

Adaptive Compilation of Multi-Level Quantum Operations

Kevin Mato* Martin Ringbauer† Stefan Hillmich‡ Robert Wille*§

*Chair for Design Automation, Technical University of Munich, Munich, Germany

†Institut für Experimentalphysik, University of Innsbruck, Innsbruck, Austria

‡Institute for Integrated Circuits, Johannes Kepler University Linz, Linz, Austria

§Software Competence Center Hagenberg (SCCH) GmbH, Hagenberg, Austria

kevin.mato@tum.de, martin.ringbauer@uibk.ac.at, stefan.hillmich@jku.at, robert.wille@tum.de

<https://www.cda.cit.tum.de/research/quantum/>

Abstract—Quantum computers have the potential to solve some important industrial and scientific problems with greater efficiency than classical computers. While most current realizations focus on two-level qubits, the underlying physics used in most hardware is capable of extending the concepts to a multi-level logic—enabling the use of qudits, which promise higher computational power and lower error rates. Based on a strong theoretical backing and motivated by recent physical accomplishments, this also calls for methods and tools for compiling quantum circuits to those devices. To enable efficient qudit compilation, we introduce the concept of an energy coupling graph for single-qudit systems and provide an adaptive algorithm that leverages this representation for compiling arbitrary unitaries. This leads to significant improvements over the state-of-the-art compilation scheme and, additionally, provides an option to trade-off worst-case costs and run-time. The developed compiler is available via github.com/cda-tum/qudit-compilation under an open-source license.

Index Terms—quantum computing, qudits, compilation

I. INTRODUCTION

Quantum computers promise to solve problems of industrial and scientific interest with better resource and algorithmic efficiency. State-of-the-art quantum devices, subject of current research, are referred to as noisy intermediate-scale quantum (NISQ) devices [1]. These devices host up to hundreds of noisy quantum bits (qubits) and support a limited number of logical operations on these qubits. Such NISQ devices have now been realized in a number of technological platforms, including superconducting circuits [2], trapped ions [3], and single photons [4]. Notably, while these devices almost exclusively work with two-level qubits, the underlying hardware almost always natively supports encoding multiple-valued logic in *qudits* (quantum digits).

The research on qudit design and computation has a long history, with efforts primarily focusing on conceptual studies of algorithms for idealized qudits and their comparison to qubits [5]. Fundamentally, a qudit can not only store and process more information per quantum particle, but also features a richer set of logical operations [6] that make information processing more efficient. As a consequence, it has been shown that qudit computation enables improvements in algorithmic and circuit complexity [5] for a wide class of problems. These results have inspired the proposals for and

demonstration of basic qudit control in numerous physical platforms, from trapped ions [7], [8], to photonic systems [9]–[12], superconducting circuits [13], [14], Rydberg atoms [15], nuclear spins [16], cold atoms [17] and nuclear magnetic resonance systems [18]. More recently, efforts have intensified with the demonstration of a universal qudit quantum processor with error rates that are competitive to qubit systems [8].

One key component to utilize the full potential of qudit hardware, however, is a method to translate abstract quantum algorithms given as unitaries to the elementary logical operations of the specific hardware platform. Due to additional structural constraints, complex cost functions, and a large number of different but valid decompositions of any given quantum operation [8], [19], compilation for qudits is significantly more complex than for binary quantum computers. The state-of-the-art static QR decomposition [8], [19]–[24] fails to exploit the full potential of the underlying technology thus far (discussed and illustrated in Section II-B).

In this work, we develop an efficient adaptive compiler for single qudit operations. To this end, we first introduce a corresponding representation of underlying physical options and constraints (i.e., the available potential) in terms of an *energy coupling graph*. Based on this, we afterwards propose adaptive methods that aims at leveraging this representation for efficiently compiling arbitrary unitaries. In particular, these methods exploit free placement of logical information on the underlying physical information carriers to achieve a significant improvement in complexity of the compiled operations.

To showcase the benefits of the proposed methods, we compiled large sets of unitaries of dimensions 3, 5, and 7 to different target architectures—confirming the flexibility to adapt to arbitrary coupling structures of different multi-level systems as well as experimental constraints. While the used cost-function is inspired by a recently presented qudit hardware platform [8], the presented methods are general and apply to any physical platform and cost function. Finally, a comparison to the state-of-the-art QR decomposition showed an significant improvement in the resulting costs on average; with further potential to trade-off worst-case costs and run-time. The resulting tool has been made available via github.com/cda-tum/qudit-compilation under an open-source license.




Figure 1: Example of local operations on entangling operation.

The remainder of the paper is structured as follows: Section II provides the required background to keep the paper self-contained and motivates the problem. In Section III, we introduce the proposed algorithm alongside the required cost function. Section IV evaluates the proposed approach. Finally, Section V concludes the paper.

II. BACKGROUND AND MOTIVATION

In this work, we provide the basis for efficient compilation of qudit operations. To this end, we first briefly review the basics of quantum information with a particular focus on multi-level quantum logic. Afterwards, we discuss the currently applied scheme of compiling corresponding operations and their drawbacks—providing the motivation of this work.

