Wafer-Scale Engines Boost Scientific Productivity and Save Energy at the Same Time

Wafer-Scale Engines (WSEs) provide a compact solution for high-performance computing, consuming ~1000 times less energy than traditional processing units while increasing speed by as much as 470 times.

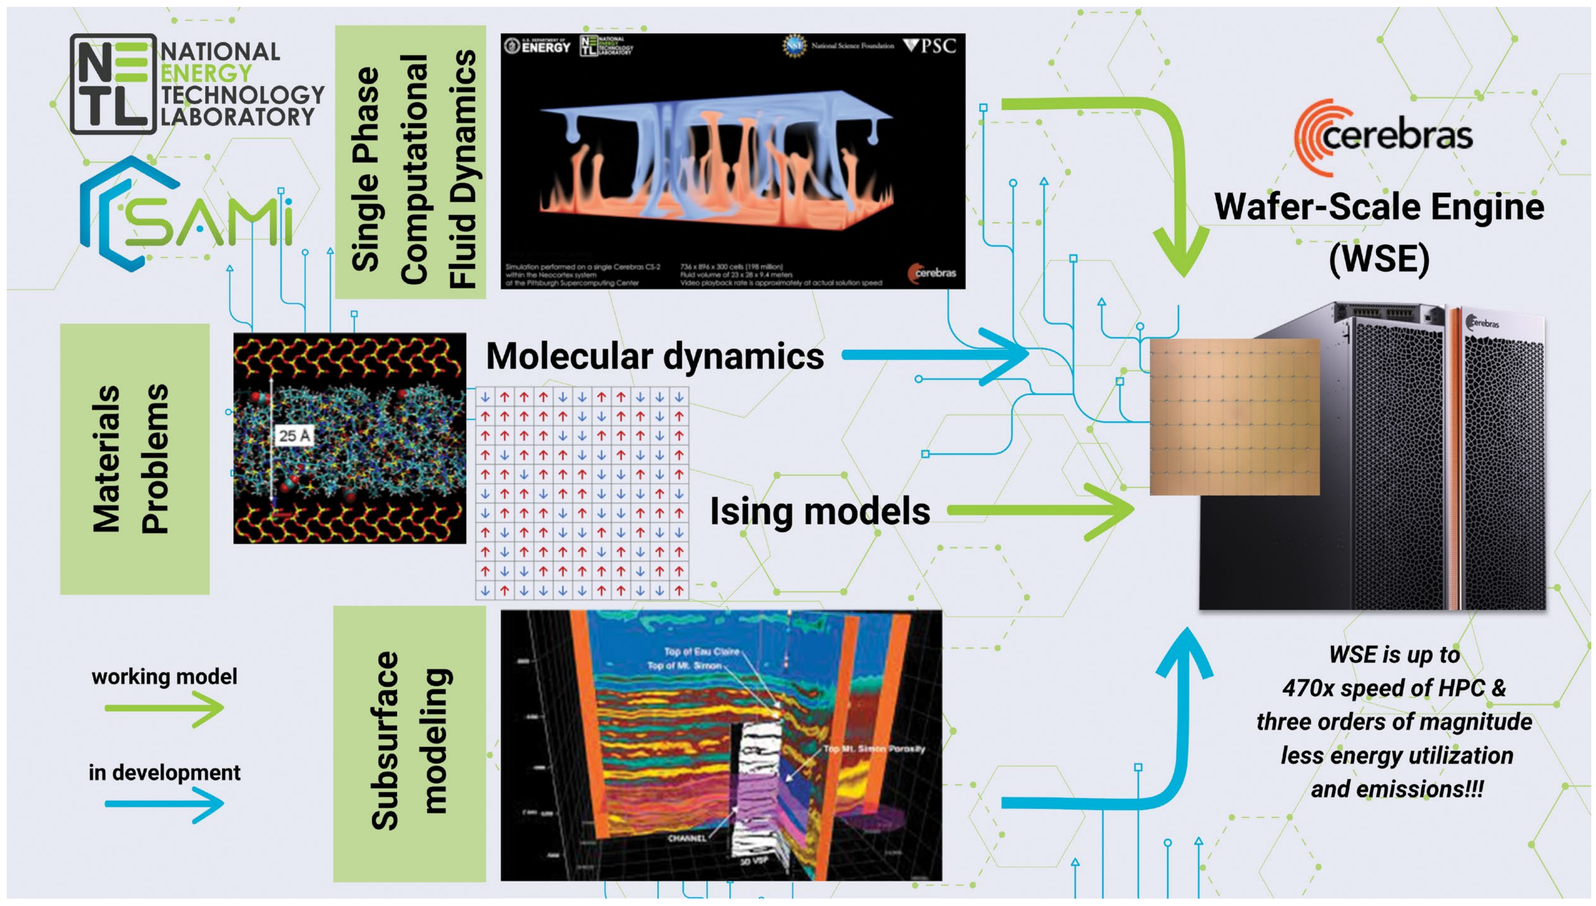

Wafer-scale engines are especially useful for high-intensity computational modeling relevant to DOE/NETL research and are able to perform such modeling in a faster and more energy efficient manner.

Applications of national importance require increasingly advanced computing to model complex phenomena and manage extensive data.

- NETL and Cerebras have collaborated to use the unique capabilities of the WSE for research relevant to FECM applications.
- NETL has advanced the technology by developing compilers and libraries within the WSE Field Equation application programming interface that enable researchers to easily develop models for complex phenomena, which can be computed at rapid speeds and with high energy efficiency.
- In 2020, a 200x speedup was achieved for the 7pt stencil structured grid BiCgStab compared to 16,386 central processing unit cores on Joule 2.0: a world record.
- By 2023, a 540x speedup was reached for a single-phase flow problem on an NVIDIA graphics processing unit.
- In 2024, NETL achieved an 88x speedup for the 2D Ising model, marking the first application in material science and demonstrating WSE efficiency for low-data-intensity workloads.

DOE PROGRAM

Hydrogen with Carbon Management

NETL PARTNERS

cerebras

PSC

ANSYS®

U.S. Department of ENERGY | NATIONAL ENERGY TECHNOLOGY LABORATORY | 25 Years Powering Energy Innovation

NETL-R&D-FECM-HCM-000017