Composable Computing Ecosystems for Varied Scientific Approaches

Bronson Messer
Scientific Computing, National Center for Computational Sciences, Oak Ridge National Laboratory

The use of simulation across scientific fields is ubiquitous: computational science has been firmly ensconced as the third pillar of scientific enquiry, along with theory and experiment. It is important to note, however, that these pillars are not co-equal: experiment and observation retain primacy. If theory or computation do not connect to experiment, they are relegated from being science. Nevertheless, this conceptual structure will become less segregated over the next decade. New capabilities that enable the integration of computation, theory, and experiment promise new avenues to learn about nature, as data analytics and new methods for reasoning about data will lead to not only discovery, but to new hybrid modes of investigation.

As this new synthesis occurs, it is important to realize that the particular uses for simulation and the details of its connection to experiment vary across disciplines. There are classes of theory and simulation in which the tie to experiment is immediate and proximate. Examples of this type of science include phenomenological models in condensed matter physics and molecular dynamics. For these models, one can easily imagine a tight feedback loop between simulation and running experiment, where simulation is used to interpret and even steer the experiment. This would ideally happen in as close to real-time as possible.

Other theories are fundamental in nature and attempt to explain physical phenomena via the expression of first-principles models. Examples include nuclear structure physics and plasma physics. Here, the tie to experiment is also quite close, but is not as immediate—in either time or space—as in the previous examples. These calculations must agree with experimental results lest they be abandoned as not useful, but they are ultimately designed to provide a deeper understanding than is provided by the immediately useful phenomenological models.

Lastly, multi-physics simulations—like those used in earth system modeling, stellar astrophysics, and various mechanical engineering applications—are the most removed from informing (or being informed by) experiment "on the fly." These examples do have a multitude of physical observables that must be compared against simulation results, and such multi-physics models are, in every case, concerned with becoming predictive at the exascale.

These differences must be accommodated in any future computational ecosystem. The most flexible (and perhaps the most ambitious) way to enable this diversity is to allow practitioners themselves to compose the necessary computational ecosystems required from component parts, in both hardware and software. Heterogeneous compute, memory, storage, bandwidth (both internal and external), and the software stack could be provisioned dynamically to best serve the needs of the various modalities described above. Multi-physics simulations could describe memory and FP-intensive constructions, as opposed to configurations to enable effective streaming of data for the approaches tightly coupled to experiments at federated facilities. This dynamic provisioning will require (a) high-bandwidth, latency-hiding approaches to data movement (b) system software and middleware layers that can interoperate, and (c) abstractions powerful enough to allow the composers to make rational and straightforward decisions. Importantly, in addition to providing an essential software capability for many of these putative compositions, AI could also be used to help make these compositions performant. Having been trained on historical data for performance, such an AI system would allow practitioners to create new workflows that include both brand-new implementations and legacy software, all mapped to appropriate hardware to ensure performance.