Compression of Qubit Circuits: Mapping to Mixed-Dimensional Quantum Systems

Kevin Mato* Stefan Hillmich†‡ Robert Wille*†
*Chair for Design Automation, Technical University of Munich, Munich, Germany
†Software Competence Center Hagenberg (SCCH) GmbH, Hagenberg, Austria
‡Institute for Integrated Circuits, Johannes Kepler University Linz, Linz, Austria
kevin.mato@tum.de, stefan.hillmich@scch.at, robert.wille@tum.de
https://www.cda.cit.tum.de/research/quantum/

Abstract*—Quantum computers are becoming a reality thanks to the accomplishments made in recent years. The quantum computers available today offer hundreds of qubits but are still limited in the number of operations they can perform before errors accumulate and the quantum state decays. In regard to the error accumulation, non-local operations such as *CX* or *CZ* are main contributors. One promising solution to reduce the number of required non-local operations is to make a more efficient use of the quantum hardware by exploiting the inherent high-dimensional capabilities of quantum systems. In a process called *circuit compression*, non-local operations between qubits are mapped to local operations in *qudits*, i.e., higher-dimensional systems. In this work, we present a strategy for enabling quantum circuit compression with the aim of mapping clusters of qubits in a given circuit to the mixed-dimensional qudits of the target hardware. Further, we discuss the principles of circuit compression as well as the physical structure of qubits and qudits, before introducing a new representation that captures the essence of quantum operations, affecting the different logical levels in the quantum states in nodes and edges of a graph. Based on this, we propose an automated approach for mapping qubit circuits of arbitrary gate sets to mixed-dimensional quantum systems, lowering the number of non-local operations. Empirical evaluations confirm the effectiveness of the proposed approach, reducing the number of non-local operations by up to 50% for almost half of the cases. Finally, the corresponding source code is available freely at github.com/cda-tdum/qudit-compression.

***Index Terms*—quantum computing, circuit compression, qudits**

I. Introduction

The emergence of quantum computers promised a way of solving relevant problems in both industry and academia that are intractable on classical computers. We are currently in the era of *Noisy Intermediate-Scale Quantum* (NISQ) devices [1]. This kind of quantum computer hosts up to several hundreds of qubits and supports a limited set of quantum operations that work on the qubits. There are multiple technologies that may be used to realize quantum hardware, such as superconducting circuits [2], single photons [3], neutral atoms [4], and trapped ions [5]. Despite the impressive scaling, NISQ devices suffer from noise in the qubits and the operations. A key observation for improving the performance is that, while the implementations of quantum computers so far mostly use two-level qubits only, the underlying physics support exploiting with a degree of efficiency arbitrary d logical levels in the form of *qudits* (quantum digits).

Research on qudit design and qudit computation has a long history. So far the efforts are primarily focusing on conceptual studies of algorithms for ideal qudits and their comparison to qubits [6]. A qudit system not only enables storage and processing of more information per quantum particle, but it also features a richer set of logical operations [7] that give the potential to make information processing more efficient. Consequently, the improvements of qudit computation algorithmic and circuit complexity have been shown for a wide class of problems [6]. Additionally, demonstration of basic qudit control has been demonstrated in numerous physical platforms, from trapped ions [8], [9], to photonic systems [10]–[13], superconducting circuits [14], [15], Rydberg atoms [16], nuclear spins [17], cold atoms [18], nuclear magnetic resonance systems [19] and molecular spin [20]. More recently, a demonstration of a universal qudit quantum processor with error rates that are competitive to qubit systems has shown not only conceptual advantages but also practical advantages in the implementations [9].

However, as of today most algorithms are designed with qubits in mind. This trend is doomed to change as recent developments in the field of quantum algorithms have proved that a more natural architecture for implementing complex applications [21], [22] are based on multi-level logic. Furthermore, simulations of models representing fermion-boson interactions, on qubit-qudit based quantum computers, would be a significant step toward real time simulations of quantum electrodynamics and other field theories with continuous or larger symmetry groups [23]–[26], due to more efficient encodings on mixed-dimensional platforms.

Gate decompositions present a reduced complexity on mixed-dimensional systems due to the temporal expansion of the Hilbert space with auxiliary levels. This is similar to the effect of using ancilla qubits, but with a simpler circuit complexity and hardware design [27]. This leaves a large potential to create algorithms and circuits untapped, since exploiting higher dimensions to store and manipulate information leads to smaller circuits that have a higher chance of succeeding due to the smaller noise accumulated.

Notably, just increasing the dimension of all qudits in a quantum circuit still leaves room for improvement [28], [29]. Although, enabling the usage of qudits of different dimensions optimizes, the compactness of the circuit, finally at a competitive error rate and experimental control [9].

