An Educational and Training Perspective on Integrating Hybrid Technologies with HPC Systems for Solving Real-World Commercial Problems

Stefano Mensa
The Hartree Centre, STFC
Warrington, Cheshire
stefano.mensa@stfc.ac.uk

Emre Sahin
The Hartree Centre, STFC
Warrington, Cheshire
emre.sahin@stfc.ac.uk

George Williamson
The Hartree Centre, STFC
Warrington, Cheshire
george.williamson@stfc.ac.uk

Robert J. Allan
The Hartree Centre, STFC
Warrington, Cheshire
robert.allan@stfc.ac.uk

ABSTRACT

Delivering training and education on hybrid technologies (including AI, ML, GPU, Data and Visual Analytics including VR and Quantum Computing) integrated with HPC resources is key to enable individuals and businesses to take full advantage of digital technologies, hence enhancing processes within organisations and providing the enabling skills to thrive in a digital economy. Supercomputing centres focused on solving industry-led problems face the challenge of having a pool of users with little experience in executing simulations on large-scale facilities, as well as limited knowledge of advanced computational techniques and integrated technologies. We aim not only at educating them in using the facilities available, but to raise awareness of methods which have the potential to increase their productivity. In this paper, we provide our perspective on how to efficiently train industry users, and how to engage with them about wider digital technologies and how these, used efficiently together, can benefit their business.

KEYWORDS

Education, Training, HPC, Integrated Technologies, Customer Success, Quantum Computing Training, GPU Training, Digital Twinning

1 INTRODUCTION

The Hartree Centre (HC) is part of the Science and Technology Facilities Council (STFC) - one of UK Research and Innovation's research councils - building on the rich established scientific heritage and a network of international expertise to support the UK's continued leadership in computational science and digital innovation [1, 4, 6]. HC supports businesses and organisations of any size in the UK in exploring and implementing technologies such as supercomputing (HPC), data analytics and artificial intelligence (AI) for increased productivity, cleverer innovation, and economic growth. The centre is home to some of the most cutting-edge digital technologies and experts in the UK, supported by sizeable UK government funding and strategic partnerships with industry leaders. In 2021, the Hartree National Centre for Digital Innovation (HNCDI) [2] programme was established to provide a safe and supportive environment for UK businesses and public sector organisations to acquire the skills needed to adopt AI, develop proofs-of-concept and de-risk investment into emerging digital technologies such as quantum computing.

The Hartree Centre's in-house skills set is key to helping industrial partners deliver solutions to real-world challenges. True customer success, however, is achieved when customers fully understand how to apply acquired knowledge and can adapt it to their own business needs. Dealing with industrial customers as end users of HPC facilities can present unique challenges. A typical user coming from an industrial background is remarkably knowledgeable in a specific domain area, however, they often lack the knowledge required to perform numerical simulations on large-scale computing facilities. Furthermore, the adoption of hybrid computational technologies (that is the use of computational techniques such as ML, AI etc. in combination with classical HPC) is hindered by the lack of detailed understanding of the functionality.

These two issues have often three negative outcomes. The first casualty due to lack of "operational knowledge", is productivity. Users that do not know their way in an HPC infrastructure usually end-up in using the facility in a sub-optimal way, hence resulting in loss of productivity and ultimately financially impacting the project itself. Indeed, poor usage of computational resources will unavoidably drain paid project compute time allocation. Second, inexperienced, and non-self-sufficient users impact data-centres operations, opening tickets and incidents that take time to solve, diverting staff time into non-critical troubleshooting. Finally, lack of understanding of the low level functionality of new technologies limits their uptake and hinders digital innovation in the business. Hartree Centre staff aim to assist in all three areas.

In this paper, we provide our perspective on how to efficiently train users with an industrial background, not only on how to use HPC systems but also on how to engage with them about wider digital technologies and how these, used efficiently together, can

https://doi.org/10.22369/issn.2153-4136/14/1/6

benefit their business. Here, we describe our training and education strategy for users with a core industrial background using the following three stages. The first stage is about building a confident and self-sufficient cohort by providing consistent and systematic training on how to use HC supercomputing facilities. These users can pass on knowledge to their colleagues. The second stage is designed to build digital innovation awareness, where we engage with customers by showcasing successful examples of integration of hybrid technologies within a business pipeline, especially by means of visualisations aids. The final stage is a more specialised training and education, where customers already aware of the benefits of digital innovation for their business can gain in-depth knowledge via the HNCDI Explain programme.

