

Qubit Routing Using Graph Neural Network Aided Monte Carlo Tree Search

Animesh Sinha*, Utkarsh Azad*, Harjinder Singh

Center for Computational Natural Sciences and Bioinformatics, International Institute of Information Technology, Hyderabad.
Center for Quantum Science and Technology, International Institute of Information Technology, Hyderabad.
{animesh.sinha, utkarsh.azad}@research.iiit.ac.in, laltu@iiit.ac.in

Abstract

Near-term quantum hardware can support two-qubit operations only on the qubits that can interact with each other. Therefore, to execute an arbitrary quantum circuit on the hardware, compilers have to first perform the task of qubit routing, i.e., to transform the quantum circuit either by inserting additional SWAP gates or by reversing existing CNOT gates to satisfy the connectivity constraints of the target topology. The depth of the transformed quantum circuits is minimized by utilizing the Monte Carlo tree search (MCTS) to perform qubit routing by making it both construct each action and search over the space of all actions. It is aided in performing these tasks by a Graph neural network that evaluates the value function and action probabilities for each state. Along with this, we propose a new method of adding mutex-lock like variables in our state representation which helps factor in the parallelization of the scheduled operations, thereby pruning the depth of the output circuit. Overall, our procedure (referred to as *QRoute*) performs qubit routing in a hardware agnostic manner, and it outperforms other available qubit routing implementations on various circuit benchmarks.

Introduction

The present-day quantum computers, more generally known as Noisy Intermediate-Scale quantum (NISQ) devices (Preskill 2018) come in a variety of hardware architectures (IBM 2021; Arute et al. 2019; Karalekas et al. 2020; Bourassa et al. 2021), but there exist a few problems pervading across all of them. These problems constitute the poor quality of qubits, limited connectivity between qubits, and the absence of error-correction for noise-induced errors encountered during the execution of gate operations. These place a considerable restriction on the number of instructions that can be executed to perform useful quantum computation (Preskill 2018). Collectively these instructions can be realized as a sequential series of one or two-qubit gates that can be visualized more easily as a quantum circuit as shown in Fig. 1a (Childs, Schoute, and Unsal 2019).

To execute an arbitrarily given quantum circuit on the target quantum hardware, a compiler routine must transform it

to satisfy the connectivity constraints of the topology of the hardware (Cowtan et al. 2019). These transformations usually include the addition of SWAP gates and the reversal of existing CNOT gates. This ensures that any non-local quantum operations are performed only between the qubits that are nearest-neighbors. This process of circuit transformation by a compiler routine for the target hardware is known as qubit routing (Cowtan et al. 2019). The output instructions in the transformed quantum circuit should follow the connectivity constraints and essentially result in the same overall unitary evolution as the original circuit (Pozzi et al. 2020).

In the context of NISQ hardware, this procedure is of extreme importance as the transformed circuit will, in general, have higher depth due to the insertion of extra SWAP gates. This overhead in the circuit depth becomes more prominent due to the high decoherence rates of the qubits and it becomes essential to find the most optimal and efficient strategy to minimize it (Cowtan et al. 2019; Herbert and Sen Gupta 2018; Pozzi et al. 2020). In this article, we present a procedure that we refer to as *QRoute*. We use Monte Carlo tree search (MCTS), which is a look-ahead search algorithm for finding optimal decisions in the decision space guided by a heuristic evaluation function (Kocsis and Szepesvari 2006; Munos 2014). We use it for explicitly searching the decision space for depth minimization and as a stable and performant machine learning setting. It is aided by a Graph neural network (GNN) (Wang et al. 2018), with an architecture that is used to learn and evaluate the heuristic function that will help guide the MCTS.

Qubit Routing

In this section, we begin by defining the problem of qubit routing formally and discussing the work done previously in the field.

Describing the Problem

The topology of quantum hardware can be visualized as a qubit connectivity graph (Fig. 2). Each node in this graph would correspond to a physical qubit which in turn might correspond to a logical qubit. The quantum instruction set, which is also referred to as quantum circuit (Fig. 1a), is a sequential series of single-qubit and two-qubit gate operations that act on the logical qubits. The two-qubit gates such as CNOT can only be performed between two logical qubits

*Animesh Sinha and Utkarsh Azad contribute equally to this work and share first authorship. Corresponding author is Utkarsh Azad. Copyright © 2022, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

