FIGURE 2: An example of a section near the rich galaxy cluster RXJ2248 from a deep stacked co-add tile color image using the *g,r,i* filters and that was generated using the DESDM pipeline.

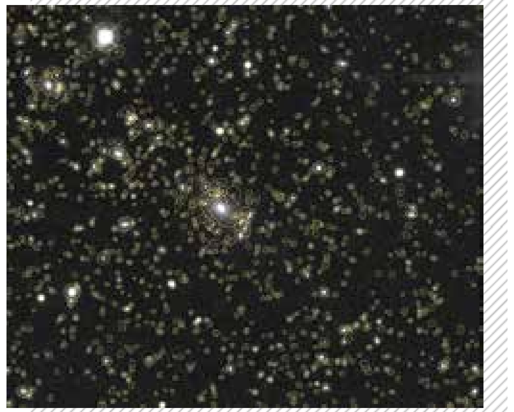

FIGURE 3: The same tile section showing ellipses for the objects detected and catalogued by the DESDM pipeline.

scale galaxy clustering (including baryon acoustic oscillations), and Type Ia supernovae. DES comprises two multi-band imaging surveys; a 5,000 square degree g,r,i,z,Y wide-survey of the southern sky to approximately 24th mag and a deeper time-domain 30 square-degree g, r, i, z deep DES Supernova (SN) Survey with a cadence of approximately five days.

DES uses the state-of-the-art three square-degree Dark Energy Camera (DECam), a 570 megapixel camera installed at the prime focus of the Blanco four meter telescope at the Cerro Tololo Interamerican Observatory (CTIO) in Northern Chile. DECam consists of 62 fully depleted, 250 micron thick 2048×4096 CCDs combined with four 2048×2048 guider and eight 2048x2048 autofocus CCDs.

For 525 nights, between 2013-2018, DES is scanning the sky to perform a 5,000 square-degree wide field survey. Over five observing seasons, DES will measure shapes, positions, fluxes, and colors for approximately 300 million galaxies and will discover and measure light curves for 3,500 supernovae and use these measurements to deliver powerful, new constraints on cosmic acceleration and dark energy. Each image arrives from CTIO in Chile to NCSA within minutes of being observed and it is usually processed by the nightly processing pipeline within 24 hours.

METHODS & RESULTS

One of main goals of the DESDM Project within the DES collaboration is the operation of the data reduction pipelines using high-performance computing (HPC) facilities to generate the survey data products. The DESDM data reduction process is composed of several pipelines or workflows that, starting from raw images, remove known instrumental signatures and sky background, masks defects, and detects and measures properties of objects to produce catalogs and calibrated images.

In 2014, we requested a Blue Waters allocation to explore the feasibility of running DESDM software and workflows on a shared Track-1 system. After nearly a year of investigations, where several software and network improvements were made by the Blue Waters team to accommodate our workflows, we were able to successfully deploy our production framework in the fall of 2015.

After each observing DES season, all DECam exposures meeting the survey data-quality criteria are systematically reprocessed using the FINALCUT and COADD pipelines. During the first trimester of 2016, we needed more computing resources to process the single-epoch FINALCUT campaign for the year three annual release (Y3A1) and used the remainder of our allocation to process 10,814 DECam exposures on the XE Compute Nodes. This corresponds to 15% of the total data volume (over 70k exposures) that DESDM processed for the Y3A1 data release. In Figure 1, we show an example of a DECam exposure processed using Blue Waters. In Figures 2 and 3, we show a section of a deep stacked co-add DES tile near the center of the cluster RXJ2248.

The DESDM system relies on the HTCondor software (CHTC UW-Madison) to manage jobs within a directed acyclic graph-based workflow. Porting this to Blue Waters represents a new and novel use of this system. Lessons from our implementation could serve as a model for other astronomy projects and pipeline that are not currently using Blue Waters.

WHY BLUE WATERS

The yearly re-processing campaigns of DESDM impose large seasonal variations in the demand for computing resources. The petascale size of Blue Waters, coupled with the perfectly parallel nature of our pipelines, allows DESDM to elastically expand its working pool of compute nodes to accommodate the burst in demand arising from the year annual processing for the collaboration.

DESDM is led by NCSA, where all images are archived and served in the community. Data processing on Blue Waters at NCSA is more robust and preformat than distributing workloads to remote sites. Moreover, the proximity between DESDM scientists and Blue Waters staff enables rapid feedback and clear communication, which are important for the success of complex implementations.

NEXT GENERATION WORK

We would like to continue production of our yearly annual releases on Blue Waters, which will continue to increase in size with each observing season. Our last data release is scheduled for 2021.

