

# A High Throughput Workflow Environment for Cosmological Simulations

Brandon M. S. Erickson  
Physics Department  
University of Michigan  
Ann Arbor, MI 48109  
bmse@umich.edu

Raminderjeet Singh  
Pervasive Technology Institute  
Indiana University  
Bloomington, IN 47408  
ramifnu@indiana.edu

August E. Evrard  
Physics Department  
University of Michigan  
Ann Arbor, MI 48109  
evrard@umich.edu

## ABSTRACT

The next generation of wide-area sky surveys offer the power to place extremely precise constraints on cosmological parameters and to test the source of cosmic acceleration. These observational programs will employ multiple techniques based on a variety of statistical signatures of galaxies and large-scale structure. These techniques have sources of systematic error that need to be understood at the percent-level in order to fully leverage the power of next-generation catalogs. Simulations of large-scale structure provide the means to characterize these uncertainties. We are using XSEDE resources to produce multiple synthetic sky surveys of galaxies and large-scale structure in support of science analysis for the Dark Energy Survey. In order to scale up our production to the level of fifty  $10^{10}$ -particle simulations, we are working to embed production control within the Apache Airavata workflow environment. We explain our methods and report how the workflow has reduced production time by 40% compared to manual management.

## Categories and Subject Descriptors

D.2 [Software Engineering]: Programming Environments—*Aravata*; D.2.11 [Software Architectures]: Domain-specific architectures

## General Terms

Science

## Keywords

Airavata, Astronomy, Astrophysics, Cosmology, Dark Energy, DES, OGCE, Scientific Workflows, XBay, XSEDE

## 1. INTRODUCTION

A decade, and a Nobel Prize,<sup>1</sup> after its discovery, the nature of cosmic acceleration remains a mystery. Evidence

<sup>1</sup> [http://www.nobelprize.org/nobel\\_prizes/physics/](http://www.nobelprize.org/nobel_prizes/physics/)

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.

XSEDE12, July 16 - 20 2012, Chicago, Illinois, USA

Copyright 2012 ACM 978-1-4503-1602-6/12/07 ...\$15.00.

continues to favor the simplest form of dark energy, so-called  $\Lambda$ CDM models with a constant vacuum energy density (or cosmological constant) having a fixed ratio of pressure to energy density (equation of state parameter),  $w = p/\rho = -1$  [11, 25]. Testing for departures from this canonical model requires large, sensitive astronomical surveys capable of delivering percent-level statistical constraints on  $w$  and the dark energy density,  $\Omega_{DE}$  [4]. Departures from  $\Lambda$ CDM expectations may signal a time-varying equation of state parameter anticipated by specific theoretical models [13] or may indicate that gravity departs from general relativity on large scales [20, 8].

Realizing the full statistical power of upcoming surveys requires addressing all potential sources of systematic error associated with applying tests of cosmic acceleration based on the large-scale distribution of galaxies and clusters of galaxies. We are performing a suite of simulations that will allow us to address a range of sources of systematic error for the upcoming Dark Energy Survey (DES).<sup>2</sup> The simulations can also be used to improve theoretical calibration of the clustered matter distribution, including the abundance and clustering of massive halos<sup>3</sup> that host galactic systems. The particular set of simulations we are performing on XSEDE in 2012 will form the basis of a Blind Cosmology Challenge for the DES collaboration.

## 2. COSMOLOGICAL SIMULATIONS FOR DES

The DES [24, 6, 5] is a Stage III<sup>4</sup> dark energy project jointly sponsored by DoE and NSF that is on track to see first light in the fall of 2012. The project will use a new panoramic camera on the Blanco 4-m telescope at the Cerro Tololo Inter-American Observatory in Chile to image  $\approx 5000$  square degrees of the sky in the South Galactic Cap in four optical bands, and to carry out repeat imaging over a smaller area to identify distant type Ia supernovae and measure their lightcurves. In addition, the main imaging area of the DES overlaps the South Pole Telescope<sup>5</sup> sub-mm survey that will identify galaxy clusters via the Sunyaev-Zel'dovich effect as well as the VISTA<sup>6</sup> infrared survey of galaxies, which will

<sup>2</sup> <http://www.darkenergysurvey.org>

<sup>3</sup>The term halos refers to self-bound, quasi-equilibrium structures that emerge via gravitational collapse of initial density peaks.

<sup>4</sup>In the language of the Dark Energy Task Force, see [4]

<sup>5</sup> <http://pole.uchicago.edu/>

<sup>6</sup> <http://www.vista.ac.uk>

provide additional information on galaxy photometric redshifts and on the properties of galaxy clusters at large cosmological redshift,  $z > 1$ .<sup>7</sup> Roughly three hundred scientists across nearly thirty institutions comprise the DES collaboration.

