# Quantum Optimization of Worldwide LHC Computing Grid data placement QUOG-DP

Mircea-Marian Popa[1,2], Federico Carminati[2], Dr. Sofia Vallecorsa[2,*], Costin Grigoras[2], Dr. Latchezar Betev[2], Si. dr. ing. Mihai Carabas[1], Andrei Tanasescu[1], Conf. dr. ing. Pantelimon G. Popescu[1]

[1]PUB, Splaiul Independenţei 313, Bucharest 060042, Romania; [2]CERN, Espl. des Particules 1, 1211 Meyrin, Switzerland
*Sofia.Vallecorsa@cern.ch.

***ABSTRACT***

The Worldwide LHC Computing Grid is the infrastructure enabling the storage and processing of the large amount of data generated by the LHC experiments, and in particular the ALICE experiment among them. Currently, the data placement at the ALICE Grid computing sites is optimised via heuristic algorithms. Optimisation of the data storage could yield substantial benefits in terms of efficiency and time-to-result. This has however proven to be arduous due to the complexity of the problem. In this work we propose a modelisation of the behaviour of the system via principal component analysis, time series analysis and machine learning, starting from the detailed data collected by the MonALISA monitoring system. Based on our modelisation, we then explore how to optimise data placement in order to maximise throughput via quantum reinforcement learning . We show that it is possible to analyse and model the throughput of the ALICE Grid to a level that has not been possible before by a recurrent neural network. We also present a possible quantum agent-based schema for the optimisation of the ALICE data distribution in order to maximise data processing throughput .

*Keywords: quantum computing; reinforcement learning; computing grid; data placement optimization.*

## 1. INTRODUCTION

In the year 2000, the HEP community decided to adopt the concept of computing Grid [1] to satisfy the computational and storage needs of the experiments at the Large Hadron Collider at CERN (LHC) [2]. It was a successful decision, as the Worldwide LHC Computing Grid (WLCG) [3] is used to store and process all LHC data and it is the computing platform that has allowed the discovery of the Higgs Boson, so far the crowning achievement of LHC physics research. One of the major problems facing the designers of the Grid was the optimal usage of its resources. Attempts at simulating this complex and highly non-linear environment did not yield practically usable results, and the development of a central “optimiser” for payload and space management never achieved the expected performance. For payload submission and management, the adopted solution was based on a local optimisation of the workload, which led to a satisfactory overall utilisation of the resources.

The central concept that drives the Grid is the payload (often referred to as a ‘job’). It runs on computing farms distributed across the globe, analyses, transforms or generates data (stored in files) and writes the results to different types of storage elements (SE), e.g. tape, HDDs, SDDs. On average, there are 120.000 jobs active on the Grid at any given moment in time. Regarding data operation, a job can read or write files. When it comes to reading, a job may have multiple options where to read from, as most files are present in multiple copies on different SEs to prevent data loss in case a SE fails or to prevent SE overload for frequently accessed data. Reads make up ~90% of the data operations. Write operations are directed to a single or, more frequently, multiple SEs, depending on the number of desired copies. The management of the read and write operations greatly influences the performance of the jobs on the Grid, which this project aims to improve.

Quantum computing is already widely explored as a tool to solve optimization problems. The main objective of this project is to demonstrate a working prototype of a quantum application that can be integrated with the current Grid. This would pave the way to a usable quantum solution.

This study presents two noteworthy results. First, the Grid I/O throughput was successfully predicted with a mean relative error of ~4% for about 3 weeks of activity. Second, a simulated quantum agent was created capable of modifying and improving the current heuristic approach to data placement optimization.

**STATE OF THE ART**

The current data placement management is performed by the MonALISA [4] framework (Monitoring Agents using a Large Integrated Services Architecture). Services located at every Grid computing centre send data to a set of central aggregator nodes. The data contain information on SE availability, network topology and throughput, computing performance etc. This information is used to direct individual jobs where to read or write files.

Two data structures contain heuristic information on the state of the Grid:

- A "distance" matrix, linking a computing farm to a SE, whose ~6700 elements are produced by computing network measurements .
- A "demotion" matrix, describing the health, availability and occupancy of each SE, .

