

# LARGE-SCALE SIMULATION OF NUCLEAR REACTORS: ISSUES AND PERSPECTIVES

**Elia Merzari, Aleks Obabko, Paul Fischer, Noah Halford, Justin Walker, Andrew Siegel**

Mathematics and Computer Science Division

Argonne National Laboratory

9700 S. Cass Avenue, Argonne, IL 60439, USA

emerzari@anl.gov; obabko@mcs.anl.gov, fischer@mcs.anl.gov; siegel@mcs.anl.gov

**Yiqi Yu**

Nuclear Engineering Division

Argonne National Laboratory

9700 S. Cass Avenue, Argonne, IL 60439, USA

yyu@anl.gov

## ABSTRACT

Numerical simulation has been an intrinsic part of nuclear engineering research since its inception. In recent years a transition is occurring toward predictive, first-principle-based tools such as computational fluid dynamics. Even with the advent of petascale computing, however, such tools still have significant limitations. In the present work some of these issues, and in particular the presence of massive multiscale separation, are discussed, as well as some of the research conducted to mitigate them.

Petascale simulations at high fidelity (large eddy simulation/direct numerical simulation) were conducted with the massively parallel spectral element code Nek5000 on a series of representative problems. These simulations shed light on the requirements of several types of simulation: (1) axial flow around fuel rods, with particular attention to wall effects; (2) natural convection in the primary vessel; and (3) flow in a rod bundle in the presence of spacing devices. The focus of the work presented here is on the lessons learned and the requirements to perform these simulations at exascale. Additional physical insight gained from these simulations is also emphasized.

## KEYWORDS

Large-Scale Simulations, Large Eddy Simulation, Exascale, CFD

## 1. INTRODUCTION

Computational fluid dynamics (CFD) is increasingly used to simulate nuclear reactor flows. Most CFD analysis, especially within industry, relies on the Reynolds averaged Navier-Stokes (RANS) approach and traditional two-equation turbulence models. Higher-fidelity approaches to the simulation of turbulence such as wall-resolved large eddy simulation (LES) and direct numerical simulation (DNS) remain limited to smaller applications or to large supercomputing platforms [1]. In fact, since the Reynolds number dictates the local resolution, large machines are currently necessary to simulate engineering systems with turbulence-resolving techniques.

Nonetheless, continued advances in supercomputing are enabling the simulation of physical systems of increasing size and complexity. Current supercomputers can accommodate grids that reach tens to hundreds of billions of points, enabling the simulation of entire rod bundles with wall-resolved LES using

CFD algorithms with good scalability properties. Examples of such simulations are discussed in this paper. These simulations can be used to gain unprecedented insight into the physics of turbulence in complex flows, and will become more widespread as petascale architectures become more accessible.

As the scale and size of LES and DNS simulations increase, however, the limitations of current algorithms become apparent. For larger systems, more temporal and spatial scale must be resolved, thus increasing the time-scale separation. While the smaller time scales dictate the size and the computational cost associated with each time step, the larger time scales dictate the length of the transient. An increased time-scale separation leads to smaller time steps and longer transients, eventually leading to simulations that are impractical or infeasible.

In practice the presence of multiple and strongly separated time scales limits the effectiveness of CFD algorithms for LES and DNS applied to large nuclear systems. Similar conclusions, although time-scale separation was not identified as the culprit, were reached in the case of a recent OECD benchmark [2]. In fact, in the PANDA experiment, the time-scale separation can be approximated as eight orders of magnitude, making turbulence-resolving techniques impractical despite the relatively contained size of the problem in terms of degrees of freedom. Even very efficient algorithms running at the strong-scaling limit will require years to compute this class of transients (we estimate eight years in the case of PANDA). Moreover, the situation is likely to become worse on future computer architectures, as even larger systems will be simulated, thus increasing the size and length of transients. At the same time transients currently simulated on petascale architectures are unlikely to become any faster on exascale architectures (Section 3.5).

One of the goals of The Center for Exascale Simulation for Advanced Reactors (CESAR) is to develop a coupled, next-generation nuclear reactor core simulation tool capable of efficient execution on exascale computing platforms. In the present manuscript we discuss a series of large-scale calculations performed with the thermal-hydraulics component, based on the massively parallel spectral element code Nek5000: (1) simulation of the axial flow around fuel rods, with particular attention to wall effects; (2) simulation of natural convection in the primary vessel; and (3) simulation of the flow in a rod bundle in the presence of spacing devices.

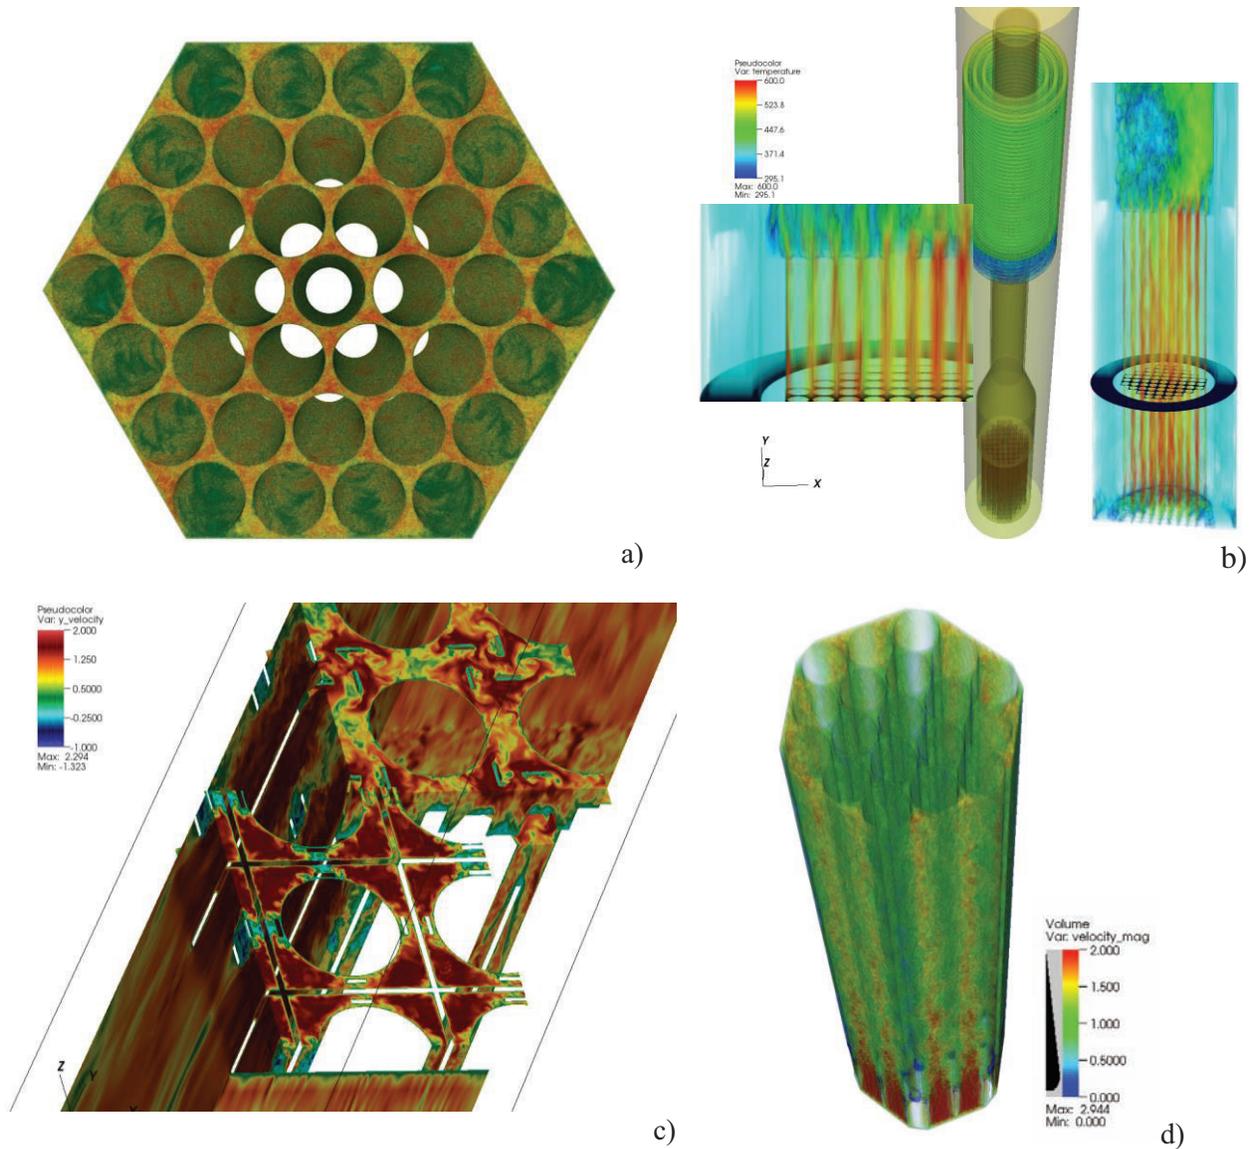
The focus of the work is on the lessons learned, the requirements to perform these simulations at exascale, and in particular the constraint posed by time-scale separation. Additional physical insight gained from these simulations is also emphasized. While a considerable verification and validation effort has been conducted for each class of simulations [3, 4, 5], these are not discussed in depth because of space limitations, but we direct the reader to relevant publications on this topic.