The DES will be the first project to combine four different methods to probe the properties of the dark sector (dark matter and dark energy) and test General Relativity gravity via evolution of the Hubble expansion parameter and the linear growth rate of structure. The methods—baryon acoustic oscillations in the matter power spectrum, the abundance and spatial distribution of galaxy groups and clusters, weak gravitational lensing by large-scale structure, and type Ia supernovae—are quasi-independent. Each has sources of systematic error associated with it, some of which are unique to the method but many of which are shared. Examples of the latter are the accuracy of photometric redshift estimates,<sup>8</sup> the form of the non-linear matter clustering power spectrum, and shape measurement errors for galaxy images that affect cosmic shear and galaxy cluster mass estimates. DES will thus be the first survey to address joint systematics in multiple methods probing accelerating expansion of the universe. N-body simulations provide key support for the analysis of systematics in the three methods associated with cosmic large-scale structure (all but supernovae in the above list). To validate science analysis codes, the DES Simulation Working Group is coordinating a Blind Cosmology Challenge (BCC) process, in which a variety of sky realizations in different cosmologies will be analyzed, in a blind manner, by DES science teams.

## 2.1 Blind Cosmology Challenge

The Blind Cosmology Challenge (BCC) process will require generating multiple galaxy catalogs to the full photometric depth across the full 5000 square degrees of the DES survey. The effort will require roughly 6M SUs and generate 300 TB of output.

Competing effects drive our simulation requirements. On one hand, the dark matter distribution needs to be modeled within a large cosmic volume. On the other hand, galaxy surveys also sample the nearby population of dwarf galaxies, implying a need for high spatial and mass resolution. We have developed an approach that generates a set of discrete N-body sky realizations of dark matter structure spanning a range in resolution and volume. We then dress this dark matter distribution with galaxies brighter than the DES limiting magnitudes in each passband.

To implement the BCC, we plan to produce fifty  $10^{10}$ -particle N-body runs on XSEDE resources over the next two years. The collaboration has requested that BCC models explore a variety of cosmologies, parameters of which are known only to the Simulation Working Group members. After a blind processing period, constraints on cosmological parameters from the science teams will be compared against their true values, gauging the validity of the processing pipelines. We detail ongoing and proposed simulations

<sup>7</sup>A redshift,  $z$ , derived from spectroscopy measures both distance and look-back time to the source. A galaxy at  $z = 1$  emitted its light when the universe was 6.2 Gyr old and it lies at a distance of 3.3 Giga-parsecs from the Milky Way.

<sup>8</sup>Photometric redshifts are distance estimates that use the multi-color fluxes measured in broad optical-IR bands as essentially a (very) low-resolution galaxy spectrum.

along with our strategy for producing synthetic sky surveys of sufficient area and depth in the next section.

## 2.2 Computational requirements

The dark matter structure from N-body simulations of a given cosmology forms the basis for galaxy catalog expectations. The N-body models we run on XSEDE resources store particle configurations,  $\{x_i(t), v_i(t)\}$ , in two forms: *snapshots* of the positions and velocities of all particles  $i$  in the simulation volume at a fixed time  $t$ , and *lightcones* [12] that hold kinematic information for particles lying on the past lightcone of a virtual observer located at a fixed position,  $x_{\text{cen}}$ , in the computational volume.<sup>9</sup> Once the N-body steps are completed, we proceed with additional processing using a combination of resources, including XSEDE, SLAC and other collaboration institutions. The post-processing aims to create a synthetic DES catalog of galaxy properties derived from the N-body lightcone output. Figure 1 shows a schematic representation of the post-processing workflow.

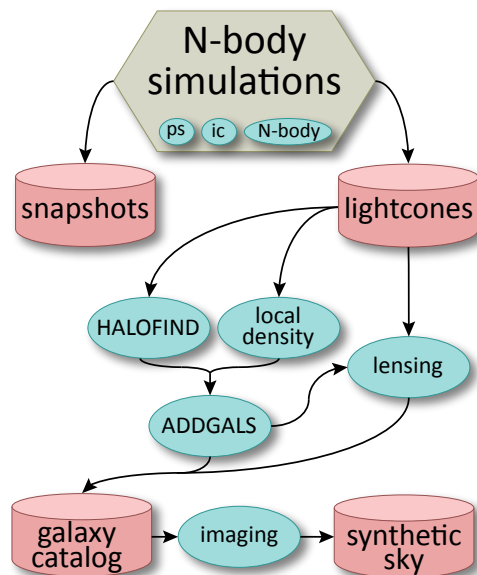


Figure 1: Processing steps to build a synthetic galaxy catalog are illustrated here and described in the text. The XBayo workflow currently controls the top-most element (N-body simulations) which consists of methods to sample a cosmological power spectrum (ps), generating an initial set of particles (ic) and evolving the particles forward in time with Gadget (N-body). The remaining methods are run manually on distributed resources.