For both read and write operations, the steps are similar. For a read, a job sends a query to the central nodes (**CN**s) requesting directions on where to read a certain file. The CNs use a catalogue to determine at which SEs the file copies are located. Next, they assign to each SE a heuristic score which is the sum of values obtained from the two previously mentioned matrices. SEs are then sorted from the lowest to the highest score and returned to the job as a sorted list. Based on this list, the job decides where to read. If the read from the the top-score SE fails, for instance because it has become unavailable by the time the job gets the answer, the job tries the next SE in the list and so on. The process is identical for a WRITE operation with the modification that the returned list of sorted SE contains all available SEs, since a job can write a file to any of them.

## 2. BREAKTHROUGH Character of the Project

With the foreseen increase in the computing requirements of the future LHC experiments, a data placement strategy which increases the efficiency of the WLCG computing infrastructure becomes extremely relevant for the scientific success of the LHC scientific programme. For example, balancing access to local or remote storage is a hard problem with currently no satisfactory solution. The access to remote copies affects about 10% of the ALICE jobs, a total amount of 50PB/year in terms of data volume. With a proper optimisation of the data placement, this percentage may be reduced up to a factor 2 with a clear positive effect on the network occupancy. It is difficult to calculate the gain in financial terms, as the network cost changes from one country to another. However, there is a clear efficiency benefit, as the saved network bandwidth occupancy will decrease, reducing the need for faster and hence more expensive connections. The potential benefit will also extend to better CPU utilization: At the moment, the analysis tasks use 48% of the allocated CPU cores, with an efficiency (CPU time over real time) of only 50%, since half the time is spent in I/O operations. Increasing the CPU utilization will have a twofold effect: decrease the turnaround time for physics analysis, thus allowing for more analysis and better results, and recovery of presently unused CPU power, making it available for other tasks (raw data reconstruction or MonteCarlo simulations).

We studied the possibility to optimize the two matrices used to control the job data using recent advancements in Reinforcement Learning and Quantum Computing. The resulting system should be able to adapt its response to the status of the Grid system in real time.

Deep learning has seen an increase in use in a variety of fields ranging from self-driving cars, up to coherent natural language models, thanks to the availability of large amounts of training data. This situation is analogous to the problem we are working on - the data collected over the past 10 years of ALICE Grid operation can be used to train the deep learning algorithms. A reinforcement learning approach consists in a real time optimisation agent capable of self-adaptation, following the new "industry 4.0" automation approaches. Paired with the training speedup potentially provided by Quantum Computing, this project aims at demonstrating a promising perspective on how the future of the data placement functionality might look like.

## 3. Project RESULTS

This project achieves two important results:

- a **high accuracy prediction** of the data placement performance;
- a prototype optimisation algorithm based on **quantum reinforcement learning (QRL)**.

A **high level of accuracy** in the prediction of the data placement performance is achieved through the use of time series analysis and recursive neural networks (**RNN**) . I/O throughput, measured by MonALISA, has been chosen as the main performance metric. It is defined by how much data is read from a SE and it is collected by services running locally on the SE.

Unfortunately, there are a lot of fluctuations in the measured throughput as shown in Fig. 1. In order to

distinguish between fluctuations due to the job submission pattern and overall system performance, we have applied additive time series decomposition to the data. We have chosen a frequency of the seasonal component that provides a noise component with a gaussian distribution centered around 0. Unsurprisingly, its value is ~5.91 days (a little more than a working week).
Seasonal and Noise components of the raw throughput are discarded as we aim at optimising first the throughput trend.