We also discuss novel algorithmic modifications that may help improve the current limitations of CFD algorithms for LES/DNS. The focus is on the use of accelerators, based on reduced-order models or other surrogates.

## **2. Nek5000**

The proposed method has been implemented in the Argonne-based open source fluid/thermal simulation code Nek5000 [6, 7], which has been designed specifically for transitional and turbulent flows in complex domains. Nek5000 is based on the spectral-element method (SEM) [8], a high-order weighted residual technique that combines the geometric flexibility of finite elements with the rapid convergence and tensor-product efficiencies of global spectral methods. Globally, the SEM is based on a decomposition of the domain into  $E$  smaller subdomains (elements), which are assumed to be curvilinear hexahedra (bricks) that conform to the domain boundaries. Locally, functions within each element are expanded as  $N$ th-order polynomials cast in tensor-product form, which allows differential operators on  $n^3$  gridpoints per element

to be evaluated with only  $O(n^4)$  work and  $O(n^3)$  storage. The principal advantage of the spectral-element method is that convergence is exponential in  $n$ , which implies that significantly fewer gridpoints per wavelength are required to accurately propagate a signal (or turbulent structure) over the extended times associated with flow simulations at high Reynolds number.



**Figure 1. Large-scale simulations performed with Nek5000. (a) 37-pin rod bundle at  $Re \sim 66,000$ ; (b) MASLWR simulation with temperature distribution [K] and volume rendering of the velocity field [m/s]; (c) contour plot of the velocity magnitude (normalized by bulk) for the flow through a periodic 2x2 pin bundle array with an AREVA spacer grid; (d) volume rendering of the velocity magnitude (normalized by bulk) 19-pin rod bundle with honeycomb spacer (SFR).**

In addition to its high-order foundation, Nek5000 has several other features that make it suited for large-scale parallel simulations. Temporal discretization is based on a high-order splitting that is third-order accurate in time and reduces the coupled velocity-pressure Stokes problem to four independent elliptic solves per time step: one for each velocity component and one for the pressure. The velocity problems are diagonally dominant and thus easily solved by using Jacobi-preconditioned conjugate gradient iteration.

The pressure substep requires a Poisson solve at each step, which is effected through multigrid-preconditioned GMRES iteration coupled with temporal projection to find an optimal initial guess. Particularly important components of Nek5000 are its scalable coarse-grid solvers that are central to parallel multigrid. The code features a fast direct solver that is optimal up to processor counts of  $p \approx 10^4$  and a fast algebraic multigrid solver for  $p = 10^5$  and beyond. Counts of 15 GMRES iterations per time step for billion-gridpoint problems are typical with the current pressure solver. Nek5000 scales extremely well on the BG/P and BG/Q architectures [7]. Time stepping is performed by using a characteristic scheme [9]. All LES calculations performed here use explicit filtering as a form of sub-grid scale (SGS) modeling, equivalent to deconvolution models [3, 10, 11]. We note that this class of SGS models is spectrally vanishing as the mesh resolution increases or in the limit of laminar flow [10,11].

### 3. LARGE-SCALE SIMULATION OF INTERNAL REACTOR FLOWS

In this section we describe a series of large-scale calculations of reactor internals performed as part of the CESAR program.

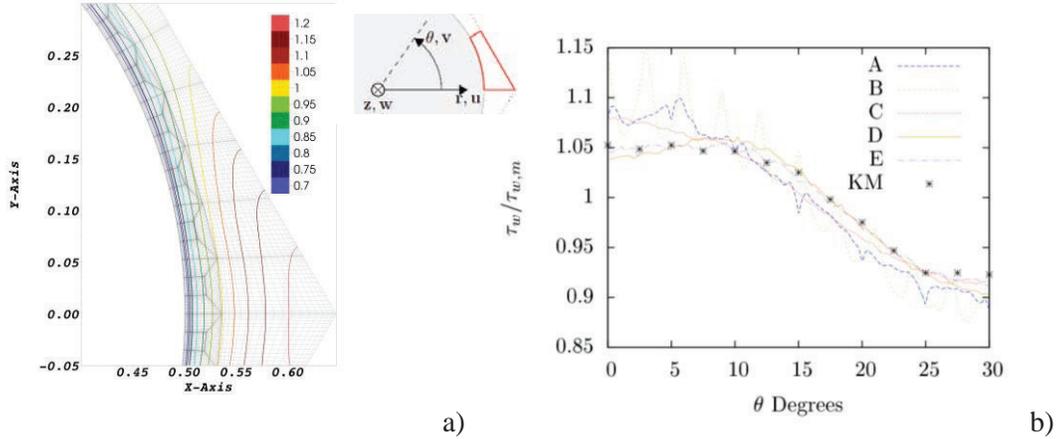
- Axial flow with focus on wall effects and turbulence. The axial flow in rod bundles has been the object of several investigations, since it is relevant for a variety of industrial applications (e.g., heat exchangers, nuclear reactor cores). For tight configurations (pitch-to-diameter ratio smaller than 1.1) the large difference in velocity within the cross section creates the possibility of a Kelvin-Helmholtz instability and the generation of a vortex street in the gap (“gap instability” [12]). This leads to a great enhancement in turbulent mixing but also to the introduction of strong, low-frequency modes that may induce fluid-structure interaction. This work aims at improving our understanding of the gap instability in a large bundle. We examine a large-scale wall-resolved LES simulation of a 37-pin hexagonal lattice rod bundle (Fig. 1a) at a pitch-to-diameter ratio of 1.12, typical of current advanced reactor designs. Although this configuration does not typically generate a vortex street in the center of the domain, it is proven here that the vortices can be observed in the proximity of the canister wall enclosing the rod bundle if the edge subchannels are sufficiently small.
- Core simulations with natural convection. Nuclear reactor cores are not isolated components; they are part of the primary system. Flow through the core is determined by the surrounding geometry, in particular when the flow is generated by natural circulation as proposed in some recent small modular reactor (SMR) designs. However, the simulation of a full primary system at high fidelity is prohibitive at present. In order to assess the feasibility at exascale, the simulation of an integral test facility, representative of an SMR, is conducted, namely, the MASLWR facility [13] (Fig. 1b).
- Flow simulations with spacing devices. Grid spacers are important components of nuclear reactor fuel assemblies because they allow separating the rods while enhancing turbulence. An LES has been conducted for a 19-pin SFR rod bundle with simplified spacers and for a 2x2 LWR subassembly. The data for the SFR rod bundle is compared to bare rod bundles to analyze the gap vortex network in this geometry. The LWR data is compared with that of the RANS simulations.

#### 3.1. Simulation of a 37 pin bare rod bundle

Examples of the grids used are shown in Fig. 2 (37-pin bundle). The figure depicts the elements and the actual collocation points for the coarse mesh. The discretization is designed to allow for at least one point near the wall at  $y^+ < 1$  and five points within  $y^+ < 10$ . Periodic boundary conditions are applied in the streamwise direction. The length in the streamwise direction is set to 10 diameters (D) for the 37-pin rod bundle case. This distance is deemed sufficient to describe the coherent structures, while limiting the computational cost. The importance of the streamwise length in these simulations has been discussed in several publications; see, for example, [14]. No-slip boundary conditions are applied at the wall.

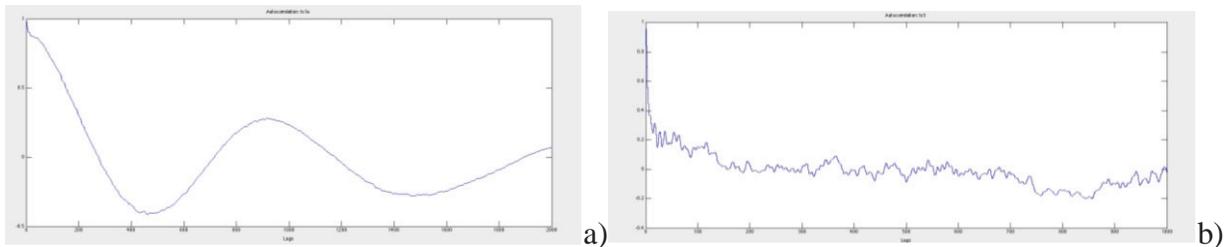
The azimuthal mesh requirements are finer in rod bundles when compared with channel flow because of the need to capture the azimuthal variation of the wall shear stress. This issue has been discussed at length

in [3], demonstrating that using higher-order polynomials as well as a higher resolution near the pin walls leads to significantly better results. This is further shown in Fig. 2b where different polynomial orders and meshes are compared for both the velocity field and the wall shear stress. One can see that with careful discretization, by clustering the mesh near the wall (Fig. 2a), results can be in excellent agreement with experimental data [15] – Fig. 2b.  $\tau_w$  is the wall shear stress and  $\tau_{w,m}$  is the mean value.



