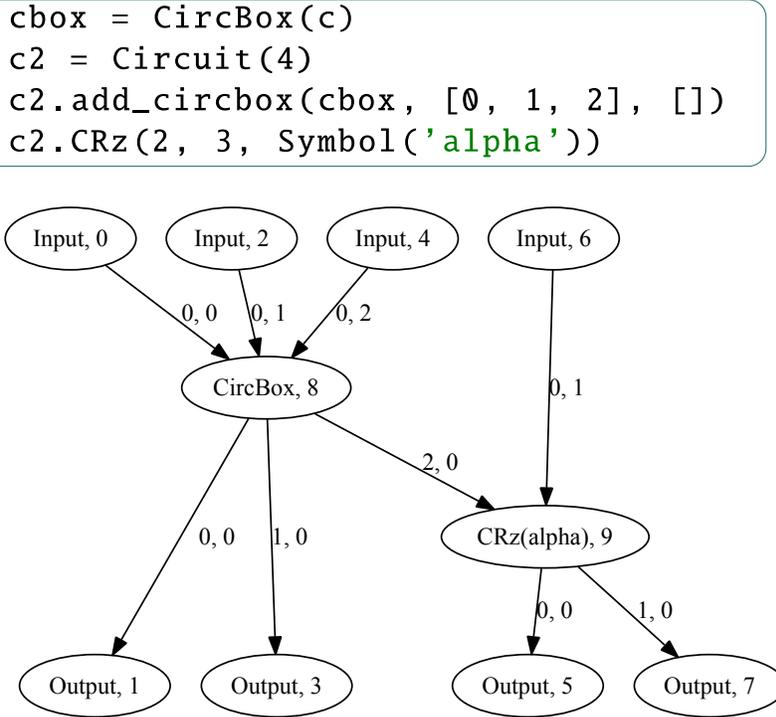


Figure 7: t|ket> circuit with a box containing the circuit from Figure 6. A parameterised controlled-Rz gate is also shown, with a symbolic parameter.

artefact of the simplified classical computational model of QASM [46] and devices that have adopted it, requiring explicit allocation of classical registers (both for classical input to execution and for result retrieval) and forbidding dynamic allocation of scratch space.

At the input and output boundaries of a circuit, the resource units – which we may identify with storage locations – are partitioned into *registers*. A circuit can contain arbitrarily many registers, and resource units are represented within a register by identifiers that are unique within the circuit. The registers specify an ordering of the resource units §. This ordering allows circuits to be composed sequentially and in parallel, and act as “port labels” for entire circuits, just as individual operations have ports within a circuit. This means that the process of composing circuits is identical to the composition of individual operations.

4.1. Gate Types

There are a wide array of allowed logic gates in t|ket>, covering the native gates of the platforms that t|ket> can interface with. An overview of the kinds of supported gates is given in Table 1. The most common quantum gates are one- and two-qubit gates,

§ In category-theoretic terms, the triple (I,O,G) of an input ordering, output ordering and graph corresponds to a structured cospan [52], where G is the apex.

Class of operation	Example
Basic single-qubit gate	Hadamard
Parameterized single-qubit gate	$Rz(\alpha)$
Basic two-qubit gate	CX
Parameterized two-qubit gate	$CRz(\alpha)$
Basic multi-qubit gate	C_nX
Parameterized multi-qubit gate	$C_nRy(\alpha)$
Classical output gate	Measure
Meta-operation	Barrier
Hierarchical vertex	CircBox

Table 1: Classes of operations available for circuits in t|ket)

reflecting the native gates on physical superconducting and ion-trap hardware, but some gates with arbitrary quantum controls are allowed; these must eventually be decomposed to hardware-native gates by the transform engine. All quantum gates in t|ket) can have arbitrary classical control, and primitive classical logic gates are supported. However, adding classical control to gates can limit the ability of the rewrite engine to optimise the circuit. An enumeration of all the allowed operation types in t|ket) can be found in the documentation at <https://cqcl.github.io/pytket/build/html/optype.html>.

Boxes are a special class of operations in t|ket). A box vertex is a container which encapsulates a whole circuit. In Figure 7, the circuit from Figure 6 is put into a box within another circuit. Boxes allow for front-ends to take in high-level descriptions with subroutines. As this subcircuit can also contain box vertices, a single circuit can contain a hierarchy of arbitrary rank. The hierarchy must be decomposed at the kernel generation stage, but this decomposition is trivial because of the compositional structure whereby circuits are equivalent to individual operations. As well as explicit circuits, box vertices can also contain other representations, which can be useful for optimising certain classes of quantum circuits. Because boxes are opaque, the parent circuit is undisturbed by the optimisation procedure acting on the subroutine. The next release of t|ket) will include such optimisations acting directly on boxes.

4.2. Gate Parameters

Unlike large-scale, fault-tolerant quantum computers, NISQ devices generally allow arbitrary angles on parameterised gates. Accordingly, t|ket) allows arbitrary angles on all parameterised gates, up to IEEE 754 double-precision [53].

In Section 1 we briefly described the variational hybrid quantum–classical algorithms proposed for NISQ devices. To enable the efficient compilation of this class of algorithms, t|ket) supports *symbolic* parameters. This allows the compilation of a parameterised circuit corresponding to an entire variational algorithm without requiring repeated compilation from scratch at each iteration of the classical optimiser. The circuit in

Figure 7 contains a parameterised controlled-Rz gate with a symbolic parameter. This class of circuits is handled using *partial compilation*: the circuit is precompiled with unknown, symbolic parameters using an expressive symbolic manipulation library. The result can be used as a template circuit and, after parameter values at a given iteration have been substituted, further simple circuit rewriting can be performed before the resulting kernel is sent to a backend to be run. This minimises the computation required between iterations of the classical optimiser, reducing the overall runtime of a variational algorithm while still using the rewrite engine of t|ket> to minimise the resource costs of the circuits. The implementation of the circuit class uses local adjacency lists at each vertex to allow near constant-time edge and vertex insertion and removal.

5. The t|ket> Transform System

In general, a quantum algorithm can be expressed in multiple ways using a given gate set; the goal is to express it in a way that minimises the gate count and circuit depth. The field of circuit optimisation is well developed, with a variety of optimisation strategies employed for different algorithms and target hardware devices [38, 54, 55]. Most commonly, a circuit can be rewritten using unitary equality between circuits, where a resource-inefficient subcircuit can be found and replaced using a closer-to-optimal one ||. The core of t|ket> is a high-performance circuit rewriting engine, referred to as the *transform system*. A function that performs rewrites using this system is called a *transform pass*. Circuit optimisation in t|ket> is described in more detail in Section 6. Aside from optimisation, the transform system has an essential role in generating circuits that are executable on the target hardware.

Each backend that t|ket> can target has associated with it a series of properties that any valid circuit must satisfy. This will include, as a minimum, the set of supported gates; for many architectures it will also include a graph representing the connectivity between the qubits. Mapping logical qubits to physical qubits also requires rewriting the circuit so that the interactions between qubits correspond only to edges in the associated connectivity graph. The transform pass that performs this rewriting is described in Section 7. Other properties may also be required, depending on the platform, and these also require transform passes.

Transform passes are composed sequentially; the resulting function is also a transform pass. For instance, a typical compiler flow will consist first of some optimisation on the circuit that has no regard for connectivity graph or gate set, followed by passes that bring the circuit closer to satisfying all of the constraints. Only if a circuit satisfies all these properties can it be executed on the target hardware.

To document and constrain the composition of transforms, the t|ket> transform engine implements a simple expression language, which follows the same principles as “contracts” in object-oriented programming [57].

|| The process of finding and replacing subgraphs in this manner is called *double pushout rewriting*. See Ehrig *et al* [56].

The functions that verify that properties are satisfied are called *predicates*. Each predicate is a function from a circuit to a boolean value: *true* if the circuit satisfies the corresponding property and *false* otherwise. These functions can incorporate some external information about the target hardware, such as connectivity graph and desired gate set; when external information is required the predicates are generated by higher-order functions. For example, to verify that the connectivity graph of a specific architecture is satisfied, a higher-order function will take in a connectivity graph and return the corresponding predicate. The full list of allowed predicates is documented at <https://cqcl.github.io/pytket/build/html/predicates.html>.

Each transform pass has a *precondition* and a *postcondition*, so that the resulting compiler pass is a Hoare triple. This is illustrated in Figure 8a. The compiler pass may be used on a circuit that satisfies the precondition, and will guarantee that afterwards the circuit satisfies the postcondition. Both the precondition and the postcondition are sets of predicates. For example, a peephole optimisation may require that the circuit be presented in a certain gate set before it can be applied. This gate-set predicate forms the precondition of the pass. The optimisation can then guarantee to the user that the rewrite rule will return the circuit in a different gate set; this guarantee is the postcondition.

These Hoare triples may then be composed, so that a custom rewriting sequence can be generated, as shown in Figure 8b. If the triples are correctly matched, so that no intermediate conditions are conflicting, the custom sequence is valid.

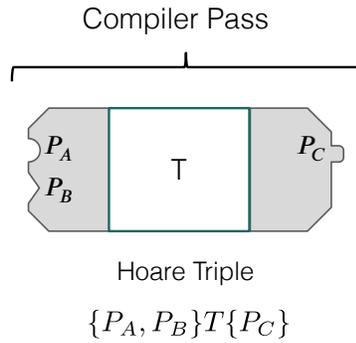
More sophisticated combinators, such as loops, can be useful for optimisation passes. For example, a user may wish to continue applying a sequence of rewrite rules until no further rewrites are possible. These combinators may be composed in the same way as sequences. When looping combinators are used, termination of the resulting pass is not guaranteed.

The full list of compiler passes and combinators can be found at <https://cqcl.github.io/pytket/build/html/passes.html>.

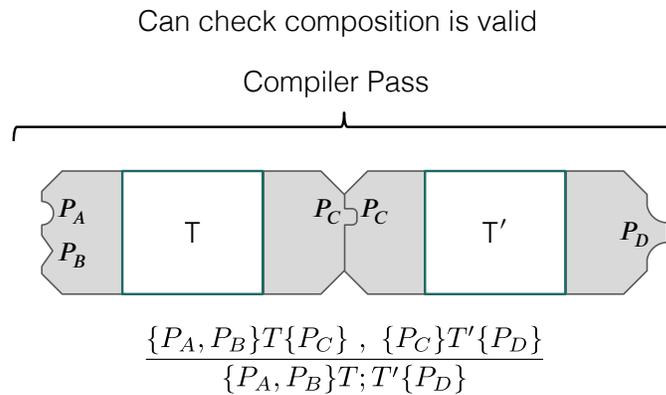
6. Circuit Optimisation Methods

With the limited fidelity available on NISQ devices, effective circuit optimisation is essential in order to extract all available performance out of the machines. The goal is to identify equivalent circuits that will accumulate less noise when run on a real device.

Circuit optimisations in t|ket) are provided as compiler passes, which can be composed into larger routines. High performance is obtained by optimising at each stage in the compilation pipeline, so it is beneficial to have both powerful optimisations that can yield better results when not constrained by qubit connectivity or gate set and procedures targeted at specific architectures. t|ket) contains some methods that are architecture-agnostic and others that are architecture-aware (parametrised over the properties of the device). Many of the architecture-agnostic passes will additionally preserve any connectivity already satisfied by the inputs, allowing them to be applied



(a) A compiler pass is a transform pass with associated pre- and post-conditions.



(b) Composition of compiler passes. The resulting Hoare triple is the standard sequential execution schema for two programs.

after routing. Designing optimisations in this way provides retargetability without sacrificing performance.

6.1. Circuit metrics

Attempting to use the actual fidelity as a cost function would require accurate simulation of the quantum circuit with realistic noise models, which is both computationally expensive and highly dependent on a specific target architecture. Further, because real devices have noise sources that are complex and hard to characterise, simple extrapolation from single-gate performance can significantly overestimate the actual performance of the device, necessitating more sophisticated, holistic measures [58, 59, 60]. However, simpler metrics can give good, device-independent approximations to noise.

Naively optimising for gate count acknowledges the key fact that all gates will introduce some degree of noise. However, NISQ devices tend to provide fast, high-fidelity single-qubit rotations, with the error rates of multi-qubit operations being an order of magnitude worse [61]. The primary focus for most optimisations in $t|ket\rangle$ is to minimise the *two-qubit gate count*, which penalises the use of these slower and less accurate

operations.

Definition 6.1. The *two-qubit gate count* of a circuit is the number of maximally-entangling two-qubit gates used in the circuit.

This is often referred to as CX-count, since any other maximally-entangling two-qubit operation (such as CZ or $e^{iXX\pi/4}$) is equivalent to a single CX up to local unitaries. This is analogous to the *T-count metric* used at the error-correcting level.

Omitting single-qubit gates entirely from consideration improves device-independence, since the number of gates required varies significantly with the gate set (for example, a single U3 gate from the IBM specification can capture any rotation, while up to three are needed if decomposed into the underlying Rz and Ry gates).

The short coherence times of qubits strongly correlates the fidelity of the circuit with the time taken to execute. An ideal device will be able to parallelise gates acting on disjoint qubits to mitigate this. We can obtain a good approximation to the time taken on such a device by considering the *depth* of the circuit.

