

PDACS - A Portal for Data Analysis Services for Cosmological Simulations

Ryan Chard
Victoria University of
Wellington
Wellington, New Zealand
ryan@ecs.vuw.ac.nz

Saba Sehrish
Fermilab
Batavia, IL 60510
ssehrish@fnal.gov

Alex Rodriguez
The University of Chicago
Chicago, IL 60637
arodri7@uchicago.edu

Ravi Madduri
Argonne National Laboratory
Lemont, IL 60435
madduri@anl.gov

Thomas D. Uram
Argonne National Laboratory
Lemont, IL 60435
turam@anl.gov

Marc Paterno
Fermilab
Batavia, IL 60510
paterno@fnal.gov

Katrin Heitmann
Argonne National Laboratory
Lemont, IL 60435
heitmann@anl.gov

Shreyas Cholia
NERSC, LBNL
Berkeley, CA 94720
scholia@lbl.gov

Jim Kowalkowski
Fermilab
Batavia, IL 60510
jbk@fnal.gov

Salman Habib
Argonne National Laboratory
Lemont, IL 60435
habib@anl.gov

ABSTRACT

Accessing and analyzing data from cosmological simulations is a major challenge due to the prohibitive size of cosmological datasets and the diversity of the associated large-scale analysis tasks. Analysis of the simulated models requires direct access to the datasets, considerable compute infrastructure, and storage capacity for the results. Resource limitations can become serious obstacles to performing research on the most advanced cosmological simulations. The Portal for Data Analysis services for Cosmological Simulations (PDACS) is a web-based workflow service and scientific gateway for cosmology. The PDACS platform provides access to shared repositories for datasets, analytical tools, cosmological workflows, and the infrastructure required to perform a wide variety of analyses. PDACS is a repurposed implementation of the Galaxy workflow engine and supports a rich collection of cosmology-specific datatypes and tools. The platform leverages high-performance computing infrastructure at the National Energy Research Scientific Computing Center (NERSC) and Argonne National Laboratory (ANL), enabling researchers to deploy computationally intensive workflows. In this paper we present PDACS and discuss the process and challenges of developing a research platform for cosmological research.

1. INTRODUCTION

Cosmology has made rapid strides in the last two decades based on the remarkable results from a number of large-scale sky surveys carried out across multiple wavebands. Investigations of the mysteries of dark energy and dark matter increasingly rely on large, high-resolution simulations of cosmological structure formation to provide the theoretical underpinnings of cosmological research. These advanced simulations require enormous amounts of compute resources and are generated on leadership computing facilities [3]. The datasets generated are characteristically large and can be hundreds of terabytes in size, extending into the petabytes. Due to the size of these datasets and the complexity of many cosmology simulations, on-demand regeneration of datasets is not feasible. When combined with the complexity of the associated analysis, the number of analyses that can be performed and therefore science that can potentially be achieved from the simulations is limited.

The Portal for Data Analysis Services for Cosmological Simulations (PDACS) project is designed to increase the value of cosmological simulations by improving the accessibility of the resulting datasets. The project aims to enrich the entire research community through shared repositories for tools, workflows, and datasets, and encourages researchers to contribute their own. PDACS is a workflow engine and scientific gateway for cosmology research. The platform exposes advanced cosmological models as well as the tools and computational resources required to analyze them. PDACS also provides the necessary storage services and manages the output of analyses for researchers.

The PDACS project is a collaboration between Argonne National Laboratory (ANL), Fermi National Accelerator Laboratory (FNAL), and the National Energy Research Scientific Computing Center (NERSC). The platform repurposes

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright 20XX ACM X-XXXXX-XX-X/XX/XX ...\$15.00.

Galaxy [4], a workflow engine for biomedical research, to create a research environment for cosmology. PDACS platforms are currently in operation at NERSC and ANL, providing a preliminary set of test users access to flexible, large-scale computational resources and storage services through NERSC’s HPC infrastructure and ANL’s Magellan cloud.

The main contributions of this work include the repurposing of Galaxy as a cosmology-specific workflow service and research gateway. We provide a scalable platform for cosmological analytical tools and facilitate parallel job submission over HPC infrastructure at both NERSC and ANL. A large set of cosmology datasets and tools are provided through the platform. New datatypes have been developed to enable the propagation of metadata through workflows and provide fine-grained provenance. We have also integrated interactive plotting capabilities for displaying and analyzing results.

In this paper we present PDACS and discuss the challenges associated with creating a research gateway for cosmology. We first discuss the role of Galaxy in Section 2. PDACS is then described in more detail in Section 3, including the role of tools, datatypes, and the infrastructure supporting the instances. Finally we conclude and present future work in Section 4.

