

**INDEPENDENT PROJECT BASED LEARNING:
APPLYING KNOWLEDGE BY CREATING A LOW-COST SUPERCOMPUTER**

Johnathan Yerby, Middle Georgia State College, johnathan.yerby@mga.edu

ABSTRACT

This article was written to detail the process of one-on-one project based learning and the benefits to the student. Students are increasingly seeking opportunities to gain experience with developing technologies that interest them personally, but may be beyond the planned curriculum. This paper explores a single independent project of a student creating a supercomputer by using the parallel processing power of ten Raspberry pi devices. The value of this paper lies in the interpretations of how independent projects may be used for students to gain outside of the existing curriculum to increase their cognitive presence in the field, how to foster an inclusive learning community for the immediate society, provide a means for students to earn course credit while not incurring cost for the college, and lastly this specific project of creating a super-computer that was named Celer-Pi.

Keywords: Project-based learning, Information Technology, Supercomputing, Qualitative

INTRODUCTION

What a student learns and what an instructor teaches is fairly well defined through committees, meetings, and approved curriculum to ensure that an academic program is meeting the goals of the institution. Meeting goals, missions, accreditation standards, retain and progress students, are very important outcomes of a university, but there is room to do all of these important things, while finding independent projects that may spark a student's creativity and foster their desire to know more than the existing list of courses. Philanthropist and founder of "Pencils for Promise," Adam Braun, once said "There are many challenges in the global education ecosystem: from top-down systemic issues in how educational services are organized and delivered, to bottom-up issues of curriculum effectiveness, accountability, and human resource allocation." [7]. The issue explored in this paper does not extend to the magnitude of challenging the global education system that Braun spoke about, but it does offer a set of ideas to think about alternatives to the way that college curriculum must be completed. The point that Braun makes in his message is that education should be designed for purposes of each person maximizing value for themselves and others.

The Celer-Pi project was started by a student that was enrolled in a Windows Systems Administration course and the desire to make a \$35 computer do much more than it was intended to do. The student contacted his instructor to discuss how the Enterprise version of Microsoft Windows Server Enterprise edition would be able to best leverage multiple Raspberry-Pi units to make them work together. After a short conversation, the idea to use Windows Server was abandoned in favor of other open-source options that would fit the needs of the student much better.

This study used ten Raspberry Pi computers that cost \$35 each to create a low cost supercomputer capable of distributing the processing and storage across the devices. The Raspberry Pi is a device the size of a credit card that plugs into a television and keyboard. It is designed to perform simple tasks such as run a word processor, spreadsheets, or games that are usually designed to run on an Android smartphone. The creators of Raspberry Pi intentionally made the device very hackable by providing the source code and several resources on their website to encourage users to think beyond the initial design [8]. The cost and time required to create high performance parallel computing devices is one of the largest obstacles that prevent educational institutions from taking on projects like this one [6]. To mitigate expenses and find areas to use tools most suitable for this scenario, free open source software such as OS Debian, Wheezy, and Raspbian were used. NginX was used to manage bandwidth and latency. MPICH2 was used to manage the cluster, coordinate the processing and distribute storage. An external hard drive, switch, router, cabling, power supply, and a custom built case for the creation were the remaining costs. The Raspberry Pi Supercomputer, which was named Celer-Pi, had a total cost of \$900. The name Celer-Pi was created from the Latin word "celer" which means fast, combined with "Pi" from the Raspberry Pi device. One of the lowest cost supercomputers available for retail starts at \$500,000. Although the processing and storage capabilities of the

\$500,000 machine from Cray exceeds the Celer-Pi, the concept and operations are comparable [3]. The creation reached performance of at least 900MHz while running a data-driven website [2].

GOALS

The goal of this study was to examine the benefits and challenges to conducting courses in a one-on-one student to teacher setting in an information technology related degree program. The study explored how project-based learning can work within the program while also describing how this model can be economically feasible for students and the university. Project based learning can be an effective way to engage students beyond the predefined curriculum and find additional pathways towards retention, progression, and graduation.

