Offline data processing in the First JUNO Data Challenge

Software optimization and performance

Tao Lin, Weiqing Yin, Guofu Cao, Ziyan Deng, Xuefeng Ding, Wenxing Fang, Maxim Gonchar, Kaixuan Huang, Xingtao Huang, Yixuan Jiang, Rebin Karaparambil, Teng Li, Weidong Li, Wuming Luo, Qiumei Ma, Liangjian Wen, Chengxin Wu, Diru Wu, Jiayang Xu, Zhengyun You, Haosen Zhang, Yongpeng Zhang, Jiaheng Zou (on the behalf of JUNO Collaboration)

Institute of High Energy Physics (IHEP), CAS

E-mail: lintao@ihep.ac.cn

1 Introduction

1.1 The JUNO experiment

The Jiangmen Underground Neutrino Observatory (JUNO) is currently under construction in southern China, with the primary goals of the determining the neutrino mass ordering and the precise measurement of oscillation parameters.

1.2 The offline data processing

The raw data is transferred from the JUNO onsite to the IHEP data center via a dedicated network. The RAW data is preprocessed and converted to the RTRAW (ROOT-based). Then both data are replicated to the other data centers, including CC-IN2P3, INFN-CNAF and JINR. There are several critical components in the data processing, such as data quality monitoring (DQM), keep-up reconstruction (KUP) and physics production (PP).

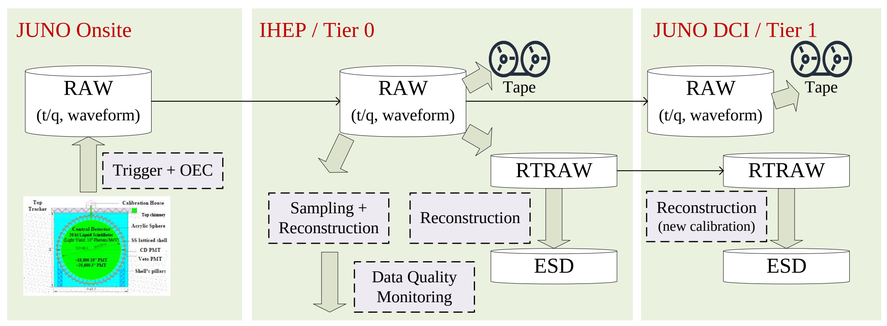

Figure 1: Baseline scheme of offline data processing

1.3 JUNO Data Challenge (DC)

A series of JUNO DC are proposed to evaluate and validate the complete data processing chain in advance. The JUNO DC not only tests the software for vertex/energy/muon reconstructions, but also acts as a system test for database, KAFKA-based data pipeline, DQM, KUP and DCI etc. The estimation of computing and storage is validated in JUNO DC.

2 The first JUNO Data Challenge (DC-1)

2.1 Workflow

There are two major steps in the JUNO DC-1. The first step involves using simulation software to produce RTRAW files. The second step involves reconstructing the RTRAW files in the different systems and producing ESD files. The conditions database is also used for testing.

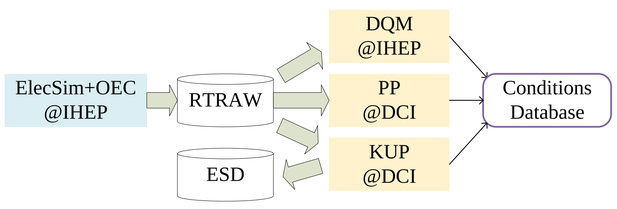

Figure 2: Workflow in DC-1

2.2 Simulation software for producing RTRAW

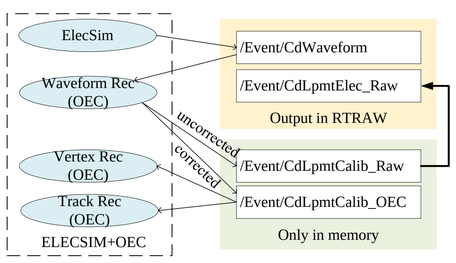

Figure 3: Workflow in ElecSim and OEC

The simulation software consists detector simulation based on Geant4 and electronics simulation with online event classification (OEC). As shown in the figure, the results from electronics simulation are used by the OEC, which reconstructs the events first and decides whether the waveform or time and charge information (t/q) should be stored in the RTRAW file. For the t/q stream, only the uncorrected data is stored into file, allowing for offline correction later.

