Engineering Educators Bringing the World Together
2025 ASEE Annual Conference & Exposition
Palais des congrès de Montréal, Montréal, QC • June 22–25, 2025
Paper ID #49430

BOARD # 65: Bring Your Own Cluster to the Classroom

Dr. Chafic Bousaba, Guilford College

* Joined Guilford College in January 2008 * Serves as Associate Professor in the Computing Technology and information Systems.

Bring Your Own Cluster to the Classroom (BYOCC): Enhancing Learning Through Raspberry Pi 5 Cluster Computing

Abstract

Bringing and utilizing innovative technology solutions in the classroom plays a crucial role in enhancing the learning experience, applying theoretical knowledge, and providing students with significant hands-on practice. While the concept of Bring Your Own Device (BYOD) or Bring Your Own Technology (BYOT) has been widely implemented, it has predominantly focused on personal devices for work-related tasks. In contrast, cluster computing, a technology gaining momentum among developers, researchers, and data scientists, is often impractical to implement in classroom settings due to its resource-intensive nature. This paper introduces the pedagogical approach of Bring Your Own Cluster to the Classroom (BYOCC), which combines the portability and affordability of personal devices with the functionality of cluster computing, offering an innovative learning solution.

Specifically, this paper explores the application of BYOCC through the use of Raspberry Pi clusters, which enable students to gain practical experience in cloud computing, cybersecurity, and current IT trends. The study compares two distinct Raspberry Pi 5 cluster architectures, detailing their building process, their use cases, technologies, and classroom applications. The first architecture utilizes the Turing Pi 2 Board, a compact 4-node ARM cluster board that incorporates a built-in Ethernet switch, offering a secure and scalable solution for edge computing. The second architecture implements a four-node high availability cluster using Docker swarm and a single managed switch. The comparison of these two architectures highlights their benefits in delivering a comprehensive, hands-on learning experience for students, fostering deeper engagement with key concepts in computing and IT infrastructure.

By exploring these cluster architectures, this paper demonstrates the potential of BYOCC in making cluster computing accessible and practical for educational environments, promoting technical learning while encouraging innovation in cloud technologies and IT practices.

1. Introduction

The integration of technology in education has transformed traditional, outdated learning environments, to provide more interactive, hands-on, and engaging experiences for students. Innovative solutions such as Bring Your Own Device (BYOD) and Bring Your Own Technology (BYOT) have become the norm in modern classrooms. These strategies encourage students to bring their personal devices to school, and primarily use them for work, communication, and learning purposes. However, the concept of utilizing computing clusters which are essential for cloud computing, container orchestration, data science, and high-performance computing (HPC) [1][2], is still relatively underexplored in educational settings, despite its growing relevance and rapid growth in today's technology.

Some of the challenges associated with teaching cluster computing include its typically high costs and resource demands, which can make it an impractical solution for many classrooms. However, this paper introduces the concept of Bring Your Own Cluster to the Classroom (BYOCC) as a viable and innovative approach to teaching cloud computing, networking, and system administration skills. The BYOCC approach could be incorporated into the Operating Systems or Computer Networks courses, which are traditionally core junior-level courses for Computer Science majors and other related Information Technology fields. By leveraging affordable and portable devices like the Raspberry Pi 5 [3], a popular low-cost, low-power single-board computer (SBC), students can create and manage their own clusters or join Pi clusters, enabling them to explore foundational concepts in cluster computing, distributed systems, and current IT trends.

This paper examines the application of BYOCC in classroom environments, focusing on the use of Raspberry Pi 5 clusters [4] to teach essential technologies such as cloud computing, container orchestration, cybersecurity, and IT infrastructure management. This paper presents two distinct Raspberry Pi cluster architectures, each offering unique advantages and challenges for students. These architectures are designed to introduce students to the world of distributed computing while allowing them to interact with real-world technologies in a low-cost, hands-on manner.

2. Background: Cluster Computing and Its Role in Education

2.1 Cluster Computing Overview

Cluster computing is a complex and valuable technology for developing large-scale, distributed systems. While traditionally associated with large data centers, there are many accessible tools, platforms, and resources that enable individuals and educational institutions to engage with cluster computing [1][2]. In simpler terms, cluster computing refers to the practice of connecting multiple computing devices together to form and work as a unified system [1]. The advantages of cluster computing include enhanced processing power, scalability, redundancy, and fault tolerance. It is widely used in fields such as cloud computing, scientific simulations, data analysis, and high-performance computing.

