



Bootable Cluster CD Cascadia 2011

Andrew Fitz Gibbon

Shodor Foundation

`fitz at cs.earlham.edu`

Tom Murphy

Contra Costa College

`tmurphy at contracosta.edu`

Skylar Thompson

University of Washington

`skylar at cs.earlham.edu`

Sam Leeman-Munk

Shodor Foundation

`sleemanmunk at gmail.com`

Charlie Peck

Earlham College

`charliep at cs.earlham.edu`

Aaron Weeden

Earlham College

`amweeden06 at cs.earlham.edu`



Sponsors

Sponsors

Terminology

What is the BCCD?

BCCD early history

BCCD today

What is the problem?

What is the problem?
(cont'd)

What is the problem?
(cont'd)

How is BCCD an
answer?

General usage of the
BCCD

Parallel methods

BCCD can use

Visualizations with the
BCCD

Nitty-gritty details of
BCCD

LittleFe

LittleFe Picture

More information

Live Demo!

References

Questions (and
answers)

- Intel Corporation
- Supercomputing Conference
- Earlham College
- Shodor Foundation
- Contra Costa College



Terminology

Sponsors

Terminology

What is the BCCD?

BCCD early history

BCCD today

What is the problem?

What is the problem?
(cont'd)

What is the problem?
(cont'd)

How is BCCD an
answer?

General usage of the
BCCD

Parallel methods

BCCD can use

Visualizations with the
BCCD

Nitty-gritty details of
BCCD

LittleFe

LittleFe Picture

More information

Live Demo!

References

Questions (and
answers)

AUFS Another Union FileSystem

BCCD Bootable Cluster CD (not just for CDs anymore!)

CUDA Compute Unified Device Architecture

CSE Computational science education

DHCP Dynamic Host Configuration Protocol

GPGPU General Purpose computation on Graphics Processor Units

HPC High performance computing

MPI Message Passing Interface

STEM Science, Technology, Engineering, and Mathematics



What is the BCCD?

- Sponsors
- Terminology
- What is the BCCD?
- BCCD early history
- BCCD today
- What is the problem?
- What is the problem? (cont'd)
- What is the problem? (cont'd)
- How is BCCD an answer?
- General usage of the BCCD
- Parallel methods BCCD can use
- Visualizations with the BCCD
- Nitty-gritty details of BCCD
- LittleFe
- LittleFe Picture
- More information
- Live Demo!
- References
- Questions (and answers)

- A software tool for domain educators for teaching CSE
- A stable platform for the writing of curricula for CSE
- One part of a larger CSE community
 - Shodor Foundation
 - National Computational Sciences Institute (NCSI)
 - CSE Reference Desk (CSERD)
- “Teaching Parallelism Made Easy”



BCCD early history

- Sponsors
- Terminology
- What is the BCCD?
- BCCD early history**
- BCCD today
- What is the problem?
- What is the problem?
(cont'd)
- What is the problem?
(cont'd)
- How is BCCD an
answer?
- General usage of the
BCCD
- Parallel methods
- BCCD can use
- Visualizations with the
BCCD
- Nitty-gritty details of
BCCD
- LittleFe
- LittleFe Picture
- More information
- Live Demo!
- References
- Questions (and
answers)

■ BCCDv1 – BCCDv2

Begun in 2003 by Paul Gray and students

Based on the LNX-BBC distribution

Custom build and packaging mechanism

Assumed to be live CD, fixed cluster difficult to configure



BCCD today

- Sponsors
- Terminology
- What is the BCCD?
- BCCD early history
- BCCD today**
- What is the problem?
- What is the problem?
(cont'd)
- What is the problem?
(cont'd)
- How is BCCD an answer?
- General usage of the BCCD
- Parallel methods
- BCCD can use
- Visualizations with the BCCD
- Nitty-gritty details of BCCD
- LittleFe
- LittleFe Picture
- More information
- Live Demo!
- References
- Questions (and answers)

■ BCCDv3

Work started in 2006 - Motivated by “liberation”

Debian-based

Standard Debian methods for build and configuration

Fixed clusters an assumed state for the BCCDv3

Backwards-compatible with BCCDv2

■ Five current active developers

■ Current work includes:

Automated test suite

Live upgrades

Accessibility



What is the problem?

