

# China-EU scientific cooperation on JUNO distributed computing

Giuseppe Andronico<sup>1,2,\*</sup>

<sup>1</sup>Istituto Nazionale di Fisica Nucleare (INFN), Catania division

<sup>2</sup>on behalf of JUNO collaboration

**Abstract.** The Jiangmen Underground Neutrino Observatory (JUNO) is an underground 20 kton liquid scintillator detector being built in the south of China. Targeting an unprecedented relative energy resolution of 3% at 1 MeV, JUNO will be able to study neutrino oscillation phenomena and determine neutrino mass ordering with a statistical significance of 3-4 sigma within six years running time.

These physics challenges are addressed by a large Collaboration localized in three continents. In this context, key to the success of JUNO will be the realization of a distributed computing infrastructure to fulfill foreseen computing needs.

Computing infrastructure development is performed jointly by the Institute for High Energy Physics (IHEP) (part of Chinese Academy of Sciences (CAS)), and a number of Italian, French and Russian data centers, already part of WLCG (Worldwide LHC Computing Grid).

Upon its establishment, JUNO is expected to deliver not less than 2 PB of data per year, to be stored in the data centers throughout China and Europe. Data analysis activities will be also carried out in cooperation.

This contribution is meant to report on China-EU cooperation to design and build together the JUNO computing infrastructure and to describe its main characteristics and requirements.

## 1 Introduction

The JUNO collaboration is working to fulfill an ambitious physical goal realizing an underground neutrino observatory able to reach very interesting results, as reported in Section 2. These goals are related to a very sophisticated detector producing huge amount of data whose analysis will require data storage and computing power, as reported in Section 3. The solution developed to meet these requirements is described in Section 4.

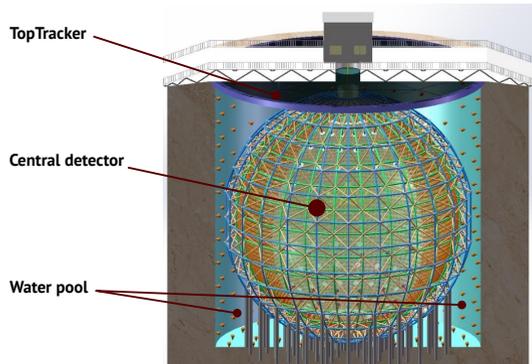
## 2 Jiangmen Underground Neutrino Observatory

The Jiangmen Underground Neutrino Observatory [1] is a 20 kton liquid scintillator (LS) detector currently under construction in the south of China close to Jiangmen exploring neutrino properties, by means of electron anti-neutrinos emitted from two nuclear power complexes, Taishan and Yangjian, with a final thermal power of  $35.8 \text{ GW}_{th}$  at a baseline of

---

\*e-mail: [giuseppe.andronico@ct.infn.it](mailto:giuseppe.andronico@ct.infn.it)

about 53km. The LS will be contained in a 35.4 m diameter acrylic sphere and monitored by  $\sim 18000$  20" photomultiplier tubes (PMTs), allowing for an unprecedented energy resolution of  $3\%/\sqrt{E(\text{MeV})}$ . A complementary system of  $\sim 25000$  3" PMTs facilitates to use the concept of double calorimetry [2]. The Central Detector will be submerged in a cylindrical water pool filled with ultra-pure water, shielded from natural radioactivity from the surrounding rock and air and equipped with  $\sim 2000$  20" PMTs to detect Cherenkov light from cosmic muons, acting as a veto detector. On top of water pool, there will be the Top Tracker, a muon detector to accurately measure the muon tracks. An artist view of the JUNO detector is shown in Figure 1.