An important aspect of realizing the impact of high-mixed-dimensional systems is the need for new automated methods, software frameworks, and theory [30]–[34]. Therefore, in this paper, we propose an efficient approach to optimize non-local operations in qubit circuits (such as *CX* or *CZ*) by mapping them to corresponding circuits that operate on qudits of mixed-dimensions. To this end, we first introduce a method of abstracting the interactions between qubits of a circuit by a graph-based representation. Based on this graph, we propose to create a clustering over the state space of the qubits—grouping together qubits with frequent interactions. The dimensions of the individual qudits are determined under consideration of a user-defined number for the preferred dimensionality that reflects the capability of the targeted hardware.

Overall, the proposed solution will, for the first time, enable circuit compression for mixed-dimensional systems. Evaluations show a drastic reduction in the number of non-local gates that have to be used to realize a given functionality—substantially reducing a main contributor to errors in quantum circuits. The source code used in the evaluation is available under the MIT license at github.com/cda-tdum/qudit-compression as part of the *Munich Quantum Toolkit* (MQT).

The remainder is organized as follows: Section II gives a brief background on quantum computations with qubits and qudits. Section III describes the considered problem, reviews the state of the art, and summarizes the contribution of this paper. Section IV introduces the proposed approach of compressing qubit circuits into qudit circuits in detail. Section V evaluates the proposed approach. Finally, Section VI concludes the paper.

II. Background

In this section, briefly review the basics of quantum information processing with a focus on mixed-dimensional quantum logic.

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 $|\psi\rangle = \alpha_0 \cdot |0\rangle + \alpha_1 \cdot |1\rangle + \ldots + \alpha_{d-1} \cdot |d-1\rangle$, or simplified as vector $|\psi\rangle = \begin{bmatrix} \alpha_0 & \alpha_1 & \cdots & \alpha_{d-1} \end{bmatrix}^{\mathrm{T}}$, where $\alpha_i \in \mathbb{C}$ are the amplitudes relative to the orthonormal basis of the Hilbert space—given by the vectors $|0\rangle, |1\rangle, |2\rangle, ..., |d-1\rangle$. 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* $\sqrt{1/3} \cdot \begin{bmatrix} 1 & 1 & 1 \end{bmatrix}^{\mathrm{T}}$.

In a similar fashion, quantum systems of mixed dimensions *can be constructed. Extending the previous qutrit state by a qubit enables represention of the following state* $|\psi'\rangle = \sqrt{1/3} \cdot |0\rangle|0\rangle + \sqrt{1/3} \cdot |1\rangle|1\rangle + \sqrt{1/3} \cdot |2\rangle|0\rangle$*—equivalently represented by the vector* $\sqrt{1/3} \cdot \begin{bmatrix} 1 & 0 & 0 & 1 & 1 & 0 \end{bmatrix}^{\mathrm{T}}$

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 $U^\dagger U = UU^\dagger = I$. For a quantum circuit consisting of multiple mixed-dimensional qudits, the unitary matrix will be of dimension $\prod_i d_i \times \prod_i d_i$ with d_i denoting the dimension of each qudit. The state after the application of U can be determined by multiplying the corresponding input state from the left with the matrix U.

III. Motivation

In this section, we discuss the concept of circuit compression, which aims at reducing the amount of resources required by a computation. The use case considered in this work is enabling the reduction in number of non-local operations by mapping a given qubit circuit to a mixed-dimensional-qudits architecture—enabling efficient usage of the available resources.

A. Considered Problem

A common way to quantify the cost of a qubit circuit is to count the number of non-local gates—commonly controlled phase-rotations (*CZ*) or controlled negations (*CX*, also referred to as *CNOT*). This is motivated by the higher error rates due to more complicated experimental controls as the number of affected qubits grows. In this regard, *circuit compression* aims at improving the quality of computation by reducing the number of operations in a sequence, with particular focus on non-local operations, and, subsequently, the number of qubits (or qudits) in the circuit.

More precisely, the process of circuit compression combines sets of qubits in the given circuit to qudits of suitable dimension, and translates the local and non-local operations in the given circuit to local multi-level operations in the resulting compressed circuit. Since the compressed circuit may use qudits of varying dimensions, it is also referred to as *mixed-dimensional circuit*. Importantly, the routine preserves the original computation, while optimizing the usage of resources provided by the quantum hardware.

