Dr. Bishop working with Muhetaer Tuerhong and Cary Randazzo on the LittleFe computer

Little Debbie and Her NanoCar

by Christopher Rodriguez, Cyber Engineering Senior

Computer science is often focused on finding algorithms that complete assignments in a timely manner regardless of the size of the computation. Some problems simply require complex algorithms. Other problems must manipulate extremely large datasets. In most cases, computational time can be reduced by bigger or faster hardware (like moving from a laptop to a desktop), but for large datasets, even the fastest processor would take years to finish. So, what can be done when facing problems like protein folding, database searching, physics simulations or deep learning on millions of examples?

High-performance distributed computing can be used to solve problems in just days which might otherwise take years to run on a normal computer. This speedy computation is achieved by breaking the solution into smaller problems that are distributed to computational nodes that are combined into a single computer cluster. The typical home computer might have two or even four processors, but super computers can have thousands to tens of thousands of nodes, each containing multiple processors with multiple computing cores. All of the nodes can then be utilized to solve the overarching problem. Still, constructing and using a super computer is not as simple as just hooking up more processors to a motherboard and clicking "run." Each node in a cluster functions as an autonomous computer with its own memory, hard drive and operating system. Making sure that they all work together efficiently is not a trivial problem.

Most students never get the chance to work with computer clusters. The few lucky students that get experience with cluster computers usually only get to interact with them remotely over a secure shell (ssh) or similar protocols. While this is the normal way to interact with a cluster, the hardware, operating systems and computer administration tasks are purposefully hidden from the user to simplify the user experience and provide a secure computing environment. This lack of transparency often means that students don't get a chance to really understand what goes into making a cluster computer function well.

During the Spring Quarter of 2018, Dr. Tom Bishop, associate professor of chemistry, molecular science and nanotechnology, nanosystems engineering and physics, led a class of seven students to explore high-performance cluster computing using a Little Fe computer. ("Fe" is the chemical symbol for iron, and supercomputers are often referred to as Big Iron.) Little Fe is a computer cluster that consists of six Netbook computer nodes combined together within a single frame and connected to a network switch. Each of the nodes can be connected to a display unit, have their power toggled, have USB devices connected to them and be connected and disconnected from the network. These features make Little Fe an ideal platform for students to get a hands-on learning experience in high-performance cluster computing and all the ways such a system can fail. The skills learned are in high demand.

The students reconfigured Little Fe with a Debian variant of the Linux operating system. Little Debbie is more user friendly and has a more extensive software stack than Little Fe. They then learned how to install the Nanoscale Molecular Dynamics (NAMD) software to run molecular dynamics simulations. The students first experimented with molecular simulations of Bovine pancreatic trypsin inhibitor (BPTI), a small protein found in cattle, using different configurations of Little Debbie. After developing a strong understanding of the computer architecture and simulation methods, the team began developing a NanoCar simulator. They created a computational model of a chemical structure made primarily of carbon atoms and placed the molecule on a fixed gold surface. When they applied heat to the system to reach a temperature of 310 K, the NanoCar randomly traveled over the surface at speeds on the order of a nanometer per nanosecond or approximately 2mph.

In the future, Dr. Bishop hopes to expand this class so that more undergraduates can be exposed to molecular dynamics and high-performance computing using the NanoCar and Little Fe. Hands-on opportunities like this offer Louisiana Tech students educational opportunities that cannot be found everywhere. The class, Applications of Parallel Computing in Bio and Nano Molecular Modeling, was the result of a COES Innovations in Education Grant. Dr. Bishops seeks to grow the Little Fe into a family of Littles (Little Susie, Little Red, Little Mints, to name a few) corresponding to different Linux distributions, and he wishes to continue exploring the modeling and simulation of NanoCars and Biomolecules.

Sources:
https://foresight.org/Conferences/MNT05/Abstracts/Michabst.php
http://littlefe.net/
https://www.top500.org
http://dna.engr.latech.edu/~zli007/Nanocar/
http://dna.engr.latech.edu/~zli007/Magic-Snake/
http://dna.engr.latech.edu/~gdash/