Project-Based Learning

Real world problems that learners are personally interested in or intrinsically motivated to learn about are great opportunities to allow students to learn and apply new knowledge to complex tasks. Project-based learning (PBL) is a model that allows students to explore, inquire, design, and make unique decisions to create their best possible solution. PBL uses authentic driving questions, a community of inquiry, and the use of technological based tools to create the atmosphere where students are more engaged and have the possibility to learn more than a traditionally structured course [4]. Constructivism plays an important role in independent project based learning that was used in the Raspberry Pi Supercomputer project because it allowed the student to use prior knowledge learned through his previous course-work and his experiences using technology and apply what he already knew to the complex challenge of creating a machine capable of parallel processing and storage. Constructivism and PBL allow learners to learn new basic skills that fill in gaps or tie together their prior knowledge and experiences [1]. Project-based Learning affords a unique opportunity where the student that is intrinsically motivated to work autonomously over extended periods of time and have frequent but short directed meetings with the instructor overseeing the student's teaching [11].

In this project, the student, we will give him a pseudonym of Frank Simon, asked the professor for some real world advice for a project that he intended to do out of curiosity. Simon was able and willing to fund the project to purchase what eventually came to a total of \$900. The instructor recognized that this would be a fascinating project and he had the most important component to make this a success, the student's interest and dedication. In the Celer-Pi project the student was interested in creating a machine capable of spreading the processing instructions across multiple devices. The student found a 2012 article where a team of engineers at the University of Southampton, UK built a supercomputer called, "Iridis-Pi" for approximately \$4031.00 plus switches and cables. The Southampton project utilized 64 Raspberry Pi devices and Lego building blocks to create the [9, 14].

One of the challenges of universities engaging in independent study or special project courses is the required resources from the instructor's time and the access to materials that may be needed for the project. Usually a special project involves only a few to as little as a single student, which is not cost effective for a university to support. Project-based Learning affords a unique opportunity where the student that is intrinsically motivated to work autonomously over extended periods of time and have frequent but short directed meetings with the instructor overseeing the student's learning. [11].

Inclusive Learning

Research shows that there are many different ways for people to learn and those different methods of teaching have varying levels of success [5]. Students in the Bachelor of Science in Information Technology at this four-year institution in the southeastern United States are consistently seeking additional hands-on activities and real-world problems. Students are exposed to traditional lectures, group projects, and discussions of varying sizes, but it is not as common for a student to be able to work one-on-one with their professor on a large scale project for many weeks due to the student to teach ratio and time requirements.

Creating real-world problems is often a very time consuming task with a large number of risks and challenges. The degree of difficulty and relation to a specific course's objectives do not align perfectly, nor do real-world challenges

align perfectly to a 16 week semester. Professors are often focusing on not only teaching their courses, but also engaged in scholarly activities and service. Trying to add an additional class with a group of ten to twenty students becomes as much of a responsibility as teaching another course, and can take much more time to do site visits, consult with students, and project sponsors. Given all the restraints, we must not give up on providing students the hands-on learning opportunities afforded by completing a real world task, we must simply find another method that could work better for some instructors.

This paper that began as a single instructor working with a single student on a very specific issue eventually grew and created a plethora of new ideas. Other instructors working in related fields wanted to contribute and learn more about this project. There are two instructors that are currently working on creating a different version of similar projects where they will create a Beowulf cluster by using salvaged computer parts in lieu of using the Raspberry Pi design that the project created [10]. Other students in the information technology program are interested in working with the professor and the student that led this project because of the overwhelmingly positive experience from Frank Simon's point of view. Frank was able to get two presentations and a publication to go along with his new upstart business. Frank has continued to work on increasing the capabilities of the Celer-Pi; all while remaining at a low cost and openly sharing the methodology with learners that would like to become more involved.

Celer-Pi Creation

The idea of using Raspberry Pi devices for more than their intended use of being a stand-alone inexpensive computing device is very popular right now as shown in the weblog 40+ Cool Ideas for your Raspberry Pi Project [12]. Using the Pi devices would allow processing and storage to be spread across nine devices with one of the Pi devices serving as the master to coordinate the computing. The Pi units were purchased from Allied Electronics and the materials for the rack came from Home Depot. There were other supplies needed to complete the build including a switch, cabling, memory cards, and some basic hardware to complete the build, which were purchased at Newegg.com, as detailed in Table 1.

Table 1. Total Project Expense

Item Description	Quantity	Cost each
50-Pack RJ-45	1	3
Adafruit 16x2 LCD	1	20
Cisco SR2024CT 24 port 10/100/1000 Compact Switch	1	178
Cisco-Linksys WRT54G Wireless-G Router	1	40
Ethernet Cable CAT6 - 100FT	1	9
Static IP Address Lease	1	15
Raspberry Pi Model B Revision 2.0 (512MB)	10	35
Sony 32GB SDHC/SDXC Memory Card	10	15
Hardware for custom built rack		135
	Total	\$900

New parallel computing techniques are constantly being developed by skilled programmers using proprietary and Open Source material. According to Amdahl law, software must be applied in order to gain advancements in parallel computing [13]. Multi-processing software is currently being engineered with MPICH2, IBM, MIT, Hadoop, Cuda, and Debian/Linux systems.

