

# How to Efficiently Handle Complex Values? Implementing Decision Diagrams for Quantum Computing

Alwin Zulehner, Stefan Hillmich, and Robert Wille

Institute for Integrated Circuits, Johannes Kepler University Linz, Austria  
alwin.zulehner@jku.at   stefan.hillmich@jku.at   robert.wille@jku.at

**Abstract**—Quantum computing promises substantial speedups by exploiting quantum mechanical phenomena such as superposition and entanglement. Corresponding design methods require efficient means of representation and manipulation of quantum functionality. In the classical domain, decision diagrams have been successfully employed as a powerful alternative to straightforward means such as truth tables. This motivated extensive research on whether decision diagrams provide similar potential in the quantum domain—resulting in new types of decision diagrams capable of substantially reducing the complexity of representing quantum states and functionality. From an implementation perspective, many concepts and techniques from the classical domain can be re-used in order to implement decision diagrams packages for the quantum realm. However, new problems—namely how to efficiently handle complex numbers—arise. In this work, we propose a solution to overcome these problems. Experimental evaluations confirm that this yields improvements of orders of magnitude in the runtime needed to create and to utilize these decision diagrams. The resulting implementation is publicly available as a quantum DD package at [http://iic.jku.at/eda/research/quantum\\_dd](http://iic.jku.at/eda/research/quantum_dd).

## I. INTRODUCTION

Quantum computing [1] promises significant speedups for certain problems, e.g., integer factorization [2], database search [3], and quantum chemistry [4]. Unlike bits in a classical computer, which can only assume one of the two basis states 0 and 1, *qubits* in a quantum computer can be in an (almost) arbitrary superposition of both. Superposition in combination with other quantum phenomena such as entanglement and phase shifts allows for an exponential speedup compared to classical computers in the best case. Not surprisingly, this triggered a huge interest in the utilization of quantum devices which, recently, led to first realizations provided by large commercial players such as IBM and Google.

These developments are currently progressing to a point where straightforward automated approaches for, e.g., synthesis, simulation, or verification, become indispensable since corresponding quantum states as well as quantum operations are mathematically described through state vectors and unitary matrices which grow exponentially in size with respect to the number of involved qubits.

In the classical domain, the design automation community successfully addressed such challenges by introducing decision diagrams which, in many cases, allow for a compact representation of functionality. Impressive accomplishments in the '90s, e.g., with *Binary Decision Diagrams* (BDDs, [5]), *Binary Moment Diagrams* (BMDs, [6]), or *Zero-suppressed Decision Diagrams* (ZDDs, [7]), are exam-

ples of the potential of those representations. Corresponding implementations (usually denoted *DD packages*) as provided by Fabio Somenzi's CUDD package [8], the Word-level-DD package [9], or Donald Knuth's BDD package [10] affect the development of design tools and methods until today. Motivated by that, in the past years, researchers spent considerable efforts in the investigation of whether decision diagrams can also be utilized in the quantum computing realm. This led to several theoretical and mathematical concepts of decision diagrams such as *X-decomposition Quantum Decision Diagrams* (XQDDs, [11]), *Quantum Decision Diagrams* (QDDs, [12]), *Quantum Information Decision Diagrams* (QuIDDs, [13]), or *Quantum Multiple-valued Decision Diagrams* (QMDDs, [14], [15])—leading to more efficient methods for the design tasks outlined above, i.e., synthesis [16]–[18], simulation [19]–[21], or verification [11], [22], [23].

However, besides the mathematical concepts, efficient *implementations* of the corresponding decision diagrams for quantum computing are needed in order to eventually utilize them for design automation at large scale—an issue that has not explicitly been addressed thus far. Key concepts required for implementing DD packages—such as unique tables, garbage collection with reference counts, or compute tables—are already known from the classical domain (which, also in the '90s, have explicitly been investigated, e.g., in [24]–[26]). While these implementation techniques can be directly incorporated into decision diagrams addressing quantum computing, the quantum realm additionally requires an efficient handling of complex numbers. These complex numbers introduce several new problems such as how to keep numerical stability, how to efficiently store nodes in unique tables, as well as how to store reoccurring operations in compute tables. None of these issues have explicitly been considered thus far.

In this work, we provide details on how to efficiently implement a DD package for quantum computing addressing these problems. Established concepts known from decision diagrams in the classical domain are re-used where applicable, while new implementation techniques for handling complex values efficiently are described in detail. Experimental results confirm that these efficient implementation techniques yield improvements of orders of magnitude with respect to runtime compared to the best known implementations available today. An implementation of the resulting DD package is publicly available at [http://iic.jku.at/eda/research/quantum\\_dd](http://iic.jku.at/eda/research/quantum_dd).

The remainder of this paper is structured as follows: Section II briefly reviews the basics of quantum computing and applicable decision diagrams. In Section III, we recapitulate implementation techniques used for classical DD packages before discussing problems related to the extension of them into the quantum domain. Based on that, Section IV describes in detail how the corresponding problems can efficiently be handled—yielding an efficient implementation of a DD package for quantum computing. Eventually, the contributions of this work have been incorporated into an implementation of a DD package for quantum computing whose performance is evaluated and compared to the state of the art in Section V. Finally, the paper is concluded in Section VI.

