VIALACTEA Science Gateway for Milky Way Analysis

Eva Sciacca, Fabio Vitello, Ugo Becciani, Alessandro Costa

INAF-Osservatorio Astrofisico di Catania, Italy

Akos Hajnal, Peter Kacsuk, Zoltan Farkas, Istvan Marton

Laboratory of Parallel and Distributed Systems SZTAKI, Budapest, Hungary

Sergio Molinari, Anna Maria di Giorgio, Eugenio Schisano, Scige John Liu, Davide Elia

INAF-Istituto di Astrofisica e Planetologia Spaziali, Roma, Italy

Stefano Cavuoti, Giuseppe Riccio, Massimo Brescia

INAF-Osservatorio Astronomico di Capodimonte, Napoli, Italy

Abstract

This paper presents the latest developments on the VIALACTEA Science Gateway in the context of the FP7 VIALACTEA project. The science gateway operates as a central workbench for the VIALACTEA community in order to allow astronomers to process the new-generation surveys (from Infrared to Radio) of the Galactic Plane to build and deliver a quantitative 3D model of our Milky Way Galaxy. The final model will be used as a template for external galaxies to study star formation across the cosmic time. The adopted agile software development process allowed to fulfill the community needs in terms of required workflows and underlying resource monitoring. Scientific requirements arose during the process highlighted the needs for easy parameter setting, fully embarrassingly parallel computations and large-scale input dataset processing. Therefore the science gateway based on the WS-PGRADE/gUSE framework has been able to fulfill the requirements mainly exploiting the parameter sweep paradigm and parallel job execution of the workflow management system. Moving from de-

---
*Corresponding author: *Eva Sciacca*, Email address: eva.sciacca@oact.inaf.it

velopment to production environment an efficient resource monitoring system has been implemented to easily analyse and debug sources of potential failures occured during workflow computations. The results of the resource monitoring system are exploitable not only for IT experts, administrators and workflow developers but also for the end-users of the gateway. The affiliation to the STARnet Gateway Federation ensures the sustainability of the presented products after the end of the project, allowing the usage of the VIALACTEA Science Gateway to all the stakeholders, not only to the community members.

1. Introduction

The Milky Way is a complex ecosystem where a cyclical transformation process brings diffuse baryonic matter into dense unstable condensations to form stars that produce radiant energy for billions of years before releasing

5 chemically enriched material back into the InterStellar Medium in their final stages of evolution. Although considerable progress has been made in the last two decades in the understanding of the evolution of isolated dense molecular clumps toward the onset of gravitational collapse and the formation of stars and planetary systems, a lot remained still hidden.

10 The aim of the European FP7 VIALACTEA project was to exploit the combination of all new-generation surveys of the Galactic Plane to build and deliver a galaxy scale predictive model for star formation of the Milky Way. This model will be used as a template for external galaxies and studies of star formation across the cosmic time. Usually the essential steps necessary to unveil the inner

15 workings of the galaxy as a star formation engine (such as the extraction of dust compact condensations or robust reconstruction of the spectral energy distribution of objects in star-forming regions) are often carried out manually by the

2

astronomers, and necessarily, over a limited number of galactic sources or very restricted regions. Therefore, scientists required new technological solutions able
₂₀ to deal with the growing data size and quantity coming from new-generation surveys (from Infrared to Radio wavelength).

The extraction of the meaningful information contained in the available data required an entirely new approach (the new paradigm of "data-driven scientific discovery") which resulted in a novel framework[1] based on advanced
₂₅ visual analytic techniques[2], data mining methodologies[3], machine learning paradigms[4] and Virtual Observatory (VO) based data representation and retrieval standards[5]. All the underlying pipelines required by this framework (e.g. knowledge base catalogue creation or maps making for visual analytic) are available through the VIALACTEA Science Gateway (VLSG).

₃₀ The gateway (described in Section 3) is based on the WS-PGRADE/gUSE [6] portal framework that provides several ready-to-use functionalities off-the-shelf. It allows development of scientific workflows composed of "nodes", corresponding to almost any kind of application, in a convenient graphical user interface. Workflows can be executed in parallel in a wide set of Distributed Comput-
₃₅ ing Infrastructures such as grids, clusters, supercomputers, and clouds. It enables sharing, importing and exporting workflows and managing credentials (and robot certificates), or gathering workflow execution statistics. Beyond these features, the portal is extensible, in fact, WS-PGRADE/gUSE offers a number of interfaces to add new applications and portlets to its base capabilities.

₄₀ This paper presents the latest developments on the VLSG including the workflows designed for the community and the resource monitoring system. The workflows (see Section 5) are mainly focused on performing CPU and data-intensive computations: map making, i.e. the formation of sky images from the instruments data; data mining to obtain band-merged catalogues relating
₄₅ sources with associated counterparts at different wavelengths; filamentary structure detection and extraction from images.

Due to the diverse variety of software and computing capabilities required by the workflows, a novel monitoring system has been developed within the gate-

3

way to supervise the status of the whole system. Monitoring covers different levels of tests (see Section 3.6), checking gateway interoperability with the computing infrastructures and workflow submission and execution processes. These tests are performed periodically and the resulting reports are published on the gateway so end-users are also aware of any potential failure of the system avoiding waste of time in debugging their work. Furthermore e-mail alerts are sent on any errors to infrastructure administrators to promptly fix the problems.

The remainder of the paper is organized as follows. Section 2 overviews the requirements and the developed technological architecture. Section 3 introduces the VLSG. Section 4 presents the federated network of science gateways (STAR-net) of the astronomical and astrophysical community. Section 5 demonstrates four applications developed using the science gateway that perform various calculations used to analyze the star forming regions of the Milky Way. Section 6 overviews the related work; finally, Section 7 concludes the paper.

2. VIALACTEA Requirements and Technological Architecture

In order to deliver a model of our galaxy with quantitative star formation laws, it is necessary to reveal and analyse throughout the galaxy the dense filamentary clouds where star-forming clumps are found. These clumps are found in very different environments and in different evolutionary stages and their properties are characterized through detailed modelling of their Spectral Energy Distribution. Their exact location is determined using the most up to date distance estimators and all these pieces need to be assembled to get a new view of our Galaxy.

