

LittleFe + BCCD + CSERD = Acme

Computational Science Education *on the Move*

Charles Peck <charliep@earlham.edu>, David Joiner <djoiner@comcast.net>, Paul Gray <gray@cs.uni.edu>, Tom Murphy <tmurphy@contracosta.edu>
Alex Lemann <lemanal@earlham.edu>, Josh McCoy <mccojo@cs.earlham.edu>, Kevin Hunter <hunteke@earlham.edu>, Kristina Wanous <kwanous@uni.edu>, Skylar Thompson <skylar@cs.earlham.edu>

Abstract: Most colleges have no access to parallel platforms for in-class computational science or parallel/distributed computing education. Key concepts, motivated by science, are taught more effectively and memorably on an actual parallel platform. LittleFe is a complete multi-node Beowulf-style portable computational cluster. The entire package weighs less than 50 pounds, easily travels via checked baggage, and sets-up in 5 minutes. Acme leverages the Bootable Cluster CD (BCCD) project and the Computational Science Education Reference Desk (CSERD). For the price of a high-end laptop, it is a powerful, ready-to-run platform supporting computational science and High Performance Computing (HPC).

Acme, "Computational Science *on the Move*", is visceral: students participate with the science. A protein folding, visualized in real-time, provides an engaging presentation to the Xbox generation, attracting more students to STEM disciplines. We must reach students with multidisciplinary, interactive, computational science curriculum, allowing them to envision themselves in an exciting science based career.

Hardware

The hardware component of Acme is LittleFe, a portable computational cluster. LittleFe's design grew out of our work building stationary clusters and our experience teaching workshops in a variety of places that lacked parallel computational facilities. The principle design constraints for LittleFe are:

- \$3,000 (USD) total cost
- Less than 50lb (including the Pelican case)
- Less than 5 minutes of setup time
- Minimal power consumption

Currently, each LittleFe is composed of the following parts:

- 6 motherboards (mini-ITX or similar, 1GHz+ CPU, 512MB RAM, 100MB ethernet)
- 6 12VDC-ATX power supplies
- 1 320 Watt 110VAC-12VDC switching power supply
- 1 40GB 7200RPM ATA disk drive (2.5" form factor)
- 1 DVD/CD optical drive (slim-line form factor)
- 1 8 port 100MB ethernet switch
- 1 rack assembly
- 1 1610 Pelican travel case
- Fasteners, cabling, and mounting hardware

The total cost per unit is about \$3,000 in parts and labor (about 10 hours of student labor to assemble and test). The motherboards, CPUs, and RAM comprise the bulk of the cost.

With all 6 nodes idling, LittleFe draws about 80 Watts of power (about the same as a light bulb). When running a CPU-intensive molecular dynamics simulation on every node, LittleFe draws about 88 Watts of power.



<http://LittleFe.net/>

Software

For the operating platform, Acme uses the Bootable Cluster CD (BCCD), a project of Paul Gray's group at the University of Northern Iowa. The BCCD comes in two flavors, a live CD version and a liberated version. The live CD makes it possible to turn any x86-based lab into a teaching cluster in a matter of minutes. The liberated version, used by Acme, is permanently installed on LittleFe's disk drive.

The BCCD contains hundreds of software applications to support parallel and distributed computing and computational science education, this greatly reduces the friction associated with setting-up an environment for HPC-based education.

A small sample of the tools included in the BCCD include:

- gcc, g77, and development tools, editors, profiling libraries and debugging utilities
- Cluster Command and Control tools
- MPICH, LAM-MPI and PVM
- The X Window System
- OpenMosix with openmosixview and userland openMosix tools
- Torque and Maui scheduler support
- octave, gnuplot, Mozilla's Firefox and about 1,400 userland utilities
- Network configuration and debugging utilities
- Ganglia and other monitoring packages

Science curriculum modules available via the BCCD's `list-packages` include:

- GROMACS – molecular dynamics primarily for biological molecules
- mpiBLAST – parallel implementation of the basic local alignment search tool



<http://bccd.net/>

Education

Our goal is science *education*. This proves a difficult task when the majority of educators and scientists do not have a common definition of 'education,' nor is there an agreement on how, when, or where it should be enacted. Many educators and scientists think of education and science as two totally disparate worlds.

Educational needs are very similar to science needs: The software in the materials must work, the science must be correct, and the results must accurate, useful, and applicable. The Computational Science Education Reference Desk (CSERD), a Pathways project of the National Science Digital Library (NSDL) and the National Science Foundation (NSF), gives teachers and scientists alike access to resources that meet these three guidelines by reviewing them in the following context:

- **verification** – Is the model logically correct and does it follow from the physical and mathematical laws used?
- **validation** – Did the model correctly predict the modeled phenomena?
- **accreditation** – Does the model or simulation fulfill its educational purpose?

Resources in CSERD range over multiple scientific disciplines and covers materials for primary, secondary, and post-secondary education.

For HPC education, we have tutorials and example code available to illustrate different levels of parallelism: embarrassingly parallel parameter searches, domain decomposition programs that require only nearest neighbor communication, and tightly coupled N-body examples. Each of these codes shows a parallel programming example in the context of solving a science problem.



<http://www.shodor.org/cserd/>