Sponsors
Terminology
What is the BCCD?
BCCD early history
BCCD today
What is the problem?
What is the problem?
(cont'd)
What is the problem?
(cont'd)
How is BCCD an
answer?
General usage of the
BCCD
Parallel methods
BCCD can use
Visualizations with the
BCCD
Nitty-gritty details of
BCCD
LittleFe
LittleFe Picture
More information
Live Demo!
References
Questions (and
answers)

- STEM education is becoming more important than ever [3]
- Computers are a large part of any science curriculum
- Early introduction of parallel & computational methods is critical
- Schools do not have the resources or expertise

What is the problem? (cont'd)

Sponsors
Terminology
What is the BCCD?
BCCD early history
BCCD today
What is the problem?
**What is the problem?
(cont'd)**
What is the problem?
(cont'd)
How is BCCD an
answer?
General usage of the
BCCD
Parallel methods
BCCD can use
Visualizations with the
BCCD
Nitty-gritty details of
BCCD
LittleFe
LittleFe Picture
More information
Live Demo!
References
Questions (and
answers)

■ Computer labs

(+) Visual

(-) Can't take it home

(-) Not all parallel paradigms available

(-) Not stable (reboots)

(-) Not predictable or uniform (what software is installed?)

What is the problem? (cont'd)

Sponsors
Terminology
What is the BCCD?
BCCD early history
BCCD today
What is the problem?
What is the problem?
(cont'd)
What is the problem?
(cont'd)
How is BCCD an
answer?
General usage of the
BCCD
Parallel methods
BCCD can use
Visualizations with the
BCCD
Nitty-gritty details of
BCCD
LittleFe
LittleFe Picture
More information
Live Demo!
References
Questions (and
answers)

- Compute cluster (if one is available)
 - (+) Generally most parallel paradigms available
 - (+) Relatively stable & uniform
 - (-) Job scheduler interferes with immediate, visual results
 - (-) Primarily (often completely) for research, not education
 - (-) Can't take it home



How is BCCD an answer?

- Sponsors
- Terminology
- What is the BCCD?
- BCCD early history
- BCCD today
- What is the problem?
- What is the problem?
(cont'd)
- What is the problem?
(cont'd)
- How is BCCD an
answer?**
- General usage of the
BCCD
- Parallel methods
BCCD can use
- Visualizations with the
BCCD
- Nitty-gritty details of
BCCD
- LittleFe
- LittleFe Picture
- More information
- Live Demo!
- References
- Questions (and
answers)

- Live CD (or USB or VM) - runs in a computer lab environment

- No need to be a sysadmin

Networking is automatic and non-disruptive

Clustering tools pre-configured (MPI, OpenMP, Torque, etc.)

- Visualization easy because every system has a GUI
- Can be taken home or back to class



General usage of the BCCD

- Sponsors
- Terminology
- What is the BCCD?
- BCCD early history
- BCCD today
- What is the problem?
- What is the problem? (cont'd)
- What is the problem? (cont'd)
- How is BCCD an answer?
- General usage of the BCCD**
- Parallel methods
- BCCD can use
- Visualizations with the BCCD
- Nitty-gritty details of BCCD
- LittleFe
- LittleFe Picture
- More information
- Live Demo!
- References
- Questions (and answers)

- Change out software packages with `Modules`
- Copy data to all BCCD nodes with `bccd-sync-dir`
- Liberation

Make BCCD permanent using “liberation”

Post-liberation, `bccd-reset-network` will setup a diskless cluster



Parallel methods BCCD can use

- Sponsors
- Terminology
- What is the BCCD?
- BCCD early history
- BCCD today
- What is the problem?
- What is the problem? (cont'd)
- What is the problem? (cont'd)
- How is BCCD an answer?
- General usage of the BCCD
- Parallel methods BCCD can use**
- Visualizations with the BCCD
- Nitty-gritty details of BCCD
- LittleFe
- LittleFe Picture
- More information
- Live Demo!
- References
- Questions (and answers)

- Embarrassingly parallel - just use `bccd-synkdir` and `ssh`
- Shared memory (pthreads, OpenMP)
- Message passing (OpenMPI)
- GPGPU (CUDA, OpenCL in the future)
- Hybrid methods

Combinations of shared memory & message passing

Or combinations of message passing & GPGPU

Etc.

