

Dirac-based solutions for JUNO production system

Xiaomei Zhang^{1,*}

on behalf of the JUNO Collaboration

¹Institute of High Energy Physics, Beijing, China

Abstract. The JUNO (Jiangmen Underground Neutrino Observatory) Monte Carlo production tasks are composed of complicated workflow and dataflow linked by data. The paper will present the design of the JUNO production system based on the DIRAC transformation framework to meet the requirements of the JUNO Monte Carlo production activities among JUNO data centres according to the JUNO computing model. The approach allows JUNO data-driven workflow and dataflow to be chained automatically with availability of data and also provides a convenient interface for production groups to create and monitor production tasks. The functions and performance tests for evaluating the prototype system would be also presented in the paper.

1 Introduction

JUNO [1] is a multipurpose neutrino experiment. JUNO plans to take about 2 PB raw data each year, which will start from 2022 and take data for more than 10 years. The JUNO Monte Carlo (MC) production activities will be arranged and operated on the JUNO distributed computing system which can integrate heterogeneous resources from the JUNO data centres globally. The experiment data including Monte Carlo data and raw data will be stored in IHEP, while multiple copies will be replicated in the European data centres. Therefore, a production system is needed to handle MC production workflow and dataflow in an automatic way.

1.1 JUNO Monte Carlo simulation

In JUNO, Monte Carlo simulation algorithms and software are built and run on the framework called SNIPER [2]. The JUNO Monte Carlo simulation [3] is used for detector design and optimization, algorithms validation and physics studies. As shown in Figure 1, each JUNO MC simulation is composed of five parts: Physics Generator (*PhyGen*), Detector Simulation (*DetSim*), Electronics Simulation (*EleSim*), PMT Reconstruction (*PmtRec/Cal*), and Event Reconstruction (*EvtRec*). Each JUNO production task includes data processing and data replication activities. The data processing includes four steps: *DetSim*, *EleSim*, *PmtRec*, *EvtRec*, where the *PhyGen* step is combined into the *DetSim* step. Every step generates different type of event data. Except *DetSim*, the other steps need input

* Corresponding author: zhangxm@ihep.ac.cn

event data. These four steps are interconnected with data, which form the JUNO simulation workflow. These data produced by these steps are replicated between data centres and sites, which form the JUNO dataflow. The JUNO production tasks include large samples of physical processes, such as Inverse Beta Decay, backgrounds, positron and electron with different momenta, muons, etc. It is hard for data production groups to handle these large and complex production tasks manually in distributed environment.

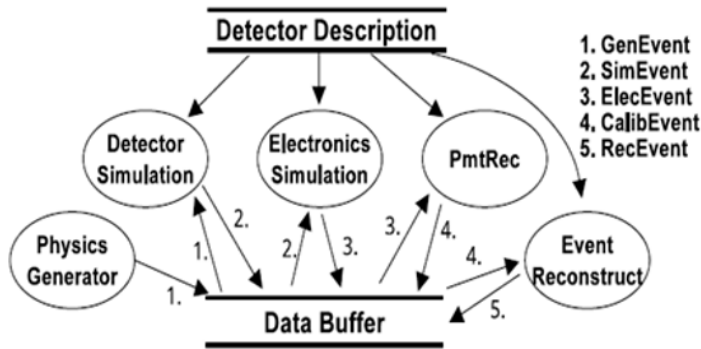


Fig. 1. Data processing steps in JUNO Monte Carlo production tasks

1.2 JUNO distributed computing and computing model

The JUNO distributed computing system has been built on DIRAC [4], which provides a complete grid solution and framework for high energy physics experiments. The resource types integrated in the JUNO distributed computing system include cluster, grid and cloud.