Halos, bound systems that host galaxies and clusters of galaxies, are identified in these outputs and their properties can be used to determine their central galaxy characteristics. Halos, as well as local density estimates in under-resolved locations, are used by the ADDGALS<sup>10</sup> algorithm to assign galaxy properties to suitably selected dark matter

<sup>9</sup>These particles satisfy the combined space-time requirement,  $|x_i(t_i) - x_{\text{cen}}| = r_{\text{cosm}}(t_i)$ , where  $r_{\text{cosm}}(t)$  is the cosmic metric distance as a function of time, a known function of the cosmological parameters.

<sup>10</sup>for Adding Density-Determined GALaxies to Lightcone Simulations

particles. The matter along the past lightcone also sets the gravitational lensing shear signal applied to these galaxies, and we are developing a new multi-grid, spherical harmonic algorithm for this computation. Finally, galaxy catalogs that include lensing shear signals are processed with telescope/instrument/noise effects to produce images expected from the back-end electronics of the Dark Energy Camera. From these images, which are generated for only 200 sq deg of sky, we can develop effective transfer functions to create full 5000 sq deg synthetic DES catalogs that contain realistic errors such as star/galaxy mis-classification, blended sources, and appropriate photometric errors.

Generating a cosmological N-body simulation has three main steps, described further below. The first two create an initial particle set consistent with the structure expected at an early time of a chosen cosmological model. The first step samples the linear perturbation spectrum of the cosmological model while the second step realizes that density and velocity field with a set of particles. For the latter, we use a second-order Lagrangian perturbation theory code (2LPTic) code that has been robustly tested in the community [10, 9]. The amount of CPU time required to make these initial conditions is small (typically 300-400 CPU hours for simulations of the scale described here), but the memory requirements are more substantial (slightly less than 2 GB per core), although easily achievable on XSEDE machines.

The final step evolves this particle distribution under its self-gravity within an expanding cosmological background. For this purpose, we are using a lean version of the Gadget cosmological N-body code [22, 23] modified by us to generate output along the past lightcone of synthetic observers in the computational volume. We have worked on XSEDE resources for the past year to modify and optimize the code. The lean version has significantly reduced memory overhead, 44 bytes per particle compared to 84 for standard Gadget. This reduction allows the simulations to fit on a smaller number of processors, affording better scaling. These simulations require 50–100k CPU hours depending on the parameters and resolution of the computation, and generate up to 10 TB of total data output.

### 3. WORKFLOW ABSTRACTIONS

The simulation codes discussed in Section 2 are executed on large scale XSEDE resources managed by batch resource managers. The heterogeneity and complexity in interfacing with these resources slow down the computational scientists in harnessing the vast amount of available computing power. The eScience workflow systems abstract out these complexities and enable the use of innovations made in computational middleware. Scientific workflows are one of the prominent abstractions that allow scientist the carry out their scientific discovery and experimentation without having to worry about the underlying complexity. These abstractions, while lowering the entry and learning curves, also become more relevant to address human inefficiency to monitor long running jobs.

To build our cosmological workflow, we leverage the experience and software developed by the Open Gateways Computing Environments project [18] facilitated by the XSEDE Extended Collaborative Support Services. The workflow infrastructure is based upon the Apache Airavata [16] framework. We are going to briefly describe the integration of the simulation codes with the workflow infrastructure and spe-

cific customizations made to the framework itself. Further details about the framework and its comparison to other workflow solutions are discussed in [16].

The Airavata workflow system is primarily targeted to support long running scientific applications on computational resources. Airavata’s XBay is a graphical workflow tool, allows composition, execution and monitoring of the workflows. The Airavata workflow engine requires these applications to be raised to a common abstraction that can be accessed using a standard protocol. The Airavata Generic Application Factory (GFac) component bridges this gap between applications and the workflow systems by providing a network accessible web service interface to the scientific application.

### 3.1 Implementation

Once the simulation codes are deployed on XSEDE computational resources, we register descriptions of these applications with the Apache Airavata registry service. These descriptions are used by the Airavata GFac component to generate the artifacts required to expose the application as a service. The workflow developer can access these wrapped application services and construct workflows and orchestrate executions on target compute resources. The resulting workflow abstractions reduce human inefficiencies by providing a uniform interface for the scientist and hiding unnecessary complexities.

To illustrate the construction of a cosmological workflow, we will describe the process in developing the N-body simulation workflow of the process illustrated in Figure 1. Firstly the nature of the applications, its execution characteristics, and its input and output data are analyzed. The application meta information, including the executable location, its nature like serial or MPI, inputs and outputs, are described and registered with Airavata registry. This process was followed for the following four applications.

#### BCC Parameter Maker

This initial setup code is written as a PYTHON script and prepares necessary configurations and parameter files for the workflow execution. This simple script is forked on the XSEDE Ranger job management nodes.

#### CAMB

The CAMB (Code for Anisotropies in the Microwave Background) [14] application computes the power spectrum of dark matter, which is necessary for generating the simulation initial conditions. This application is as a serial FORTRAN code. The output files are relatively small ASCII files describing the power spectrum.