In educational contexts, cluster computing offers a unique opportunity to engage students with complex concepts such as parallel computing, load balancing, and distributed databases. These concepts are integral to modern IT infrastructure, and hands-on experience with cluster systems is invaluable for the students who are pursuing careers in fields such as data science, cloud architecture, and cybersecurity.

2.2 Challenges of Implementing Cluster Computing in the Classroom

While cluster computing provides essential skills for students, its implementation in classrooms has historically been difficult. The primary obstacles include:

- Cost: Traditional cluster computing systems often require expensive hardware, such as high-performance servers and dedicated networking infrastructure.

- Space and Power Consumption: Clusters typically need substantial physical space, energy consumption, and specialized cooling and ventilation, making them impractical for small classrooms or environments with limited resources.
- Technical Expertise: Setting up and managing clusters requires substantial technical knowledge, both for instructors and students, making it a challenge for instructors to design and maintain cluster-based learning environments.

Despite these challenges, recent developments in affordable hardware and cloud-based services have made cluster computing more accessible to educational institutions. The Raspberry Pi 5, for example, provides a compact, low-cost, and energy-efficient alternative to traditional servers [3], making it an ideal platform for cluster computing in classrooms.

3. Introducing Bring Your Own Cluster to the Classroom (BYOCC)

3.1 The Concept of BYOCC

The BYOCC model encourages students to bring their own Raspberry Pis (preferably Pi 4 or Pi 5 with 8 GB of RAM) to the classroom and set up their own clusters. Students can also bring their Pis to join a Pi cluster as nodes and disjoin the cluster when done. This approach makes cluster computing more accessible by taking advantage of personal computing devices, which students already own or can acquire affordably. By leveraging the Raspberry Pi 5 capabilities, students can simulate the environment of a full-scale cluster, learning valuable skills in cloud technologies, system administration, and cybersecurity in a hands-on manner. The Pi cluster also requires a managed switch with a minimum of 4 ports which is also very affordable and portable. The BYOCC model emphasizes affordability, portability, and flexibility while fostering self-sufficiency for the students. By allowing students to build and manage their clusters, this approach enables a deeper understanding of theoretical concepts through real-world applications.

3.2 Benefits of BYOCC in Education

- **Affordability**: Raspberry Pi devices are affordable (typically Raspberry Pi 4 and 5 with 8 GB of RAM are priced between $70 to $100), making them a practical choice for classroom environments with limited budgets. While Raspberry Pi with lower RAMs are obviously cheaper. (Typically, Raspberry Pi 4 and 5 with 4 GB of RAM are priced between $60 to $70).
- **Portability**: The small form factor of Raspberry Pi devices allows students to transport and set up their clusters easily, making it possible to engage in practical exercises both inside and outside the classroom.
  The Raspberry Pi 5 dimensions are: 3.35 (Length) x 2.2 (Width) x 0.67 (Height) inches.
- **Scalability and Flexibility**: The modularity of the Raspberry Pi 5 or the Turing Pi 2.0 boards, which are the key component of our Pi clusters, allows educators to start small and scale up as the class progresses. New modules can be added to the cluster as needed, and different configurations can be tested. This process mimics cloud environments where resources are dynamically allocated or de-allocated, providing a learning opportunity on how cloud infrastructures work. Thus, the scalability of the Pi clusters provides a flexible learning environment.

