Interactive Knowledge Graph Supporting Self-Directed and Personalised High-Performance Computing Education

1st Eva Fernandez Amez
Department of Computer Science
Durham University
Durham, UK
0009-0007-9422-5786

2nd Thomas A. Flynn
Advanced Research Computing
Durham University
Durham, UK
0000-0001-9800-594X

3rd Mladen Ivkovic
Department of Computer Science
Durham University
Durham, UK
0000-0002-3539-3831

4th Christopher D. Marcotte
Department of Computer Science
Durham University
Durham, UK
0000-0002-3751-6970

5th Tobias Weinzierl
Department of Computer Science
Durham University
Durham, UK
0000-0002-6208-1841

***Abstract*—Across universities, compute centres, and industry partners, significant resources are put into developing educational materials for researchers and developers in High-Performance Computing (HPC). It is important to acknowledge that the range of research backgrounds in HPC is beneficial for the community in many ways but poses challenges for education. In particular, many HPC users and developers would benefit from foundational knowledge that may seem obvious to computer scientists. We propose an interactive knowledge graph of fundamental concepts of HPC, allowing learners to choose their own starting point and navigate forwards or backwards through related topics. This content is delivered through accessible, short lecture videos featuring animations that are delivered by HPC specialists. These videos are supplemented with formative assessments, and all of the concepts are structured into an interactive knowledge graph which specifies dependencies between concepts so that the learner can find their own pathways through the learning journey.**

***Index Terms*—high-performance computing, education**

I. Introduction

Inter-disciplinarity shapes High-Performance Computing (HPC) and its community. From researchers to software developers to platform engineers, HPC is used and supported by a variety of specialists. This variety helps drive the development of HPC, but it also creates challenges, e.g., in training.

HPC education has grown to include master's programs, workshops, hackathons, and many more. Across these delivery methods there is a difficulty in pitching to the audience because the users of HPC come from an increasingly diverse background.

The authors acknowledge support via funding through the UKRI Digital Research Infrastructure Programme under grants UKRI1801 (SHAREing) and UKRI293 (HAI-End). This work has also been supported by DiRAC.

Some of the authors have used OpenAI's ChatGPT for language editing and proofreading and have reviewed all edits and take full responsibility for the final content.

In this manuscript we discuss some of the challenges we see in HPC education and introduce our proposed method of delivering fundamental concepts in HPC to a wide audience.

II. HPC training

Workshops are an excellent way of diving into a topic, but often learners can feel they are stitching together a patchwork of computational knowledge. For example, a learner may attend a software optimisation workshop which will likely use lots of unfamiliar jargon. These varied backgrounds make it difficult to design learning resources that are accessible and sufficiently engaging for different users.

Researchers often first encounter HPC once they have already begun to conduct research. Hence, they not only come from different domains, but may learn concepts in HPC out of order. In this sense, the biggest challenge is not necessarily learning new concepts, but figuring out the blind spots and foundational building blocks needed for next steps. This lack of conceptual sequencing makes it difficult to build a solid foundation.

To address these challenges, this work focuses on two objectives: (1) creating accessible, engaging educational material for users and developers new to HPC, and (2) developing modular, interconnected training that supports personalised learning pathways.

III. Our knowledge graph

Rather than presenting a predefined curriculum that learners must adapt to, we give them the opportunity to design their own learning pathway through our interactive knowledge graph.

To achieve this, we have identified a set of *foundational* concepts, each represented as an individual node within the