A. Quantum Information Processing

In classical computations, the primary unit of information is the bit (binary digit), which can exclusively be observed in either the 0 or the 1 state. This concept can easily be generalized to quantum computers, with the *qubit* (quantum bit) as the corresponding unit of information for quantum computations. The crucial difference to the classical case, however, is that qubits can be in any linear combination of $|0\rangle$ and $|1\rangle$ (using Dirac’s bra-ket notation).

Qubits are usually constructed by restricting the natural multi-level structure of the underlying physical carriers of quantum information. These systems, therefore, natively support multi-level logic with the fundamental unit of information termed a *qudit* (quantum digit). A qudit is the quantum equivalent of a d -ary digit with $d \geq 2$, whose state can be described as a vector in the d -dimensional Hilbert space \mathcal{H}_d . The state of a qudit can thus be written as a linear combination $|\psi\rangle = \alpha_0 \cdot |0\rangle + \alpha_1 \cdot |1\rangle + \dots + \alpha_{d-1} \cdot |d-1\rangle$, or simplified as vector $|\psi\rangle = [\alpha_0 \ \alpha_1 \ \dots \ \alpha_{d-1}]^T$, where $\alpha_i \in \mathbb{C}$ are the amplitudes relative to the orthonormal basis of the Hilbert space—given by the vectors $|0\rangle, |1\rangle, |2\rangle, \dots, |d-1\rangle$. The squared magnitude of an amplitude $|\alpha_i|^2$ defines the probability with which the corresponding basis state i will be observed when measuring the qudit. Since the probabilities have to add up to 1, the amplitudes have to satisfy $\sum_{i=0}^{d-1} |\alpha_i|^2 = 1$.

Example 1: Consider a system of one qudit with only three energy levels (also referred to as *qutrit*). The quantum state $|\psi\rangle = \sqrt{1/3} \cdot |0\rangle + \sqrt{1/3} \cdot |1\rangle + \sqrt{1/3} \cdot |2\rangle$ is a valid state with equal probability of measuring each basis. Equivalently, the quantum state may be represented as vector $\sqrt{1/3} \cdot [1 \ 1 \ 1]^T$.

Two key properties that distinguish quantum computing from classical computing are superposition and entanglement. A qudit is said to be in a *superposition* of states in a given basis when at least two amplitudes are non-zero relative to

this basis. *Entanglement*, on the other hand, describes a form of superposition born from interactions in multi-qudit systems. Entanglement is a powerful form of quantum correlation, where the quantum information is encoded in the state of the whole system and cannot be extracted from the individual qudits anymore. The state of a single d -level qudit system can be manipulated by operations which are represented in terms of $d \times d$ -dimensional unitary matrices U , i.e., matrices that satisfy $U^\dagger U = U U^\dagger = I$. The state after the application of U can be determined by multiplying the corresponding input state from the left with the matrix U .

Example 2: Consider a three-level qudit (i.e., a qutrit) initially in the state $|0\rangle$. Applying the Hadamard operation H to it yields the output state shown before in Example 1, i.e.,

$$H \cdot |0\rangle = \frac{1}{\sqrt{3}} \begin{bmatrix} 1 & 1 & 1 \\ 1 & e^{\frac{2\pi}{3}} & e^{-\frac{2\pi}{3}} \\ 1 & e^{-\frac{2\pi}{3}} & e^{\frac{2\pi}{3}} \end{bmatrix} \cdot \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} = \frac{1}{\sqrt{3}} \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}.$$

B. Compiling Multi-Level Quantum Operations

Using the basics of quantum information processing as reviewed in the previous section, arbitrary (multi-level) quantum operations can be defined. However, the elementary operations available in state-of-the-art quantum computing hardware typically couple only two levels and may be subject to additional physical or practical constraints. As a consequence, only a small set of single-qudit operations are available natively, while the vast majority require compilation into elementary operations. The goal of a good compiler is then to not only realize the desired computations, but also to do so in the most efficient way by fully exploiting the potential of the available technology. Solutions designed for compiling unitaries on qubit systems are generally directed towards scalability issues and error mitigation. In a first step, the focus of compilation lies on single-qudit operations. The reasons for this are three-fold:

- 1) One of the main advantages of qudit systems is that complicated qubit entangling operations are traded for simpler qudit local operations. A powerful local compiler is critical for exploiting this potential.
- 2) Qudit systems present the possibility of performing computations between quantum systems of different dimensions and different characteristics [9]. This is in contrast to qubit systems where all the units are computationally identical to each other.
- 3) Compilation of entangling operations relies heavily on the use of local unitaries. Hence, a local unitary compiler, such as the one we present, is a prerequisite for compiling general multi-qudit unitaries.