- **Cost-Effective Learning**: Unlike traditional expensive server setups or cloud services, Raspberry Pi clusters offer a low-cost solution for setting up a multi-node environment. In our case, to alleviate the financial burden of the student, the computer science department supplied the students with all the necessary hardware to build and maintain the Pi clusters. Educational institutions can build these clusters for a fraction of the cost of traditional servers, making it accessible to a broader range of students. This setup can be used as a permanent lab environment, where it can be maintained and reused across multiple years of classes and even in various classes. Thus, the Pi clusters yield an efficient return-on-investment for both the educational institution and the students.
- **Hands-On Learning**: The BYOCC model allows students to create and configure their own clusters, gaining valuable experience in distributed computing, network administration, and system troubleshooting. The students practiced with real-world tools and used modern software stacks such as Docker [5], Kubernetes [6], and containerization in a collaborative and realistic setting. The students also learned valuable troubleshooting and problem-solving skills which encourages collaborative learning, as the students worked together to solve issues within their clusters. The students used technologies they might encounter in their future careers, which makes this learning experience much more relevant.
- **Job Marketability**: developing practical and desired IT skills such as setting up, managing, and scaling a cluster seamlessly aligns with key skills in high demand in the job market, especially in fields like cloud engineering, DevOps, data science, Artificial Intelligence/Machine Learning (AI/ML), distributed systems, and big data analytics. Many companies, especially in technology, finance, healthcare, and AI, are leveraging cloud computing and distributed architectures. Students who practiced and worked with these technologies will be ahead of the curve when entering the workforce.

4. Raspberry Pi Cluster Architectures for the Classroom

4.1 Turing Pi Board 2.0 Cluster Architecture

The Turing Pi 2.0 [7] is a compact 4-node ARM-based cluster board designed to facilitate cluster computing in small environments. This board incorporates a built-in Ethernet switch, simplifying the networking setup required for a cluster. Each node on the Turing Pi 2.0 runs on a Compute Module 4 (CM4) or Nvidia Jetson compute modules in any combination. The Nvidia Jetson was not considered due to its high cost compared to the CM4. The CM4 features a powerful ARM Cortex-A72 processor, multiple RAM options, and support for external storage. Our Turing Pi 2.0 cluster used 4 CM4s, with 8 GB RAM and 32GB eMMC Flash memory.
This architecture offers the following advantages:

- **Scalability**: The Turing Pi 2.0 allows for easy expansion of the cluster by adding more nodes, making it an excellent option for students to explore scalable cloud computing environments.
- **Edge Computing**: The architecture is particularly suited for exploring edge computing, where computational tasks are offloaded from centralized data centers to localized devices that are placed at the “edge” of the network.

- **Secure and Manageable**: The built-in Ethernet switch simplifies networking, while the ARM-based architecture provides security features suitable for learning about secure computing practices in distributed environments.

Turing Pi Cluster: Hardware Setup

- **Turing Pi 2.0 Board**: The Turing Pi 2.0 board acts as a central platform to house up to 4 Raspberry Pi CM4s. It also includes built-in Gigabit Ethernet and power distribution, allowing multiple Pi units to operate on a single power supply.
- **Raspberry Pi Compute Module 4 (CM4)**: The CM4 comes with a quad-core ARM Cortex-A72 CPU, up to 8GB of RAM, and an optional eMMC storage or microSD card. Unlike standard Raspberry Pi boards, the CM4 is designed for embedded applications, offering a more compact and flexible architecture for clusters.
- **Power Supply**: The Turing Pi 2.0 requires a single 12V power supply, which powers all connected CM4 nodes through the board's power distribution system.
- **Cluster Case**: A custom-designed cluster case for the Turing Pi 2.0 can house multiple CM4s and includes provisions for passive cooling or active fans.

Our Turing Pi 2.0 cluster, see Figure 1, is securely housed inside of a small Dell desktop case, uses the desktop's power supply and cables, and the desktop's hard drives for storage.

Turing Pi Cluster: Software Setup

The Turing Pi 2.0 cluster can run K3s [8], Docker [5], and other distributed computing tools. The Turing Pi 2.0 provides an environment where each CM4 can run an individual operating system instance, and software management is handled through Kubernetes or similar orchestration tools. The architecture is suitable for learning and deploying containerized applications at scale, including those using microservices and distributed databases.