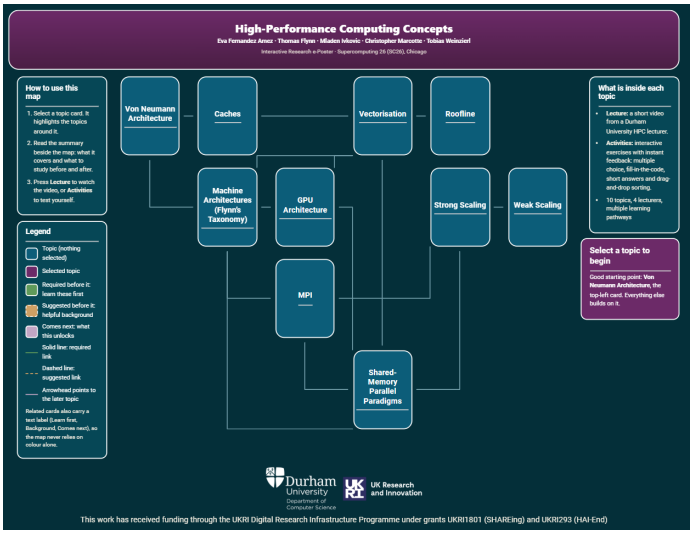


Fig. 1. Default view of the HPC graph.

graph (Figure 1). Nodes are interconnected through prerequisite and recommended relationships following pedagogical and conceptual dependencies. This allows users to choose the starting point that best suits their prior knowledge. For example, a newcomer might wish to begin with the most fundamental content – here, the Von Neumann Architecture – whereas someone with a computer science background may prefer starting with MPI.

Selecting a topic will dynamically highlight three categories of related material: (1) prerequisites that one should understand beforehand to make the most of the content; (2) suggested background concepts that provide useful context, though are nonessential; and (3) recommended follow-up nodes that can extend this material (Figure 2).

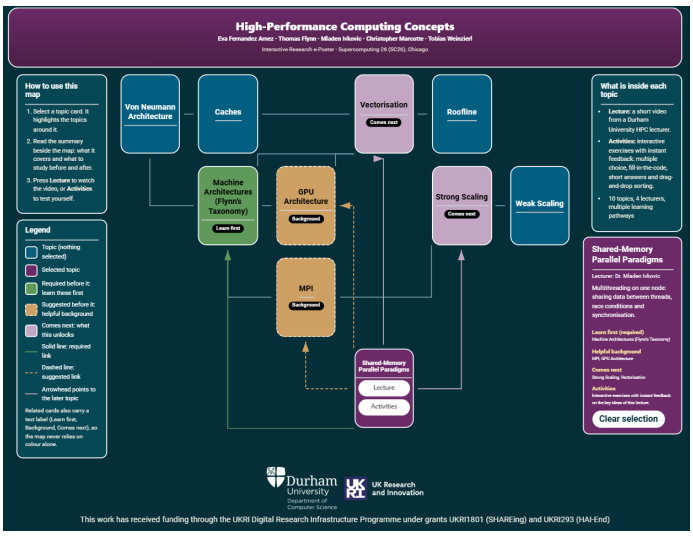


Fig. 2. Concept dependencies visualised after selecting a module.

Each topic features an educational video of approximately 10 - 15 minutes, in which HPC researchers explain technical concepts through professionally animated visualisations and real-world analogies. These videos have been designed to be accessible and engaging for learners with a wide range of disciplinary backgrounds. Furthermore, each module incorporates formative assessment activities – including multiple choice questions, true-or-false, and gap-filling exercises – that reinforce learning through feedback.

Overall, the graph combines pedagogical, technical and outreach principles to provide a resource that enables learners to situate individual topics within HPC.

IV. Impact and Future Work

This resource has already been used as review material for performance analysis workshops and HPC courses at Durham University. It is also embedded within the United Kingdom Digital Research Infrastructure portfolio.

Supported by ongoing funding, future work will focus on expanding the catalogue of topics and analysing how users navigate the knowledge graph. We are also in the process of liaising with external institutions to signpost existing training resources and avoid duplication. Overall, we aim to contribute to a more integrated and open-science approach to HPC education.

V. Benefits of the e-Poster

The interactive format of this resource provides capabilities that cannot be effectively communicated through a conventional poster.

Visitors will begin by selecting a concept within the knowledge graph and exploring its prerequisites and recommended relationships before watching the associated video. The videos include embedded reviewed captions, allowing the content to be followed without audio in the exhibition environment. For visitors who require it, downloadable transcripts and audio playback via a laptop will also be available. Visitors will also be encouraged to complete a selection of formative assessment activities.

The touchscreen interface is, therefore, an essential component of the educational model, enabling users to experience how concept dependency visualisations can support self-directed HPC learning.