SANDIA REPORT
SAND2023-12700
Printed February 2023

Sandia National Laboratories

Full Order/Reduced Order Modeling Thermal Analysis Comparison of Crash-Burn Scenario using Aria and Pressio_Aria

Flint Pierce
John Tencer
Victor Brunini
Francesco Rizzi

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico
87185 and Livermore,
California 94550

Sandia National Laboratories
U.S. DEPARTMENT OF ENERGY NNSA

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

Issued by Sandia National Laboratories, operated for the United States Department of Energy by National Technology & Engineering Solutions of Sandia, LLC.

Available to DOE and DOE contractors from
U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-Mail: reports@osti.gov
Online ordering: http://www.osti.gov/scitech

Available to the public from
U.S. Department of Commerce
National Technical Information Service
5301 Shawnee Rd
Alexandria, VA 22312

Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-Mail: orders@ntis.gov
Online order: https://classic.ntis.gov/help/order-methods/

ABSTRACT

This report is intended to describe the use of Reduced Order Modeling (ROM) via the Sierra module *Pressio_Aria* [Brunini 2022] in the context of a multiphysics application referred to here as the Crash-Burn scenario [Tan-Torres 2020]. Specifically, we focus on the accuracy and efficiency of this approach for the model in its original, unperturbed geometry relative to full order modeling (FOM) approaches as used in Sierra multiphysics module *Aria* [Aria 2022]. Our primary motivation is to generate results useful for uncertainty quantification (UQ)/verification and validation (VV) studies. As is often done for this type of investigation, we explore the variation of a simulation quantity of interest (QOI) metric, in this case the transient temperature at a specific model location, to a range of material parameter variations within expected distributional limits. The approach taken here makes use of SNL's Sierra multiphysics module *Aria* [Aria 2022], Python Pressio basis creation scripts [Blonigan 2019a/2019b/2020/2021a/2021c, Parish 2020, Rizzi 2020], the reduced order model hybrid thermal analysis code *Pressio_Aria* [Brunini 2022, Parish 2021], and Sandia National Laboratories' (SNL) UQ/VV code Dakota [Dakota 2020] for generating Latin hyper-cube (LHC) input parameter sampling, code execution, and production of metric statistics from a series of UQ/VV simulations. We describe the process of obtaining FOM calculations using *Aria*, generating ROM basis functions and sample meshes, and performing ROM calculations via *Pressio_Aria*. We then present a comparison of simulation metrics including accuracy and timings. Our findings indicate that while ROM approaches of this nature offer useful tools and significant speed increases for investigating multiphysics applications, such tools should be used with a full understanding of the inherent limitations of such approaches. This study demonstrates that ROM approaches of this type can be particularly useful for normal thermal environmental analysis and for this scenario have accurately reproduced FOM simulation results for the chosen QOI at significantly reduced computational cost in comparison to FOM calculations. For abnormal thermal environments, particularly those in which chemical reactions play a greater role at elevated temperatures, the accuracy of ROM predictions depends strongly on the fraction of the full order mesh utilized, the so-called hyper-reduction mesh fraction, as well as the specific metric of interest. Accuracy across the entire computational domain cannot be guaranteed at low mesh fractions, though in many cases analysts are primarily concerned with accuracy for a small set of QOI. We discuss this such a scenario here. Further study is needed to illuminate methods for identifying the minimum mesh fractions required to reproduce FOM results within a given tolerance in an a priori fashion. Additionally, continued efforts can help elucidate methods for choosing FOM mesh locations to be included within the hyper-reduced mesh to optimize ROM solution accuracy.

ACKNOWLEDGEMENTS

We would like to acknowledge all of the developers and applications experts that contributed to this work, including all members of the Crash-Burn milestone team [Tan-Torres 2020] and those of the *Pressio_Aria* and Pressio ROM project teams. We also want to express our thanks to the Sierra STK team for their contributions in improving the speed of hyper-reduced ROM basis production. We would additionally like to acknowledge the ASC program for its continued support of this project.

CONTENTS

LIST OF FIGURES

LIST OF TABLES

This page left blank

EXECUTIVE SUMMARY

This work summarizes the findings of a reduced order model (ROM) study performed using Sierra ROM module *Pressio_Aria* [Brunini 2022] on Sandia National Laboratories' (SNL) Crash-Burn L2 ASC milestone thermal model [Tan-Torres 2020] with pristine geometry. Comparisons are made to full order model (FOM) results for this same Crash-Burn model using Sierra multiphysics module *Aria* [Aria 2022]. Both ROM and FOM simulations have been performed in the context of a Dakota [Dakota 2020] Latin hyper-cube (LHC) sweep across a set of material parameters as was done in the original L2 milestone project. Thermal conditions explored include a normal thermal environment (NTE) at varying but near-ambient temperatures as well as an abnormal thermal environment (ATE) at an exterior temperature of 1010°C, reminiscent of a fire environment. This study demonstrates the utility of such ROM approaches, especially in the context of NTE, where excellent agreement between ROM and FOM results has been demonstrated at a significantly reduced computational cost. For ATE, agreement is more challenging, but is shown to be a function of the hyper-reduction mesh fraction of the original FOM mesh that is used. Potentially satisfactory agreement between FOM and ROM results can be obtained for a thermal quantity of interest (QOI) when a sufficiently large ROM basis set and hyper-reduction mesh fraction are utilized, again at significant reduction in computational time. Challenges in ATE versus NTE are discussed.