2. GALAXY PREVIEW

The nature of biomedical research is becoming increasingly data-intensive, with a considerable amount of time and resources being required to analyze next-generation genome sequencing data. Driven by these considerations, Galaxy is a workflow engine specifically developed for large-scale, interactive, biomedical research. Galaxy supports the assembly of complex, multi-step workflows, and then automates their execution over an abstracted infrastructure. Galaxy is a growing platform for biomedical research as it provides a user friendly method to deploy and manage large-scale analyses of genome sequences. In addition to the Galaxy Team operating a Galaxy instance themselves, the software is open source and can be customized and operated independently. Globus Genomics [6] is an example of this which provides customized Galaxy instances that scale on-demand with user requirements by leveraging Amazon’s Web Services.

The Galaxy platform incorporates a number of commonly used tools and supports the data types necessary to analyze genomes. Users are encouraged to contribute their own applications to the platform by wrapping the application as an executable Galaxy tool. When wrapped, the input and output parameters and their respective data type are encoded to be recognized by Galaxy. The notion of a tool shed provides a mechanism for users to share their own tools with the Galaxy community, reducing the siloing effect that is often associated with research. Through the standardization of data types, these tools can be included in reproducible workflows and placed in the public domain for future use.

In order for Galaxy to be easily deployed over various infrastructures, job runners have been developed. A job runner provides an abstraction from the complexity of operating with the supporting infrastructure and enables the execution of workloads over different HPC platforms. Galaxy provides a set of standard job runners to interface with common schedulers, such as HTCondor.

3. PDACS FEATURES

PDACS provides the capability to use and share existing cosmology analysis tools and link them in a workflow with coherent metadata propagation and provenance tracking. New tools, datatypes, and formats have been included for workflows. Researchers are able to use and share existing cosmology simulation data and workflows, as well as upload and share new content. Interactive plotting is also supported within PDACS, enabling results to be quickly visualized and evaluated. While many of these capabilities were present in Galaxy to meet the needs of biomedical researchers, significant changes have been made to Galaxy to make it suitable for the cosmology community. In this section, we explain the requirements for PDACS, and describe how each has been addressed.

3.1 New Cosmology Tools

Galaxy includes a set of tools that are commonly used by the biomedical community. A *Galaxy tool* is either a user defined application or script that invokes a user application written in any language. For PDACS, we required a stripped down version of Galaxy without the tools, datatypes, and formats designed for biomedical research.

To establish a lightweight, cosmology-focused, Galaxy instance, the biomedical features of the platform have been minimized and replaced with cosmology-oriented alternatives. New tool integration is accomplished by adding a *wrapper* or *tool config* XML file that is used to describe the tool interface, tool invocation, and available options. To provide analysis applications to the cosmology community through the PDACS interface, we added wrappers and tools for each application to expose them to users.

PDACS provides a set of frequently used cosmology tools, allowing users to perform tasks such as finding clusters of particles, known as halos, within simulation models. The provided tools include: Friends-of-Friends (FOF) and Spherical Overdensity (SO) halo finders, halo profile measurements, concentration measurements, binning routines to calculate, e.g., mass functions from the halo finder outputs, a concentration-mass relation emulator, a power spectrum emulator, and measurement tools for two-point statistics, such as the fluctuation power spectrum and correlation function. Computationally intensive tools, such as the halo finder, are automatically submitted as a batch job to be run on cluster and supercomputing resources, while less intensive tools, such as the emulators, are executed on the same node that manages the Galaxy instance.

Users can add new tools to PDACS by supplying the application runnable from the command line and the required wrapper code. One of the goals of PDACS is to make the introduction of new tools easier. Users can either run these tools as standalone applications or assemble these tools into pipelines through a drag and drop interface, creating sharable and reproducible analytical workflows.

3.2 New Datatypes for Metadata Propagation

Galaxy can recognize data that are introduced into it. This is accomplished by a combination of what Galaxy calls the *datatype* of the file and the *format* of the file. If the file type is recognized by Galaxy, information (the *metadata*) describing the contents of the file can be readily extracted and stored in Galaxy’s metadata database. Such metadata may be viewed in the Galaxy web interface. The format

of the file is used to ensure that a tool consumes only files that the associated scientific application understands, and thus to help prevent assembly of incorrect workflows. The format of a file is assigned when Galaxy runs a tool that produces the file in question as output. The format is itself not a part of the file; it is part of the description of the tool that produces data, and is contained in the metadata Galaxy manages for each file.

The simple datatypes provided by the Galaxy distribution are insufficient for the intended uses of PDACS. They have inadequate support for fine-grained metadata, making the extraction of the application information and the tracking of provenance difficult. The original datatypes lack the support for the forwarding of metadata along a workflow; this is required for the provenance tracking necessary within PDACS. A much richer collection of metadata is necessary to ensure that 1) the user of a given data file can identify the application that created the file, 2) all the inputs that were used to generate the file, and 3) the set of configuration parameters used by the application that created it.