Figure 2: Comparison between QR decomposition (3 steps) and optimal solution (1 step) to rotate between $|0\rangle$ and $|3\rangle$

Example 3: Figure 1 shows an example of an entangling operation that performs a controlled arbitrary rotation; there is a total of four unitaries and one entangling gate, in this case the Mølmer-Sørensen (MS) gate [25]. To ensure correct behavior of the MS gate as a *CNOT* gate, it requires additional local gates on each one of the qudits involved in the operation, applied just after the native operation. First, $d - 1$ operations per qudit are required to correct intrinsic phase shifts [8]. Second, another 4 local gates turn the corrected MS gate into the desired *CNOT* gate, resulting in a total of 10 local operations for compiling one entangling gate on two ququarts (4 levels). Evidently, the ratio of local gates per entangling gates is rather high and not being able to compile efficiently each one of the local unitaries is an obstacle to scalability in this type of systems once full applications are deployed.

The state-of-the-art qudit compilation [8], [19]–[24] relies on a straight-forward *QR decomposition algorithm*, which decomposes a matrix into an orthogonal matrix and an upper triangular matrix. In such an algorithm, the sequence of elementary two-level operations is fixed a priori and only the angles of the rotations are calculated depending on the target unitary. Importantly, this method does not take care of particular restrictions encountered in real-world systems, for example the impossibility of directly performing certain rotations and shorter paths for routing two levels close to each other. An arbitrary operation will therefore in general have a better decomposition than the one generated by the fixed sequence and, hence, the blind application of more operations could lead to greater error-rates. Overall, the state of the art does not fully exploit the potential of current multi-level architectures—as also illustrated by the following example.

Example 4: Consider a target unitary U that rotates between $|0\rangle$ and $|3\rangle$. The QR decomposition will generate three valid operations to realize the rotation, namely U will be decomposed into $R_{2,3}, R_{1,2}, R_{0,1}$. Figure 2a illustrates the fixed sequence.

Now, given the connectivity in Figure 2, a shorter (and optimal) decomposition $R_{0,3}$ exists, that immediately utilizes the existing connection, but is not found by the QR approach. Figure 2b illustrates the optimal decomposition. In comparison, the fixed sequence requires three reordering pulses, while

the optimal sequence will be composed of only one gate—highlighting the possible improvements over the state of the art even in small examples.

The objective of this work is to develop a method that is adaptive and can be generalized to arbitrary coupling structures. However, due to physical constraints, every technology has restrictions which have to be taken into account during compilation. In order to make these restrictions transparent, we propose a dedicated representation of options as well as constraints in terms of an *energy coupling graph* as well as an adaptive algorithm that enables us to control a number of trade-offs due to the technology that implements the qudit logic, as well as the influence of the sources of error. Compiling in this work is then viewed as a multi-tiered procedure that requires several techniques in order to apply a single arbitrary operation as efficiently and reliably as possible. Moreover, an adaptive algorithm often leads to decompositions that reduce the cost of application for a given unitary, or a sequence of them.

III. ADAPTIVE COMPILATION OF MULTI-LEVEL QUANTUM OPERATIONS

Based on the general idea sketched in Section II-B, this section provides a detailed description of the proposed approach. To this end, we first introduce the concept of an *energy coupling graph* that describes the possible physical connections in the architecture as well as the corresponding subset used for the actual compilation and, by this, provides a representation of both, the used coupling as well as the available potential. Based on that, an adaptive decomposition approach is described which explicitly considers (is adaptive) to the potential of underlying technology when compiling a given unitary.

A. Energy Coupling Graph

The way information is encoded in the qudits (i.e., what energy levels represent what basis state) has a major impact on the quality of the resulting compilation result. Various technologies have been employed to implement quantum operations [7]–[18], which leverage different physical platforms with different degrees of freedom for encoding information. By nature, they inherently use multi-level physical structures to represent their state, albeit currently only two of these levels are commonly used. The construction of a representation of these levels is compulsory for understanding how to efficiently access them, and relate to each other for performing operations.

The connections between physical levels are not arbitrary, but heavily dependent on the technology. To capture the possible connections, we use the term *coupling* for levels of a qudit in analogy to the coupling describing possible interactions between qubits in two-level quantum computers. More precisely, a coupling between levels of a qudit is the possibility of transitioning from a certain level to a second one, given an impulse to the system. Furthermore, we model the possible couplings between levels with a graph—which we refer to as *energy coupling graph*. The restrictions imposed by the energy coupling graph are correspondingly referred to as




Figure 3: A qutrit energy coupling graph

energy coupling constraints. The term is inspired by the fact that the nodes of the graph represent different energy states within the physical information carrier. The transitions can be mathematically represented with unitary matrix operators and associated to corresponding gates. In terms of the two-level Bloch sphere, an operation is a rotation on the sphere with the considered energy levels at the poles. By that, the resulting energy coupling graph represents possible mapping of states to levels as well as correspondingly resulting constraints as provided by recent physical realizations [19].

Example 5: Figure 3 visualizes the proposed energy coupling graph for a single three-level qudit, or qutrit. The black horizontal bars are the physical levels provided by the technology, and the black dotted lines show the theoretically possible couplings. The blue rectangles are the logical nodes of the graph mapped to a subset of the physical levels, and the blue lines are the edges of the constituted graph that are built on the theoretically available coupling. The blue lines represent the physically feasible transitions, or an efficient subset of them.