## II. BACKGROUND

In this section, we review the basics of quantum computing and decision diagrams for representing quantum functionality.

### A. Quantum Computing

Computations in the quantum realm use qubits, which can assume more states than the basis states (here, written as  $|0\rangle$  and  $|1\rangle$  using Dirac-notation) used in classical computations. A quantum state  $|\psi\rangle$  is given by  $\alpha \cdot |0\rangle + \beta \cdot |1\rangle$  with complex-valued amplitudes  $\alpha, \beta$ . Following this description,  $|\alpha|^2$  and  $|\beta|^2$  are the probabilities to measure the base state  $|0\rangle$  or  $|1\rangle$ , respectively, and, therefore, their sum  $|\alpha|^2 + |\beta|^2$  has to be equal to 1.

Manipulation of a quantum state is achieved by a sequence of “simple” quantum operations (also denoted quantum gates) that are described by *unitary matrices* acting on one or more qubits each. Such sequences of quantum operations are represented by circuit diagrams that indicate (from left to right) which operations are applied to which qubits.

**Example 1.** Common operations performed on single qubits are the *NOT* operation  $X$ , the Hadamard operation  $H$  to set a qubit into superposition, and the phase shift operation  $T$ . The corresponding unitary matrices are defined as

$$X = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \quad H = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}, \quad \text{and} \quad T = \begin{bmatrix} 1 & 0 \\ 0 & e^{i\frac{\pi}{4}} \end{bmatrix}.$$

In addition to the single qubit operations, there are also controlled operations such as the controlled *NOT* (*CNOT*) operation shown in the right-hand side of the quantum circuits depicted in Figure 1a (next to the Hadamard operation). Here, the qubit  $q_1$  will only be negated iff  $q_0$  is in the  $|1\rangle$  state. Due to the superposition introduced by the Hadamard operation, both qubits become entangled, i.e., a measurement of one qubit affects the quantum state of the other qubit as well.

The overall functionality of a quantum circuit is determined by successively multiplying all gate matrices of the circuit. Hence, all gate matrices must have a dimension of  $2^n \times 2^n$  (assuming an  $n$ -qubit system). If a gate matrix operates on a subset of the qubits only, the  $2 \times 2$  identity matrix  $I_2$  is assumed for the other qubits. Forming the Kronecker product of these matrices results again in a matrix of size  $2^n \times 2^n$  (assuming an  $n$ -qubit quantum system).




Fig. 1: Quantum circuits

### B. Decision Diagrams for Quantum Computing

Since the unitary matrices representing quantum functionality grow exponentially with respect to the number of qubits, i.e., the size of the quantum system, representations based on two-dimensional arrays quickly become infeasible. However, the structure of many explicitly used matrix instances offers the potential to exploit redundancies and, by this, allow for a drastically more compact representation (while still preserving an efficient manipulation). As in the classical realm, this is exploited by decision diagrams such as [11]–[15]. These decision diagrams represent the matrix as a directed acyclic graph, where identical sub-matrices are combined in a shared graph structure. Further potential for redundancies can be achieved by sharing (sub-)graphs which are structurally equivalent and only differ by a common factor (to be annotated as a weight to the corresponding edges). The following example illustrate the main ideas:

**Example 2.** Figure 2 shows different representations of the quantum functionality shown in Figure 1b, i.e., described by the Kronecker-product  $H \otimes I_2$ . In Figure 2a, the corresponding unitary matrix is shown, while Figure 2b shows a corresponding decision diagram structure. In this graph, the root node labeled  $q_0$  represents the whole matrix, while the four outgoing edges point to nodes that represent the top-left, top-right, bottom-left, and bottom-right sub-matrices (from left to right, hinted in the matrix by dashed lines). These  $2 \times 2$  sub-matrices (representing functionality with respect to  $q_1$  only and, hence, represented by nodes labeled  $q_1$ ) are further decomposed—yielding single complex values (or  $1 \times 1$  sub-matrices) represented by terminal nodes. As shown in Figure 2a, the top-left, top-right, and bottom-left sub-matrices are identical and, therefore, are represented by the same node in Figure 2b (the leftmost node labeled  $q_1$ ).

Furthermore, the bottom-right sub-matrix is structurally equivalent to the other sub-matrices and only differs by a common factor, namely  $-1$ . Hence, the matrix can be represented even more compactly as shown in Figure 2c. Here, additional edge weights are employed while the eventual matrix entry is determined by multiplying all edge weights from the root node to the terminal.<sup>1</sup> This does not only allow to share the node representing the bottom-right sub-matrix, but also to boil down the number of required terminals to one (the actual complex values are now determined as product of the edge weights). As an example, the matrix element highlighted bold in Figure 2a is determined by multiplying all weight on the bolded path in Figure 2c. Eventually, this yields a substantially more compact representation as originally given by the matrix.

<sup>1</sup>For the sake of readability, edge weights 1 are omitted and weights 0 are depicted as stubs.