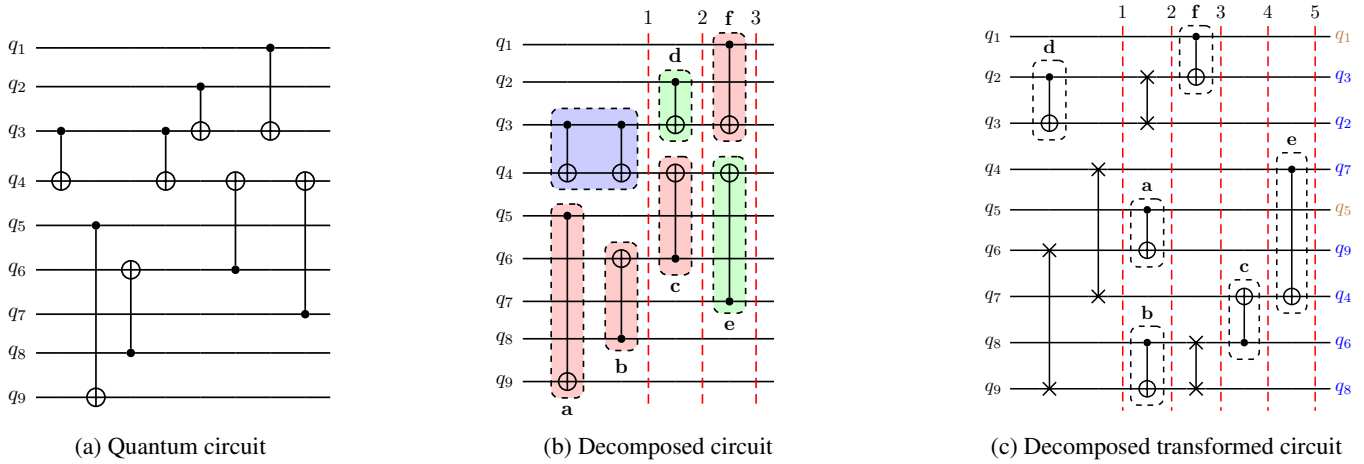


Figure 1: An example of qubit routing on a quantum circuit for 3×3 grid architecture (Figure 2a). (a) For simplicity, the original quantum circuit consists only of two-qubit gate operations. (b) Decomposition of the original quantum circuit into series of slices such that all the instructions present in a slice can be executed in parallel. The two-qubit gate operations: $\{d, e\}$ (green) comply with the topology of the grid architecture whereas the operations: $\{a, b, c, f\}$ (red) do not comply with the topology (and therefore cannot be performed). Note that the successive two-qubit gate operations on $q_3 \rightarrow q_4$ (blue) are redundant and are not considered while routing. (c) Decomposition of the transformed quantum circuit we get after qubit routing. Four additional SWAP gates are added that increased the circuit depth to 5, i.e., an overhead circuit depth of 2. The final qubit labels are represented at the end right side of the circuit. The qubits that are not moved (or swapped) are shown in brown ($\{q_1, q_5\}$), while the rest of them are shown in blue.

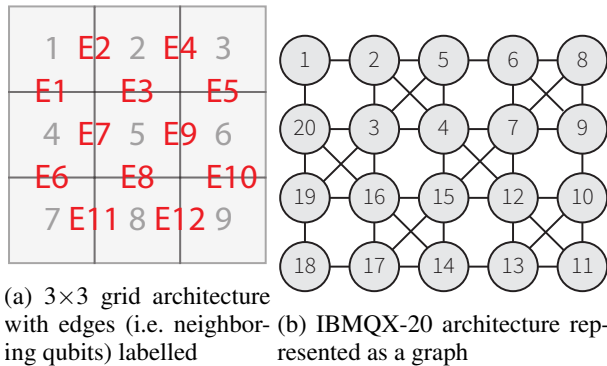


Figure 2: Examples of qubit connectivity graphs for some common quantum architectures

iff there exists an edge between the nodes that correspond to the physical qubits, (Herbert and Sengupta 2018). This edge could be either unidirectional or bidirectional, i.e., CNOT can be performed either in one direction or in both directions. In this work, we consider only the bidirectional case, while noting that the direction of a CNOT gate can be reversed by sandwiching it between a pair of Hadamard gates acting on both control and target qubits (Garcia-Escartin and Chamorro-Posada 2011).

Given a target hardware topology \mathcal{D} and a quantum circuit \mathcal{C} , the task of qubit routing refers to transforming this quantum circuit by adding a series of SWAP gates such that all its gate operations then satisfy the connectivity constraints of the target topology (Fig. 1c). Formally, for a routing algo-

rithm R , we can represent this process as follows:

$$R(\mathcal{C}, \mathcal{D}) \rightarrow \mathcal{C}' \quad (1)$$

Depth of \mathcal{C}' (transformed quantum circuit) will, in general, be more than that of the original circuit due to the insertion of additional SWAP gates. This comes from the definition of the term *depth* in the context of quantum circuits. This can be understood by decomposing a quantum circuit into series of individual slices, each of which contains a group of gate operations that have no overlapping qubits, i.e., all the instructions present in a slice can be executed in parallel (Fig. 1b). The depth of the quantum circuit then refers to the minimum number of such slices the circuit can be decomposed into, i.e., the minimum amount of parallel executions needed to execute the circuit. The goal is to minimize the overhead depth of the transformed circuit with respect to the original circuit.

This goal involves solving two subsequent problems of (i) qubit allocation, which refers to the mapping of program qubits to logic qubits, and (ii) qubit movement, which refers to routing qubits between different locations such that interaction can be made possible (Tannu and Qureshi 2019). In this work, we focus on the latter problem of qubit movement only and refer to it as qubit routing. However, it should be noted that qubit allocation is also an important problem and it can play an important role in minimizing the effort needed to perform qubit movement.

Related Work

The first major attraction for solving the qubit routing problem was the competition organized by IBM in 2018 to find

the best routing algorithm. This competition was won by Zulehner, Paler, and Wille (2019), for developing a Computer Aided Design-based (CAD) routing strategy. Since then, this problem has been presented in many different ways. These include graph-based architecture-agnostic solution by Cowtan et al. (2019), showing equivalence to the travelling salesman problem by Paler, Zulehner, and Wille (2021), machine learning based methods by Paler et al. (2020), and heuristic approaches by Venturelli et al. (2017), Baiocchi, Rasconi, and Oddi (2021), Chand et al. (2019), etc. A reinforcement learning in a combinatorial action space solution was proposed by Herbert and Sengupta (2018), which suggested used simulated annealing to search through the combinatorial action space, aided by a Feed-Forward neural network to judge the long-term expected depth. This was further extended to use Double Deep Q-learning and prioritized experience replay by Pozzi et al. (2020).

Recently, Monte Carlo tree search (MCTS), a popular reinforcement learning algorithm (Browne et al. 2012) previously proven successful in a variety of domains like playing puzzle games such as Chess and Go (Silver et al. 2016), and was used by Zhou, Feng, and Li (2020) to develop a qubit routing solution.