The Galactic distribution of Star Formation Rate (stellar mass produced per unit time) and Efficiency (stellar mass produced per unit mass of available dense gas) can be quantitatively related to the variety of physical agents that drive star formation in the Galaxy. The timely exploitation of the huge amount of data available requires new technological solutions able to overcome the current challenges pushing the envelope of the current state of the art both from
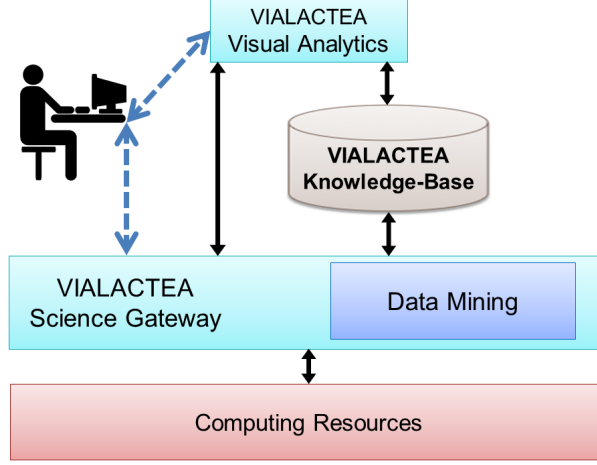
4

Figure 1: VIALACTEA integrated technological framework.

technological and scientific point of view. Therefore a novel system has been
implemented[7] based on: advanced visual analytics techniques, data mining
pipelines, VO-based standards and science gateway technologies. The frame-
work can be seen as an integrated workspace where the Visual Analytics desk-
top client, the Science Gateway embedding the Data Mining pipelines and the
VIALACTEA Knowledge Base can be employed both as independent actors
or as interacting components (see Figure 1). In the following paragraphs we
highlight the technological solutions developed to overcome the aforementioned
requirements.

Data Requirements. Data challenges required in VIALACTEA has been man-
aged through an archive named VIALACTEA Knowledge Base[5] (VLKB) which
includes a combination of storage facilities, a Relational Data Base (RDB) server
and web services on top of them. It allows easier searches and cross correlations
between data, and currently contains: 2D surveys, catalogue sources and related
band merged information; structural information such as filament structures or
bubbles; and Radio Datacubes with search and cutout services. Data-mining
and machine-learning pipelines are embedded within the Science Gateway as
workflows and employed to carry out building of Spectral Energy Distributions,

5

distance estimate and Evolutionary classification of hundreds of thousands of star forming objects on the Galactic Plane. All these produced results are then ingested to the VLKB. The VIALACTEA Visual Analytics (VLVA) tool[2] allows the interaction with the VIALACTEA data and to carry out complex tasks for multi-criteria data/metadata queries on the VLKB, subsample selection and further analysis processed over the science gateway, or real-time control of data fitting to theoretical models.

Analysis Tools. The science gateway is exploited by the scientists to configure and run the VIALACTEA workflows implementing the pipelines developed by the community (see Section 5). Furthermore the science gateway allows the VLVA to submit workflows through the usage of WS-PGRADE/gUSE Remote API [8] . This API also provides methods for checking the workflow's status, and for downloading the outputs. The scientists required easy parameter setting, fully embarrassingly parallel computations and large-scale input dataset processing. Therefore we selected the WS-PGRADE/gUSE framework that was suitable to fulfill the requirements (see Section 3).

Software Development. Due to the cross-domain scientists involved in the community (computer scientists, technologists and astronomers) agile software development approach has been adopted. This approach in fact promotes adaptive planning, evolutionary development, early delivery, and continuous improvement, and it encourages rapid and flexible response to change. Cross-disciplinary face-to-face meetings have been organized to promote an iterative, incremental and evolutionary framework based on several cycles of requirements and feedback sessions.

3. VIALACTEA Science Gateway

Usage of science gateways provides user-friendliness (intuitive user interface), efficiency (fast response time even for complex user requests), scalability

(fast response time even for a large number of simultaneous user requests), robustness (keeps working under any circumstances and recovers gracefully from exceptions) and extensibility (easy to be extended with new interfaces and functionalities).

The VIALACTEA Science Gateway[1] developed in the VIALACTEA project is based on a customized version of WS-PGRADE/gUSE.

WS-PGRADE/gUSE is a science gateway framework, which is, by default, not specialized for a certain scientific area, thus scientists from many different areas can use it, typically, after tailoring it to the needs of the given scientific community. WS-PGRADE/gUSE offers a large set of generic tools and services out-of-the-box, including certificate manegement, users management, accounting, file and data management, job submission, workflow creation and management and monitoring components in the frontend, and DCI and data storage access mechanisms in the backend, respectively. WS-PGRADE/gUSE is also extensible, custom user interface components (menus, portlets, skins, etc.) can be added to the frontend and also the backend can easily be extended (e.g. new DCI or data access components) due to its plugin architecture. Moreover, accessing WS-PGRADE/gUSE services from program code are supported by APIs.

The VLSG is a highly tailored portal based on WS-PGRADE/gUSE gateway framework technology, which aims at fulfilling the requirements of astronomers' user community, and more specifically, to reach the goals that arose during the VIALACTEA project.

This section outlines some of the characteristics of WS-PGRADE/gUSE that have been identified as key features for the VLSG.

3.1. Workflow Concept

WS-PGRADE/gUSE offers a workflow management system, in which scientific applications are designed, interpreted and enacted in the form of data-

---

[1]VIALACTEA Science Gateway: `https://vialactea-sg.oact.inaf.it`

driven workflows to access, process, filter and visualize scientific data in an automated way. Workflows are basically directed acyclic graphs, where each node corresponds to one particular standalone entity of computation (such as an executable, a web-service invocation, etc.), whereas edges (ending with so called "ports") represent potential inputs, outputs and data passing, respectively, between these basic blocks. During workflow execution, jobs are orchestrated automatically in a data-driven way, that is, a node is executed (scheduled for execution) as soon as all of its inputs become available, potentially, all those nodes in parallel having this condition fulfilled.

The workflow development process consists of three main steps. First, an "Abstract Workflow" is designed determining the basic blocks of computations and the possible flows of data, then its "Concrete Workflow" is created by specifying particular executables for nodes and configuring specific computation resources for these nodes. The concrete then can be actually executed resulting in individual "Workflow Instances".