**Figure 1.** Schematic view of the JUNO detector main components.

The main goal of JUNO is to determine the neutrino mass hierarchy by measuring the oscillations of reactor antineutrinos emitted by the two power plants. Furthermore, measurements of the oscillation parameters  $\sin^2\theta_{12}$ ,  $\Delta m_{12}^2$ , and  $|\Delta m_{ee}^2|$  can be achieved with subpercentage precision. Additional physics goals exist [1], among them we mention:

**Supernova neutrino** and other astrophysical sources looking at inverse beta decay; also  $\nu - p$  and  $\nu - e$  elastic scattering will provide useful signals;

**Atmospheric neutrino** detecting  $\nu_\mu, \bar{\nu}_\mu$  naturally produced in atmosphere [3];

**Geoneutrinos** detecting the  $\bar{\nu}_e$  emitted from  $^{238}\text{U}$  and  $^{232}\text{Th}$  will help in testing Thorium and Uranium abundance in Earth's nucleus and mantle, estimating Th/U rate [4, 5];

**Solar neutrinos** more precise determination of solar neutrinos, thanks to its high resolution and low energy threshold [6];

**Dark matter** due to JUNO excellent energy resolution we look for evidence of a possible dip in the neutrino energy spectrum due to resonant interactions between Galactic supernova neutrinos and dark matter particles.

The JUNO simulation software is used to study the detector response and optimize the detector design. It is a GEANT4 [7] based Monte Carlo simulation software, which has been carefully tuned for several characteristics as liquid scintillator light yield, charge response and energy non-linearity based on Daya Bay experiment experience.

### 3 JUNO computing model

JUNO is a complex experimental setup. The components main sources of data will be:

**CD-LPMT** the set of 20" large photomultipliers (LPMT) in the number of ~18000, all around the acrylic sphere looking at the light emitted from Liquid Scintillator (LS), and the related electronics to acquire the data;

**CD-SPMT** the set of 3" small photomultipliers (SPMT) in the number of ~25000, also them all around the acrylic sphere looking at the light emitted from LS, and the related electronics;

**WaterPool-PMT** the set of PMT looking at the water pool where the acrylic sphere is immersed in the number of ~2000, and the related electronics.

From the physical point of view, the main events able to excite the LS in the CD or produce Cherenkov effect in the water pool so to emit light are radioactivity, inverse beta decay, dark noise and muons. The components will acquire data with a rate that ranges from few Hz to 1 kHz, with waveform sampled at 1 GHz for 1000 ns and a data size of 4 bytes for each sampled point. This sums up to ~40 GB/s coming out from detector's several components, that is around 1 EB of data for year.

The number of events that JUNO is expected to observe are estimated in [1, 8], based on several models and simulators. Considering physical events, from neutrino and from other sources, around 1 million hits are a preliminary estimate, plus noise from electronics, that will be processed from trigger and on line event classification. Using the JUNO simulator an estimate was made of the event size. So we can appraise that the data rate we expect from the combination of trigger and online event classification should be of the order of 60-70 MB/s. This sums up to about 2 PB of raw data per year.

Another source of relevant data to be stored and processed are calibration data, estimated at ~20 TB per year, not impacting a lot the total amount. With this level of data rate, it is enough to relay on a link with a 1 Gb/s bandwidth.

All these data, produced from JUNO detector, have to be copied from the experimental site to IHEP in Beijing, the main storage center, where the first event reconstruction will take place.

Preliminary software testing shows that reconstructed and simulated data will amount to 600 TB per year. The same testing also shows that typical reconstruction time of raw data is ~1 second per event. Taking into account 2 data productions, related to the 2 software versions per year planned, and to 3 data re-processing for the new calibration constants, using 6000 CPU cores the computing time required to reconstruct 1 year data taking sums up to 365 days. Other activities are planned that require additional CPU power, expressed in cores per year:

**Data quality monitoring** requires 1000 CPU cores;

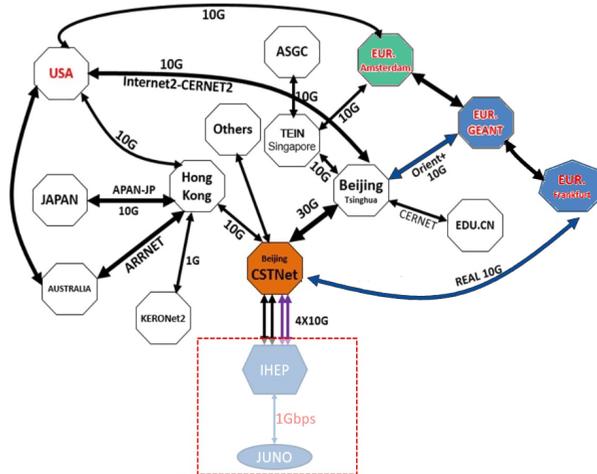
**Physics analysis and Monte Carlo simulation** require 3000 CPU cores.