#### 2LPTic

The Second-order Lagrangian Perturbation Theory initial conditions code [9] (2LPTic) is programmed using Message Passing MPI C code that computes the initial conditions for the simulation from parameters and an input power spectrum generated by CAMB. The output of this application are a set of binary files that vary in size from ~80–250 GB depending on the simulation resolution.

#### LGadget

The LGadget simulation code is MPI based C code that uses a TreePM algorithm to evolve a gravitational

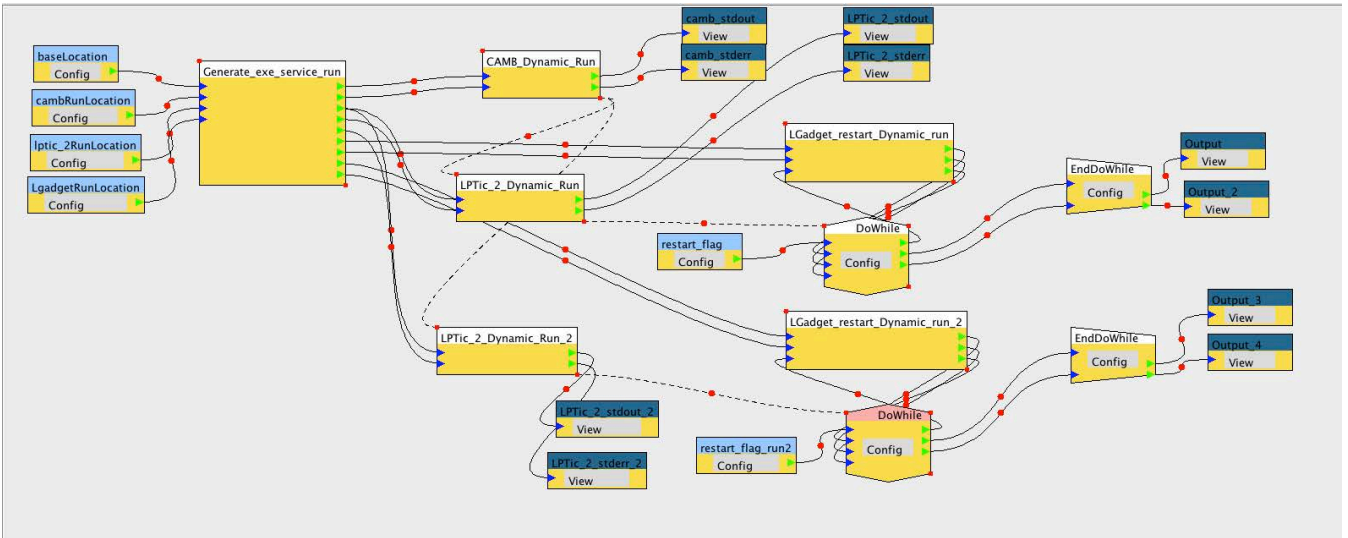


Figure 2: XBay workflow for a full cosmology with automatic queue resubmission if wallclock limits are reached and on-the-fly executable compilation.

N-body system [22, 23]. The outputs of this step are system state snapshot files, as well as lightcone files, and some properties of the matter distribution, including diagnostics such as total system energies and momenta. The total output from LGadget depends on resolution and the number of system snapshots stored, and approaches close to 10 TeraBytes for large DES simulation volumes.

After all the above applications are registered, the Blind Cosmology Challenge Workflow is constructed using Airavata XBay. The resultant workflow graph is shown in Figure 2. The workflow provides capabilities to configure the generation of initial conditions as well as the full N-body simulation components of the BCC process.

### 3.2 Workflow System Enhancements

**Iterative execution support for long running applications** The N-Body simulation requires multiple days of execution, but the XSEDE Ranger cluster limits maximum wall time of 48 hours. To mitigate this limitation, the workflow infrastructure has to allow iterative support so the job can be broken down into multiple increments of 48 hour jobs harnessing the check-point restart capabilities within the application. These capabilities required sophistication beyond the blind restarts, in order to account for application execution patterns and exception handling. These capabilities can be potentially matured into a formal Do-While construct semantics of workflow engines.

**Output Transfers** The workflow executions tend to produce terabytes of data residing on the cluster scratch file systems and have to be persisted for a longer durations. The data movement to archival systems like TACC Ranch for long term storage have to be provided. The large file data movement is non-trivial process. Even though many advancements have been made in this area, the seamless reliable data transfers are still challenging. The emerging solutions like Globus Online [1], GridFTP client API[2] and bbcp [3] are potentially viable options. We leverage the Ranch archival system mounting on Ranger and use bbcp

to copy the workflow outputs for test runs. We have yet to explore this as a solution for production executions. To transfer data to the post processing remote locations, Globus Online and GridFTP client are viable options.

## 4. RESULTS

Toward our goal of fifty cosmological simulations over two years, we have completed seven on XSEDE resources in the final quarter of 2011 and the first quarter of 2012. Most of these simulations were run ‘by hand,’ while the last was performed using the XBay workflow environment §3. We summarize the completed simulations in Table 1.