For defining abstract workflows several data-flow and control-flow patterns are available in WS-PGRADE/gUSE. Besides defining pipelines, parallel branches, joins, it is also possible to implement conditional branches, recursion (by embedding other workflows into nodes), and parameter sweep applications, where the same node may imply multiple job instances to process the given set of inputs (or different combinations of different input sets such as Cartesian or dot product).

Within the VIALACTEA project an integrated workflow editor has been developed[9, 10] that allows end-users to compose abstract workflows, configure concrete workflows and run them within the same graphical user interface using a web browser only.

3.2. Distributed Computing Infrastructures Interoperability

Solving distributed computing infrastructure (DCI) incompatibility issues in a generic way is a challenging and complex task. Most of the gateways and workflow management systems are thus tightly bound to a specific or a limited

8

number of supported DCIs, respectively. In WS-PGRADE/gUSE, interoperability between the gateway and the different types of DCIs is done using a dedicated, individual component, named "DCI Bridge", whose task is to seamlessly enable execution of workflow jobs in various DCIs.

The communication between the gateway and the different types of DCI middlewares is realized using a set of pre-implemented DCI Bridge plugins. DCI Bridge already supports most major DCI platforms out of the box such as ARC, Boinc, Globus, gLite, UNICORE, PBS, as well as web services, Google App Engine, Sun Grid Engine, cloud-brokering services (CloudBroker), or private and public cloud platforms (EC2, CloudSigma, OCCI).

The different DCI Bridge plugins take care of individual job submissions to the DCIs, maintain different queues, schedule job executions, monitor their states, make logs about the most important events and handle error with the predefined failover mechanisms. DCI Bridge plugins contain thus all the necessary program codes to communicate with the different types of DCIs using the corresponding APIs and middleware libraries, whereas DCI Bridge's frontend interface hides all these details by showing a single and uniform Open Grid Services Architecture (OGSA) Basic Execution Service (BES) compliant interface towards the gateway.

Beyond job execution, the DCI Bridge is also capable of staging in and out data files (inputs and outputs) between the internal storage of the gateway and DCI worker nodes, which enables to exchange data between workflow jobs without the need of dedicated storages installed. This solution is suitable for transferring file of size from small to medium (up to a few hundred megabytes). For larger data, WS-PGRADE/gUSE offers data bridging services, called Data Avenue, that allows accessing various storage resources from compute sites, as described in section 3.4.

Selecting the appropriate DCI as the target platform for execution of a specific job is very easy for the end-users: they simply select one of the supported middlewares connected with the portal. In some cases, users are not even required to provide credentials for accessing the selected DCI (such as usernames

and passwords) due robot certificates automatically available for all authenticated users. This simplicity however does not hide the possibility to control job submission at a very low level if needed. Within the context of the VIA-

₂₁₅ LACTEA project, it was required not merely to enable jobs to be submitted to a cluster with robot permissions, but certain jobs required to be run on specific worker nodes that meet specific criteria such as having at least a given number of processor cores, size of available memory, respectively. These fine grained control options are also available for advanced workflow developers on

₂₂₀ the graphical user interface of the science gateway.

3.3. Levels of Parallelism

Four levels of parallelism are supported by the workflow management system of WS-PGRADE/gUSE. At the lowest level, which is called node-level parallelism, the application itself is prepared to utilize multi-core processors or

₂₂₅ cluster systems, respectively. In multi-core environments, these applications are designed as multi-threaded applications. In cluster systems, applications use specific programming libraries that implement MPI specifications (such as open MPI).

At the next level, parallel execution of jobs corresponding to nodes residing

₂₃₀ on different branches in the workflow graph is supported, which is probably the most intuitive concurrent execution (branch-level parallelism).

At the third level of parallelism, the same job is executed on a large parameter field, which are called parameter study or parameter sweep (PS) applications. The series of different inputs to be processed by multiple instances of the

₂₃₅ same PS node is produced by so called "generator nodes" or "generator ports", whereas the collection of outputs of the PS job instances is done by a "collector node".

The possibility of the execution workflows concurrently means the highest level of parallelism in WS-PGRADE/gUSE. In the most simple case, parallel

₂₄₀ execution of the same workflow is initiated by the user with different configurations, but as workflows can be embedded into other, concurrent enactment of

10

multiple workflows is also possible.

As described in section 3.2, workflow jobs can be submitted to and executed in various remote computing resources. When creating concrete workflows from abstract ones, for each node a particular computing infrastructure is selected. During workflow configuration each node may be associated with a different DCI even within the same workflow; so, for example, the first node in a workflow could be configured to be executed locally (within the DCI Bridge), a subsequent node in some PBS cluster, and another node in some cloud, respectively. WS-PGRADE/gUSE automatically handles job submission to the appropriate site of computation and data staging.

Beyond this flexibility of selecting the most suitable platform for the different jobs corresponding to workflow nodes, a recent improvement of WS-PGRADE/gUSE allows to distribute job instances of the same (PS) node to different DCIs, which technique is called meta-brokering[11]. Using meta-brokering, the overall execution time can significantly be reduced by distributing and balancing the load among multiple DCIs, to avoid overload of a single computing resource. Two types of load balancing are supported: static and dynamic. In the static case, at each individual job execution one of the given DCIs is selected with probability corresponding to predefined resource weights. In the dynamic case, such resource weights are updated periodically, and so this kind of brokering considers the actual load of the set of available DCIs. In the Vialactea project, meta-brokering allows distributing the load accross multiple computing resources provided by the STARnet federation.

3.4. Remote Storage Management

Data-intensive scientific applications should preferably rely on high performance storage systems for efficiency reasons. To allow workflow jobs to access data residing on storages WS-PGRADE/gUSE provides a data bridging service called Data Avenue [12].

Numerous storage solutions have been evolved during the past decades. Besides the conventional HTTP/FTP servers, storage resources accessible over

SFTP (Secure File Transfer Protocol), GridFTP protocols, Storage Resource Management (SRM) systems, as well as logical file systems such as LCG File Catalogs, integrated Rule-Oriented Data-management Systems (iRODS) and recent cloud storages such as Amazon S3 can potentially be used from within grids and clusters. Accessing these storage resources however typically requires dedicated protocols and tools to allow users of up or downloading data. The problem is that it is not always possible to install the necessary softwares on computing resources.