To ensure data safety, all the data (raw, reconstructed, calibration, analyzed, simulated) have to be replicated. In the spirit of collaboration between China and Europe, four European data centers have given their availability to provide resources and join the storage and computing effort.

In fact, JUNO collaboration can rely on good international network connections. In Figure 2 a graph where the direct link between IHEP and JUNO site is in evidence. Also, it shows connections between IHEP and the rest of the world. In evidence, in blue, links with Europe:

**Orient+** the 10 Gb/s link between China research and education networks, based on undersea cables

**Real 10G** the recent 10 Gb/s link connecting directly Chinese Academy of Sciences, and IHEP, with Europe based on underground cables.



**Figure 2.** IHEP and international network connections.

By means of these connections and exploiting other big networks, as GEANT in Europe, all the JUNO partners can share resources. In this way the JUNO European data centers will be able to maintain at least two further copies of all the data in Europe. At the moment the working hypothesis is to have one full copy of all data (raw, simulation, calibration, etc) at INFN CNAF and another one at JINR, while CC-IN2P3 will dynamically keep a partial copy of the data. In this way the data centers are expected to contribute to the computing effort, proportionally to the installed processing power that will be available locally for JUNO experiment.

## 4 Distributed computing infrastructure

In the previous section we introduced the JUNO computing model, based on cooperation between IHEP and European data centers. The European data centers participating in JUNO efforts are:

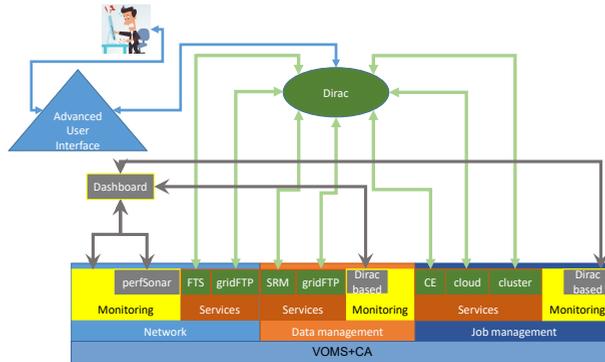
**CC-IN2P3** located in Lyon, is the french Tier1 for LHC experiments;

**INFN CNAF** located in Bologna, is the italian Tier1 for LHC experiments;

**JINR** located in Dubna, is the russian Tier1 for CMS experiment;

**MSU** located in Russia, is a data center under development.

IHEP and these data centers worked jointly, especially from January 2018, to define a common Distributed Computing Infrastructure (DCI). The basic idea is summarized in Figure 3. All the DCI is based on a system for user authentication and authorization (AAI).



**Figure 3.** Schematic view of the main components composing the JUNO Distributed Computing Infrastructure.

For the moment we are relying on VOMS [9] and certification authorities (CA), enabling administrators to manage user rights in accessing data. VOMS groups and roles enable for fine grained administration.

The DCI is made from a number of services that can be grouped, as in Figure 3, in three blocks:

**Network** refers to monitoring service; based on perfSonar [10], needed to ensure networks are properly working and no problems arise from this side in data transfer. Also this block contains the basic services needed for data transfer. In this first version we relied on gridFTP [11] and the file transfer service (FTS) [12], an efficient file transfer service based on gridFTP to move data, tested with success for a long time in WLCG, also providing a monitoring page to check file transfer status.

**Data Management** refers to services to handle data, that is copy, replicate and delete files; based on storage resource manager (SRM), providing a common interface for different storage systems called Storage Elements (SEs), and gridFTP; DIRAC owns a built in data management system able to interact with SRM and FTS, providing also a monitoring system. SEs are active at IHEP, CC-IN2P3, CNAF and JINR; all were tested by means of data transfers between these sites.