Visualizations with the BCCD

Sponsors
Terminology
What is the BCCD?
BCCD early history
BCCD today
What is the problem?
What is the problem?
(cont'd)
What is the problem?
(cont'd)
How is BCCD an
answer?
General usage of the
BCCD
Parallel methods
BCCD can use
**Visualizations with the
BCCD**
Nitty-gritty details of
BCCD
LittleFe
LittleFe Picture
More information
Live Demo!
References
Questions (and
answers)

■ GalaxSee

N-body simulation of galaxy formation

Uses MPI for communication

■ Parameter Space

Monte Carlo simulation of dart throwing

Uses MPI for communication

■ Life

Cellular automata simulation

Can use serial, OpenMP, or MPI



Nitty-gritty details of BCCD

- Sponsors
- Terminology
- What is the BCCD?
- BCCD early history
- BCCD today
- What is the problem?
- What is the problem? (cont'd)
- What is the problem? (cont'd)
- How is BCCD an answer?
- General usage of the BCCD
- Parallel methods
- BCCD can use
- Visualizations with the BCCD
- Nitty-gritty details of BCCD**
- LittleFe
- LittleFe Picture
- More information
- Live Demo!
- References
- Questions (and answers)

- Live CD and diskless nodes use AUFS
Read/write branch stored in RAM
- Live CD uses cloop back-end (from KNOPPIX)
- DHCP lock-down is accomplished by using custom options
- Node auto-detection uses custom “pkbcast”



