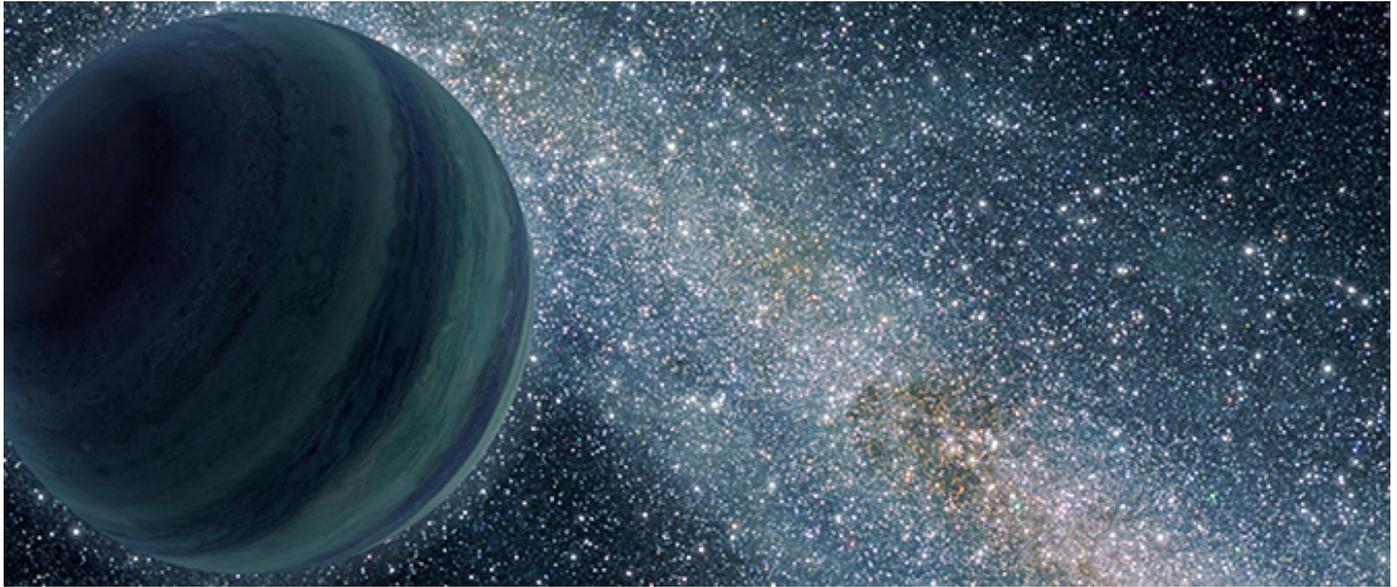


i We're gradually importing our scientific contents in French. Please forgive the visual glitches that may result from the operation.

To infinity and beyond! How HPC accelerates space exploration

By The Editorial Team | March 09, 2016



HPC's contribution to space exploration is invaluable since it allows simulation and visualization of millions, sometimes billions of data across a given frame of time. This allows achieving spectacular animated and dynamic results in a fraction of the time needed to calculate a static image only a decade before.

One of these extreme examples has been showed by the European Space Agency, ESA, showing the simulated interaction of solar winds with 67P/Churyumov-Gerasimenko, the famous comet targeted the Rosetta mission.

The simulated conditions represent those expected at 1.3 AU from the Sun, close to perihelion, where the comet is strongly active – a gas production rate of 5×10^{27} molecules/s is assumed here. The solar wind approaches from the left at ~ 400 km/s, carrying with it the embedded interplanetary magnetic field with a strength of about 5 nT. The material from the comet's nucleus forms an extensive envelope, the coma, several million km in size (not shown here). Part of the neutral gas molecules in the coma gets ionised by solar UV radiation or by charge exchange with the solar wind particles. These cometary ions are picked up by the approaching solar wind, a process known as mass loading, and cause it to slow down. In the model simulation enough ions are produced and picked up by the solar wind to slow it down from supersonic speed to sub sonic speed, causing a bow shock to form in front of the comet.

Simulating the Interior Dynamics of Stars and Giant Exoplanets

In the two decades since the first exoplanets were found in the mid-1990s, astronomers have discovered nearly 2,000 planets outside our solar system. Many of these are known as "hot Jupiters," planets that are similar in size to Jupiter but are much closer to their host stars, and therefore have faster orbits and much hotter surface temperatures.

