PROCEEDINGS OF SCIENCE

Running Fermi-LAT analysis on Cloud: the experience with DODAS with EGI-ACE Project

Isabella Mereu,[a,*] Marica Antonacci,[b] Diego Ciangottini,[a] Sara Cutini,[a] Cristina Duma,[d] Luciano Gaido,[e] Stefano Nicotri[b] and Daniele Spiga[a]

[a] *INFN sezione di Perugia*
[b] *INFN sezione di Bari*
[c] *INFN sezione di Pisa*
[d] *INFN CNAF*
[e] *INFN Sezione di Torino*
E-mail: mereuisabella@gmail.com, sara.cutini@pg.infn.it, daniele.spiga@pg.infn.it, dciangot@cern.ch

The *Fermi*-LAT long-term Transient (FLT) monitoring aim is the routine search of γ-ray sources on monthly time intervals of *Fermi*-LAT data.
The FLT analysis consists of two steps: first the monthly data sets were analyzed using a wavelet-based source detection algorithm that provided the candidate new transient sources; finally these transient candidates were analyzed using the standard *Fermi*-LAT maximum likelihood analysis method. Only sources with a statistical significance above 4σ in at least one monthly bin were listed in a catalog.
The strategy adopted to implement the maximum likelihood analysis pipeline has been based on cloud solutions adopting the Dynamic On Demand Analysis Service (DODAS) [1] service as technology enabler. DODAS represents a solution to transparently exploit cloud computing with almost zero effort for a user community. This contribute will detail the technical implementation providing the point of view of the user community.

International Symposium on Grids & Clouds 2022 (ISGC 2022)
21 - 25 March, 2022
*Online, Academia Sinica Computing Centre (ASGC), Taipei, Taiwan****

*Speaker

© Copyright owned by the author(s) under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License (CC BY-NC-ND 4.0).

PoS(ISGC2022)023

1. Introduction

The Universe is home to numerous exotic and beautiful phenomena, some of which can generate almost inconceivable amounts of energy. Supermassive black holes, merging neutron stars, pulsars are examples of phenomena that generate γ-ray radiation, the most energetic form of radiation, billions of times more energetic than the type of light visible to our eyes. The *Fermi* Gamma-ray Space Telescope permits the investigation of this high-energy world, helping to understand the origin of the γ-ray emission. *Fermi* was launched June 11, 2008 at 12:05 pm EDT. For its realization, NASA has teamed up with the U.S. Department of Energy and institutions in France, Germany, Japan, Italy and Sweden. This space observatory studies the cosmos in the photon energy range of 8,000 electron volts (8 keV) to greater than 300 billion electron volts (300 GeV).

Fermi carries two instruments: the Large Area Telescope (LAT) and the Gamma-ray Burst Monitor (GBM). The LAT is Fermi's primary instrument, and the GBM is the complementary instrument.

The Large Area Telescope (LAT) is an imaging pair-conversion detector that measures the arrival direction, energy, and time of individual γ-rays with energies from about 20 MeV to over 300 GeV. It has a large field of covering about 20% of the sky and can provide an image of the entire sky approximately every three hours, that is, two orbits around Earth. The LAT has an effective detector area of about a square meter for gamma rays with energies above 1 GeV and its angular resolution for a single gamma ray is finer than 1°, which allows localization of most gamma-ray sources to within 10 arcminutes on the sky. Its timing accuracy is 1 microsecond. The astrophysical photons of primary interest will be a tiny fraction of the particles that will penetrate into the LAT. The LAT on-board analysis reduces raw LAT trigger rate, which can approach 10 kHz, to $\sim$ 400 events per second which are sent to the ground for further analysis. Of these ~400 Hz only ~2-5 Hz are astrophysical photons [2]. After the ground data processing, the reconstructed γ-ray photon events are made available at the *Fermi* Science Support Center (FSSC) [3]. The occupied storage by these photon files in 14 years of operation is about 150 Gb.