Fig. 2. JUNO distributed computing “Tier” architecture

The JUNO distributed computing plans to use “Tier” architecture composed of three layers, as shown in Figure 2. The IHEP data centre as Tier0 will hold central Storage Element (SE), receive and store raw data from the onsite, and also store one copy of all other data types including simulation data, reconstruction data, calibration data, etc. Tier0 will be responsible for first-time full reconstruction and calibration, and also will perform simulation and user analysis. The data centres in Europe (IN2P3, JINR, CNAF) as Tier1 will hold another copy of the whole data and perform re-reconstruction, simulation, user analysis. Small and opportunistic sites without SE as Tier2 will perform some part of simulation. Small sites with cache linked to the SEs in data centres will also support user

analysis. The JUNO production tasks will be assigned by the production group through the distributed computing system to all the JUNO centres and sites. To reduce the burden of central SE, two SEs will be used to receive the output data produced from local SEs. That means the data produced in local SE will be replicated to IHEP or one of data centres in Europe, and synchronized between IHEP and this data centre, and replicated to other data centres if needed.

2 JUNO production system

The purpose of JUNO production system is to provide a convenient interface for the JUNO product groups to submit production tasks and manage the JUNO MC simulation workflow and dataflow in an easy way.

2.1 Design

2.1.1 Architecture and Functions

The JUNO production system is designed based on the DIRAC Transformation System (TS) [5]. The TS provides a framework to handle “repetitive” work and chain production workflow and dataflow in a data-driven way. As shown in Figure 3, the JUNO production system mainly comprises of four parts: production manager and transformation system, DIRAC Workload Management System (WMS) and Data Management System (DMS). The JUNO production manager allows the production group to define production requests with a steering file. Interfaced with the TS, the JUNO production manager transforms these requests into transformation instances according to the JUNO computing model. These transformation instances are interconnected with the availability of data by checking DIRAC File Catalogue (DFC) with metadata. Each type of transformation instance creates a sequence of jobs or a list of data replications for each step of production tasks. These jobs and replications created by the transformations are submitted to the WMS and DMS separately, where they will be scheduled to the related services and resources for real operations. The JUNO production system also provides an interface for the production group to monitor and control of those workflow and dataflow.

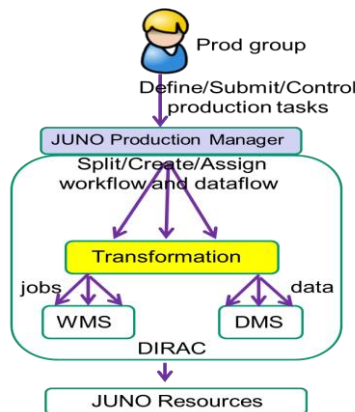


Fig. 3. Architecture of JUNO production system

2.1.2 Design of transformation modules for JUNO workflow and dataflow

The transformation is the heart of production system. Design of transformation modules according to JUNO workflow and dataflow is the most important part. Each production task includes five steps: *DetSim*, *EleSim*, *PmtRec/Cal*, *EvtRec/Rec* and replications from closest SEs to final SEs. As shown in Figure 4, accordingly each step is taken care by a transformation module. These transformation modules are chained by the metadata query. The first transformation module *DetSim* without need of inputs is launched directly by the production system. Other transformation modules are started when their inputs are found to be available by looking into DFC with metadata query. When the query returns file lists, those modules are triggered to generate jobs or data transfer tasks with these files. All the output data from jobs in last step is downloaded to closest SEs and registered in DFC with predefined metadata when arriving in SE, which can be immediately known by the next step and in turn trigger the following step.




Fig. 4. JUNO workflow and dataflow in transformation

2.2 Implementation

The whole system is mainly implemented in three parts: production manager, workflow, dataflow. The production manager is to accept JUNO production requests, transform these requests into transformation tasks and launch the production chain. This part is JUNO-specific, and closely integrate with JUNO data processing activities. The workflow and dataflow implementation are general, and can be used in other experiments as well. More details on workflow and dataflow implementations are explained in the followings.