To learn more about the interior dynamics of hot Jupiter exoplanets and their stars, astrophysicist Tamara Rogers and her team at the University of Arizona's Lunar and Planetary Laboratory ran a series of groundbreaking simulations on the Pleiades supercomputer, located at the NASA Advanced Supercomputing (NAS) facility at Ames Research Center.

"Modeling and simulation on high-performance computers are very effective tools for researching the dynamical processes that occur within stars and planets," said Rogers, now a lecturer at Newcastle University in the U.K. "Understanding these phenomena can help us learn how hot Jupiters formed and how they affect the evolution of planetary systems."

The team's simulations of hot Jupiters—which were the first to include magnetic fields—along with their massive star simulations, can help astronomers interpret data collected from space-based observatories like NASA's Kepler, Spitzer, and Hubble telescopes. For example, the team's findings may help explain some puzzling observations, such as why planets circling cool stars tend to have orbits that align with the star's spin direction while those around hot stars often have misaligned orbits; and why many hot Jupiters are bigger and less dense than expected given their mass, even accounting for their extreme temperatures. The simulation results also reveal how magnetic effects can influence winds on these planets, a finding that could provide a method for estimating the planets' magnetic fields based on observations of their atmospheres.

By studying hot Jupiters, so different from the gas giants that slowly circle our own Sun, astronomers are expanding their knowledge of planetary structure and evolution—research that is crucial to the search for rocky, Earth-like exoplanets that may support life.

Accelerating the Exoplanet Search : from Kepler to TESS

The first exoplanet orbiting another star like our sun was discovered in 1995. Exoplanets, especially small Earth-size worlds, belonged within the realm of science fiction just 21 years ago. Today, and thousands of discoveries later, astronomers are on the cusp of finding something people have dreamed about for thousands of years — another Earth.

The Kepler spacecraft has observed 200,000 stars in a patch of sky in the Cygnus and Lyra constellations over a period of four years. Using knowledge gained from Kepler discoveries on the distribution of planet populations, scientists can estimate the occurrence rates of habitable Earth-sized planets within our galaxy. As the Kepler primary mission draws to an end, a new mission—the Transiting Exoplanet Survey Satellite (TESS)—is just beginning to take shape. The purpose of this follow-up mission will be to survey the entire sky looking for exoplanets that are good candidates for performing detailed characterizations of the planets and their atmospheres.

The telescope on board the Kepler spacecraft has a 1.1-meter mirror and a 95-megapixel camera. Pixel brightness, which is used to detect planet size, is accumulated over 30-minute periods. The spacecraft is located in an Earth-trailing orbit—instead of orbiting Earth itself, the observatory trails behind Earth as it orbits the Sun. Its remote location does not permit downloading all the pixels at once, so the Kepler team selects a number of target stars representing about 5% of the total pixels and downloads them to the Kepler Science Operations Center (SOC) at NASA's Ames Research Center. These pixels are then calibrated and combined to form light curves, and scientists search for a dimming in the brightness of the stars that indicates a transiting exoplanet.

Picking up where Kepler left off, the new TESS spacecraft, planned for launch in 2017, will be in a highly elliptical orbit looking at the brightness of more than 500,000 stars during a two-year mission. Every 27 days, TESS will observe a 96-degree swath of sky (a set of observing sectors), and after two years, TESS will have observed almost the entire sky. The goal of this mission is complementary to Kepler's: While Kepler captured a statistically diverse sample of planets, TESS will focus on planets that are good candidates for further characterization by the James Webb Space Telescope.

Kepler data has yielded a trove of tremendous discoveries in the field of exoplanet science: As of November 1, 2014, there are 989 confirmed or validated planets, 4,234 planet candidates, 2,165 eclipsing binary stars, and the first Earth-sized habitable zone planet, Kepler-186f. Now that SOC scientists have a sufficiently large sample size, they can perform statistics on this population of discoveries to estimate the occurrence rates of planets. Of particular interest are Earth-sized planets in the habitable zone. By estimating the occurrence rate of these planets it may one day be possible to estimate the number of life-bearing worlds in the universe.