We want to improve the level of automation of HTCondor-based workflow on Blue Waters for the subsequent production campaigns.

We are also interested in the use of User Defined Images (UDI) for standard deployment of our software stacks to HPC systems.

AB INITIO MODELS OF SOLAR ACTIVITY

Allocation: NSF PRAC/9.01 Mnh
PI: Robert Stein[1]
Co-PIs: Mats Carlsson[2], Aake Nordlund[3], William Abbett[4], and Bart De Pontieu[5]
Collaborators: Viggo Hansteen[2]

[1]Michigan State University
[2]Oslo University
[3]Copenhagen University
[4]University of California Berkley
[5]Lockheed-Martin Space Sciences Lab

EXECUTIVE SUMMARY

The goal of this project is to understand how solar magneto-convection forms and controls solar active regions–how magnetic flux emerges through the solar surface and how that contributes to the heating of the chromosphere and corona and the acceleration of charged particles into the interplanetary medium.

INTRODUCTION

This project is motivated by society's vulnerability to harmful space weather. X-ray bursts and the high-energy particles associated with intense solar activity can harm astronauts, disable satellites, and hamper terrestrial systems for communication, guidance, and power distribution. Earth's heliospheric environment is controlled by magnetic fields produced by a subsurface convective dynamo. Some of the fields produced emerge through the solar surface into its atmosphere. Convective motions move these field foot points around in the photosphere. This tangles them and causes reconnection higher in the sun's atmosphere, heating the chromosphere and corona. Reconnection of active region fields produce X-rays and accelerate charged particles to high energies and drives them from the Sun into the heliosphere.

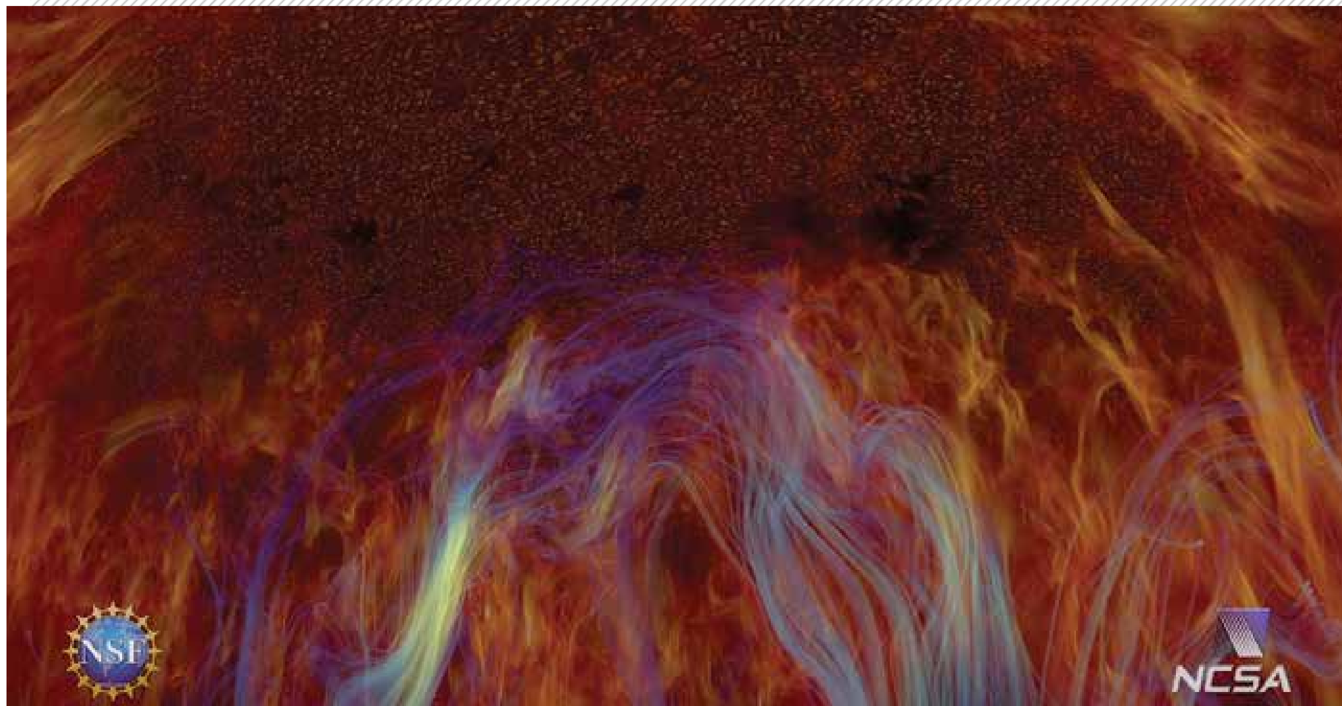

FIGURE 1: Image based on data from STAGGER calculations, from the planetarium show "Solar Superstorms" produced at NCSA showing the emergence of magnetic field lines through the solar surface.