For a given cosmology, we generate four N-body simulations in nested volumes, consisting of three large-volume realizations with  $2048^3$  particles and one smaller volume of  $1400^3$  particles. This approach allows a better match to halo mass selection imposed by the magnitude-limited nature of the DES galaxy sample. As indicated in Table 1, the mass resolution varies by nearly a factor of 60 from our smallest to largest volumes. A halo resolved by a minimum of 100 particles ranges from a mass of  $3 \times 10^{12} h^{-1} M_{\odot}$  in the near-field simulation to  $2 \times 10^{14} h^{-1} M_{\odot}$  in the far-field.<sup>11</sup> The former is roughly the mass of our Milky Way galaxy’s halo while the latter corresponds to the mass scale of clusters of galaxies.

Each simulation produces lightcone outputs centered on each of the eight corners of the computational volume. By employing the periodic boundary conditions of the computational domain, we can stitch these octants of sky into a single  $4\pi$  representation of the full past lightcone of a hypothetical observer placed at the origin of the simulation. A map of the resultant structure in a thin radial slice of synthetic sky is shown in Figure 3. Along with these lightcone outputs, we also record snapshots of the particle configuration in the full volume at 20 epochs, leading to an overall data output of 8.4 TB for the  $2048^3$ -particle runs.

<sup>11</sup>Here,  $h$  denotes Hubble’s constant,  $H_0$ , in dimensionless form,  $h = H_0/100 \text{ km s}^{-1} \text{ Mpc}^{-1}$ .






Figure 3: Full-sky image of the dark matter density in a thin radial slice (50–75  $h^{-1}$  Mpc distance from observer) taken from our 2600  $h^{-1}$  Mpc  $\Lambda$ CDM simulation. Color maps the local matter density relative to the mean value on a logarithmic scale ranging from 0.2 (blue) to roughly 500 (red).

Table 1: Summary of completed simulations, giving the side length  $L$  (in  $h^{-1}$  Gpc) of the periodic cube, number of particles, particle mass (in  $10^{12} h^{-1} M_{\odot}$ ), number of simulations run, kiloSUs used for the run as well as any completed postprocessing and the amount of data generated in Terabytes.

$L$	$N_{\text{part}}$	$M_{\text{part}}$	$N_{\text{run}}$	kSU	Data (TB)
1.05	1400 <sup>3</sup>	0.027	2	121	5.4
2.60	2048 <sup>3</sup>	0.131	2	284	16.8
4.00	2048 <sup>3</sup>	0.476	2	149	16.8
6.00	2048 <sup>3</sup>	1.650	1	95	8.4
All	–	–	7	649	47.4

The combined lightcone files for a single cosmology,  $\sim 2$  TB of data, are transferred manually via Globus Online to SLAC for the post-processing steps, illustrated in Figure 1, that create DES galaxy catalogs. We first identify dark matter halos using a new algorithm, dubbed ROCKSTAR [7], that uses a direct socket-to-socket task-scheduling approach to operate efficiently on sub-regions of large simulations. We determine a local Lagrangian density by computing the distance to the  $N$ th nearest particle, where  $N$  is chosen to enclose a mass of  $\sim 10^{13} M_{\odot}$ , roughly the transition mass above which halos host more than one bright galaxy. With the dark matter halos and Lagrangian density estimate in place, the ADDGALS algorithm creates a synthetic galaxy catalog for science analysis. Finally, gravitational lensing shear is computed from the lightcone matter distribution, and its effect on galaxy images recorded. A single post-processed galaxy catalog, with 102 parameters per galaxy, is an 0.5 TB dataset.

Table 2: Comparison of Manual and Workflow-enabled production times

Run	Total Time	CPU Time	Efficiency
Manual	8:15:33:05	4:07:24:10	50.0%
Manual	4:05:39:07	2:17:50:06	64.8%
Workflow	2:09:53:23	2:05:28:09	92.4%

#### 4.1 Efficiency gains with XBaye

The XBaye workflow shown in Figure 2 was tested and refined using smaller simulations over the period Oct 2011 to Mar 2012. We transitioned to production use for our most recent simulations. One workflow-managed simulation has run to completion, while a second job crashed because of a hardware problem on TACC Ranger.

Even with this limited information, we can compare the efficiency of running the required jobs under XBaye to those previously run manually. Jobs were submitted to the long queue at TACC Ranger, which has maximum resource limits of 1024 processors and 48 hour runtime. In Table 2, Total Time is the wallclock time interval for the entire production process, while CPU Time gives the sum of the run times of the required jobs. Times reflect the full  $N$ -body production process, from generating initial conditions all the way through to completing of the final  $N$ -body timestep. Efficiency is the ratio of CPU to Total times, with 100% representing the ideal scenario of running without interruption.