The Gamma-ray Burst Monitor (GBM) is an array of sodium iodide and bismuth germanate crystal scintillators that views everything in the sky not occulted by Earth. It's sensitive to X-rays and γ-rays with energies from 8 keV to 40 MeV. Gamma-ray burst localization for this instrument is typical 3° and its timing accuracy is 2 microseconds. The GBM produces an average of 1.4 Gbits/day, with a minimum of 1.2 Gbits/day and a maximum allocated rate of 2.2 Gbits/day.

The *Fermi* huge field of view and the "scanning mode" of operation makes possible two simultaneous approaches to γ-ray astrophysics:

- Constantly deepening exposure of the gamma-ray sky. The instruments accumulate ever more photons with each orbit, giving deeper and clearer views of persistent sources and spatially extended features as time progresses.

- Time-domain γ-ray astronomy. By continually monitoring the cosmos, the instruments are sensitive to changes in the γ-ray sky on time scales ranging from microseconds to years.

PoS(ISGC2022)023

2. Time domain astrophysics: particular Long-Term Transient Sources (FLT)

Time-domain astrophysics studies transient and variable astronomical events, i.e. phenomena changing with time. Transients characterize astronomical objects or phenomena whose duration of presentation may be from milliseconds to days, weeks, or even several years. Common targets included are violent deep-sky events such as γ-ray burst, supernovae, pulsars, novae and active galactic nuclei. Time-domain astrophysics also involve long-term studies of variable stars, including periodic, quasi-periodic, and that of changing behavior or type, and their changes on the timescale of minutes to decades.

Fermi's scanning mode, coupled with the instruments' ability to measure γ-ray arrival times, makes possible variability studies of γ-ray sources over time scales from microseconds to multiple years. The brightest extragalactic γ-ray sources seen by *Fermi* are AGNs, galaxies with cores that are much more luminous than our own over a broad range of photon energies. AGNs are thought to be powered by active accretion of matter onto their central supermassive black holes. Many AGNs produce powerful jets of photons and particles with ultrarelativistic bulk velocities that serve as strong, highly variable sources of collimated γ-rays. The largest single class of sources seen by the LAT is blazars, AGNs for which the relativistic jets happen to be beamed within a few degrees of our direction. If the AGNs flare intensity is fainter than the integrated background over years, we lose the capability to detect these sources in long-time integrated catalogs.

The *Fermi*-LAT Long Term (FLT) analysis studies γ-ray sources over monthly time scales. The monthly time scale allows us to identify transient and variable sources that may have not been reported in *Fermi*-LAT general catalogs, based on the long-term observation of sources. The FLT analysis was performed for photon energies between 0.1 and 300 GeV considering only photons with galactic latitude over than 10° to avoid confusion with low-latitude galactic diffuse emission. We analyzed till now 13 years of data dived in 156 months and also performed a 15-day shift of each month in order to not lose any flare at the edges of each time bin. The total number of months (hereafter called also Time Bins, TBINs) analyzed is therefor 312. The monthly datasets were analyzed using a wavelet-based source detection algorithm that provided the candidate new transient sources. The sky of each of the 312 monthly TBINs is divided into 192 circular regions of interest (ROIs). The detection of the candidate sources (seeds) in each ROI is performed using a wavelet transform analysis [4] using the *PGWave* tool [5]. With this method we collected roughly 1000 seeds per month. The transient candidates not-positionally coincident with known γ-ray sources ($\sim$ 200 per month) were then analyzed using the standard *Fermi*-LAT maximum likelihood analysis method. To perform the analysis we used the *Fermitools* 1.2.1 package (FSSC)[6] and the *Fermipy* 0.18.0 software package [7]. Maximum likelihood analysis is a procedure very time consuming $\sim$ 0.5 h per seed: i.e. using a single workstation it would take 1300 days of computing.