Data Avenue solves this issue by offering a uniform interface with a simple REST API through which a wide range of storage resources become available for workflow jobs. Worker nodes in this way can communicate with Data Avenue with commonly available tools such as wget or curl to transfer files, whereas Data Avenue does the necessary protocol conversion, i.e. data bridging, towards the given storage resource.

Data Avenue currently supports HTTP, HTTPS, SFTP, GSIFTP, SRM, iRODS and S3 protocols, and so workflows can use any of these storage resources to download data from or upload data to, respectively. It facilitates the handling of extremely large files as well (e.g., Amazon supports objects of size up to 5 TB). Moreover, data bridging solution can also offer even more efficient data staging, as it avoids transferring data through a chain of intermediate sites (for example, gateway storage - DCI Bridge - PBS head node - PBS worker node).

Beyond data bridging services, Data Avenue allows to transfer files between any of the supported storage systems, e.g. from FTP to S3, and allows for managing and organizing data files in storage resources (deleting, renaming, creating folders).

In the VIALACTEA project an individual Data Avenue service has been deployed close to the computing infrastructure to make such data transfers as efficient as possible.

3.5. Workflow-level Debugging

WS-PGRADE/gUSE allows for debugging workflow executions by applying so called breakpoints. Beyond merely stopping and continuing the flow of control at specific points, users can also interact with workflows at runtime: they can directly influence and change the subsequent behaviour of the workflow by enabling or prohibiting the execution of individual job instances at breakpoints.

The built-in workflow interpreter (WFI) supports job instance-level checkpointing, i.e., the interpretation of the workflow can be interrupted by the "SUSPEND" command at any time, whereas "RESUME" command allows its continuation. Jobs in transient state ("submitted" or "running") can be killed, which allows for deleting job instances in error state and so further enactment of the workflow can be continued based on successfully completed jobs.

There are two basic forms of the debugging: 1) postponing the execution of all instances of a specific job; 2) postponing the execution of all jobs followed by a specific job.

It is also possible to specify timeouts for breakpoints. When it is set, the workflow execution stops at the given breakpoints, and without user interaction, it is going to continue after the specified timeout (by default 60 minutes).

Debugging possibility is a very important feature in workflow systems. VIA-LACTEA workflows perform very complex tasks involving many steps and iterations, whereas tasks perform very diverse computations. During workflow design, development and testing, finding root causes of workflow execution failures could have been extremely tedious, costly and time consuming without a debugging tool.

3.6. Resource Monitoring

Continuous monitoring of the operational status, or briefly the "health" of the underlying distributed computing infrastructures is of high importance. Any outage of the underlying DCIs can cause breaking the flow of calculations, and in spite of built-in failover mechanisms, it can be very difficult to localize the fault without having proper information about the current and past behavior

of the given computing infrastructure. Sometimes, these errors are not even repeatable; temporary blackouts and failures (e.g. when worker nodes run out of disk space) may prevent the system to record any notice about the actual cause.

Using a DCI monitoring system, such as the one designed and implemented in the VLSG, workflow developers can make sure that all the related DCIs operate normally prior to starting long running calculations; and also, on error, they can be sure that the error is not caused by the failure of the underlying infrastructure, respectively, by revising historical data of monitoring records. System administrators can also benefit from using resource monitoring, as they can quickly overview all the systems under their supervision, and thanks to e-mail alerting, they can react to the corrupted behavior as soon as possible. Historical monitoring data of computing resources may help to plan for potential improvements, and to provide measures to prevent similar failures in the future.

At the moment, resource monitoring for Portable Batch Systems (PBS) is available. To help in better localize the errors four levels of monitoring have been designed: Level 1 (PBS cluster infrastructure head node monitoring); Level 2 (PBS cluster worker nodes environment monitoring); Level 3 (Portal PBS cluster interoperability monitoring); and Level 4 (VIALACTEA, domain-specific, workflow operational monitoring). For the different levels different "frequency" can be specified, i.e., how often and at what time they are re-tested.

Level 1 checks that whether the DCI is indeed accessible from the gateway (head node responds to ping, successful SSH connection can be established) and all the essential middleware commands (`qsub`, `qstat`, `pbsnodes`, etc.) operate as expected. Level 2 tests scan through all worker nodes of the DCI and check that whether the expected execution environment is available, such as enough disk space, necessary libraries (Java, IDL, Matlab, Python, etc.). Level 3 tests execute a probe workflow in all DCIs. Level 4 tests submit domain-specific workflows, having characteristics similar to full-fledged applications, but with parameters resulting in the least possible load to the DCI.

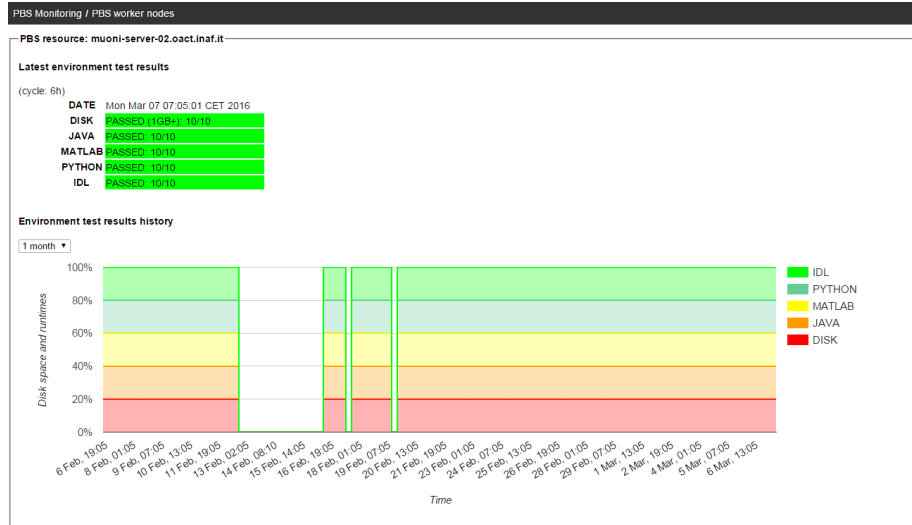Monitoring data can be viewed in a graphical user interface in the gateway,

14

Figure 2: Monitoring results of "PBS worker nodes" tests (level 2).

summarized and displayed in the form of tables and charts. Figure 2 shows
level 2 results. The table (on the top in figure 2) shows the latest results and
the re-evaluation frequency (6 hours). PASSED 10/10 means that all tests (free
disk space, Java, Matlab, Python, IDL) had been passed on all the 10 worker
nodes. The chart (on the bottom in figure 2) shows test results of 30 days,
indicating outage on dates 12–16, 17, and 19 February.