1.1 HNCDI: The Hartree National Centre for Digital Innovation

The Hartree National Centre for Digital Innovation is a collaborative programme with IBM which will enable businesses to acquire the skills, knowledge and technical capability required to adopt digital technologies like supercomputing, data analytics, artificial intelligence (AI) and quantum computing.

Through HNCDI we provide a safe and supportive environment for organisations to explore the latest digital technologies and skills, develop proofs-of-concept and apply them to industry and public sector challenges. Our dynamic and collaborative approach is driven by industry requirements and will help organisations to de-risk investment in new and emerging digital technologies.

Either at the start of their digital journey or trying to advance to the next level, we can help businesses navigate the possibilities of AI and quantum computing technologies to discover the next step for their digital development.

The HNCDI programme is divided in four work-streams.

(1) *Emerging Technology:* we are looking at the future of computing in the UK, helping businesses to identify the areas where emerging digital technologies like quantum computing might offer the most competitive advantage.
(2) *Excelerate:* through our applied industrial research, we help to turn good ideas into industry-ready solutions that address business challenges, embedding AI solutions across the industry.
(3) *Explore:* it aims to go one step further by finding ways to solve industry challenges when there isn't an existing off-the-shelf solution but there is evidence it can be solved and a business value and motivation to solve it.
(4) *Explain:* in this work-stream, HC staff works with individuals to identify learning pathways through our course catalogue that will equip their organisation with the skills needed to take advantage of digital technologies. Explain will be discussed more in-depth in the following Sections.

2 BUILDING A CONFIDENT AND SELF-SUFFICIENT USER COHORT

Successful routine use of a supercomputer in a commercial project goes hand in hand with the proficiency of its project members in making the most of the available infrastructure. This involves managing their data as well as efficiently targeting the resources in terms of processor type, number, etc. for specific simulation cases. Although HPC infrastructure across the world operates with the same principles (distribute computing over a fast network, hybrid hardware, job scheduler to orchestrate the workload, distributed file system etc.) and use pretty much the same family of operating systems (e.g. Linux based clusters), each data centre is different and only rarely will HPC systems have identical features. Thus, even experienced users will have somewhat to re-learn and adapt when moving onto a new supercomputer. The time taken to adapt to a new machine depends on the proficiency of the user. There are several educative and training perspectives as mentioned in [11, 14, 26–29].

We believe that customer on-boarding plays a crucial role for businesses' journey towards integration of hybrid technologies. For this reason, each and every new HC user undergoes an on-boarding process that we call "driving license", a training course delivered as a two hour lecture in which the users are expected to learn the fundamentals of our supercomputer, Scafell Pike (Top 500 list). To complement the lecture, a hand-book is also made available, see [19]. Topics covered in the course span both hardware and software of the machine, and the message we want to share is that there is no efficient usage of a supercomputer if first it is not clearly understood how the machine works, in terms of it's hardware, the nature of the file-system, the job-scheduler and the overall software stack. Finally, practical examples of job submissions are provided, also useful as a starting template for customised submission scripts. Before users receive their machine accounts, a "driving license" test needs to be passed, in which the users demonstrate competency in the usage of the machine, by being quizzed on a number of question regarding our HPC facility usage.

The HC on-boarding process has a number of benefits. First, the training provided guarantees a consistent minimum working knowledge across users in the centre, meaning that even a complete novice possess the relevant knowledge to comfortably move around the system. Second, the amount of downtime that users experience due to unexpected issues at submission time, execution time and so forth is significantly reduced, as understanding of the system provides the user with basic diagnostic skills (e.g. the job did not produce any output because it was submitted from the wrong base directory). Third, data-centre operations also benefits a trained user workforce a, generally speaking, trained users tend to raise less tickets for issues, thus reducing the amount of staff time spent troubleshooting basic issues.

3 BUILDING DIGITAL INNOVATION AWARENESS

The second fundamental stage of training and educating users with an industrial background is to build digital innovation awareness, that is to understand which and how digital technologies can enhance business productivity and profits. However, learning about the power of computational methods and, in general, novel technologies to address business critical objectives is hard if approached under a purely theoretical perspective, and a practical, tangible example would be a more effective learning tool. The HC has 10 years of experience in enhancing businesses profitability through the application of advanced technologies, with a large portfolio of projects and case studies available to share. Thanks to our portfolio, we built a number of demonstrations showcasing projects outcomes obtained with hybrid technologies (HPC, AI/ML etc.), using visual computing strategies aiming at immersing the customer in an engaging visual and virtual environment, to help promoting and understanding the impact of a knowledgeable use of such hybrid technologies. We found out that visual computing is a good enabler towards digital innovation awareness.

