THE UNIVERSITY OF TEXAS AT AUSTIN
TEXAS ADVANCED COMPUTING CENTER

From broken bones to black holes, researchers use TeraGrid supercomputers to advance

TeraGrid Provided 1.75 Billion Hours of Computational Time to 1300 users in 2010

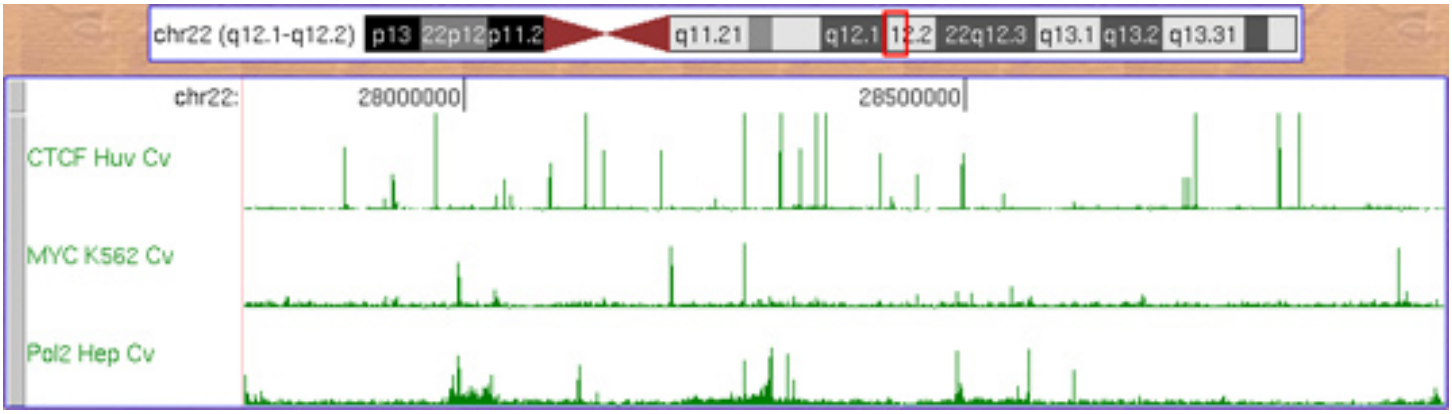

This image shows some of the computational resources that were allocated at the latest TeraGrid Resource Allocation Committee meeting, including (clockwise from top left): Blacklight (Pittsburgh Supercomputing Center); Lonestar 4 (TACC); Ember (National Center for Supercomputing Applications); and Kraken (National Institute of Computational Sciences).

TeraGrid is the world's largest, most comprehensive distributed cyberinfrastructure for open scientific research. Funded by the National Science Foundation, TeraGrid combines leadership-class supercomputing resources at 11 partner sites to create an integrated, persistent computational resource, helping thousands of scientists each year perform critical research.

The TeraGrid Resource Allocation Committee (TRAC) held their final quarterly meeting of 2010 in Austin, Texas in December. With 45 of the top minds in computational science debating who will be given free access to the nation's supercomputers, the peer review process is impressive to witness.

"I am always awed by the collective expertise we have in the room," said Kent Milfeld, a research associate at TACC and the TeraGrid allocations coordinator. "Bringing these researchers together and coordinating the allocation award process has been the highlight of my career."

TeraGrid provides researchers across the country 1.75 billion service units (SUs), or computer processing hours each year – the equivalent of 200,000 years on a single processor. A total of 125 requests were awarded this quarter. These projects span nearly every field of science, and range from clinical investigations of broken bones to simulations of the core collapse of stars to new materials for industry.

TRAC members are stewards of the computational time, storage, and expertise that are provided by TeraGrid's resource provider sites. Each request is reviewed by a team of researchers who are familiar with the science that is represented by the request. Their mission is to make sure that the resources are used efficiently and effectively.

Three new systems were allocated for the first time at the December meeting: Lonestar 4, a 302 teraflops Dell cluster that will go into production at the Texas Advanced Computing Center (TACC) in February; Trestles, a 100 teraflops system at the San Diego Supercomputing Center that is expected to come online in January; and Blacklight, at the Pittsburgh Supercomputing Center (PSC), an Altix UV1000 system that came online in October. Kraken, at the National Institute for Computation Sciences, was recently upgraded, providing additional 144 teraflops to the system for a peak performance of 1,174 teraflops. The three new systems, plus Kraken's upgrade, add 350 million hours annually – or approximately 87 million hours each quarter.