4. STARnet Affiliation

The VLSG is affiliated with the STARnet Gateway Federation [13]. STAR-
net is a unique example of a federated network of Science Gateways based
on WS-PGRADE/gUSE technologies, and explicitly designed and tuned to the
needs of the Astronomical and Astrophysical (A&A) community in Europe. The
use of a federated gateway infrastructure allows exploring new collaboration op-
portunities for advancing the scientific research activity within A&A. STARnet
envisages sharing a set of services for authentication, a common and distributed
computing infrastructure, data archives and workflow repositories. Each STAR-
net gateway provides access to specialized applications via customized work-

15

flows.

As required from the STARnet architecture, the gateway has been developed by means of virtual machines containing the WS-PGRADE/gUSE gateway installation and proper configuration for VIALACTEA. The Liferay and WS-PGRADE/gUSE databases, the local WS-PGRADE/gUSE storage and the local WS-PGRADE/gUSE application repository are configured on the hosting machine to facilitate the regular upgrading procedure of the new virtual machine containing bug fixes and latest WS-PGRADE/gUSE gateway releases.

The affiliation to the STARnet Gateway Federation ensures the sustainability of the products after the end of the VIALACTEA project. This will allow the usage of the science gateway by all the future possible stakeholders, and not only by the VIALACTEA community.

5. Milky Way Analysis Through the Gateway

The applications and workflows developed on the gateway for the analysis of star forming regions within the Milky Way have been mainly devoted for: map making, i.e. the production of high quality images from the raw instruments data; data mining to obtain band-merged catalogues, whose entries consist of sources with associated counterparts at different wavelengths; filamentary structure detection and extraction from images.

5.1. MOSAIC

The MOSAIC workflow employs Unimap [14] as map maker software to produce high quality mosaic images from raw instruments data of the infrared imaging photometers onboard of the ESA Herschel satellite that observed the complete 360° of the Galactic Plane in 5 photometric bands (70, 160, 250, 350 and 500 μm). Unimap requires a complex setup, input data images should follow a pre-defined naming and format convention and input parameters employed for the different pipeline algorithms (to make the Time Ordered Pixels, signal pre-processing, glitch and drift removal, pointing correction etc.) are more than 50.

16

405 Therefore, the workflow has been implemented to facilitate these tasks, fixing some parameters to the most reliable values and letting the user change only the most sensitive ones; and pre-processing the input data images as required from the Unimap processing.
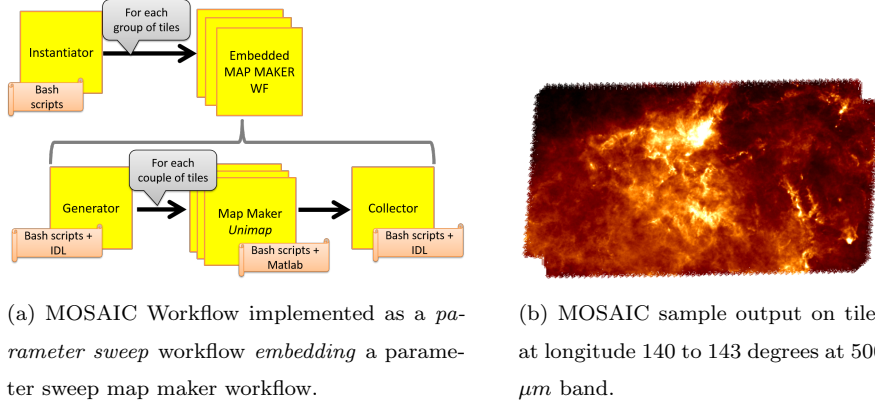

(a) MOSAIC Workflow implemented as a *parameter sweep* workflow *embedding* a parameter sweep map maker workflow.

(b) MOSAIC sample output on tiles at longitude 140 to 143 degrees at 500 μm band.

Figure 3: MOSAIC Workflow schema and sample results.

The employed applications are coded in IDL (for pre-processing input files
410 for Unimap), Matlab (for the Unimap application) and Bash scripting language (for setting up input parameters). The workflow has been implemented as a parameter sweep workflow [15] embedding a parameter sweep map maker workflow. This allows for full parallelization of the processes to be executed. See Figure 3a for the schema of the workflow.

415 The input, given as plane text file, specifies the tiles to be processed (longitude and wavelength) and a sub-set of the parameters of the Unimap application. The workflow automatically imports the required data from the Herschel infrared Galactic Plane Survey (Hi-GAL) [16, 17] stored into the computational resource storage available on the Data Avenue. The "Instantiator" job prepares
420 the different group of input tiles to be processed in parallel by the map maker embedded workflow, which computes each group of tiles separately. The "Generator" job prepares the tiles in couples to be processed in parallel by the Map Maker job (Unimap). Finally, the output is given by the "Collector" job of the map maker embedded workflow, which contains the maps in FITS (Flexible Im-

17

age Transport System) file format (see Fig. 3b for a sample output) for further studies of dust structures to discover nascent clump-forming filaments or sites of massive star formations.

5.2. PPMAP

The PPMAP workflow executes a Point Process MAPping (PPMAP) [18], which is a Bayesian procedure that uses images of dust continuum emission at multiple wavelengths to produce resolution-enhanced image cubes of differential column density as a function of dust temperature and position. Application of PPMAP to a filamentary complex shows that the decomposition into different temperatures facilitates the separation of different physical components along the line of sight and has the potential to provide insights into the mechanisms associated with column density Probability Density Functions of molecular clouds providing key information on the initial conditions for star formation.

(a) PPMAP Workflow implemented as a *parameter sweep* workflow *embedding* a parameter sweep map maker workflow.