3.1 Visual Computing as an Educational Tool

Visual computing and digital twinning have unquestionable advantage to bring insight and deeper understanding for scientific, industrial and educational fields. Exploiting such emerging technologies can allow us to bring more awareness and interest to applications of HPC.

Here at the Hartree Centre we have a range of visual computing facilities suitable for demonstration and training purposes as well as use in a project work. Our two showcase rooms house large main displays, and both of these displays are stereo 3D capable and both are equipped with high-end visualisation workstations and HD audio systems. All the displays and devices throughout the Visual Computing suit are connected through a high-speed multi-cast network utilising Crestron NVX transducers and switches such that any computer in the system can be connected to any display. We typically connect to our HPC facilities through VNC [5] for interactive access and remote visualisation. We have a number of licensed and open source software tools available for our scientists, engineers and partners to train and utilise for project work. Combining tools such as these with our super-computing systems can allow us to manipulate, control and visualise data on a massive scale for educational and training purposes.

Ranging from partners to students, we demonstrate a number of visually exciting case studies from past projects to help raise awareness of how we can engage with them and help them realise their potential. Our visualisation systems are instrumental in improving the quality of demonstrations and presentations, and this helps us to enable better experience during training and education processes.

One of the projects we demonstrate in this way is our Virtual Wind Tunnel project (VWT) [30]. We visualised a project investigating airflow over a prototype car body, and it produced some near photo-realistic renders of streamlined data from computational fluid dynamics (CFD) simulations. We ran the CFD simulation on one of our supercomputer systems and the data produced was post-processed and overlaid onto the CAD model in the rendering software also running on the supercomputer. This led to the development of a supercomputer CFD workflow, which we call our Virtual Wind Tunnel (VWT) and an application used to display and investigate CFD data overlaid on CAD models in a more realistic and human-relatable 3D/VR environment. The app simulates an actual wind tunnel and allows the user to move around the test object and view it from all angles while displaying streamlined data. Particles can be added to the streams and animated along the streamline trajectories to show the development of flow. This work-in-progress app when shown on our larger displays gives a more realistic and relatable feel when visualising data. This in turn demonstrates to our customers how the correct visualisation can give insight into what the data is showing us by putting the human back in the loop.

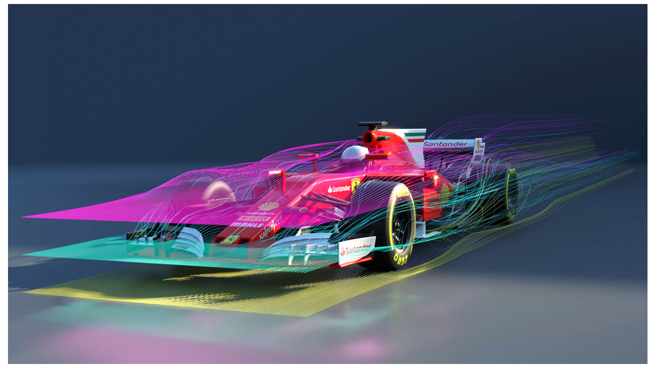

Figure 1: Photorealistic rendering of a F1 Ferrari overlapped with CFD streamline data simulated on our HPC facility.

Another example of visualising big data is our Dengue Virus demonstration. Here we show a stereo 3D visualisation of the viral protein that causes dengue fever on our 4.7-metre display pointing out to people that this is just one way to view the 1.1 million atoms that make up the protein. Such kind of protein visualisation is helpful to intuitively grasp the structure of the virus, potentially identifying hidden pockets that could be targeted by drugs.

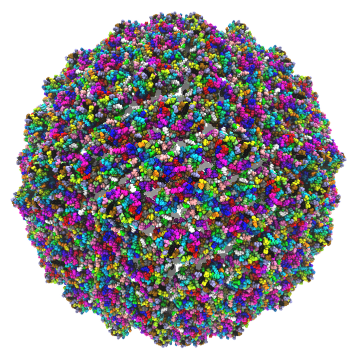

Figure 2: Dengue virus envelope protein rendered with VDW drawing method using VMD [16].

Demonstrating techniques for visualising big data sets in this way will hopefully show a good visualisation can bring potential benefits to educational, scientific and industrial projects and how it can help to make big data more accessible to more people.