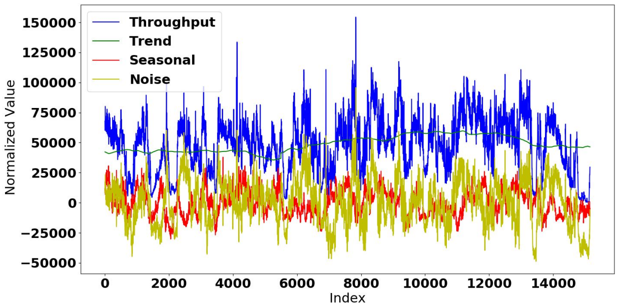


**Fig. 1** Measured throughput and the noise, seasonal and trend components.

In a normal optimisation problem, a set of independent parameters are modified according to the value of a dependent objective cost function in order to find the set of values that corresponds to an extreme of the function. In our problem, the MonALISA heuristic matrices represent the independent variables: by choosing the matrices configurations that maximise the I/O throughput trend we can make sure that the MonALISA system provides the optimal answer to each READ/WRITE query.
In our case, it was unfortunately impossible to modify these matrixes and measure the corresponding variation of the ALICE Grid throughput. The ALICE production system is the beating heart of the experiment physics programme, and any disruption of it would have serious consequences.
We decided therefore to simulate the behaviour of the ALICE Grid throughput as a function of the values of the heuristic matrixes. This was no small task in itself, since the simulation of the Grid behaviour has been tried many times in the past with not reliable results. Thanks to the recent development of Machine / Deep Learning and the specific expertise in this area in our group, we tried an approach based on AI to reproduce the Grid behaviour.
In order to reduce the dimensionality of the problem (the heuristic matrices contain, together thousands of elements) we used and compared two alternative approaches: a principal components analysis (PCA) applied to the input matrixes followed by a recurrent network (RNN), we call this approach Principle Component Network (PCN), and an encoder-decoder network (EDN) architecture coupled to a RNN to output the corresponding throughput trend.
For both approaches multiple hyperparameters have been optimised, ranging from the time sequence window to the regularization approach (a 10% dropout strategy) for the EDN architecture. Finally, the optimal dimension of the downsized matrix space, expressed by the number of principal components for the PCA is 11 (yielding up to 95% of the total variance and the latent variables in the case of the EDN is 10. Fig. 2 and Fig. 3 show the results on the validation set for the PCN and EDN respectively. In both cases the total accuracy is around 4%. The accuracy is calculated using an independent validation extracted from the MonALISA monitoring systems. In Fig. 4 and Fig. 5 the PCN and EDN architectures are summarised.

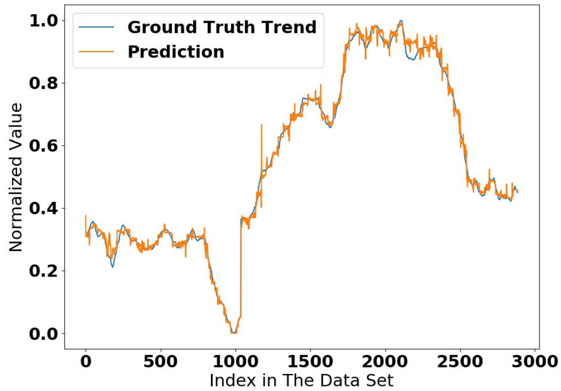


**Fig. 2** Overlay between PCA+RNN prediction and ground truth on validation data set

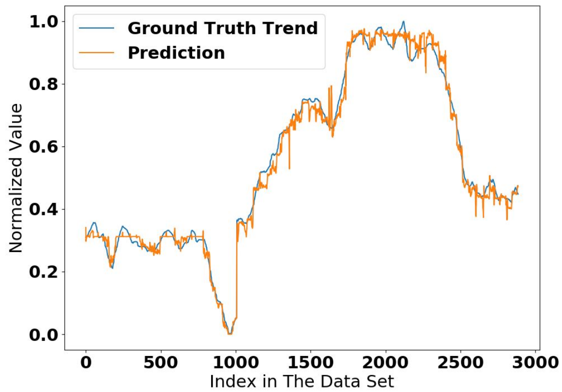


**Fig. 3** Overlay between EDN+RNN prediction and ground truth on validation data set

| RNN Layers | Output Dimension |
|---|---|
| Time Distributed Dense | (40,12) |

| | |
|---|---|
| BachNormalization | (40,12) |
| Bidirectional LSTM | (40,24) |
| BatchNormalization | (40,24) |
| TimeDistributed Dense | (40,10) |
| BatchNormalization | (40,10) |
| TimeDistributed Dense | (40,5) |
| BatchNormalization | (40,5) |
| Bidirectional LSTM | (2) |
| BatchNormalization | (2) |
| Dense | (1) |