ACRONYMS AND DEFINITIONS

Abbreviation	Definition
API	Application Programming Interface
FOM	Full Order Model
ROM	Reduced Order Model
VV	Verification/Validation
UQ	Uncertainty Quantification
LHC	Latin Hyper-Cube
Al6061	Aluminum 6061
SS304	Stainless Steel 304
CG	Conjugate Gradient
NW	Nuclear Weapons
SNL	Sandia National Laboratories
FEM	Finite Element Method
NTE	Normal Thermal Environment
ATE	Abnormal Thermal Environment
GMRES	Generalized Minimum RESidual method
BDF2	Backward Differentiation Formula version 2 (time stepping algorithm)
SVD	Singular Value Decomposition
ODE	Ordinary Differential Equation
LSPG	Least Squares Petrov Galerkin
ϕ	Mesh fraction for ROM hyper-reduction
ASC	Advanced Scientific Computing
L2	Level 2

1. INTRODUCTION

A primary analysis component of all nuclear weapon (NW) qualification work performed at Sandia National Laboratories (SNL) is thermal analysis. Due to the high cost of evaluating system response to different experimental environments using real-world NW systems, subsystems, and components, there continues to be an increased focus on computer simulation for the evaluation of both normal (NTE) and abnormal (ATE) thermal environments as well as a variety of mechanical environments for these systems. This approach has been facilitated by the Moore's Law-type growth in computational power over the last several decades as well as accelerated investment in high performance computing (HPC) resources at the national laboratories, in industry, and in academic settings. Sophistication of thermal and other physics models has also improved the credibility of such approaches. The focus on verification/validation (VV) and uncertainty quantification (UQ) techniques has aided analysts in understanding the accuracy and credibility of modeling approaches as well as enabled them to predict solution uncertainties based on known variances in material properties and other system parameters, including solution timestep, boundary conditions, numerical solver type and tolerances, etc. [Tan-Torres 2020, Dakota 2020]

Despite the increasing computational power available to the thermal analyst, specific scenarios for a variety of weapons systems have inherently high computational cost, a result of increasingly complex physics and/or stringent spatial-temporal resolution requirements, with finite element method (FEM) mesh counts reaching into the many millions of elements/nodes in standard FEM studies. Notably, in certain cases, large physical time scales are needed while simultaneously requiring relatively small computational timesteps for the resolution of physics and chemistry models. This can pose significant challenges to the analyst. Additionally, since VV/UQ investigations often require an ensemble of independent simulations run over a large set of varied parameters, the actual cost of such a study in CPU time can become prohibitively high. As a result, analysts and developers continue to seek methods for improving computational efficiencies.

Here we propose a novel method for reducing the cost in simulation time for generating accurate simulation predictions using ROM methods. Specifically, in this report we describe the use of *Pressio_Aria*, a hybrid code built from Sierra multimechanics module *Aria* [Aria 2022], coupled with the reduced order modeling (ROM) capabilities found in the Pressio C++ library [Blonigan 2019a/2019b, Rizzi 2020]. *Aria* is the primary thermal module for SNL's SIERRA multiphysics application code suite. Pressio is a C++ ROM library that includes projection-based methods for solving large, potentially nonlinear, dynamical systems. Pressio presents a potential performance improvement over full order model (FOM) approaches and promises to aid the developer/analyst in solving large-scale physics problems, especially in the context of multiple parameter variants as seen in VV/UQ studies.

This report focuses on an analysis metric measured during a series of Full Order Model (FOM) and Reduced Order Model (ROM) simulations. FOM simulations utilize Sierra Aria for fully-described, traditional FEM thermal analysis simulations on a complete computational mesh. ROM simulations make use of the ROM code *Pressio_Aria* to generate solution predictions based on SVD bases determined from FOM simulations. Here ROM simulations are solved on a hyper-reduced mesh, which represents a fraction of the original FOM mesh. The metric or quantity of interest (QOI) is the transient temperature measured at a representative location in the simulation model of what we term here as the Crash-Burn scenario. For this study we use the pristine, pre-crash system geometry only. This scenario was the focus of a recent L2 milestone project at SNL [Tan-Torres 2020] which focused on a characteristic nuclear weapon (NW) device both in pristine and mechanically damaged forms experiencing thermal insult from a fire-type environment. For the current study, we focus on

the pristine geometry under two distinct thermal conditions. The first is that of an NTE for a typical warm summer day with lows ~19°C and highs ~ 30°C. The second is an ATE, focusing on the response of the device to the conditions of an external fire temperature of 1010°C, representative of a fully engulfing hydrocarbon fuel fire.