PDACS requires both more sophisticated datatypes and a complex hierarchy of formats to support cosmological workflows. In addition, since a core goal of PDACS is to simplify the introduction of new tools, PDACS also requires facilities to aid the developer in creating new datatypes and formats.

To efficiently capture and support different tools that can be assembled into pipelines we introduce a new datatype based on SQLite. SQLite is a self-contained database, where the entire state of the database is typically stored within a single file. For each tool, we capture metadata including the application that created the file, along with input and configuration parameters. The new datatype has been registered and integrated into Galaxy, allowing users to easily view metadata from a particular application.

Each tool generates one or more output files which typically consist of tabular data. These are recognized by Galaxy as a known datatype and can be viewed in the browser. Some tools generate binary data which can be then converted to ASCII format using a converter tool within PDACS.

3.3 Data Access - Upload and Share

The data-intensive nature of cosmological research requires maintaining and processing large volumes of simulation data. In order to enrich the value of cosmological simulations, PDACS provides access to the generated datasets and the infrastructure to analyze them. Currently, PDACS exposes the set of 37 cosmological models known as the Coyote Universe [5]. The metadata regarding the generation of each model and the parameters used during the simulation are captured in our SQLite datatype when the model is loaded for a tool. The metadata is closely bound to the input data as its datatype and is passed to subsequent tools, propagating it through a workflow.

Users are also capable of uploading and sharing their own datasets with the PDACS community. To provide this functionality we have integrated Globus Transfer [1] into the PDACS platform. Globus provides reliable, secure and efficient data-transfer for large datasets in an asynchronous manner. Transfers can be started from within the PDACS instance and uploaded to the supported storage services.

3.4 Job Submission

The PDACS service operates as a science gateway and




Figure 1: The process of using NERSC’s PDACS platform to analyze cosmological models.

can be easily deployed on different compute infrastructures. Currently, two distinct PDACS platforms are operational and utilize NERSC’s computing infrastructure and ANL’s scientific cloud, Magellan, respectively. These instances are currently available to a preliminary group of test users and will in future be open to a wider range of cosmology researchers. Job runners have been developed to transparently dispatch jobs over the HPC resources available at NERSC and Magellan. Authenticated users are able to run smaller jobs on the node serving PDACS, and large parallel jobs on the associated HPC resources that support the platform

In order to deploy jobs to NERSC’s computing resources, a RESTful API, called NEWT is used [2]. NEWT exposes a range of NERSC functionality and resources, providing the ability to authenticate users, monitor system status, and schedule workloads over each of the HPC resources maintained at NERSC. PDACS leverages NEWT for user authentication and dispatching jobs. Authenticating users with NERSC credentials enables jobs and workflows that are dispatched through PDACS to be associated with a NERSC user, allowing resource accountability and the enforcement of compute and data restrictions.

A NEWT runner has been developed to enable seamless integration of PDACS with the NERSC infrastructure. The runner manages job deployment and workload executions. The NEWT runner is transparent to users and allows tools to be included within PDACS without knowledge of the compute infrastructure and submission process.

Figure 1 depicts the complete process of utilizing PDACS to analyze cosmological models on NERSC resources. Users are authenticated through NEWT and gain access to a wide range of tools and datasets through the PDACS platform. Pipelines can be assembled through the PDACS interface and deployed to compute infrastructure. The resulting output files can then be processed by the powerful library of visualization tools that are exposed through PDACS.

The PDACS platform operating on the Magellan cloud authenticates users with single sign-on through Shibboleth [7] using Argonne credentials. Job executions are tied to a user’s Galaxy account to support accountability. Job submissions require instances to be provisioned on the cloud

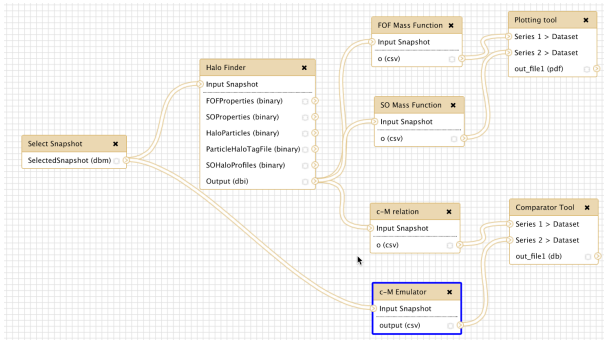


Figure 2: A PDACS workflow that analyzes the predictive performance of the concentration-mass relation emulator.

before they can be executed. Once a job is completed the instances are terminated and released to the available pool of resources. A persistent pool of running instances will be available for PDACS jobs on Magellan in the future.