**Fig. 4** Summary of the best PCA+RNN architecture. The output dimensions of the RNN layers are listed.

| AutoEncoder Layers | Output Dimension |
|---|---|
| Input Layer | (40,6768) |
| TimeDistributed Dense | (40,10) |
| BatchNormalization | (40,10) |
| Time Distributed Dense | (40,6768) |

**Fig. 5** Summary of the AutoEncoder layers for the best EDN+RNN architecture. .

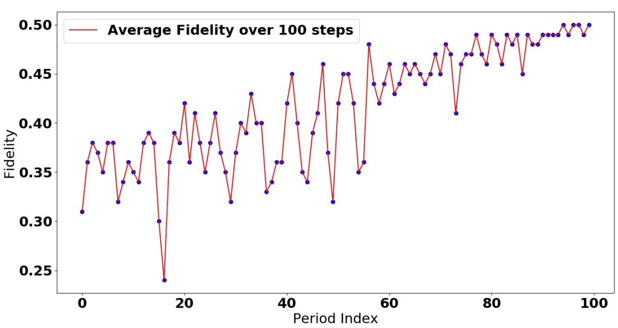


**Fig. 6** Average fidelity of one agent on the Grid World problem

The possibility to use any of the two RNN-based architectures to accurately predict the throughput trend is the starting point to design a quantum reinforcement learning (QRL) agent capable of optimising data placement by maximising the I/O throughput. In Reinforcement Learning, the optimisation process happens through the interaction of an agent with the external environment (in this case represented by the RNN networks) which provides a realistic feedback for every action the agent takes (described by a change in the matrices configuration, the 11 PCN components or 10 EDN latent space size). In classical terms, this function is an abstraction which can be implemented by either a matrix (number of states times number of actions)s x # actions) or by a deep quantum network.

We have implemented a QRL framework using Quantum Boltzmann Machine (QBM) as a free energy-based Reinforcement Learning agent [5]. This implementation uses the D-WAVE Quantum Annealing framework[6]..

In [5], a clamped QBM is used to get Q values. First off, the state space is discretized (the average read size and the matrix components) in 32 levels per component. An action is defined as a +1/-1/0 change to each component level. Even in the case of the reduced space we obtained via PCA, still this generates a huge space of $3^{11}$ total actions. Therefore, for our initial prototype we have restricted the optmisation to the first 3 PCA components (collecting 94% of the total variance) obtaining only 27 possible actions per step.

In order to verify the correctness of the QBM training scheme and the state-action representation in quantum circuits we have run a set of simplified tests using the D-WAVE Ocean Software Suite [6]. Fig. 6 shows the QBM agent convergence for a standard RL benchmark: ta simplified grid-world problem, where an agent learns the shortest past to reach a specified point on a 2D grid. Performance is measured in terms of fidelity, defined as the distance between the agent answer and the optimal action policy ( $\text{fidelity}_{\text{period\_i}} = \sum_{t = period\_i * 100}^{t = (period\_i+1) * 100} 1\{action_t \in optimal\ policy\}$ )

## 4. FUTURE PROJECT VISION

There are two immediate tasks that would largely benefit from further development: the Deep Learning based throughput prediction and the optimisation of the QRL agent performance.

A 4% accuracy on the throughput trend prediction is already an impressive result: further optimisation of the EDN approach, would prove the benefits of an end-to-end deep learning solution independent of preliminary feature extraction processes (such as the

PCA analysis) and full automatization of highly non-linear simulations.

On the quantum side, the next natural step would be careful tuning of the QRL training hyperparameters, comparison to a classical deep learning agent.
and finally, benchmark on real quantum hardware.

Furthermore, in order to extend the approach to a larger sized problem, such as the full ALICE Grid, alternative RL actions-state representations could be considered, such as policy-based algorithms, in order to overcome the problems related to the dimension of a quantized action space. In this kind of approach, the agent holds a stochastic variable that spans the action space and it queries it for an action choice. Since, by their nature, qubits act as stochastic variables, they seem particularly suitable for a policy-based agent.

### 4.1. Technology Scaling

There are a few steps that can be performed in order to make the prototype production-ready. The first step would be to gather more data in order to improve the environment network prediction and develop an end-to-end deep learning approach to downsizing the heuristic matrices (eliminate the principal components step entirely) as explained above. Also a policy-based agent would allow for scaling up the solution to larger realistic systems. Finally, creating a hybrid quantum classical integrated computing system would pave the way to the solution of numerous highly-non linear, large scale optimisation problems.