Our Contributions

Our work demonstrates the use of MCTS on the task of Qubit Routing and presents state of the art results. Following are the novelties of this approach:

- We use an array of mutex locks to represent the current state of parallelization, helping to reduce the depth of the circuits instead of the total quantum volume, in contrast to previous use of MCTS for qubit routing in Zhou, Feng, and Li (2020).
- The actions in each timestep (layer of the output circuit) belong to an innumerable large action space. We phrase the construction of such actions as a Markov decision process, making the training stabler and the results better, particularly at larger circuit sizes, than those obtained by performing simulated annealing to search in such action spaces (Herbert and Sengupta 2018; Pozzi et al. 2020). Such approach should be applicable to other problems of a similar nature.
- Graph neural networks are used as an improved architecture to help guide the tree search.

Finally, we provide a simple python package containing the implementation of QRoute, together with an easy interface for trying out different neural net architectures, combining algorithms, reward structures, etc.

Method

The QRoute algorithm takes in an input circuit and an injective map, $\mathcal{M} : Q \rightarrow N$, from logical qubits to nodes (physical qubits). Iteratively, over multiple timesteps, it tries to schedule the gate operations that are present in the input circuit onto the target hardware. To do so, from the set of unscheduled gate operations, \mathcal{P} , it takes all the current operations, which are the first unscheduled operation for both

the qubits that they act on, and tries to make them into local operations, which are those two-qubit operations that involve qubits that are mapped to nodes connected on the target hardware.

In every timestep t , QRoute starts by greedily scheduling all the operations that are both current and local in \mathcal{P} . To evolve \mathcal{M} , it then performs a Monte Carlo tree search (MCTS) to find an optimal set of SWAPs by the evaluation metrics described in the Section such that all operations in the current timestep put together form a parallelizable set, i.e., a set of local operations such that no two operations in the set act on the same qubit. The number of states we can encounter in the action space explodes exponentially with the depth of our search, therefore an explicit search till the circuit is done compiling is not possible. Therefore we cut short our search at some shallow intermediate state, and use a neural network to get its heuristic evaluation.

The following subsections describe in greater detail the working of the search and the heuristic evaluation.

State and Action Space

Definition 0.1 (State) *It captures entire specification of the state of compilation at some timestep t . Abstractly, it is described as:*

$$s_t = (\mathcal{D}, \mathcal{M}_t, \mathcal{P}_t, \mathcal{L}_t) \quad (2)$$

where, \mathcal{D} is the topology of target hardware, and \mathcal{M}_t and \mathcal{P}_t represents the current values of \mathcal{M} and \mathcal{P} respectively. \mathcal{L}_t is the set of nodes that are locked by the gate operations from the previous timestep and therefore cannot be operated in the current timestep.

Definition 0.2 (Action) *It is a set of SWAP gates (represented by the pair of qubits it acts on) such that all gates are local, and its union with the set of operations that were scheduled in the same timestep forms a parallelizable set.*

We are performing a tree search over state-action pairs. Since the number of actions that can be taken at any timestep is exponential in the number of connections on the hardware, we are forced to build a single action up, step-by-step.

Definition 0.3 (Move) *It is a single step in a search procedure which either builds up the action or applies it to the current state. Moves are of the following two types:*

1. **SWAP**(n_1, n_2): *Inserts a new SWAP on nodes n_1 and n_2 into the action set. Such an insertion is only possible if the operation is local and resulting set of operations for the timestep form a parallel set.*
2. **COMMIT**: *Finishes the construction of the action set for that timestep. It also uses the action formed until now to update the state s_t (schedules the SWAP gates on the hardware), and resets the action set for the next step.*

In reality, different gate operations take different counts of timesteps for execution. For example, if a hardware requires SWAP gate to be broken down into CNOT gates, then it would take three timesteps for complete execution (Garcia-Escartin and Chamorro-Posada 2011). This means, operations which are being scheduled must maintain mutual exclusivity with other other operations over the nodes