3.5 Graphing/Plotting

Visualization is generally the final step of a workflow. A requirement for PDACS is the ability to plot multiple columns from the same, or multiple, datasets against one another. It is also beneficial for users to view all the column names and a subset of the data before plotting.

Figure 2 presents a commonly used workflow that is used to compare the results from the analysis of a large simulation (upper part of the workflow, carried out on the back-end of a computing cluster) with the prediction of a fast emulator tool (lower part). In addition, the workflow shows the comparison of two different mass function definitions (FOF for friends-of-friends and SO for overdensity). Both comparisons are visualized as part of the workflow (the comparator and plotting tools are used for this). Setting up the workflow as shown in PDACS ensures consistent use of input parameters for both simulation analysis and prediction tool. PDACS takes care of submitting the parallel analysis job to the supercomputer, synchronizes the inputs, and provides the results readily via its plotting routines. This eases the analysis task immensely for the user.

Multiple visualization tools have been integrated into the PDACS platform including a JavaScript-compiled Gnuplot, providing powerful graphical support to users. In addition to these, we have added a Shiny based application as an interactive visualization tool. Shiny is a web application framework for R [8]. Our Shiny server runs alongside the PDACS instance on the host node permitting immediate access to data. The server allows users to plot their tabular data and combine multiple columns from the same or multiple datasets within the same graph, while supporting the dynamic selection and modification of datasets.

4. CONCLUSIONS AND FUTURE WORK

PDACS provides a workflow gateway, exposing a large set of cosmological datasets and composable analysis tools that can operate on large datasets. PDACS workflows can be easily shared via a provenance capture mechanism. PDACS provides access to back-end clusters, high throughput computing and supercomputing resources, and flexible cloud re-

sources as well as to services for data transfer and storage.

PDACS aims to increase the value of large-scale cosmological simulations by providing researchers with access to the datasets and the computational infrastructure required to analyze them. PDACS is soon to be released to a small subgroup of expert users in order to further evaluate the platform and to guide future work. Our immediate direction involves incorporating a wider range of analysis tools, extending the capabilities of the workflow engine, leveraging Galaxy’s tool shed concept to improve collaboration, adding a new command-line interface, and further developing visualization capabilities. In future, we envision the addition of observational datasets with associated pipelines and methods for joint analysis of simulation and observational data.

5. ACKNOWLEDGMENTS

We thank Juliana Kwan and Suman Bhattacharya of the Cosmic Frontier Theory Group at Argonne National Laboratory for providing us with some of the cosmological tools used in PDACS. Initial funding for PDACS was provided by the DOE High Energy Physics Computing Program. We also acknowledge the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-06CH11357.

6. REFERENCES

- [1] B. Allen, J. Bresnahan, L. Childers, I. Foster, G. Kandaswamy, R. Kettimuthu, J. Kordas, M. Link, S. Martin, K. Pickett, et al. Software as a service for data scientists. *Communications of the ACM*, 55(2):81–88, 2012.
- [2] S. Cholia, D. Skinner, and J. Boverhof. NEWT: A RESTful service for building high performance computing web applications. In *Gateway Computing Environments Workshop (GCE)*, 2010, pages 1–11. IEEE, 2010.
- [3] A. Connolly, S. Habib, A. Szalay, J. Borrill, G. Fuller, N. Gnedin, K. Heitmann, D. Jacobs, D. Lamb, T. Mezzacappa, et al. Snowmass computing frontier: Computing for the cosmic frontier, astrophysics, and cosmology. *arXiv preprint arXiv:1311.2841*, 2013.
- [4] J. Goecks, A. Nekutenko, J. Taylor, et al. Galaxy: a comprehensive approach for supporting accessible, reproducible, and transparent computational research in the life sciences. *Genome Biol*, 11(8):R86, 2010.
- [5] K. Heitmann, D. Higdon, M. White, S. Habib, B. J. Williams, E. Lawrence, and C. Wagner. The coyote universe. ii. cosmological models and precision emulation of the nonlinear matter power spectrum. *The Astrophysical Journal*, 705(1):156, 2009.
- [6] R. K. Madduri, D. Sulakhe, L. Lacinski, B. Liu, A. Rodriguez, K. Chard, U. J. Dave, and I. T. Foster. Experiences building globus genomics: a next-generation sequencing analysis service using galaxy, globus, and amazon web services. *Concurrency and Computation: Practice and Experience*, 26(13):2266–2279, 2014.
- [7] R. Morgan, S. Cantor, S. Carmody, W. Hoehn, and K. Klingenstein. Federated security: The shibboleth approach. *Educause Quarterly*, 27(4):12–17, 2004.
- [8] R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2014.