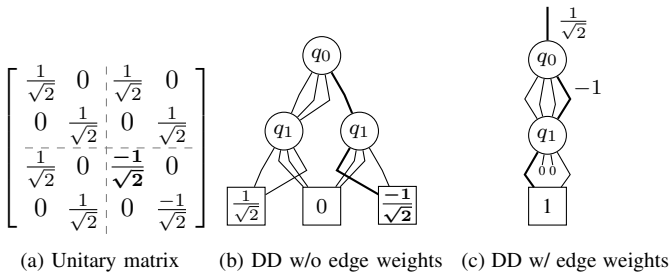


Fig. 2: Representations for  $U = H \otimes I_2$ .

The compaction can be further improved by *normalizing* the nodes of the decision diagram (such as described in [27]). This is commonly achieved by using the leftmost non-zero weight of the outgoing edges as a normalization factor. All outgoing edges of the corresponding node are then divided by this factor, which is propagated towards the incoming edges to ensure products on all paths do not change. This normalization scheme further provides a canonical, i.e., unique, representation of unitary matrices (as proven in [15]).

### III. IMPLEMENTING DECISION DIAGRAMS FOR QUANTUM COMPUTING

In this section, we discuss how to implement DD packages for the quantum domain in an efficient fashion. To this end, we utilize concepts such as *unique tables*, dedicated *garbage collection*, or *compute tables* which are already taken for granted in DD packages for the classical domain [8]–[10]. In addition to that, we show that just adopting those concepts is not sufficient for a highly efficient quantum DD package and additional implementation techniques are required which allow for efficiently handling complex numbers and sub-factors thereof (essential for the edge weights discussed in the previous section).

#### A. Established Implementation Techniques

The key concepts required for implementing DD packages for the classical domain in an efficient fashion have been introduced in the '90s [24]–[26] and leveraged the development of efficient packages such as [8]–[10] (which served as basis for efficient methods for verification or synthesis). Many of these concepts can be directly re-used decision diagrams in the quantum domain, namely:

- *Unique tables* which store the nodes of the decision diagram and allow to efficiently detect redundancies in the structure. The unique table is realized as two-dimensional hash table (one for each variable in the decision diagram), where each bucket of the table contains a linked list of DD nodes with an equal associated hash value for each variable. Before any new DD node is inserted, it is checked whether the node already exists in the unique table.<sup>2</sup> A DD node is returned by means of a pointer to the respectively allocated memory to allow dereferencing in constant time. Hence, it can be denoted as *strong canonical form* [25], i.e., a unique representation.

<sup>2</sup>The hash function should distribute the nodes equally to the buckets to reduce the number of collisions at lookup.

- Dedicated *garbage collection* which frequently removes unused DD nodes from the unique table and, by this, allows for a fast insertion of new nodes as well as for keeping the memory usage low. Here, *reference counting* is used to keep track which node shall be kept in the unique table. In order to avoid recurring memory allocations and deallocations, free DD nodes (i.e., allocated nodes that are currently not stored in the unique table) are stored in a list and corresponding memory is allocated en block.<sup>3</sup> Nodes that are removed from the unique table are appended to this list. To reduce the overhead caused by the lists (the list of free nodes as well as the lists in the buckets of the unique table), DD nodes themselves contain a pointer to the next element in the list.
- *Compute tables* which cache the results of operations that are repeatedly conducted on the same DD nodes. Since decision diagrams exploit redundancies by using shared nodes, this proves very beneficial as many operations are frequently repeated and, hence, do not have to be recomputed. Compute tables are realized as hash tables, where the DD nodes denoting the operands (and possibly the type of operation) are used to determine the hash key [25]. The entries in the hash table might contain a single result that may be overwritten, or linked lists to store all computations with the same hash key to increase the *hit rate* (e.g., with an LRU strategy as proposed in [26]). Note that the compute tables have to be (partially) invalidated/cleared after garbage collection, since operands (referenced by a pointer) may be overwritten after the respective nodes have been removed from the unique table.

In fact, all these concepts can be (and have been) directly incorporated into implementations of decision diagrams for quantum computing reviewed in Section II-B. However, DD packages for quantum computing additionally have to handle the frequently occurring complex numbers. This constitutes *the* major obstacle towards an efficient implementation of a fully-fledged DD package for the quantum domain as discussed in the next section.

#### B. Handling Complex Numbers

As discussed in Section II, weights that are attached to DD nodes offer the possibility for further compaction. In the classical domain, having such edge weights does not constitute an issue from the implementation perspective, since they are (tuples of) integers [5]–[10]—again, a strong canonical form (a unique representation). This allows to use the efficient concepts outlined above, since computing unambiguous hash keys for DD nodes (containing weights attached to outgoing edges) is still possible.