Example 2. *Figure 1 illustrates the process of circuit compression. The input to the routine is a qubit (i.e., two-level) circuit with a multitude of local operations, interleaved with entangling operations between all the qubits. It is possible to observe that the first two qubits from the top interact with two-qubit operations with a high frequency. Accordingly, a* compressed *circuit (as shown on the right-hand side) could*

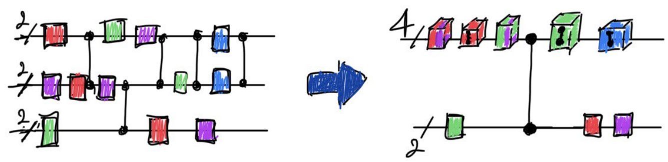

Fig. 1: Example of a 3 qubit circuit compressed into a ququart and qubit circuit. The squares represent qubit local operations, while the cubes represent multi-dimensional local operations. The black lines between the black dots represent non-local operations (*CZ*). The operation of compression shows how going from a qubit circuit to a qudit circuit allows to reduce the number of entangling operation at the cost of local high-dimensional ones.

combine these two qubits into a ququart (i.e., four-level qudit). This enables all the operations applied on the original first two qubits to be translated into local operations on the ququart.

This circuit compression results in a drastic reduction of non-local operations in the resulting circuit: only one non-local operation is required in between the ququart and the remained qubit.

B. State of the Art

The topic of quantum circuit compression has recently moved into the focus of researchers. It has been timidly explored, with outcomes that shed some light on the conceptual understanding on the relation between qubit logic and qudit logic, and on the experimental realization of qubit circuits compression. In the following, we discuss two recent works that represent the state of the art on circuit compression.

In Ref. [35], the authors present a proof-of-concept of a workflow for the compression of qubit circuits into a qudit platform, which comprises transpilation, simulation, and post-processing. Crucially, the presented workflow works on the assumption that qudit circuits are described only by a given set of high-level operations such as multi-controlled Toffoli gates.

The procedure described by Ref. [35] starts with the mapping of qubit circuits to a qudit platform with a fixed dimensionality per qudit in order to reduce qubits and entangling operations. This is accomplished through the enumeration of *all possible mappings*, which is only computable for small instances. The problem of finding mapping methods that can scale polynomially with the size of the qubit circuit remains an open one. The transpilation of non-local operations is in linear time complexity, but this is again only possible due to the restriction of the work to a specific multi-qubit entangling gate set.

Ref. [36] analyses the problem of circuit compression of qubit circuits on qudit circuits, again with qudits of a fixed dimension, with a focus on multi-qubit entangling gates. The emphasis is on the circuit complexity bounds after performing a compression, which is shown to be combinatorial in the dimensions of the considered systems. There is a discussion on the further improvement that can be gained by expressing non-local operations in the qudit space, with specific experimental realizations on different technologies. The introduction of algorithmic methods and heuristics that could define an efficient mapping for quantum circuit compression guided by the physical capabilities of the qudit platform remains an open problem.

Ref. [37] proposes a compression scheme for superconducting ququarts, consisting of a new gate-set for performing qubit-ququart gates and relative mapping procedure to assign qubits to ququarts. Although being complete, the approach is extremely technology specific and the number of levels is constrained to maximum four, while several technology platform can outperform the target architecture studied.

C. Contribution

In this work we investigate quantum circuit compression from arbitrary qubit circuits into mixed-dimensional quantum systems, to enable the potential benefits in terms of improved circuit efficiency. To this end, we present a new set of design automation tools that enable automatic and efficient quantum circuit compression.

The proposed approach of mapping qubit circuits to qudit circuits (quantum circuit compression) consists of two steps:

1) Encoding the qubit logic into a Hilbert space generated by the combination of suitable higher dimensional carriers and
2) translating the non-local qubit operations into either (i) corresponding non-local operations between mixed-dimensional qudits or (ii) local operations within a qudit.

The encoding in the context of the first task can be achieved by constructing the tensor product of the single qudits' state spaces. The crucial part of this step is to find suitable dimensions for the individual qudits. We propose an efficient algorithm to cluster the qubits in the given circuit to reduce the number of non-local operations between the clusters that will subsequently compress into qudits. The mapping method improves the efficiency of compilation algorithms by finding an efficient logic encoding, that will lead to shorter circuits compared to a naive one, such as the ladder-type coupling [31]. Further, the mixed dimensionality avoids wasting efforts in controlling information not necessary for the computation.

