Chapter 4

LittleFe

LittleFe is a complete multi-node Beowulf-style [Brown, 2008] portable computational cluster designed as an "educational appliance" for substantially reducing the friction associated with teaching high performance computing (HPC) and computational science in a variety of settings. The entire package weighs less than 50 pounds, easily travels via checked baggage, and sets-up in 5 minutes. Working with colleagues Paul Gray, Thomas Murphy, and David Joiner, I am jointly responsible for the design and primarily responsible for the engineering and production of *LittleFe*.

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. Once we had some gear and some experience moving it around we worked through three different approaches before arriving at the system described here. The principle design constraints for *LittleFe* are:

- $3,000USD total cost
- Less than 50lb (including the Pelican travel case)
- Less than 5 minutes to setup

- Minimal power consumption; less than 100 Watts peak, 80 Watts average

The current production *LittleFe* design is composed of the following major components:

- 6 mainboards (Mini-ITX, 1–2GHz CPU, 512MB–1GB RAM, 100Mb/1Gb 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/1Gb ethernet switch
- 1 rack assembly
- 1 1610 Pelican travel case
- Fasteners, cabling, and mounting hardware

The $3,000USD cost per unit includes about 10 hours of student labor to assemble and test each unit. This includes liberating the Bootable Cluster CD image onto the disk drive and configuring the users. The mainboards, 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 an incandescent light bulb). When running a CPU-intensive molecular dynamics simulation using every node *LittleFe* draws about 88 Watts of power. See Appendix A *LittleFe - Parts Manifests* for a detailed parts manifests, cost estimates and sources.

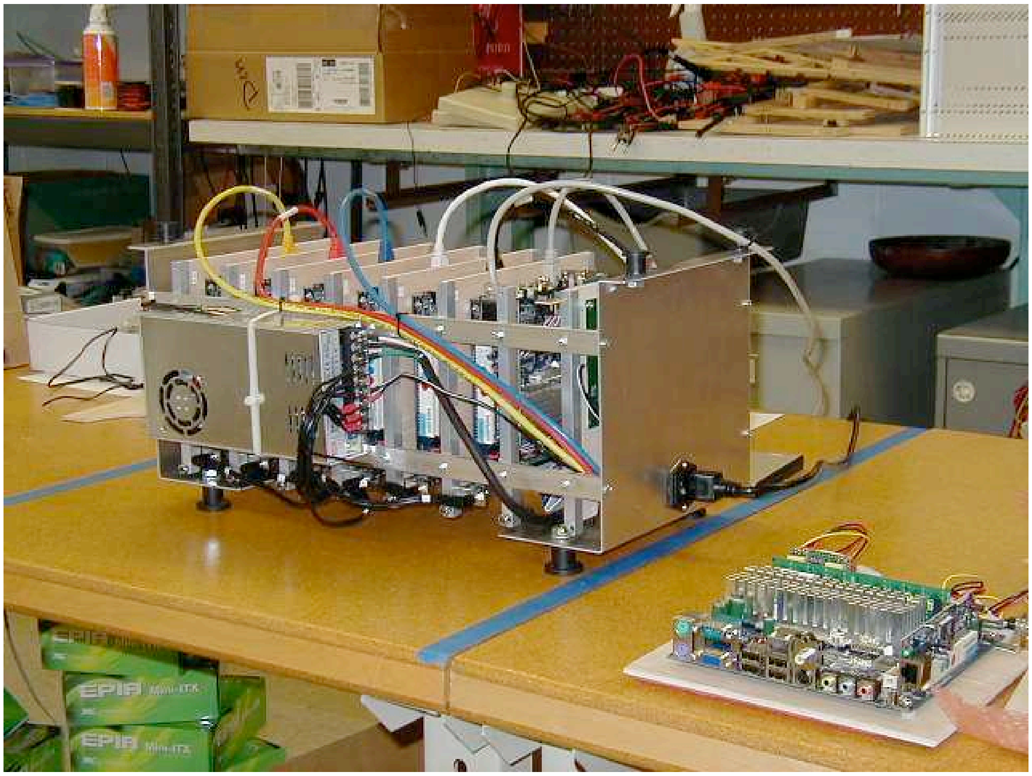
Figure 4.1: An early version 3 production unit almost ready for deployment.

4.1 Motivation

One of the principle challenges to computational science and HPC education is that many institutions do not have access to HPC platforms for demonstrations and laboratories. Paul Gray's Bootable Cluster CD (BCCD) project [Gray, 2004] has made great strides in this area by making it possible to non-destructively, and with little effort, convert a computer lab of Windows or Macintosh computers into an *ad hoc* cluster for educational use.

LittleFe takes that concept one step further by merging the BCCD with an inexpensive design for an 6-8 node portable computational cluster. The result is a machine that weighs less than 50 pounds, easily and safely travels via checked baggage on the airlines, and sets-up in 5 minutes wherever there is a 110VAC outlet and a wall to project an image on. The BCCD's package management feature supports curriculum modules in a variety of natural science disciplines, making the combination of *LittleFe* and the BCCD a ready-to-run solution for computational science and HPC education.