**Figure 2. 37-pin bare rod bundle case. (a) Typical mesh for a portion of single pin and averaged streamwise velocity distribution (normalized by bulk velocity); (b) azimuthal distribution of the wall shear stress at different resolutions as a function of the azimuthal angle on a single pin. Case E is equivalent to the case used here for the full bundle. KM is the experimental data [15].**

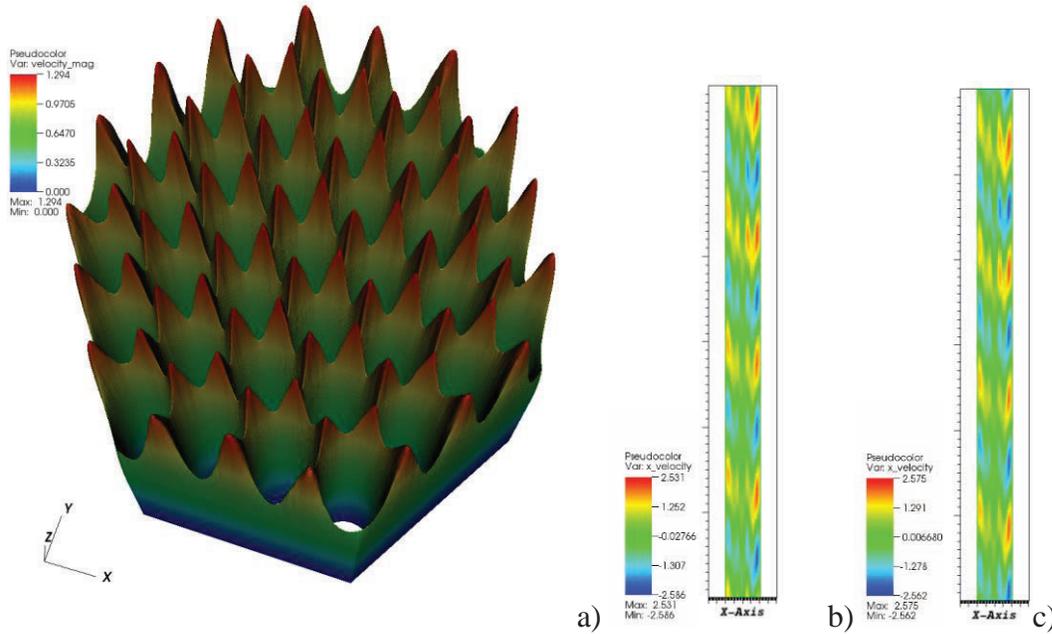
In the LES of rod bundles, the cost of reproducing accurately the wall dynamics and wall shear stress is high and requires optimization. In fact it constitutes the primary cost in wall-resolving LES and supersedes the cost of geometric complexity, as discussed below. For the 37-pin bundle case, run at a Reynolds number of approximately 66,000, over 8 billion collocation points (at 19th polynomial order) were judged necessary in order to achieve the accuracy described in Fig. 3. More than 500,000 MPI ranks were used to carry out the simulations (costing approximately 60,000,000 CPU-hours). The machine used was Mira, an IBM Blue Gene/Q system. The parallel efficiency for this problem proved excellent, exceeding 75% at 500,000 MPI ranks. The simulation was run for over 10 turnover times (ToTs) before collecting turbulence statistics (which included higher-order momentums such as skewness and kurtosis) for 10 additional ToTs.



**Figure 3. Time series analysis for 37-pin rod bundle. (a) Autocorrelation of cross velocity in a point near the canister wall; (b) autocorrelation of cross velocity for point in the middle of the domain.**

The simulation shows that for sufficiently small W/D (W is the distance between outer pins and canister walls) a strong difference exists between the dynamic behavior in the interior of the channel and the exterior. This has been observed for the first time in the simulations presented here; it has been discussed extensively in [16] and indicates the complexity of the flow in rod bundles. In this case, it is highlighted in Fig. 3, where the autocorrelation of the cross-velocity signal is compared in different parts of the

domain. Near the canister wall the autocorrelation exhibits the behavior typical of a gap vortex street, while in the center of the domain it presents the typical profile of incoherent turbulence (low correlation). The averaged velocity is shown in Fig. 4a.



**Figure 4. Averaged velocity data for the 37-pin rod bundle case. Values normalized by bulk velocity. (a) Velocity magnitude 3D plot, typical POD mode pair, for 37-pin rod bundle case with regional clipping, velocity in the x direction. Cross section at  $y=3.42 D$ : (b) Mode 3, (c) mode 4.**

One of the key methods to identify the gap vortex street is proper orthogonal decomposition (POD) [17]. It allows for the recognition of the most energetic modes of turbulence by performing an eigenvalue decomposition of the covariance matrix. Such a matrix is computed according to Sirovich [18]. A set of  $M$  snapshots of the flow field is collected, and the following eigenvalue-eigenvector problem is solved to obtain the POD eigenmodes:

$$C\vec{a}_k = \lambda_k \vec{a}_k, \quad (1)$$

where  $\vec{a}_k = a_n^k$  are the real coefficients that define the eigenmodes and  $\lambda_k$  are eigenvalues. The matrix  $C$  (size  $M \times M$ ) is defined as

$$C = C_{mn} = \frac{1}{M} (\vec{u}_m, \vec{u}_n), \quad (2)$$

where  $\vec{u}_m$  are snapshots of the velocity field. Each eigenmode can be then be computed by using

$$\vec{\sigma}_k(\vec{x}) \cong \sum_n a_n^k \vec{u}_n(\vec{x}), \quad (3)$$

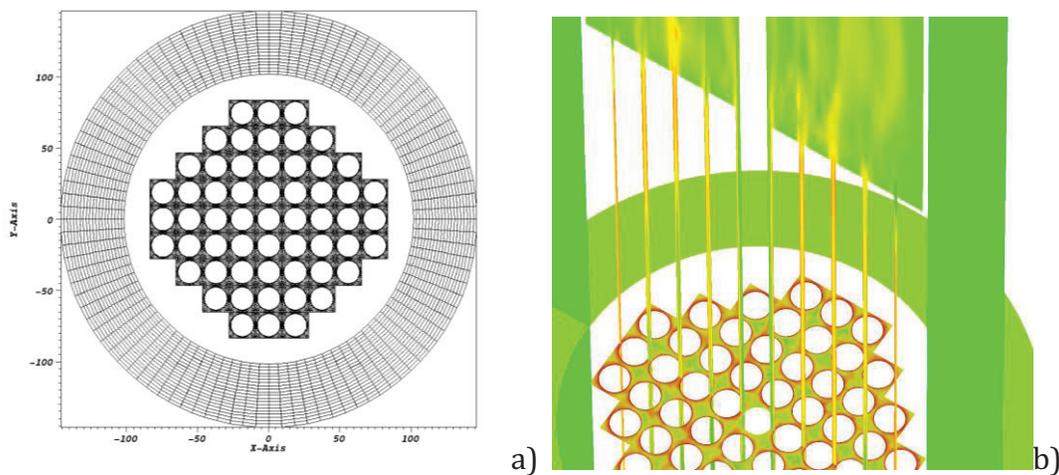
In order to examine the local structures, a regional clipping operator  $\hat{L}$  can be applied. The focus can be any region in the domain. Equation (2) becomes

$$C = C_{mn} = \frac{1}{M} (\hat{L}\vec{u}_m, \hat{L}\vec{u}_n). \quad (4)$$

where  $\hat{L}$  will generally be zero, the velocity vector where clipping is being applied. Over 3,000 snapshots (for 600 Tbyte of data, one of the largest POD ever performed) have been collected, and local clipping has been applied to them. This approach has allowed localizing the coherent structure in the outer subchannels. Eigenvectors in this region exist in pairs (Fig. 4b and Fig 4c), a characteristic of traveling waves [14]. POD applied without clipping produces a less clear structure and is of limited use in identifying and studying the gap vortex street.

### 3.2. Simulation of the MASLWR facility

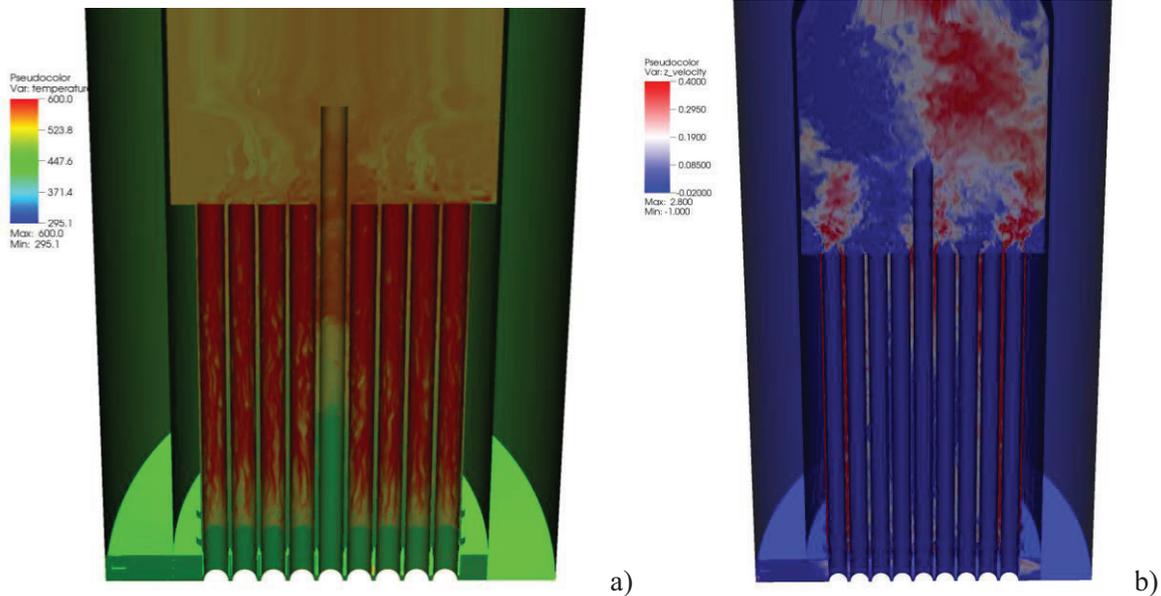
Besides expanding the number of pins of reactor core simulations, an additional expansion of high-fidelity computations may extend beyond the reactor core to investigate the fully primary system. These calculations may be instrumental in studying three-dimensional stability effects in SMRs. Although studying the full primary system for an SMR with DNS and LES may still be impossible, the investigation of integral test facilities is useful for examining the requirements and cost of these calculations at scale.