In the quantum domain, however, weights are formed out of complex numbers. From a mathematical perspective, this does not cause problems since complex numbers also provide a strong canonical form. However, it introduces severe

<sup>3</sup>Whenever the list is empty, we allocate enough memory to instantiate several new DD nodes.

challenges from an implementation perspective where machine accuracy is limited and, hence, complex numbers are approximated—yielding to numerical errors in computations. In fact, changing one complex number attached as weight to an outgoing edge by a tiny fraction (e.g., by flipping the least significant bit of the mantissa of the real or imaginary part) may yield a completely differently computed hash key. Accordingly, redundancies might remain undetected by the DD package since the node is searched in the wrong bucket of the unique table—causing a substantially larger decision diagrams even though redundancies are actually present. Additionally considering that weights represent sub-factors of complex numbers further increase the possibilities of corresponding numerical instabilities.<sup>4</sup>

To overcome this issue, one can represent complex numbers as two quadratic irrational numbers of the form  $\frac{a+b\sqrt{2}}{c}$  with  $a, b, c \in \mathbb{Z}$  for the real and imaginary part [29], [30]—severely limiting the number of possible complex numbers and, thus, not allowing for representing arbitrary quantum functionality. Alternatively, one can use an additional table to store complex numbers, where an edge weight is then represented by an index of this table holding the corresponding number [30], [31]. The table is then maintained in a fashion that numbers that do not differ by more than a tolerance value  $\epsilon$  share an entry in this table.<sup>5</sup>

However, while using a lookup table for complex numbers indeed allows for representing arbitrary quantum functionality, it introduces several new problems when aiming for an efficient implementation. Thus, rather straightforward implementations of the respective DD concepts are available only. In the remainder of this paper, we identify arising difficulties and propose new implementation techniques which, for the first time, explicitly allow for implementing complex-valued edge weights in an efficient fashion.

#### IV. EFFICIENT HANDLING OF COMPLEX EDGE WEIGHTS

As discussed above, handling complex-valued edge weights (and complex numbers in general) in an efficient fashion is a key to a fully-fledged DD package for quantum computing. To this end, we propose such techniques in this section—covering arising issues like numerical instabilities caused by  $\epsilon$ , how to realize an efficient lookup of complex numbers that considers  $\epsilon$  (thus, providing a strong canonical form), as well as how operations on DDs can be handled efficiently.

##### A. Obtaining Numerical Stability

Design tasks for quantum computing (like synthesis, verification, or simulation) heavily rely on multiplying unitary matrices either with each other or with vectors. From a numerical perspective, these operations are not critical, since the multiplication with a unitary matrix is a well conditioned operation—even when the individual entries of the matrix are determined as product of several factors (as done in an edge-valued decision diagram). However, this changes when

introducing the tolerance value  $\epsilon$  as discussed above to detect redundancies. More precisely, some factors of an entry in the unitary matrix might be significantly rounded.

This problem becomes evident when using a normalization scheme (to gain canonicity of DD nodes) as described in [32]. Here, all weights of outgoing edges are simply divided by the leftmost non-zero outgoing edge weight while propagating this extracted factor to the parent nodes (cf. Section II)—edge weights in the decision diagrams are likely to become either rather large or rather small. If the real and the imaginary part of an edge weight are now close to 0 (in the interval  $[-\epsilon, \epsilon]$ ), the weight is rounded to 0. By this, a sub-tree of the decision diagram is possibly pruned by setting several entries in the matrix to zero—ending up with a huge round-off error and numerical instabilities.

A much better numerical stability is reached by changing the utilized normalization scheme. We propose to divide all weights of the outgoing edges of a DD node by the weight with the largest absolute value. If several outgoing edges have attached weights with equal absolute values (as can be seen in Figure 2c for the node labeled  $q_0$ ), we divide by the leftmost of these weights to preserve canonicity. With this normalization scheme, it is guaranteed that all edge weights in the decision diagram have an absolute value between 0 and 1 (and, thus, also the absolute values of the real and imaginary parts are between 0 and 1)—making it less likely that a factor unintendedly rounds to 0. While this indeed helps to increase numerical stability, one can additionally exploit the knowledge that all occurring complex numbers are either on or inside the unit circle to store them efficiently. This is discussed in the next section.

##### B. Looking up Complex Numbers

As discussed in Section III-B, hashing complex numbers is usually not possible for many quantum functions due to rounding errors caused by the limited machine accuracy. Hence, different methods are required that allow for a unique and efficient lookup of complex numbers while still considering the tolerance value  $\epsilon$ .

The general idea for an efficient lookup is to exploit the fact that numbers can be sorted. Since there does not exist a total order for complex numbers, we split them into their real and imaginary parts. These real-valued parts are then stored separately in a lookup-table.<sup>6</sup> A complex number is then represented by a pair of pointers to elements in the lookup-table—a strong canonical form. Moreover, we store only the absolute value of the real and imaginary part and “hide” the sign bit in the pointer to the respective entry of the table. This additionally allows to conduct operations like multiplication with constants like  $-1$ ,  $i$ , or  $-i$ , as well as computing the complex conjugate just by flipping bits and/or swapping pointers.

<sup>4</sup>A more detailed discussion of this issue is provided in [28].

<sup>5</sup>Note that, in [30], this  $\epsilon$  is a relative tolerance value, while it is absolute in [31].

<sup>6</sup>Note that a separate insertion of the real and imaginary part may also reduce the overall numbers to be stored since the same number might occur as real/imaginary part in several complex numbers.