The second step encompasses the translation of the non-local qubit operations into genuine entangling two-qudit operations. This step is important to ensure that the resulting qudit circuit can be executed on a physical quantum computer with the appropriate non-local operations in the higher dimensions. In this regard, while the mapping routine is going to be a necessary and useful preliminary step, a dedicated compiler for high and mixed-dimensional quantum systems is the suitable candidate to produce close-to-optimal results. The simple translation of the qubit entangling operations to qudit ones would just require the construction of correct pulses by optimization, however the strength of mixed-dimensional

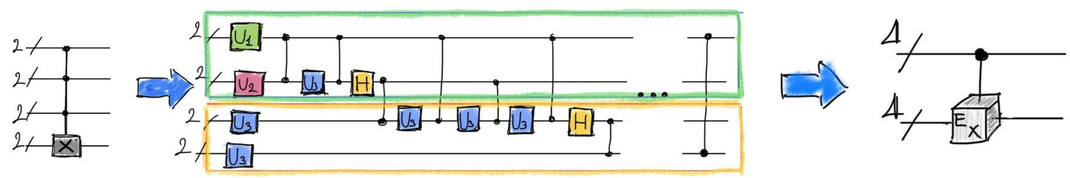

Fig. 2: On the left the multi-controlled Toffoli, *MCX* gate. On the right, a qubit circuit representing the transpiled version of the *MCX* gate to two-qubit interactions and local gates (U_1, U_2, U_3 and Hadamard). The two colored outlines represent the individual qudits, where sub-portions of the qubit circuit are assigned to, depending on the results of the mapping.

systems is the richer entangling gate set given by the choice of a specific technology.

An example illustrates the idea behind both the proposed steps.

Example 3. *Figure 2 demonstrates the mapping from a four qubit circuit with a single three-controlled Toffoli (MCX) gate. The implementation of the MCX is rather expensive in terms of number of operations. After transpiling the original gate into single-qubit gates and CZ gates, the physical interactions between the qubits are observable and can be used to cluster the qubits. In this case, the clusters are on the first two and last two qubits, which subsequently get compressed into two ququarts (qudits of dimension four). The translation step then results in only one gate, the controlled-exchange [9] (CEX) gate.*

In the following section, we introduce an algorithm that efficiently compresses qubit circuit into qudit circuits. The proposed solution will, for the first time, enable circuit compression for mixed-dimensional systems.

IV. Mapping to Mixed-Dimensional Systems

In this section, we describe a method to compress qudit circuits into mixed-dimensional qudit circuits. More precisely, this section describes the proposed approach to cluster qubits in a given circuit based on the local and non-local operations and translate these clusters into a circuit with qudits of corresponding dimensionality. Before, however, we briefly review the benefits in exploiting different qudit platforms and entanglement structures since they play an important role in determining appropriate mappings.

A. Rationale for Qudit Systems

Qudit technology platforms present several challenges, but promise to show near-term advantages over equivalent qubit realizations. In fact, qudit systems have lower decoherence rates [6], they can use higher energy states for implementing error correcting codes [38] and more efficient encoding schemes [39].

In Ref. [40] is presented a comparative study between all the most prominent qudit technologies, in terms of gate efficiency for systems composed of sets of qubits or qudits. This allows to understand the trade-offs required between the scaling of information density inside a quantum circuit and noise error rates. The efficiency of qudit gates must always be larger than the multi-qubit one by a factor $\mathcal{O}(d^2/\log_2(d))$. Given equivalent Hilbert space dimensions, viable qudit platforms

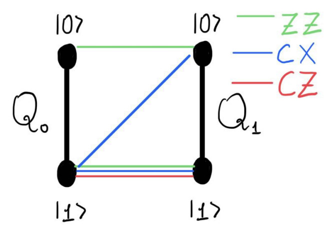

Fig. 3: The black solid lines are local operations between basis states, the black dots. The colored lines are the entangling structure of *CZ*, *CX*, and *CZ* on two qubits.

are capable of outperforming equivalent state-of-the-art multi-qubit ones in gate fidelity for pure dephasing. Moreover, this performance could be extended to qudits with d as large as 40 or more. The reflections that come naturally with this analysis is that qubit circuit compression can be performed with different performances depending on the structure of the quantum circuit, but especially on the capabilities of the technology platform chosen for implementing d-leveled circuits. Performing the most efficient compression of a specific qubit circuit will require the choice of the most suitable dimensionality of the target qudits, and consequently the right technology that can efficiently realize that circuit.

B. Entanglement Structures

Entanglement is a powerful quantum mechanical effect and a key component in the efficiency of quantum computations. Two or more qubits are in an *entangled* state, if their state cannot be written as a product of states of the individual qubits. Moreover, the *entangling power* or *structure* of non-local operations corresponds to the amount of entanglement generated by an operation, accompanied by the analysis of how the constituent basis states of an entangled state are affected by it.