**Figure 5. Model MASLWR facility. (a) Mesh in the core region; (b) instantaneous temperature distribution.**

We focus here on the MASLWR integral test facility at Oregon State University [13], which is, however, not a CFD-oriented validation experiment. It comprises an integral PWR at scale with 59 heated pin rods in the core as well as a helical steam generator. The flow is driven entirely by natural convection. Nek5000 is characterized by  $\sim 850,000$  elements. The mesh is designed for a polynomial order of 11, for a total of roughly 1.5 billion collocation points. A low-order polynomial mesh is shown in Fig. 5a (it is much finer at higher polynomial orders). The mesh was generated by using the following guidelines:  $y^+ \sim 1$  for the first node and more than 5 points within the boundary layer,  $y^+ \sim 30/40$  spacing in the spanwise direction, and  $y^+ < 100$  spacing in the streamwise direction. The steam generator was simplified after several tests were conducted. We consider here a case at a fraction of roughly 20% of the nominal “core” power (80 kW).

The Reynolds number is widely changing (Fig. 5b), with highest value in the riser (the lowest value is in the core region,  $\sim 3,000$ ). The incompressible formulation and Boussinesq/full density change in the momentum equation. Very long integration times are required, given that ToT is 120 seconds, with a time step  $\sim 5 \times 10^{-5}$  s (driven by wall resolution). Since 1 second costs up to 1,000,000 CPU-hours, potentially the total of a transient could be in fractions of billions of CPU-hours (trillions of CPU-hours for the full SMR).



**Figure 6. Simulation on MASLWR facility. (a) Temperature distribution during a “frozen” velocity transient [K]; (b) velocity distribution during a “frozen” temperature transient [m/s].**

How can one speed convergence? As a first step, the initial condition is key in natural circulation. A mapping routine was written to load a solution from a commercial RANS code (STAR-CCM+) and map it to Nek5000. Even starting from a “good initial condition” is insufficient, however, since turbulence is not established. The following procedure has been used to accelerate convergence to a steady state:

- Alternate between “frozen velocity” and “frozen temperature” modes. The frozen velocity mode is about 20 times faster (huge savings); see Fig. 6.
- Continue with fully coupled system.

This led to a converged solution relatively fast—and to a substantially good agreement with available experimental data (Table I) for steady state conditions.

**Table I. Averaged Quantities for the 80 kW core power case.**

Quantity	Experiment	Calculation (Avg. over 10 s)
Flow rate	65.6 l/min	65.1 l/min
Inlet core T	489 K	489 K
Outlet core T	506 K	506.2 K

The cost of large-scale calculations of integral facilities is still prohibitive. But the problem is not the mesh size. Long integrals and slow convergence are the big issues; and they will remain so even at exascale. Techniques that speed convergence or time integration are critical. Here we demonstrated a crude technique. A possible improvement of this scheme is discussed in Section 4. The model developed as part of this work is being used to conduct very large stability calculations.

### 3.3. Simulation of a 19-pin SFR rod bundle

In addition to wall resolution and the complexity of simulating a primary, further complexity can be driven by the presence of reactor internals with an inherently complex geometry. This poses a challenge

in terms of computational requirements because the geometry needs to be resolved. In this section and the next we examine two representative calculations that include spacer grids for SFRs and LWRs.

The geometry investigated here is a triangular lattice SFR 19-pin bundle with a pitch-to-diameter ratio of 1.08. This value is consistent with values that are anticipated for long-life reactor cores. Typically SFR bundles are spaced by using wire wrappers. However, a minority of these reactors are spaced through grid spacers. These consist of honeycomb grids 3.125 pin diameters in length, in contact with each pin through a cosine-shaped dimple. In the present case the grids are separated by 25-pin diameters (Fig. 7a), and periodic conditions are applied in the streamwise direction at that distance [19]. In the cross section the blockage ratio reaches 40% of the flow area in the region of the spacers.

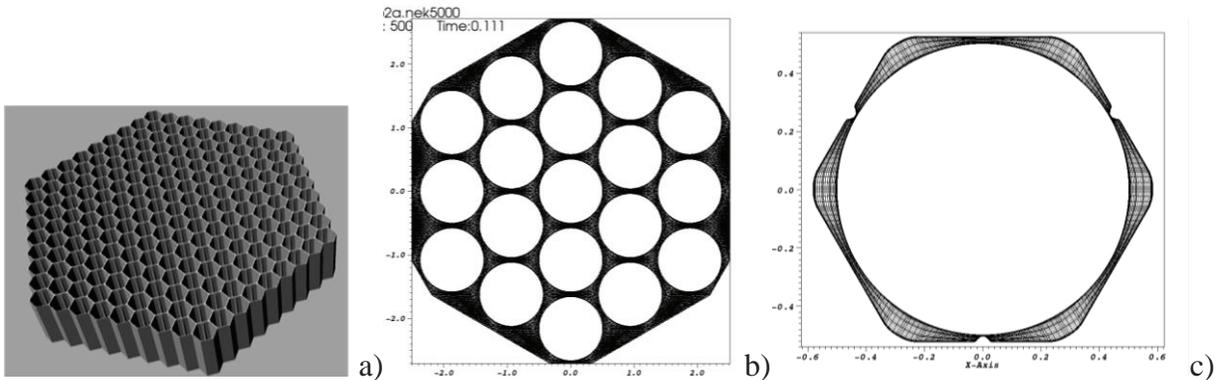


Figure 7. Mesh for the 19-pin SFR bundle. (a) Detail of spacers; (b) mesh far from spacers; (c) mesh in the region of the spacers (single-pin detail).