4 SPECIALIST TRAINING AND EDUCATION WITH THE HNCDI EXPLAIN PROGRAMME

To meet some of the above goals, we offer application-focused training via the HNCDI Explain program, designed to enable individuals and businesses to take full advantage of digital technologies to enhance processes within their organisation and provide the skills that enable them to thrive in a digital economy.

Our training model is flexible and built with industry in mind. Whether learning the basics or searching for new tools and techniques to apply, a combination of self-directed online learning and face-to-face practical sessions can be used with certification.

The HC team will work with individuals to identify learning pathways through our course catalogue that will equip their organisation with the skills needed to take advantage of digital technologies. These skills can then be passed on to colleagues who will also have free access to the training materials. In general, four different levels of training are provided:

- *Introductory:* trainees from a non-related background with very little knowledge of the subject area;
- *Learner:* trainees with some theoretical or practical knowledge within the relevant subject area;
- *Independent user:* trainees who can work independently within the subject area but would require guidance for solving complex problems;
- *Practitioner:* trainees actively working in the subject field, looking to investigate emerging technology developments, and new techniques and/or develop collaborative multidisciplinary applications with higher levels of complexity.

Below two examples of specialist training we provide. A more comprehensive list of training courses offered can be found here [3].

4.1 Computational use of GPUs

Much research software, particularly open source, is nowadays developed to work using GPUs for the acceleration of critical numerical components. This is particularly true in fields such as machine learning and AI, bio-informatics and chemistry, solution of linear systems of equations for engineering applications and so forth. Despite the technological relevance, good software engineering practice for GPU accelerated software is somewhat limited to a small portion of specialised software engineers, and it's only being taught in specific academic degree courses, leaving interested users to rely on online open-source material or resorting to a self-taught strategy.

HC has a core specialism in GPU software development, as well as owning large GPU-based resources. Such specialism and hardware are exploited by the Centre to train users, aiming at sharing good GPUs software engineering practices. Teaming with partners such as NVIDIA, we offer hands-on training in GPU accelerated computing to solve real-world and industry-relevant problems, getting much-needed practical experience and earning a certificate of competency to support professional growth. As mentioned in [12, 15], this significantly improves the interest of users in the field and encourages more people to accelerate their codes. Furthermore, supervising skilled and interested students across Europe in projects with state-of-art topics using GPUs during PRACE Summer of HPC, brought more interest and possible workforce for the area.

4.2 Quantum Computing

The rise of Quantum Computing (QC) as the next mainstream computing paradigm for code acceleration is gaining momentum in the scientific computing community, promising to change the way we solve real-world challenges. Based on completely different physical rules compared to traditional, classical computing, QC requires a different mind-set and a different approach in the way code is written. This is true due to a number of factors: completely new hardware to be interfaced with classical facilities, classical codes needs to be ported to be suitable to work on quantum hardware and, overall, basic understanding of quantum mechanical rules and how these affects computation. As such, the need for structured and rigorous training for QC is very much needed.

Structured training is a strong challenge for super-computing centres to adapt and prepare materials for emerging technologies such as QC. There are readily-available materials for quantum computing such as:

- *Introductory and Learner level:* Michael Nielsen's book [25] and tutorials [17, 18, 23, 24] are one of the most welcoming when it comes to introduction for QC. As previously mentioned by [20–22], games and interactive environments have significant importance in education and training in HPC, and there are also QC-related games available [7, 32]. There are several very well-prepared materials such as Qiskit textbook [8, 31] Tensorflow Quantum [10], Cirq [13], Pennylane Tutorials [9], however, they are tailored for specific SDKs and hardware. Furthermore, installing a QC environment and using it can be challenging especially for end-users. In addition, SDKs and QC environments are not stable and getting updated regularly.
- *Practitioner level:* Previously mentioned materials also have advanced levels for people who wants to specialise. In addition, there are many articles available and published every day about quantum computing.

In HC, *via* the Explain program we aim at building the next generation of quantum software engineers from bottom up. First, we provide training for the basics of quantum computing with respect to introductory applied quantum mechanics, and afterwards, for specific hardware (e.g., quantum annealing, universal gate-based etc.) and their SDKs. Despite we commonly use universal gate-based systems for hands-on training using simulators, we can tailor the materials according to the specific needs of a project or users. Furthermore, we also encourage and support individuals to take extensive courses from external sources and join related events to be specialists in the area.