(b) PPMAP sample output of differential column density on tiles at longitude from 136 to 138 degrees.
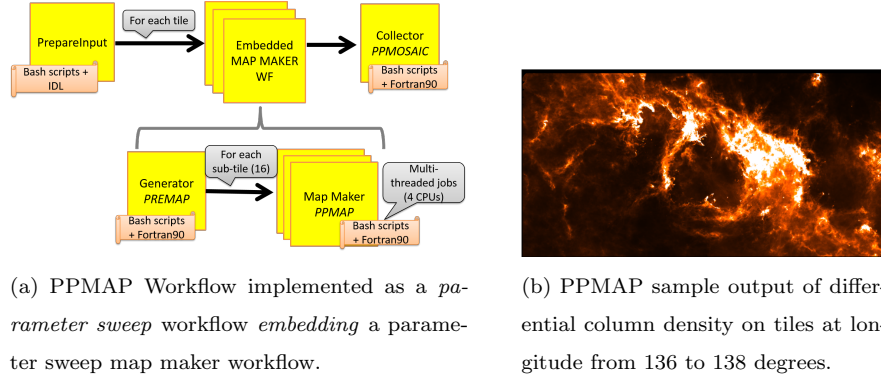
Figure 4: PPMAP Workflow schema and sample results.

The employed applications are coded in Fortran90 (for the PREMAP, PPMAP and PPMOSAIC tools), IDL (for tiles pre-processing) and Bash scripting language (for setting up input parameters). As for MOSAIC, this workflow has been implemented using the parameter sweep submission schema as shown in Figure 4a.

Inputs, uploaded as plane text files, specify the tiles to be processed and the parameters (one for each input tile) to be sent to the PPMAP application. The

18

workflow automatically imports the required data from the Hi-GAL Survey at all available wavelengths (from 70 to 500 μm). The "PrepareInput" job prepares the different tiles to be processed in parallel by the map maker embedded workflow, which processes each tile separately. The "Generator" job prepares the tiles in 16 smaller sub-tiles (typically 40 x 40 spatial pixels) to be processed in parallel by the Map Maker job (PPMAP) using a multi-thread OpenMP approach. The output is given by the "Collector" job of the workflow that does the mosaic of all the sub-tiles resulting in the final maps in FITS file format (see Fig. 4b for a sample output).

5.3. Q-FULLTREE

The Q-FULLTREE workflow performs compact source identification through band-merging. It allows to correlate the Hi-Gal catalogue with sources coming from other surveys employed for the ViaLactea project. The application is based on the positional cross-match among sources at different wavelengths. According to basic positional order relationships imposed by the variable spatial resolution and variation of position angle at different wavelengths, the background guideline is always the search for higher resolution sources within an elliptical region centred on the lower resolution counterpart and dimensionally limited by its two Full Width at Half Maximum values, defining the two axes of the ellipse.

The "Q-FULLTREE" job is configured as a multi-threading splitting the input catalogues in a given wavelength into a user-chosen number of small sub-catalogues, with a user-selected percentage of overlapping sources in order to avoid the loss of merged sources related to borderline sources. The post-processing jobs, "QualityRank", "QualityFitness" and "FT-Recap" are executed in parallel (see Fig. 5 for the workflow schema). In particular, the first two jobs calculate the quality rank and quality fitness of the processed sources; while FT-Recap (FullTree-Recap), is submitted to re-organize the output of the Q-FULLTREE in order to be ingested into the VLKB for the visualization of the Spectral Energy Distribution via the VLVA.

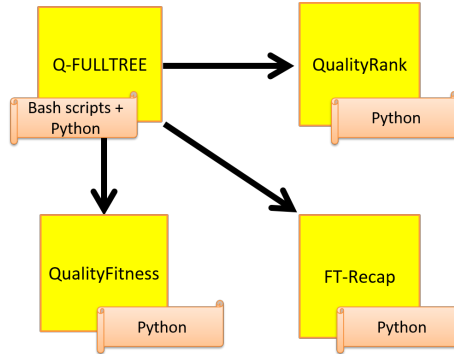The employed applications are coded in Python and make internal use of

19

Figure 5: Q-FULLTREE Workflow schema.

the STILTS public library [19]. The inputs of the workflow, provided by the user, are: a TAR archive containing the sources at different wavelengths in CSV format and two text files specifying the setup and the configuration for the application.

5.4. Filamentary Structure Detection

This workflow has been designed to perform filament extraction. The underlying application [20] identifies filamentary-like extended structures on astronomical images and determines their morphological and physical parameters. Filaments are defined as extended structures with elongated, cylindrical-like shapes, that present a relatively brighter contrast with respect to their surrounding.

The workflow is developed as a three steps processing: i) feature detection, ii) filament extraction and iii) filter artifacts and create the final catalogue. All these applications are implemented in IDL. The first step performs the detection of candidates through advanced image analysis techniques based on mapping of eigenvalues of the local Hessian Matrix computed from the input map. The second step analyses the region of interests with the support of morphological operators that decompose the initial binary mask into simpler units. Finally, the third step analyses the candidate list and filters out low elongated structures and possible artifacts building up the final candidate filamentary catalogue.

20

6. Related Work

To deal with the data deluge that the Astrophysics community is facing, different science gateways and workflow technologies are being exploited. Apart from the WS-PGRADE/gUSE framework that has been extensively employed by the authors, see e.g. [21, 22, 23], different approaches have been followed to allow the end users to easily interact with the applications ported on the DCIs.

In [24], the authors present an approach based on the Taverna Workbench[2] [25] and the Astrotaverna plugin[3] [26] to perform kinematical modelling of galaxies as an example of analysis task required by the SKA project (which aims to build an instrument that will be the worlds largest radio interferometer, able to reach data rates in the exa-scale). Apache Airavata[4] [27] environment on XSEDE[5] resources have been used in [28] to produce multiple synthetic sky surveys of galaxies and large-scale structure in support of Dark Energy Survey analysis. The underlying technologies described in those works are well suited to be ported into a science gateway such as VLSG, but requires time and extra IT effort for coding web services (as wrapper) on top of each application of interest of the astronomers.

The Kepler[6] scientific workflow system [29] has been employed in [30] to implement automatic data reduction pipelines. This approach could have been very useful within the VIALACTEA project but again it requires IT effort to build the required Kepler actors for each application.

Finally, to our knowledge, none of the above solutions included a resource monitoring system able to check the status of the overall gateway interacting components, including the required runtime libraries, as required by the VIALACTEA community.

