Accelerating Discovery and Innovation through Advanced Computing:

Perspective of a High-Performance Computing Scientist in Costa Rica

Esteban Meneses
Advanced Computing Laboratory, Costa Rica National High Technology Center

A new generation of computational scientists is currently being trained across Costa Rica, mainly at research laboratories in public universities. This new wave of scientists is expected to further increase the impact of advanced computing in conceiving creative, powerful ways to understand the universe and solve complex problems.

Since its creation in 1999, the Costa Rica National High Technology Center (CeNAT; http://www.cenat.ac.cr/en/) has fulfilled a fundamental role in fostering interaction among the triple helix—academia, government and industry—through the use and development of high-technology products and services. The Advanced Computing Laboratory belongs to CeNAT. Our laboratory' philosophy is that advanced computing catalyzes scientific discovery and technological innovation. We truly believe algorithms, data structures, and computer architectures have a transformative power to bring about groundbreaking solutions to the most pressing problems in science and engineering.

I am a computer scientist by training and an expert in high-performance computing (HPC). My research concentrates on parallel programming models and fault-tolerance strategies for HPC systems. Yet, my interests extend to all aspects of advanced computing: modeling, simulation, and massive data processing (see http://www.emeneses.org/ for a good collection of my research work).

I began my research career as an instructor in the School of Computing at the Costa Rica Institute of Technology (ITCR) in 2001. I still hold a partial appointment at ITCR as an associate professor of graduate studies. I obtained a PhD in computer science from the University of Illinois at Urbana-Champaign in 2013. After graduation, I was hired as an assistant research professor in the Center for Simulation and Modelling (SaM) of the University of Pittsburgh. In 2015, I

Copublished by the IEEE CS and the AIP
1521-9615/18/$33 ©2018 IEEE

returned to Costa Rica and resumed my academic career as a professor at ITCR. A year later, I took over responsibility as director of the Advanced Computing Laboratory at CeNAT.

Leading the Advanced Computing Laboratory at CeNAT has been an exciting experience. Not a week has passed without finding new colleagues, collaborations, funding opportunities for high-potential projects, or promising research ideas. The CeNAT is in a strategic position since it is part of (and funded by) the national system of public universities. This means that we have a direct connection to scientists and engineers with concrete needs in terms of how to accelerate their simulations or data processing tasks. Our lab at CeNAT employs approximately 15 people, half of them permanent staff, the other half student assistants (see Figure 1). Our members have either electrical engineering or computer science backgrounds. All staff have solid computer programming skills that they use to enrich science with computational tools through a side-to-side collaboration with the scientists and engineers who utilize those tools. We maintain our own computer cluster, participate in research projects, write papers, look for funding opportunities, teach workshops (R, Python, MPI, OpenMP, and OpenACC), and organize academic events. We also very much enjoy working with colleagues from areas as diverse as biology and political science. To learn more about our research projects and services, you can check out our webpage (http://cluster.cenat.ac.cr/).

Figure 1. Staff members of the Advanced Computing Laboratory at CeNAT.

As far as I know, our lab is the only unit in Costa Rica's academic environment devoted to the research and development of advanced computing. However, there are several other laboratories and centers working with computational science and engineering tools. Some of them also own computer clusters. In fact, we have identified more than 20 such computer clusters in the country. Those infrastructures belong mainly to academic units in physics departments across the national system of public universities. There are a few efforts on the government side as well. As for industry, we maintain contact with a couple of groups.

One key goal in our lab is to grow the community of users of our advanced computing services. We organize an annual High-Performance Computing School to train people on HPC programming languages and tools. Approximately 50 people attended the most recent event. In addition, we lead the Scientific Computing Research Network (RICC), which serves as a platform to interact with people interested in advanced computing tools and services. The RICC organizes workshops, talks, and a scholarship program to build bridges among scientific domains. Thanks to these initiatives, our user base is growing stronger.

There are three domains that stand out for burning cycles on our computer infrastructure: oceanography, machine learning, and bioinformatics. Computational oceanography helps in predicting

waves on both coasts of the country. Machine learning is used to tackle signal processing challenges, including those associated with cancer research. Bioinformatics is mostly used to assemble and analyze genomic sequences from a wide range of species (from viruses to trees). Other domains represented in our user base are computational seismology, physics simulations, computational bioacoustics, image analysis, big data processing, and computational neuroscience. Figure 2 presents a visualization of a quake on a volcano in Costa Rica that was created by our collaborator Yuen Law-Wan from RWTH Aachen University. This is just one example of the type of projects being carried out using our advanced computing tools and services.

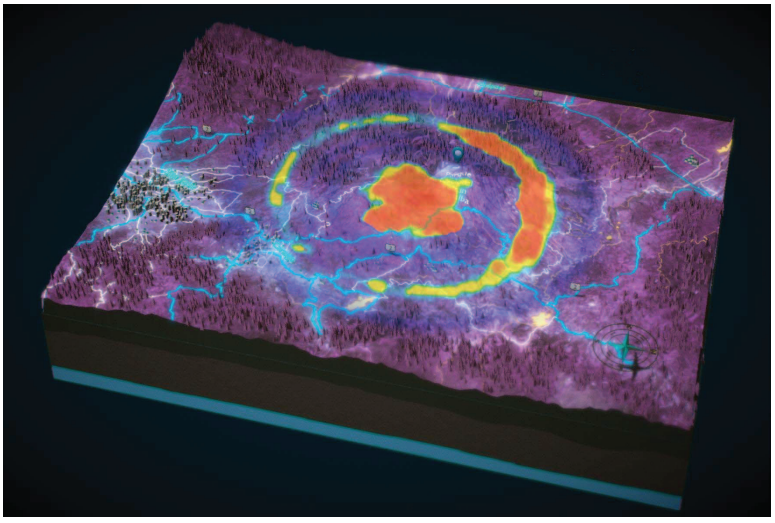

Figure 2. Visualization of seismic activity on a volcano in Costa Rica created by CeNAT collaborator Yuen Law-Wan of RWTH Aachen University.

Other initiatives complement our efforts in promoting the use and development of advanced computing. Master's and doctorate programs provide a platform to develop theses on computational science. Courses on simulation, parallel computing, and computational science are regularly offered as part of those programs, providing a venue for training in the area. There is, however, no specific computational science program in any university in Costa Rica.

A new generation of computational scientists is currently being trained at several institutions across the country, including the Advanced Computing Laboratory at CeNAT, some research laboratories at the University of Costa Rica, the eScience program at the Costa Rica Institute of Technology, and others. We expect this new wave of scientists will further increase the impact of advanced computing in conceiving creative and powerful ways to understand the universe, and solve complex problems.

ABOUT THE AUTHOR

Esteban Meneses is director of the Advanced Computing Laboratory at the Costa Rica National High Technology Center (CeNAT). His research interests include resilience techniques for HPC systems and parallel programming models. Meneses received a PhD in computer science from the University of Illinois at Urbana-Champaign. Contact him at esteban.meneses@acm.org.