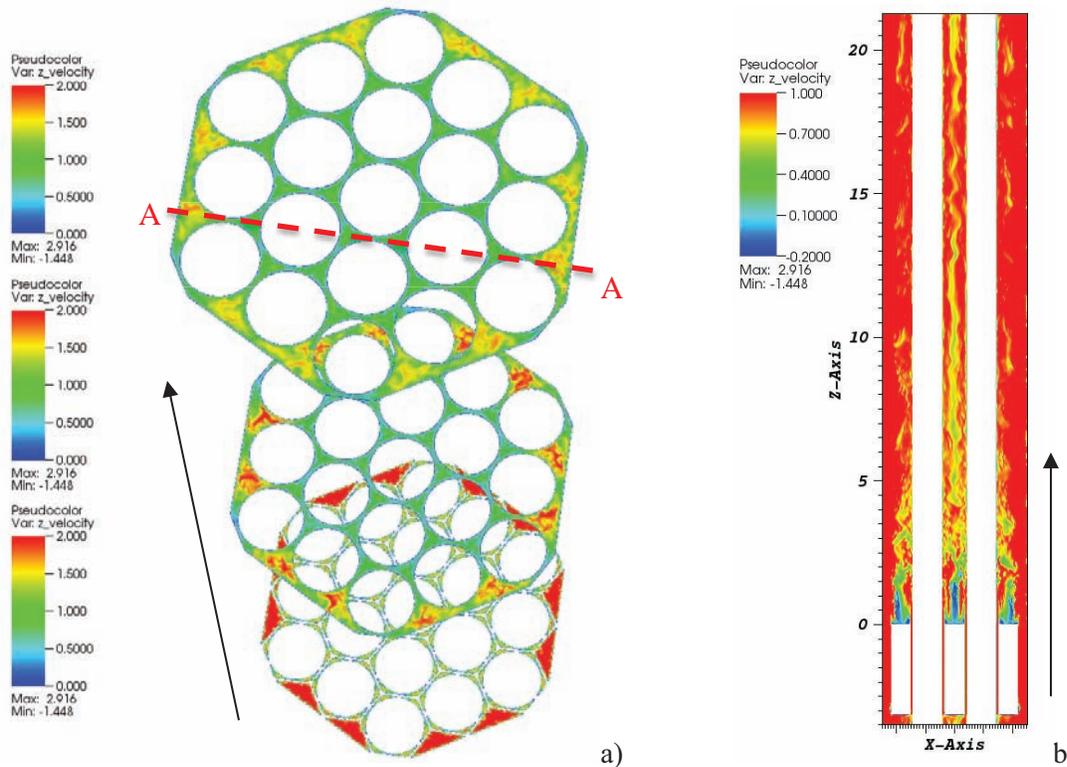
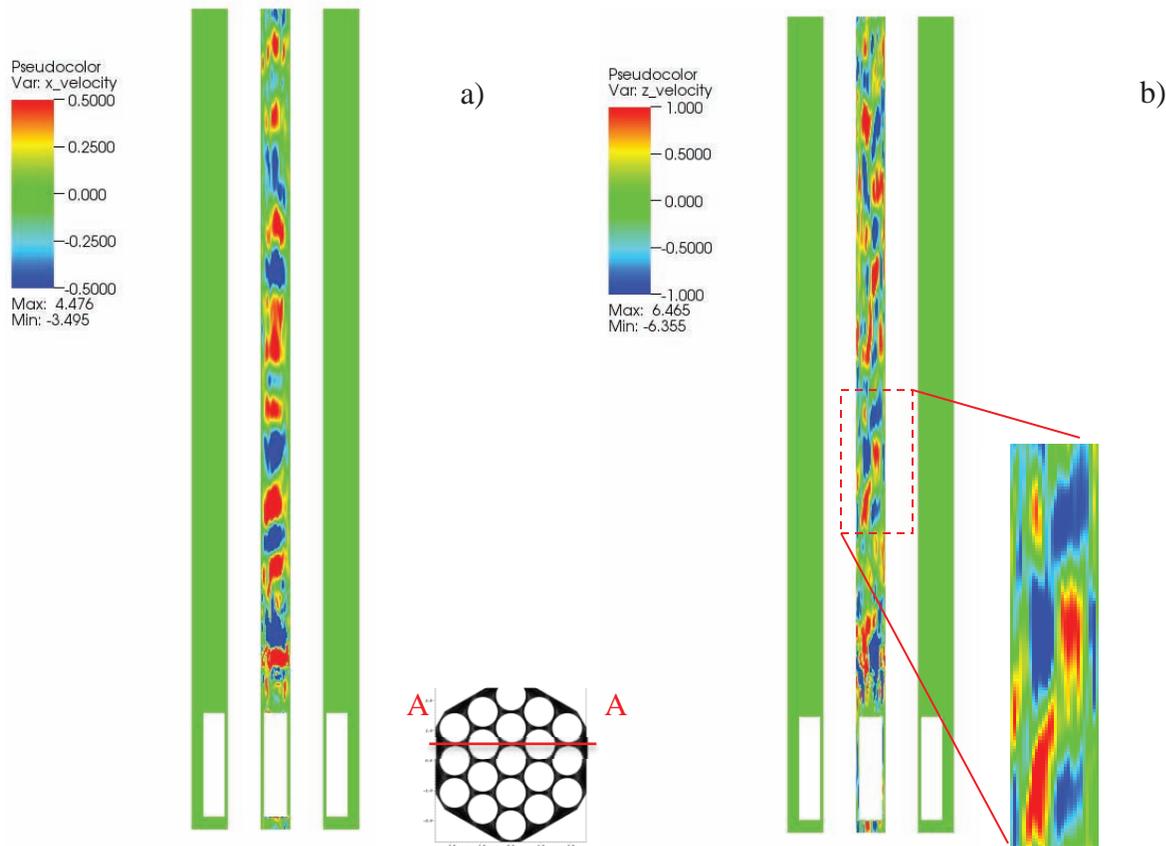


Figure 8. 19-pin SFR rod bundle, streamwise velocity normalized by bulk velocity. (a) Sections in the streamwise direction; (b) Section A-A.

The grid used is shown in Fig. 7b and Fig. 7c (single-pin section in the region of the spacers) for streamwise-normal cross sections. The figures depict the elements and the actual collocation points for the coarse mesh. The discretization is designed to allow for at least one point near the wall at  $y+\Delta y < 1$  and five points within  $y+\Delta y < 10$ . The corner channels have been removed in the 19-pin mesh to avoid excessive bypass flow. The simulations were carried out with up to 250,000,000 collocation points (40,000 elements/pin), which were found to be sufficient to describe the flow field accurately through a mesh refinement study. The computations were performed on up to 16,384 processors and required approximately 10 million CPU-hours to complete.

The Reynolds number considered here is 15,000. The simulation was conducted to examine the dynamic properties of the flow in this geometry. The flow is clearly edge-channel dominated, with higher flow in the outer channels (Fig. 8a). A gap vortex network has been observed (in Fig. 8b one can see the “sinusoidal path” in the streamwise direction). However, the vortex street is not as dominant as in the single-channel calculation because of the presence of the dominant edge-channels with no gap vortex street. An important consequence is that traditional POD did not lead to useful results for the present simulation. Regional clipping, defined in (4), is necessary. Time series analysis confirms the presence of the gap vortex network, with strong two-dimensional behavior in the center of the narrow gap (in the central channels); the effect of the spacers dies out quickly, within approximately 12 D.



**Figure 9. 19-pin SFR rod bundle, principal POD mode. Values normalized by bulk velocity. (a) Cross velocity; (b) streamwise velocity with detail.**

Figure 9 shows the most energetic POD mode obtained by processing approximately 1,000 snapshots (separated by 0.0004 ToT each). Regional clipping has been applied, and the focus is the region around the central pin. The mode presents a strong similarity with the principal mode observed in previous

analysis [19]. It appears to be less regular, however, because of a combination of the effect of regional clipping and the additional complexity brought by the gap vortex network. Additional POD, focused on larger or different regions of the domain, will need to be performed in order to examine in more detail the dynamics of the vortex network. An additional conclusion is that although spacers introduce complexity into the gap vortex network dynamics, they do not preempt its development. In fact, honeycomb spacers cannot be used to prevent the gap vortex street from developing, as shown in recent studies [20].

### 3.4. Simulation of a 2x2 LWR periodic rod bundle array

In this section we discuss the simulation of the flow in a rod bundle with a spacer grid and mixing vanes of the swirl type. The primary objective here is to examine the computational requirements for a calculation of this type, and in particular the additional burden imposed by the complexity of the geometry. Only minimal approximations were made in simulating the springs and mixing vane. A 2x2 section of rod bundle was simulated by using periodic boundary conditions in the cross section. Inlet-outlet simulations were used in the streamwise direction. Turbulence at the inlet was generated using recirculating boundary conditions [5]. The Reynolds number based on the pin diameter chosen is 15,000.

No detailed validation data was available for this geometry, but the same simulation practices were used to simulate the MATIS-H experiment [4], where the Nek5000 submission ranked highly, especially for turbulence predictions. The methods have been adapted to a spacer grid with additional complexity (springs). Approximately 300,000 elements were used at a polynomial order of 12, giving at least one point near the wall at  $y^+ < 1$  and at least seven points within  $y^+ < 10$  everywhere. Lower-order results (up to the 8th order) yielded similar results. The CFL in the simulations was kept between 1 and 2 using a characteristic scheme [8]. The computational cost is 3 seconds per time step on 32,768 processors (Blue Gene/Q).

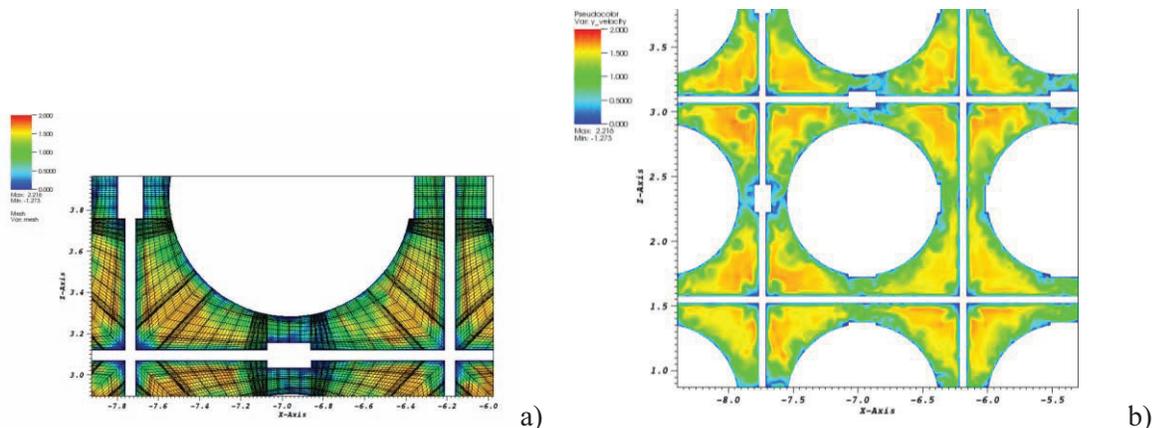
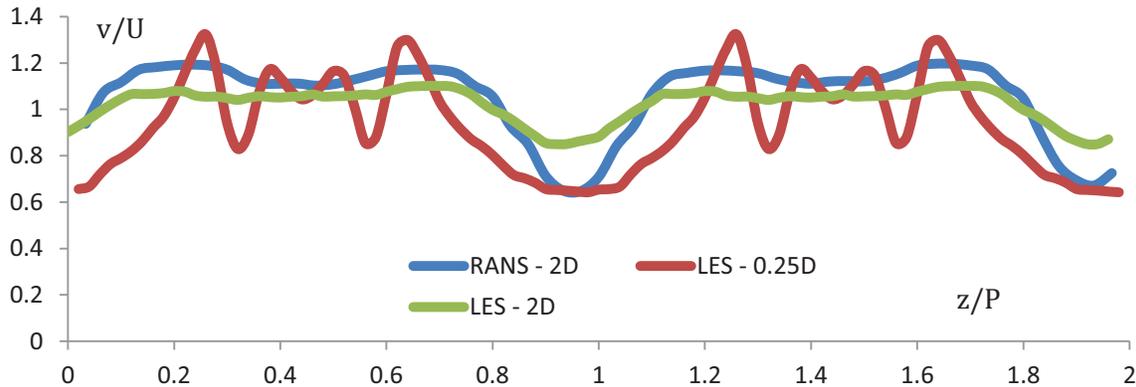


Figure 10. 2x2 LWR rod bundle case. (a) Detail of mesh. (b) Instantaneous snapshot of the streamwise velocity normalized by bulk velocity.