The first two rows of In Table 2 list different manually-processed simulations. The first is a large simulation that needed four total submissions to the queue: one for initial conditions, and three for  $N$ -body computation. The second row was a smaller job requiring three submission, one for ini-

tial conditions and two for N-body. These runs are relatively inefficient because, each time the wallclock limit is reached, the user is notified via email, and then must log back into the cluster to submit the next job. If the wallclock limit is reached at an inconvenient time, considerable time can elapse before the next submission. More submissions tends to drive up the inefficiency.

The last row shows the efficiency of running the full production process via the XBay workflow environment. By enabling immediate submission of jobs when the preceding job finishes, the efficiency improves to 92%, well above the 50–60% found for manual processing. Under the workflow, the only time spent not computing is spent waiting in the queue time on the cluster. That is, our job production becomes limited only by the instantaneous compute resources available on TACC Ranger.

We found that the workflow can also help prevent errors in simulation set-up. Our first production level workflow was designed to have the same parameters as a pair of manually completed simulations. We soon found that the latter, ‘by hand’ simulations were inconsistent with the workflow simulations. Investigating the source of the inconsistency, we found that the workflow was correct, and that a parameter had been mistakenly set to a wrong value in the original simulations. This shows the workflow’s value in reducing the risk of simple human run-time errors.

## 4.2 DES science

Synthetic surveys from the first set of simulations are currently in use by DES science groups. We provide here two examples of analyses from cluster and weak lensing science groups.

Members of the Galaxy Cluster Working Group are using the synthetic galaxy catalogs to evaluate different methods for identifying the massive halos that host clusters. Because 5-band optical photometry provides relatively crude distance information for each single galaxy, the ability to identify localized spatial clusters is compromised by poor depth resolution. Different cluster finding algorithms have methods to mitigate this loss of information, but none has been applied to a large survey with the depth of DES.

Model galaxies are assigned to specific, unique dark matter halos, so an ideal cluster finder would return the list of original halos with all members intact. In practice, the distance errors, incorrect choice of cluster location and other effects imply that the set of clusters found does not perfectly match the input set of halos. However, a correspondence map between cluster and halo sets can be derived using joint galaxy membership criteria. The utility of a cluster finder can then be characterized by two parameters: i) purity, the fraction of clusters that correspond to genuinely massive halos, and; ii) completeness, the fraction of halos that are found by the cluster finder. The ideal value for both parameters is one.

Figure 4 shows recent purity and completeness measurements from the redMaPPer cluster finder [19]. The redMaPPer cluster finder works on so-called red-sequence galaxies, galaxies with evolved stellar populations that tend to be found in massive halos. The algorithm is a Bayesian method that assigns galaxies a cluster membership probability based on the galaxy’s color, and the density of nearby galaxies of a similar color. Clusters are identified as clumps of similar high-likelihood member galaxies. Figure 4 shows that





Figure 4: Purity as a function of the observed mass proxy (top) and completeness as a function of true halo mass (bottom) of the redMaPPer galaxy cluster finding algorithm applied to the first synthetic DES sky catalog. Colors denote the redshift intervals shown in the legend.

redMaPPer is  $> 80\%$  pure and has similarly high completeness above a redshift-dependent minimum halo mass.

In a related project, members of the Weak Lensing Working Group have been investigating how best to estimate the mass of an observed cluster based on the weak lensing distortion of background galaxies. Light from faint background galaxies is bent by the gravitational field of a massive halo, distorting the shape (shear) and density (magnification) of the galaxies. Using a maximum likelihood estimator [21] to fit the synthetic observations, and comparing to the true halo masses from the simulation, the group can understand how best to estimate masses and also provide feedback on the model used to generate the synthetic observations. A recent calibration plot comparing estimated and true masses is shown in Figure 5.

## 4.3 Future Directions

Implementing the Airavata workflow for this project has




Figure 5: Comparison of the mass inferred from weak lensing shear analysis ( $y$ -axis) to the true halo mass ( $x$ -axis) for several thousand galaxy clusters identified in the first synthetic DES sky survey. The dotted line is the identity relation.

entailed some overhead. Scripts that set up the input parameter files needed to be developed, and new features were added to the existing codebase so that our applications would integrate more effectively with the workflow framework. Interaction between the co-authors of this document—domain scientists along with TeraGrid AUSS team members—was essential to achieving a production-level service. The effort invested has been worthwhile, with a significant gain in realized efficiency (Table 2) for our first production run.

In the near term, we have applied for continued XSEDE ECS support aimed at integrating the postprocessing and catalog production stages into the Airavata workflow (see Figure 1). We also want to integrate data movement into the workflow. In addition, we have requested time at two XSEDE facilities (TACC Ranger and UCSD Trestles), and would like to expand the workflow logic to choose execution location in real time, based on queue loads.