Overall *Fermi*-LAT does not require a extremely large computing and storage resources however a dynamic and elastic system would be perfectly suitable to implement a model to support short processing period with relatively huge amount of resources to satisfy peak or resource requests and possibly to allow portability. The latter is a key to exploit any provider, even the opportunistic one. Those represent the main motivation that lead us to try to port our system in cloud.

PoS(ISGC2022)023

3. EGI Advanced Computing for EOSC Project (EGI-ACE)

EGI-ACE G.A 101017567 is a project coordinated by the EGI Foundation with a mission to empower researchers from all disciplines to collaborate in data and compute intensive research through free-at-point-of-use services. Building on the distributed computing integration in the EOSC-hub project, EGI-ACE delivers the EOSC Compute Platform and contributes to the EOSC Data Commons through a federation of cloud compute and storage facilities, PaaS services and data spaces with analytic tools and federated access services. The consortium of the project builds on the expertise and assets of the EGI federation members, key research communities, data providers and collaborating initiatives. The EGI consortium is composed by 33 partners and 23 third-party beneficiaries from 29 countries worldwide. Consortium at glance: 16 partners to deliver cloud, HTC and HPC resources, 12 user access and platform solution providers, 21 providers from 13 communities to deliver data spaces, 12 federation service providers.

The main EGI-ACE project objectives are:

- Deliver the European Open Science Cloud Compute Platform and expand the supply-side
- Contribute to the implementation of the EU Data Strategy and the EOSC Data Commons to support the Green Deal, Health and Fundamental Research
- Integrate the EOSC Compute Platform with the EOSC Portal and the EOSC Core
- Contribute to the realization of a global Open Science Cloud
- Expand the demand-side of EOSC across sectors and disciplines

A key aspect of the project is that it offers compute and storage resources, compute platform services, data management services and related user support and training. The total capacity that EGI-ACE makes available through the call between 2021-2023 is: 80 million of CPU hours, 250,000 GPU hours and 20 petabyte of storage. In order to offer access to infrastructure and platform services as well as to a dedicated user support and training, EGI-ACE makes periodic calls for use cases. Applicants are requested to provide technical and scientific information so that use cases that will be selected through the call. This represented a great opportunity and thus we decided to prepare the application for *Fermi*-LAT analysis workflow. Based on our previous experiences and looking at service portfolio of EGI-ACE already during the submission proposal we decided to exploit Dynamic On Demand Analysis Service (DODAS) [1].

4. Dynamic On Demand Analysis Service: DODAS

Dynamic On Demand Analysis Service (DODAS) is a Platform as a Service tool built combining several solutions and products developed by the INDIGO-DataCloud H2020 project[8] and now part of the EOSC-hub H2020 Project[9]. DODAS allows to instantiate on-demand complex infrastructures over any cloud with almost zero effort and with very limited knowledge of the underlying technical details. In particular DODAS provides the end user with all the support to deploy from scratch a variety of solution dedicated (but not limited) to scientific data analysis. For instance, something that we identified as a key feature, is that DODAS allow to create cluster of resources

PoS(ISGC2022)023

with almost zero effort represented for us a very important aspect because of the very reduced effort we can allocate to computing specific activities. In particular we identified the DODAS provides baselines solutions ready to be used and to be possibly extended as completely suitable for us:

- an HTCondor batch system
- a Caching on demand system based on XRootD
- Spark+Jupyter cluster for interective and big-data analysis

While the first two have been identified as the top priorities for us the third is a nice to have and maybe a something for the future.