Definition 6.2. For a gate g with predecessors $P(g)$, we define $\text{depth}(g)$ by:

$$\text{depth}(g) \triangleq \begin{cases} 0 & P(g) = \emptyset \\ 1 + \max_{p \in P(g)} \text{depth}(p) & P(g) \neq \emptyset \end{cases}$$

The *depth* of a circuit is the maximum value of $\text{depth}(g)$ over all gates g . For any gate type G , the *G-depth* of the circuit is obtained by considering only the contribution from G gates.

Again, given the characteristics of multi-qubit operations on current hardware, CX-depth (or depth with respect to any other maximal two-qubit gate) gives a device-agnostic estimate of the time cost of a circuit.

6.2. Peephole optimisations

Circuit optimisations in t|ket) can broadly be categorised into *peephole optimisations* and *macroscopic analysis*. Peephole optimisations are analogous to their namesake in classical compilers, where a sliding window traverses the instruction graph, looking for specific small patterns or classes of subcircuits and substituting equivalent subcircuits (with lower gate counts or depth) in place. Basic examples include the elimination of redundant gates such as identities, gate-inverse pairs, and diagonal gates before measurements. Local gate commutation rules can be considered at the point of pattern identification, or as standalone passes to enable further optimisations.

These techniques are generic, in the sense that they are not tuned for particular applications. The majority are written for best performance in the intermediate gate set of CX, Rz, and Rx, though when they can be expressed more naturally in a different gate set (such as a set of Clifford gates), rebase passes can be applied to convert between them.

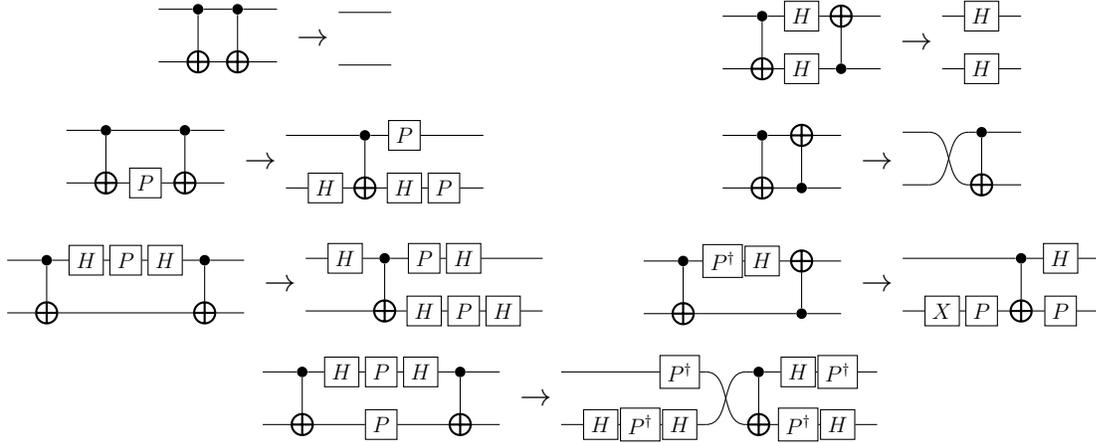


Figure 9: Clifford identities that can be recognised and applied in t|ket) to reduce the CX count. $Rz(\frac{\pi}{2})$ (phase) gates are represented by P in the diagrams. The two identities that would introduce SWAPs can invalidate any connectivity, so can optionally be disabled.

Clifford circuits are defined as the class generated by CX, Hadamard, and $Rz(\frac{\pi}{2})$ gates. These are known to be efficiently simulable [62, 63], and there is a wide literature on simplification techniques [64, 65, 66]. In particular, there are several useful small identities for reducing the CX-count of a circuit, which t|ket) can recognise and apply: these are summarised in Figure 9.

The Cartan decomposition [67] specifies a way to synthesise arbitrary n -qubit unitaries into sequences of local unitaries on fewer qubits and a small number of entangling operations between them. This decomposition gives the following upper bounds for small instances:

Theorem 6.3. *Any single-qubit unitary can be decomposed into a sequence of at most three rotations using any choice of R_x , R_y , and R_z gates. The angles of rotation are given by the Euler-angle decomposition of the combined rotation on the Bloch sphere.*

Theorem 6.4. *Any two-qubit unitary can be synthesised using at most three CX gates and 15 parametrised single-qubit gates (from any choice of R_x , R_y , and R_z), given by the KAK decomposition [68, 69].*

t|ket) implements the Euler and KAK decompositions by scanning the circuit graph for long sequences of gates over one or two qubits and replacing them whenever this helps to reduce CX count or overall gate count. Known closed-form expressions for manipulating Euler angles allow the single-qubit reduction to be performed on symbolic circuits. t|ket) does not currently support performing the KAK decomposition with symbolic gate parameters, or a generic Cartan decomposition for more than two qubits.

6.3. Macroscopic optimisations

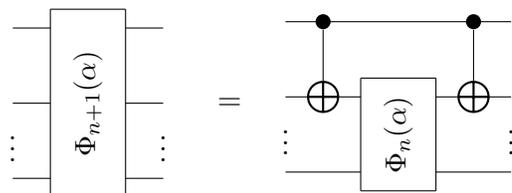
Other optimisation procedures aim to identify high-level macroscopic structures in the circuit or alternative mathematical representations of different classes of circuits that are

easier to manipulate than individual gates. The general procedure here is to recognise these structures or subcircuits of the appropriate class and treat them as first-class gates. The algebra of the structure or representation can identify non-local optimisations on the original circuit. Efficient synthesis methods can then be applied to reduce these back down to primitive gates in a way that uses fewer CX gates, parallelises them better, or restructures the circuit to enable more peephole optimisations.

As simulation of molecular systems is a leading candidate application for NISQ devices, t|ket> implements a novel technique for optimising a new class of multi-qubit subcircuits, called *Pauli gadgets*, which occur frequently in chemistry circuits designed for this purpose.

Definition 6.5. The *phase gadget* $\Phi_n(\alpha)$ is a canonical representation of a multi-qubit operator of the form $e^{\frac{1}{2}i\alpha Z^{\otimes n}}$.

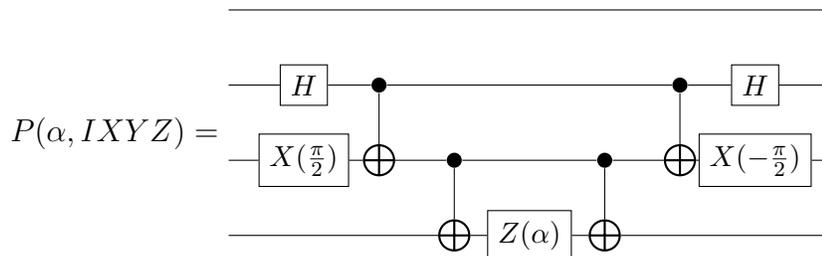
$$\Phi_1(\alpha) := Z(\alpha) \quad \Phi_{n+1}(\alpha) := (CX \otimes 1_{n-1}); (1_1 \otimes \Phi_n(\alpha)); (CX \otimes 1_{n-1})$$



Definition 6.6. The *Pauli gadget* $P(\alpha, s) := U(s); \Phi_{|s|}(\alpha); U(s)^\dagger$ is a canonical representation of a multi-qubit operator of the form $e^{\frac{1}{2}i\alpha s}$, where s is a string (tensor product) of Pauli operators and the unitary $U(s)$ is defined recursively:

$$U(Zs') := I \otimes U(s') \quad U(Ys') := X\left(\frac{\pi}{2}\right) \otimes U(s') \quad U(Xs') := H \otimes U(s')$$

Example 6.7. The simplest construction of a Pauli gadget is a single parameterised rotation gate conjugated by a cascade of CX gates and some single-qubit Clifford gates.



The authors have previously written a comprehensive account of Pauli gadgets and their use in t|ket> [70]. Such gadgets enjoy a powerful equational theory, giving rules for commutation, merging, and interaction with Clifford gates, which are easily proven using the ZX-calculus [71]. By recognising these structures in the input circuit, optimising the sequences of gadgets, and efficiently transforming them back to a standard gate set, we can achieve depth reductions greater than 50%. See Figure 20 for a summary of results comparing this technique in t|ket> to other compiler stacks for optimising circuits relating to electronic structure problems.

The first step when optimising with macroscopic structures is to identify good candidates in the circuit. It is obviously preferable to work with circuits that are already built from the structures to simplify this step. The integration of t|ket) with application software can make this possible by, for instance, allowing users to directly insert Pauli gadgets into the circuit using the corresponding box type.

Future versions of t|ket) will expand on this area of optimisations to cover other useful intermediate representations, including phase polynomials [38, 72], ZX-diagrams [73, 65, 74], Clifford tableaus [63, 75], and linear-reversible functions [76].

6.4. Example procedure

Each of these methods gives rise to a compiler pass that can either be invoked on its own or composed (as described in Section 5) into more effective routines. t|ket) comes with some predefined passes combining several of these optimisations. Each backend has a default compilation pass, which guarantees (as far as possible) that the output will be compatible with the backend’s hardware or simulator requirements; these passes include a small selection of the peephole optimisations for fast, basic gate reduction.

Figure 10 demonstrates how to compose the basic passes in `pytket`, recreating the effect of the pre-built `SynthesiseIBM` pass. Starting with `RebaseIBM` will decompose multi-qubit gates into a consistent gate set that is easier to manipulate. The `RemoveRedundancies` pass covers a handful of optimisations based on removing different types of redundant gates. Applying commutation rules can potentially uncover more candidates for removal, so the `simplify` routine is repeated until the gate count stops decreasing.

7. Mapping to Physical Qubits

Quantum computing devices have different constraints on possible operations between their physical qubits. Some hardware allow two-qubit (or higher order) operations between any set of physical qubits, while others do not. We define some physical qubits as being connected if the hardware’s primitive multi-qubit operations can be executed between them. The connectivity constraints of a hardware can be specified by an undirected graph $G_D = (V_D, E_D)$, where the vertices V_D are the physical qubits and edges E_D connect physical qubits which can interact. Figure 12 shows an example connectivity graph. As logical quantum circuits are usually written without considering these connectivity constraints, they typically must be modified before execution on a hardware to ensure that every logical multi-qubit operation is mapped to connected physical qubits. We define logical qubits and operations as those present in the logical quantum circuit and state the routing problem as finding a mapping of logical operations to allowed physical operations.

The routing problem is solved by permuting the mapping of logical qubits to physical qubits throughout a circuit’s execution, which is achieved by adding SWAP operations.

```

1  from pytket.passes import (
2      CommuteThroughMultis, RebaseIBM,
3      RemoveRedundancies, USquashIBM, SequencePass,
4      RepeatWithMetricPass)
5
6  simplify = SequencePass([
7      RemoveRedundancies(),
8      CommuteThroughMultis(),
9      USquashIBM()
10 ])
11 mypass = SequencePass([
12     RebaseIBM(),
13     simplify,
14     RepeatWithMetricPass(simplify, lambda circ : circ.n_gates)
15 ])

```

Figure 10: Code example showing how individual optimisation passes can be composed into a more complex routine

CliffordSimp Recognises patterns from Figure 9 to reduce CX count.

CommuteThroughMultis Commutes single-qubit gates through multi-qubit gates towards the front of the circuit.

KAKDecomposition Reduces two-qubit subcircuits to at most three CX gates using the *KAK* decomposition [68, 69].

OptimiseCliffordsZX Represents a Clifford circuit as a ZX-diagram, reduces it to a canonical form, and resynthesises it as a circuit.

OptimisePhaseGadgets Identifies phase gadgets [70] in the circuit and resynthesises them in a shallow manner, attempting to align CX gates between adjacent gadgets for further simplification.

PauliSimp Represents a circuit as a sequence of Pauli gadgets and a Clifford circuit, then resynthesises them pairwise [70].

RemoveRedundancies Removes redundant gates, including gate-inverse pairs, identity rotations, diagonal gates before measurements, and adjacent rotations that can be merged.

USquashIBM Merges adjacent U1, U2, and U3 gates into a single U1 or U3 gate by Euler-angle decomposition.

Figure 11: Some of the elementary optimisation passes available in t|ket)

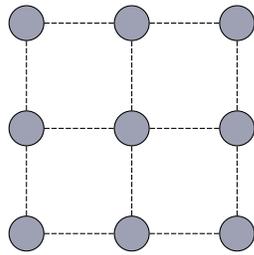


Figure 12: An undirected graph showing connectivity constraints for a hypothetical 9 qubit device.

Sometimes gate decompositions can be used to convert non-adjacent multi-qubit gates to distance 1 implementations. An example of this for the CX gate is shown in Figure 13. Finding an optimal solution to the routing problem in this manner is NP-complete in general [77]. Note that SWAP operations have different implementations on different hardware; ion trap devices have physical SWAPs while superconducting devices require the logical states to be transferred between two physical qubits through three CX gates, shown in Figure 14. A solution is reached when the logical circuit is modified such

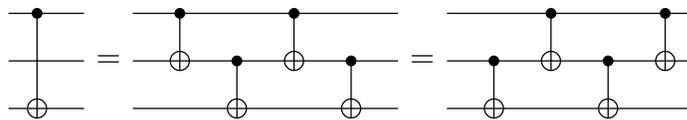


Figure 13: Distance 2 distributed CX gate and decompositions to distance 1 CX gates

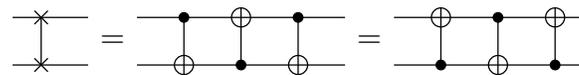


Figure 14: SWAP gate and decompositions to CX gates

that there is an injective map of logical qubits to physical qubits, or *placement*, p where for every two-qubit gate acting on logical qubits (q, q') , the mapped physical qubits are connected on the connectivity graph G_D , or that $(p(q), p(q')) \in E_D$.