2.3 Development of multi-threaded reconstruction

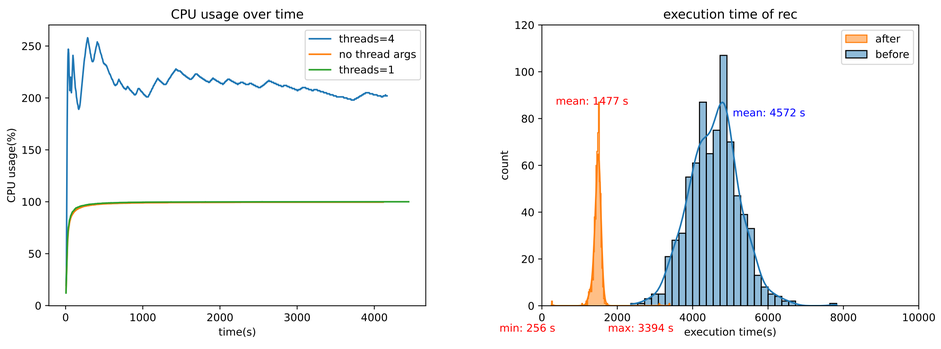

Figure 4: (a) Low CPU usage due to an issue in multi-threaded track reconstruction. (b) Execution time before and after issue fixed.

JUNO Hackathon is organized to migrate the software from a serial version to a multi-threaded version. Issues encountered during testing have been addressed and resolved. Profiling tools are used to identify bottlenecks. The figures show one of the issues related to low CPU usage and the result after optimization.

2 The first JUNO Data Challenge (DC-1)

2.4 Computing resources

The JUNO DQM cluster is used for the testing. This cluster comprises 36 computing nodes with a total of 2304 cores. A total of 576 job slots are allocated, each with 4 cores and 15 GB memory per job slot.

3 Software performance

3.1 RTRAW production

One week of RTRAW data have been generated in DC-1. Each RTRAW file contains about 851 events within 6-second interval. One of the challenges in simulation is memory consumption, as both electronics simulation and OEC run together while multiple events are cached in the memory. These tasks are executed on the large memory computing nodes.

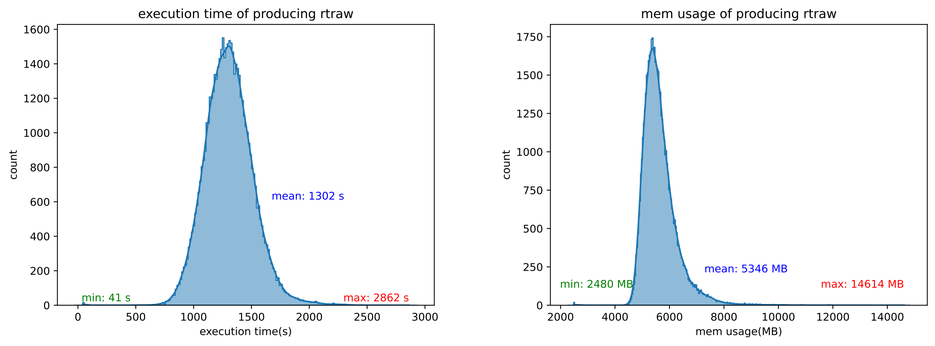

Figure 5: Performance of generating 6s RTRAW data. (a) Execution time of the simulation. (b) The maximum memory usage during simulation.

