ESTEEM
THE OU CENTRE FOR
STEM PEDAGOGY

The Open University

Teaching distributed computing using Raspberry Pi clusters at a distance
Final Report

Executive Summary

A Parallel and distributed computing (PDC) is now considered a threshold concept for computing, and is embedded in computing curricula across the globe. While the costs of traditional computing clusters have made developing practical activities challenging, the rise of low-cost computers, particularly the Raspberry Pi, has led to an exploration of how PDC can be taught to students using Raspberry Pi clusters. Building on this work, we report our experiences from developing a series of low-cost Raspberry Pi clusters for use with open distance university students. Based on survey results from 484 students, we argue that our work demonstrates the benefits that remote practical activities can have for teaching PDC concepts, as well as engaging students. We conclude with a discussion of two key challenges: supporting active learning through student-led programming on the clusters, and supporting lower-performing students at a distance.

Aims and scope of your project

The original aim of the project was to explore the impact new Raspberry Pi clusters have on TM129 students. We will be making the clusters available as optional activities in two presentations to assess their suitability, before making them a compulsory part of the assessment from 20J. We will be evaluating these clusters for:

1) increased engagement and interest of the students due to the novelty of the technology;

2) that the activities will increase students' knowledge of PDC; and

3) that the increased engagement and interest will improve retention.

Activities

We experienced few difficulties in deploying our project. After constructing the clusters, and working with the OpenSTEM Lab team in making them accessible for students, the main concerns of this project was in assessing the impact on students. To do so we deployed an optional questionnaire within the VLE, and conducted follow-up interviews with 8 students.

The clusters were provided to all students in the October 2020 and February 2021 cohorts, covering 993 and 1082 students respectively at the time of use. Completing the questionnaire was a completely optional activity. 251 students completed the questionnaire across the two presentations.

Eight students from the February cohort were interviewed to try and elicit information that would complement the questionnaire results. The analysis of the interview transcripts did not provide any insight beyond that gleaned from the questionnaires, and for clarity we thus focus on reporting the questionnaire results.

Findings

Tables 2 through 4 show the median and *(interquartile range)* for responses to the Likert-scale questions in the questionnaire.

Focusing first on the clusters themselves, we asked students on a scale of 1 to 5, where 1 is strongly disagree and 5 is strongly agree, how much do you agree with the statements in Table 2.

	October 2020 cohort	**February 2021 cohort**
After I completed the first cluster activity, I had difficulty solving the other cluster activities.	1 *(2)*	1 *(2)*
I found the challenge of solving the cluster activities motivating.	3 *(1)*	4 *(1)*
I enjoyed solving the cluster activities.	4 *(1)*	4 *(2)*
I am interested in learning more about parallel and distributed computing.	4 *(2)*	4 *(2)*
I am interested in learning more about computer science.	4 *(1)*	5 *(1)*

Table 2. Responses to the questions broadly relating to engagement with the discipline and the clusters.

We then asked students about their skills confidence as per the questions in Table 3; rated on a scale of 1 to 5, where 1 is not at all confident and 5 is very confident.

	October 2020 cohort	**February 2021 cohort**
Solve computer software problems.	3 *(1)*	4 *(1)*
Develop algorithms that work in parallel.	3 *(1)*	3 *(2)*
Solve problems using the clusters.	3 *(2)*	3 *(1)*

Table 3. Responses to the questions relating to student skill confidence

The final set of questions focussed on study interests as per Table 4; rated on a scale of 1 to 5, where 1 is not at all and 5 is very much.

	October 2020 cohort	**February 2021 cohort**
Use the Raspberry Pi clusters in other OU modules	4 *(1.25)*	4 *(2)*
Take future classes in parallel and distributed computing.	4 *(1)*	4 *(1)*
Take future classes in computer science.	4 *(1)*	4 *(1)*

Table 4. Responses to the questions relating to future study interests

Our interpretation of these figures is that the cluster activities appear to have been motivational and enjoyable. However, the efficacy of the clusters in terms of making the students competent in solving software problems or developing parallel algorithms is more questionable, potentially due to the passive nature of operating them (i.e. through a fixed interface, rather the programming the clusters directly).

We report the three key themes that emerged from our thematic analysis of the questionnaire results: exposure to real hardware; engagement with the concept of PDC and Practicalities of accessing the remote hardware. Please see our paper for full details (10.1080/02680513.2022.2118573).

Impact

The use of physical hardware was a novelty to many of these students, and clearly engaged a large proportion of the students who used the system. We are pleased with the feedback we received, including general praise from the student body: *"Genuinely some of the best activities I've done so far through [year 1]. The direct interaction with remote hardware felt engaging and is something I'd like to see again"*. One student was so enthused by the concepts taught that they contacted the module team to discuss their work in simulating their own cluster using VirtualBox on their local machine. Our experience through the project has demonstrated the need for C&C to more heavily engage in producing materials which simulate real-world activities.

As a final observation, when reflecting on the project, we note that there are specific challenges of working with open distance students which can slow the adoption of novel teaching practices. The construction of the clusters and developing the activities took time: academics also need to be prepared for the time committed to getting materials in front of students, particularly working through the development of a suitable infrastructure, considering accessibility needs and adapting the materials appropriately, and the time needed for long-term maintenance requirements across multiple presentations a year, each with high numbers of students. Academics and instructors need to be aware of these needs before beginning to deploy such systems, to ensure they have suitable capacity to see projects through to fruition.

List of deliverables

The clusters remain accessible through the OpenSTEM Labs as a resource for TM129 students. We have also published our work at:

Daniel Gooch, Jon Rosewell, Douglas Leith & Mike Richards (2022) Passive or active learning: the challenges of teaching distributed computing using Raspberry Pi clusters to open distance university students, Open Learning: The Journal of Open, Distance and e-Learning, DOI: 10.1080/02680513.2022.2118573

References

Please see publication above (10.1080/02680513.2022.2118573)

University approval processes

SRPP were consulted but stated "We've reviewed your application and supllementary material, and determined that it won't be necessary for you to make an SRPP application on this occasion due to the operational/functional nature of the proposed evaluation and its small scale. If you should find you want to explore the usefulness of the PIs as a learning tool in more detail, or go out to a wider population, that would be a more appropriate subject for an application to our panel"