Despite NSF's decades-long commitment to the growth of our national cyberinfrastructure, demand exceeds supply. Over the past few years, emerging fields have begun to leverage high performance computing resources (HPC) and veteran HPC researchers require increasingly more. Consequently, at times it is necessary for the TRAC process to reduce, or deny requests that have not provided

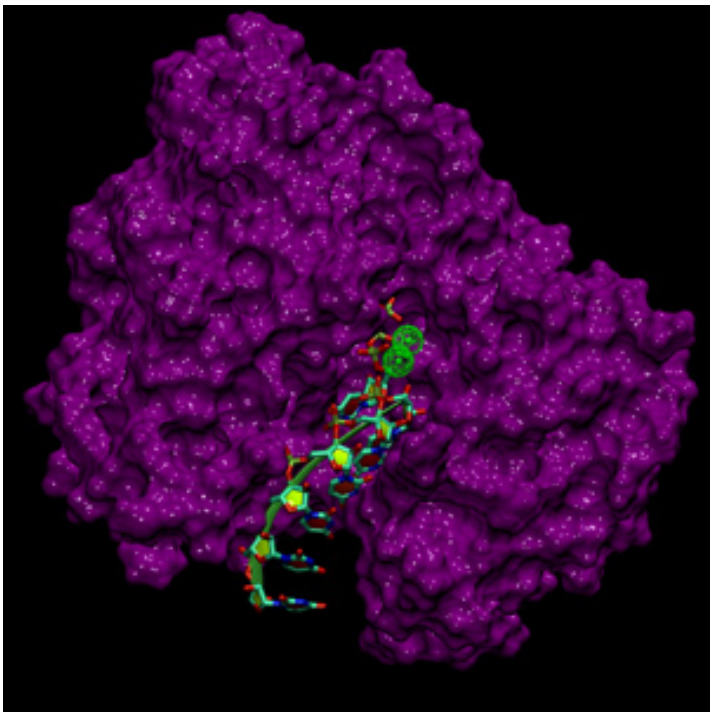

Targeting the mitochondrial machinery for therapeutic potential: TUTase binds mRNA. (credit to Drs. Ozlem Demir and Rommie Amaro, University of California, Irvine; image created with VMD.)

an appropriate justification for the use of a desired resource. The process ensures that the computing needs of each project are carefully debated and allocated to enable varied and promising research.

“The diversity of the group and their enthusiasm for research and computers assures that a fair and thorough evaluation occurs, and that meritorious requests receive the compute time, storage, and support that is needed,” Milfeld said.

The committee allocated four awards of more than 5 million computing hours to projects leveraging TeraGrid resources at TACC. These projects, led by Rommie Amaro (University of California, Irvine), Jeffrey Grossman (MIT), Chris Van de Walle (University of California, Santa Barbara), and Thomas DeGrand (University of Colorado, Boulder), will explore drug design, solar photovoltaic cells, new materials for semiconductors, and particle physics, respectively.

For DeGrand, the opportunity to use Ranger, one of TeraGrid's largest supercomputers, will allow him to study "new physics" beyond the standard model, which might be observed at the Large Hadron Collider in Geneva. "One of the ways we study new theories is via their response to system size, and the largest systems require supercomputers," DeGrand said. "TeraGrid resources will be absolutely essential to the success of our project."

Rommie Amaro, an assistant professor of pharmaceutical sciences, primarily uses TeraGrid resources to run large-scale simulations for biomolecular systems that are involved in disease. The information resulting from these simulations helps identify novel "druggable" target sites on the protein.

"We've been successful in identifying new lead compounds for several infectious diseases, including African sleeping sickness and influenza, as well as cancer," Amaro said. "TeraGrid machines enable us to perform these large-scale simulations easily and efficiently, helping us drive discovery in collaboration with experimental labs."

Not all of the allocations were for extremely large computations. Of the 67 allocations awarded to TACC systems, 35 were for projects that required 1 million SUs or less, which in today's computing environment encompasses everyone from experienced, medium-scale users to new users computing on the TeraGrid for the first time.

"We support the entire computational community, helping researchers grow, evolve, and do their best work," said Milfeld. "In the end, it's about empowering scientific discovery."

Many institutions have Campus Champions — individuals who serve as the local source of knowledge about high-performance computing opportunities and resources. Champions will help those who are new to the process prepare a competitive request. See https://www.teragrid.org/web/eot/campus_champions for more information.

To learn more about TeraGrid's allocation process, visit: https://www.teragrid.org/web/user-support/allocations.

Published January 12, 2011

Rommie Amaro, an assistant professor of pharmaceutical sciences at the University of California, Irvine