4.1 Porting Fermi-LAT on Cloud with DODAS

The main objective for us was to exploit DODAS as a platform for the FLT project to provide a all-in-one system, with HTCondor batch on-demand and possibly extendable toward additional type of workflow such as the interactive data analysis. Any type of workflow must easily support customization in terms of software packages and software dependencies. In principle it is a requirement that any user can possibly define a specific run-time environment. Another requirement identified was to enable a effective remote data read and access. Given the relatively small amount of data, many options are open for us: from active data replication up to data caching. In general we are looking for a flexible integration with storage and cache. Finally in view of future extensions in terms of data to analyze and type of workflows we identified two aspects: being able to seamlessly federate resources provided by multiple clusters, particularly important for us to satisfy peaks of usages, finally we look for a solution which allow to easily interact with cloud system from our local work station. Starting from the basic DODAS features the implemented architectures is explained more in details in sec. 4.2 and sec. 4.3.

4.2 Compute integration

We focused primarily on enabling the batch on demand for *Fermi*-LAT by meaning instantiating HTCondor batch system and we worked on generating our community specific run-time environment. DODAS manages container based applications and thus it relies on container orchestration for application deployment and management. Technically this means that we were allowed to integrate a user tailored run-time environment by using our own images. Moreover DODAS integrates CVMFS and thus we have been provided with a private CVMFS server which grant us the possibility to deploy the analysis software even on the worker-nodes of the HTCondor cluster, with this we completely satisfied our needs. Regarding the deployment of the applications (the batch system in this case) DODAS leverages the templating capabilities of Helm [10]. This has been seen as a very important features for us because allow to concentrate on specific configuration (i.e. how many worker-nodes, how much disc size, RAM etc) instead of spending time in installing software etc. Moreover, being k8s based system, DODAS offers a kubeapps[11] integration which make the management of the system a really lightweight system admin activity. Finally the DODAS support of Token based Authentication represents a secure way to implement the remote job submission which is actually a very useful feature we aimed to exploit.

PoS(ISGC2022)023

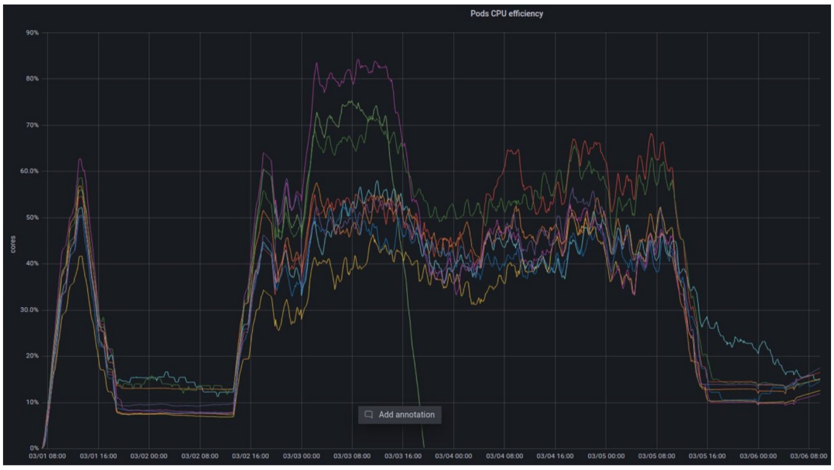

(a)

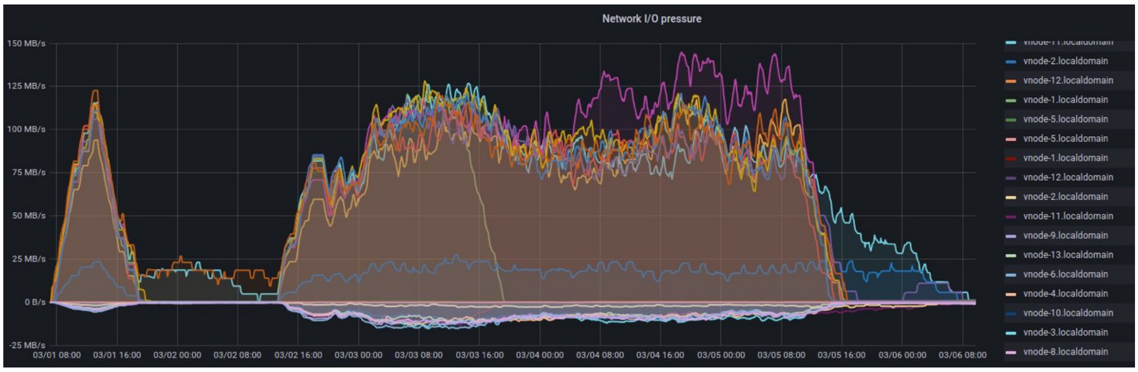

(b)

