Current Trends in the Large-Scale Simulation of Nuclear Reactor Flows

Elia Merzari, Paul Fischer, Andrew Siegel
Mathematics and Computer Science Division
Argonne National Laboratory

Computational fluid dynamics (CFD) is increasingly used to simulate nuclear reactor flows. Most CFD analysis, especially in industry, relies on the Reynolds-averaged Navier-Stokes (RANS) approach and traditional two-equation turbulence models. Higher-fidelity approaches for the simulation of turbulence, such as wall-resolved large eddy simulation (LES) and direct numerical simulation (DNS), can be done only on large supercomputing platforms. In fact, since the Reynolds number dictates the local resolution in wall-bounded flows, large machines are currently necessary to simulate engineering systems with turbulence-resolving techniques.

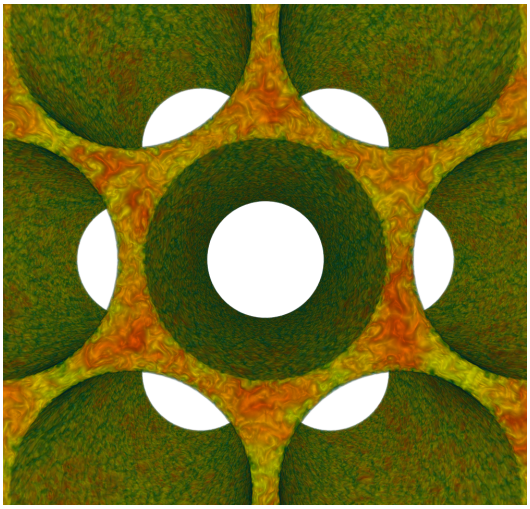

Fig. 1. Flow in a 37-pin rod bundle.

Nonetheless, petascale architectures are enabling the simulation of physical systems of increasing size and complexity. Current supercomputers have been used to simulate nuclear systems with grids that reach tens of billions of points, enabling the simulation of entire rod bundles (Fig. 1) with wall-resolved LES.

Codes that have achieved this grid sizes include the spectral-element code Nek5000 [1], PHASTA [2], and CODE_SATURNE [3]. Codes that have been used for petascale simulations of nuclear systems include HYDRA-TH, Open-FOAM, and DREKAR.

These codes share several characteristics on extreme-scale architectures:

1. Meshing is a bottleneck in the workflow at large scales.
2. Algorithms are latency bound at the strong-scaling limit.
3. The on-node performance is memory-bandwidth bound

4. Transients in turbulence-resolved simulations are characterized by massive time-scale separation due to the large system size and inherent transient scales (see Table 1).

Each of these performance challenges is expected to be amplified on future architectures. **Issues of scaling related to multiphysics coupling are largely unexplored.**

Table 1. Simulation data of four representative cases [4]. U is the bulk velocity; D is a characteristic spatial scale.

Simulation	Points/ (Pin*10D)	Reynolds Number	Time Scale [Lower]	Time Scale [Higher]
37-pin bare rod bundle	216 million	66,000	$2\ 10^{-4}$ D/U	10^{2} D/U
MASLWR core (59 pins)	1 million	~5000	$3\ 10^{-3}$ D/U	10^{6} D/U
19-pin SFR bundle	6 million	15,000	10^{-3} D/U	10^{2} D/U
2x2 LWR bundle	16 million	15,000	10^{-3} D/U	10^{2} D/U

The above-mentioned codes have been used to simulate – at ever increasing Reynolds numbers – part of the fuel assemblies that constitute the reactor core or other nuclear components (e.g., T-junctions, portions of the upper plenum) by using DNS/LES, as well as large portions of nuclear reactor vessels by using RANS or unsteady RANS (URANS). In both cases the trend has been toward simulating larger and larger systems rather than making problems run faster. Yet in the case of DNS/LES in particular, there are fundamental limitations to the margin for accelerating current turbulent transients.

In addition to the formidable range of temporal scales needed to be resolved, a fundamental computational science issue exists. Because of power constraints, high-performance architectures are being designed to support extreme concurrency. Unfortunately, little can be done to reduce internode latency, which sets the node-level granularity of simulations and, ultimately, the rate at which work can be done. Extreme concurrency, however, provides an avenue to solve larger problems rather than to solve today's problems faster (assuming, as in the present case, that we are already running at the strong-scale limit).

Unfortunately, this situation has consequences also at the exascale. The presence of these two constraints means that when running cases orders of magnitude larger (at higher Reynolds numbers or for larger domain size) for longer integration times, the time-scale separation will increase. Consequently, accelerating transients will increasingly become an imperative on larger architectures [4] – not only because solving today's problems faster will be of increasing importance to draw industry into using supercomputing platforms, but also because **without accelerating transients, the rate of return of increasing computational power will diminish.**

However, given the sheer scale of nuclear systems and the limited reach of current high-fidelity simulations, **weak scaling** likely will still play an important role in the near future, although probably at a lower rate of return. Foreseeable future applications include the following.

1. Full-core fluid calculations aimed at better predicting the steady-state performance. These will likely be conducted with hybrid RANS/LES. LES may be used to simulate a portion of a core, while the rest will be handled by RANS. Results of these calculations may provide power spectral densities for structural calculations.
2. DNS/LES for smaller sizes to benchmark RANS/hybrids, and potential multiscale applications involving LES informing RANS/design-level models dynamically.
3. Multiphase simulations of fuel assemblies, increasing from the current single-sub-channel simulations.
4. System-level coupling between petascale models (e.g., two petascale models of a assemblies coupled by a system-level algebraic constraint). This may involve integration with system (1D) models.
5. Conjugate heat transfer and coupling with neutronics (related to cases 1 and 4) and other multiphysics aspects (structural mechanics) at increasingly larger scales.
6. Uncertainty quantification of current problems, which would benefit from nearly perfect weak scaling.
7. Shape optimization of nuclear components (related to case 6).

These directions will provide enhanced methods to study a subset of advanced reactor designs and to aid the design of new reactors. **There is no obvious theoretical limit to achieving these goals**. Some algorithmic work will be necessary in order to achieve good scaling in cases 5, 6, and 4. Porting and optimization of codes to new architectures will be necessary. All these applications will benefit from natural and smooth transition from petascale to exascale because the level of resources can dictate the amount of resolution employed. For instance, in case 1 all assemblies will need to be simulated with RANS in order to model a full core on the Argonne Blue Gene/Q "Mira" supercomputer, while on Aurora already a few assemblies can be simulated by using LES.

Besides weak scaling, some important **additional applications merit special attention.**

1. Nuclear transients at full-core level to improve the accuracy of safety predictions and reduce margins. At present, safety analysis is performed only with massive simplifications (1D, porous media).
2. Inverse problems in safety analysis.
3. Acceleration of current petascale-level simulations.

All these will benefit from algorithmic advances in the area of reduced-order modeling to accelerate transients (although RANS may already be used today). Multiple ensembles, or ensemble averaging, should also be considered as a way to accelerate the collection of turbulent statistics as a means to achieve case 4. While these applications remain more challenging and less certain, **they provide a higher potential to affect the nuclear industry and bridge the gap between supercomputing and engineering practice.**

[1] http://nek5000.mcs.anl.gov
[2] https://secure.cci.rpi.edu/wiki/index.php/PHASTA
[3] http://www.code-saturne.org/cms/
[4] E. Merzari, A. Obabko, P. Fischer, N. Halford, J. Walker, A. Siegel, Y. Yu, "Large-Scale Simulation of Nuclear Reactors: Issues and Perspectives," in Proceedings of NURETH-16, August 30 – September 4, 2015, Chicago, IL