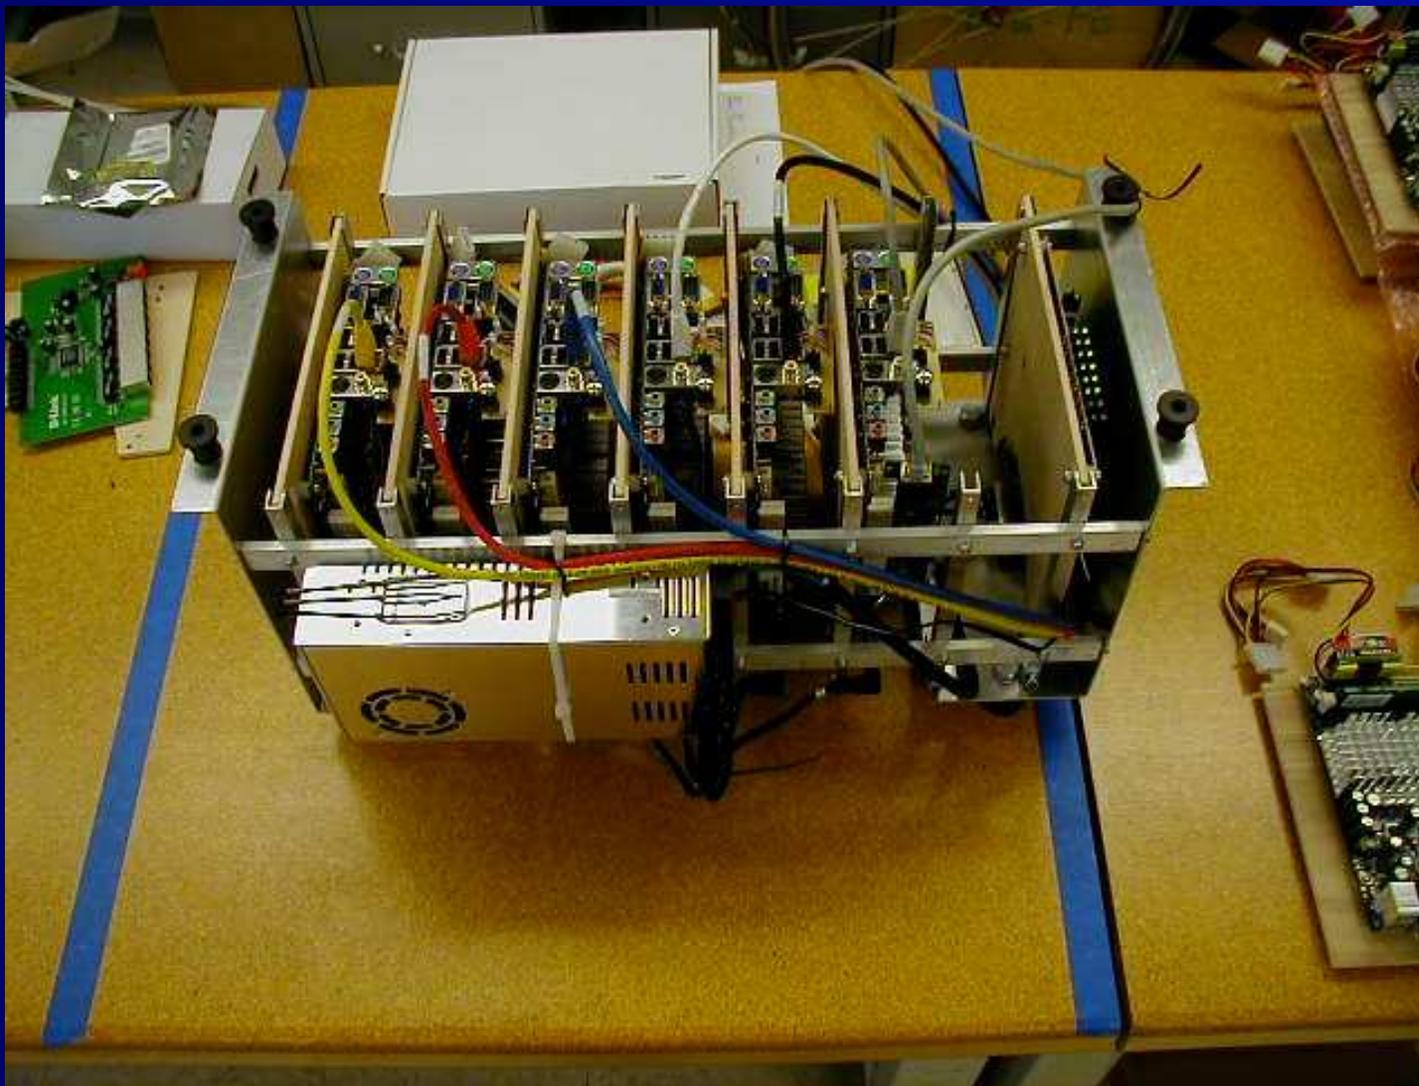
LittleFe

- Sponsors
- Terminology
- What is the BCCD?
- BCCD early history
- BCCD today
- What is the problem?
- What is the problem?
(cont'd)
- What is the problem?
(cont'd)
- How is BCCD an
answer?
- General usage of the
BCCD
- Parallel methods
BCCD can use
- Visualizations with the
BCCD
- Nitty-gritty details of
BCCD
- LittleFe**
- LittleFe Picture
- More information
- Live Demo!
- References
- Questions (and
answers)

- Solution for people who don't want to build their own cluster [2]
- BCCD designed to run on LittleFe
- Designed to be portable
- Hardware
 - 6x dual-core Intel Atom boards
 - NVIDIA/CUDA support
 - Ethernet switch
 - AC/DC power supply
 - Custom aluminum case

LittleFe Picture

- Sponsors
- Terminology
- What is the BCCD?
- BCCD early history
- BCCD today
- What is the problem?
- What is the problem?
(cont'd)
- What is the problem?
(cont'd)
- How is BCCD an answer?
- General usage of the BCCD
- Parallel methods
- BCCD can use
- Visualizations with the BCCD
- Nitty-gritty details of BCCD
- LittleFe
- LittleFe Picture**
- More information
- Live Demo!
- References
- Questions (and answers)





More information

- Sponsors
- Terminology
- What is the BCCD?
- BCCD early history
- BCCD today
- What is the problem?
- What is the problem? (cont'd)
- What is the problem? (cont'd)
- How is BCCD an answer?
- General usage of the BCCD
- Parallel methods BCCD can use
- Visualizations with the BCCD
- Nitty-gritty details of BCCD
- LittleFe
- LittleFe Picture
- More information**
- Live Demo!
- References
- Questions (and answers)

■ BCCD

Main website: <http://bccd.net>

Email: [bccd-developers at bccd.net](mailto:bccd-developers@bccd.net)

■ LittleFe

Main website: <http://littlefe.net>

Email: [littlefe at cs.earlham.edu](mailto:littlefe@cs.earlham.edu)

Build-out event: <http://littlefe.net/buildout>

■ *Teaching high performance computing to undergraduate faculty and undergraduate students [1]*

Teragrid 2010 Best Paper in Education

Available to ACM members

Live Demo!