METHODS & RESULTS

A finite-difference code is used to numerically solve the conservation equations for mass, momentum, internal energy. The induction equation for the magnetic field and the equation for non-gray radiation transport are used to model the magneto-convection and the behavior of the overlying chromosphere and corona.

The bulk of this project year was spent finding and removing an incompatibility introduced between the interior calculation scheme and the bottom boundary conditions. It was finally possible to start the relaxation of a 192 Megameter (Mm)—one million meter—by 20 Mm deep model of the outer solar convection zone using a grid of 4,032 by 4,032 by 500. This large size make it possible to use data from interior dynamo simulations to control the bottom boundary conditions for the study of active region emergence.

WHY BLUE WATERS

Blue Waters was particularly well suited for this task. Although it is up to four times slower on a per core basis than the newer Pleiades supercomputer, it was possible to run efficiently on 2,016 nodes or 64,512 integer cores on Blue Waters, whereas only 4,032 cores can be used on Pleiades because of excessive queue wait times for larger jobs.

NEXT GENERATION WORK

A new code, DISPATCH, is being developed which uses task based scheduling of small Cartesian patches that evolve with local time stepping. This type of code is ideal for obtaining **unprecedented** performance and scaling the current petascale and the near-future exascale computing eras, and with its large number of nodes and emphasis on a small number of large-scale projects, the Blue Waters system is particularly well suited for DISPATCH.

FIGURE 1: A large galaxy from our cosmological simulations is modeled with different active galactic nucleus (AGN) feedback and dynamics. The morphology of the galaxy is significantly different depending on whether we use (from left to right) our new AGN model, no AGN feedback, poor black hole dynamics or a less sophisticated accretion model.

EVOLUTION OF THE SMALL GALAXY POPULATION FROM HIGH REDSHIFT TO THE PRESENT

Allocation: NSF PRAC/11.9 mnh
PI: Thomas Quinn[1]
Co-PI: Fabio Governato[1]
Collaborators: Lauren Anderson[1], Michael Tremmel[1], Sanjay Kale[2], and Harshitha Menon[2]

[1]University of Washington
[2]University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Creating robust models of the formation and the evolution of galaxies requires the simulation of a cosmologically significant volume with sufficient resolution and subgrid physics to model individual star-forming regions within galaxies. This project aims to perform such modeling with the specific goal of interpreting Hubble Space Telescope observations of high redshift galaxies. We are using the highly scalable N-body/Smooth Particle Hydrodynamics code, ChaNGa, based on the Charm++ runtime system on Blue Waters to perform a simulation of a 25 Mpc cubed volume of the universe with a physically motivated star formation/supernovae feedback model. This past year's accomplishments include incorporating realistic black hole formation, dynamics, and feedback into our model which allows us to create and advance more realistic modeling of the knee in the galaxy luminosity function at high redshift. Comparisons with Hubble data show that we can reproduce the high redshift galaxy population with models that also reproduce the morphologies of present-day galaxies.

INTRODUCTION

The cold dark matter (CDM) paradigm for structure formation has had many successes over a large range of scales, from cosmic microwave background fluctuations on the scale of the horizon to the formation and clustering of individual galaxies. However, at the low end of the galaxy luminosity function, the CDM theory and observations are somewhat at odds. In particular, the existence of bulgeless, cored small galaxies is not a natural prediction of CDM. However, these are the scales where the baryonic physics of gas cooling, star formation, and feedback can significantly impact the overall mass of the galaxy. Furthermore, accurately modeling the star formation process requires a spatial resolution of order 100 parsecs to resolve the molecular star-forming regions of the interstellar medium. On the other hand, survey volumes addressing small galaxies, including recently approved Hubble Space Telescope (HST) programs, are over 10,000 cubic Mpc. Only with large simulations can we perform proper comparisons with these programs to address the following basic issues of the CDM model:

- Does the standard ΛCDM model produce the correct number densities of galaxies as a function of mass or luminosity?
- What role to these galaxies play in the evolution of the baryons in the Universe?
- How do these galaxies relate to the galaxies we can study in detail in the local universe?

METHODS & RESULTS

We used the highly scalable N-body/smooth particle hydrodynamics code ChaNGa to simulate the volumes surveyed by HST with sufficient resolution to make robust predictions of the luminosity function, star formation rate, and morphologies appropriate for these surveys. The results of the simulations were processed by our parallel data reduction pipeline that creates simulated observations. These results can be directly compared with results from observational programs.