With the energy coupling graph introduced, we can graphically inspect the three main differences between the theoretically possible connectivity in a qudit and the realistic connections in a quantum system. Given the graph of all the possible levels available from the technology, we start by mapping the logical states to the physical levels. Typically, not all the physical levels in the energy coupling graph are used to encode logical states—some of them may be used for ancillary tasks, such as routing two distant logic states or as a cache memory, temporarily storing information. In theory, i.e., in the ideal qudit, there are transitions in an all-to-all fashion between all the logic states. However, in the non-ideal real world, only certain physical transitions can be effectively used—enabling only certain rotations in the qudit state-space. Additionally, there is a chance that logically contiguous states are placed on distant nodes of the graph, and it may happen when optimizing some particular applications on the qudit architecture. In this work, we lay the foundations for a new layer of abstraction that works as an interface between a physical qudit and a corresponding compiler.

Example 6: Consider Figure 4a showing a realistic qudit, which has three main differences compared to the ideal case depicted in Figure 4b.

- Figure 4a has fewer connections than the ideal case, which may be due to physical (e.g. atomic selection rules), or practical constraints.




Figure 4: A realistic qudit, limited by physical complexity, compared to an ideal qudit, inspired by Ref. [8]

- In Figure 4a, the logic states are not placed in order, for example the logic state $|2\rangle$ is connected to $|0\rangle$, but not $|1\rangle$ and $|3\rangle$.
- Both the graphs are connected, but Figure 4a has an ancillary state $|a\rangle$ to ensure connectivity.

The energy coupling graph as proposed offers designers and automatic tools a means to properly reflect the potential as well as the restrictions of the underlying technology. In the following, the energy coupling graph is used to guide the compilation.

B. Decomposition

With the energy coupling graph introduced in the previous section, the next step in the multi-tiered compilation process is a *decomposition* of the target unitary U . In the scope of this work, the decomposition algorithm is intended to output a sequence V_i of two-level rotations between adjacent physical levels in the energy coupling graph and a diagonal matrix Θ of arbitrary phases [19] as in Eq. (1).

$$U = V_k \cdot V_{k-1} \cdot \dots \cdot V_1 \cdot \Theta \quad (1)$$

The objective of the algorithm is to output a sequence of operations that realize the target unitary U , taking into account the capabilities and restrictions of the underlying technology.

Example 7: Consider performing a rotation between two levels, that are not directly connected. Some naive solutions for improving the flexibility of the QR decomposition to different graphs, could be either to permute the unitary to match the encoding on the graph, with a reordered basis set *or* to fix the encoding and add single-qudit reordering pulses to the sequence to bring logic states in the energy coupling graph closer, perform the operation and then bringing them back. Both variants work in a *static* fashion: potentially introducing unnecessary operations and, thus, unnecessary costs.

To alleviate the drawbacks of static decomposition approaches, we propose an *adaptive algorithm* based on a recursive depth-first search and using a tree structure to keep track of the exploration of the possible decomposition. An overview is provided in Algorithm 1. By adaptively changing the energy coupling graph, this algorithm avoids the problems the static solutions in Example 7 have. The algorithm starts by allocating only the root of the tree, where the only information stored is the cost limits imposed by the fixed QR decomposition,

the target unitary to compile and the initial energy coupling graph. The cost limits for the adaptive approach are the costs of the QR decomposition applied on the same target unitary, and the summed costs of each single rotation of the adaptive decomposition will not exceed these limits. The complexity for calculating this cost bound is asymptotically quadratic in the number of dimensions. In every step, the algorithm applies a two-level Givens-rotation [8], [21] to the unitary and checks if the resulting matrix is diagonal. If it is diagonal, the algorithm returns the nodes in the tree that are on the path of the best decomposition and terminates. Otherwise, the code enters three nested loops: one over all the columns, and two over indices that keep track of two different rows, where the innermost loop's index is always greater than the first loop over the rows. At every step of the execution of these loops the entries of the matrix indexed by the column counter and the two rows indices are selected and they are compared against two constants (indicated as 0 and $threshold$). Since the angle of rotation is calculated as a ratio between the two entries and the phase is calculated similarly with a ratio on the two components of each of these complex numbers, the thresholds guarantee that the coefficients will not go to infinity due to the usage of floating point numbers and possible division by a number close to zero. Moreover, the threshold should be chosen in relation to the experimental performance of the fundamental two-level couplings, such that the resulting parameter fluctuations are negligible compared to the inherent imperfections of the quantum operations. Once the integrity of the entries is verified, the angle and phase of rotation are calculated as shown in the pseudo-code. These two parameters are then used to create the Givens-rotations of the decomposition as elementary blocks. These matrices represent the rotation on the Bloch sphere of two energy levels, but they are also the rotation of a subspace of the larger Hilbert space of the qudit.