Sponsors

Terminology

What is the BCCD?

BCCD early history

BCCD today

What is the problem?

What is the problem?

(cont'd)

What is the problem?

(cont'd)

How is BCCD an answer?

General usage of the BCCD

Parallel methods

BCCD can use

Visualizations with the BCCD

Nitty-gritty details of BCCD

LittleFe

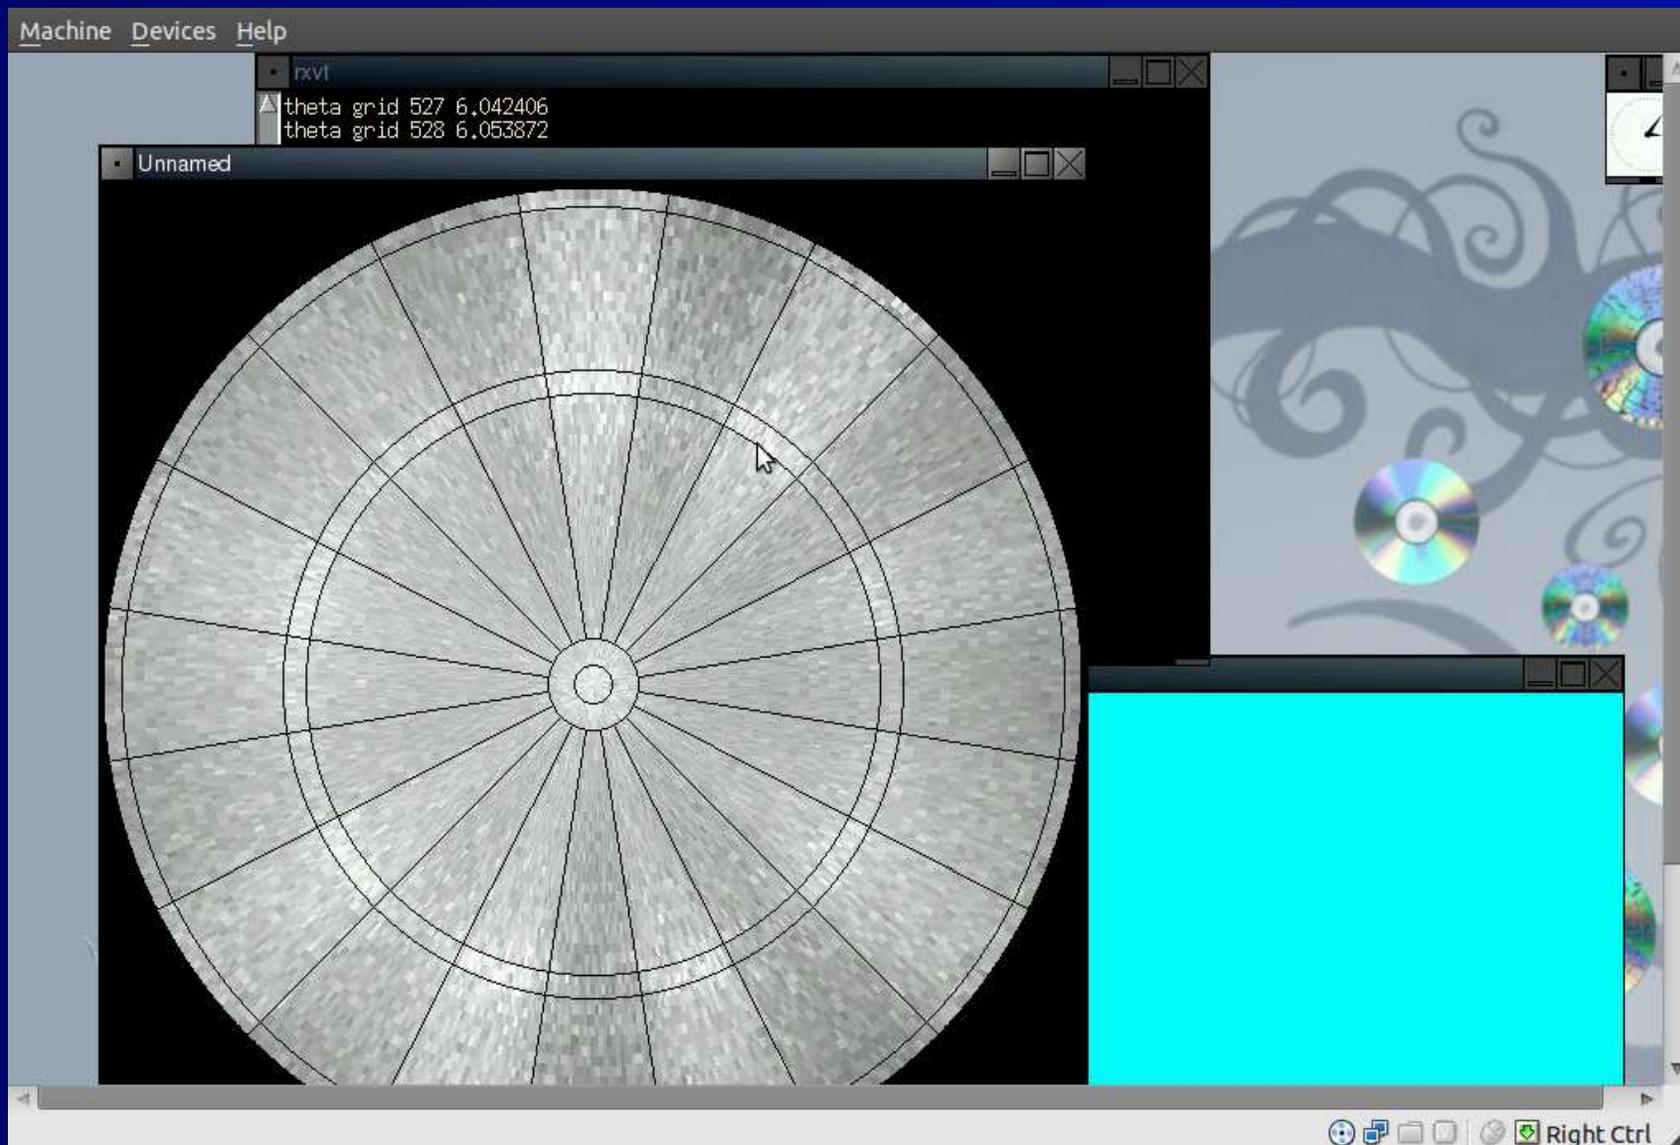
LittleFe Picture

More information

[Live Demo!](#)

References

Questions (and answers)





References

Sponsors
Terminology
What is the BCCD?
BCCD early history
BCCD today
What is the problem?
What is the problem?
(cont'd)
What is the problem?
(cont'd)
How is BCCD an
answer?
General usage of the
BCCD
Parallel methods
BCCD can use
Visualizations with the
BCCD
Nitty-gritty details of
BCCD
LittleFe
LittleFe Picture
More information
Live Demo!
References
Questions (and
answers)

- [1] Andrew Fitz Gibbon, David A. Joiner, Henry Neeman, Charles Peck, and Skylar Thompson. Teaching high performance computing to undergraduate faculty and undergraduate students. In *Proceedings of the 2010 TeraGrid Conference, TG '10*, pages 7:1–7:7, New York, NY, USA, 2010. ACM.
- [2] Tom Murphy, Charlie Peck, and Paul Gray. Little-fe: A portable, educational pc cluster. *HPC Wire*, 14(42), October 2005.
- [3] Jeannette M. Wing. Computational thinking. *Commun. ACM*, 49:33–35, March 2006.



Questions (and answers)

Sponsors

Terminology

What is the BCCD?

BCCD early history

BCCD today

What is the problem?

What is the problem?

(cont'd)

What is the problem?

(cont'd)

How is BCCD an
answer?

General usage of the
BCCD

Parallel methods

BCCD can use

Visualizations with the
BCCD

Nitty-gritty details of
BCCD

LittleFe

LittleFe Picture

More information

Live Demo!

References

Questions (and
answers)