

What do supercomputers and supermarket cashiers have in common?

A supercomputer performance is light years away from what we can find in consumer devices. Not even the best gaming computers come close to what a supercomputer can do. Technically speaking, supercomputers can generally perform 10^{17} floating-point operations per second (FLOPS), or 100 PFLOPS (peta FLOPS) compared to a few hundred GFLOPS (giga FLOPS) of a regular desktop computer. These machines are capable of simultaneously processing tremendous amounts of information from various applications.

Supercomputing plays an important role in Computer Science, specifically in high-demanding tasks such as quantum mechanics, weather forecasting, climate research, oil and gas exploration, molecular modelling, physical simulations of the early moments of the universe, airplane and spacecraft aerodynamics, and nuclear fusion.

These infrastructures usually perform many high-demanding tasks simultaneously and the increasing volumes of data led to the need for the prioritization of these tasks. “A supercomputer that has hundreds to thousands of tasks running typically stores and accesses data from a shared file system. This results in a competition for resources so fierce that some tasks might get delayed indefinitely”, clarified João Paulo, a Computer Science researcher at the [INESC TEC - Institute for Systems and Computer Engineering, Technology and Science](#) and the University of Minho (Portugal).

Supercomputers vs Supermarkets

To address this issue, the [PAIO framework](#) came together, joining researchers from INESC TEC, the [University of Texas at Austin](#) (USA), and the [AIST - National Institute of Advanced Industrial Science and Technology](#) (Japan) to collaborate on the [BigHPC project](#). “The PAIO framework has been developed to ensure fair data access and storage for applications using shared storage resources. We can achieve this by attributing the same priority to all applications, or even prioritizing the critical ones so they can access data faster and conclude their studies as soon as possible”, explained Ricardo Macedo, a Computer Science researcher at the INESC TEC and the University of Minho.

Supercomputers running multiple tasks are like supermarkets. At the supermarket’s opening time, customers would choose a cashier they believed would clear the queue faster. However, this is not linear, and sometimes choosing the “wrong” cashier would lead to considerable delays. Now, imagine the open cashiers are the supercomputer’s shared storage resources and the costumers are the programs trying to complete their tasks as soon as possible. What the supermarket industry eventually adopted is a single-line solution and automatic cashier attribution that reduce waiting times, leading to a more efficient checkout experience. This is one of the principles of the PAIO framework: efficiently managing supercomputer storage resources.

Deep Learning

Among the tasks supercomputers are used for, a particular one has stirred the interest of INESC TEC researchers: deep learning. Deep learning is one category of machine learning that can achieve state-of-the-art accuracy, being the backbone of technologies like virtual assistants, facial recognition, and computer vision.

The team of researchers is also proposing a new framework called [Monarch](#), whose main goal is training deep learning models faster - currently being developed in collaboration with the [Texas Advanced Computing Center](#) and the [Hood College](#) (USA) within the scope of the [PASTor project](#). Using local storage resources dedicated to deep learning applications when performing these model-training tasks usually leads to faster results when compared with using shared storage, which is used by many applications at the same time.

“The dedicated local storage resources are not always known to the supercomputer user and choosing which information should be stored in these resources is a hard and time-consuming manual task” explained João Paulo. “The Monarch system automates this process and makes it possible to use both local and shared storage in order to improve the training of deep learning models”, he adds.

Deep learning vs Supermarkets

The introduction of self-checkout in supermarkets was implemented to improve customer checkout speed. However, the adoption of these automatic solutions is dependent on how comfortable customers are with using them. They might be faster, but if they are also too complicated, the customers will probably stick with the classic cashier and human interaction.

What Monarch aims to achieve is that supercomputer users (*the supermarket customers*) can transparently use locally available resources (*self-checkout solutions*) taking advantage of its improved performance without having to change the way they traditionally execute the applications (*without having to change the way they checked out on a classic cashier*).

Even though the frameworks have only been tested in controlled scenarios, INESC TEC and its partners, namely TACC and AIST, aim to continue the development of the technologies and implement them in their supercomputers. Supercomputers push scientific and technological knowledge to new levels, and improving their efficiency means improving the speed the world learns of new ways to solve global societal issues: from climate change to cancer and energy crisis.

Acknowledgments

The project BigHPC – A Management Framework for Consolidated Big Data and HPC (reference POCI-01-0247-FEDER-045924), leading to this work, is co-financed by the ERDF - European Regional Development Fund through the Operational Program for Competitiveness and Internationalisation - COMPETE 2020, the Lisbon Portugal Regional Operational Program - Lisboa 2020 and by the Portuguese Foundation for Science and Technology - FCT under UT

Austin Portugal. The project PASTor was funded through Portuguese funding by FCT – Fundação para a Ciência e a Tecnologia, I.P.