The Micro-Cluster Showcase: 7 Inexpensive Beowulf Clusters for Teaching PDC

Joel C. Adams
(Organizer) Calvin College
adams@calvin.edu

Jacob Caswell
St. Olaf College
caswell@stolaf.edu

Suzanne J. Matthews
U.S. Military Academy
suzanne.matthews@usma.edu

Charles Peck
Earlham College
charliep@cs.earlham.edu

Elizabeth Shoop
Macalester College
shoop@macalester.edu

David Toth
Centre College
david.toth@centre.edu

James Wolfer
Indiana University, S. Bend
jwolfer@iusb.edu

CCS Concepts

• **Computing methodologies** ~ Parallel computing methodologies • **Computer systems organization** ~ Distributed architectures • **Social and professional topics** ~ Computer science education

Keywords

Architecture; Beowulf; clusters; computer science; computing; distributed; education; GPU; heterogeneous; hardware; inexpensive; multicore; parallel; processes; teaching; threads.

1. SUMMARY

Just as a micro-computer is a personal, portable computer, a *micro-cluster* is a personal, portable, Beowulf cluster. In this special session, six cluster designers will bring and demonstrate micro-clusters they have built using inexpensive single-board computers (SBCs). The educators will describe how they have used their clusters to provide their students with hands-on experience using the shared-memory, distributed-memory, and heterogeneous computing paradigms, and thus achieve the parallel and distributed computing (PDC) objectives of CS 2013 [1].

2. SESSION OBJECTIVE

Most CS educators are aware of the *Raspberry Pi* [8], but many have never heard of more powerful SBCs, such as Adapteva's *Parallella* [3], the Nvidia *Jetson* [6], or the *ODROID C1+* or *XU4* [5]. In this session, our primary objective is to give CS educators the opportunity to:

1. See clusters that their peers have built using SBCs.
2. Hear how their peers are using their clusters to teach PDC.
3. Interact with those peers on a personal basis.

A secondary objective is to inspire other CS educators to create their own clusters. By showcasing these clusters and exploring how they are being using in the classroom, we hope to catalyze a wave of innovation in PDC education.

SIGCSE '16, March 2–5, 2016, Memphis, TN, USA.
ACM 978-1-4503-3685-7/16/03.
http://dx.doi.org/10.1145/2839509.2844670

3. SESSION OUTLINE

Some SBCs are as small as a credit card, allowing our presenters to bring their clusters to the conference. Prior to the session, the presenters will set up their clusters for display and demonstration on tables at the front of the room.

In the first 32 minutes of the session, each presenter will give a 5-minute, 5-slide "lightning talk" on his or her cluster, describing its specifications and how they use it to teach PDC.

The middle 18 minutes will be a "Group Q&A" in which audience members can direct questions to the presenters as a group.

The last 25 minutes will be a "Show&Tell" session in which the presenters will move to their clusters and audience members can approach and interact with the presenters individually.

Table 1 presents a concise overview of the session.

What	Who	Duration
Lightning Talks	Joel Adams	2 min.
- *StudentParallella*	Suzanne Matthews	5 min.
- *PIsToGo*	Jacob Caswell	5 min.
- *HSC-5* and *HSC-6*	David Toth	5 min.
- *Rosie*	Elizabeth Shoop	5 min.
- *LittleFe*	Charlie Peck	5 min.
- *Cu-T-Pi*	James Wolfer	5 min.
Group Q&A	Audience	15 min.
Show&Tell	Everyone	28 min.

Table 1. Session Overview

This 75-minute session thus has three parts: lightning talks, an all-group Q&A, and a Show&Tell.

In the remainder of this section, we provide descriptions of our presenters, the clusters they will present, and how they use those clusters to teach PDC topics. For each of these clusters, the total cost includes everything needed to build the cluster (Gigabit Ethernet interconnect, parts, shipping, etc.), except for a monitor, mouse, and keyboard. The presenters' slides and the disk images for their clusters will be archived at *csinparallel.org* [4].

3.1 Joel Adams (Organizer)

Dr. Adams is a professor of computer science at Calvin College. He has designed and built four Beowulf clusters, including Microwulf [2], a portable sub-$2500 system that was the first cluster to break the $100/GFLOP barrier. He will introduce each presenter and moderate the session.

3.2 Suzanne Matthews

Dr. Matthews is an assistant professor of computer science at the United States Military Academy. To incorporate parallel concepts into her department's curriculum, she has created two courses: (i) an elective *Parallel Computing* course that exposes students to a variety of parallel libraries and architectures, including Adapteva's Parallella, an 18-core SBC; and (ii) a required *Computer Organization* course that introduces students to concurrent and shared memory parallel concepts using the POSIX API and x86 architecture. Dr. Matthews also uses two SBC clusters (*StudentPi* and *StudentParallella*) to engage her students in undergraduate research. Both cost about $700 to build.

3.3 Jacob Caswell

Jacob is a junior CS and Physics major at St. Olaf College. He will present *PIsToGo*, a complete, five-node, passively cooled, Raspberry Pi-2 cluster-in-a-briefcase that cost $300. He will also describe *Remora*, a virtual cluster-on-a-flash-drive that can be run on any network of workstations for $3 per node.