In some cases a placement p can be found that solves the routing problem without adding SWAP operations. Treating logical qubits as vertices and two-qubit interactions between them in the circuit as edges, we can form an interaction graph for a logical circuit $G_I = (V_I, E_I)$. If there is a subgraph monomorphism $p : V_I \rightarrow V_D$ which respects $(q, q') \in E_I \Rightarrow (p(q), p(q')) \in E_D$, then only a relabelling of logical qubits to physical qubits is required.

In Figure 15, a short example CX circuit is shown. This circuit can be mapped

to the connectivity graph from Figure 12 without the addition of additional gates, as shown in Figure 16.

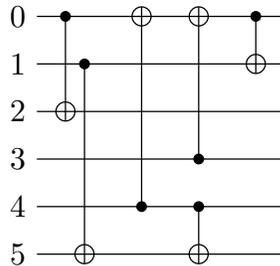


Figure 15: An example 6 qubit circuit with only CX gates.

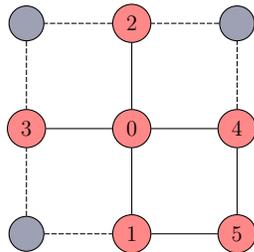


Figure 16: An example mapping of logical qubits in Figure 15 to physical qubits in Figure 12. This example satisfies the routing problem without logical circuit modification. The solid lines between red nodes represent physical interactions performed by the circuit; grey nodes and dashed lines are unused by the circuit.

t|ket> solves the problem in two steps: finding an initial placement of logical qubits to physical qubits and subsequent addition of SWAP operations to the circuit. We consider this to be a *dynamic* approach, in contrast to *static* approaches [78, 79, 80], that partition circuits into parallelised slices of two-qubit interactions, and then use SWAP networks to permute logical qubits between placements that satisfy these slices.

7.1. Noise aware graph placement

The initial placement p is chosen to maximise the fidelity of the circuit implementation on the device, using both a proxy heuristic which tries to minimise additional gate overhead from routing, and an error heuristic which uses device error characteristics.

Indirectly, knowing only the connectivity graph G_D , candidate placements are chosen to minimise the number of gates the subsequent routing procedure will need to add, as these additional operations are most likely error prone two qubit gates. Routing will add gates dynamically as it proceeds through a circuit, and so it is in general not possible to predict which placements will correspond to the fewest gates added. As a

heuristic, placements are found such that a maximum number of two-qubit operations at the beginning of the circuit can be completed with no SWAP gates added.

First this problem is cast as finding a subgraph monomorphism $p : V_I \rightarrow V_D$ which respects $(q, q') \in E_I \Rightarrow (p(q), p(q')) \in E_D$, for the interaction graph $G_I = (V_I, E_I)$ and device graph $G_D = (V_D, E_D)$. If a monomorphism cannot be found, the routine removes an edge from G_I belonging to the latest circuit slice and attempts the graph matching again. This iterates until a monomorphism is found.

Logical qubits $q \in V_I$ which no longer have any edges are removed from G_I , thus the subgraph monomorphism routine in practice produces a set of candidate *partial* placements: maps which only act on a subset of the logical qubits in the circuit. The subsequent routing procedure can accept this as input, and will naively place unmapped qubits near those they next interact with as it proceeds. As device architecture graphs G_D currently have, and will likely continue to have, large regular subgraphs, the set of matches can be large, especially when V_I is small compared to V_D .

If gate fidelity information for individual qubits is available for the target hardware, these candidate placements are scored for maximum expected overall fidelity and the highest scoring is chosen. In NISQ devices qubits and primitive gates often have highly heterogeneous error characteristics, using this to choose from the possible equivalent graph matches can result in a higher fidelity implementation of a given circuit on a given device. Section 9.2 compares the performance of different placement methods available in t|ket>.

7.2. Routing

Given an initial partial placement p , the routing algorithm adds SWAP operations until all logical operations satisfy connectivity constraints. As SWAP operations are added, p is permuted, and so we define a temporary placement p' which is the permutation of p from added SWAP operations up to some slice of circuit gates S .

Two-qubit gates in the circuit are iterated through in time order (via a topological sort of the DAG), finding the first set of two-qubit interactions (q, q') such that $(p(q), p(q'))$ does not respect G_D and no q is in multiple interactions. We call this set the first slice S_0 and log the permutation of p, p' , up to S_0 . The routing algorithm then aims to pick the optimal edge $e \in E_D$ of the connectivity graph to implement a SWAP operation on, given interacting logical qubits in S_0 and p' .

A set of candidate placements $\{p''_0, \dots, p''_n\}$ is constructed by permuting instances of p' with SWAP operations on edges in E_D . If an edge has no qubits in S_0 it is ignored. Each candidate placement is scored and the winning placement is chosen, with the scoring function based on the distance between interactions in S_0 given p'' . If there is no winning placement for S_0 then tied placements are scored for a new slice S_1 , where S_1 is the next set of set of two-qubit interactions (q, q') in the circuit such that $(p(q), p(q'))$ does not respect G_D and no q is in multiple interactions. If there is no winning placement for S_1 then tied placements are scored for a new slice S_2 . This is repeated until there is

a winning placement p'' for some S_n .

The winning placement p'' is produced from p' via a permutation along its associated winning edge e . In most cases a SWAP operation is inserted along e directly before S_0 and a new first slice S_0 is found.

In some cases a distributed CX is considered instead: at least one of the logical qubits associated with e is in an interaction (a two-qubit gate g) in S_0 . If g is a CX gate and its logical qubits are at distance two on the device graph G_D (for the temporary placement p') then a distributed CX may be added instead. A new two element set of candidate placements is constructed comprised of p' and p'' and a similar scoring process is implemented, comparing p' and p'' over multiple slices (S_0, S_1, S_2 and so on). If p' wins g is replaced with a distributed CX and no SWAP operation is added. Else, if p'' wins the SWAP operation is added and p' replaced with p'' .

This whole process is then repeated, finding new first slices S_0 and new winning placements p'' . The algorithm terminates when S_0 is returned empty.

The algorithm employs a high performance heuristic, which when coupled with an efficient C++ implementation results in fast runtime. t|ket> routing typically performs at least as well as other software solutions when comparing circuit size and depth [77].

Heterogeneity in NISQ device noise means a routed circuit with minimal SWAP overhead may not always prove best though. Some solutions consider device noise [36, 81] when routing, using gate fidelity information to produce a routed circuit with best execution fidelity. This motivated a routing solution we implemented that used a fidelity-based heuristic approach to score and pick SWAP operations, for which the scoring method used an estimation of the noise accumulated in realising all interactions in a slice S_0 . The estimate was produced by finding SWAP paths required to permit interactions in S_0 , and then calculating potential error accrued by each logical qubit in realising these paths. In practice we found that the fidelity heuristic could not accurately determine when diverging from adding the minimal number of SWAP gates would improve circuit fidelity, and so in general aiming to minimise SWAP operation overhead provided better results.

8. Applications

Quantum chemistry simulations are performed using a supplementary software package called Eumen. This provides an interface between traditional quantum chemistry problems and various hybrid classical–quantum algorithms, enabling effective chemistry simulations on NISQ hardware using t|ket>. For such simulations, Eumen accepts a range of input parameters, such as the molecular or lattice geometry, system charge, multiplicity restrictions, type of simulation, ansatz, optimisations, hardware backend, and qubit mapping. t|ket> mediates between Eumen and the hardware on which the quantum-algorithmic part of the chemistry simulation runs.

Eumen can compute optimal geometries and properties of the ground state or excited states. For example, ground-state energies may be calculated using VQE or imaginary-time evolution methods. For excited-state calculations, one can use methods such as quantum subspace expansion, reduced density matrix approximation, penalty functions, and symmetry constraints [82]. These methods require the measurement of either the expectation value of many-body operators or the overlap of two different states; these measurements are performed by t|ket> using results from the quantum hardware. The states may be prepared with hardware-efficient ansätze or the approximated circuit representation of various physically motivated ansätze, such as UCCSD, k-UpCCGSD, or the time evolution operator.

The depth of the state preparation circuit is significantly reduced by t|ket>'s ansatz-specific optimisation methods, which can identify specific structures in circuits and reduce the gate count required for their execution, as described in Section 6. Future optimisation methods specific to QAOA instances are also planned. Construction of such structures and subsequent optimisation is aided by boxes: for example, a Pauli gadget can be added to a circuit in abstract form via a `PauliExpBox` operation. These operations, and variational circuits in general, can make use of symbolic parameters and compilation to simplify their use and speed up compilation. Finally, variational algorithms, and other applications that use Hamiltonian estimation via calculation of multiple terms of the Hamiltonian, can benefit from t|ket>'s back-end methods, which allow compilation and submission of multiple circuits. Circuit execution on the back-end can occur asynchronously, with results being retrieved for processing when execution is complete.

9. Benchmarks

In this section we provide some benchmarks of t|ket> compiler performance. First we conduct benchmarks of end-to-end compiler performance, including comparisons with other available quantum compilation tools. Secondly, we perform experiments on a publicly available quantum device to determine whether noise-aware placement offers a benefit. The full datasets and scripts used for generating these results can be found at https://github.com/CQCL/tket_benchmarking.

9.1. End-to-end compilation

We present a series of benchmarks of end-to-end compiler performance on a set of circuits, and compare the performance of t|ket> with that of two widely-used alternatives, Qiskit and the Quilc compiler, which are able to do both general circuit optimisation and routing.

We define end-to-end compilation as the process of translating a circuit, presented in OpenQASM, and outputting an equivalent circuit that has been optimised and has the relevant device constraints satisfied, i.e. has been routed and converted to the correct

gate set. We do not include high-level algorithm design or low-level pulse optimisation which are respectively beyond the scope of a compiler and only in its infancy as a research topic.

9.1.1. Benchmark circuits The benchmark set is made up of three test sets, with circuits of at most 10^4 initial gates. (This threshold was chosen as circuits larger than this are several orders of magnitude too large for near-term devices. In addition, the runtimes of Qiskit and Quilc already reach several minutes per circuit at this range, making it impractical to benchmark against the larger circuits.)

- (i) The IBM test set is a series of circuits published as part of the Qiskit Developer Challenge, a public competition to design a better routing algorithm. These circuits are not amenable to significant peephole optimisation to restrict the impact this can have and focus on the efficiency of the routing algorithm. However, they were designed to be easily verified for correctness by mapping the $|0\rangle^{\otimes n}$ state into some other computational basis state; as such, these tests could be circumvented by applying state preparation optimisations. The IBM circuit set in OpenQASM can be found at https://github.com/iic-jku/ibm_qx_mapping.
- (ii) The UCCSD test set is a series of circuits for electronic-structure calculations. They correspond to VQE circuits for estimating the ground state energy of small molecules by the Unitary Coupled Cluster approach [83], using some choice of qubit encoding (Jordan–Wigner, parity mapping, or Bravyi–Kitaev [84]). These circuits are very amenable to optimisation, as well as requiring routing. They are representative of algorithms that have been proposed as suitable for application on NISQ devices [83], and were generated using Qiskit Aqua. The set used here updates and extends that used by Cowtan *et al* [70], whose OpenQASM files can be found at <https://github.com/CQCL/pytket>.
- (iii) The Product Formula test set is a series of circuits for Hamiltonian simulation. These circuits are thought to be candidates for quantum advantage [85], and were used as a test case for the circuit optimizer by Nam *et al* [38]. They are given in the ASCII format of the Quipper language, and each is formed of a repeated subroutine. We convert this subroutine to a quantum circuit in OpenQASM. We included circuits both before and after optimisation by Nam *et al* [38], since they still require mapping to the architecture and have potential for further optimisation. We had to further edit the circuits to ensure the rotation angles of gates exceeded Qiskit’s very high cutoff (10^{-5}), below which the rotations would be treated as identities, making these circuits almost trivial. These circuits contain some Pauli gadgets, but also have large regions which are not amenable to this kind of optimisation. These circuits can be found at <https://github.com/njross/optimizer>.

The full collated benchmark set can be found at https://github.com/CQCL/tket_benchmarking.

9.1.2. *Experiments* We compare compilation for four different architectures:

- 1) The fully-connected graph, for which no routing is required.
- 2) The connectivity graph of IBM Rochester, a 53-qubit device.
- 3) The connectivity graph of Google’s 53-qubit Sycamore device.
- 4) The Rigetti Aspen 16-qubit architecture. For this case, all circuits with more than 16 qubits were discarded.

