Reproducibility Report for ACM SIGMOD 2023 Paper: "Ready to Leap (by Co-Design)? Join Order Optimisation on Quantum Hardware"

SEBASTIAN SCHMIDL, Hasso Plattner Institute, University of Potsdam, Germany
YANGSHEN DENG, Southern University of Science and Technology, China
PINGCHUAN MA, Hong Kong University of Science and Technology, Hong Kong SAR
MANUEL SCHÖNBERGER, Technical University of Applied Sciences Regensburg, Germany
WOLFGANG MAUERER, Technical University of Applied Sciences Regensburg and Siemens AG, Corporate Research, Germany

The paper "Ready to Leap (by Co-Design)? Join Order Optimisation on Quantum Hardware" proposes the first approach to solve the problem of join order optimization on quantum hardware. The authors characterize the applicability and limitations of current state-of-the-art quantum hardware, i. e. gate-based quantum computing and quantum annealing, for join ordering and recommend key improvements to the physical hardware to reach practical utility. Based on the provided database queries and QPU system processing data, we have been able to reproduce the original paper's key insights and quantum problem characteristics reported in its experimental section. The authors provided a self-contained and fully automated reproduction package, including data (database queries, statistics, and collected QPU processing data), experiment scripts, and plotting routines that allowed the identical reconstruction of the three main figures in the paper.

1 INTRODUCTION

This reproducibility report concerns the paper "Ready to Leap (by Co-Design)? Join Order Optimisation on Quantum Hardware" [2], which presents a quantum implementation of join ordering on gate-based quantum computing and quantum annealing approaches. The paper is joint work between Manuel Schönberger (Technical University of Applied Sciences Regensburg), Stefanie Scherzinger (Passau University), and Wolfgang Mauerer (Technical University of Applied Sciences Regensburg, and Siemens Technology, Munich). The experiments demonstrate the speed-up of their methods compared to the best known classical join ordering approaches, in addition to the limitations of current QPU architectures. The recommended experimental settings using the provided QPU processing data could successfully reproduce the experimental results reported in the paper.

2 SUBMISSION

The paper provides detailed instructions and a fully automated reproduction package for readers to deploy and run the experiments in the form of a Docker image. The reproduction code, data, and instructions are hosted in the GitHub repository https://github.com/lfd/sigmod23-reproduction.

The review process engendered multiple refinements of the reproduction code base. Following the reviewers' suggestions, the authors updated the submission files to improve the usability and guarantee long-term availability and feasibility of the reproduction process through the following changes:

(1) Initially, some users could experience file permission errors when calling the dispatcher scripts. The repository files were updated to inherently feature the correct set of permissions.
(2) The reproduction package documentation was updated to include the steps necessary to access the generated result figures, which are stored within the Docker container. The reviewers suggested two options. To avoid further complicated steps and non-intuitive behavior from mounting the results' directory (option 1), the package authors decided to

include a manual instruction to copy the results from the Docker container to the host file system (option 2).

(3) The authors archived a complete Docker image with the recommended reproduction settings on Zenodo. This avoids any risk incurred by breaking changes resulting from external package updates, which may otherwise jeopardize the reproduction process.

3 HARDWARE AND SOFTWARE ENVIRONMENT

Table 1 summarizes the hardware environments used in the paper and in the reproducibility. The container image includes all software libraries; thus, the software environment is identical for all experiment executions, and is independent of external resources (note that this is the case for configurations (ii) and (iii), since configuration (i) requires access to the different QPUs which are external resources).