Why HPC Matters

This year, for the first time, NASA was able to recalibrate all the pixels ever collected by Kepler, using the Pleiades supercomputer at the NASA Advanced Supercomputing facility at NASA Ames. This will be the cleanest dataset ever produced by the Kepler team as we incorporate the results of our best algorithms for removing noise from various sources (instruments, cosmic rays, stars). We will then search this new dataset for additional planet candidates. The Kepler science pipeline searches 200,000 stars observed by Kepler; during this search the Transiting Planet Search algorithm performs $1E11$ statistical tests. This is only possible with the high-performance computing available from the NAS facility. In addition, NAS visualization experts helped fine-tune the Kepler SOC software infrastructure and algorithms for use on Pleiades. The TESS mission will utilize NAS supercomputing capabilities for nearly all its computational needs to produce a new dataset every 27 days for use by the scientific community.

Simulating the Interior Rotation and Dynamics of Stars and Giant Planets

In order to understand the structure and evolution of stars and planets, and accurately interpret our observations, we must find ways to learn more about the dynamic processes that occur within them. Modeling and simulation on high-performance computers are very effective tools for researching these phenomena. We have developed a 3D, spherical magnetohydrodynamic (MHD) code that can simulate the interior of giant planets, the Sun, and other stars. These studies can help us interpret various astronomical observations and make predictions that can help guide future observations.

Using the MHD code on the Pleiades supercomputer at the NASA Advanced Supercomputer (NAS) facility, Nasa's Ames Research Center has recently carried out two groundbreaking simulation studies. They performed the first MHD simulations of a "hot Jupiter" planet—a Jupiter-sized planet close to its host star—that included variable magnetic diffusivity and self-consistently modeled Ohmic heating.

They also carried out 2D and 3D self-consistent simulations that coupled the convective and radiative regions in a massive star in order to simulate the star's convectively driven internal gravity waves (IGWs). The results showed that the waves could alter the surface rotation of these stars and can induce large-scale mean flows that vary in time and are potentially observable.

Results and Impact

Nasa's hot Jupiter simulations show dynamics that are related solely to the presence of a magnetic field. The magnetic effects could influence winds on these exoplanets, which may be observable through temperature maps. Our results may lead to a method of inferring the planets' magnetic fields from observations made using NASA's Spitzer and Hubble space telescopes—and other currently available observatories. In addition, our results show that Ohmic heating had been overestimated in earlier, simplified calculations—and therefore likely cannot explain the observed inflated radii of these planets, as previously thought.

The flows predicted by the simulations of massive stars—in effect, momentum transport by IGWs—may explain a host of outstanding mysteries, including chemical anomalies that have been observed in such stars, the mechanism that drives the intermittent accretion disks of Be-class stars, and the origin of macroturbulence in intermediate and massive stars. These flows may also explain the observed dichotomy of obliquity between hot-Jupiter planets orbiting hot and cool stars—in other words, why the orbits of these planets around cool stars tend to align with the star's spin direction, while those around hot stars tend to have misaligned orbits.

Why HPC Matters

These simulations required substantial computing resources. The relatively low-resolution hot Jupiter calculations required hundreds of thousands of processor hours on Pleiades; each single 2D massive star calculation also required hundreds of thousands of processor hours. The 3D massive star calculations each required approximately one million processor hours, depending on complexity. In order to analyze the simulation results Nasa's Research Center needed to store vast amounts of data—approximately 60 terabytes—on the Lou mass storage system at NAS. Without NASA's high-performance computing resources, these simulations and data analysis would not be possible.

Test-Driving the "Sun in a Box" with NASA's IRIS Solar Observatory

The Sun is the closest star to Earth, and it impacts us in a variety of ways, many of which are not well understood. The radiation from the Sun plays a role in Earth's climate, and its violent plasma storms directly impact the space environment around our planet, which can lead to significant disruptions to our technology-dependent society (GPS reception, power grid stability). To better understand how the Sun's atmosphere is shaped and heated, and how it impacts Earth, our team is simulating a small piece of the Sun in a computational box that could hold the Earth six times over, and comparing our results with high-resolution observations of the Sun from NASA's recently launched Interface Region Imaging Spectrograph (IRIS) Mission.