3.4 David Toth

Dr. Toth is an assistant professor of computer science at Centre College. He will present his latest "Half-Shoebox Clusters," (*HSC-5* & *HSC-6*) which are portable clusters inexpensive enough for each student to purchase one as lab equipment.

- *HSC-5* has two ODROID C1+ nodes, each with an ARM quad core CPU, at a total cost of $150.
- *HSC-6* has two ODROID XU4 nodes, each with two ARM quad core CPUs and an 8 core MP6 GPU; total cost: $220.

Dr. Toth will discuss his use of HSC clusters to teach OpenMP, MPI, and OpenCL in his *Parallel Computing* course, comparing them to courses in which the students used shared equipment.

3.5 Elizabeth Shoop

Dr. Shoop is a professor of computer science at Macalester College. She will present *Rosie*, a cluster with six Nvidia Jetson-TK1 nodes. Each node has an ARM quad core CPU and a 192 CUDA core Kepler GPU. The nodes share a 500GB hard disk. The total cost was $1350. She will also describe her use of Rosie to teach shared-memory, distributed-memory, and GPU/CUDA computing concepts in two courses: *System Fundamentals* and *Parallel Computing*.

3.6 Charles Peck

Dr. Peck is a professor of computer science at Earlham College. He will present *LittleFe* [7], the original cluster-in-a-suitcase. LittleFe has six custom-built nodes. Each has a dual-core Intel CPU and a 16 CUDA-core Nvidia GPU. The total cost is $2500 (but free for *LittleFe Buildout* participants!), including a shock- and water-proof carrying case. He will discuss his use of LittleFe to teach shared-memory, distributed-memory, and GPU/CUDA concepts in his *Parallel & Distributed Computing* course.

3.7 James Wolfer

Dr. Wolfer is a professor of computer science at Indiana University, South Bend. He will present *CU-T-Pi*, a heterogeneous cluster consisting of one Nvidia Jetson-TK1 and four Raspberry Pi nodes [9]. The total cost of the system was about $500. He has used his cluster to illustrate communication overhead and heterogeneous benchmarking considerations in a *Parallel Programming* course. He has also used it generate random numbers for an *Operating Systems* class by harvesting entropy from a Geiger counter interface, during which it supported remote access by 31 students for sixty 24/7 days.

4. EXPECTATIONS

We believe that hands-on exercises in which students must apply abstract concepts improve their understanding of those concepts. However, the hardware needed for students to experience some PDC topics (e.g., scalability) has traditionally been expensive.

The clusters featured in this session range in price from $150 to $2500, making them much more affordable than traditional Beowulf clusters. People attending this session will thus learn about a range of clusters at a variety of prices, will learn how CS educators are using those clusters, and will have the opportunity to ask detailed questions to those educators.

As a result, this session will be of interest to anyone who is interested in PDC or is teaching a course that covers PDC topics. Such courses include: *Algorithms*, *Computer Architecture*, *Computer Organization*, *Graphics*, *Operating Systems*, *Parallel Computing, Software Engineering*, *System Fundamentals*, the *Senior Capstone* course, and others. This special session will thus be of interest to instructors from a wide variety of CS courses.

5. SPECIAL SESSION SUITABILITY

The presenters will bring, display, and demonstrate their clusters, making this 75-minute session a "show and tell" style session. During the first part, the audience will receive an overview of the clusters and how they are being used to teach PDC. The "Group Q&A" part will allow the audience to engage with the presenters in a format that lets everyone hear all questions and answers. In the "Show&Tell" part, presenters will concurrently demonstrate their clusters to interested audience members, who can move from cluster to cluster, view the clusters up close, and ask detailed questions. This session's structure will thus be quite different from that of a typical paper or a panel session.

6. REFERENCES

Parts of this work were supported by NSF DUE #1225739.

7. REFERENCES

[1] ACM/IEEE CS Joint Task Force. "Computer Science Curriculum 2013". Online. Accessed 2015-08-20. https://www.acm.org/education/CS2013-final-report.pdf.

[2] J. Adams and T. Brom. "Microwulf: A Beowulf Cluster for Every Desk", *Proc. of the 39th SIGCSE Technical Symposium*, March, 2008, p. 121-125, DOI= 10.1145/1352135.1352178.

[3] Adapteva. "Parallella". Online, accessed 2015-08-20: http://www.adapteva.com/parallella/.

[4] R. Brown, L. Shoop, J. Adams. "CSinParallel: Parallel Computing in the CS Curriculum". Online, accessed 2015-08-20. http://csinparallel.org.

[5] Hardkernel. "ODROID-XU4 and ODROID-C1+". Online, accessed 2015-08-20: http://hardkernel.com.

[6] Nvidia. "Jetson-Tk1". Online, accessed 2015-08-20: https://developer.nvidia.com/jetson-tk1.

[7] C. Peck, et al. "LittleFe: Parallel and Cluster Computing On The Move". Online, accessed 2015-08-20: http://littlefe.net.

[8] Raspberry Pi Foundation. "Raspberry Pi 2". Online, accessed 2015-08-20: https://www.raspberrypi.org/products/.

[9] J. Wolfer, "A Heterogeneous Supercomputer Model for High-Performance Parallel Computing Pedagogy", *IEEE Global Engineering Education Conference*, March, 2015.