Figure 1: In picture (a) we can see the cpu efficiency of *Fermi* jobs. The value is referred to the first implementation and not yet optimized. Picture (b) represents the throughout generated by the worker-nodes that analyze downtime data present on the MinIO storage.

4.3 Data Integration

For what concern the data access we decided to deploy a layer of cache close to the compute part. In this respect we exploited the DODAS strategy to serve data archives and since we aimed at a POSIX based solution we decided for the MinIO[12] integration. As stated DODAS is based on k8s and in turn this allow to use the k8s operator provided by MinIO allows for seamless and on-demand provisioning of S3 compatible object storage on top of Kubernetes[13] any volume provider. The operator has been instantiated via the official HELM chart, and afterward, a dedicated object storage has been created over 16 volumes distributed on 4 Virtual Machines with a declarative approach, simply indicating the requested features as for any other Kubernetes resources. The MinIO operator will then take care of instantiating the services as asked for, and eventually update them when any changes occur in the user configuration. Furthermore, users can access their data not only via S3 APIs but also via a POSIX-like filesystem mounted inside the containers via RClone[14] mount technology; enabling a user-friendly experience for sharing and inspecting experiment libraries and code. In Figure 1 can be seen the CPU efficiency and the throughout generated by the worker-nodes.

PoS(ISGC2022)023

5. Physics achievements: FLT Results

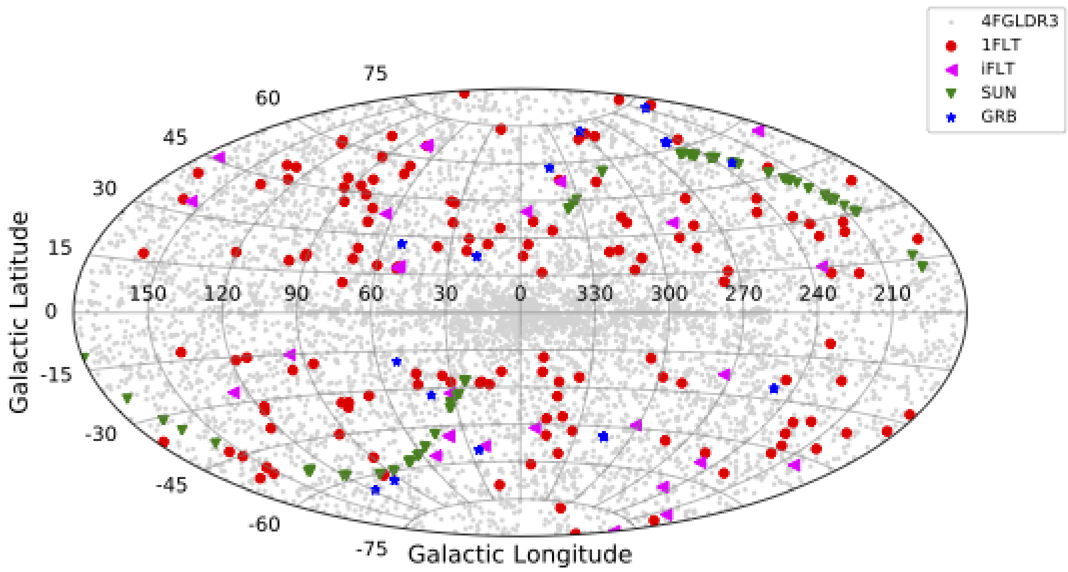

Figure 2: Aitoff projection of 1FLT sources represented in Galactic coordinates (red filled circles), iFLT sources (magenta filled triangles), Sun detections (green filled triangles) and GRB detections (blue filled stars) superimposed on 4FGL-DR3 sources represented in gray.