LittleFe's principle edge is resource availability for computational science education. To teach a realistic curriculum in computational science, there must be guaranteed and predictable access to HPC resources. There are currently two common barriers to this access. Local policies typically allocate HPC resources under a "research first, pedagogy second" prioritization scheme, which often precludes the use of "compute it now" science applications in the classroom. The second barrier is the capital and ongoing maintenance costs associated with owning an HPC resource. This affects most mid-size and smaller educational institutions and is particularly acute in liberal arts environments, community colleges, and K-12 settings.

While relatively low-cost Beowulf style clusters have improved this situation somewhat, HPC resource ownership is still out of reach for many educational institutions. *LittleFe's* total cost is less than $3,000USD, making it easily affordable by a wide variety of K-16 schools. This is particularly important for institutions which serve traditionally under-served groups; typically they have access to fewer technology resources than other schools.

LittleFe's second important feature is ease of use, both technically and educationally. Our adoption of the BCCD as the software distribution toolkit makes it possible to smoothly and rapidly advance from bare hardware to science. Further, we have minimized ongoing maintenance since both hardware and software are standardized. Paul Gray, from the University of Northern Iowa, and a number of our student research assistants have successfully maintained the BCCD for many years now via a highly responsive listserv and well organized web presence, `http://bccd.net`.

Portability is useful in a variety of settings, such as workshops, conferences, outreach events and the like. It is also useful for educators, whether illustrating principles in the K-12 arena or being easily passed from college classroom to college classroom.

4.2 Overall Design

The first *LittleFe* consisted of eight Travla Mini-ITX VIA computers placed in a nearly indestructible Pelican case. To use it one would take all the nodes, networking gear, power supplies, *etc.* out of the case and set it up on a table. Each node was a complete computer with its own hard drive. While this design met the portability, cost, and low-power design goals, it was overweight and deployment was both time-consuming and error-prone.

Successive versions of *LittleFe* have moved to a physical architecture where the compute nodes are bare Mini-ITX mainboards mounted in a custom designed frame, which in turn is housed in a Pelican traveling case. To accomplish this we stripped the Travla nodes down, using only their mainboards, and replaced their relatively large power supplies with daughter-board style units that mount directly to the mainboard's ATX power connector. These changes saved both space and weight. Current *LittleFe's* use disk-less compute nodes, only the head node has a disk drive. Removing seven disk drives from the system reduced power consumption considerably and further reduced the weight and packaging complexity.

4.3 Hardware

4.3.1 Mainboard

Basing our design around the Mini-ITX mainboard form factor standard has served us well. Currently there is significant demand in the industry for this size system, which yields rapid evolution, a significant number of choices from multiple vendors such as VIA, Intel, and Advanced Micro Devices (AMD), and consequently low price points.

Smaller boards such as the PC-104 form factor, while using less electrical power, lack the computational power to be useful. Larger boards, while offering much more computational power, would be impractical in terms of both the physical packaging and electrical power consumption.

When specifying the mainboard, care should be taken to ensure that the overall height is less than the inter-board spacing in the frame (see the diagrams in Appendix B *LittleFe - Assembly Instructions*.

The head node should have a minimum of 1GB of RAM since it will be a file server for itself and all the compute nodes. Compute nodes should have at least 512MB of RAM.

4.3.2 Storage

All of the persistent storage devices are attached to the head node. Due to packaging and weight constraints the disk drive must be in a 2.5" form factor. The disk drive can be ATA or SATA: spindle speed, buffer size, and overall transfer rate are the most important criteria. Since there is only a single disk in each *LittleFe*, it should be a fast one. For most applications a 60GB disk drive is sufficient.

The speed of the CD/DVD is not particularly important as it is usually only used to load software. The CD/DVD drive must be a slimline form factor to fit with the disk drive on the frame.

4.3.3 Network Fabric

While 100Mb networking is sufficient for some applications, the availability and cost of Mini-ITX mainboards with 1Gb NICs makes them very attractive now for most new units.

Any small unmanaged 8–12 port network switch which uses 9–12VDC for line-in voltage can usually be mounted in the frame. Some *LittleFe* units sport a small WiFi access point, allowing a group of people to interact with a simulation from their laptops.

4.3.4 Power

120-240VAC line-input is brought to a frame mounted fused switch and then routed to a 320 Watt 12VDC regulated power supply. This minimizes the amount of high-voltage wiring in the system and provides the source for powering the 120 Watt daughter-board ATX power supplies located on each mainboard. With a peak draw of about 100 Watts, the primary power supply is generously sized for cool operation and increased reliability and longevity. For the head node, the daughter-board ATX power supply provides Molex connectors for the disk drive and CD/DVD drive.