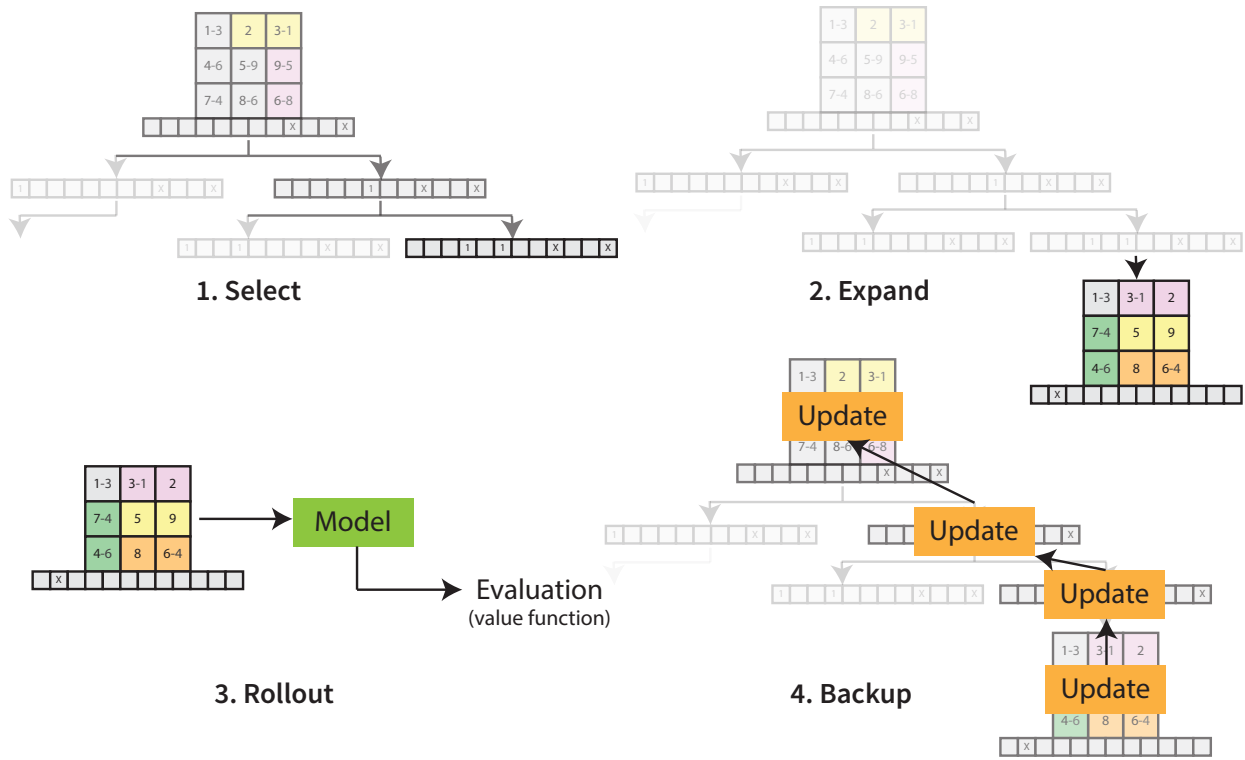


Figure 3: Iteration of a Monte Carlo tree search: (i) select - recursively choosing a node in the search tree for exploration using a selection criteria, (ii) expand - expanding the tree using the available actions if an unexplored leaf node is selected, (iii) rollout - estimating long term reward using a neural network for the action-state pair for the new node added to search tree, and (iv) backup - propagating reward value from the leaf nodes upwards to update the evaluations of its ancestors.

which participates in them. This is essential to minimizing the depth of the circuit since it models parallelizability of operations.

However, constructing a parallelizable set and representing the state of parallelization to our heuristic evaluator is a challenge. But an analogy can be drawn here to the nodes being thought of as “resources” that cannot be shared, and the operations as “consumers” (Dijkstra 2002). This motivates us to propose the use of Mutex Locks for this purpose. These will lock a node until a scheduled gate operation involving that node executes completely. Therefore, this allows our framework to naturally handle different types of operations which take different amounts of time to complete.

For every state-action pair, the application of a feasible move m on it will result in a new state-action pair: $(s, a) \xrightarrow{m} (s', a')$. This is a formulation of the problem of search as a Markov Decision Process. Associated with each such state-action-move tuple $((s, a), m)$, we maintain two additional values that are used by MCTS:

1. *N-value* - The number of times we have taken the said move m from said state-action pair (s, a) .
2. *Q-value* - Given a reward function \mathcal{R} , it is the average long-term reward expected after taking said move m over all iterations of the search. (Future rewards are discounted by a factor γ)

$$Q((s, a), m) = \mathcal{R}((s, a), m) + \gamma \frac{\sum_{m'} N((s', a'), m') \cdot Q((s', a'), m')}{\sum_{m'} N((s', a'), m')} \quad (3)$$

Monte Carlo Tree Search

Monte Carlo tree search progresses iteratively by executing its four phases: select, expand, rollout, and backup as illustrated in Fig. 3. In each iteration, it begins traversing down the existing search tree by selecting the node with the maximum UCT value (Eq. 4) at each level. During this traversal, whenever it encounters a leaf node, it expands the tree by choosing a move m from that leaf node. Then, it estimates the scalar evaluation for the new state-action pair and back-propagates it up the tree to update evaluations of its ancestors.

To build an optimal action set, we would want to select the move m with the maximum true Q-value. But since true Q-values are intractably expensive to compute, we can only approximate the Q-values through efficient exploration. We use the Upper Confidence Bound on Trees (UCT) objective (Kocsis and Szepesvari 2006) to balance exploration and exploitation as we traverse through the search tree. Moreover, as this problem results in a highly asymmetric tree, since

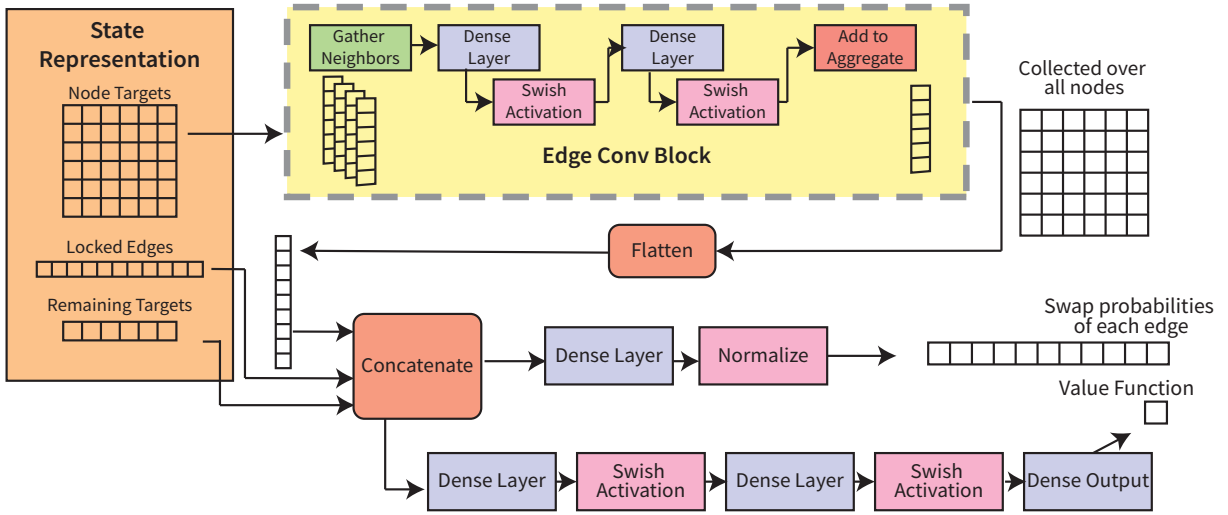


Figure 4: Graph neural network architecture that approximates the value function and the policy function.