The rotation between two levels i and j are expressed as

$$R_{i,j}(\theta, \phi) = \exp\left(\frac{-i\theta}{2}\left(\cos(\phi)\sigma_x^{i,j} + \sin(\phi)\sigma_y^{i,j}\right)\right) \quad (2)$$

where $\sigma_{x,y}$ are the Pauli- $\{X, Y\}$ matrices, describing the two-level coupling, θ is the rotation angle, and ϕ is the phase of the rotation.

The cost of each rotation is calculated by a method that takes the necessary routing of the states on the graph of the qudit into account, as detailed in Section III-E. Additionally, the reordering sequence that enables the rotation, and the modified graph are returned alongside the calculated cost. If the total current costs of the decomposition is smaller than the cost limit, the rotations, the cost, the new unitary rotated, and the new graph are saved as new child of the current node in the tree. The function recursively calls itself with each child of the current node as parameter. As the implemented method is in the category of backtracking algorithms, the worst case complexity is asymptotically exponential [26]. However, the experimental results show that a suitable choice of the cost function allows the algorithm to perform similarly to the fixed QR decomposition.

Algorithm 1 Adaptive Decomposition

Require: Tree.root: n , U , Energy coupling Graph: g , Cost Limit: cl

Ensure: Decomposition, Best Cost

```

if checkDiagonal( $n.U$ ) then return
end if
father =  $n$ 
for  $c$  do
  for  $r \geq c$  do
    for  $r_2 > r$  do
      if  $U_{r,c} \neq 0$  and  $U_{r_2,c} > threshold$  then
         $\theta = 2 \cdot \arctan[|U_{r_2,c}/U_{r,c}|]$ 
         $\phi = -(\pi/2 + \arg[U_{r,c}] - \arg[U_{r_2,c}])$ 
         $c', \pi gates, g' = Cost(U, g, R)$ 
         $U' = R_{r,r_2}(\theta, \phi) \cdot U$ 
        if  $father.cost + c' < cl$  then
           $father.addNode(gates, U', g', c', cl)$ 
        end if
      end if
    end for
  end for
end for
for child in father.children do
  Adaptive(child)
end for

```

C. Satisfying Energy Coupling Constraints

The decomposition from the previous section led to a sequence of operations which are elementary on the respective technology, but not necessarily in-line with the level-to-state mapping, i.e., transitions between states might not be immediately possible. This problem of satisfying the energy coupling constraints is addressed by inserting reordering gates into the operation sequence and introducing rules for correct patterned sequences.

1) *Routing*: The decomposition is driven by the cost for the possible logical rotations on the transition of two logical states. Here, the energy coupling graph is used to track the position of each logical state. Consider a Givens-rotation on the logic states i and j . The returned routing sequence is a list of reordering gates that will bring the logical state j adjacent to the logical state i in the energy coupling graph. The reordering gates are constructed as Givens-rotations with default values $\theta = \pi, \phi = -\pi/2$ and indices of rotation depending on the nodes on the shortest path from i to j . The energy coupling graph is adapted accordingly. Algorithm 1 takes the energy coupling constraints into account.

2) *Graph rules*: Another challenge arises from the fact that the reordering pulses are not permutation matrices, but carry additional sign flips. This is a result of the mathematical form of the generators of the primitive physical rotations. In the following we outline a set of simple rules that have to be applied to the reordered gate sequence. First, rotation matrices are defined to rotate from a lower level to a higher level. If

during compilation we encounter a situation where this ordering is reversed, this leads to an inversion of the sign of the phase ϕ for that operation. Second, whenever a gate in the sequence is preceded by a reordering pulse whose higher level coincides with any of the levels of the gate in question, the sign of θ is flipped. The last rule leverages a particular feature of the graph, where each node is able to store a phase value that the physical level has accumulated during computation. The feature is useful for solving two challenges: the first one arises when compiling several unitaries in a row where the Z -rotations that are not expressed in the form of a gate can be just recorded in the graph's nodes, allowing for even shorter sequence and avoiding the problem of propagating these gates explicitly. The second challenge is unique to the adaptive algorithm where routing sequences are not uncomputed as in the QR decomposition case. Consequently, the sign flips from the reordering pulses do not cancel and need to be tracked explicitly as an accumulation of π phases due to the routing gates. Hence, the third rule adds to the the phase value of a gate the phase stored on the higher energy node and subtracts the phase stored on the lower energy node.

3) *Ancilla states*: A further possibility enabled by the proposed approach is the inclusion of ancillary states in the graph. Use of ancillary states is known to enable improvements in gate complexity [9] and simplify the application of established quantum gates in qudit Hilbert spaces [8]. After a state is declared as ancillary, the algorithm tracks them like regular states, but gains the option of exploiting them for routing purposes or caching.