5 CONCLUSIONS

In this paper we provide our perspective on how to effectively train industry users, and how to engage with them about wider digital technologies and how these, used efficiently together, can benefit their business. Specifically, we have discussed our three stages education plan. In the first stage, we provide each and every user with a core training on how to use efficiently our HPC system, building a confident and self-sufficient user cohort that can productively use the machine. In the second stage, we engage with industry users building digital innovation awareness. This stage is key to provide businesses with concrete examples of how applied digital strategy can bring benefits to their business. The third and last stage is to provide specialist training tailored to match the business needs of the users, as well as training on novel emerging technologies via our HNCDI Explain work stream. This allows us to stay competitive meanwhile building the future's work force.

To conclude, through the approaches described in this paper we have introduced businesses to new ways of working, and incorporating new technologies into their research pipeline through enhanced service offerings. Demonstrating the capabilities based on each kind of technology leads to an increased demand to use them. As specific examples, this includes using remote interactive access and visualisation, use of GPUs and quantum computing. Data analysis and AI solutions are also being included in software development work, e.g. using AI models alongside more traditional mathematical models.

ACKNOWLEDGMENTS

This work was supported by the Hartree National Centre for Digital Innovation, a UK Government-funded collaboration between STFC and IBM.

REFERENCES

[1] 2022. The Hartree Centre. https://www.hartree.stfc.ac.uk/.
[2] 2022. Hartree National Centre for Digintal Innovation. https://www.hartree.stfc.ac.uk/Pages/Hartree-National-Centre-for-Digital-Innovation-(HNCDI).aspx.
[3] 2022. HNCDI Explain courses. https://www.eventbrite.com/cc/hartree-centre-explain-training-programme-259399.
[4] 2022. Science and Technology Facilities Council. https://www.ukri.org/councils/stfc/.
[5] 2022. Turbo VNC. https://www.turbovnc.org/.
[6] 2022. UKRI - UK Research and Innovation. https://www.ukri.org/.
[7] TOPTICA Photonics AG. 2022. Quantum Quiz. https://www.toptica.com/quantumquiz
[8] Gadi Aleksandrowicz et al. 2019. Qiskit: An open-source framework for quantum computing. (2019).
[9] Ville Bergholm, Josh Izaac, Maria Schuld, Christian Gogolin, M Sohaib Alam, Shahnawaz Ahmed, Juan Miguel Arrazola, Carsten Blank, Alain Delgado, Soran Jahangiri, et al. 2018. Pennylane: Automatic differentiation of hybrid quantum-classical computations. *arXiv preprint arXiv:1811.04968* (2018).
[10] Michael Broughton, Guillaume Verdon, Trevor McCourt, Antonio J Martinez, Jae Hyeon Yoo, Sergei V Isakov, Philip Massey, Ramin Halavati, Murphy Yuezhen Niu, Alexander Zlokapa, et al. 2020. Tensorflow quantum: A software framework for quantum machine learning. *arXiv preprint arXiv:2003.02989* (2020).
[11] Manuel Caeiro-Rodríguez, Mario Manso-Vázquez, Fernando A Mikic-Fonte, Martín Llamas-Nistal, Manuel J Fernández-Iglesias, Hariklia Tsalapatas, Olivier Heidmann, Carlos Vaz De Carvalho, Triinu Jesmin, Jaanus Terasmaa, et al. 2021. Teaching soft skills in engineering education: An European perspective. *IEEE Access* 9 (2021), 29222–29242.
[12] Xi Chen, Gregory S. Gutmann, and Joe Bungo. 2018. Deep Learning by Doing: The NVIDIA Deep Learning Institute and University Ambassador Program. *CoRR* abs/1812.08671 (2018). arXiv:1812.08671 http://arxiv.org/abs/1812.08671
[13] Cirq Developers. 2022. Cirq. https://doi.org/10.5281/zenodo.6599601 See full list of authors on Github: https://github .com/quantumlib/Cirq/graphs/contributors.