2.2.1 Workflow

The workflow setup aims to create jobs by the workflow transformations and assign them to the WMS. Three systems are involved: DFC, TS and WMS. In TS, three agents are used to create and submit data-driven jobs: *InputData* agent, *Workflow Task* agent, *Transformation* agent. The *InputData* agent queries the DFC with metadata to see if the files as inputs to jobs are available in SE to create jobs. As shown in Figure 5, when the files are ready in SE, the *Transformation* agent creates jobs and the *Workflow Task* agent is responsible to submit the jobs to the DIRAC WMS and also keep track of the status of jobs to report to the monitoring part. When the jobs arrived in the *Task Queue*, the DIRAC WMS will schedule jobs to sites.

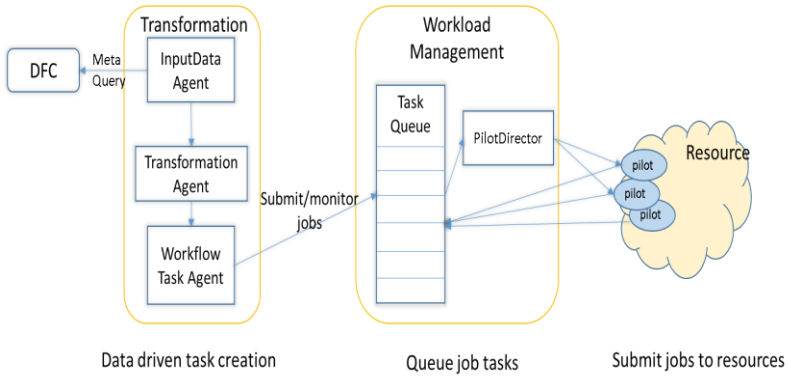


Fig. 5. JUNO workflow implementation

2.2.2 Dataflow

The dataflow setup creates and assigns data replication tasks by transformations to the FTS (File Transfer System) [6] which can take care of file-by-file transfers between SEs. The dataflow part can also be acted as an independent data replication system which accepts only data replication requests. As shown in Figure 6, five systems are involved to complete dataflow work: DFC, TS, RMS (Request Management System), DIRAC FTS service and FTS. Just as what does in the workflow part, first the *InputData* Agent queries DFC to check the availability of data. When data is ready, the *Request Task* Agent creates data replication tasks and puts them into queues of the RMS. The RMS submits tasks to FTS service which is interfaced with external FTS to do real replications between SEs.

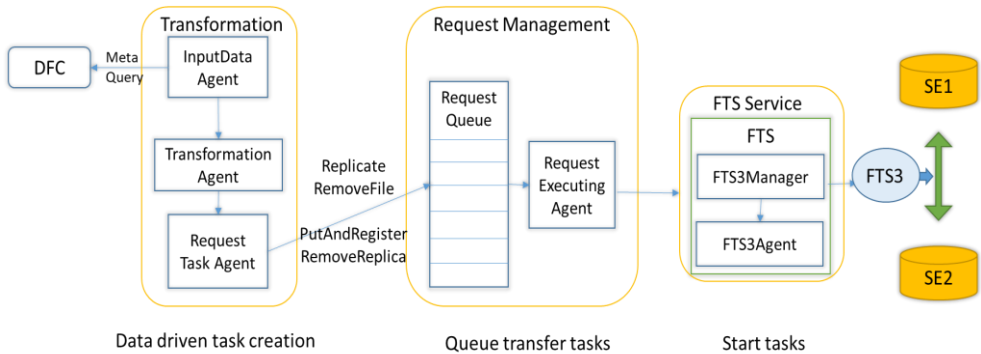


Fig. 6. JUNO dataflow implementation

2.3 Tests

The JUNO software is deployed through CVMFS. The JUNO software version used in the tests is J17v1r1. The production tasks for testing are to create samples of positron at different momenta which includes 0.0 MeV, 1.398 MeV, 4.460 MeV and 6.469 MeV. For each momentum, eight transformation instances are created, including four workflow transformation types (*detsim*, *elecsim*, *cal*, *rec*) and four replication transformation types (*detsim-replication*, *elecsim-replication*, *cal-replication*, *rec-replication*). Each type of

workflow transformation instance generates 100 jobs and each job processes 1000 events. Figure 7 shows two plots of these tests. All the jobs and replications were successfully completed.