The FLT analysis investigated the *Fermi*-LAT sky on monthly timescales to search for sources which are not detected when integrating on timescales of years and/or that are not already hosted in the *Fermi*-LAT general catalogs, based on the long-term observation of sources. The FLT analysis found ~ new γ-ray sources per years of which ~ 70% are associated with soft AGN-type counterparts. Approximately one third of FLT sources remain unassociated. This fraction is similar to the percentage of unassociated sources found in the *Fermi*-LAT general catalogs. The FLT method allowed us to collect also 72 detections associated with the Sun and 27 associated with 14 of the brightest *Fermi*-LAT γ-ray bursts (GRBs).

The discovery of new γ-ray sources shows that integrating over different time intervals does not reproduce the same γ-ray sky and confirms that if the flare intensity is fainter than the integrated background over years, we lose the capability to detect these sources in long-time integrated catalogs.

Porting our analysis in Cloud via DODAS is opening new opportunities either to exploit the distributed resources in a transparent manner but also to reduce our computation time.

The FLT results are available distinguished in 1FLT, a catalog based on the first then years (August 2008-August 2018) of *Fermi* operations [15], [16], [17] and iFLT, the incremental list from August 2018, updated yearly [18],[17]. The sky locations of the 1FLT and iFLT sources together with Sun, GRB and 4FGL-DR3 (the latest *Fermi*-LAT general catalog data release [19]) detections are shown in Figure 2.

PoS(ISGC2022)023

6. Summary and Conclusions

The current experience has been very positive and we did not discovery any show stopper and this is a motivation for further developments. We successfully completed the porting our analysis in Cloud via DODAS and this is opening new opportunities either in the way we can exploit the distributed resources in the close future which in turn means to reduce our computation time. The possibility to easily deploy storage such as MinIO offer, represent an opportunity for our needs of synchronize multiple distributed instances. As an next step we will extend what we have done so far toward multiple providers belonging the EGI-ACE project such as the INFN-CLOUD-CNAF. For the future we are looking for a integration of the developed system in the context of the INFN Cloud project [20] were technically wise there is a full compliance with the solutions we have integrated in this project.

References

[1] D. Spiga et al., *EPJ Web Conf.* **245** (2020) 07033.

[2] W. B. Atwood et al., *ApJ* **697** (2009) 1071.

[3] `https://fermi.gsfc.nasa.gov/ssc/data/`.

[4] F. Damiani et al., *ApJ* **483** (1997) 370.

[5] S. Ciprini et al., in *AIP Conference Proceedings*, vol. 921, pp. 546–547, AIP, 2007, DOI.

[6] `https://fermi.gsfc.nasa.gov/ssc/data/analysis/`.

[7] M. Wood et al., *PoS (ICRC2017)* **824** (2017) .

[8] `https://www.indigo-datacloud.eu`.

[9] `https://www.eosc-hub.eu`.

[10] `https://helm.sh/`.

[11] `https://kubeapps.com/`.

[12] `https://min.io/`.

[13] `https://kubernetes.io/`.

[14] `https://rclone.org`.

[15] L. Baldini et al., *ApJS* **256** (2021) 13.

[16] `http://www.ssdc.asi.it/fermi1flt`.

[17] I. Mereu et al., *PoS (ICRC2021)* **641** (2022) .

[18] `https://www.ssdc.asi.it/fermi_iflt/`.

PoS(ISGC2022)023

[19] S. Abdollahi et al., *arXiv preprint arXiv:2201.11184* (2022) .

[20] D. Salomoni et al., *PoS (ISGC2022)* (2022) .

PoS(ISGC2022)023