We also would like to use Airavata to improve our provenance practices. The workflow system can capture provenance, including information such as when the data set was created, by whom, where, and with what application version and which input parameters. Improved provenance can enable broader forms of sharing, reuse, and long-term preservation of our simulations and the resultant galaxy catalogs. Additional development could include standardizing an API for simulation parameter input and output, so that other codes could be easily implemented in the workflow, such as  $N$ -body models that employ modified gravity [17, 15].

In the longer term, we could also expand our scope, generalizing our galaxy catalog construction process into a science gateway that would support broader classes of astrophysical studies. The optical catalogs we create could be augmented by synthetic surveys at other wavelengths, from radio to X-ray, and our focus on galaxies could be expanded to include quasars, galactic stars, and other astrophysical objects.

## 5. SUMMARY

To meet the challenge of interpreting Big Astronomical Data for cosmological and astrophysical knowledge, new modes of study that incorporate model expectations derived from sophisticated simulations will be required. Growing demand for simulated data products motivates the automation of simulation production methods within grid-aware workflow environments. This work represents a first step in that direction.

We are using XSEDE resources to produce multiple synthetic sky surveys of galaxies and large-scale structure in support of science analysis for the Dark Energy Survey. To scale up our production to the level of fifty  $10^{10}$ -particle simulations, we have embedded production control within the Apache Airavata workflow environment, resulting in a significant increase in production efficiency compared to manual management.

## 6. ACKNOWLEDGMENTS

This work is partially supported by the Extreme Science and Engineering Discovery Environment (XSEDE) Extended Collaborative Support Service (ECSS) Program funded by the National Science Foundation through the award: OCI-1053575, and by NSF OCI Award 1032742, SDCI NMI Improvement: Open Gateway Computing Environments—Tools for Cyberinfrastructure-Enabled Science and Education. AEE acknowledges support from NSF AST-0708150 and NASA NNX07AN58G. We thank Jeeseon Song, Jörg Dietrich, Eli Rykoff, and Eduardo Rozo for providing assistance with Figures 4 and 5.

## 7. ADDITIONAL AUTHORS

Matthew R. Becker (Department of Physics, Kavli Institute for Cosmological Physics The University of Chicago, Chicago, IL 60637 email: [beckermr@uchicago.edu](mailto:beckermr@uchicago.edu)) and Michael T. Busha (Institute for Theoretical Physics, University of Zürich, Zürich, Switzerland and Physics Division, Lawrence Berkeley National Laboratory, Berkeley, CA 94720 email: [mbusha@physik.uzh.ch](mailto:mbusha@physik.uzh.ch)) and Andrey V. Kravtsov (Kavli Institute for Cosmological Physics, Department of Astronomy and Astrophysics, The University of Chicago, Chicago, IL 60637 email: [andrey@oddjob.uchicago.edu](mailto:andrey@oddjob.uchicago.edu)) and Suresh Marru (Pervasive Technology Institute, Indiana University, Bloomington, IN 47408 email: [smarru@cs.indiana.edu](mailto:smarru@cs.indiana.edu)) and Marlon Pierce (Pervasive Technology Institute, Indiana University, Bloomington, IN 47408 email: [mpierce@cs.indiana.edu](mailto:mpierce@cs.indiana.edu)) and Risa H. Wechsler (Kavli Institute for Particle Astrophysics and Cosmology, Physics Department, Stanford University, Stanford, CA 94305 email: [rwechsler@stanford.edu](mailto:rwechsler@stanford.edu))

## 8. REFERENCES

- [1] [online] Available from: <https://www.globusonline.org/>.
- [2] [online] Available from: <http://www.globus.org/toolkit/docs/latest-stable/gridftp/rn/#gridftpRN>.
- [3] [online] Available from: <http://www.slac.stanford.edu/~abh/bbcp/>.
- [4] A. Albrecht, G. Bernstein, R. Cahn, et al. Report of the Dark Energy Task Force. *ArXiv Astrophysics e-prints*, Sept. 2006. [arXiv:astro-ph/0609591](https://arxiv.org/abs/astro-ph/0609591).