3.2 Serial reconstruction

The performance of serial reconstruction is measured as a reference.

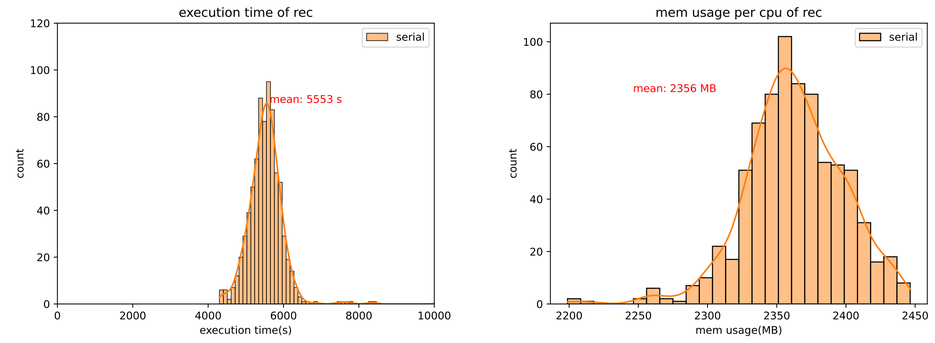

Figure 6: Performance of serial reconstruction. (a) Execution time. (b) Memory usage.

3.3 Multi-threaded reconstruction

Because the processing times for different types of events vary, the output could be blocked if one event is not yet finished. The software offers two output modes: in the 'global output' mode, events are cached in memory with the correct time order, and then data is saved into the file. In the 'output in thread' mode, events from different threads are saved into separate files first and then merged and sorted at the end of the jobs. The following figures show the two cases. The time consumption in "output in thread" mode is less than the other mode. Therefore, the "output in thread" mode is chosen in the official data production.

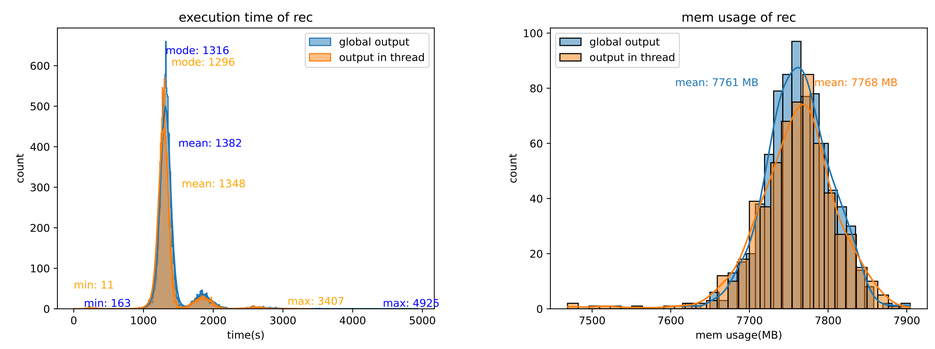

Figure 7: Performance of multi-threaded reconstruction with 4 CPU cores. (a) Execution time. (b) Memory usage. Note about the second peak in Figure 7 (a): some jobs are scheduled to other computing nodes with different CPU models.

4 Conclusions and plans

The JUNO DC-1 is the first time to test the data processing chain starting from RTRAW, mimicking the real data processing. Multi-threaded algorithms have been developed and tested. The database is used as well in the local cluster and DCI. All major goals have been successfully achieved, and ongoing checks on the produced data are in progress.

Some parts are still missing and need further improvement in DC-1. For example, only the reconstruction algorithms for the central detector have been tested. This work will be addressed in the upcoming rounds of the JUNO DC.

Acknowledgements

This work is supported by the Strategic Priority Research Program of the Chinese Academy of Sciences (Grant No. XDA10010900), National Natural Science Foundation of China (12375195), the Youth Innovation Promotion Association, CAS.