Example 4. *Figure 3 illustrates three different examples of entangling structures for two qubits.*

- *ZZ (green line) entangles $|0\rangle$ and $|1\rangle$ of Q_0 with $|0\rangle$ and $|1\rangle$ of Q_1, respectively.*
- *CX (red line) entangles $|1\rangle$ of Q_0 with $|0\rangle$ and $|1\rangle$ of Q_1.*
- *CZ (red line) entangles $|1\rangle$ of Q_0 with $|1\rangle$ of Q_1.*

The graph is derived from the matrix representations of the operations and provides an insight on the physical application of the pulses representing these operations. The three operations present three completely different entangling structures;

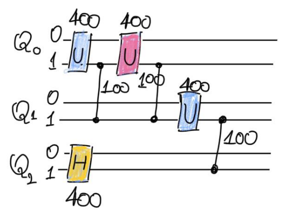

Fig. 4: A transpiled circuit where the operations are applied to the single levels, the colored gates are respectively U gates and Hadamard. There are weights assigned to each operation.

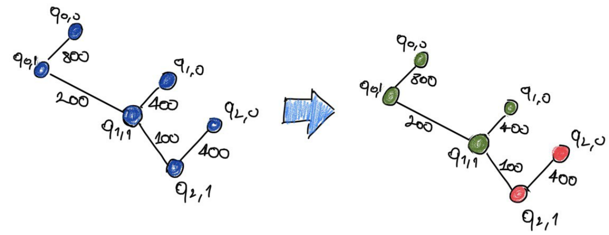

Fig. 5: Initial and clustered graph. The left-hand side of shows the corresponding initial (unlabeled) graph of levels and weights between the levels. The right-hand side shows the clustered graph with a ququad (green nodes, four dimensions) and a qubit (red nodes, two dimensions).

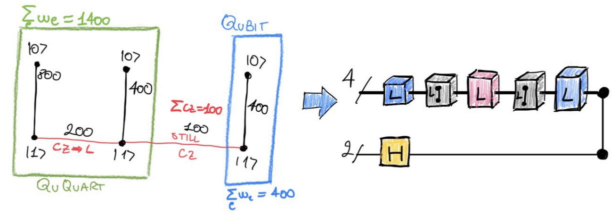

Fig. 6: Compiling into qudit circuit–illustrates a clustering on a graph as well as the correspondingly compiled circuit.

although the CX and CZ have the same amount of entropy or entanglement, these physical measures are spread in different ways, as graphically represented. The ZZ has double the amount of entanglement, and we can visually represent it as the application of two CZ on the upper and lower nodes of the graph.

C. Mapping Qubit Circuits by Clustering Weighted Graphs

Recall from the previous sections that the objective of mapping qubit circuits to mixed-dimensional quantum systems to suitable qudits is to reduce the number of qudits and non-local operations. In this section, we utilize a graph representation of circuits, where nodes correspond to levels in the qubit states and edges correspond to non-local operations between them, if the nodes belong to two different qubits, or local operations if two nodes belong to the same qubit.

Since the single levels of the qubit have now an increased relevance, we will represent quantum circuits with each qubit split into level 0 and level 1, as in Figure 4. In the graph, edges with a high weight signify many local or non-local operations between the corresponding nodes and, as such, promising candidates of nodes to be included in a cluster.

To transform a given circuit into the graph representation, it is first transpiled to only unitary local operations and controlled-Z (CZ) gates. This enables an easy discrimination between local and non-local operations when more sophisticated gates are used. The choice of the CZ gate also simplifies the resulting graph since it is only a one-edge connection between the $|1\rangle$ of two different qubits (see Figure 3). Moreover, it is a well studied gate with useful properties for application between two qudits [36], [41].

Every operation in the quantum circuit is assigned a weight, and the corresponding edge will store the sum of all the weights of all the operations applied to it. The weight represents the attraction between two distinct levels in the graph. Further, the weight of local operations is 4 times higher than that of CZ gates.

Although the error rates of entangling operations being one order of magnitude higher than the local ones, the cost of implementation of a qubit local operations transpiled into a non-local one in the qudit circuit is roughly estimated to be at least 4 times higher, as it supposed that a local operation transpiled to an entangling operation will require several swap operations to make the operation feasible. In this way, the graph allows also to split the logic of a single qubit between two qudits, when necessary. This can be useful when a qubit has no local operations but is frequently used as control or as an ancilla. The control level can be mapped together with the target qubit inside the same qudit, giving a local connotation to the previous entangling relation.

Once the weighted graph is retrieved, levels with high interaction should ideally be mapped to the same qudit. This is achieved with an adaption of the K-means algorithm. We refer to the GCLU software that implements the K-means and M-algorithm for graph clustering [42] for details on this algorithm. The algorithm effectively has the objective of assembling together nodes that maximize the total sum of the weights on the edges internal to the cluster, while it minimizes the total sum of the weights of the edges connecting two different clusters. Since the mapping to qudit systems requires the minimization of entangling operations, in this way the algorithm reduces the number of qubit operations and also the number of entangling operations in-between qudits that need to be transpiled.