some move block a lot of other moves, while others block fewer moves, we use the formulation of UCT adapted for asymmetric trees (Moerland et al. 2018):

$$\text{UCT}((s, a), m) = Q((s, a), m) + c \frac{\sqrt{\sum_m N((s, a), m)}}{N((s, a), m)} \times p(m|(s, a)) \quad (4)$$

Here, the value $p(m|(s, a))$ is the prior policy function, which is obtained by adding a Dirichlet noise to the policy output of the neural network (Silver et al. 2017). As MCTS continues probing the action space, it gets a better estimate of the true values of the actions. This means that it acts as a policy enhancement function whose output policy (Eq. 5) can be used to train the neural network’s prior (π), and the average Q-value computed can be used to train its scalar evaluation (Eq. 6).

$$\pi(m|(s, a)) \propto N((s, a), m) \quad (5)$$

$$\mathcal{V}((s, a)) = \frac{\sum_m Q((s, a), m)}{\sum_m N((s, a), m)} \quad (6)$$

The details of how MCTS progresses have been elaborated in the supplementary. Once it gets terminated, i.e., the search gets completed, we go down the tree selecting the child with the maximum Q-value at each step until a COMMIT action is found, we use the action set of the selected state-action pair to schedule SWAPs for the current timestep, and we re-root the tree at the child node of the COMMIT action to prepare for the next timestep.

Neural Network Architecture

Each iteration of the MCTS requires evaluation of Q-values for a newly encountered state-action pair. But these values are impossible to be computed exactly since it would involve an intractable number of iterations in exploring and

expanding the complete search tree. Therefore, it is favorable to heuristically evaluate the expected long-term reward from the state-action pair using a Neural Network, as it acts as an excellent function approximator that can learn the symmetries and general rules inherent to the system.

So, once the MCTS sends a state-action pair to the evaluator, it begins by committing the action to the state and getting the resultant state. We then generate the following featured representation of this state and pass this representation through the neural-network architecture as shown in Fig. 4.

1. *Node Targets* - It is a square boolean matrix whose rows and columns correspond to the nodes on a target device. An element (i, j) is true iff some logical qubits q_x and q_y are currently mapped to nodes i and j respectively, such that (q_x, q_y) is the first unscheduled operation that q_x partakes in.
2. *Locked Edges* - It is a set of edges (pairs of connected nodes) that are still locked due to either of its qubits being involved in an operation in the current timestep or another longer operation that hasn’t yet terminated from the previous timesteps.
3. *Remaining targets* - It is a list of the number of gate-operations that are yet to be scheduled for each logical qubit.

The SWAP operations each qubit would partake in depends primarily on its target node, and on those of the nodes in its neighborhood that might be competing for the same resources. It seems reasonable that we can use a Graph Neural Network with the device topology graph for its connectivity since the decision of the optimal SWAP action for some node is largely affected by other nodes in its physical neighborhood. Therefore, our architecture includes an edge-convolution block (Wang et al. 2018), followed by some fully-connected layers with Swish (Ramachandran, Zoph, and Le 2017) activations for the policy and value heads. The

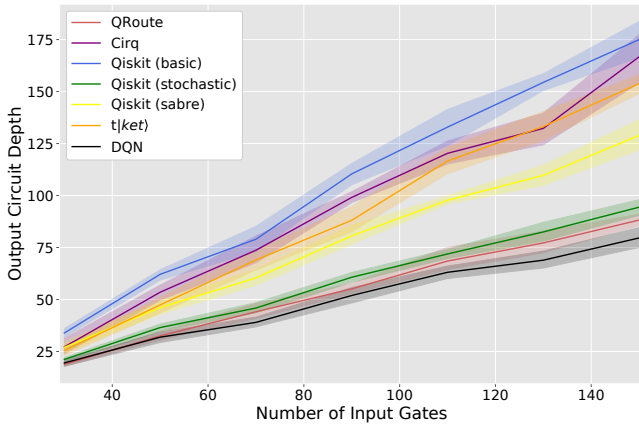


Figure 5: Comparative performance of routing algorithms on random circuits as a function of the number of two-qubit operations in the circuit.

value function and the policy function computed from this neural network are returned back to the MCTS.

Results

We compare QRoute against the routing algorithms from other state-of-the-art frameworks on various circuit benchmarks: (i) Qiskit and its three variants (Anis et al. 2021): (a) basic, (b) stochastic, and (c) sabre, (ii) Deep-Q-Networks (DQN) from (Pozzi et al. 2020), (iii) Cirq (Gidney et al. 2021), and (iv) t|ket from Cambridge Quantum Computing (CQC) (Sivarajah et al. 2020). Qiskit’s transpiler uses gate commutation rules while perform qubit routing. This strategy is shown to be advantageous in achieving lower circuit depths (Itoko et al. 2019) but was disabled in our simulations to have a fair comparison. The results for DQN shown are adapted from the data provided by the authors Pozzi et al. (2020).

Random Test Circuits

The first benchmark for comparing our performance comprises of random circuits. These circuits are generated on the fly and initialized with the same number of qubits as there are nodes on the device. Then two-qubit gates are put between any pair of qubits chosen at random. In our simulations, the number of such gates is varied from 30 to 150 and the results for assessing performance of different frameworks are given in Fig. 5. The experiments were repeated 10 times on each circuit size, and final results were aggregated over this repetition.