This project uses advanced multi-dimensional radiative magnetohydrodynamic simulations to take into account the complex physical processes that power the Sun's atmosphere. We perform detailed comparisons of our simulations with high-resolution observations of the low solar atmosphere from NASA's IRIS solar observatory. This combination of numerical modeling and observations is focused on revealing:

How the outer solar atmosphere is heated to much greater temperatures than those at the Sun's surface through the dissipation of magnetic field energy.

How the complex interactions of magnetic fields, hydrodynamics, and radiation fields shape the Sun's atmosphere.

By comparing the results of our simulations with images taken by the IRIS spacecraft, we have been able to explain some of the mission's puzzling findings. The simulations have helped reveal how small-scale magnetic fields impact the lower solar atmosphere (chromosphere) and how friction between ionized and neutral particles dissipates magnetic energy and helps heat the chromosphere. Understanding the properties of the chromosphere is important, since this part of the Sun is the main source of the ultra-violet radiation that impacts Earth's upper atmosphere.

Our simulations have provided critical insight into the physical mechanisms that drive some of the violent events observed by IRIS, as well as much needed advances in methods for interpreting the complex radiation data recorded by the solar observatory.

Why HPC Matters

Numerical modeling of the complex radiative transfer and physical processes of the Sun, combined with the enormous contrasts of density, temperature, and magnetic field within the solar atmosphere, demands powerful supercomputing resources. The simulations utilize the massively parallel capabilities of the Pleiades supercomputer at the NASA Advanced Supercomputing facility at NASA's Ames Research Center. Approximately 2 million processor-hours (48 days on 1,728 processors) were required to compute one hour of solar time.

As NASA's scientists learn more from the discrepancies between the observations and simulations, it allows working on new solar simulations with much higher spatial resolution to better resolve some of the dominant spatial scales in the solar atmosphere, and on larger computational volumes to capture the large-scale behavior of the Sun's magnetic field. This is expected to lead to better agreement with IRIS observations, which will in turn provide a better

understanding of the complex mechanisms at work in our closest star.

ChaNGA : Unprecedented Simulations of Galaxy Formation

Understanding the formation of galaxies like our own Milky Way, and the tiny dwarf galaxies around it, is key to furthering our understanding of how cosmic structures formed and the nature of dark matter and black holes. To follow the formation of even one galaxy over the lifetime of the universe requires an accurate physical model that includes many different processes which act on both large and small scales. A new code, ChaNGA, is being run on NASA high-performance computers to produce realistic galaxy simulations that capture gravity and gas hydrodynamics, and describe how stars form and die and how black holes evolve. The simulations resolve galaxy structures at unprecedented resolution—down to several hundred light years (about 100 parsecs). Simulation results are being used to interpret observations gathered by NASA missions, such as the Hubble Space Telescope, to further NASA's goal in astrophysics: "Discover how the universe works, explore how it began and evolved, and search for life on planets around other stars."

ChaNGA (Charm N-body GrAvity solver), was developed at the University of Washington and the University of Illinois to perform N-body plus hydrodynamics simulations. A unique load-balancing scheme, based on the CHARM runtime system, allows us to obtain good performance on massively parallel systems such as Pleiades supercomputer located at NASA's Ames Research Center. On Pleiades, we are running high-fidelity simulations of dozens of individual galaxies, spanning from the mass of the Milky Way down to those 1,000 times less massive, with force resolutions under 100 parsec (1 parsec = 3.26 light years). Examples of ongoing projects with these simulations include: quantifying the redistribution of matter in galaxies when supernova energy is deposited; exploring the growth of black holes and the impact of active galactic nuclei (AGN) on galaxy evolution; and determining whether the ultraviolet light from stars in galaxies can "escape" to re-ionize the universe.

The high-resolution simulations already produced have revolutionized scientists' view of galaxy formation. Scientists have discovered that when supernovae occur in the high-density regions where stars are born, their energy can be transferred to dark matter, pushing the dark matter out of the center of galaxies. This process cannot occur in lower-resolution simulations, and thus evaded detection for over a decade, despite other attempts to produce realistic galaxy simulations. These new results explain several long-standing observational challenges to Lambda Cold Dark Matter (CDM) galaxy formation theory, and open new paths of inquiry.