- [5] J. Annis, S. Bridle, F. J. Castander, et al. Constraining Dark Energy with the Dark Energy Survey: Theoretical Challenges. *ArXiv Astrophysics e-prints*, Oct. 2005. [arXiv:astro-ph/0510195](#).
- [6] J. Annis, F. J. Castander, A. E. Evrard, et al. Dark Energy Studies: Challenges to Computational Cosmology. *ArXiv Astrophysics e-prints*, Oct. 2005. [arXiv:astro-ph/0510194](#).
- [7] P. S. Behroozi, R. H. Wechsler, and H.-Y. Wu. The Rockstar Phase-Space Temporal Halo Finder and the Velocity Offsets of Cluster Cores. *ArXiv e-prints*, Oct. 2011. [arXiv:1110.4372](#).
- [8] E. Bertschinger. Simulations of Structure Formation in the Universe. *Annu. Rev. Astron. Astrophys.*, 36:599–654, 1998. [doi:10.1146/annurev.astro.36.1.599](#).
- [9] M. Crocce, S. Pueblas, and R. Scoccimarro. Transients from initial conditions in cosmological simulations. *Mon. Not. R. Astron. Soc.*, 373:369–381, Nov. 2006. [arXiv:astro-ph/0606505](#), [doi:10.1111/j.1365-2966.2006.11040.x](#).
- [10] M. Crocce and R. Scoccimarro. Renormalized cosmological perturbation theory. *Phys. Rev. D*, 73(6):063519, Mar. 2006. [arXiv:astro-ph/0509418](#), [doi:10.1103/PhysRevD.73.063519](#).
- [11] J. Dunkley, E. Komatsu, M. R. Nolta, et al. Five-Year Wilkinson Microwave Anisotropy Probe Observations: Likelihoods and Parameters from the WMAP Data. *Astrophys. J. Supp.*, 180:306–329, Feb. 2009. [arXiv:0803.0586](#), [doi:10.1088/0067-0049/180/2/306](#).
- [12] A. E. Evrard, T. J. MacFarland, Couchman, et al. Galaxy Clusters in Hubble Volume Simulations: Cosmological Constraints from Sky Survey Populations. *Astrophys. J.*, 573:7–36, July 2002. [arXiv:astro-ph/0110246](#), [doi:10.1086/340551](#).
- [13] J. A. Frieman, M. S. Turner, and D. Huterer. Dark Energy and the Accelerating Universe. *Annu. Rev. Astron. Astrophys.*, 46:385–432, Sept. 2008. [arXiv:0803.0982](#), [doi:10.1146/annurev.astro.46.060407.145243](#).
- [14] A. Lewis, A. Challinor, and A. Lasenby. Efficient Computation of Cosmic Microwave Background Anisotropies in Closed Friedmann-Robertson-Walker Models. *Astrophys. J.*, 538:473–476, Aug. 2000. [arXiv:astro-ph/9911177](#), [doi:10.1086/309179](#).
- [15] B. Li, G.-B. Zhao, R. Teyssier, and K. Koyama. ECOSMOG: an Efficient Code for Simulating MOdified Gravity. *J. Cosmol. Astropart. Phys.*, 1:51, Jan. 2012. [arXiv:1110.1379](#), [doi:10.1088/1475-7516/2012/01/051](#).
- [16] S. Marru, L. Gunathilake, C. Herath, et al. Apache airavata: a framework for distributed applications and computational workflows. In *Proceedings of the 2011 ACM workshop on Gateway computing environments*, pages 21–28. ACM, 2011.
- [17] H. Oyaizu. Nonlinear evolution of  $f(R)$  cosmologies. I. Methodology. *Phys. Rev. D*, 78(12):123523, Dec. 2008. [arXiv:0807.2449](#), [doi:10.1103/PhysRevD.78.123523](#).
- [18] M. Pierce, S. Marru, R. Singh, A. Kulshrestha, and K. Muthuraman. Open grid computing environments: advanced gateway support activities. In *Proceedings of the 2010 TeraGrid Conference*, pages 16:1–16:9. ACM, 2010.
- [19] E. S. Rykoff, B. P. Koester, E. Rozo, et al. Robust Optical Richness Estimation with Reduced Scatter. *Astrophys. J.*, 746:178, Feb. 2012. [arXiv:1104.2089](#), [doi:10.1088/0004-637X/746/2/178](#).
- [20] F. Schmidt, A. Vikhlinin, and W. Hu. Cluster constraints on  $f(R)$  gravity. *Phys. Rev. D*, 80(8):083505, Oct. 2009. [arXiv:0908.2457](#), [doi:10.1103/PhysRevD.80.083505](#).
- [21] P. Schneider, L. King, and T. Erben. Cluster mass profiles from weak lensing: constraints from shear and magnification information. *Astron. Astrophys.*, 353:41–56, Jan. 2000. [arXiv:astro-ph/9907143](#).
- [22] V. Springel. The cosmological simulation code GADGET-2. *Mon. Not. R. Astron. Soc.*, 364:1105–1134, Dec. 2005. [arXiv:astro-ph/0505010](#), [doi:10.1111/j.1365-2966.2005.09655.x](#).
- [23] V. Springel, S. D. M. White, A. Jenkins, et al. Simulations of the formation, evolution and clustering of galaxies and quasars. *Nature*, 435:629–636, June 2005. [arXiv:astro-ph/0504097](#), [doi:10.1038/nature03597](#).
- [24] The Dark Energy Survey Collaboration. The Dark Energy Survey. *ArXiv Astrophysics e-prints*, Oct. 2005. [arXiv:astro-ph/0510346](#).
- [25] A. Vikhlinin, A. V. Kravtsov, R. A. Burenin, et al. Chandra Cluster Cosmology Project III: Cosmological Parameter Constraints. *Astrophys. J.*, 692:1060–1074, Feb. 2009. [arXiv:0812.2720](#), [doi:10.1088/0004-637X/692/2/1060](#).