Amongst the frameworks compared, QRoute ranks a very close second only to Deep-Q-Network guided simulated annealer (DQN). Nevertheless, QRoute still does consistently better than all the other major frameworks: Qiskit, Cirq and t|ket, and it scales well when we increase the number of layers and the layer density in the input circuit. QRoute shows equivalent performance to DQN on smaller circuits, and on the larger circuits it outputs depths which are on average ≤ 4 layers more than those of DQN. Some part of this can be at-

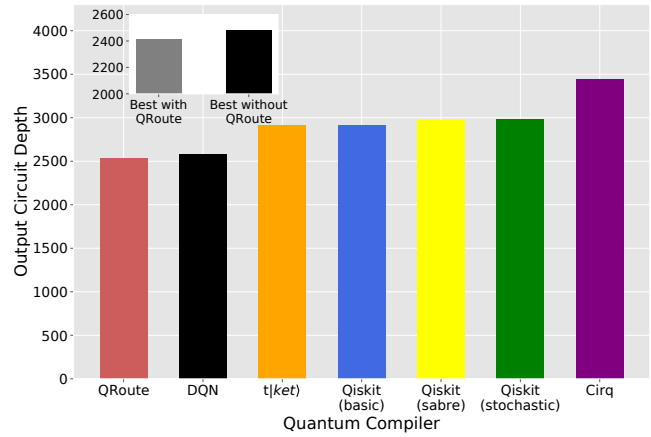


Figure 6: Plots of output circuit depths of routing algorithms over small realistic circuits (≤ 100 gates), summed over the entire dataset. The inset shows the results on the same data comparing the best performing scheduler excluding and including QRoute on each circuit respectively.

tributed to MCTS, in it’s limited depth search, choosing the worse of two moves with very close Q-values, resulting in the scheduling of some unnecessary SWAP operations.

Small Realistic Circuits

Next we test on the set of all circuits which use 100 or less gates from the IBM-Q realistic quantum circuit dataset used by Zulehner, Paler, and Wille (2018). The comparative performance of all routing frameworks has been shown by plotting the depths of the output circuits summed over all the circuits in the test set in Fig. 6. Since the lack of a good initial qubit allocation becomes a significant problem for all pure routing algorithms on small circuits, we have benchmarked QRoute on this dataset from three trials with different initial allocations.

The model presented herein has the best performance on this dataset. We also compare the best result from a pool of all routers including QRoute against that of another pool of the same routers but excluding QRoute. The pool including QRoute gives on average 2.5% lower circuit depth, indicating that there is a significant number of circuits where QRoute is the best routing method available.

On this dataset also, closest to QRoute performance is shown by Deep-Q-Network guided simulated annealer. To compare performances, we look at the average circuit depth ratio (CDR), which is defined by (Pozzi et al. 2020):

$$\text{CDR} = \frac{1}{\#\text{circuits}} \sum_{\text{circuits}} \frac{\text{Output Circuit Depth}}{\text{Input Circuit Depth}} \quad (7)$$

The resultant CDR for QRoute is 1.178, where as the reported CDR for the DQN is 1.19. In fact, QRoute outperforms DQN on at least 80% of the circuits. This is significant because in contrast to the random circuit benchmark, the realistic benchmarks consist of the circuits that are closer to the circuits used in useful computation.

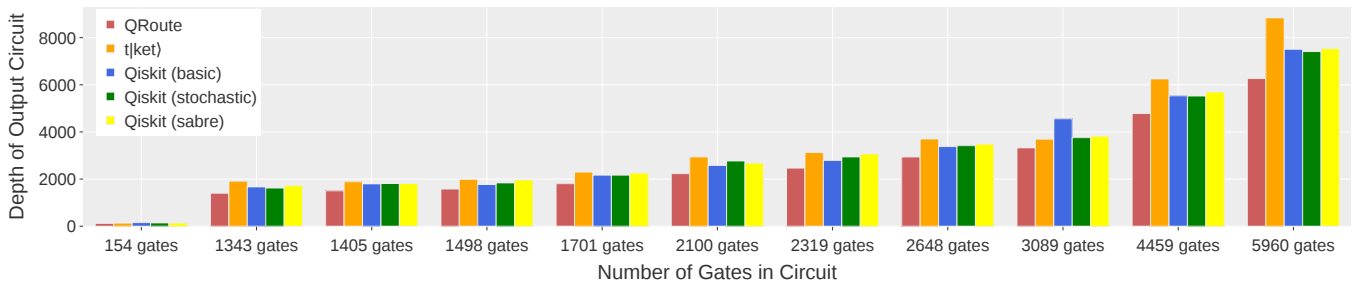


Figure 7: The results over eight circuits sampled from the large realistic dataset benchmark, the outputs of each routing algorithm are shown for every circuit.

Large Realistic Circuit

For final benchmark, we take eight large circuits ranging from 154 gates to 5960 gates in its input from the IBM-Quantum realistic test dataset (Zulehner, Paler, and Wille 2018). The results are plotted in Fig. 7. QRoute has the best performance of all available routing methods: Qiskit and t|ket), on every one of these sampled circuits with on an average 13.6% lower circuit depth, and notable increase in winning difference on the larger circuits.

The results from DQN and Cirq are not available for these benchmarks as they are not designed to scale to such huge circuits. In case of DQN, the CDR data results were not provided for the circuits over 200 gates, mainly because simulated annealing used in it is computationally expensive. Similarly, for Cirq, it takes several days to compile each of the near 5000 qubit circuits. In contrast, QRoute is able to compile these circuits in at most 4 hours, and its compilation process can be sped up by reducing the depth of the search. Spending more time, however, helps MCTS to better approximate the Q-values leading to circuits with lower resulting depth.