To realize a lookup table for real-valued entries, we exploit the knowledge of the normalization scheme discussed above. In fact, it is guaranteed that all numbers of the table are within the interval  $[0, 1]$ .<sup>7</sup> To allow for an efficient lookup, we split this interval into  $N$  equally distributed chunks.<sup>8</sup> These chunks are represented by entries in an array of size  $N$ , where each entry initially contains an empty list for occurring numbers in the respective interval (similar to the buckets in a hash table). When a new real number  $r$  shall be inserted, the corresponding bucket is traversed. If one of the numbers in this list is equal to  $r$  (considering the tolerance value  $\epsilon$ ), a pointer to the respective number is returned. Otherwise,  $r$  is inserted into the bucket.

However, one has to be careful since, by allowing a tolerance value of  $\epsilon$ , a sufficiently close number might be located in one of the neighboring buckets. More precisely, if  $r - \epsilon$  or  $r + \epsilon$  exceeds the border of the considered interval (represented by a bucket), one has to additionally look for a number sufficiently close to  $r$  in the the corresponding neighboring bucket. For performance reasons, we propose to return the first value found which deviates less than  $\epsilon$  from  $r$ . Alternatively, to improve numerical stability, one can continue searching for a value that is even closer to  $r$ .<sup>9</sup>

We employ the implementation techniques used for the unique table (for storing DD nodes) also for the lookup table (for storing real numbers) to allow for an efficient realization. This includes reference counting to keep track which entries are still required to be in the lookup table, a list of “free” numbers, as well as a procedure to allocate memory for several new entries at once. Whenever garbage collection is conducted on the unique table, we additionally run a garbage collection routine on the lookup table to remove entries (appending them to the list of free entries) with a reference count of 0. Moreover, numbers can easily be dereferenced, since pointers are returned.

By using a lookup table as discussed above, inserting a value  $r$  has complexity  $\mathcal{O}(1)$  if  $N$  is chosen suitably and the (inserted) numbers are distributed equally in the interval  $[0, 1]$  (avoiding long collision chains). However, like a hash table, a worst-case complexity of  $\mathcal{O}(n)$  results when all  $n$  entries in the table are stored in the same bucket. As alternative to a table, one can also use a self-balancing binary tree like an AVL tree [33] or a red-black tree [34] to store the real numbers. Then, the lookup of a number  $r$  requires  $\mathcal{O}(\log n)$ , where  $n$  is the number elements in the tree. This is again possible, since real numbers can be totally ordered, but additionally requires overhead for re-balancing the tree. Eventually, a combination of both—a lookup table where all entries stored in one bucket are realized as a self-balancing binary tree—provides a compromise of both ideas from a complexity-theoretic consideration.

However, even with an efficient implementation of a lookup

<sup>7</sup>Note that this also holds for the real and imaginary part of the complex weight attached to the edge pointing to the root node, since we deal with unitary matrices.

<sup>8</sup>Note that other strategies for splitting the interval  $[0, 1]$  are possible.

<sup>9</sup>Note that there exist at most two numbers in the lookup table that are closer than  $\epsilon$  to  $r$ .

for complex (real) numbers, several other issues have to be considered when conducting operations on decision diagrams. These are discussed in the following section.

### C. Conducting Operations on Decision Diagrams

In this section, we discuss how to efficiently handle complex-valued edge weights when conducting operations on decision diagrams. To this end, we first analyze arising issues (from using a lookup table for complex numbers) that might impair efficiency, namely:

- Each intermediate computation on complex numbers requires to perform a lookup in the table and, additionally, may cause rounding. Such intermediate values occur before normalizing a DD node.
- Sub-results might contain complex numbers with absolute value larger than 1 (before normalizing a DD node).

The first issue may affect the efficiency since the number of entries in the lookup table grows significantly (by inserting intermediate values that are not used anymore afterwards). Moreover, intermediate computations are not conducted as accurately as possible (by performing a lookup that considers  $\epsilon$ ). The second issue even prohibits the use of a lookup table as introduced above, since it is not guaranteed that all entries are in within interval  $[0, 1]$ . However, both issues can be resolved by introducing a cache for complex numbers that are used for storing intermediate results. Entries are taken from this cache whenever intermediate results are computed, and fed back when normalizing a newly computed DD node (before looking it up in the unique table).

In general, the efficiency of operations on decision diagrams results from a recursive formulation. Repeatedly re-evaluating the same (sub-)operation is avoided by using a compute table. Since each recursive call returns an already normalized DD node, cached complex numbers are only required for the current recursion level. Moreover, also the recursion levels above the current one may hold some cached complex numbers representing sub-results from other recursive calls (at most one for each other outgoing edge). From this, we can infer that a cache size linear to the number of variables is sufficient. Even more, by fixing the maximum number of variables in the decision diagram beforehand, a cache with fixed size for complex numbers (which is allocated at initialization of the package) is sufficient.