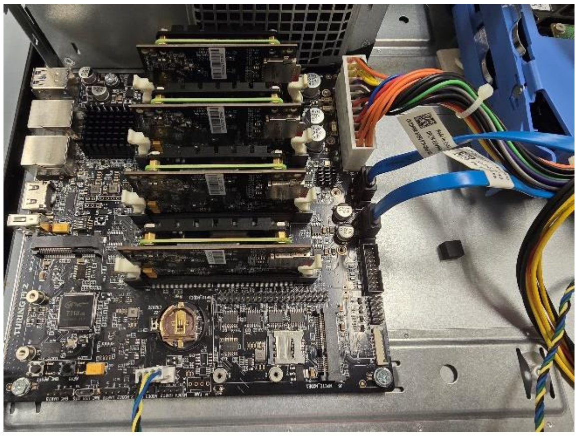

Figure 1 – Our Turing Pi 2.0 Cluster

4.2 High Availability Cluster with Docker Swarm

The second cluster architecture utilizes a three-node high availability (HA) cluster set up using Docker Swarm [9][10], a container orchestration platform that enables easy deployment and management of containerized applications. In this setup, each Raspberry Pi node runs Docker and is connected to a managed Ethernet switch, allowing students to explore high availability, containerization, and fault tolerance.
Key features and benefits of this architecture include:

- **Containerization**: Students gain hands-on experience with Docker, one of the most popular containerization technologies used in modern IT infrastructures.
- **High Availability**: Docker Swarm enables students to configure a fault-tolerant system with automatic failover, providing experience with high availability and load balancing.

- **Networking**: The managed switch ensures reliable network connectivity between nodes, enabling students to explore distributed networking and container communication.

HA Pi 5 Cluster: Hardware Setup

- **Raspberry Pi 5 Nodes**: Each node in the cluster is a Raspberry Pi 5 board. Each comes with a quad-core ARM Cortex-A76 CPU, options for 4GB or 8GB of RAM, Gigabit Ethernet, and USB 3.0 ports, providing substantial performance for distributed tasks.
- **Ethernet Switch**: A Gigabit Ethernet switch connects all Raspberry Pi units to each other, facilitating communication between the nodes in the cluster.
- **Power Management**: Individual power supplies (or USB power hubs) are required for each Raspberry Pi 5 node.
- **Storage**: Each Raspberry Pi uses a microSD card or eMMC storage (optional) for the operating system and data storage.
- **Cluster Case**: A stackable cluster case holds and organizes the Raspberry Pi boards, providing ventilation and physical security.

Our HA Pi cluster, see Figure 2, is totally housed inside a clear cluster case with an ethernet switch and USB power switch. It also has the ability to add more Raspberry Pi 5 nodes to the cluster as needed.

HA Pi 5 Cluster: Software Setup

The Raspberry Pi 5 cluster can run a variety of **distributed systems software** such as **Kubernetes**, **Docker Swarm**, **MPI**, or **Hadoop** for parallel computing. The lightweight **K3s Kubernetes** distribution is particularly suited for the resource constraints of Raspberry Pi devices. The cluster setup typically involves installing an OS (e.g., **Raspberry Pi OS Lite**) on each Raspberry Pi, configuring network communication, and setting up the orchestration software to manage containerized workloads.

Figure 2 - High Availability Pi Cluster

5. Comparing the Two Architectures

Both the Turing Pi 2.0 cluster and the Docker Swarm-based high availability cluster offer unique advantages and use cases in the classroom. While the Turing Pi 2.0 is a more compact and scalable solution with built-in networking features, the Docker Swarm setup provides more flexibility in terms of containerized applications and fault tolerance. The choice of architecture depends on the specific learning objectives, with the Turing Pi 2.0 being more suitable for edge computing and the Docker Swarm cluster better suited for teaching modern containerization and IT infrastructure concepts.

5.1. Cost

- **Raspberry Pi 5 Cluster**:
  - **Advantages**: The Raspberry Pi 5 is inexpensive (around $70–$100 per unit), and using individual Raspberry Pi units means that additional components such as an Ethernet switch or a power switch can be scaled as needed. The Raspberry Pi is highly accessible for educational purposes, making it an affordable solution for students and educational institutions.
  - **Disadvantages**: The cost of additional components like a power hub, Ethernet switch, and cluster cases can add up, especially as the number of nodes increases. Each Raspberry Pi requires a separate microSD card or eMMC storage, which can increase total costs. The cost of the microSD card varies depending on the storage capacity and transfer speed.
- **Turing Pi Board 2.0 Cluster**:
  - **Advantages**: The Turing Pi 2.0 offers a highly integrated solution with 4 slots that can accept either the Raspberry Pi Compute Module 4 CM4 or Nvidia Jetson Compute Modules in any combination. The Pi CM4s may be more cost-effective for managing multiple nodes but the Nvidia Compute Modules cost more than $200 which is very prohibitive for our project. The power distribution and networking are handled by the Turing Pi board itself, reducing the need for additional networking components.
  - **Disadvantages**: The Turing Pi 2.0 board seems to be unavailable for sale but the newer Turing Pi 2.5 costs $279 [7], and each CM4 module costs approximately $45–$70 depending on the RAM and storage configuration. Thus, the total upfront cost could be higher for smaller-scale clusters compared to the Raspberry Pi 5 cluster.