4.3.5 Cooling

LittleFe's open frame, vertically mounted board design promotes a significant amount of natural cooling. This reduces the need for additional fans, further reducing the power consumption profile and the amount of system noise. *LittleFe* is quiet enough to use even when sitting in the arrivals lounge at an airport.

The Pelican traveling case, with the lid open and nothing packed around *LittleFe*, provides enough circulation that the unit can be run without removing it from the case. This is a particularly useful feature when using *LittleFe* in the field; for example when collecting, analyzing, and visualizing data from an attached water parameter probe.

4.4 Packaging

4.4.1 Frame

The frame is made of .080" smooth plate aluminum with punches for the rail mounts, case mounts, and line-input power switch. The rails and board guides are also made of standard aluminum stock with pre-drilled mounting holes. See Appendix B *LittleFe - Assembly Instructions* for a diagram of the frame.

The mainboards are mounted on 1/4" AA luan plywood plates using 1/8" nylon standoffs. We explored many other materials for the mainboard plates, particularly aluminum and a variety of plastics and polymers. None could match the strength to weight ratio, cost, or ease of use associated with high quality plywood.

4.4.2 Traveling Case

The system is shock-mounted in a Pelican 1610 traveling case using a two-part rubber cup and plug system. The cups are mounted on the floor and lid of the Pelican case. The plugs are mounted on the top and bottom of each frame's end-plates. When you place the frame in the case the plugs nestle in the cups on the floor. When the lid is closed those cups encase the plugs on top. This gives the system support and shock resistance in each direction. We have tested this system extensively both in field trials and by examining the results of approximately 25 commercial flights where *LittleFe* travelled as checked baggage. While there is occasional distortion of an end-plate if an excessive amount of baggage is placed on top of the Pelican case, on the whole the system appears to functional adequately.

In terms of overall weight, *LittleFe*, traveling case, and any accessories, can be no more than 50lbs. This is the maximum allowed by airlines before heavy baggage surcharges apply. Practically, it is also about the maximum amount that most people can safely maneuver around an airport or school building. One advantage of the Pelican 1610 is that is has built-in wheels and a retractable tow handle.

4.5 Assembly and Testing

Assembling *LittleFe* consists of the following steps:

1. Assembling the frame and rails
2. Mounting the regulated 110/240 VAC power supply to the frame
3. Mounting the network switch and installing the network cabling
4. Mounting the mainboards to the cards
5. Mounting the power supplies and switches to the mainboards
6. Installing the up-link NIC on the head-node
7. Installing the mainboards in the cage
8. Cabling the power supplies
9. Mounting the disk drive and CD/DVD drive to the frame and installing the power and data cables

10. Plugging in the monitor, keyboard, and mouse
11. Performing the initial power-up tests
12. Configuring the BIOS on 5 of the mainboards to boot via the LAN and PXE

Basic hand tools: screwdrivers, pliers, wire cutters, adjustable wrench, drill, and a soldering iron are all the tools which are needed to fully assemble a unit. Most people budget a full day to do a complete assembly, test, and software installation. With practice, it has been shown that if nothing goes wrong, a single unit can be assembled in about 4 hours. See Appendix B *LittleFe - Assembly Instructions* for detailed step-by-step instructions and the URL of a video that illustrates the assembly of a unit.

4.6 Software

Early versions of *LittleFe* used the Debian Linux distribution as the basis for the system software. This was augmented by a wide variety of system, communication, and computational science packages, each of which had to be installed and configured on each of the nodes. Even with cluster management software such as the C3 tools, this was still a time-consuming process.

One of the primary goals of this project has been to reduce the friction associated with using HPC resources for computational science education. This friction is made up of the time and knowledge required to configure and maintain HPC resources. To this end, *LittleFe's* system software was re-designed to use Paul Gray's Bootable Cluster CD distribution [Gray, 2004]. The BCCD comes ready-to-run with many of the system and scientific software tools necessary to support a wide range of computational science education.

4.7 Status

Funding from TeraGrid, the SuperComputing Conference, and private sources has enabled my group to put about 15 *LittleFe* units into production as of this writing. *LittleFe* units are used in a variety of contexts: undergraduate computer science education at Earlham and other colleges and universities, K-12 science outreach and engagement programs, the SuperComputing Education Program's workshops and conference program, and the Dine'h Grid project of the Navajo Nation in Crownpoint, New Mexico.

LittleFe is very much a work in progress. Over the past three years, my colleagues and I have done extensive work in this area prototyping and testing fundamental design considerations, developing power and cooling solutions within a narrow design envelope, and porting and developing software laboratories for education, outreach, and training. As Moore's "Law" [Moore, 1965] continues to hold true, we reconsider design choices in an effort to make *LittleFe* smaller, cheaper, more powerful computationally, and lower in power consumption.

For more information about *LittleFe* see the chapter *Results and Future Work*.

Figure 4.2: A demonstration of *LittleFe* at the Oklahoma Supercomputing Symposium in 2006.