We implement this cache as list of real numbers and thereby utilize the same data structure as for entries in the lookup table discussed above. This allows to use entries from the lookup table and the cache interchangeably when computing, e.g., the product of two complex numbers. Cached complex numbers are allocated by taking two real numbers from the front of the list representing the cache. Before feeding them back to the cache when normalizing the computed DD node, we insert them into the lookup table (if no suitable number has been inserted yet). This is necessary since the lookup in the unique table requires a strong canonical form for the weights attached to the outgoing edges to detect redundancies. Since we look up complex numbers as late as possible, all computations are conducted with maximal precision. Intermediate results

are not inserted into the lookup table for complex numbers, which keeps the number of entries in this table low and avoids repeated rounding (caused by  $\epsilon$ ) during computations.

On the downside, having a cache for complex numbers affects the compute tables that store sub-results. In fact, depending on the considered operation, an operand’s weight (i.e., a DD edge weight) is either stored in the cache or in the lookup table. For complex numbers stored in the cache, we again have the problem that they do not necessarily preserve a strong canonical form. Hence, a completely different hash key may result from two slightly different complex numbers. This is not as critical as in the unique table, since it only may decrease the hit rate of the compute table. However, our internal evaluation has shown that the hit rate is hardly affected when rounding the complex numbers before computing the hash key.

## V. RESULTING DD PACKAGE

In this section, we describe the resulting DD package for quantum computing when utilizing the implementation techniques introduced above. As representative of a particular decision diagram type, we chose QMDDs as introduced in [15], since this type incorporates all state-of-the-art concepts of decision diagrams for quantum computing reviewed in Section II-B.

As basis for our implementation served the QMDD package provided at <http://informatik.uni-bremen.de/agra/eng/qmdd.php>, which already utilizes the implementation techniques common for DD packages in the classical domain such as unique and compute tables, as well as garbage collection with reference counters. However, complex numbers are handled in a straightforward and inefficient fashion, namely:

- Complex numbers are stored within an array of fixed size.
- The array is linearly traversed at each lookup, i.e., inserting a new number requires traversing the complete array.
- Complex numbers are inserted into the array at each step (causing rounding of intermediate results which may affect numerical stability).
- The normalization scheme does not consider numerical implications caused by  $\epsilon$ .

All these shortcomings make the original QMDD package a proof of concept implementation (handling only small instances efficiently) rather than a fully-fledged DD package for quantum computing.

The contributions of this paper, i.e., the implementation techniques for an efficient handling of complex values, address these shortcomings and, hence, have been implemented on top of this package in C. The resulting DD package is now publicly available at [http://iic.jku.at/eda/research/quantum\\_dd](http://iic.jku.at/eda/research/quantum_dd).

To demonstrate the improved efficiency, we took established quantum functionality and generated the corresponding decision diagrams with the improved package. More precisely, we considered the *Quantum Fourier Transform* (QFT [35]) and the functionality of quantum circuits proposed by Google for quantum supremacy experiments [36]. Since building the decision diagrams representing the respective functionalities is

TABLE I: Efficiency of the proposed techniques

name	$q$	$\#op$	size	$\#complex$	$t_{original}$	$t_{proposed}$
Supremacy	16	50	80	238	0.00	0.00
Supremacy	16	70	6278	18 803	4.04	0.06
Supremacy	16	75	36 161	225 560	1 147.14	0.57
Supremacy	16	80	1 195 979	–	>3 600.00	65.18
Supremacy	20	70	455	3534	0.06	0.01
Supremacy	20	80	899	8010	0.26	0.02
Supremacy	20	89	71 105	97 942	1 453.63	2.52
Supremacy	20	95	1 742 795	–	>3 600.00	912.11
Supremacy	25	100	912	6151	0.15	0.02
Supremacy	25	110	1751	12 851	0.46	0.03
Supremacy	25	119	58 939	96 647	1 282.28	2.63
Supremacy	25	120	365 643	–	>3 600.00	3.51
QFT	15	120	32 767	32 785	3.98	0.39
QFT	16	136	65 535	65 554	17.48	0.86
QFT	17	153	131 071	131 091	67.02	1.64
QFT	18	171	262 143	262 164	295.48	3.50
QFT	19	190	524 287	524 309	1 269.35	8.38
QFT	20	210	1 048 575	–	>3 600.00	22.96

$q$ : number of qubits     $\#op$ : number of operations  
size: size of the QMDD     $\#complex$ : number of occurring complex values  
 $t_{original}$ : time when using the original QMDD package [15]  
 $t_{proposed}$ : time when incorporating the techniques proposed in this paper

conducted by successively multiplying the individual quantum operations (also represented in terms of decision diagrams), this serves as a representative case study on the efficiency of the package (after all, design automation methods heavily rely on matrix multiplication).