[14] Bence Ferdinandy, Ángel Manuel Guerrero-Higueras, Éva Verderber, Ádám Miklósi, and Vicente Matellán. 2019. Analysis of users' first contact with High-Performance Computing: first approach with ethology researchers. In *Proceedings of the Seventh International Conference on Technological Ecosystems for Enhancing Multiculturality*. 554–557.
[15] Liqiang He. 2010. Computer architecture education in multicore era: Is the time to change. In *2010 3rd International Conference on Computer Science and Information Technology*, Vol. 9. IEEE, 724–728.
[16] William Humphrey, Andrew Dalke, and Klaus Schulten. 1996. VMD – Visual Molecular Dynamics. *Journal of Molecular Graphics* 14 (1996), 33–38.
[17] Andy Matuschak and Michael A Nielsen. 2019. How does the quantum search algorithm work? *URL: https://quantum.country/search* (2019).
[18] Andy Matuschak and Michael A Nielsen. 2019. Quantum computing for the very curious. *URL: https://quantum. country/qcvc (cit. on p. 8)* (2019).
[19] Stefano Mensa. 2022. Scafell Pike Customer Onboarding. https://www.hartree.stfc.ac.uk/Pages/The_Hartre_Centre_Scafell_Pike_Introductory_Manual_V1_6.pdf.
[20] Julia Mullen, Chansup Byun, Vijay Gadepally, Siddharth Samsi, Albert Reuther, and Jeremy Kepner. 2017. Learning by doing, High Performance Computing education in the MOOC era. *J. Parallel and Distrib. Comput.* 105 (2017), 105–115.
[21] Julia Mullen, Lauren Milechin, and Dennis Milechin. 2021. Teaching and learning HPC through serious games. *J. Parallel and Distrib. Comput.* 158 (2021), 115–125.
[22] Julia Mullen, Albert Reuther, William Arcand, Bill Bergeron, David Bestor, Chansup Byun, Vijay Gadepally, Michael Houle, Matthew Hubbell, Michael Jones, Anna Klein, Peter Michaleas, Lauren Milechin, Andrew Prout, Antonio Rosa, Siddharth Samsi, Charles Yee, and Jeremy Kepner. 2018. Lessons Learned from a Decade of Providing Interactive, On-Demand High Performance Computing to Scientists and Engineers. In *High Performance Computing*, Rio Yokota, Michèle Weiland, John Shalf, and Sadaf Alam (Eds.). Springer International Publishing, Cham, 655–668.
[23] Michael Nielsen. 2010. Quantum computing for the determined. (2010).
[24] Michael A Nielsen. 2010. The Quantum Computing for the Determined. *URL: https://www.youtube.com/* (2010).
[25] Michael A Nielsen and Isaac Chuang. 2002. Quantum computation and quantum information.
[26] Rajendra K Raj, Carol J Romanowski, Sherif G Aly, Brett A Becker, Juan Chen, Sheikh Ghafoor, Nasser Giacaman, Steven I Gordon, Cruz Izu, Shahram Rahimi, et al. 2020. Toward High Performance Computing Education. In *Proceedings of the 2020 ACM Conference on Innovation and Technology in Computer Science Education*. 504–505.
[27] Rajendra K Raj, Carol J Romanowski, John Impagliazzo, Sherif G Aly, Brett A Becker, Juan Chen, Sheikh Ghafoor, Nasser Giacaman, Steven I Gordon, Cruz Izu, et al. 2020. High performance computing education: Current challenges and future directions. In *Proceedings of the Working Group Reports on Innovation and Technology in Computer Science Education*. 51–74.
[28] Felipe Augusto Lara Soares, Cristiane Neri Nobre, and Henrique Cota de Freitas. 2019. Parallel programming in computing undergraduate courses: A systematic mapping of the literature. *IEEE Latin America Transactions* 17, 08 (2019), 1371–1381.
[29] Leonardo BA Vasconcelos, Felipe AL Soares, Pedro Henrique MM Penna, Max V Machado, Luís Fabrício W Góes, Carlos Augusto PS Martins, and Henrique C Freitas. 2019. Teaching parallel programming to freshmen in an undergraduate computer science program. In *2019 IEEE Frontiers in Education Conference (FIE)*. IEEE, 1–8.
[30] George Williamson. 2019. Virtual Wind Tunnel. *URL: https://www.hartree.stfc.ac.uk/Pages/Virtual-Wind-Tunnel.aspx* (2019).
[31] James R. Wootton, Francis Harkins, Nicholas T. Bronn, Almudena Carrera Vazquez, Anna Phan, and Abraham T. Asfaw. 2021. Teaching quantum computing with an interactive textbook. In *2021 IEEE International Conference on Quantum Computing and Engineering (QCE)*. 385–391. https://doi.org/10.1109/QCE52317.2021.00058
[32] Iftact Yakar. 2022. Quantle. https://github.com/deduckproject/quantle.