These connectivity graphs are shown in Figure 17. Two end-to-end comparisons were made:

- (i) To compare t|ket) to other available compilation software, the default compiler passes of Qiskit (optimisation level 3) and Quilc, and the recommended generic pass for t|ket) (`FullPeepholeOptimise`, followed by the default qubit mapping pass, `SynthesiseIBM`, and the rebase pass into the desired gateset, henceforth referred to as ‘FullPass’) were applied to all available circuits.
- (ii) To demonstrate the necessity of appropriate usage of situational compiler optimisations, we compare t|ket)’s UCCSD-specific pass (the `PauliSimp` pass followed by the ‘FullPass’ routine, henceforth dubbed ‘ChemPass’) against the default Qiskit and Quilc passes, on the UCCSD circuits (test set (ii) above). These circuits contain adjacent Pauli gadgets, and we demonstrate that the reduction in two-qubit gate count and depth can be substantial compared to optimising naively.

The benchmarks were performed using t|ket) v0.4.1, Quilc v1.16.3 and Qiskit Terra v0.12.0. All results were obtained using a machine with a 2.3 GHz Intel Core i5 processor and 8 GB of 2133 MHz LPDDR3 memory, running MacOS Mojave v10.14.

9.1.3. *Metric* The figures of merit for these benchmarks are two-qubit gate count and depth. As described in Section 6, two-qubit gates have error rates an order of magnitude higher than single-qubit gates for existing architectures [61], and so the counts and depths are reasonable proxies for the overall expected error rates of a circuit run on a NISQ device.

We defined end-to-end compilation earlier, including the conversion to the device’s native gate set. Google’s Sycamore device can accept CZ gates natively as a two-qubit operation, whereas IBM Rochester only supports CX gates. As only single-qubit Hadamard gates are required for conversion between CX and CZ gates, we discount the gate-conversion step, and accept either gate set for the two-qubit gate-count and depth metrics. Thus, unlike total gate count, for these backends the two-qubit gate count should be independent of final basis set chosen, meaning the comparison between architectures is purely based on connectivity. The exception here is the Rigetti Aspen device which can use both CZ gates and the XY family (including the iSWAP gate) natively. Since these can be obtained with similar fidelities, the simple two-qubit gate count is justified in weighting their costs equally.

9.1.4. Results Two-qubit gate count and depth benchmark results for default compilation are shown in Figures 18 and 19. The corresponding results for chemistry-specific compilation are shown in Figures 20 and 21 respectively. Each figure includes results for all three compilers and all four target backends. The figure of merit in each is multiplicative overhead of two-qubit gate count or depth, i.e. the ratio of the values before and after compilation.

The *FullConnectivity* case shows the results without the effects of routing. For default compilation passes the majority of benchmark circuits show little to no change in the number of two-qubit gates, demonstrating the difficulty of entangling-gate reduction with generic optimisation passes. In cases where gains can be made, the optimisations in t|ket) are able to make larger reductions than the other compilers, and are able to make reductions in more instances.

By contrast, as shown in Figures 20 and 21, adding the `PauliSimp` pass leads to more significant reductions on the UCCSD circuits, even after mapping onto a device with restricted connectivity. However, on other classes of circuits that don't resemble the UCCSD set, adding this optimisation can cause a drastic drop in performance, by trying to fit them to this model, and potentially making them less amenable to routing.

Across the end-to-end compilation results targeting restricted-connectivity architectures, the general ranking of performance is: t|ket), followed by Quilc (with special note of their performance for the Aspen device), with Qiskit consistently introducing a very high gate overhead. The results are sufficiently spread that for circuits viable for execution on near-term devices, the choice of compiler and pass sequence makes a significant difference to the size of the final circuit. By comparing to the *FullConnectivity* case, we can see that these differences are dominated by differences in routing performance.

9.2. Noise-aware placement

In Section 7.1 we outlined *graph placement* (GP), a subgraph-monomorphism-based method for finding initial qubit mappings, and *noise-aware graph placement* (NAGP), a fidelity-aware heuristic for scoring those placements. As the effectiveness of these methods depends strongly on the error characteristics of physical devices and the fidelity with which they execute the circuit, we assess and compare their performance by running benchmark circuits on a device with and without the mapping methods applied.

9.2.1. Benchmark circuits At the time of writing, it is difficult to implement many common algorithms on publicly-available quantum devices and extract a signal from the noise; this is why implemented algorithms are usually variational. In order to test a large enough data set to have confidence in measured differences, we instead choose to implement random circuits of constrained sizes.

Sets of random circuits are parameterised by number of qubits and total number of gates. Each circuit is generated by uniformly sampling gates from $\{X, Y, Z, H, T, S, CX\}$, and uniformly sampling from all qubits for single-qubit gates and from the set of pairs of

all qubits in the case of a CX. As the two-qubit operations are the most error-prone [61], samples that included no CX gates were excluded. Circuit sets of qubit numbers of 4 and 8, and gate counts of 20, 40, 60 and 80, were generated. Circuits over 4 qubits with 80 gates were too deep and therefore noisy for effective comparison of methods, so were omitted, leaving a total of 7 sets each with 90 samples.

9.2.2. Experiments As described in Section 7, *placement* is the task of finding initial maps from logical qubits of the circuit to physical qubits of the device, and *routing* is the addition of two-qubit gates to satisfy the connectivity constraints of the device. For these experiments, each benchmark circuit was compiled in three different ways, corresponding to three methods of calculating an initial partial placement: “None” (corresponding to no qubits placed, therefore relying on default on-the-fly placement performed by routing), “Graph Placement”, and “Noise-Aware Graph Placement”. Each placed circuit was then compiled with identical routing and post-routing optimisation passes from t|ket>.

All circuits were run on the publicly-available `ibmq_16_melbourne` device via the IBM Q Experience [88]. Correspondingly, compilation also included translating the circuits to the IBM Q gate set of $\{U_1, U_2, U_3, CX\}$. Programs for execution on an IBM Q device are sent via the API as “jobs”, with a maximum of 75 circuits in each job. All compiled circuits that corresponded to the same initial circuit were evaluated consecutively and within the same job, to mitigate the effects of device-characteristic deviations between jobs on the method comparisons. Each compiled circuit was evaluated with the maximum 8192 shots.

9.2.3. Metric The ideal metric for comparing the placement methods is overall fidelity of the implemented circuits. However, measurement of this involves a number of circuit measurements that scales exponentially with qubit number and is infeasible for large experiment sets [89]. Instead, we choose a metric that can quantify the distance between the distribution of measurements in the computational basis and the same distribution generated from an ideal simulation. This methodology requires classical resources that scale exponentially with qubit number, as it involves simulation of all circuits. However, as the techniques proposed are only relevant for NISQ-era error rates and device heterogeneity, we expect the number of qubits to remain low enough for applicability of the proposed methods while the same levels of heterogeneity also hold.

The Kullback–Leibler (KL) divergence has been used for comparing measured and simulated distributions [90]. However, it has some shortcomings. The KL divergence between two distributions P, Q over values x_i is defined as:

$$D_{\text{KL}}(P, Q) = \sum_i P(x_i) \log \frac{P(x_i)}{Q(x_i)}$$

This is asymmetric between P and Q . More importantly, it is defined to be infinite if $\text{support}(P) \not\subseteq \text{support}(Q)$. A standard technique to account for this is padding zeroes in the distribution with small values and renormalising. Although this would still show

qualitative difference, the absolute value would depend on the free parameter of tuning, so would not be a useful quantitative measure.

We instead use the Jenses–Shannon divergence D_{JS} [91], which is closely related to D_{KL} . For distributions P, Q and $M = \frac{1}{2}(P + Q)$ it is defined by:

$$D_{\text{JS}}(P, Q) = \frac{1}{2}D_{\text{KL}}(P, M) + \frac{1}{2}D_{\text{KL}}(Q, M)$$

This is a symmetric, finite function with bounds $0 \leq D_{\text{JS}}(P, Q) \leq 1$ when the base-2 logarithm is used.

9.2.4. Results Figure 22 plots the mean value of D_{JS} over each benchmark circuit set for the three placement methods, each benchmark set parameterised by qubit number and gate count. In general, GP is seen to reduce the mean D_{JS} when compared to no initial placement; this can be explained by GP reducing the number of error-prone two-qubit gates that need to be added to map the circuit. We also see that scoring of these placements using device-reported error information by NAGP is able to make further significant reductions to D_{JS} , suggesting that such exploitation of device heterogeneity is a worthwhile avenue of exploration for maximising near-term device use.

Comparing 4-qubit and 8-qubit results, D_{JS} means are higher in the 8-qubit case, as expected, as more qubits are entangled together in general and so the system is more prone to error. In the 8-qubit case, D_{JS} is also seen to monotonically increase with gate count, also matching expectations. The peak for D_{JS} mean at 40 gates for 4 qubits, for all placement methods, is unexpected and warrants further investigation.

10. Conclusions and Future Work

In this paper we have described CQC’s compiler system t|ket>, with particular emphasis on its transformation engine and qubit mapping routine. We showed that t|ket> offers significant improvement in terms of gate count and gate depth over other comparable compiler systems when evaluated on realistic quantum circuits and real quantum architectures. Further, for devices with heterogeneous gate and qubit error rates t|ket> can use the component-level fidelity information to appreciably improve overall device performance. For NISQ-era quantum computing such performance differences may be the margin between success and failure.

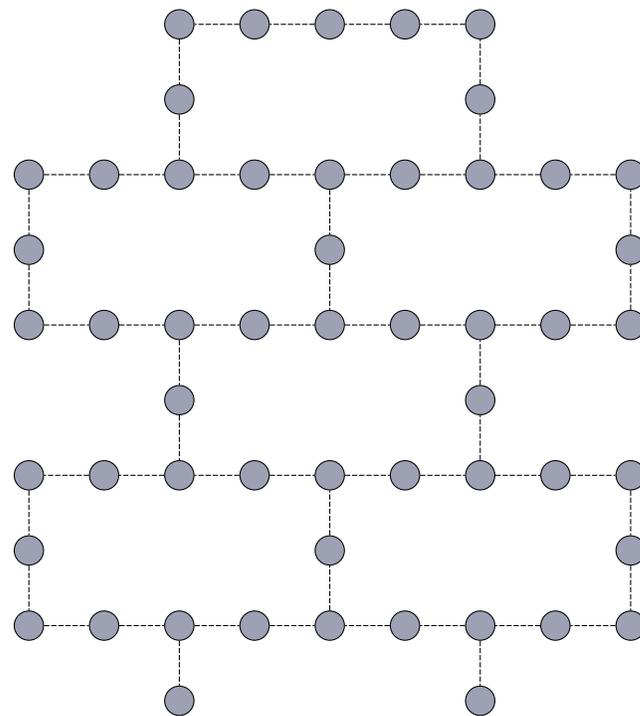
The flexible design of t|ket> presents many possibilities for future improvements. Here we sketch three promising directions.

For all-to-all connectivity, the `PauliSimp` pass achieved staggering depth reductions on the chemistry benchmark set, which is not totally surprising because it was designed to exploit the recurring structures found in UCCSD ansätze. However there is a lot of scope to improve this method, particularly if the Pauli Gadgets are treated as multi-qubit gates and synthesised by the architecture-aware phase of the compilation process, solving the problem mentioned in Section 9.1.4. More generally we expect the use of higher

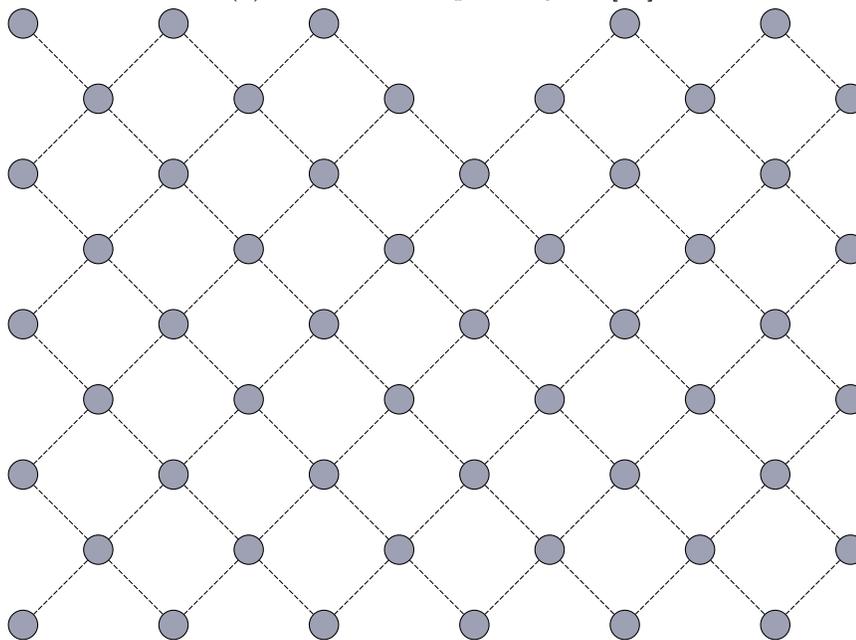
level “big gates”, equipped with their own equational theory, and tuned to particular algorithms or ansätze (for example QAOA instances) to yield similar improvements. This kind of application-specific optimisations cannot be discovered by working with random circuits and require real use cases.

Recent results on quantum volume [58] suggest that available qubit numbers already surpass the gate fidelity by such a margin that a large fraction of the qubits cannot effectively be exploited. This suggests that, in the near term at least, extremely shallow circuits will be required. One possible route to such depth reduction is to exploit large numbers of ancillary qubits, combined with techniques from measurement-based quantum computation [92] to effectively trade time for space. The ZX-calculus [71] is already incorporated into t|ket> and has proven an effective tool for MBQC calculations in the past [73, 93, 94].