The obtained results are provided in Table I which lists the number of qubits  $q$ , the number of quantum operations  $\#op$ , the size (i.e., the number of DD nodes) of the resulting decision diagram, as well as the number of occurring  $\#complex$  values throughout the computation.<sup>10</sup> Moreover, we list the runtime for building the decision diagram representing the respective functionality when using the originally available implementation (denoted  $t_{original}$ ) as well as when using the improved DD package utilizing the implementation techniques proposed in this work (denoted  $t_{proposed}$ ).<sup>11</sup>

The obtained results clearly show that the proposed techniques allow to handle complex numbers much more efficiently than the original implementation. For example, consider the functionality of quantum circuits proposed for quantum supremacy experiments (proposed by researchers from Google [36]) with 16 qubits. Building the decision diagram representing the quantum functionality given by the first 50 operations can be handled by both packages in a fraction of a second. Considering the first 70 operations already results in a decision diagram with more than 6 000 nodes and requires to deal with close to 20 000 complex numbers. Even though these are not large numbers for decision diagrams, the original package already requires several seconds to build the decision diagram. When considering the first 75 operations, dealing with more than 36 000 DD nodes and with approximately 225 000 complex values requires almost 20 minutes—a task that can be solved in less than a second by the DD package resulting from this work. In fact, the limiting factor is not only the number of nodes, but also the number of complex values to deal with. Hence, the proposed implementation allows to build

<sup>10</sup>Note that the number of occurring complex values can only be determined for the original package since all of them remain in an array during the computation.

<sup>11</sup>Note that we had to slightly adjust the original package to allow storing more than 10 000 different complex values. This, however, did not affect the runtime performance of the package.

up the functionality of the first 80 operations—a rather large decision diagram with almost 1.2 million nodes—in a bit more than one minute, whereas the original implementation fails to do that within an hour.

Overall, improvements in the performance of several orders of magnitude can be observed. Hence, the evaluation demonstrates that complex values (especially when used as edge weights) can be handled much more efficiently by using the implementation techniques described in this paper. It is assumed that this performance improvement can directly be utilized in design automation methods relying on decision diagrams (such as those presented in [11], [16]–[23]) just by replacing the currently used DD packages with the package proposed here (which is publicly available at [http://iic.jku.at/eda/research/quantum\\_dd](http://iic.jku.at/eda/research/quantum_dd)).

## VI. CONCLUSIONS

In this paper, we proposed implementation techniques for efficiently handling complex numbers in decision diagrams for quantum computing—especially when occurring as edge weights of the respective decision diagram. By this, we leverage—in joint consideration of implementation techniques for decision diagrams in the classical domain developed decades ago—the development of a fully-fledged DD package for the quantum domain (publicly available at [http://iic.jku.at/eda/research/quantum\\_dd](http://iic.jku.at/eda/research/quantum_dd)). The experimental evaluation showed that this package is indeed capable of handling complex numbers much more efficiently and allows constructing decision diagrams for established quantum functionality in significantly less runtime (up to several orders of magnitude). This performance boost is expected to be can passed to design automation methods utilizing decision diagrams just by incorporating this new package.

## ACKNOWLEDGMENTS

This work has partially been supported by the LIT Secure and Correct System Lab funded by the State of Upper Austria.

