ACHIEVE

Co-funded by
the European Union

Modern Computing Infrastructures:

Why Students Must See Beyond the Software

Authors: Prof. Gianluca Palermo, Prof. Danilo Ardagna, Prof. Manuel Roveri

Institution: Politecnico di Milano – Dept. of Electronics, Information
and Bioengineering

Modern Computing Infrastructures:

Why Students Must See Beyond the Software

Authors: Prof. Gianluca Palermo, Prof. Danilo Ardagna, Prof. Manuel Roveri

Institution: Politecnico di Milano – Dept. of Electronics, Information and Bioengineering

> *When we think about computing today, we can easily imagine apps, online services, or artificial intelligence models. However, behind the software that shapes our digital lives, there are large physical infrastructures consisting of large-scale data centers and supercomputing facilities. These sites are the backbone of the modern digital world, enabling everything from video streaming and online banking to cutting-edge climate simulations and drug discovery.*

The Rise of Large Data Centers and Supercomputing Facilities

Large-scale data centers are the factories of the digital age. They host thousands, or sometimes hundreds of thousands, of servers, all working together to provide services to millions of users. Companies such as Amazon, Microsoft, and Google operate huge facilities around the world where hardware, networking, and software are integrated to ensure that cloud services are available 24/7. These infrastructures are highly efficient technological solutions in which servers are packed into racks and connected by high-speed networks, all supported by sophisticated cooling systems that prevent overheating. Their power consumption is enormous. A single hyperscale data center can consume the same amount of electricity as a medium-sized city, making energy efficiency and sustainability critical aspects of data center design. Even worse is if we consider the energy needs of artificial intelligence services that now far exceed those of most conventional digital services. Reports estimate that running generative AI models can require up to ten times more electricity than standard cloud applications. To meet this demand, new AI-oriented data centers are being planned on a scale comparable to power plants (gigawatt-class), highlighting the unprecedented size and infrastructure required to support future workloads.

While data centers power the modern digital economy, supercomputing centers drive scientific discovery and innovation. Rather than serving billions of lightweight requests, supercomputers are designed to solve some of the most demanding computational problems, such as modelling the Earth's climate, simulating nuclear reactions or molecular interactions, designing aircraft engines, and training artificial intelligence models with billions of parameters. These facilities operate at the cutting edge of performance. The current generation of supercomputers has entered the exascale era, with the ability to perform more than 10^{18} calculations per second. Such machines are not only fast, but also carefully engineered systems in which processors, accelerators, memory, storage, and interconnects are tuned to work together with minimal latency and maximum throughput.

More Than Software: The Importance of Hardware Awareness

When entering the world of computing, students and young professionals naturally tend to focus on software, such as programming applications for the cloud, parallelizing algorithms for supercomputers, and leveraging frameworks such as TensorFlow, MPI, and Kubernetes. While these skills are essential, they only tell half the story. Understanding the 'hard' side of computing (i.e. the infrastructure) is just as important as understanding the 'soft' side (i.e. the applications). Every program runs on physical machines housed in racks, powered by electrical systems, and kept alive by advanced cooling and monitoring technologies. The performance, reliability, and cost of computing are all shaped by these physical factors. A good HPC or cloud engineer needs to know not only how to code, but also where the code runs. Understanding of the physical constraints helps in writing more efficient, realistic applications, and fosters collaboration with the hardware specialists who keep alive the whole infrastructure.

Bridging the Two Worlds

Modern computing education must bring together knowledge of both software and hardware. Students who focus solely on programming may fail to recognize the real bottlenecks, such as memory bandwidth, network topology, and cooling constraints. Conversely, if we focus solely on infrastructure, we overlook the creative power of algorithms and applications. The most capable professionals are those who can move across both domains, understanding how a parallel program scales across nodes, while also recognizing the energy cost of each computation, or the trade-offs between GPUs and CPUs in terms of power density and cooling. As computing evolves into new frontiers, such as quantum technologies, AI accelerators, and energy-efficient designs, the boundary between ‘soft’ and ‘hard’ will blur even further. Breakthroughs will not only come from writing better code, but also from designing systems in which hardware and software advance together. This dual perspective will make them better engineers, innovators, and problem solvers in an increasingly complex technological landscape.

Conclusions

Modern computing infrastructures are the new cathedrals of the digital age. Data centers keep the cloud running while supercomputers push forward scientific research. At Politecnico di Milano, students on the double degree Master in Cloud, Networking Infrastructure and HPC, supported by the ACHIEVE project, are required to take a dedicated course on Computing Infrastructures. Here, students learn that software and AI innovation cannot exist without solid hardware foundations (racks, power, cooling, and reliability), which support and shape everything else. Understanding hardware and software layers is essential in preparing for a future where performance, sustainability, and innovation are closely interconnected.