5.2. Scalability

- **Raspberry Pi 5 Cluster**:
  - **Advantages**: The Raspberry Pi 5 cluster can be easily scaled up by adding more Raspberry Pi units and an additional Ethernet switch. This provides a flexible, modular approach to building clusters.
  - **Disadvantages**: The scalability of Raspberry Pi 5 clusters can become cumbersome as the number of nodes increases. For large-scale clusters, managing multiple power adapters, Ethernet cables, and network configurations may become difficult.
- **Turing Pi Board 2.0 Cluster**:
  - **Advantages**: The Turing Pi 2.0's modular design allows for easy addition of up to 4 CM4 units per board. For larger clusters, multiple Turing Pi 2.0 boards can be stacked, providing an efficient way to scale vertically while maintaining compactness and simplicity.
  - **Disadvantages**: While the architecture is scalable within a single board, scaling horizontally (i.e., adding more physical boards) requires more complex infrastructure. Moreover, each additional Turing Pi 2.0 board increases the total cost.

5.3. Performance

- **Raspberry Pi 5 Cluster**:
  - **Advantages**: The Raspberry Pi 5 offers a substantial performance boost over its predecessors, with a quad-core Cortex-A76 processor and up to 8GB RAM, making it suitable for most distributed applications, including web servers, data analytics, and machine learning.
  - **Disadvantages**: While the Raspberry Pi 5 is powerful, it still has limitations in terms of CPU and RAM when compared to larger, traditional server architectures. For highly parallel workloads, the performance of each node may be a bottleneck.
- **Turing Pi Board 2.0 Cluster**:
  - **Advantages**: The **CM4s** used in the Turing Pi 2.0 have similar performance to the Raspberry Pi 5 but benefit from better resource management and integration. The Turing Pi board centralizes networking and power distribution, improving overall cluster efficiency.
  - **Disadvantages**: The overall performance of the cluster may still be limited by the compute power of each CM4, particularly for memory- and CPU-intensive tasks, as the ARM cores are not as powerful as enterprise-grade servers or cloud instances.

5.4. Flexibility

- **Raspberry Pi 5 Cluster**:
  - **Advantages**: Raspberry Pi 5 nodes are more flexible in terms of hardware selection. Users can choose from various RAM and storage options, including using external drives or network storage.
  - **Disadvantages**: Managing multiple Raspberry Pi boards can introduce complexity in terms of networking and power management, especially in larger clusters.
- **Turing Pi Board 2.0 Cluster**:
  - **Advantages**: The Turing Pi 2.0 provides a more integrated solution, reducing the number of components required to build the cluster. Its compact design allows for easier management and can fit into spaces where individual Raspberry Pi units might be impractical.
  - **Disadvantages**: The use of CM4 modules limits flexibility compared to the Raspberry Pi 5 since the CM4's peripheral options (e.g., USB ports) are less accessible for projects requiring external hardware.

5.5 Curriculum Alignment, Teaching Use Cases and Applications

The Pi clusters were implemented as hands-on labs for a select group of students enrolled in the Operating Systems course, a junior-level core class for Computer Science and Cybersecurity majors. These students were eager to learn about high-availability cluster computing, container orchestration, and microservices architectures. Each team of students first learned about the hardware and software for their respective projects, then implemented the solutions (HA R-Pi cluster and Turing Pi CM4 cluster), and finally troubleshot and made improvements. These hands-on labs will be used by all students enrolled in the Operating Systems course during the Fall 2025 semester.