## REFERENCES

- [1] M. Nielsen and I. Chuang, *Quantum Computation and Quantum Information*. Cambridge Univ. Press, 2000.
- [2] P. W. Shor, “Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer,” *SIAM Jour. of Comp.*, vol. 26, no. 5, pp. 1484–1509, 1997.
- [3] L. K. Grover, “A fast quantum mechanical algorithm for database search,” in *Symp. on Theory of Computing*, 1996, pp. 212–219.
- [4] B. P. Lanyon, J. D. Whitfield, G. G. Gillett, M. E. Goggin, M. P. Almeida, I. Kassal, J. D. Biamonte, M. Mohseni, B. J. Powell, M. Barbieri *et al.*, “Towards quantum chemistry on a quantum computer,” *Nature chemistry*, vol. 2, no. 2, p. 106, 2010.
- [5] R. Bryant, “Symbolic manipulation of boolean functions using a graphical representation,” in *Design Automation Conf.*, 1985, pp. 688–694.
- [6] R. E. Bryant and Y.-A. Chen, “Verification of arithmetic circuits with binary moment diagrams,” in *Design Automation Conf.* IEEE, 1995, pp. 535–541.
- [7] S. Minato, “Zero-suppressed BDDs for set manipulation in combinatorial problems,” in *Design Automation Conf.*, 1993, pp. 272–277.
- [8] F. Somenzi, “CUDD: CU decision diagram package release 3.0.0,” <http://vlsi.colorado.edu/~fabio/>, 2015.
- [9] M. Herbstritt, “wld: A C++ library for decision diagrams,” <https://ira.informatik.uni-freiburg.de/software/wld/>, 2004.
- [10] D. E. Knuth, “The art of computer programming: Binary decision diagrams,” <https://www-cs-faculty.stanford.edu/~knuth/programs.html>, 2011.
- [11] S.-A. Wang, C.-Y. Lu, I.-M. Tsai, and S.-Y. Kuo, “An XQDD-based verification method for quantum circuits,” *IEICE Trans. Fundamentals*, vol. 91-A, no. 2, pp. 584–594, 2008.
- [12] A. Abdollahi and M. Pedram, “Analysis and synthesis of quantum circuits by using quantum decision diagrams,” in *Design, Automation and Test in Europe*. European Design and Automation Association, 2006, pp. 317–322.
- [13] G. F. Viamontes, M. Rajagopalan, I. L. Markov, and J. P. Hayes, “Gate-level simulation of quantum circuits,” in *Asia and South Pacific Design Automation Conf.*, 2003, pp. 295–301.
- [14] D. M. Miller and M. A. Thornton, “QMDD: A decision diagram structure for reversible and quantum circuits,” in *Int’l Symp. on Multi-Valued Logic*, 2006, p. 6.
- [15] P. Niemann, R. Wille, D. M. Miller, M. A. Thornton, and R. Drechsler, “QMDDs: Efficient quantum function representation and manipulation,” *IEEE Trans. on CAD of Integrated Circuits and Systems*, vol. 35, no. 1, pp. 86–99, 2016.
- [16] P. Niemann, R. Datta, and R. Wille, “Logic synthesis for quantum state generation,” in *Int’l Symp. on Multi-Valued Logic*. IEEE, 2016, pp. 247–252.
- [17] A. Zulehner and R. Wille, “One-pass design of reversible circuits: Combining embedding and synthesis for reversible logic,” *IEEE Trans. on CAD of Integrated Circuits and Systems*, vol. 37, no. 5, pp. 996–1008, 2017.
- [18] P. Niemann, R. Wille, and R. Drechsler, “Improved synthesis of Clifford+T quantum functionality,” *Design, Automation and Test in Europe*, 2018.
- [19] G. F. Viamontes, I. L. Markov, and J. P. Hayes, “High-performance QuIDD-based simulation of quantum circuits,” in *Design, Automation and Test in Europe*. IEEE Computer Society, 2004, p. 21354.
- [20] A. Zulehner and R. Wille, “Advanced simulation of quantum computations,” *IEEE Trans. on CAD of Integrated Circuits and Systems*, 2018.
- [21] —, “Matrix-vector vs. matrix-matrix multiplication: Potential in DD-based simulation of quantum computations,” in *Design, Automation and Test in Europe*, 2019, pp. 90–95.
- [22] S. Yamashita, S.-i. Minato, and D. M. Miller, “DDMF: An efficient decision diagram structure for design verification of quantum circuits under a practical restriction,” *IEICE Trans. Fundamentals*, vol. 91, no. 12, pp. 3793–3802, 2008.
- [23] P. Niemann, R. Wille, and R. Drechsler, “Equivalence checking in multi-level quantum systems,” in *Int’l Conf. of Reversible Computation*, 2014, pp. 201–215.
- [24] R. E. Bryant, “Graph-based algorithms for Boolean function manipulation,” *IEEE Trans. on Computers*, vol. 35, no. 8, pp. 677–691, 1986.
- [25] K. S. Brace, R. L. Rudell, and R. E. Bryant, “Efficient implementation of a BDD package,” in *Design Automation Conf.*, 1990, pp. 40–45.
- [26] T. Stornetta and F. Brewer, “Implementation of an efficient parallel BDD package,” in *Design Automation Conf.*, 1996, pp. 641–644.
- [27] P. Niemann, R. Wille, and R. Drechsler, “On the ‘Q’ in QMDDs: Efficient representation of quantum functionality in the QMDD data-structure,” in *Int’l Conf. of Reversible Computation*, 2013, pp. 125–140.
- [28] A. Zulehner, P. Niemann, R. Drechsler, and R. Wille, “Accuracy and compactness in decision diagrams for quantum computation,” in *Design, Automation and Test in Europe*, 2019.
- [29] D. M. Miller, M. A. Thornton, and D. Goodman, “A decision diagram package for reversible and quantum circuit simulation,” in *IEEE World Congress on Computational Intelligence*, 2006, pp. 8597–8604.
- [30] G. F. Viamontes, I. L. Markov, and J. P. Hayes, *Quantum Circuit Simulation*. Springer, 2009.
- [31] P. Niemann, A. Zulehner, R. Wille, and R. Drechsler, “Efficient construction of QMDDs for irreversible, reversible, and quantum functions,” in *Int’l Conf. of Reversible Computation*. Springer, 2017, pp. 214–231.
- [32] D. M. Miller, D. Y. Feinstein, and M. A. Thornton, “QMDD minimization using sifting for variable reordering,” *Multiple-Valued Logic and Soft Computing*, vol. 13, no. 4-6, pp. 537–552, 2007.
- [33] R. Sedgewick, *Algorithms*. Addison-Wesley, 1983, ch. 15, p. 199.
- [34] L. J. Guibas and R. Sedgewick, “A dichromatic framework for balanced trees,” in *Symp. on Foundations of Computer Science*. IEEE, 1978, pp. 8–21.
- [35] A. Ekert and R. Jozsa, “Quantum computation and Shor’s factoring algorithm,” *Reviews of Modern Physics*, vol. 68, no. 3, p. 733, 1996.
- [36] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, R. Babbush, N. Ding, Z. Jiang, M. J. Bremner, J. M. Martinis, and H. Neven, “Characterizing quantum supremacy in near-term devices,” *Nature Physics*, vol. 14, no. 6, p. 595, 2018.