There are several resource monitoring tools available, such as Ganglia[31],

---

[2]Taverna web site: `http://www.taverna.org.uk`

[3]AstroTaverna plugin: `http://amiga.iaa.es/p/290-astrotaverna.htm`

[4]Apache Airavata web site: `http://airavata.apache.org`

[5]XSEDE web site: `https://www.xsede.org`

[6]Kepler project web site: `https://kepler-project.org`

Nagios[7], Zabbix[8], Prometheus[9] to mention a few, shipped with numerous probes out-of-the-box to monitor typical host and service metrics such as availability, CPU, network utilization, memory, disk space usage, service checks, etc. In our case, however, where worker nodes in PBS clusters are inaccessible from outside (they reside in private network), these general tools proved to be inadequate, as their monitoring was possible only by submitting dedicated PBS jobs. Also, verifying the results of workflow execution, which can only be done using the "remote API" of the portal, cannot be done using such general tools. Our implementation, and its integration into the portal has other advantages as well: it uses the same monitoring source (host of the gateway) and mechanisms (software libraries, SSH connections, PBS commands) as the portal, so it tests resources from identical environment. Nevertheless, we connected our tool to Zabbix to record workflow execution time metric, and we used Zabbix triggers, notifications, and chart visualization.

7. Conclusions and Outlook

In this paper we have introduced a new framework that allows astronomers to process the new-generation surveys of the Galactic Plane to build and deliver a quantitative model of Milky Way Galaxy. The presented science gateway operates as a central workbench for the VIALACTEA community allowing to deal with the growing data size and quantity coming from new-generation surveys. The extraction of the meaningful information contained in the available data required an entirely new approach (the new paradigm of data driven scientific discovery), which resulted in a novel framework based on advanced visual analytics techniques, data mining methodologies, machine learning paradigms and Virtual Observatory based data representation and retrieval standards.

The focus of the presented workflow applications is on map making, i.e.

---

[7]Nagios web site: `http://www.nagios.org`

[8]Zabbix web site: `http://www.zabbix.com`

[9]Prometheus web site: `https://prometheus.io`

the formation of sky images from the instruments data; data mining to obtain band-merged catalogues relating galactic sources with associated counterparts at different wavelengths; and filamentary structure detection and extraction from sky images. Furthermore, we have highlighted how the usage of WS-PGRADE/gUSE framework have been able to fulfill the project requirements thanks to its key features: user-friendliness, efficiency, scalability, robustness and extensibility.

This paper also described a novel resource surveillance component integrated into WS-PGRADE/gUSE portal capable of checking operational status of the employed computational infrastructures based on Portable Batch Systems (PBS). The monitoring covers different levels of tests checking the gateway interoperability with the computing infrastructures and the workflow submission and execution processes. These tests are performed periodically and the resulting reports are published on the gateway so that also end-users are aware of any failure of the system avoiding waste of time in debugging their work.

Amongst the things deserving further studies is the evaluation of MetaBrokering service of WS-PGRADE/gUSE which is capable of distributing and balancing the load among different distributed computing infrastructures. This will be exploited for parameter sweep jobs, such as the map making computations, avoiding excessive load of one resource with respect to other having higher capacity.

Acknowledgment

References

[1] S. Molinari, R. Butora, S. Cavuoti, M. Molinaro, G. Riccio, E. Sciacca, F. Vitello, U. Becciani, M. Brescia, A. Costa, et al., Integrated data access,

23

575 visualization and analysis for Galactic Plane surveys: the VIALACTEA case, Proceedings of the International Astronomical Union 12 (S325) (2016) 291–298.

[2] U. Becciani, et al., Visual Analytics in Astrophysics: a novel tool integrated into VisIVO, Astronomical Data Analysis Software and Systems XXVI.

580 [3] G. Riccio, M. Brescia, S. Cavuoti, A. Mercurio, A. M. di Giorgio, S. Molinari, C 3, A Command-line Catalog Cross-match Tool for Large Astrophysical Catalogs, Publications of the Astronomical Society of the Pacific 129 (972) (2016) 024005.

[4] G. Riccio, S. Cavuoti, E. Schisano, M. Brescia, A. Mercurio, D. Elia, M. Benedettini, S. Pezzuto, S. Molinari, A. M. Di Giorgio, Machine learning based data mining for Milky Way filamentary structures reconstruction, in: Advances in Neural Networks, Springer, 2016, pp. 27–36.

[5] M. Molinaro, R. Butora, M. Bandieramonte, U. Becciani, M. Brescia, S. Cavuoti, A. Costa, A. M. Di Giorgio, D. Elia, A. Hajnal, et al., VIA-LACTEA knowledge base homogenizing access to Milky Way data, in: SPIE Astronomical Telescopes+ Instrumentation, International Society for Optics and Photonics, 2016, pp. 99130H–99130H.

[6] P. Kacsuk, Z. Farkas, M. Kozlovszky, G. Hermann, A. Balasko, K. Karoczkai, I. Marton, WS-PGRADE/gUSE Generic DCI Gateway Framework for a Large Variety of User Communities, Journal of Grid Computing 10 (4) (2012) 601–630.

[7] U. Becciani, et al., Advanced Environment for Knowledge Discovery in the VIALACTEA Project, Astronomical Data Analysis Software and Systems XXV.

600 [8] A. Balasko, Z. Farkas, P. Kacsuk, Building Science Gateways by Utilizing the Generic WS-PGRADE/gUSE Workflow System, Computer Science 14 (2) (2013) 307–325.

[9] G. A. McGilvary, M. Atkinson, S. Gesing, A. Aguilera, R. Grunzke, E. Sciacca, Enhanced Usability of Managing Workflows in an Industrial Data Gateway, in: e-Science (e-Science), 2015 IEEE 11th International Conference on, IEEE, 2015, pp. 495–502.

[10] F. Vitello, E. Sciacca, U. Becciani, A. Costa, P. Massimino, É. Takács, B. Szakál, Mobile application development exploiting science gateway technologies, Concurrency and Computation: Practice and Experience 27 (16) (2015) 4361–4376.

[11] K. Karoczkai, A. Kertesz, P. Kacsuk, A Meta-Brokering Framework for Science Gateways, Journal of Grid Computing 14 (4) (2016) 687–703.

