

Supercomputing Shouldn't Be Rocket Science

A piece for Asian Scientist arguing that the hard part of supercomputing is not the computing but the gap between a working laptop simulation and code that survives at scale.

march 1, 2019 · Press

a piece for asian scientist, written while running archanan, arguing that the hard part of supercomputing is not the computing.

picture a researcher with a simulation that works on their laptop and an allocation on a national machine. between those two things sits weeks of work that has nothing to do with their science: mpi topologies, schedulers, queue etiquette, and the first genuinely large run is usually also the first time anyone finds out whether the code scales. that run is expensive, and when it fails it usually fails for a reason that has nothing to do with the physics.

the argument is that this gap is an engineering problem, not a fact of nature, and that treating it as a rite of passage keeps the field closed to anyone who does not already know the folklore. domain scientists should be able to spend their allocation on science, not on discovering that their halo exchange deadlocks above four thousand ranks.

[read the full article](#)

https://alexnodeland.com/timeline/190301_hpc-accessibility/