

NumPEX, building a software stack for computing at Exascale

Julien Bigot

May 6, 2025

With the advent of Exascale supercomputers, capable of performing every second up to a billion of billion operations on floating-point numbers, it becomes possible to develop new numerical simulations and data processing workflows at scale and accuracy never envisioned before. This power does however not come without difficulty. With the end of Moore's law, sequential processing power stopped increasing and new sources of improvement had to be found, such as the extreme parallelism of GPUs that power all Exascale supercomputers; this makes their programming a challenge. Computing power also increases much faster than other aspects of supercomputers, such as communication and storage latency or size; as a result, disks do not have the ability to store all data produced anymore, one has to find solutions to reduce it before storage. In this presentation, I will introduce the NumPEX programs that aims to produce the software stack of the French Exascale supercomputer to come, by proposing transformative solutions to solve these bottlenecks, implement them in publicly available open source tools and libraries, and validate them at scale in collaboration with application developers. I will especially focus on the new C++-based GPU programming models and on solution for transparent in situ data processing that leverage machine learning approaches.