Applied knowledge is only correct if the system researched is understood and correct standards are applied. Due to this great technological separation, educated users must continue to cipher through material that supports the validity of certain tasks they are trying to accomplish with computers. Being able to reference valuable material and deciding cognitively why it is needed is essential to implementing the correct process. Applied knowledge is then correctly administered using commands initiated in standard operating procedures format.

The Raspbian software chosen has direct support from the Debian Wheezy library. Apache2 and Nginx were both considered and implemented into the Celer-Pi. Apache2 contains more modules than Nginx yet, both work well with

large-scale clusters. Apache is seen more in major transactions and Nginx is light on resources when dealing with quick small file transfers. The main development computer had sufficient 300GB SSD memory and Duo AMD Turion X2 64 processing capability and received Apache2.

To respond and deliver information from client requests Nginx 1.4 server software was implemented into the Pi cluster. The web server is lightweight yet robust containing powerful code that operates popular sites such as Facebook, Hulu, Pinterest, Dropbox, and Word Press to name a few. Since Nginx is a relatively newcomer in the market the web servers software requires knowledge to be obtained daily and immediate implementation of security of data applied.

Open SSH version 6.0 was initialized to place the main computer into master mode and allow connection to the network. SSH applies proven encryption methods transferring data securely and protecting the integrity data. Open SSL was installed to use the cryptography tool kit manager on the Session layer of the OSI model. SSL sits between http and TCP providing security protection. The software provides many key cryptographic functions such as Time Stamp requests that support generation and verification.

The distributed file system is accomplished by the MPICH2 version of the MPI standard. MPICH2 was chosen for its interoperability with a wide range of devices. It is a high performance and lightweight alternative to the original MPICH software build from National Argonne Labs. MPICH2 software adds capability of controlling multiple processors from one master location through the combined use of cluster architecture and networked nodes. Using MPICH2 software the total environment can be controlled by the master Pi or independent environments can be controlled separately.

Once the initial build configuration was complete, initiation of communications using MPICH2 build data was calculated based on number of subsystems. Using the master Pi a calculation command was then sent to the subordinate Pi systems. Data received at the subordinate Pi cluster is included instructions for portions of data sets to calculate and send results back to the master Pi. The master Pi receives information from the ongoing processes of each sub machine and displays the data.




Figure 1. Cluster design




Figure 2. Celer-Pi completed

The system provides the ability to develop, administer, and secure an online website for the recurring cost of a \$12 domain name and \$15 static IP lease each year. Each Pi device is able to support 8-10 websites with PHP5, MySQL, and email servers in the cluster. The Celer-Pi, shown above in Figure 2, has a small physical footprint, uses a small amount of power, and does not require additional cooling beyond that of a normal office environment. The design of the Celer-Pi allows for scalability to add additional Pi devices with minimal configuration, as well as the ability to quickly and inexpensively replace any single Pi device that may fail.

At the hardware level each computer system comes complete with its own Mac/IP, hardware, core, Solid Drive, and Ethernet 10/100 network connectivity. Only the master Pi, as shown in Figure 1, is able to reach the outside network. Clients are allowed to communicate with the master Pi through http port 80. The independent environments can only be accessed through direct communication from the master Pi. The number of users that can be responded to or data that can be analyzed by the master is limited only by the bandwidth and the amount of solid-state memory available to the system.

The master Pi to Pi stack relationship allows for the separation of MySQL, PHP5, Nginx, and email servers on individual Pi. This relationship increases response and up time in websites through increased availability of resources and data communications, effectively removing single bottleneck points of failure.

The current system setup allows for quick addition of new Pi and Solid Drive expansion to the distributed file system network. Main memory and backups can be safely stored off the network using Nginx and Solid State Memory. If any computer in the cluster or even the whole cluster crashes Wheezy software images, Nginx web servers, and user data from MySQL will be refreshed using offsite Solid Drive card backups from the main computer. The complete Pi system could be reformatted in 10-15 minutes.

Economics

Conducting one-on-one independent projects can be cost prohibitive due to the amount of time and resources that could be needed. In this case we were able to overcome financial obstacles in the following ways. The student agreed to treat this independent project as an upper level information technology elective credit. Frank Simon would pay tuition and agree that there would be at least 50 contact hours to ensure that there would be no issues with the project counting as a course. Treating the project as a course for the student not only allowed him to learn more about a real-world problem that he was interested in exploring, but it also had a positive impact on progression as it moved him one course closer to graduation.

