AALBORG UNIVERSITY
DENMARK

SAPIENZA
UNIVERSITÀ DI ROMA

NEURAL INFORMATION
PROCESSING SYSTEMS

RL-Guided Contraction of Symbolic Tensor Networks for Quantum Circuit Equivalence

Suhaib Al-Rousan[1], Christian Schilling[1], Max Tschaikowski[2], Kim Guldstrand Larsen[1]

[1]Aalborg University, Denmark

[2]Sapienza University of Rome, Italy

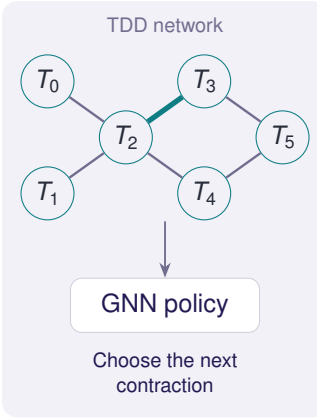

Do the two circuits do the same thing?

Reference circuit C vs. compiled circuit C'

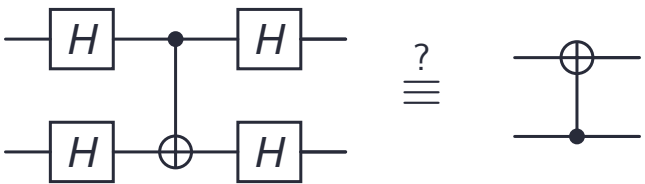

The bottleneck

An n-qubit unitary has $2^n \times 2^n$ entries.

$$C \equiv C' \quad \Longleftrightarrow \quad U_C U_{C'}^{\dagger} = e^{i\theta} I_{2^n}$$

Do the two circuits do the same thing?

Combine C with the inverse implementation $C'^{\dagger}$

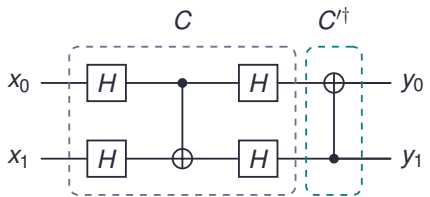

$$C \equiv C' \iff U_C U_{C'}^{\dagger} = e^{i\theta} I_{2^n}$$

Check whether the combined circuit is the identity, up to global phase.

Here C' is self-inverse.

Do the two circuits do the same thing?

Gates become tensors; wires become shared indices

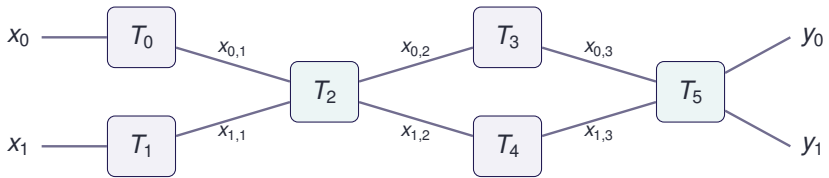

T_0, T_1, T_3, T_4: Hadamard gates. T_2, T_5: the two controlled gates.

$$C \equiv C' \iff U_C U_{C'}^{\dagger} = e^{i\theta} I_{2^n}$$

Contraction sums shared indices: $R_{ik} = \sum_j A_{ij} B_{jk}$.
Our route: represent tensors as compact tensor decision diagrams (TDDs), then contract.

Contraction order controls the peak

TDDs compress tensors by sharing repeated subgraphs.