The results were compared with steady-state RANS results ( $k-\epsilon$  realizable with two-layer wall modeling) obtained with STAR-CCM+, version 8. As expected, the RANS results underestimate the mixing and turbulence levels downstream, whereas in the LES simulation the turbulence levels higher. Additional standard RANS simulations performed with different turbulence models ( $k-\omega$  SST, standard  $k-\epsilon$  with Low-Reynolds modeling) show similar trends. This is by no means an exhaustive test, however, and other models are expected to perform better. Figure 11 shows the streamwise velocity distribution on Section A-A (Fig. 10) as a function of the curvilinear position. The LES results two diameters downstream of the mixing vane predict a more uniform profile. The RANS results are obtained used a

trimmed mesh and were confirmed with a mesh refinement study using three grids with increasing mesh resolution.



**Figure 11. 2x2 LWR rod bundle case. Streamwise velocity at 0.25 D and 2D downstream of the spacer. LES and RANS calculations. P is the pitch, U is the bulk velocity, v is the velocity in the stream-wise direction.**

We reiterate that extensive work has been done on assessing CFD methods for this class of geometries, including an OECD/NEA benchmark [21]. This work has little to add to the physical understanding of the class of flows to those studies. It highlights, however, the potential to use high fidelity simulation to inform or improve turbulence models at low to medium Reynolds numbers (of interest in some accident conditions) especially when detailed experimental data is not available.

### 3.5. Discussion

Table II summarizes the time scales and resolution of the four cases investigated. The most stringent resolution requirement is driven by the need to resolve the near-wall dynamics in wall-resolving simulations. Although more complex geometries tend to have inherently higher element counts, the requirements associated with resolving the geometry (and not the wall dynamics) do not scale with the Reynolds number with the same steep relationship.

**Table II. Summary of minimum mesh resolution and time-scale range. U is the bulk velocity.**

Simulation	Points/ (Pin*10D)	Reynolds Number	Time Scale [Lower]	Time Scale [Higher]
37-pin bare rod bundle	216 million	66,000	$2 \cdot 10^{-4}$ D/U	$10^2$ D/U
MASLWR core (59 pins)	1 million	~5000	$3 \cdot 10^{-3}$ 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 smallest scale decreases in value with the smallest Reynolds number. The overall number of degrees of freedom determines the cost per time step, but the time-scale separation represents an independent cost. The largest time scale is determined by external factors (domain size, inherent transient). We note that for

MASLWR, the simulation of nuclear transients (time scale in the order of  $\sim$ days) would increase the time scale separation by 3 orders of magnitude.

At exascale, in addition to this formidable range of temporal scales a fundamental computational science constraint is likely to be at work. Because of power constraints, high-performance architectures are being designed to support extreme concurrency. Unfortunately, there is little that 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 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). That will mean running cases orders of magnitude larger (at higher Reynolds numbers or for larger domain size) for longer integration times, as the time scale separation will increase. As a consequence, accelerating transients will become an imperative on larger architectures.

#### 4. ACCELERATING TRANSIENTS

For problems such a transient for MASLWR, even at the strong-scaling limit, simulations may require several continuous years of calculation ( $\sim$ 1-10 years), or potentially more than the lifetime of available machines. On the other hand, if only heat transport were to be simulated in the problem as suggested in Section 3.2, assuming an imposed flow field, the cost would be dramatically lower, and the same calculation could be performed on higher processor counts (the advection-diffusion equation is more scalable than Navier-Stokes), leading to a projected simulation time of 15 days. Savings can be even more dramatic if the problem allows for coarser grids for heat transfer. Unfortunately, the velocity is not known a priori, and it is time-varying. However simple, this approach is the most rudimentary form of multiscale decoupling: the most expensive time scale is removed and replaced by a surrogate.

Many choices for surrogate are possible. We consider here a surrogate produced by solving a reduced-order model comprising a set of  $N$  nonlinear ordinary differential equations. As a first case of interest we examine ergodic flows with constant properties and no buoyancy. Future work will be dedicated to relaxing these assumptions [22]. An error is associated with the introduction of a surrogate. Given the impossibility of performing long transients at full high fidelity, such an error may be unavoidable. An important continuation of the present work will be systematic verification of the errors introduced by the surrogates against DNS or wall-resolved LES data for a set of benchmark cases. The objective of this strategy is to construct a computationally inexpensive, verified surrogate.

The reduced-order model here is based on POD [23, 17]. Historically the main concern for POD-based reduced-order modeling is that it tends to lead to poor results *when extrapolating out of its parameter space*. In fact, although a POD series is optimal in energy content, closures are necessary because of series truncation. Nevertheless, the error introduced by truncation, although typically growing in time, can usually be bounded both a priori [24] or a posteriori; and a considerable amount of research has been done in this direction. Moreover, an increase in dimensionality of the ROM system leads to lower errors.

The ROM in the method proposed here is not computed once and for all and used to extrapolate out of its parameter space, but it is generated at run time for each specific set of parameters (e.g., the Reynolds number). An optimization stage may also be included in which the surrogate is tuned to perform the best possible results. Several optimization strategies are under consideration, but they will not be discussed here. The power of POD lies in the possibility of using POD modes to reduce the Navier-Stokes equations for mass and momentum conservation,

$$NS(\vec{x}, \vec{u}) = 0, \quad (5)$$

to a set of ODEs by using the expansion  $\vec{u} = \sum_i c_i \vec{\sigma}_i$  ( $c_i$  are the coefficients) and exploiting the orthonormality of the POD eigenfunctions.

$$(\vec{\sigma}_k, NS(\vec{x}, \sum_i c_i \vec{\sigma}_i)) = 0 \quad (6)$$

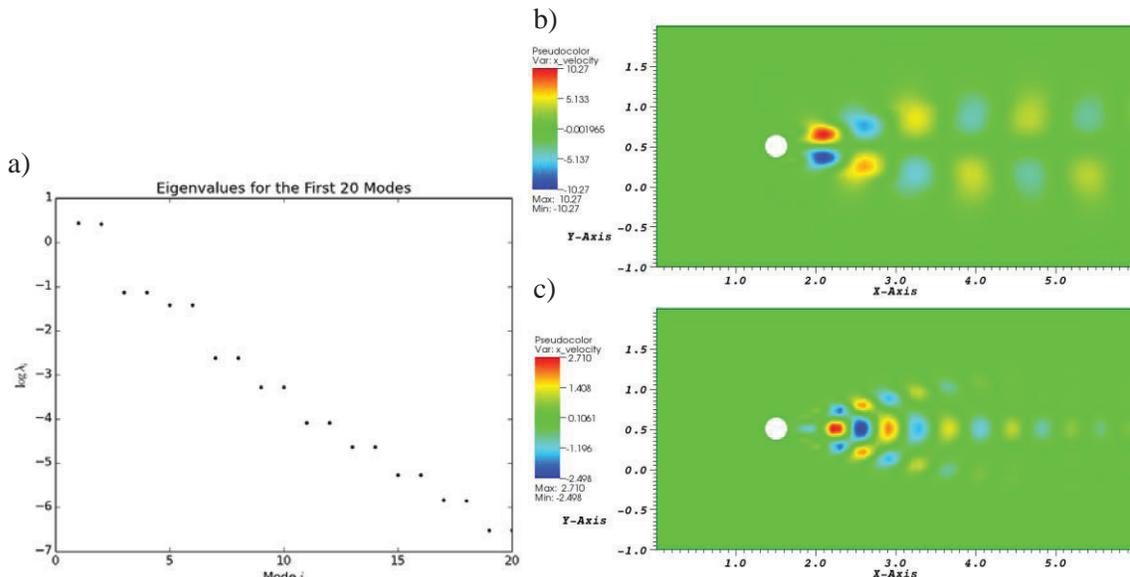
Solving such a system would be considerably simpler numerically than solving the Navier-Stokes equations directly. Since the POD eigenfunctions are divergence free (they are a linear combination of divergence-free functions), the mass conservation equation becomes superfluous. The momentum conservation equation is reduced to a set of ODEs of the following type.

$$\frac{dc_k}{dt} = A_{ijk} c_i c_j + B_{ik} c_i + C_k \quad (7)$$

The nonlinear coefficients  $A$  represent higher-order convective transport between modes. The linear terms  $B$  represent the effect of viscosity and the energy transfer with the mean flow field. The forcing term  $C$  represents the coupling between modes and the Reynolds stresses. The coefficients are computed by using equation (6). Other formulations of (7) are possible, depending on the treatment of the forcing term and the addition of closure models to compensate for the effect of truncation. This is a rich research area where several models have been proposed in the past two decades [25, 26]. Moreover, we note here that if LES data, and not DNS data, is used to generate the reduced basis in (6), (7) will need to be modified to account for additional terms. Equation (7) could then be used to reconstruct the flow field and solve the transport equation for a scalar  $s$  on longer time scales.

$$\frac{\partial s}{\partial t} + \sum_k c_k \vec{\sigma}_k \cdot \nabla s + \langle \vec{u} \rangle \cdot \nabla s = D \Delta s \quad (8)$$

where  $D$  is the diffusion coefficient for the scalar  $s$  and  $\langle \vec{u} \rangle$  is the mean velocity.