The study of the **whole WLCG** use case would be a first natural extension within the realm of High Energy Physics. However this work also represents a proof of concept (PoC) for large scale solutions to production planning, logistics management and workload scheduling in real-world industrial scenarios: the optimisation of planning of production orders with finite capacity, job interdependencies (multiple orders sharing production phases) and technical constraints.
**This PoC could be extended and optimised to include** multi-step production orders**,** multiple departments **and** tens of thousands of concurrent orders, **with daily granularity, and over multiple month time-spans. For example, we are currently exploring the possibility to build** a realistic use case using a steel factory production data**.**

### Project Synergies and Outreach

As mentioned above any problem requiring an adaptive algorithm, with real time optimisation capabilities, would benefit from the experience and insight gained through this project on Reinforcement Learning.

Potential synergies are numerous. For example, collaboration with AI and engi**neering** consulting companies**, providing solutions for large scale** industrial partners**, would bring us the necessary** field expertise **needed to adjust our PoC to the specific industrial use case and it would simplify access to** realistic datasets**.**
**Collaboration with** cloud service providers **would get access to dedicated computing resources and design the initial prototypes of QCaaS (Quantum Computing as a Service), similar to D-Wave Leap and Leap2 systems.**

### 4.2. Technology application and demonstration cases

First and foremost, demonstrating the viability of a quantum agent as a solution to such a complex problem as managing data placement in a worldwide computing Grid would be another proof of the potential in using quantum computing. As already mentioned, the same steps could be taken to solve similar problems (e.g. traffic management, production lines, etc).

### 4.3. Technology commercialization

Both the Deep Learning and Quantum Reinforcement Learning aspects of our projects would find easy applications in the commercial world. As an example, social networks come to mind as having to manage a large amount of data as well as content distribution networks (e.g. Youtube).

In addition, any optimisation of complex distribution computational tasks involving the analysis of Big Data and Deep Reinforcement Learning (DRL). DRL represents one of the most promising technologies in the field of decision-making and self-learning and it is at the core of many autonomous control systems. RL solves a classical optimization problem by introducing a feedback to the system, which slowly changes the system itself and converges to an optimal solution in the configuration space. This is a flexible approach allowing agents to be trained in quasi-real time fashion and adapt to changing environment conditions. Moreover, decision rules are not implemented by human experts, they are automatically extracted **from the data by Deep Neural Networks (Steel production lines and other industrial processes** , **"**industry 4.0"**).**

### 4.4. Envisioned risks

This project represents one of the first applications of a QRL approach to a realistic use case. Its novelty represents its main risk and, at the same time, its greatest asset. On the one hand, the complexity of the problem to solve, with its large dimensionality might require quantum resources that current hardware cannot yield. This hindrance of course would be solved, in time, by the development of more and more powerful quantum computers. At the same time, the lessons learned through the development project would be invaluable towards the implementation of hybrid classical-quantum systems of production-level quality.

**ACKNOWLEDGEMENT**

This project has received funding from the ATTRACT project funded by the EC under Grant Agreement 777222.

## 7. References

[1] Foster I, Kesselman C (1998). The Grid: Blueprint for a New Computing Infrastructure (The Elsevier Series in Grid Computing). Morgan Kaufmann; 1st edition.

[2] https://home.cern/topics/large-hadron-collider. Last read on August 28, 2020.

[3] http://wlcg.web.cern.ch. Accessed August 28, 2020.

[4] I.C.Legrand, H.B.Newman, California Institute of Technology, Pasadena, CA 91125, USA
R.Voicu, European Center for Nuclear Research . CERN, Geneva, Switzerland
C.Cirstoiu, C.Grigoras, M.Toarta, C. Dobre, Polytechnic University Bucharest, Romania
CHEP 2004, Interlaken, Switzerland, September 2004

[5] Levit, A., Crawford, D., Ghadermarzy, N., Oberoi, J. S., Zahedinejad, E., & Ronagh, P. (2017). Free energy-based reinforcement learning using a quantum processor.

[6] https://www.dwavesys.com/software. Accessed August 28, 2020