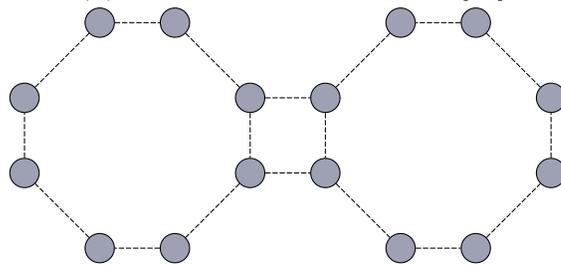
Finally, we will look at techniques to attack the noisiness of NISQ devices head on. Incorporating a very naive noise model into t|ket>’s qubit placement algorithm (Section 9.2) made a noticeable difference our results. However it is well known that the noise channels in real devices are much more complex and more difficult to characterise [95, 60, 96]. Incorporating better analysis of the device errors into the compilation process, and techniques to suppress and mitigate errors [25, 23, 22] surely have a role to play in the compilation process for NISQ devices for the foreseeable future.



(a) Rochester 53-qubit layout [86].

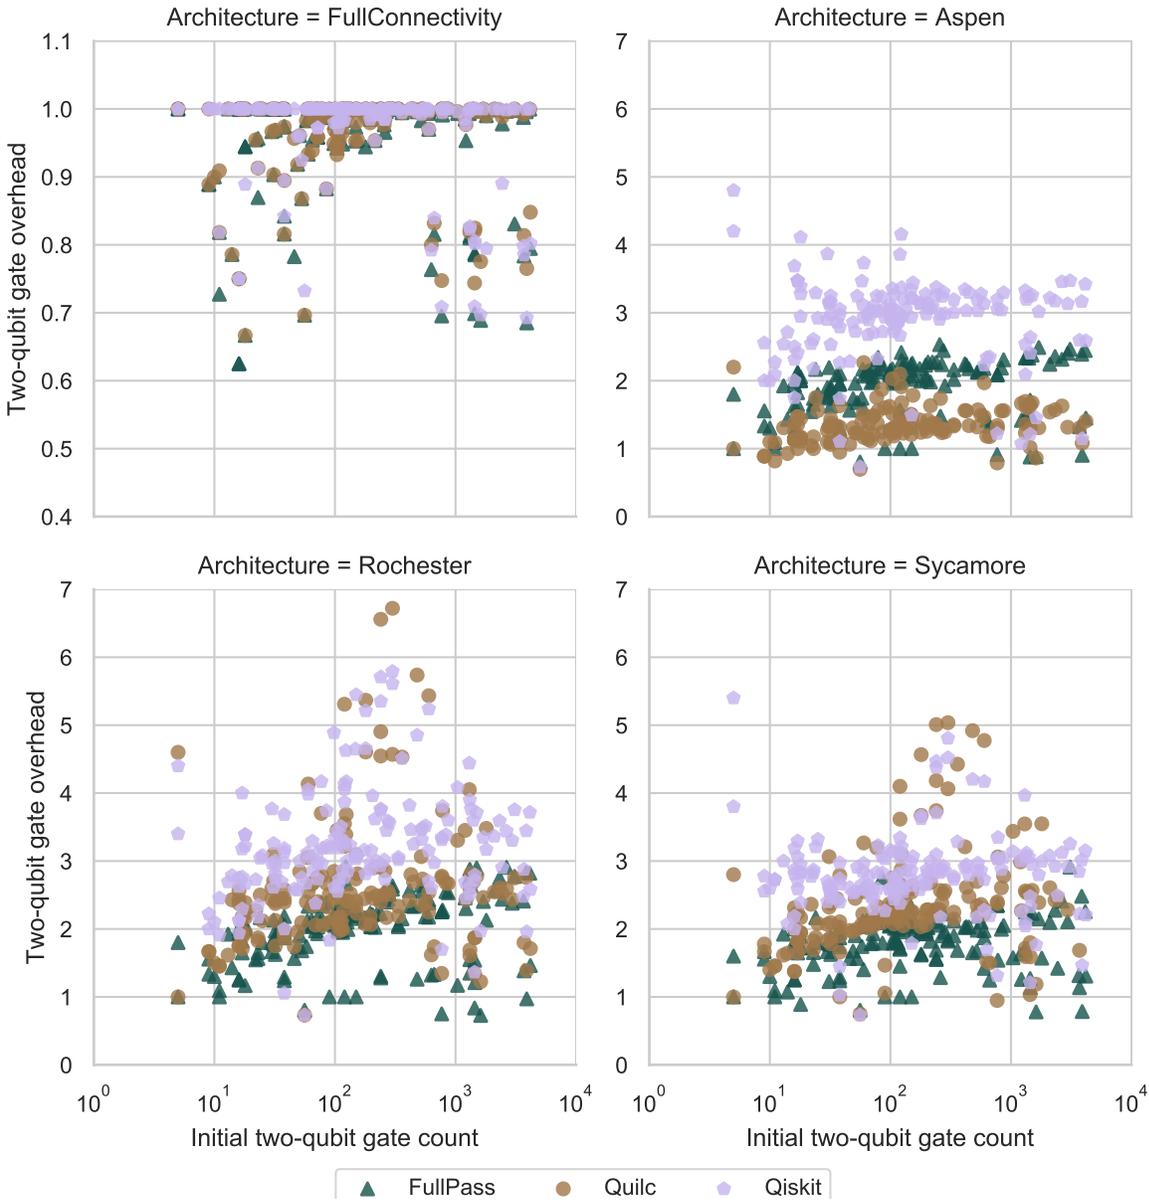


(b) Sycamore 53-qubit layout [61].



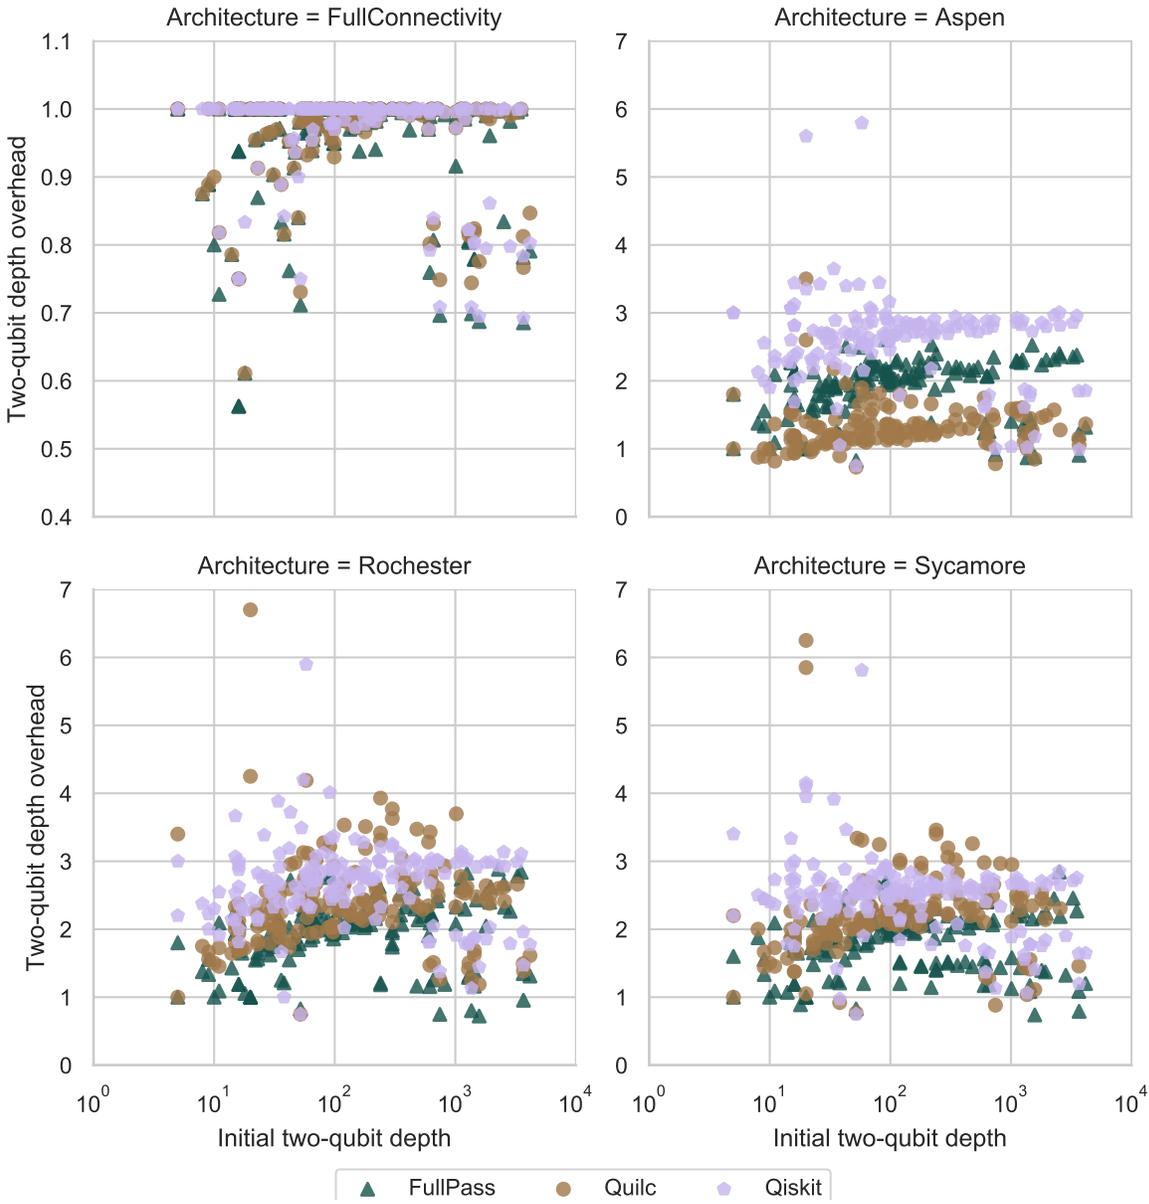
(c) Rigetti Aspen 16-qubit layout [87].

Figure 17: Device connectivity layouts used in end-to-end compilation benchmarks.



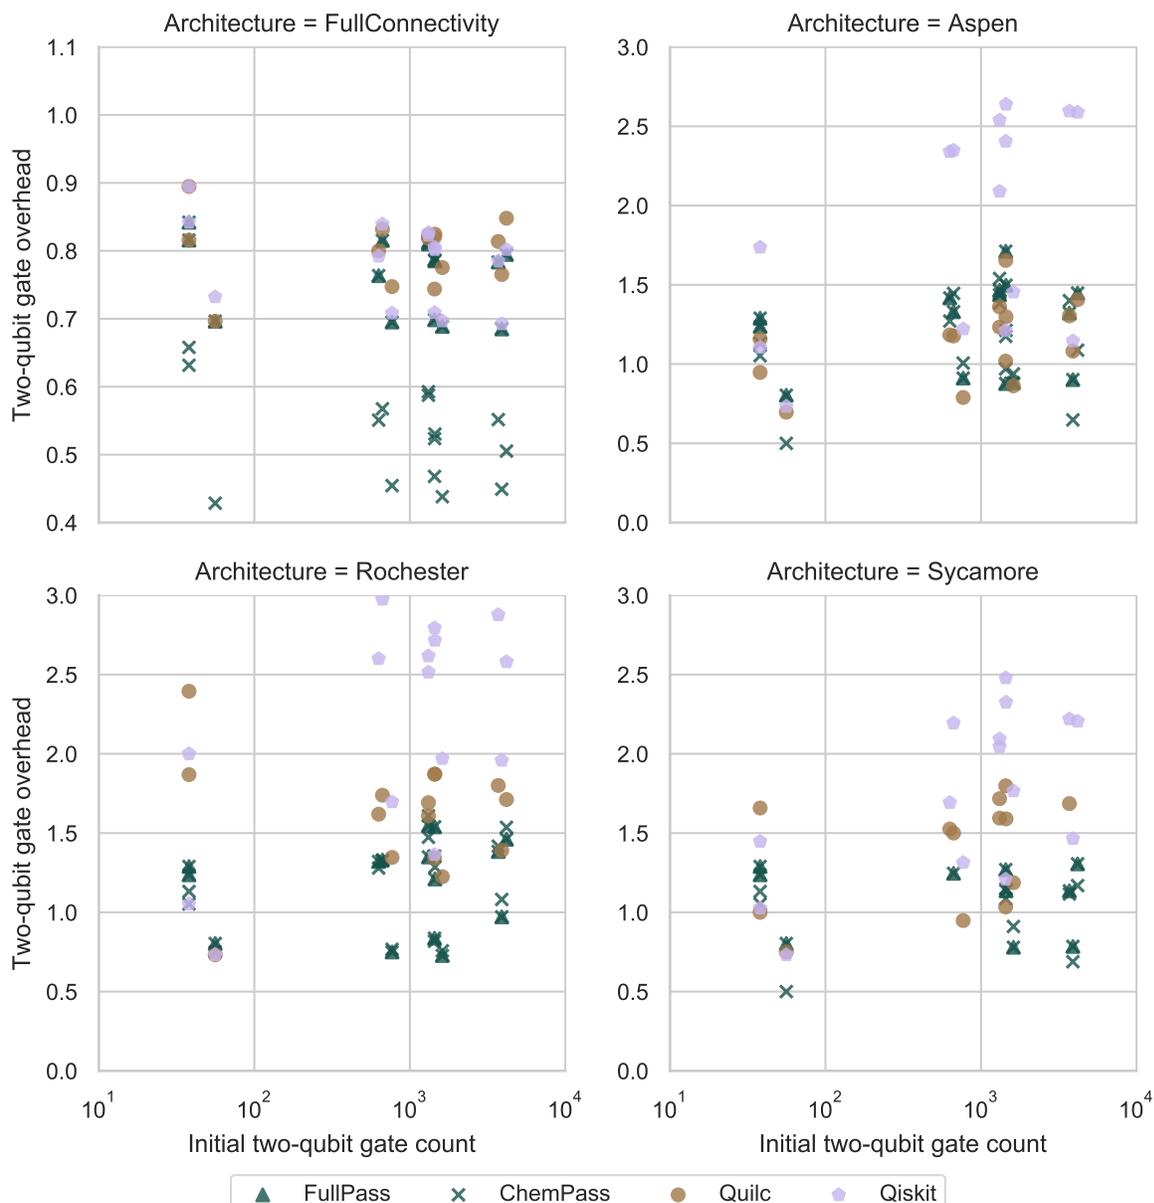
	FullConnectivity	Aspen	Rochester	Sycamore
FullPass	0.939 ± 0.011	1.864 ± 0.039	1.975 ± 0.042	1.773 ± 0.034
Quilc	0.949 ± 0.010	1.362 ± 0.022	2.595 ± 0.074	2.343 ± 0.061
Qiskit	0.958 ± 0.010	2.896 ± 0.052	3.201 ± 0.065	2.789 ± 0.047