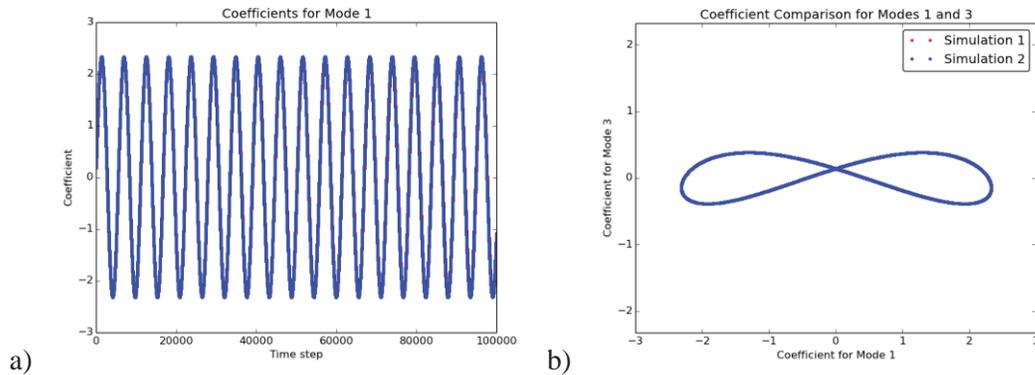


**Figure 12. 2D flow past a cylinder. (a) Eigenvalue structure; (b) most energetic mode; (c) second most energetic mode.**

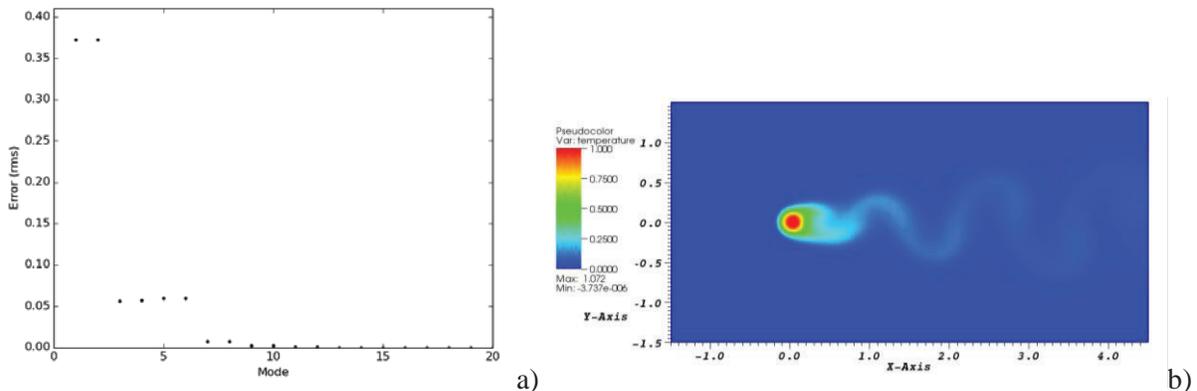
Let us examine a simple transient conjugate heat transfer problem: two-dimensional flow past a cylinder at  $Re=100$  (Fig. 12). The domain size is approximately equal to 30 times the cylinder diameter in the

streamwise direction and 10 times the hydraulic diameter in the other. In the streamwise direction inlet/outlet boundary conditions are used with constant velocity on the inlet and a Neumann boundary on the outlet. Periodic boundary conditions are assumed in the stream-normal direction.

A series of snapshots is collected, and POD is then performed to recognize the most important modes (the most energetic eigenmodes are shown in Fig. 12). The first 20 modes (>99% of the total energy) are used to reconstruct the flow field by using equation (7); the coefficients are computed by projecting the Navier-Stokes equations on each mode. The time histories of the eigenmodes can then be compared with the projections (which are computed by projecting DNS data on the POD modes). Figure 13 shows the phase space between different eigenmodes for both projections and coefficients ( $m$  represents the order of the mode). One can see that the agreement is nearly perfect (blue lines refer to the projections, obtained).



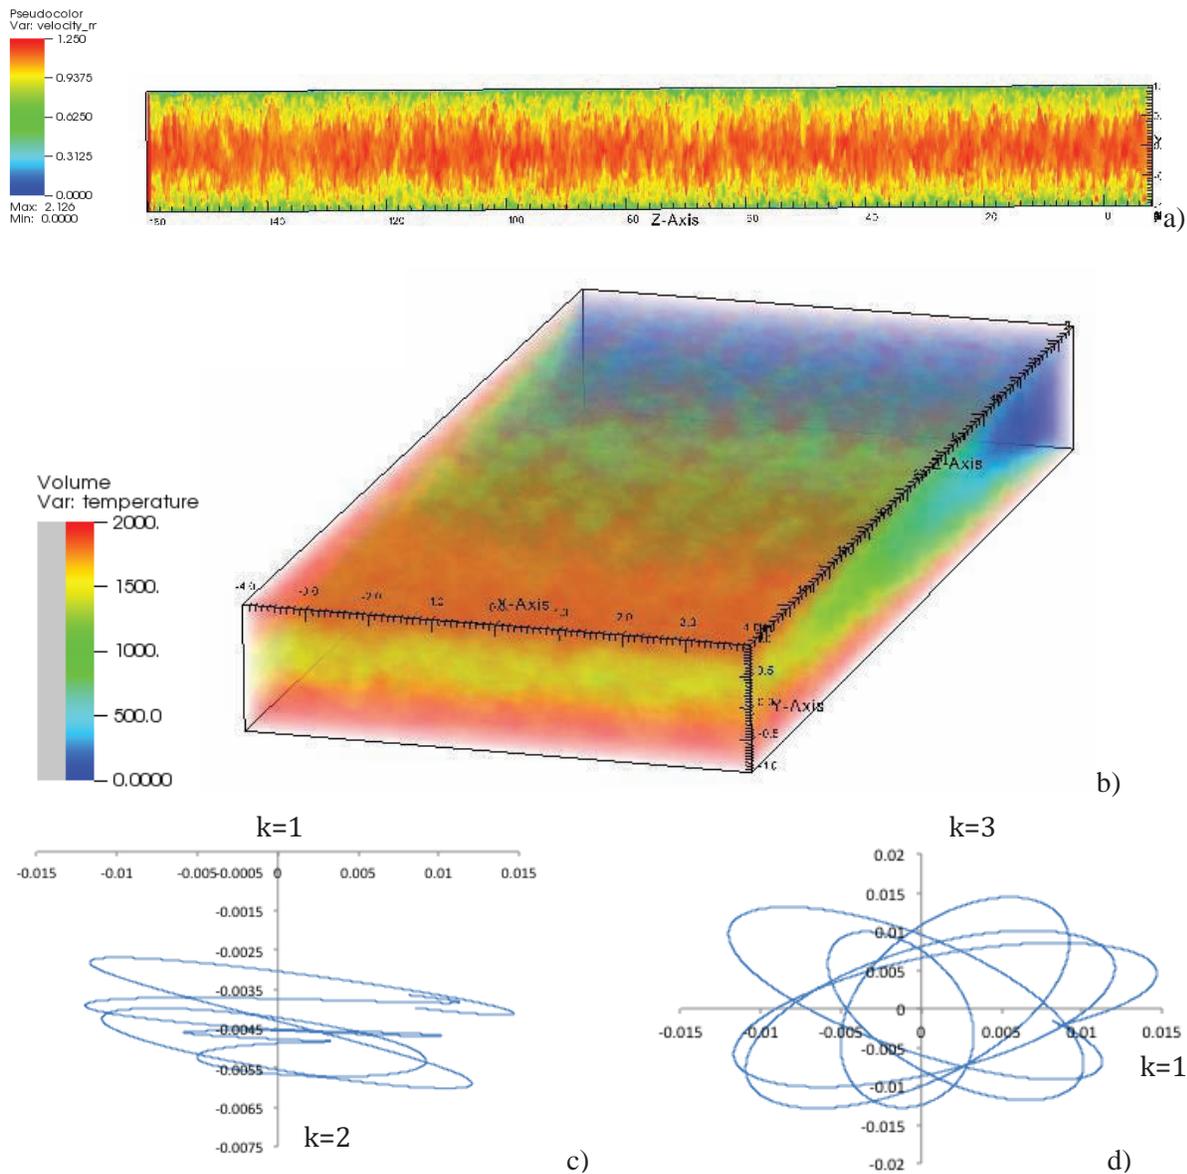
**Figure 13. 2D flow past a cylinder, with projections (blue) and time history of the coefficients (red). (a) Time history for the first mode; (b) phase plot between different modes.**



**Figure 14. Flow past a cylinder. (a) Error as a function of the number of modes included in the reduced-order model. (b) Transient temperature reconstructed by using equation (8).**

Since the dynamics is so well reproduced (Fig. 13), it is not surprising that the snapshot is nearly undistinguishable from the reconstructed flow field. Figure 14a shows the sum of the total error (root mean square) of the coefficient at time step 100,000 (200 flow-through times) as a function of the number of modes included in the reduced-order model. The error decreases dramatically if at least four modes are included. The reconstructed flow field is used to provide the advection term in the energy equation. A conjugate heat transfer transient is simulated where the power density in the cylinder presents an oscillatory behavior summed to a ramp. A snapshot of the temperature distribution is shown in Fig. 14b. The results are virtually identical to the results obtained with a traditional conjugate heat transfer calculation, at a cost lower than one-tenth. Higher savings are expected in three-dimensions, due to the

larger dimensionality of the system, and typical higher cost associated to the hydrodynamic solver in three-dimensional turbulence cases.