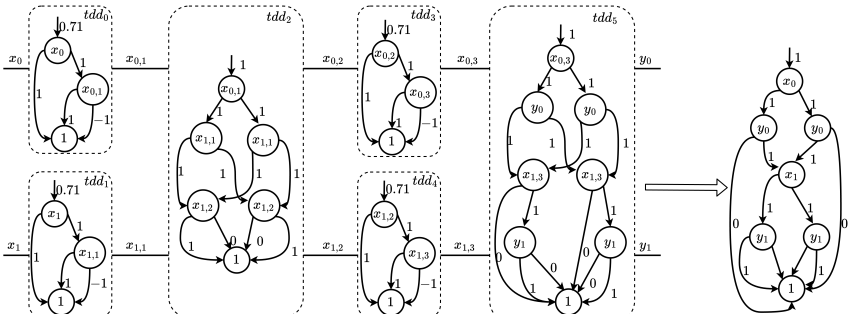

Any contraction order leads to the same result but different intermediate sizes.
A contraction can destroy subgraph sharing.

TDD network for the example circuit. Six TDDs contract to the identity TDD on the right.

Learn one contraction at a time

Observe the TDD network, select an edge, contract, and repeat.

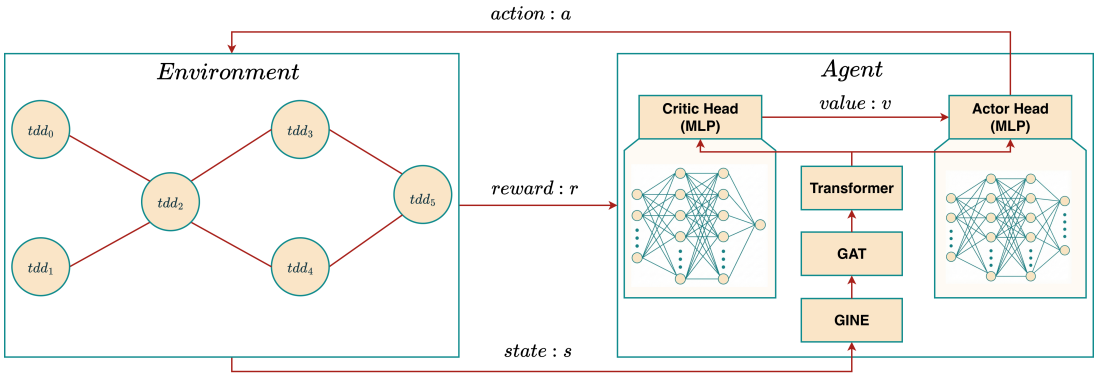

Goal: minimize peak intermediate TDD size (node count).

Evaluation

Benchmarks: DJ, GHZ, W-State, QNN, and unseen-family BV (5–23 qubits).

Baselines

Counting: TDD-specific heuristic.
Greedy / **Random Greedy:** cost-based planning.
Betweenness: graph-centrality ordering.
KaHyPar: hypergraph partitioning.

Measurements

Peak TDD node count
Total time: planning + contraction.
All planners use the same preprocessing and TDD executor.

Paper's filtered evaluation set: held-out / otherwise unseen circuits. Identical preprocessing and TDD execution across planners.

Evaluation

Benchmarks: DJ, GHZ, W-State, QNN, and unseen-family BV (5–23 qubits).

RL Policy | Counting | Greedy | Random Greedy | Betweenness | KaHyPar

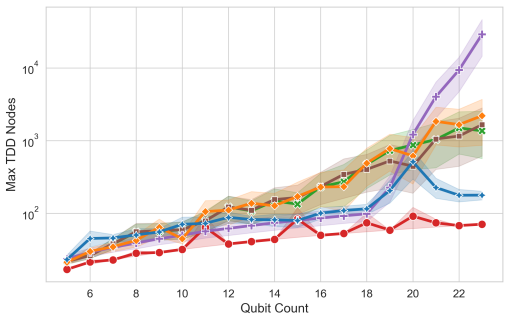

Peak TDD node count (log scale).

Measurements

Peak TDD node count
Total time: planning + contraction.
All planners use the same preprocessing and TDD executor.

Paper's filtered evaluation set: held-out / otherwise unseen circuits. Identical preprocessing and TDD execution across planners.

Evaluation

Benchmarks: DJ, GHZ, W-State, QNN, and unseen-family BV (5–23 qubits).