Table 1. Hardware environment

	Paper [2]	Reviewer #1	Reviewer #2	Reviewer #3
CPU	AMD	AMD	Intel	Intel
Model	Ryzen 7 PRO 5850U	Ryzen 3970X	Xeon Gold 5318Y	Xeon E5-2630 v4
Cores	8 (16)	32 (64)	2×24 (48)	10 (20)
GHz	4.40	4.50	2.10	2.20
RAM	32 GB	256 GB	504 GB	64 GB
QPU	D-Wave Advantage	–	–	–
Qubits	5000			
QPU	IBM Q Auckland	–	–	–
Qubits	127			

4 REPRODUCIBILITY EVALUATION

4.1 Process

We followed the recommended process provided in Section 5 of the reproducibility instructions [1], which consists of only four steps: (i) configuring the experiments using the file `base/config.py`, (ii) building the Docker image from the repository, (iii) starting the Docker container to run the experiments, and (iv) retrieving the figures from the Docker container.

The reproduction package supports three types of configurations for the experiments: Running on (i) cloud QPU (requires commercial licenses), (ii) local simulator (requires academic licenses), and (iii) collected QPU statistics of the authors' experiment executions. We tested configuration (ii) and (iii). However, configuration (ii) simulates a QPU on a CPU and did not finish the experiments in one month. All three reviewers could successfully execute configuration (iii) within a couple of hours.

4.2 Results

Our local results faithfully match the reported results in the original paper using configuration (iii); the Figures 1a, 1b, and 1c are identical to the corresponding ones in the paper. However, the experiments using configuration (i) were not executed and configuration (ii) did not finish in a reasonable amount of time (one month).

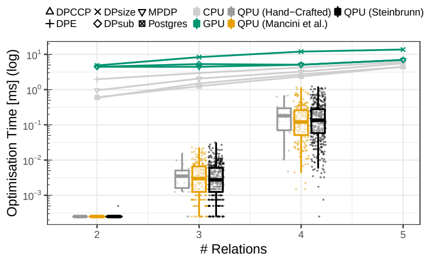

(a) Figure 2 in [2]: Potential optimization time on QPUs compared to CPU/GPU approaches.

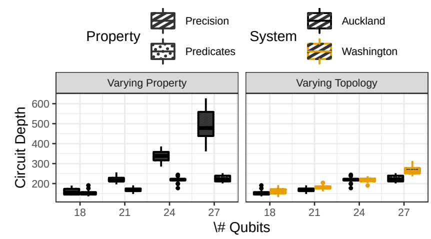

(b) Figure 3 in [2]: Join ordering problem circuit depths for different scenarios.

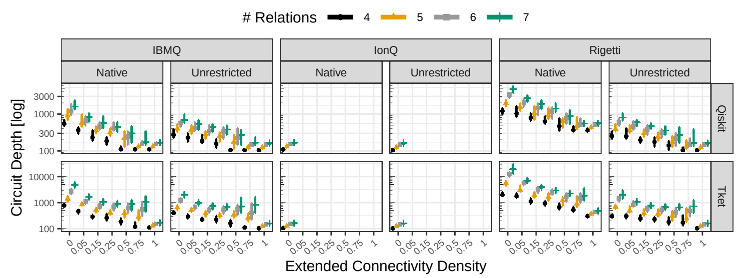

(c) Figure 5 in [2]: Potential circuit depths for future QPUs on varying scenarios for join order problems.

Fig. 1. Reproduced Figures 2, 3, and 5 in [2].

5 SUMMARY

The core thesis of the paper is reproduced and during the reproducibility evaluation suggestions for improvement of the artifacts were taken into account.

REFERENCES

[1] Manuel Schönberger and Wolfgang Mauerer. 2023. Reproduction Package for Ready to Leap (by Co-Design)? Join Order Optimisation on Quantum Hardware. https://github.com/lfd/sigmod23-reproduction (last accessed Dec. 2023).

[2] Manuel Schönberger, Stefanie Scherzinger, and Wolfgang Mauerer. 2023. Ready to Leap (by Co-Design)? Join Order Optimisation on Quantum Hardware. *Proc. ACM Manag. Data* 1, 1, Article 92 (May 2023), 27 pages. https://doi.org/10.1145/3588946