**Figure 15. Transient heat transfer in a channel. (a) Velocity field for benchmark case (normalized by bulk velocity); (b) volume rendering of the temperature distribution for benchmark case (arbitrary units). (c-d) phase plots for principal POD modes (modes  $k=1$ ,  $k=2$ ,  $k=3$ ).**

## 5. CONCLUSIONS

Several large-scale simulations have been conducted in representative simulations in nuclear engineering. Numerous issues have been discussed, but the most significant issue has been identified in the time-scale separation. The separation between the smallest and largest scales increases with the Reynolds numbers and the overall dimensions of the system. This situation hinders the capacity of supercomputing to address nuclear engineering problems of great impact.

A methodology to accelerate transients has been discussed, implemented, and partially tested on 2D and 3D flows. The 3D testing (Fig. 15) is being conducted on a proposed benchmark that includes transient heat transfer with a linearly increasing heat flow in time [27]. The focus has been limited to ergodic steady turbulent flows where the heat or scalar source is transient in time. By employing reduced-order models to create a surrogate for the solution of the Navier-Stokes equations, significant acceleration has been observed, ranging from one order of magnitude in the case of 2D flows to two orders of magnitude for 3D flows.

We note that RANS and multiscale decoupling are not uncorrelated. The solution of the RANS equations can be used as a surrogate; however, the error cannot be easily bound, since unfortunately no general satisfactory closure relationship is available, despite decades of research. In fact, RANS models developed for internal flows do not usually fare well for mixing problems. In general, RANS and the approach described have different applicability and optimal areas of application. However, the potential exists to combine RANS with the methods described in this work, and this possibility will be carefully considered.

## ACKNOWLEDGMENTS

This material was based upon work performed at the DOE Office of Science User Facility ALCF (Argonne Leadership Computing Facility) and funded by the U.S. Department of Energy, Office of Science, under contract DE-AC02-06CH11357, as part of the CESAR program.

## REFERENCES

1. G. Grötzbach and M. Wörner, “Direct numerical and large eddy simulations in nuclear applications.” *International Journal of Heat and Fluid Flow*, **20**, pp. 222–240 (1999).
2. M. Andreani, A. Badillo, and K. Kapulla, “Synthesis of the OECD/NEA-PSI CFD benchmark exercise.” In Proceedings of CFD4NRS5, Zurich, Sept. 9-11, 2014.
3. J. Walker, E. Merzari, A. Obabko, P. Fischer, and A. Siegel, “Accurate prediction of the wall shear stress in rod bundles with the spectral element method at high Reynolds numbers.” *International Journal of Heat and Fluid Flow*, **50**, pp. 287-299 (2014).
4. A. V. Obabko, A. E. Aksenova, V. V. Chudanov, P. F. Fischer, V.M. Goloviznin, A.A. Kanaev, S. Karabasov, E. Merzari, S. Parker, V. A. Pervichko, W. D. Pointer, T. Tautges, M. A. Zaytsev, “Cross-verification and validation simulations of Matis benchmark,” ANL/TM-341, 2015.
5. A. V. Obabko, P. F. Fischer, T. J. Tautges, V. M. Goloviznin, M. A. Zaytsev, V. V. Chudanov, V. A. Pervichko, A. E. Aksenova, and S. Karabasov. Large Eddy Simulation of Thermo-Hydraulic Mixing in a T-Junction. In Donna Guillen, editor, Nuclear Reactor Thermal Hydraulics and Other Applications. November 2012
6. P. Fischer *et al*, Nek5000 website. <https://nek5000.mcs.anl.gov> (2014)
7. P. F. Fischer, J. Lottes, W. D. Pointer, and A. Siegel, “Petascale algorithms for reactor hydrodynamics.” *Journal of Physics Conference Series* (2008).
8. Y. Maday, A. T. Patera, and E. M. Ronquist, “An operator-integration-factor splitting method for time-dependent problems: Application to incompressible fluid flow.” *Journal of Scientific Computing*, **5**, pp. 263–292 (2009).
9. P. Fischer, “Implementation considerations for the OIFS/characteristics approach to convection problems.” <http://www.mcs.anl.gov/~fischer/nek5000/oifs.pdf>
10. S. Stolz, P. Schlatter and L. Kleiser, “High-pass filtered eddy-viscosity models for large-eddy simulations of transitional and turbulent flow”, *Phys. Fluids*, **17**, 065103 (2005)
11. P. Schlatter, S. Stoltz and L. Kleiser “LES of transitional flows using the approximate deconvolution model”, *International Journal of Heat and Fluid Flow*, **25**, pp. 549–558 (2004)

12. G. H. Choueiri and S. Tavoularis, "Experimental investigation of flow development and gap vortex street in an eccentric annular channel. Part 1. Overview of the flow structure." *Journal of Fluid Mechanics*, **752**, pp. 521- 542 (2014).
13. J. N. Reyes, J. Groome, B. G. Woods, E. Young, K. Abel, Y. Yao, and Y. J. Yoo, "Testing of the multi-application small light water reactor (MASLWR) passive safety systems." *Nuclear Engineering and Design*, **237**, pp.1999–2005 (2007).
14. E. Merzari and H. Ninokata, "Proper orthogonal decomposition of the flow in a tight-lattice rod-bundle." *Nuclear Engineering and Design*, **241**, 4621-4632 (2011).
15. T. Krauss and L. Meyer, "Experimental investigation of turbulent transport of momentum and energy in a heated rod bundle." *Nuclear Engineering and Design*, **180**, 185-206 (1998).
16. E. Merzari, J. Walker, and P. Fischer "Large-scale simulation of rod bundles: Coherent structure recognition and stability analysis." In Proceedings of AJK2015, Seoul, July 2015.
17. G. Berkooz, P. Holmes, and J. L. Lumley, "The proper orthogonal decomposition in the analysis of turbulent flow." *Annual Reviews of Fluid Mechanics*, **25**, pp. 539–575 (1993).
18. L. Sirovich, "Turbulence and the dynamics of coherent structures, Part I: Coherent structures." *Quarterly of Applied Mathematics*, **45**, 561–571 (1987).
19. E. Merzari, W. D. Pointer, P. Fischer, and H. Ninokata, "Numerical simulation of the flow in a tight-lattice SFR rod-bundle with grid spacers." In Proceedings of NURETH-14, Toronto, Canada, Septe.2011.
20. E. Merzari and P. Fischer, "Controlling turbulence in tight rod bundles." In Proceedings of FEDSM2013, Lake Tahoe, July 2013.
21. Ryong Lee, Jungwoo Kim, and Chul-Hwa Song, "Synthesis of the OECD/NEA-KAERI Rod bundle CFD benchmark exercise." In Proceedings of CFD4NRS-4, Daejeon, S. Korea, Sept. 10-12, 2012.
22. S. Deparis and G. Rozza, "Reduced basis method for multi-parameter-dependent steady Navier-Stokes equations: Applications to natural convection in a cavity." *Journal of Computational Physics*, **228**, pp. 4359–4378. (2009).
23. E. Merzari, W. D. Pointer, and P. Fischer, "A POD-based solver for the advection-diffusion equation." In Proceedings of AJK2011, Hamamatsu, Japan, 2011.
24. C. Homescu, L. R. Petzold, and R. Serban, "Error estimation for reduced-order models of dynamical systems." *SIAM Journal of Numerical Analysis*, **43**, pp. 1693–1714 (2006).
25. N. Aubry, P. Holmes, J. L. Lumley, and E. Stone, "Reduced basis approximation and a posteriori error estimation for the time-dependent viscous burgers equation." *Journal of Fluid Mechanics*, **192**, pp. 115– 173 (1998).
26. M. Bergmann, C. Bruneau, and A. Iollo, "Enablers for robust POD models." *Journal of Computational Physics*, **228**, pp. 516–538 (2009).
27. E. Merzari, N. Halford, P. Fischer, and A. Siegel, "Surrogate-based methods for the acceleration of high-fidelity transient simulations." In E-Proceedings of the 36th IAHR World Congress (2015).

The submitted manuscript has been created by UChicago Argonne, LLC, Operator of Argonne National Laboratory ("Argonne"). Argonne, a U.S. Department of Energy Office of Science laboratory, is operated under Contract No. DE-AC02-06CH11357. The U.S. Government retains for itself, and others acting on its behalf, a paid-up nonexclusive, irrevocable worldwide license in said article to reproduce, prepare derivative works, distribute copies to the public, and perform publicly and display publicly, by or on behalf of the Government.