Discussion and Conclusion

In this article, we have shown that the problem of qubit routing has a very powerful and elegant formulation in Reinforcement Learning (RL) which can surpass the results of any classical heuristic algorithm across all sizes of circuits and types of architectures. Furthermore, the central idea of building up solutions step-by-step when searching in combinatorial action spaces and enforcing constraints using mutex locks, can be adapted for several other combinatorial optimization problems (Mazyavkina et al. 2020; Xing and Tu 2020; Xu and Lieberherr 2019; Abe et al. 2019; Larterre et al. 2018). Our approach is flexible enough to compile circuits of any size onto any device, from small ones like IBMQX20 with 20 qubits, to much larger hardware like Google Sycamore (results provided in supplementary) with 53 qubits (the Circuit Depth Ratio for small realistic circuits on Google Sycamore was 1.64). Also, it intrinsically deals with hardware having different primitive instruction set, for example on hardware where SWAP gates are not a primitive and they get decomposed to 3 operations. QRoute enjoys significant tunability; hyperparameters can be changed easily to alter the tradeoff between time taken and optimality of

decisions, exploration and exploitation, etc.

QRoute is a reasonably fast method, taking well under 10 minutes to route a circuit with under 100 operations, and at most 4 hours for those with upto 5000 operations, when tested on a personal machine with an i3 processor (3.7 GHz) and no GPU acceleration. Yet more can be desired in terms of speed. However, it is hard to achieve any significant improvement without reducing the number of search iterations and trading off a bit of performance. More predictive neural networks can help squeeze in better speeds.

One of the challenges of methods like DQN, that use Simulated Annealing to build up their actions is that the algorithm cannot plan for the gates which are not yet waiting to be scheduled, those which will come to the head of the list once the gates which are currently waiting are executed (Pozzi et al. 2020). QRoute also shares this deficiency, but the effect of this issue is mitigated by the explicit tree search which takes into account the rewards that will be accrued in the longer-term future. There is scope to further improve this by feeding the entire list of future targets directly into our neural network by using transformer encoders to handle the arbitrary length sequence data. This and other aspects of neural network design will be a primary facet of future explorations. Another means of improving the performance would be to introduce new actions by incorporating use of BRIDGE gates (Itoko et al. 2019) and gate commutation rules (Garcia-Escartin and Chamorro-Posada 2011) alongside currently used SWAP gates. The advantage of former is that it allows running CNOT gates on non-adjacent qubit without permuting the ordering of the logical qubits; whereas, the latter would allow MCTS to recognize the redundancy in action space, making its exploration and selection more efficient.

Finally, we provide an open-sourced access to our software library. It will allow researchers and developers to implement variants of our methods with minimal effort. We hope that this will aid future research in quantum circuit transformations. For review we are providing, the codebase and a multimedia in the supplementary.

On the whole, the Monte Carlo Tree Search for building up solutions in combinatorial action spaces has exceeded the current state of art methods that perform qubit routing. Despite its success, we note that QRoute is a primitive implementation of our ideas, and there is great scope of improvement in future.

Acknowledgements

A. S. and U. A. have contributed equally to the work presented in this manuscript. A. S. and U. A. would like to acknowledge the assistance from their colleague Bhuvanesh Sridharan for his help in implementing and refining MCTS. Additionally, they would also like to thank their colleagues Jai Bardhan, Kalp Shah for their suggestions.