Example 8: Consider the matrix in Eq. (3), which is composed of two different portions. The upper-left part is the Hadamard gate that we want to compile on the first 3 logic levels. The lower-right part is the matrix that will operate on two ancillary states. It is therefore possible to compile two separate operations in one unitary.

$$\begin{bmatrix} 1 & 1 & 1 & 0 & 0 \\ 1 & e^{\frac{2\pi}{3}} & e^{-\frac{2\pi}{3}} & 0 & 0 \\ 1 & e^{-\frac{2\pi}{3}} & e^{\frac{2\pi}{3}} & 0 & 0 \\ 0 & 0 & 0 & a_{11} & a_{12} \\ 0 & 0 & 0 & a_{21} & a_{22} \end{bmatrix} \quad (3)$$

D. Phase Shift Propagation

After decomposition and insertion of reordering gates from the previous sections, the phase shifts have to be taken into account. The phase shifts are encoded in the diagonal matrix Θ (see Eq. (1)) and have to be decomposed into individual phase shifts on single levels to ensure correctness of the overall decomposition, albeit the phase shift are not physically realized. These individual phase-shifts are formulated as diagonal matrices, with only the i -th entry on the diagonal set to a phase shift of ϕ of the logic state, and 1 otherwise, with $i \in \{0, \dots, d-1\}$.

Example 9: Consider a three-level qudit system and a rotation of ϕ on the second entry on the diagonal. The corresponding matrix is shown in Eq. (4).

$$R_{Z,1}(\phi) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & e^{i\phi} & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad (4)$$

These phase shifts can be introduced at zero cost in the sequence, since they can be propagated to the end of the circuit, where a phase shift is then negligible, as changes in phase immediately before a measurement do not change the outcome probabilities. We refer to these rotations as “virtual Z -gates”¹ since they are not executed on the quantum computer. The commutation of phase gates and Givens-rotations is exemplified in Eq. (5).

$$M \cdot R_{0,1}(\theta, \alpha) = R_{0,1}(\theta, \alpha - \phi + \gamma) \cdot M \quad (5)$$

with $M = \begin{bmatrix} e^{i\phi} & 0 & 0 \\ 0 & e^{i\gamma} & 0 \\ 0 & 0 & e^{i\delta} \end{bmatrix}$

The equation is given for a rotation on the coupling $R_{0,1}$ but it can be generalized to any coupling. The application of the commutation can lead to a linear cost reduction, but is potentially limited once entangling gates are introduced into the circuit.

E. Cost function

The previous sections described the steps to decompose the target unitary, satisfy the energy coupling constraints, and propagate the phase shifts. The decisions made in these steps are guided by a cost function, that is provided in this section. While this cost function is inspired by the trapped-ion system recently presented in [8] and the numerical results derived from benchmarking the platform in the same reference, it is primarily determined by considerations that are found in a similar form also in other technologies. More precisely:

- The cost of each rotation scales linearly with the rotation angle.
- Each two-level coupling is calibrated for a fixed rotation angle, typically $\pi/2$. Angles that differ from $\pi/2$ will incur additional cost due to non-ideal performance of physical devices. We assume this cost to increase linearly with the distance from the calibrated value.
- The system is initialized almost perfectly in a fixed initial state. Those couplings that do not involve the initial state will incur a larger cost (here we assume a factor, function of the distance of two logic states in the graph), since they require a more complicated multi-step calibration procedure.

Taking into account these effects, we model the cost of the single rotation as

$$C_1 = 10^{-4} \cdot d(i, j) \cdot \left(4\theta + \left| \text{mod}\left(\theta + \frac{1}{4}, \frac{1}{2}\right) - \frac{1}{4} \right| \right) \quad (6)$$

¹The term virtual Z -gate does not refer to the Pauli operation.

Table I: Evaluation of costs for the QR and the proposed (adaptive) decomposition

Dim	Unitaries	QR decomp. [8]			Proposed decomp.		
		min	avg	max	min	avg	max
3	333	8	25.47	140	4	21.69	36
	333	4	21.19	28	4	9.49	12
	333	12	35.89	44	4	14.93	20
5	2985	60	166.34	200	64	113.36	124.63
	2985	36	101.27	128	16	82.86	91.05
	2985	88	240.62	280	16	121.97	133.7
7	6438	132	662.35	756	44	255.58	279.08
	6438	192	518.23	620	48	389.94	425.51
	6438	88	392.35	508	44	330.3	361.51

Cost is multiplied by 10^4 to aid legibility.

for $\theta \geq 1/4$ (in units of π). Here the overall factor 10^{-4} comes from the fundamental cost of a π rotation [8].

Besides the bare cost of logical rotations, there is an additional cost associated with reordering the logical states on the coupling graph, since it is, in general, not fully connected. This is done with a sequence of π -rotations, with a cost again governed by Eq. (6). These two terms then constitute the full *experimental cost* of a decomposition.

IV. EVALUATION

The proposed method of adaptive compilation utilizes the energy coupling graph to find more efficient decompositions. To evaluate the method, we developed an implementation and compared it against the state-of-the-art fixed-sequence algorithm based on the QR factorization [8]. The implementation is available under a free license via github.com/cda-tum/qudit-compilation. To account for different scenarios, we compile sets of random Clifford-unitaries [27], which are central to quantum computing, for dimensions 3, 5 and 7 to three target architectures each. The architectures are different in the number of nodes, connectivity, and mapping of the logical states on the physical levels—resulting in nine different target architectures. Diagonal unitaries are excluded from the sets, as they are decomposable in a sequence of virtual phase-gates and the cost of the decomposition in each algorithm is zero. The evaluation was performed on a computer with an Intel i7-10750H and 32 GB of memory, running Ubuntu 20.04 and Python 3.8.