Since running K-means on the graph enables results only where the final qudits always have local operations, we propose a second construction, that considers the possibility of having qudits without local operations—referred to as *full connectivity*. The initial construction of the graph remains identical, but in an additional step, all pairs of nodes without an edge between them are connected by an edge with a small but non-zero weight. Figure 5 shows how the results of clustering a graph following only the circuit's connectivity.

Figure 6 shows how once the final mapping is produced, the transpilation step can take place. For this step, the efficient realization of the entangling operations between qudits is assumed to be solved externally from the tool by a compiler, and that the experimental control is up to the task.

V. Evaluation

In this section, we evaluate the method proposed in this paper. To this end, we consider a set of established algorithms

TABLE I: Evaluation of the proposed approach comparing the number of required non-local *CZ* gates

Benchmark			Full Connectivity			Circuit Connectivity		
Name	Qubits	#*CZ*	Qudits	#*CZ*	Ratio	Qudits	#*CZ*	Ratio
Amplitude Est.	10	90	[8x1,2x1,4x3]	78	0.867	[8x1,2x1,4x3]	78	0.867
	31	918	[8x2,16x1,2x5,4x8]	878	0.956	[128x1,2x6,4x9]	858	0.935
GHZ State	10	9	[4x5]	4	0.444	[4x5]	4	0.444
	15	14	[2x1,4x7]	7	0.500	[2x1,4x7]	7	0.500
	31	30	[2x1,4x15]	15	0.500	[8x1,2x2,4x13]	15	0.500
Graph State	10	10	[8x1,2x1,4x3]	4	0.400	[16x1,1x1,4x3]	6	0.600
	15	15	[2x1,4x7]	8	0.533	[2x1,4x7]	8	0.533
	31	31	[2x1,4x15]	16	0.516	[2x1,4x15]	16	0.516
Grover no-Ancillas	10	39 032	[32x1,2x3,4x1]	36 924	0.946	[32x1,2x3,4x1]	36 924	0.946
QAOA (Max-Cut)	10	40	[8x1,2x1,4x3]	16	0.400	[8x1,2x1,4x3]	16	0.400
	15	60	[2x1,4x7]	32	0.533	[2x1,4x7]	32	0.533
Quantum Phase Est. Inexact	10	102	[8x1,2x1,4x3]	81	0.794	[8x1,2x1,4x3]	81	0.794
	15	231	[16x1,2x3,4x4]	202	0.874	[16x1,2x3,4x4]	202	0.874
	31	963	[64x1,2x5,4x10]	886	0.920	[32x1,2x4,4x11]	891	0.925
VQE (Max-Cut with TwoLocal ansatz)	10	90	[4x5]	80	0.889	[4x5]	80	0.889
	15	210	[2x1,4x7]	196	0.933	[2x1,4x7]	196	0.933
W-State	10	18	[4x5]	8	0.444	[4x5]	8	0.444
	15	28	[2x1,4x7]	14	0.500	[2x1,4x7]	14	0.500
	31	60	[2x1,4x15]	30	0.500	[2x1,4x15]	30	0.500

"Ratio" is the fraction of *CZ* gates in the compressed and initial circuits.

working on qubits and compress them to circuits working on a mixed-dimensional set of qudits.

The implementation is available freely under the MIT license at github.com/cda-tdum/qudit-compression as part of the *Munich Quantum Toolkit* (MQT). It is completely written in Python 3, with exception of the external dependencies Qiskit [43] and GCLU Software [42]. This section provides benchmarks for the problem of circuit compression to mixed-dimensional systems, solved through the realization of a new mapping solution for qubit circuits towards high and mixed-dimensional systems.

The evaluations were performed on a server running GNU/Linux using an Intel(R) Xeon(R) W-1370P (at 3.6 GHz) and 128 GiB main memory. The execution time is dependent on Qiskit [43] and Python version under which the initial parsing and analysis of the qubit circuit is performed, and on the number of iterated runs of the clustering algorithm. The clustering procedure is run 20 000 times, with 4 as preferable dimensions of the clusters. The clustered physical dimensions are converted in terms of dimensions of the single qudit spaces. Each of the considered benchmarks has a total execution time in the order of seconds.

The results are presented in Table I. The first group of columns contains the basic information on the considered benchmark: The name of the algorithm, the number of qubits, and the number of *CZ* gates after decomposition. The following two column groups contain the information on the resulting qudit circuits for "Circuit Connectivity" (edges in the graph correspond to the connections of the states according to the circuit) and "Full Connectivity" (all nodes in the graph are pair-wise connected). The "Qudit" column lists the dimensions of the qudits in the resulting circuit in the form *Dimension x Count*, e.g., 4x3 denoting three qudits of dimension four. Following, the columns "#*CZ*" give the number of non-local *CZ* operations in the circuits—with the "Ratio" between the compressed circuit and the initial qubit circuit. In many of the benchmark instances, the compressed circuits need less than 50 % of the non-local gates compared to the initial qubit circuit. The method has always success even if with a small percentage of improvement.