References

- Abe, K.; Xu, Z.; Sato, I.; and Sugiyama, M. 2019. Solving NP-Hard Problems on Graphs with Extended AlphaGo Zero. arXiv:1905.11623.
- Anis, M. S.; Abraham, H.; AduOffei; Agliardi, G.; Aharoni, M.; Akhalwaya, I. Y.; and et al. 2021. Qiskit: An Open-source Framework for Quantum Computing. *Zenodo*.
- Arute, F.; Arya, K.; Babbush, R.; and et al. 2019. Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779): 505–510.
- Baiocchi, M.; Rasconi, R.; and Oddi, A. 2021. A Novel Ant Colony Optimization Strategy for the Quantum Circuit Compilation Problem. In *EvoCOP*, 1–16.
- Bourassa, J. E.; Alexander, R. N.; Vasmer, M.; Patil, A.; Tzitrin, I.; Matsuura, T.; Su, D.; Baragiola, B. Q.; Guha, S.; Dauphinais, G.; Sabapathy, K. K.; Menicucci, N. C.; and Dhand, I. 2021. Blueprint for a Scalable Photonic Fault-Tolerant Quantum Computer. *Quantum*, 5: 392.
- Browne, C. B.; Powley, E.; Whitehouse, D.; Lucas, S. M.; Cowling, P. I.; Rohlfshagen, P.; Tavener, S.; Perez, D.; Samothrakis, S.; and Colton, S. 2012. A Survey of Monte Carlo Tree Search Methods. *IEEE Transactions on Computational Intelligence and AI in Games*, 4(1): 1–43.
- Chand, S.; Singh, H. K.; Ray, T.; and Ryan, M. 2019. Rollout based heuristics for the quantum circuit compilation problem. In *2019 IEEE Congress on Evolutionary Computation (CEC)*, 974–981. IEEE.
- Childs, A. M.; Schoute, E.; and Unsal, C. M. 2019. Circuit Transformations for Quantum Architectures. In *14th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2019)*, volume 135 of *Leibniz International Proceedings in Informatics (LIPIcs)*, 3:1–3:24. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. ISBN 978-3-95977-112-2.
- Cowtan, A.; Dilkes, S.; Duncan, R.; Krajenbrink, A.; Simmons, W.; and Sivarajah, S. 2019. On the Qubit Routing Problem. In *14th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2019)*, volume 135 of *Leibniz International Proceedings in Informatics (LIPIcs)*, 5:1–5:32. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. ISBN 978-3-95977-112-2.
- Dijkstra, E. W. 2002. *Cooperating Sequential Processes*, 65–138. New York, NY: Springer New York. ISBN 978-1-4757-3472-0.
- Garcia-Escartin, J. C.; and Chamorro-Posada, P. 2011. Equivalent Quantum Circuits. arXiv:1110.2998.
- Gidney, C.; Bacon, D.; Pato, B.; Strain, D.; Neeley, M.; Zalcman, A.; and et al. 2021. Cirq. *Zenodo*.
- Herbert, S.; and Sengupta, A. 2018. Using Reinforcement Learning to find Efficient Qubit Routing Policies for Deployment in Near-term Quantum Computers. arXiv:1812.11619.
- IBM. 2021. IBM Quantum. <https://quantum-computing.ibm.com/>.
- Itoko, T.; Raymond, R.; Imamichi, T.; and Matsuo, A. 2019. Optimization of Quantum Circuit Mapping using Gate Transformation and Commutation. arXiv:1907.02686.
- Karalekas, P. J.; Tezak, N. A.; Peterson, E. C.; Ryan, C. A.; da Silva, M. P.; and Smith, R. S. 2020. A quantum-classical cloud platform optimized for variational hybrid algorithms. *Quantum Sci. Technol.*, 5: 024003.
- Kocsis, L.; and Szepesvari, C. 2006. Bandit Based Monte-Carlo Planning. In *European Conference on Machine Learning*, 282–293. Springer.
- Latterre, A.; Fu, Y.; Khalil Jabri, M.; Cohen, A.-S.; Kas, D.; Hajjar, K.; Dahl, T. S.; Kerkeni, A.; and Beguir, K. 2018. Ranked Reward: Enabling Self-Play Reinforcement Learning for Combinatorial Optimization. arXiv:1807.01672.
- Mazyavkina, N.; Sviridov, S.; Ivanov, S.; and Burnaev, E. 2020. Reinforcement Learning for Combinatorial Optimization: A Survey. arXiv:2003.03600.
- Moerland, T. M.; Broekens, J.; Plaats, A.; and Jonker, C. M. 2018. Monte Carlo Tree Search for Asymmetric Trees. *arXiv e-prints*.
- Munos, R. 2014. From Bandits to Monte-Carlo Tree Search: The Optimistic Principle Applied to Optimization and Planning. *Found. Trends Mach. Learn.*, 7: 1–129.
- Paler, A.; Sasu, L. M.; Florea, A.; and Andonie, R. 2020. Machine Learning Optimization of Quantum Circuit Layouts. arXiv:2007.14608.
- Paler, A.; Zulehner, A.; and Wille, R. 2021. NISQ circuit compilation is the travelling salesman problem on a torus. *Quantum Science and Technology*, 6(2): 025016.
- Pozzi, M. G.; Herbert, S. J.; Sengupta, A.; and Mullins, R. D. 2020. Using Reinforcement Learning to Perform Qubit Routing in Quantum Compilers. arXiv:2007.15957.
- Preskill, J. 2018. Quantum computing in the NISQ era and beyond. *Quantum*, 2: 79.
- Ramachandran, P.; Zoph, B.; and Le, Q. V. 2017. Searching for Activation Functions. arXiv:1710.05941.
- Silver, D.; Huang, A.; Maddison, C. J.; Guez, A.; Sifre, L.; van den Driessche, G.; and et al. 2016. Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587): 484–489.
- Silver, D.; Hubert, T.; Schrittwieser, J.; Antonoglou, I.; Lai, M.; Guez, A.; Lanctot, M.; Sifre, L.; Kumaran, D.; Graepel, T.; Lillicrap, T. P.; Simonyan, K.; and Hassabis, D. 2017. Mastering Chess and Shogi by Self-Play with a General Reinforcement Learning Algorithm. arXiv:1712.01815.
- Sivarajah, S.; Dilkes, S.; Cowtan, A.; Simmons, W.; Edgington, A.; and Duncan, R. 2020. t|ket>: A retargetable compiler for NISQ devices. *Quantum Science and Technology*, 6(1): 014003.

- Tannu, S. S.; and Qureshi, M. K. 2019. Not All Qubits Are Created Equal: A Case for Variability-Aware Policies for NISQ-Era Quantum Computers. In *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*, ASPLOS '19, 987–999. New York, NY, USA: Association for Computing Machinery. ISBN 9781450362405.
- Venturelli, D.; Do, M.; Rieffel, E. G.; and Frank, J. 2017. Temporal Planning for Compilation of Quantum Approximate Optimization Circuits. In *IJCAI*, 4440–4446.
- Wang, Y.; Sun, Y.; Liu, Z.; Sarma, S. E.; Bronstein, M. M.; and Solomon, J. M. 2018. Dynamic Graph CNN for Learning on Point Clouds. arXiv:1801.07829.
- Xing, Z.; and Tu, S. 2020. A Graph Neural Network Assisted Monte Carlo Tree Search Approach to Traveling Salesman Problem. *IEEE Access*, 8: 108418–108428.
- Xu, R.; and Lieberherr, K. 2019. Learning Self-Game-Play Agents for Combinatorial Optimization Problems. arXiv:1903.03674.
- Zhou, X.; Feng, Y.; and Li, S. 2020. A Monte Carlo Tree Search Framework for Quantum Circuit Transformation. In *Proceedings of the 39th International Conference on Computer-Aided Design, ICCAD '20*. New York, NY, USA: Association for Computing Machinery. ISBN 9781450380263.
- Zulehner, A.; Paler, A.; and Wille, R. 2018. IBM Qiskit Developer Challenge. <https://www.ibm.com/blogs/research/2018/08/winners-qiskit-developer-challenge/>. Accessed on: 2021-03-23.
- Zulehner, A.; Paler, A.; and Wille, R. 2019. An Efficient Methodology for Mapping Quantum Circuits to the IBM QX Architectures. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 38(7): 1226–1236.