Figure 18: Default compilation benchmarks over all benchmark circuits. The multiplicative overhead in two-qubit gate count from input circuit to output circuit is plotted against input two-qubit gate count. The table shows means across the circuit sets and associated standard error. In general, routing induces a larger-than-1 overhead, but for *FullConnectivity* when no routing is required it is 1 or below. ‘FullPass’ refers to the recommended t|ket) routine, consisting of the `FullPeepholeOptimise` pass, followed by the corresponding qubit mapping pass, `SynthesiseIBM` and rebasing to the final gate set.



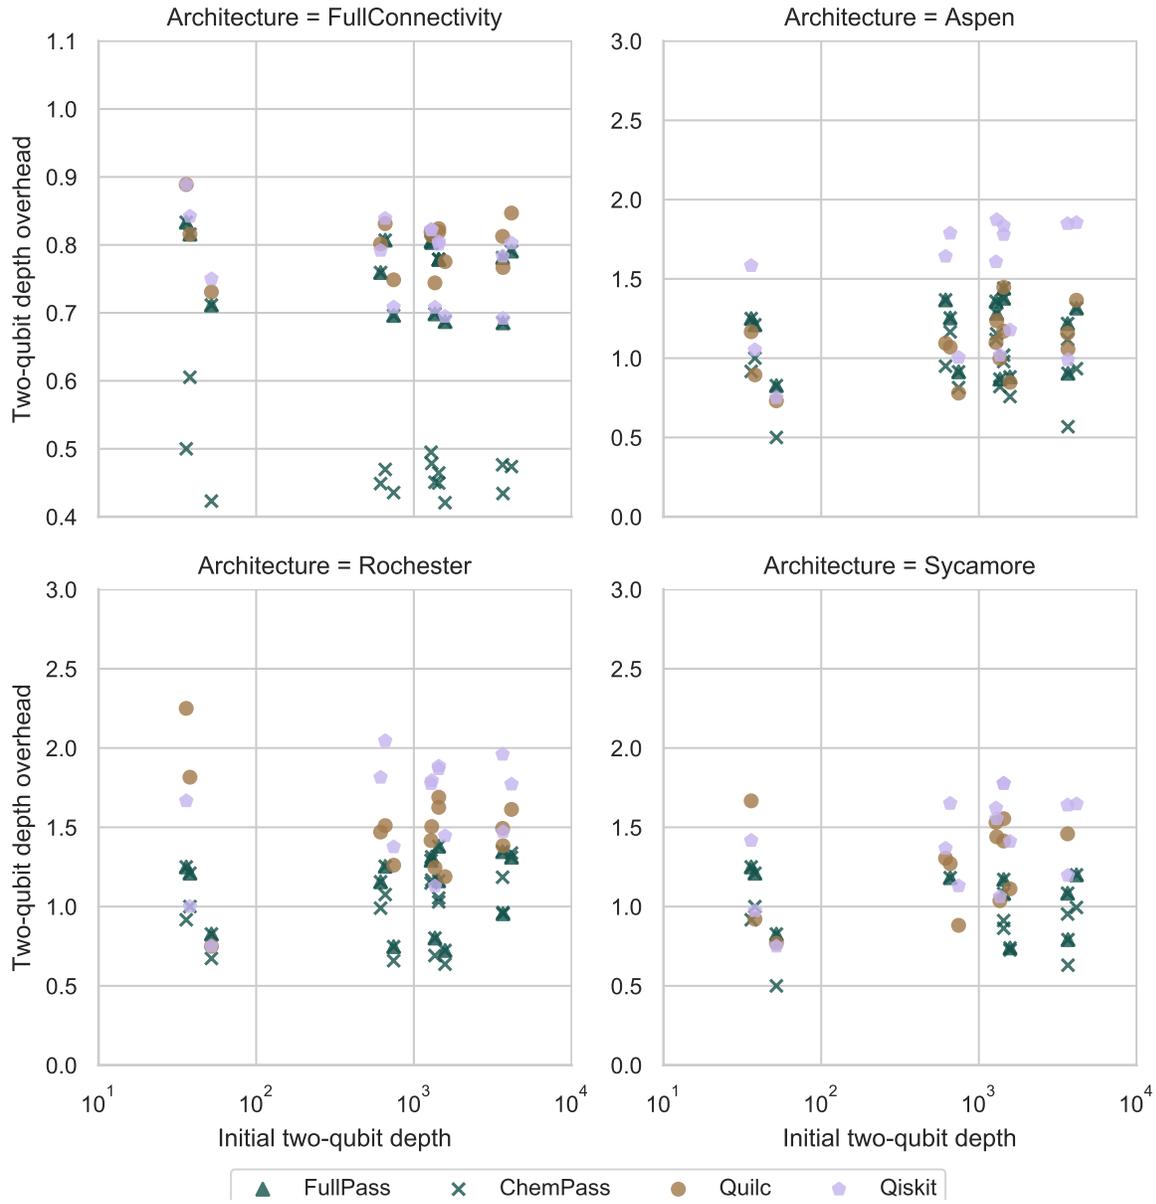
	FullConnectivity	Aspen	Rochester	Sycamore
FullPass	0.937 ± 0.011	1.844 ± 0.040	1.924 ± 0.041	1.780 ± 0.038
Quilc	0.950 ± 0.010	1.312 ± 0.027	2.418 ± 0.063	2.223 ± 0.053
Qiskit	0.955 ± 0.010	2.631 ± 0.065	2.748 ± 0.081	2.473 ± 0.045

Figure 19: Default compilation benchmarks over all benchmark circuits. The multiplicative overhead in two-qubit gate depth from input circuit to output circuit is plotted against input two-qubit depth. The table shows means across the circuit sets and associated standard error. In general, routing induces a larger-than-1 overhead, but for *FullConnectivity* when no routing is required it is 1 or below. ‘FullPass’ refers to the recommended t|ket⟩ routine, consisting of the FullPeepholeOptimise pass, followed by the corresponding qubit mapping pass, SynthesiseIBM and rebasing to the final gate set.



	FullConnectivity	Rigetti	Rochester	Sycamore
FullPass	0.765 ± 0.014	1.236 ± 0.074	1.184 ± 0.074	1.099 ± 0.070
ChemPass	0.647 ± 0.025	1.176 ± 0.052	1.180 ± 0.052	1.042 ± 0.053
Quilc	0.801 ± 0.013	1.145 ± 0.065	1.614 ± 0.098	1.384 ± 0.097
Qiskit	0.783 ± 0.016	1.877 ± 0.173	2.163 ± 0.180	1.747 ± 0.136

Figure 20: Chemistry-specific compilation benchmarks over the UCCSD test set. The multiplicative overhead in two-qubit gate count from input circuit to output circuit is plotted against input two-qubit gate count. The table shows means across the circuit sets and associated standard error. ‘ChemPass’ refers to application of the t|ket) PauliSimp pass, followed by the ‘FullPass’ routine.



	FullConnectivity	Rigetti	Rochester	Sycamore
FullPass	0.762 ± 0.013	1.165 ± 0.056	1.115 ± 0.061	1.054 ± 0.061
ChemPass	0.615 ± 0.029	1.043 ± 0.044	1.042 ± 0.043	0.949 ± 0.049
Quilc	0.803 ± 0.011	1.074 ± 0.052	1.481 ± 0.085	1.258 ± 0.080
Qiskit	0.783 ± 0.016	1.454 ± 0.104	1.584 ± 0.098	1.398 ± 0.081

Figure 21: Chemistry-specific compilation benchmarks over the UCCSD test set. The multiplicative overhead in two-qubit depth from input circuit to output circuit is plotted against input two-qubit depth. The table shows means across the circuit sets and associated standard error. ‘ChemPass’ refers to application of the $t|ket\rangle$ PauliSimp pass, followed by the ‘FullPass’ routine.

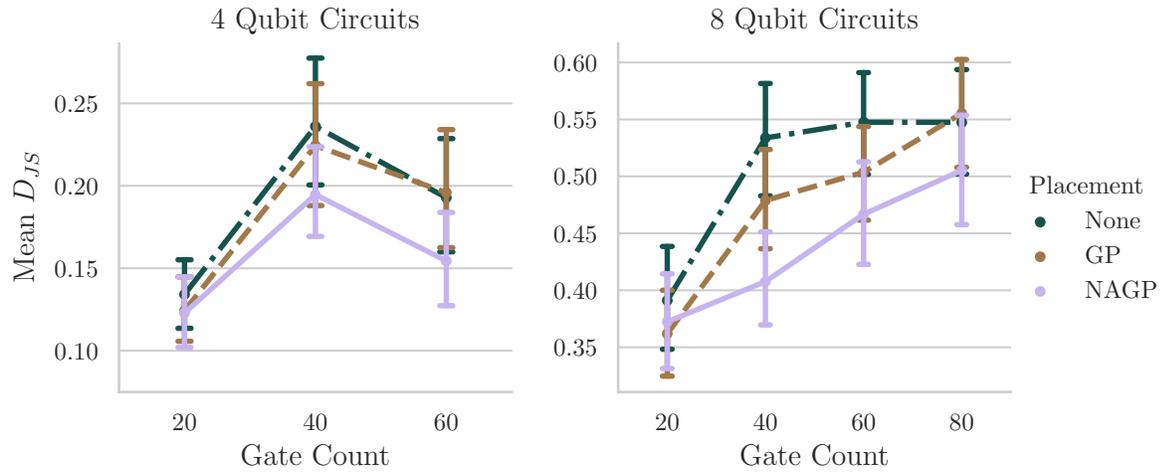


Figure 22: Experimental comparison of placement methods. “None” refers to the case of no initial placement of qubits, “GP” to partial placement via graph placement, and “NAGP” to noise-aware graph placement. Each data point corresponds to a mean Jensen–Shannon divergence D_{JS} of measured distribution from ideal distribution, over the 90 random input circuit samples of the given size, compiled and executed on the device with the three different placement methods.