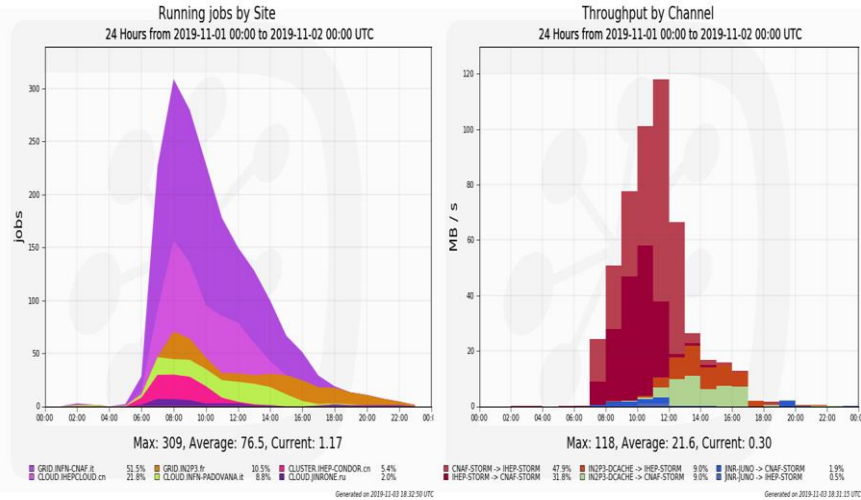


Fig. 7. Tests done for JUNO production system. The left plot shows the jobs running in sites and the right plot shows the replication speed between sites

3 Summary and Plans

The JUNO production system have been designed and tested for the JUNO Monte Carlo production, and also can be extended to other activities such as data reconstruction if needed. The tests with real JUNO production tasks have proved that the system is working well as planned. The system was also successfully applied to replicate raw and reconstruction data from IHEP to other sites. In the near future, heavier tests with more resource involved will be deployed, to check the whole system for possible bottleneck and tune its performance.

The authors would like to thank the JUNO offline software team for their help with the JUNO software, the DIRAC team for their help with the DIRAC transformation system, and colleagues at the IHEP computing centre for their support. This work was funded by the National Natural Science Foundation of China (NSFC) under grant no. 11775246.

References

1. F.P. An, et al., Neutrino Physics with JUNO, *J. Phys. G* 43 030401 (2016)
2. J.H. Zou, SNI_{PER}: an offline software framework for non-collider physics experiments, *J. Phys.: Conf. Ser.* 664 072053 (2015)
3. T.Lin, J.H.Zou, W.D.Li, Z.Y.Deng, X.Fang, Application of SNI_{PER} to the JUNO Simulation, *J. Phys.: Conf. Ser.* 898 042029 (2017)
4. A.Tsaregorodtsev, M.Bargiotti, N.Brook, A.C.Ramo, G.Castellani, P.Charpentier, C. Cioffi, J.Closier, G.DiazR, G.Kuznetsov, Y.Y.Li, R.Nandakumar, S.Paterson, R.Santinelli, A.C.Smith, M.S.Migueluez and G.JimenezS, *J. Phys.: Conf. Ser.* 119 062048 (2008)

5. F. Stagni, P. Charpentier, *J. Phys.: Conf. Ser.* 368 012010 (2012)
6. A.A.Ayllon, M.Salichos, M.K.Simon, O.Keeble, *J. Phys.: Conf. Ser.* 513 032081(2014)