**Job management** refers to services to handle jobs, that is to submit jobs, monitor job status, retrieve output and more; what is required in this block is to be able to efficiently manage the workload across distributed resources, as grid computing elements (CE), clusters or clouds. These functions are provided by DIRAC.

A central role in DCI is reserved to DIRAC [13], a software framework for the management of distributed computing. Main duty of DIRAC is to distribute submitted jobs to sites where resources and needed data are available to execute jobs as soon as possible, trying to use at the best available resources. More details about DCI, the prototype and production system can be found in [14].

Another important component is **CernVM File System** [15] providing a scalable, reliable software distribution service, used also for data with a small size, very slowly changing in time and frequently accessed.

Once IHEP and European data centers started to strictly to cooperate, all the DCI components were discussed, fixed, deployed and tested. From the beginning of September 2018 a

DCI prototype is up and running. In 2019 last 6 months some general tests were performed. First of all a functional data transfer test, where about 100 TB of data were copied from IHEP to JINR and then to CNAF and CC-IN2P3 just to test network connections. Some problems with configurations were evidenced and resolved and more recent tests worked quite well. Just after we started to test job submission. Also in this case the DCI was able to support a total of about 20 thousand jobs assigned and run in all the involved sites.

## 5 Conclusions

In JUNO we take advantage from the experience the data centers involved in the computing effort have in supporting WLCG experiments. We exploited such an experience in designing and deploying a pilot distributed computing infrastructure in a common effort between China and Europe. Results obtained are very positive, letting us to be confident that the designed DCI is well suited for JUNO computing needs and will be ready for the next-coming data taking.

## References

- [1] F. An, et al., *Journal of Physics G: Nuclear and Particle Physics* **43**, 030401 (2016)
- [2] M. He, on behalf of the JUNO collaboration, *Radiation Detection Technology and Methods* **1**, 21 (2017)
- [3] W. Guo, *Journal of Physics: Conference Series* **888**, 012205 (2017)
- [4] M. Reguzzoni, et al., *Journal of Geophysical Research: Solid Earth* **124**, 4231 (2019)
- [5] R. Gao, Z. Li, R. Han, A. Wang, Y. Li, Y. Xi, J. Liu, X. Mao, Y. Sun, Y. Xu, *Physics of the Earth and Planetary Interiors* **299**, 106409 (2020), 1903.11871
- [6] G. Salamanna, *Solar neutrinos with the JUNO experiment* (2019), pp. 329–344, arXiv:1809.03821
- [7] S. Agostinelli, et al., *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **506**, 250 (2003)
- [8] G. Andronico, X. Zhang, W. Li, *PoS LeptonPhoton2019*, 122 (2019)
- [9] R. Alfieri, et al., *VOMS, an Authorization System for Virtual Organizations*, in *Grid Computing* (Springer Berlin Heidelberg, Berlin, Heidelberg, 2004), pp. 33–40, ISBN 978-3-540-24689-3
- [10] A. Hanemann, et al., *PerfSONAR: A Service Oriented Architecture for Multi-domain Network Monitoring*, in *Service-Oriented Computing - ICSOC 2005* (Springer Berlin Heidelberg, Berlin, Heidelberg, 2005), pp. 241–254, ISBN 978-3-540-32294-8
- [11] F. Ould-Saada, *Gridftp: Protocol extensions to ftp for the grid* (2013), <https://doi.org/10.5281/zenodo.6828>
- [12] Á. Frohner, et al., *Journal of Physics: Conference Series* **219**, 062012 (2010)
- [13] A. Tsaregorodtsev, et al., *Journal of Physics: Conference Series* **219**, 062029 (2010)
- [14] X. Zhang, X. Zhao, T. Lin, *Dirac-based solutions for jun0 production system* (2019), <https://doi.org/10.5281/zenodo.3599689>
- [15] J. Blomer, C. Aguado-Sánchez, P. Buncic, A. Harutyunyan, *Journal of Physics: Conference Series* **331**, 042003 (2011)