References

- [1] P. W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM J.Sci.Statist.Comput.*, 26(5), 1997.
- [2] Lov K. Grover. Quantum computers can search arbitrarily large databases by a single query. *Phys. Rev. Lett.*, 79(23):4709–4712, 1997.
- [3] Aram W. Harrow, Avinandan Hassidim, and Seth Lloyd. Quantum algorithm for linear systems of equations. *Phys. Rev. Lett.*, 103:150502, Oct 2009.
- [4] I. M. Georgescu, S. Ashhab, and Franco Nori. Quantum simulation. *Rev. Mod. Phys.*, 2013, arXiv:1308.6253.
- [5] Grace Murray Hopper. The education of a computer. In *Proceedings of the 1952 ACM National Meeting (Pittsburgh)*, ACM '52, pages 243–249, New York, NY, USA, 1952. Association for Computing Machinery.
- [6] Matt Godbolt. Optimizations in C++ compilers. *Queue*, 17(5):69–100, October 2019.
- [7] E. Knill. Conventions for quantum pseudocode. Technical Report LAUR-96-2724, Los Alamos National Laboratory, 1996.
- [8] IBM Research. Qiskit Aqua. <https://qiskit.org/aqua/>.
- [9] Ville Bergholm, Josh Izaac, Maria Schuld, Christian Gogolin, Carsten Blank, Keri McKiernan, and Nathan Killoran. Pennylane: Automatic differentiation of hybrid quantum-classical computations. *arXiv.org*, 2018, arXiv:1811.04968.
- [10] IBM Quantum Experience. <https://quantum-computing.ibm.com>.
- [11] Rigetti Computing. <https://www.rigetti.com/systems>.
- [12] Google Quantum. <https://research.google/teams/applied-science/quantum/>.
- [13] K. Wright, K. M. Beck, S. Debnath, J. M. Amini, Y. Nam, N. Grzesiak, J. S. Chen, N. C. Pisenti, 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, and J. Kim. Benchmarking an 11-qubit quantum computer. *Nature Communications*, 10(1):5464, 2019.
- [14] Honeywell quantum systems. <https://www.honeywell.com/en-us/company/quantum>, 2020.
- [15] Xanadu Hardware. <https://www.xanadu.ai/hardware/>.
- [16] Terry Rudolph. Why i am optimistic about the silicon-photon route to quantum computing. *arXiv preprint*, 2016, arXiv:1607.08535.
- [17] John Preskill. Quantum Computing in the NISQ era and beyond. *Quantum*, 2:79, August 2018.
- [18] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.
- [19] Thomas E O’Brien, Brian Tarasinski, and Barbara M Terhal. Quantum phase estimation of multiple eigenvalues for small-scale (noisy) experiments. *New Journal of Physics*, 21(2):023022, feb 2019, arXiv:1809.09697.
- [20] Alberto Peruzzo, Jarrod McClean, Peter Shadbolt, Man-Hong Yung, Xiao-Qi Zhou, Peter J. Love, Alán Aspuru-Guzik, and Jeremy L. O’Brien. A variational eigenvalue solver on a photonic quantum processor. *Nature Communications*, 5, 07 2014, arXiv:1304.3061.
- [21] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. A quantum approximate optimization algorithm. *arXiv.org*, 2014, arXiv:1411.4028.
- [22] Abhinav Kandala, Kristan Temme, Antonio D. Corcoles, Antonio Mezzacapo, Jerry M. Chow, and Jay M. Gambetta. Error mitigation extends the computational reach of a noisy quantum processor. *Nature*, 567:491–495, 2019, arxiv:1805.04492.
- [23] Joel J. Wallman and Joseph Emerson. Noise tailoring for scalable quantum computation via randomized compiling. *Phys. Rev. A*, 94:052325, Nov 2016, arXiv:1512.01098.
- [24] Bibek Pokharel, Namit Anand, Benjamin Fortman, and Daniel A. Lidar. Demonstration of fidelity improvement using dynamical decoupling with superconducting qubits. *Phys. Rev. Lett.*, 121:220502, Nov 2018, arXiv:1807.08768.

- [25] Harrison Ball, Michael J. Biercuk, Andre Carvalho, Rajib Chakravorty, Jiayin Chen, Leonardo A. de Castro, Steven Gore, David Hover, Michael Hush, Per J. Liebermann, Robert Love, Kevin Nguyen, Viktor S. Perunicic, Harry J. Slatyer, Claire Edmunds, Virginia Frey, Cornelius Hempel, and Alistair Milne. Software tools for quantum control: Improving quantum computer performance through noise and error suppression. *arXiv.org*, 2020, arXiv:2001.04060.
- [26] Thomas Häner, Damian S. Steiger, Krysta Svore, and Matthias Troyer. A software methodology for compiling quantum programs. *arXiv.org*, (1604.01401), 2016, arXiv:1604.01401.
- [27] Alexander S. Green, Peter LeFanu Lumsdaine, Neil J. Ross, Peter Selinger, and Benoît Valiron. Quipper: A scalable quantum programming language. In *Programming language design and implementation (PLDI'13)*, volume 48 of *ACM SIGPLAN Notices*, pages 333–342, 2013, arXiv:1304.3390.
- [28] Ali JavadiAbhari, Shruti Patil, Daniel Kudrow, Jeff Heckey, Alexey Lvov, Frederic T. Chong, and Margaret Martonosi. Scaffcc: Scalable compilation and analysis of quantum programs. *Parallel Computing*, 45:2 – 17, 2015, arXiv:1507.01902. Computing Frontiers 2014: Best Papers.
- [29] Krysta Svore, Alan Geller, Matthias Troyer, John Azariah, Christopher Granade, Bettina Heim, Vadym Kliuchnikov, Mariia Mykhailova, Andres Paz, and Martin Roetteler. Q#: Enabling scalable quantum computing and development with a high-level dsl. In *Proceedings of the Real World Domain Specific Languages Workshop 2018*, RWDSL2018, pages 7:1–7:10, New York, NY, USA, 2018. ACM.
- [30] Nathan Killoran, Josh Izaac, Nicolás Quesada, Ville Bergholm, Matthew Amy, and Christian Weedbrook. Strawberry Fields: A Software Platform for Photonic Quantum Computing. *Quantum*, 3:129, Mar 2019.
- [31] Rigetti Computing. Forest - rigetti. <http://rigetti.com/forest>.
- [32] Héctor Abraham, Ismail Yunus Akhalwaya, Gadi Aleksandrowicz, Thomas Alexander, Gadi Alexandrowics, Eli Arbel, Abraham Asfaw, Carlos Azaustre, AzizNgoueya, Panagiotis Barkoutsos, George Barron, Luciano Bello, Yael Ben-Haim, Daniel Bevenius, Lev S. Bishop, Samuel Bosch, David Bucher, CZ, Fran Cabrera, Padraic Calpin, Lauren Capelluto, Jorge Carballo, Ginés Carrascal, Adrian Chen, Chun-Fu Chen, Richard Chen, Jerry M. Chow, Christian Claus, Christian Clauss, Abigail J. Cross, Andrew W. Cross, Simon Cross, Juan Cruz-Benito, Cryoris, Chris Culver, Antonio D. Córcoles-Gonzales, Sean Dague, Matthieu Dartiailh, DavideFrr, Abdón Rodríguez Davila, Delton Ding, Eric Drechsler, Drew, Eugene Dumitrescu, Karel Dumon, Ivan Duran, Eric Eastman, Pieter Eendebak, Daniel Egger, Mark Everitt, Paco Martín Fernández, Paco Martín Fernández, Axel Hernández Ferrera, Albert Frisch, Andreas Fuhrer, MELVIN GEORGE, IAN GOULD, Julien Gacon, Gadi, Borja Godoy Gago, Jay M. Gambetta, Luis Garcia, Shelly Garion, Gawel-Kus, Juan Gomez-Mosquera, Salvador de la Puente González, Donny Greenberg, Dmitry Grinko, Wen Guan, John A. Gunnels, Isabel Haide, Ikko Hamamura, Vojtech Havlicek, Joe Hellmers, Łukasz Herok, Stefan Hillmich, Hiroshi Horii, Connor Howington, Shaohan Hu, Wei Hu, Haruki Imai, Takashi Imamichi, Kazuaki Ishizaki, Raban Iten, Toshinari Itoko, Ali Javadi-Abhari, Jessica, Kiran Johns, Naoki Kanazawa, Kang-Bae, Anton Karazeev, Paul Kassebaum, Knabberjoe, Arseny Kovyrshin, Vivek Krishnan, Kevin Krsulich, Gawel Kus, Ryan LaRose, Raphaël Lambert, Joe Latone, Scott Lawrence, Dennis Liu, Peng Liu, Panagiotis Barkoutsos ZRL Mac, Yunho Maeng, Aleksei Malyshev, Jakub Marecek, Manoel Marques, Dolph Mathews, Atsushi Matsuo, Douglas T. McClure, Cameron McGarry, David McKay, Srujan Meesala, Antonio Mezzacapo, Rohit Midha, Zlatko Minev, Michael Duane Mooring, Renier Morales, Niall Moran, Prakash Murali, Jan Müggenburg, David Nadlinger, Giacomo Nannicini, Paul Nation, Yehuda Naveh, Nick-Singstock, Pradeep Niroula, Hassi Norlen, Lee James O’Riordan, Oluwatobi Ogunbayo, Pauline Ollitrault, Steven Oud, Dan Padilha, Hanhee Paik, Simone Perriello, Anna Phan, Marco Pistoia, Alejandro Pozas-iKerstjens, Viktor Prutyaynov, Daniel Puzzuoli, Jesús Pérez, Quintiii, Rudy Raymond, Rafael Martín-Cuevas Redondo, Max Reuter, Diego M. Rodríguez, Mingi Ryu, Tharmashastha SAPV, SamFerracin, Martin Sandberg, Ninad Sathaye, Bruno Schmitt, Chris Schnabel, Travis L. Scholten, Eddie Schoute, Ismael Faro Sertage, Nathan Shammah, Yunong

- Shi, Adenilton Silva, Yukio Siraichi, Iskandar Sitdikov, Seyon Sivarajah, John A. Smolin, Mathias Soeken, Dominik Steenken, Matt Stypulkoski, Hitomi Takahashi, Charles Taylor, Pete T aylour, Soolu Thomas, Mathieu Tillet, Maddy Tod, Enrique de la Torre, Kenso Trabling, Matthew Treinish, TrishaPe, Wes Turner, Yotam Vaknin, Carmen Recio Valcarce, Francois Varchon, Desiree Vogt-Lee, Christophe Vuillot, James Weaver, Rafal Wieczorek, Jonathan A. Wildstrom, Robert Wille, Erick Winston, Jack J. Woehr, Stefan Woerner, Ryan Woo, Christopher J. Wood, Ryan Wood, Stephen Wood, James Wootton, Daniyar Yeralin, Jessie Yu, Christopher Zachow, Laura Zdanski, Zoufal, anedumla, azulehner, bcamorrison, brandhsn, chlorophyll zz, dime10, drholmie, elfrocampeador, faisaldebouni, fanizzamarco, gruu, kanejess, klinvill, kurarr, lerongil, ma5x, merav aharoni, mrossinek, neupat, ordmoj, sethmerkel, strickroman, sumitpuri, tigerjack, toural, willhbang, yang.luh, and yotamvakninibm. Qiskit: An open-source framework for quantum computing, 2019.
- [33] Damian S. Steiger, Thomas Häner, and Matthias Troyer. ProjectQ: An open source software framework for quantum computing. *arXiv.org*, (1612.08091), 2016, arXiv:1612.08091.
- [34] The Cirq Developers. Cirq: A python library for nisq circuits. <https://cirq.readthedocs.io/en/stable/>.
- [35] Alexander McCaskey, Eugene Dumitrescu, Dmitry Lyakh, M. Chen, W. Feng, and Travis Humble. A language and hardware independent approach to quantum–classical computing. *SoftwareX*, 7:245–254, Jan 2018.
- [36] Prakash Murali, Norbert Matthias Linke, Margaret Martonosi, Ali Javadi Abhari, Nhung Hong Nguyen, and Cinthia Huerta Alderete. Full-stack, real-system quantum computer studies: Architectural comparisons and design insights. In *Proc. ICISA 2019*, 2019, arXiv:1905.11349.
- [37] Robert S Smith, Eric C Peterson, Erik J Davis, and Mark G Skilbeck. quilc: An optimizing quil compiler, February 2020.
- [38] Yunseong Nam, Neil J. Ross, Yuan Su, Andrew M. Childs, and Dmitri Maslov. Automated optimization of large quantum circuits with continuous parameters. *npj Quantum Information*, 4(1):23, 2018.
- [39] Davide Venturelli, Minh Binh Do, Bryan O’Gorman, Jeremy Frank, Eleanor G. Rieffel, Kyle E. C. Booth, Thành Nhut Nguyen, P. P. S. Narayan, and Sasha Nanda. Quantum circuit compilation : An emerging application for automated reasoning. In *12th International Scheduling and Planning Application Workshop (SPARK)*, 2019.
- [40] Prakash Murali, Jonathan M. Baker, Ali Javadi Abhari, Frederic T. Chong, and Margaret Martonosi. Noise-adaptive compiler mappings for noisy intermediate-scale quantum computers. *arXiv.org*, 2019, arXiv:1901.11054.
- [41] Prakash Murali, David C. McKay, Margaret Martonosi, and Ali Javadi-Abhari. Software mitigation of crosstalk on noisy intermediate-scale quantum computers. *arXiv*, 2020, 2001.02826.
- [42] Eric C. Peterson, Gavin E. Crooks, and Robert S. Smith. Fixed-depth two-qubit circuits and the monodromy polytope. *arXiv.org*, 2019, 1904.10541.
- [43] Nelson Leung, Mohamed Abdelhafez, Jens Koch, and David Schuster. Speedup for quantum optimal control from automatic differentiation based on graphics processing units. *Phys. Rev. A*, 95:042318, Apr 2017, arXiv:1612.04929.
- [44] Pranav Gokhale, Yongshan Ding, Thomas Propson, Christopher Winkler, Nelson Leung, Yunong Shi, David I. Schuster, Henry Hoffmann, and Frederic T. Chong. Partial compilation of variational algorithms for noisy intermediate-scale quantum machines. In *Proceedings of the 52Nd Annual IEEE/ACM International Symposium on Microarchitecture, MICRO ’52*, pages 266–278, New York, NY, USA, 2019. ACM, arXiv:1909.07522.
- [45] The LLVM compiler infrastructure. <http://www.llvm.org>.
- [46] Andrew W. Cross, Lev S. Bishop, John A. Smolin, and Jay M. Gambetta. Open quantum assembly language. <https://github.com/IBM/qiskit-openqasm/blob/master/spec/qasm2.pdf>, January 2017.
- [47] Aleks Kissinger and John van de Wetering. Pyzx: Large scale automated diagrammatic reasoning.