There is an instance of the Max-Cut problem, solved with two different algorithms, namely QAOA and VQE. One set of solutions is squarely set around 40 % to 50 %, while the other group is around 90 %. Circuit structure plays a major role in mapping to a mixed-dimensional, albeit the study of the reasons is beyond the scope of this work.

Overall, the results confirm the drastic reduction in the number of non-local operations as mixed-dimensional qudits are used to realize the circuit. This reduction directly translates to a reduction of one of the main contributors for errors in quantum operations for the circuits.

VI. Conclusion

In this work, we proposed an approach to compress qubit circuits into qudit circuits. While qudits offer a richer state-space, most algorithms as of today are developed with qubits in mind. The proposed method enables utilizing the best of both worlds: The algorithms can be based on two-level qubits but during the eventual execution the hardware can efficiently exploit higher-dimensions. The evaluations showed a significant reduction in the number of non-local quantum operations and, therefore, a reduction in one of the main contributions of undesired noise due to errors in quantum operations. Possible future work comprises a stronger consideration of the hardware to create clusters that are more compliant to specific gate sets imposed by the qudit platform.

Acknowledgments

The authors want to thank Martin Ringbauer for his feedback. This work has received funding from the European Union's Horizon 2020 research and innovation programme under the ERC Consolidator Grant (agreement No 101001318) and the NeQST Grant (agreement No 101080086). 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).

References

[1] J. Preskill, "Quantum computing in the NISQ era and beyond," *Quantum*, vol. 2, p. 79, 2018.

[2] F. Arute, K. Arya, *et al.*, "Quantum supremacy using a programmable superconducting processor," *Nature*, vol. 574, no. 7779, pp. 505–510, 2019.

[3] F. Flamini, N. Spagnolo, and F. Sciarrino, "Photonic quantum information processing: A review," *Reports Prog. Phys.*, vol. 82, p. 016 001, 2019.

[4] D. S. Weiss and M. Saffman, "Quantum computing with neutral atoms," *Physics Today*, vol. 70, no. 7, pp. 44–50, 2017.

[5] I. Pogorelov, T. Feldker, *et al.*, "Compact Ion-Trap Quantum Computing Demonstrator," *PRX Quantum*, vol. 2, p. 020 343, 2021.

[6] Y. Wang, Z. Hu, B. C. Sanders, and S. Kais, "Qudits and high-dimensional quantum computing," *Front. Phys.*, vol. 8, 2020.

[7] M. Huber and J. I. de Vicente, "Structure of multidimensional entanglement in multipartite systems," *Phys. Rev. Lett.*, vol. 110, p. 030 501, 2013.

[8] X. Zhang, M. Um, *et al.*, "State-independent experimental test of quantum contextuality with a single trapped ion," *Phys. Rev. Letters*, vol. 110, p. 070 401, 2013.

[9] M. Ringbauer, M. Meth, *et al.*, "A universal qudit quantum processor with trapped ions," 2021.

[10] B. P. Lanyon, M. Barbieri, *et al.*, "Simplifying quantum logic using higher-dimensional Hilbert spaces," *Nature Physics*, vol. 5, pp. 134–140, 2008.

[11] M. Ringbauer, T. R. Bromley, *et al.*, "Certification and quantification of multilevel quantum coherence," *Phys. Rev. X*, vol. 8, p. 041 007, 2018.

[12] X.-M. Hu, Y. Guo, *et al.*, "Beating the channel capacity limit for superdense coding with entangled ququarts," *Sci. Adv.*, vol. 4, no. 7, eaat9304, 2018.

[13] M. Malik, M. Erhard, *et al.*, "Multi-photon entanglement in high dimensions," *Nature Photonics*, vol. 10, pp. 248–252, 2016.

[14] M. Kononenko, M. Yurtalan, *et al.*, "Characterization of control in a superconducting qutrit using randomized benchmarking," *Phys. Rev. Res.*, vol. 3, no. 4, p. L042007, 2021.

[15] A. Morvan, V. V. Ramasesh, *et al.*, "Qutrit randomized benchmarking," *Phys. Rev. Letters*, vol. 126, p. 210 504, 2021.

[16] J. Ahn, T. Weinacht, and P. Bucksbaum, "Information storage and retrieval through quantum phase," *Science*, vol. 287, no. 5452, pp. 463–465, 2000.

