

Languages, APIs and Compilers In Service of Parallelism

Máté Ferenc Nagy-Egri

GPU Lab, Wigner Research Centre for Physics

Since the dawn of Graphics Processing Units, architectural development has never stopped for an instance, neither hardware-wise nor software-wise. The capabilities of GPUs have become so versatile that the domains of General Purpose GPU usage, highly parallel CPU and Field Programmable Gate Array programming can mostly be accounted for with a single notation. New capabilities brought about in the last few years have spawned a myriad of compilers and APIs that are hard to navigate about, thus standardization efforts have multiplied to keep the number of tools at bay.

The following overview aims to serve as a compass to those who have not actively followed recent advancements in the field of massively parallel computing.