The college developed a new model to encourage instructors to take on these types of independent learning opportunities. The economic model provides an instructor a single payment of \$3000, after the professor has completed 15 special projects or internships, which equals out to \$200 per project. The flexibility of the independent project based-learning model allows professors to work with multiple students on projects that do not always have to directly align with the college calendar.

The student had previously purchased five of the \$35 Raspberry Pi units and was looking for a way to create a supercomputer without spending much more money. The total cost eventually reached \$900, but the student was able to fund the entire hardware expenses. To recoup some of his expenses, Frank was able to immediately begin using the newly created Super-computer, Celer-Pi, to host websites for his pressure washing business as well as building and hosting websites for other local businesses [2].

CHALLENGES AND LIMITATIONS

A static IP was leased to allow any user to find the network that is browsing the internet and wants to view a site being hosted by the Celer-Pi. The router allows for the look up of the master Pi that then finds the subordinate Pi addresses on the personal network and requests data. Under the current hardware setup the Cisco switch can support up to 24 connections simultaneously to individual Raspberry Pi computers. The Pi devices are mounted to the 4mm all thread wire with added plastic spacers separating each individual Pi. This is accomplished using pre-drilled holes from the factory. Each Pi in the cluster is therefore suspended allowing for greater heat exhaustion if needed.

Time and collaboration was an issue addressed by setting scheduled Skype sessions during the evenings and weekends when the instructor and student could both dedicate one to two hours at a time to meet. In total there were about 25 hours of video conferencing spread across 25 weeks. The project spanned two semesters, but the student only registered for the course during the second semester of the project. There were also 20-30 short in person meetings that would last 15-45 minutes each to discuss the project and determine next steps. The project garnered the attention from other instructors within the school that would also provide guidance and input during the frequent short meetings. Frank, the student, would take the responsibility to implement each set of changes and report his

results back. The professor acting as a sponsor and mentor for this project benefited in multiple ways including; adding this to one of the special projects courses completed and scholarship by publishing a poster session as research in progress, then a paper detailing the technical aspects of the Celer-Pi, and finally a paper explaining the pedagogical benefits of projects like this. This project started a group of students that took a deeper interest in a topic that would have otherwise not have been in their curriculum. The sponsor and other professors that contributed to the project had the benefit of applying concepts to a project that they would have otherwise not have taken on. Overall this project has been an exceptional example of how independent real-world problems can benefit many people in several ways.

REFERENCES

1. Diehl, W., Grobe, T., Lopez, H., & Cabral, C. (1999). Project-based learning: A strategy for teaching and learning. *Boston: Center for youth development and education.*
2. AUTHOR CITATION
3. Gohring, N. (2013, May 7). Own your own cray supercomputer for a mere \$500,000. *CNN Money*. Available: <http://money.cnn.com/2013/05/07/technology/enterprise/cray-supercomputer/>.
4. Krajcik, J. S., Blumenfeld, P. C., Marx, R. W., & Soloway, E. (1994). A collaborative model for helping middle grade science teachers learn project-based instruction. *The elementary school journal*, 483-497.
5. Lage, M. J., Platt, G. J., & Treglia, M. (2000). Inverting the classroom: A gateway to creating an inclusive learning environment. *The Journal of Economic Education*, 31(1), 30-43.
6. Pounds, A. J., Nalluri, R., & Coleman, B. L. (2005). The development of a Tri-Use cluster for general computer education, high performance computing education, and computationally intensive research. *ACM Southeast Regional Conference*, 345. doi:10.1145/1167350.1167446.
7. Raman, R. (2011, March 19). Adam Braun Makes a Penciled Promise to Children. *Raman Media Network*.
8. Raspberry Pi Foundation. (2014, January). *Faqs - Raspberry Pi*. Available: <http://www.raspberrypi.org/faqs>.
9. Spector, D. (2000, December 1). Building Your Own Beowulf Cluster. *Wired Magazine*, 8.12.
10. Thomas, J. W. (2000). A review of research on project-based learning. Retrieved from http://www.ri.net/middletown/mef/linksresources/documents/researchreviewPBL_070226.pdf.
11. Tom, R. (2012, December 3). [Web log message]. Retrieved from <http://pingbin.com/2012/12/30-cool-ideas-Raspberry-Pi-project/>.