Why HPC Matters

Achieving the high-resolution simulations that have revolutionized our theory of galaxy formation requires billions of particles in a given simulation, and the high-density regions where stars form require small time steps. A single galaxy simulation can take up to 1 million processor hours. Our newly updated version of ChaNGA allows us to scale to hundreds of thousands of cores. Tests performed on Pleiades have produced science-ready galaxy simulations, and more state-of-the-art simulations will be run on Pleiades over the next year.

Europe's Leading Public-Private Partnership for Cloud

The Helix Nebula Initiative is a partnership between industry, space and science to establish a dynamic ecosystem, benefiting from open cloud services for the seamless integration of science into a business environment. Today, the partnership counts over 40 public and private partners and four of Europe's biggest research centres, CERN, EMBL, ESA and PIC, charting a course towards sustainable cloud services for the research communities – the Science Cloud.

CERN has recently published a paper which outlines the establishment of the European Open Science Cloud that will enable digital science by introducing IT as a Service to the public research sector in Europe. CERN use its own use case to highlight the benefits of this approach :

- Evaluating the use of cloud technologies for LHC data processing
- Transparent integration of cloud computing resources with ATLAS distributed computing software and services
- Evaluation of financial costs of processing, data transfer and data storage

Service Level Agreements and Governance model

As outlined by CERN, physics and computing lead to a huge challenge :

- Billions of events are delivered to the experiments from proton-proton and proton-lead collisions in the Run 1 period (2009-2013)
- Collisions every 50 ns = 20 MHz crossing rate
- ~35 interactions per crossing at peak luminosity
- ~1600 charged particles produced in every collision
- ~ 5PB/year/experiment

CERN's Worldwide LHC Computing Grid

The WLCG is an international collaboration to store, process and analyse data produced from the LHC. It integrates computer centres worldwide that provide computing and storage resources into a single infrastructure.

Several R&D initiatives started by the experiment collaborations to investigate and exploit cloud resources, in order to utilize private and public clouds as an extra computing resource and set up a mechanism to cope with peak loads on the Grid.

PanDa : ATLAS' workload management system

PanDA, short for the ATLAS' Production ANd Distributed Analysis (PanDA) system is an homogeneous processing system layered over heterogeneous resources used for job submission, and to pilot jobs for acquisition of processing resources while supporting both managed production and analysis.

Experiment workflow has been successfully tested with Monte Carlo jobs. The Geant4 based simulation of the particles propagation through the ATLAS detector is as follows :

- Long (~4h), very intensive CPU usage, low I/O usage.
- Input: MC generator 4 vector files
- Output: ~50 MB/file of 50 events

The concept of open science and a European open science cloud seems to be widely accepted, but how to implement it remains largely unknown and Helix Nebula and the European Science cloud are examples that can be used to define the future structure...

Big science teams up with big business : the Helix Nebula Initiative welcomes HNSciCloud – a European Cloud computing partnership

The HNSciCloud is a European pre-commercial procurement (PCP) project co-funded by the European Commission Horizon 2020 Work Programme, which kicked-off in January 2016. Driven by the Pre-Commercial Procurement (PCP) commitment of leading research organisations from across Europe, HNSciCloud creates a competitive marketplace of innovative cloud services serving scientific users from a wide range of domains.

The marketplace builds on a hybrid cloud platform including commercial cloud service providers, publicly funded e-infrastructures and procurers' in-house resources. HNSciCloud will launch a tender call for innovative cloud services. The tender will be presented at the Pre-Commercial Procurement Open Market Consultation (OMC) event, on 17 March 2016, in Geneva, Switzerland.

Once completed, the planned aims of the infrastructure marketplace are to :

- provide access to worldwide and world class resources through a dynamic and sustainable marketplace.
- be built on public and commercial assets, will cover the entire scientific workflow
- offer the broadest range of services

- ensure use of open standard and interoperability of service providers while adhering to European policies, norm and requirements.

It is interesting to see how far and how much industry and research have teamed up in long term projects covering the whole lifecycle, from definition to resource consumption and execution. As research advances, the processes needed to sustain these joint efforts will be identified accordingly and resolved within the ecosystem. The timing is perfect, and gives the consortium a 4/5-year development window before reaching the Exascale Supercomputing era.

© HPC Today 2017 - All rights reserved.

Thank you for reading HPC Today.