The results are presented in Table I. The first column gives the dimensionality, followed by the number of considered unitaries in the second column. The next three columns show the costs of the QR decomposition for the considered unitaries, listing the minimal, average, and maximal costs according to the cost function in Section III-E. The same listing on costs are contained in the next three columns for the adaptive compilation.

Comparing the costs in Table I shows a clear advantage of the proposed adaptive algorithm. In every case, the maximal cost of the considered set of unitaries is lower compared to that of the QR decomposition—in some cases there is an improvement of up to a factor of 3. The same conclusion can be drawn

from the average costs as well, which is considerably lower in each set of considered unitaries. Regarding the run-time, each decomposition of a unitary in this evaluation took less than 1 s for the QR decomposition and the adaptive approach, respectively. Hence, by spending the same computational efforts, we get significantly better compilation results.

Moreover, the proposed method can be configured to get even better results. In the evaluations summarized above, we set the cost limit (see Algorithm 1) to 1.1 of the cost of the QR decomposition to have an acceptable run-time for the thousands of unitaries. For a more realistic scenario, where a designer wants to realize an individual unitary, decreasing the cost limit will push the proposed algorithm to get even smaller costs—at the expense of longer run-time. Furthermore, the improvement attainable by the proposed adaptive method of course depends on the unitary that is considered, but also on particularities of the underlying technology. One important factor is how dense the connectivity of the energy levels is, where a denser (and more physically unrealistic) connectivity leads to less improvement on the cost by the adaptive algorithm. With the implementation of the proposed method being publicly available, every designer can accordingly use these features to get better results than the ones summarized in the extensive evaluations covering thousands of unitaries from above.

Overall, the results confirm that, using the proposed method, more efficient decompositions can be determined and, additionally, further option exists to trade-off worst-case costs and run-time—providing a solution that aims at fully exploiting the potential provided by the given technology.

V. CONCLUSION

In this work, we proposed an *adaptive compilation* algorithm to attain efficient decompositions of arbitrary local unitaries for multi-level quantum systems. Compared to binary quantum systems, qudits feature a much richer state-space, allowing for different placements of logical information under consideration of practical constraints, such as the energy coupling constraints. With the constraints and options represented in an *energy coupling graph*, the proposed algorithm exploits the available potential by flexibly adapting the mapping between logical states and physical energy levels to enable a much more efficient compilation. Indeed, turning the argument around, the proposed algorithm provides a way to compare different placements of logical states by means of the average cost to decompose random Clifford unitaries. This flexibility of the proposed algorithm to adapt to realistic constraints in the possible couplings as well as the cost of different rotations enable significant improvements in gate costs.

ACKNOWLEDGEMENTS

This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No. 101001318). It is part of the Munich Quantum Valley, which is supported by the Bavarian state government with funds from the Hightech Agenda Bayern Plus and was partially supported

by the BMK, BMDW, and the State of Upper Austria in the frame of the COMET program (managed by the FFG). This project has received funding from the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 840450.