The Turing Pi 2.0 cluster is ideal for teaching students about edge computing, distributed computing, and scalable cloud environments. It enables the exploration of both hardware and software aspects of cluster management. On the other hand, the Docker Swarm cluster is focused on teaching container orchestration, high availability, and microservices architectures. It is well-suited for lessons on cloud computing, and fault-tolerant system design.
Here are some suggested classroom hands-on lab activities that are incremental in difficulty and complexity:

1. Students can start by setting up a Raspberry Pi Cluster or a Turing Pi 2.0 Cluster. They'll learn how to network multiple nodes, configure IP addresses, and set up shared storage.
2. Students can deploy Docker containers on the cluster and orchestrate them using Kubernetes. They can create a microservices-based architecture, deploy a simple web application, and scale it up or down based on resource needs.
3. Students can implement load balancing across the cluster using NGINX [11] or HAProxy [12]. Students can learn how to distribute traffic among nodes and ensure high availability for web applications.
4. Students can run distributed frameworks like Apache Spark [13] or Hadoop [14] to analyze large datasets across multiple Raspberry Pi nodes. Students can learn how to set up and manage distributed data processing pipelines.
5. Students can simulate cloud computing in the classroom by setting up a private cloud using OpenStack [15] or Docker Swarm [10]. This can demonstrate how workloads are distributed and managed in the cloud.

The suggested time for each of the above hands-on labs is between 3 and 6 hours, and students should be allowed to access their Pi clusters and work outside of designated lab hours if needed.

6. Conclusion

The Bring Your Own Cluster to the Classroom (BYOCC) model offers an innovative and practical solution for teaching cluster computing and related technologies in educational environments. By leveraging affordable and portable devices such as Raspberry Pi, students can gain hands-on experience with cloud-native and cutting-edge technologies, distributed systems, and IT infrastructure. The comparison of two Raspberry Pi cluster architectures, the Turing Pi 2 and a Docker Swarm high availability cluster, demonstrates the versatility and potential of the BYOCC model in supporting various learning outcomes. The Docker Swarm high availability Pi 5 cluster is a flexible, cost-effective solution for smaller clusters and provides a hands-on experience with individual nodes, whereas the Turing Pi Board 2.0 cluster offers a more integrated, scalable, and compact solution with improved power and network management. Essentially, the choice between the two cluster architectures depends on the specific educational goals, budget, and scalability needs of the classroom environment. Ultimately, BYOCC empowers students to develop critical technical skills while fostering deeper engagement with the concepts and practices driving modern computing. By implementing BYOCC, educators can make complex technologies more accessible and affordable, allowing students to engage with real-world challenges in a practical and meaningful way.

References

1. D. Kalavikatte, Cluster Computing, Grid Computing, Cloud and Virtualization, Kindle ed., DSK Publisher, 2020.
2. T. Sterling, M. Brodowicz, and M. Anderson, High Performance Computing: Modern Systems and Practices, 2nd ed., Morgan Kaufmann, 2024.
3. Raspberry Pi, "Raspberry Pi," https://www.raspberrypi.com/, accessed Jan. 19, 2025.
4. A. K. Dennis, Raspberry Pi Super Cluster, Packt Pub Ltd, 2013.
5. Docker, "Docker," https://www.docker.com/, accessed Jan. 19, 2025.
6. Kubernetes, "Kubernetes," https://kubernetes.io/, accessed Jan. 19, 2025.
7. Turing Pi, "Turing Pi," https://turingpi.com/, accessed Jan. 19, 2025.
8. K3s, "K3s," https://k3s.io/, accessed Jan. 19, 2025.
9. G. N. Schenker, The Ultimate Docker Container Book - Third Edition: Build, Test, Ship, and Run Containers with Docker and Kubernetes, 3rd ed., 2023.
10. BetterStack, "Scaling Docker: High Availability with Docker Swarm," https://betterstack.com/community/guides/scaling-docker/ha-docker-swarm/, accessed Jan. 19, 2025.
11. NGINX, "NGINX," https://nginx.org/ , accessed Jan. 19, 2025.
12. HAProxy, "HAProxy," https://www.haproxy.org/, accessed Jan. 19, 2025.
13. Apache Spark, "Apache Spark," https://spark.apache.org/, accessed Jan. 19, 2025.
14. Apache Hadoop, "Apache Hadoop," https://hadoop.apache.org/, accessed Jan. 19, 2025.
15. OpenStack, "OpenStack," https://www.openstack.org/, accessed Jan. 19, 2025.