RL Policy Counting Greedy Random Greedy Betweenness KaHyPar

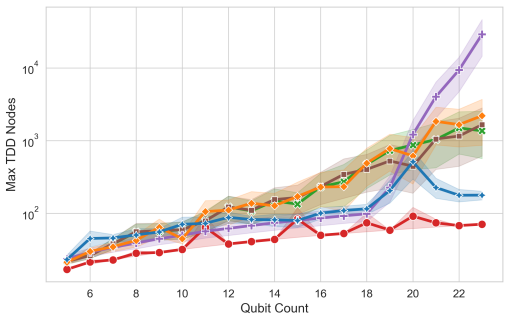

Peak TDD node count (log scale).

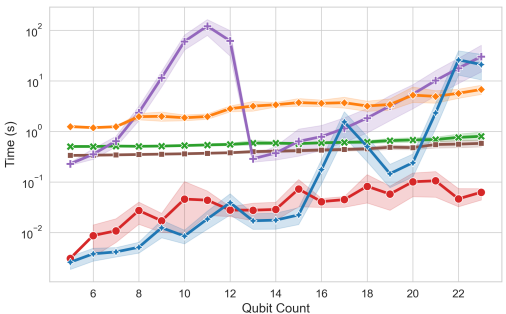

Planning + contraction time (log scale).

RL-Policy mean total time: 0.047 s. Low aggregate peak TDD size.

Paper's filtered evaluation set: held-out / otherwise unseen circuits. Identical preprocessing and TDD execution across planners.

Scaling beyond the training qubit range

Two regimes

Pure RL Policy: 45-qubit DJ, GHZ, QNN, and BV, when their graphs fit.

Hybrid: Start with the Counting method, then switch to RL Policy.

8.4 s
at **50 qubits**, with **15,289** peak TDD nodes.
All tested baselines exceed 120 s.

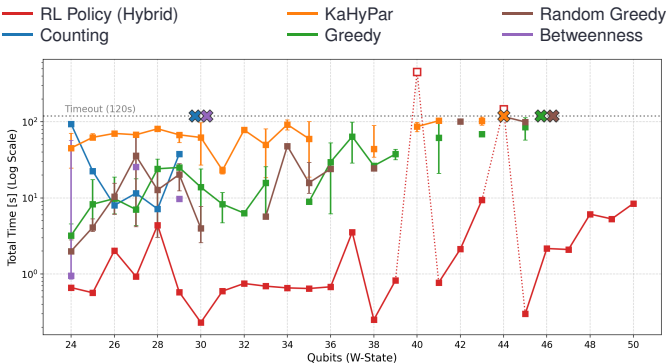

25 of 27 W-State sizes finish within 120 s.

24–50 qubits; log scale. Hollow RL Policy points at 40/44 qubits exceeded 120 s and were re-measured with a 900 s limit.

Transfer in qubit count does not imply transfer beyond the graph-size bounds.

Takeaway

Learn contraction orders to reduce peak TDD size.

- TDD-aware features guide a graph policy.
- Adapt contraction choices as the network evolves.
- The final TDD determines equivalence.

Current scope: structured circuit families; large graphs use a hybrid.
Transfer to unstructured random circuits remains limited.



Code + data

Thank you

Selected references

Hong et al. (2022). A Tensor Network based Decision Diagram for Representation of Quantum Circuits. *ACM TODAES*, 27(6).

Zhang et al. (2024). Quantum circuit simulation with fast tensor decision diagram. *ISQED*.

Larsen et al. (2024). Contraction heuristics for tensor decision diagrams. *Entropy*, 26(12).

Meirom et al. (2022). Optimizing Tensor Network Contraction Using Reinforcement Learning. *ICML*.

This presentation: Al-Rousan, Schilling, Tschaikowski, and Larsen. *RL-Guided Contraction of Symbolic Tensor Networks for Quantum Circuit Equivalence.*

Code and data: github.com/srousan18/rl_tddn_contraction_order