[17] C. Godfrin, A. Ferhat, *et al.*, "Operating quantum states in single magnetic molecules: Implementation of Grover's quantum algorithm," *Phys. Rev. Letters*, vol. 119, p. 187 702, 2017.

[18] B. E. Anderson, H. Sosa-Martinez, *et al.*, "Accurate and robust unitary transformations of a high-dimensional quantum system," *Phys. Rev. Letters*, vol. 114, p. 240 401, 2015.

[19] Z. Gedik, I. A. Silva, *et al.*, "Computational speed-up with a single qudit," *Sci. Rep.*, vol. 5, p. 14 671, 2015.

[20] F. Tacchino, A. Chiesa, *et al.*, "A proposal for using molecular spin qudits as quantum simulators of light–matter interactions," *J. Mater. Chem. C*, vol. 9, pp. 10 266–10 275, 32 2021.

[21] A. S. Nikolaeva, E. O. Kiktenko, and A. K. Fedorov, "Generalized toffoli gate decomposition using ququints: Towards realizing grover's algorithm with qudits," *Entropy*, vol. 25, no. 2, 2023.

[22] S. Das and F. Caruso, "A hybrid-qudit representation of digital rgb images," 2022.

[23] E. J. Gustafson, "Prospects for simulating a qudit-based model of $(1+1)$D scalar qed," *Phys. Rev. D*, vol. 103, p. 114 505, 11 2021.

[24] D. M. Kurkcuoglu, M. S. Alam, *et al.*, "Quantum simulation of ϕ^4 theories in qudit systems," *TBD*, 2021.

[25] J. Casanova, L. Lamata, *et al.*, "Quantum simulation of quantum field theories in trapped ions," *Phys. Rev. Lett.*, vol. 107, p. 260 501, 26 2011.

[26] L. Lamata, A. Mezzacapo, J. Casanova, and E. Solano, "Efficient quantum simulation of fermionic and bosonic models in trapped ions," *EPJ Quantum Technology*, vol. 1, p. 9, 2014.

[27] B. Lanyon, M. Barbieri, *et al.*, "Quantum computing using shortcuts through higher dimensions," *Nature Physics*, vol. 5, 2008.

[28] A. Litteken, J. M. Baker, and F. T. Chong, "Communication trade offs in intermediate qudit circuits," pp. 43–49, 2022.

[29] L. Seifert, J. Chadwick, *et al.*, "Time-efficient qudit gates through incremental pulse re-seeding," pp. 304–313, 2022.

[30] T. J. Stavenger, E. Crane, *et al.*, "Bosonic qiskit," 2022.

[31] K. Mato, M. Ringbauer, S. Hillmich, and R. Wille, "Compilation of entangling gates for high-dimensional quantum systems," 2023.

[32] D. Volya and P. Mishra, "Quantum data compression for efficient generation of control pulses," pp. 216–221, 2023.

[33] K. Mato, M. Ringbauer, S. Hillmich, and R. Wille, "Adaptive compilation of multi-level quantum operations," pp. 484–491, 2022.

[34] B. Poór, Q. Wang, *et al.*, "Completeness for arbitrary finite dimensions of zxw-calculus, a unifying calculus," 2023.

[35] A. S. Nikolaeva, E. O. Kiktenko, and A. K. Fedorov, *Efficient realization of quantum algorithms with qudits*, 2021.

[36] X. Gao, P. Appel, *et al.*, "On the role of entanglement in qudit-based circuit compression," *arXiv preprint arXiv:2209.14584*, 2022.

[37] A. Litteken, L. M. Seifert, *et al.*, "Qompress: Efficient compilation for ququarts exploiting partial and mixed radix operations for communication reduction," pp. 646–659, 2023.

[38] A. Chiesa, F. Petiziol, *et al.*, "Embedded quantum-error correction and controlled-phase gate for molecular spin qubits," *AIP Advances*, vol. 11, no. 2, p. 025 134, 2021.

[39] F. Petiziol, A. Chiesa, *et al.*, "Counteracting dephasing in molecular nanomagnets by optimized qudit encodings," *npj Quantum Information*, vol. 7, no. 1, 2021.

[40] D. Janković, J.-G. Hartmann, M. Ruben, and P.-A. Hervieux, "Noisy qudit vs multiple qubits : Conditions on gate efficiency," 2023.

[41] S. Clark, "Valence bond solid formalism for d-level one-way quantum computation," *Journal of Physics A: Mathematical and General*, vol. 39, no. 11, pp. 2701–2721, 2006.

[42] S. Sieranoja and P. Fränti, "Adapting k-means for graph clustering," *Knowl. Inf. Syst.*, vol. 64, no. 1, pp. 115–142, 2022.

[43] A-tA-v, M. S. ANIS, *et al.*, *Qiskit: An open-source framework for quantum computing*, 2021.