- arXiv.org*, 2019, arXiv:1904.04735.
- [48] Jarrod R. McClean, Kevin J. Sung, Ian D. Kivlichan, Yudong Cao, Chengyu Dai, E. Schuyler Fried, Craig Gidney, Brendan Gimby, Pranav Gokhale, Thomas Häner, Tarini Hardikar, Vojtěch Havlíček, Oscar Higgott, Cupjin Huang, Josh Izaac, Zhang Jiang, Xinle Liu, Sam McArdle, Matthew Neeley, Thomas O’Brien, Bryan O’Gorman, Isil Ozfidan, Maxwell D. Radin, Jhonathan Romero, Nicholas Rubin, Nicolas P. D. Sawaya, Kanav Setia, Sukin Sim, Damian S. Steiger, Mark Steudtner, Qiming Sun, Wei Sun, Daochen Wang, Fang Zhang, and Ryan Babbush. Openfermion: The electronic structure package for quantum computers. *arXiv.org*, 2017, arXiv:1710.07629.
- [49] Mingyu Sun and Michael R. Geller. Efficient characterization of correlated SPAM errors. *ArXiv.org*, 2018, arXiv:1810.10523.
- [50] W. Wootters and W. Zurek. A single quantum cannot be cloned. *Nature*, 299:802–803, 1982.
- [51] A.K. Pati and S. L. Braunstein. Impossibility of deleting an unknown quantum state. *Nature*, 404:164–165, 2000.
- [52] John C. Baez and Kenny Courser. Structured cospans. *arXiv.org*, 2019, arXiv:1911.04630.
- [53] Ieee standard for floating-point arithmetic. *IEEE Std 754-2019 (Revision of IEEE 754-2008)*, pages 1–84, 2019.
- [54] Dmitri Maslov. Basic circuit compilation techniques for an ion-trap quantum machine. *New Journal of Physics*, 19(023035), 2017, arXiv:1603.07678.
- [55] Vadym Kliuchnikov and Dmitri Maslov. Optimization of clifford circuits. *Phys. Rev. A*, 88:052307, Nov 2013.
- [56] Hartmut Ehrig, Karsten Ehrig, Ulrike Prange, and Gabriele Taentzer. *Fundamentals of Algebraic Graph Transformation*. Monographs in Theoretical Computer Science. Springer Berlin Heidelberg, 2006.
- [57] Richard Mitchell, Jim McKim, and Bertrand Meyer. *Design by Contract, by Example*. 0201634600. Addison Wesley Longman Publishing Co., Inc., USA, 2001.
- [58] Andrew W. Cross, Lev S. Bishop, Sarah Sheldon, Paul D. Nation, and Jay M. Gambetta. Validating quantum computers using randomized model circuits. *arXiv.org*, 2018, arXiv:1811.12926.
- [59] Robin Blume-Kohout and Kevin C. Young. A volumetric framework for quantum computer benchmarks. *arXiv preprint*, 2019, arXiv:1904.05546.
- [60] Alexander Erhard, Joel James Wallman, Lukas Postler, Michael Meth, Roman Stricker, Esteban Adrian Martinez, Philipp Schindler, Thomas Monz, Joseph Emerson, and Rainer Blatt. Characterizing large-scale quantum computers via cycle benchmarking. *arXiv.org*, 2019, arXiv:1902.08543.
- [61] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C. Bardin, Rami Barends, Rupak Biswas, Sergio Boixo, Fernando G. S. L. Brandao, David A. Buell, Brian Burkett, Yu Chen, Zijun Chen, Ben Chiaro, Roberto Collins, William Courtney, Andrew Dunsworth, Edward Farhi, Brooks Foxen, Austin Fowler, Craig Gidney, Marissa Giustina, Rob Graff, Keith Guerin, Steve Habegger, Matthew P. Harrigan, Michael J. Hartmann, Alan Ho, Markus Hoffmann, Trent Huang, Travis S. Humble, Sergei V. Isakov, Evan Jeffrey, Zhang Jiang, Dvir Kafri, Kostyantyn Kechedzhi, Julian Kelly, Paul V. Klimov, Sergey Knysh, Alexander Korotkov, Fedor Kostritsa, David Landhuis, Mike Lindmark, Erik Lucero, Dmitry Lyakh, Salvatore Mandrà, Jarrod R. McClean, Matthew McEwen, Anthony Megrant, Xiao Mi, Kristel Michielsen, Masoud Mohseni, Josh Mutus, Ofer Naaman, Matthew Neeley, Charles Neill, Murphy Yuezhen Niu, Eric Ostby, Andre Petukhov, John C. Platt, Chris Quintana, Eleanor G. Rieffel, Pedram Roushan, Nicholas C. Rubin, Daniel Sank, Kevin J. Satzinger, Vadim Smelyanskiy, Kevin J. Sung, Matthew D. Trevithick, Amit Vainsencher, Benjamin Villalonga, Theodore White, Z. Jamie Yao, Ping Yeh, Adam Zalcman, Hartmut Neven, and John M. Martinis. Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779):505–510, 2019.
- [62] Daniel Gottesman. The heisenberg representation of quantum computers. In S. P. Corney, R. Delbourgo, and P. D. Jarvis, editors, *Proceedings of the XXII International Colloquium on Group Theoretical Methods in Physics*, pages 32–43. International Press, 1999, arXiv:quant-

- ph/9807006.
- [63] S. Aaronson and D Gottesman. Improved simulation of stabilizer circuits. *Phys. Rev. A*, 70(052328), 2004, arXiv:quant-ph/0406196v5.
 - [64] Matthew Amy, Jianxin Chen, and Neil J Ross. A finite presentation of cnot-dihedral operators. *arXiv preprint arXiv:1701.00140*, 2016.
 - [65] Andrew Fagan and Ross Duncan. Optimising clifford circuits with quantomatic. In Peter Selinger and Giulio Chiribella, editors, Proceedings of the 15th International Conference on *Quantum Physics and Logic*, Halifax, Canada, 3-7th June 2018, volume 287 of *Electronic Proceedings in Theoretical Computer Science*, pages 85–105. Open Publishing Association, 2019, arXiv:1901.10114.
 - [66] Peter Selinger. Generators and relations for n-qubit clifford operators. *Logical Methods in Computer Science*, 11(2:10), 2015, arXiv:1310.6813.
 - [67] Navin Khaneja and Steffen J. Glaser. Cartan decomposition of $su(2n)$ and control of spin systems. *Chemical Physics*, 267(1-3):11–23, 2001.
 - [68] M Blaauboer and RL De Visser. An analytical decomposition protocol for optimal implementation of two-qubit entangling gates. *Journal of Physics A: Mathematical and Theoretical*, 41(39):395307, 2008.
 - [69] Guifre Vidal and Christopher M Dawson. Universal quantum circuit for two-qubit transformations with three controlled-not gates. *Physical Review A*, 69(1):010301, 2004.
 - [70] Alexander Cowtan, Silas Dilkes, Ross Duncan, Will Simmons, and Seyon Sivarajah. Phase gadget synthesis for shallow circuits. In *Proceedings of QPL2019 (to appear)*, 2019, arXiv:1906.01734.
 - [71] Bob Coecke and Ross Duncan. Interacting quantum observables: Categorical algebra and diagrammatics. *New J. Phys*, 13(043016), 2011, arXiv:0906.4725.
 - [72] Matthew Amy, Dmitri Maslov, and Michele Mosca. 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):1476–1489, 2014.
 - [73] Ross Duncan, Aleks Kissinger, Simon Perdrix, and John van de Wetering. Graph-theoretic simplification of quantum circuits with the zx-calculus. *arXiv.org*, 2019.
 - [74] Aleks Kissinger and John van de Wetering. Reducing t-count with the zx-calculus. *arXiv.org*, 2019, arXiv:1903.10477.
 - [75] Dmitri Maslov and Martin Roetteler. Shorter stabilizer circuits via bruhat decomposition and quantum circuit transformations. *IEEE Transactions on Information Theory*, (7):4729–4738, 2017, arXiv:1705.09176.
 - [76] D. Maslov, G.W. Dueck, and D.M. Miller. Toffoli network synthesis with templates. *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions*, 24(6):807–817, June 2005.
 - [77] Alexander Cowtan, Silas Dilkes, Ross Duncan, Alexandre Krajenbrink, Will Simmons, and Seyon Sivarajah. On the qubit routing problem. In Wim van Dam and Laura Mancinska, editors, *14th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2019)*, volume 135 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 5:1–5:32, 2019.
 - [78] Andrew M. Childs, Eddie Schoute, and Cem M. Unsal. Circuit transformations for quantum architectures. In Wim van Dam and Laura Mancinska, editors, *14th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2019)*, volume 135 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 3:1–3:24, 2019, arXiv:1902.09102.
 - [79] Alwin Zulehner and Robert Wille. Compiling $su(4)$ quantum circuits to ibm qx architectures. *arXiv.org*, 2018, arXiv:1808.05661.
 - [80] Alwin Zulehner, Alexandru Paler, and Robert Wille. An efficient methodology for mapping quantum circuits to the ibm qx architectures. *arXiv.org*, 2017, arXiv:1712.04722.
 - [81] Swamit S. Tannu and Moinuddin K.Qureshi. A case for variability-aware policies for nisq-era quantum computers. *arXiv.org*, 2018, arXiv:1805.10224.

- [82] Jarrod R. McClean, Mollie E. Kimchi-Schwartz, Jonathan Carter, and Wibe A. de Jong. Hybrid quantum-classical hierarchy for mitigation of decoherence and determination of excited states. *Phys. Rev. A*, 95:042308, Apr 2017, arXiv:1603.05681.
- [83] Jonathan Romero, Ryan Babbush, Jarrod R McClean, Cornelius Hempel, Peter J Love, and Alán Aspuru-Guzik. Strategies for quantum computing molecular energies using the unitary coupled cluster ansatz. *Quantum Science and Technology*, 4(1):014008, oct 2018.
- [84] Mark Steudtner and Stephanie Wehner. Fermion-to-qubit mappings with varying resource requirements for quantum simulation. *New Journal of Physics*, 20(6):063010, jun 2018.
- [85] Andrew M. Childs, Dmitri Maslov, Yunseong Nam, Neil J. Ross, and Yuan Su. Toward the first quantum simulation with quantum speedup. *Proceedings of the National Academy of Sciences*, 115(38):9456, 09 2018.
- [86] Doug Finke. <https://quantumcomputingreport.com/news/ibm-opens-new-quantum-data-center-introduces-new-53-qubit-machine/>. <https://quantumcomputingreport.com/news/ibm-opens-new-quantum-data-center-introduces-new-53-qubit-machine/>.
- [87] Chad Rigetti. <https://medium.com/rigetti/the-rigetti-128-qubit-chip-and-what-it-means-for-quantum-df757d1b71ea>. <https://medium.com/rigetti/the-rigetti-128-qubit-chip-and-what-it-means-for-quantum-df757d1b71ea>.
- [88] IBM Q. <https://www.ibm.com/quantum-computing/>.
- [89] Steven T. Flammia and Yi-Kai Liu. Direct fidelity estimation from few pauli measurements. *Phys. Rev. Lett.*, 106:230501, Jun 2011, arXiv:1104.4695.
- [90] Shin Nishio, Yulu Pan, Takahiko Satoh, Hideharu Amano, and Rodney Van Meter. Extracting success from ibm’s 20-qubit machines using error-aware compilation. *arXiv.org*, 2019.
- [91] A. P. Majtey, P. W. Lamberti, and D. P. Prato. Jensen-shannon divergence as a measure of distinguishability between mixed quantum states. *Phys. Rev. A*, 72:052310, Nov 2005.
- [92] Michael A. Nielsen. Cluster-state quantum computation. *Reports on Mathematical Physics*, 57(1):147 – 161, 2006, arXiv:quant-ph/0504097.
- [93] Ross Duncan. A graphical approach to measurement-based quantum computing. In Chris Heunen, Mehrnoosh Sadrzadeh, and Edward Grefenstette, editors, *Quantum Physics and Linguistics: A Compositional, Diagrammatic Discourse*, chapter 3. Oxford University Press, 2013, arxiv:1203.6242.
- [94] R. Duncan and S. Perdrix. Rewriting measurement-based quantum computations with generalised flow. In S. Abramsky, C. Gavaille, C Kirchner, F. Meyer auf der Heide, and P. G. Spirakis, editors, *Automata, Languages and Programming, 37th International Colloquium, ICALP 2010, Proceedings Part II*, volume 6199 of *Lecture Notes in Computer Science*, pages 285–296. Springer, 2010.
- [95] Robin Harper, Steven T. Flammia, and Joel J. Wallman. Efficient learning of quantum noise. *arXiv.org*, 2019, arXiv:1907.13022.
- [96] Youngkyu Sung, Félix Beaudoin, Leigh M. Norris, Fei Yan, David K. Kim, Jack Y. Qiu, Uwe von Lüpkke, Jonilyn L. Yoder, Terry P. Orlando, Simon Gustavsson, Lorenza Viola, and William D. Oliver. Non-gaussian noise spectroscopy with a superconducting qubit sensor. *Nature Communications*, 10(1):3715, 2019, arXiv:1903.01043.