REFERENCES

- [1] J. Preskill, “Quantum computing in the NISQ era and beyond,” *Quantum*, vol. 2, p. 79, 2018.
- [2] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. S. L. Brandao, *et al.*, “Quantum supremacy using a programmable superconducting processor,” *Nature*, vol. 574, no. 7779, pp. 505–510, 2019.
- [3] I. Pogorelov, T. Feldker, C. D. Marciniak, L. Postler, G. Jacob, O. Krieglsteiner, V. Podlesnic, M. Meth, V. Negnevitsky, *et al.*, “Compact Ion-Trap Quantum Computing Demonstrator,” *PRX Quantum*, vol. 2, p. 020343, 2021.
- [4] F. Flamini, N. Spagnolo, and F. Sciarrino, “Photonic quantum information processing: a review,” *Reports Prog. Phys.*, vol. 82, p. 016001, 2019.
- [5] Y. Wang, Z. Hu, B. C. Sanders, and S. Kais, “Qudits and High-Dimensional Quantum Computing,” *Front. Phys.*, vol. 8, 2020.
- [6] M. Huber and J. I. de Vicente, “Structure of Multidimensional Entanglement in Multipartite Systems,” *Phys. Rev. Lett.*, vol. 110, p. 030501, 2013.
- [7] X. Zhang, M. Um, J. Zhang, S. An, Y. Wang, D.-l. Deng, C. Shen, L.-M. Duan, and K. Kim, “State-Independent Experimental Test of Quantum Contextuality with a Single Trapped Ion,” *Phys. Rev. Lett.*, vol. 110, p. 070401, 2013.
- [8] M. Ringbauer, M. Meth, L. Postler, R. Stricker, R. Blatt, P. Schindler, and T. Monz, *A universal qudit quantum processor with trapped ions*, 2021. arXiv: 2109.06903.
- [9] B. P. Lanyon, M. Barbieri, M. P. Almeida, T. Jennewein, T. C. Ralph, K. J. Resch, G. J. Pryde, J. L. O’Brien, A. Gilchrist, *et al.*, “Simplifying quantum logic using higher-dimensional Hilbert spaces,” *Nature Physics*, vol. 5, pp. 134–140, 2008.
- [10] M. Ringbauer, T. R. Bromley, M. Cianciaruso, L. Lami, W. Y. S. Lau, G. Adesso, A. G. White, A. Fedrizzi, and M. Piani, “Certification and Quantification of Multilevel Quantum Coherence,” *Phys. Rev. X*, vol. 8, p. 041007, 2018.
- [11] X.-M. Hu, Y. Guo, B.-H. Liu, Y.-F. Huang, C.-F. Li, and G.-C. Guo, “Beating the channel capacity limit for superdense coding with entangled ququarts,” *Sci. Adv.*, vol. 4, no. 7, eaat9304, 2018.
- [12] M. Malik, M. Erhard, M. Huber, M. Krenn, R. Fickler, and A. Zeilinger, “Multi-photon entanglement in high dimensions,” *Nature Photonics*, vol. 10, pp. 248–252, 2016.
- [13] M. Kononenko, M. Yurtalan, S. Ren, J. Shi, S. Ashhab, and A. Lupascu, “Characterization of control in a superconducting qutrit using randomized benchmarking,” *Physical Review Research*, vol. 3, no. 4, p. L042007, 2021.
- [14] A. Morvan, V. V. Ramasesh, M. S. Blok, J. M. Kreikebaum, K. O’Brien, L. Chen, B. K. Mitchell, R. K. Naik, D. I. Santiago, *et al.*, “Qutrit Randomized Benchmarking,” *Phys. Rev. Lett.*, vol. 126, p. 210504, 2021.
- [15] J. Ahn, T. Weinacht, and P. Bucksbaum, “Information Storage and Retrieval Through Quantum Phase,” *Science*, vol. 287, no. 5452, pp. 463–465, 2000.
- [16] C. Godfrin, A. Ferhat, R. Ballou, S. Klyatskaya, M. Ruben, W. Wernsdorfer, and F. Balestro, “Operating Quantum States in Single Magnetic Molecules: Implementation of Grover’s Quantum Algorithm,” *Phys. Rev. Lett.*, vol. 119, p. 187702, 2017.
- [17] B. E. Anderson, H. Sosa-Martinez, C. A. Riofrío, I. H. Deutsch, and P. S. Jessen, “Accurate and Robust Unitary Transformations of a High-Dimensional Quantum System,” *Phys. Rev. Lett.*, vol. 114, p. 240401, 2015.
- [18] Z. Gedik, I. A. Silva, B. Çakmak, G. Karpat, E. L. G. Vidoto, D. O. Soares-Pinto, E. R. DeAzevedo, and F. F. Fanchini, “Computational speed-up with a single qudit,” *Sci. Rep.*, vol. 5, p. 14671, 2015.
- [19] P. J. Low, B. M. White, A. A. Cox, M. L. Day, and C. Senko, “Practical trapped-ion protocols for universal qudit-based quantum computing,” *Phys. Rev. Res.*, vol. 2, no. 3, 2020.
- [20] G. K. Brennen, S. S. Bullock, and D. P. O’Leary, *Efficient circuits for exact-universal computations with qudits*, 2005. arXiv: quant-ph/0509161.
- [21] G. Brennen, D. O’Leary, and S. Bullock, “Criteria for exact qudit universality,” *Phys. Rev. A*, vol. 71, no. 5, 2005.
- [22] D. P. O’Leary, G. K. Brennen, and S. S. Bullock, “Parallelism for quantum computation with qudits,” *Phys. Rev. A*, vol. 74, no. 3, 2006.
- [23] S. Bullock, D. O’Leary, and G. Brennen, “Asymptotically optimal quantum circuits for d-level systems,” *Phys. Rev. Lett.*, vol. 94, no. 23, 2005.
- [24] E. O. Kiktenko, A. S. Nikolaeva, P. Xu, G. V. Shlyapnikov, and A. K. Fedorov, “Scalable quantum computing with qudits on a graph,” *Phys. Rev. A*, vol. 101, p. 022304, 2020.
- [25] K. Mølmer and A. Sørensen, “Multiparticle entanglement of hot trapped ions,” *Phys. Rev. Lett.*, vol. 82, pp. 1835–1838, 1999.
- [26] D. M. Nicol, “Expected performance of m-solution backtracking,” *SIAM Journal on Computing*, vol. 17, no. 1, pp. 114–127, 1988.
- [27] S. Clark, “Valence bond solid formalism for d-level one-way quantum computation,” *J. Phys. A*, vol. 39, no. 11, pp. 2701–2721, 2006.