[12] Á. Hajnal, Z. Farkas, P. Kacsuk, T. Pintér, Remote storage resource management in WS-PGRADE/gUSE, in: Science Gateways for Distributed Computing Infrastructures, Springer, 2014, pp. 69–81.

[13] U. Becciani, E. Sciacca, A. Costa, P. Massimino, F. Vitello, S. Cassisi, A. Pietrinferni, G. Castelli, C. Knapic, R. Smareglia, et al., Creating gateway alliances using WS-PGRADE/gUSE, in: Science Gateways for Distributed Computing Infrastructures, Springer, 2014, pp. 255–270.

[14] L. Piazzo, L. Calzoletti, F. Faustini, M. Pestalozzi, S. Pezzuto, D. Elia, A. di Giorgio, S. Molinari, UNIMAP: a generalized least-squares map maker for Herschel data, Monthly Notices of the Royal Astronomical Society 447 (2) (2015) 1471–1483.

[15] P. Kacsuk, K. Karoczkai, G. Hermann, G. Sipos, J. Kovacs, WS-PGRADE: Supporting parameter sweep applications in workflows, in: Workflows in Support of Large-Scale Science, 2008. WORKS 2008. Third Workshop on, Ieee, 2008, pp. 1–10.

[16] S. Molinari, B. Swinyard, J. Bally, M. Barlow, J.-P. Bernard, P. Martin, T. Moore, A. Noriega-Crespo, R. Plume, L. Testi, et al., Hi-gal: The

630   herschel infrared galactic plane survey, Publications of the Astronomical Society of the Pacific 122 (889) (2010) 314.

[17] D. Elia, S. Molinari, Y. Fukui, E. Schisano, L. Olmi, M. Veneziani, T. Hayakawa, M. Pestalozzi, N. Schneider, M. Benedettini, et al., THE FIRST Hi-GAL OBSERVATIONS OF THE OUTER GALAXY: A LOOK
635   AT STAR FORMATION IN THE THIRD GALACTIC QUADRANT IN THE LONGITUDE RANGE 216. 5   225. 5, The Astrophysical Journal 772 (1) (2013) 45.

[18] K. Marsh, A. Whitworth, O. Lomax, Temperature as a third dimension in column-density mapping of dusty astrophysical structures associated with
640   star formation, Monthly Notices of the Royal Astronomical Society 454 (4) (2015) 4282–4292.

[19] M. Taylor, STILTS-A Package for Command-Line Processing of Tabular Data, in: Astronomical Data Analysis Software and Systems XV, Vol. 351, 2006, p. 666.

645 [20] E. Schisano, K. Rygl, S. Molinari, G. Busquet, D. Elia, M. Pestalozzi, D. Polychroni, N. Billot, S. Carey, R. Paladini, et al., The identification of filaments on far-infrared and submillimiter images: Morphology, physical conditions and relation with star formation of filamentary structure, The Astrophysical Journal 791 (1) (2014) 27.

650 [21] U. Becciani, E. Sciacca, A. Costa, P. Massimino, C. Pistagna, S. Riggi, F. Vitello, C. Petta, M. Bandieramonte, M. Krokos, Science gateway technologies for the astrophysics community, Concurrency and Computation: Practice and Experience 27 (2) (2015) 306–327.

[22] E. Sciacca, M. Bandieramonte, U. Becciani, A. Costa, M. Krokos, P. Mas-
655   simino, C. Petta, C. Pistagna, S. Riggi, F. Vitello, VisIVO Science Gateway: a Collaborative Environment for the Astrophysics Community, in: 5th International Workshop on Science Gateways, IWSG 2013, CEUR Workshop Proceedings, 2013.

26

[23] A. Costa, P. Massimino, M. Bandieramonte, U. Becciani, M. Krokos, C. Pistagna, S. Riggi, E. Sciacca, F. Vitello, An Innovative Science Gateway for the Cherenkov Telescope Array, Journal of Grid Computing 13 (4) (2015) 547–559.

[24] S. Sanchez Exposito, P. Martin, J. E. Ruiz, L. Verdes-Montenegro, J. Garrido, R. S. Pardell, A. Ruiz Falco, R. Badia, Web Services as Building Blocks for Science Gateways in Astrophysics, in: Science Gateways (IWSG), 2015 7th International Workshop on, IEEE, 2015, pp. 80–84.

[25] K. Wolstencroft, R. Haines, D. Fellows, A. Williams, D. Withers, S. Owen, S. Soiland-Reyes, I. Dunlop, A. Nenadic, P. Fisher, et al., The Taverna workflow suite: designing and executing workflows of Web Services on the desktop, web or in the cloud, Nucleic acids research (2013) gkt328.

[26] J. Ruiz, J. Garrido, J. Santander-Vela, S. Sánchez-Expósito, L. Verdes-Montenegro, AstroTaverna-Building workflows with Virtual Observatory services, Astronomy and Computing 7 (2014) 3–11.

[27] M. E. Pierce, S. Marru, L. Gunathilake, D. K. Wijeratne, R. Singh, C. Wimalasena, S. Ratnayaka, S. Pamidighantam, Apache Airavata: design and directions of a science gateway framework, Concurrency and Computation: Practice and Experience 27 (16) (2015) 4282–4291.

[28] B. Erickson, R. Singh, A. E. Evrard, M. R. Becker, M. T. Busha, A. V. Kravtsov, S. Marru, M. Pierce, R. H. Wechsler, Enabling dark energy survey science analysis with simulations on xsede resources, in: Proceedings of the Conference on Extreme Science and Engineering Discovery Environment: Gateway to Discovery, ACM, 2013, p. 16.

[29] B. Ludäscher, I. Altintas, C. Berkley, D. Higgins, E. Jaeger, M. Jones, E. A. Lee, J. Tao, Y. Zhao, Scientific workflow management and the Kepler system, Concurrency and Computation: Practice and Experience 18 (10) (2006) 1039–1065.

[30] W. Freudling, M. Romaniello, D. Bramich, P. Ballester, V. Forchi, C. García-Dabló, S. Moehler, M. Neeser, Automated data reduction workflows for astronomy-The ESO Reflex environment, Astronomy & Astrophysics 559 (2013) A96.

[31] M. L. Massie, B. N. Chun, D. E. Culler, The ganglia distributed monitoring system: design, implementation, and experience, Parallel Computing 30 (7) (2004) 817–840.

690

28