We investigate the accuracy of *Pressio_Aria* (ROM) model temperature predictions relative to *Aria* (FOM) calculations for both environments, commenting on the agreement and differences across the VV/UQ sampling set used for each as a function of the ROM hyper-reduction mesh fraction [Brunini 2022, Blonigan 2019a/2019b/2020/2021a/2021c, Parish 2020, Rizzi 2020]. Additionally, we discuss the increased computational efficiency observed when using *Pressio_Aria* over *Aria*, especially in the context of hyper-reduction techniques, a ROM method that allows the use of a sample(reduced) mesh that is a fraction of the size of the corresponding FOM FEM mesh. In such cases, the ROM solves for solution variables on this sample mesh at a potentially significantly reduced computational cost, with cost reduction being a function of the mesh fraction employed.

2. CRASH-BURN SCENARIO

For this study, we chose to apply ROM techniques to the Crash-Burn scenario, the subject of a recent ASC L2 milestone effort at SNL [Tan-Torres 2020]. As stated above, this scenario involves an object that is characteristic of an NW-type device, with primary geometry and components similar to those found in many NW systems that SNL regularly models, both thermally and mechanically. The computational thermal model describes a multiphysics set which includes heat conduction, thermal radiation transport via the enclosure radiation assumption, species reaction chemistry (ChemEq), and level sets. This physics set, or its close likeness, is commonly encountered in a variety of NW qualification studies performed at SNL. For simplicity, we chose to use the pristine, mechanically undeformed geometry of the Crash-Burn scenario. While specific crash geometries could also be analyzed with the current approach, a sweep over such variants is beyond the scope of the current investigation.

The original Crash-Burn geometry is show in cross-section in Figure 2-1 below. The quantity of interest that the current study focuses on, the mass center temperature, is indicated.

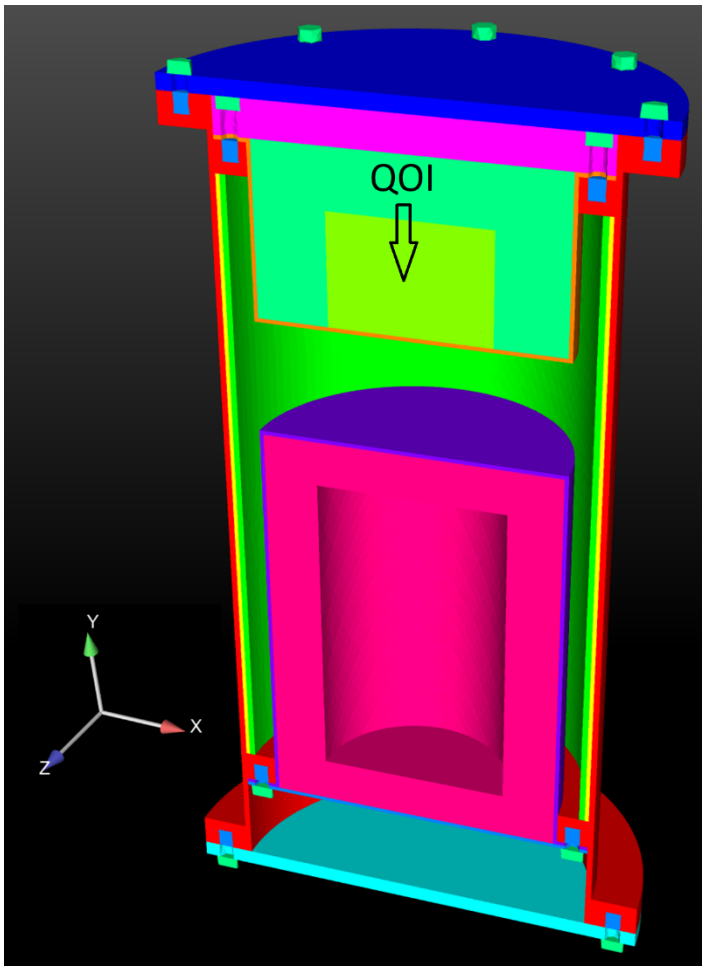

Figure 2-1. Cross-section of Crash-Burn geometry with individual material blocks identified by color. Quantity of Interest (mass center temperature) is indicated.

Figure 2-2 displays the same Crash-Burn model geometry disassembled into its various material blocks/components.

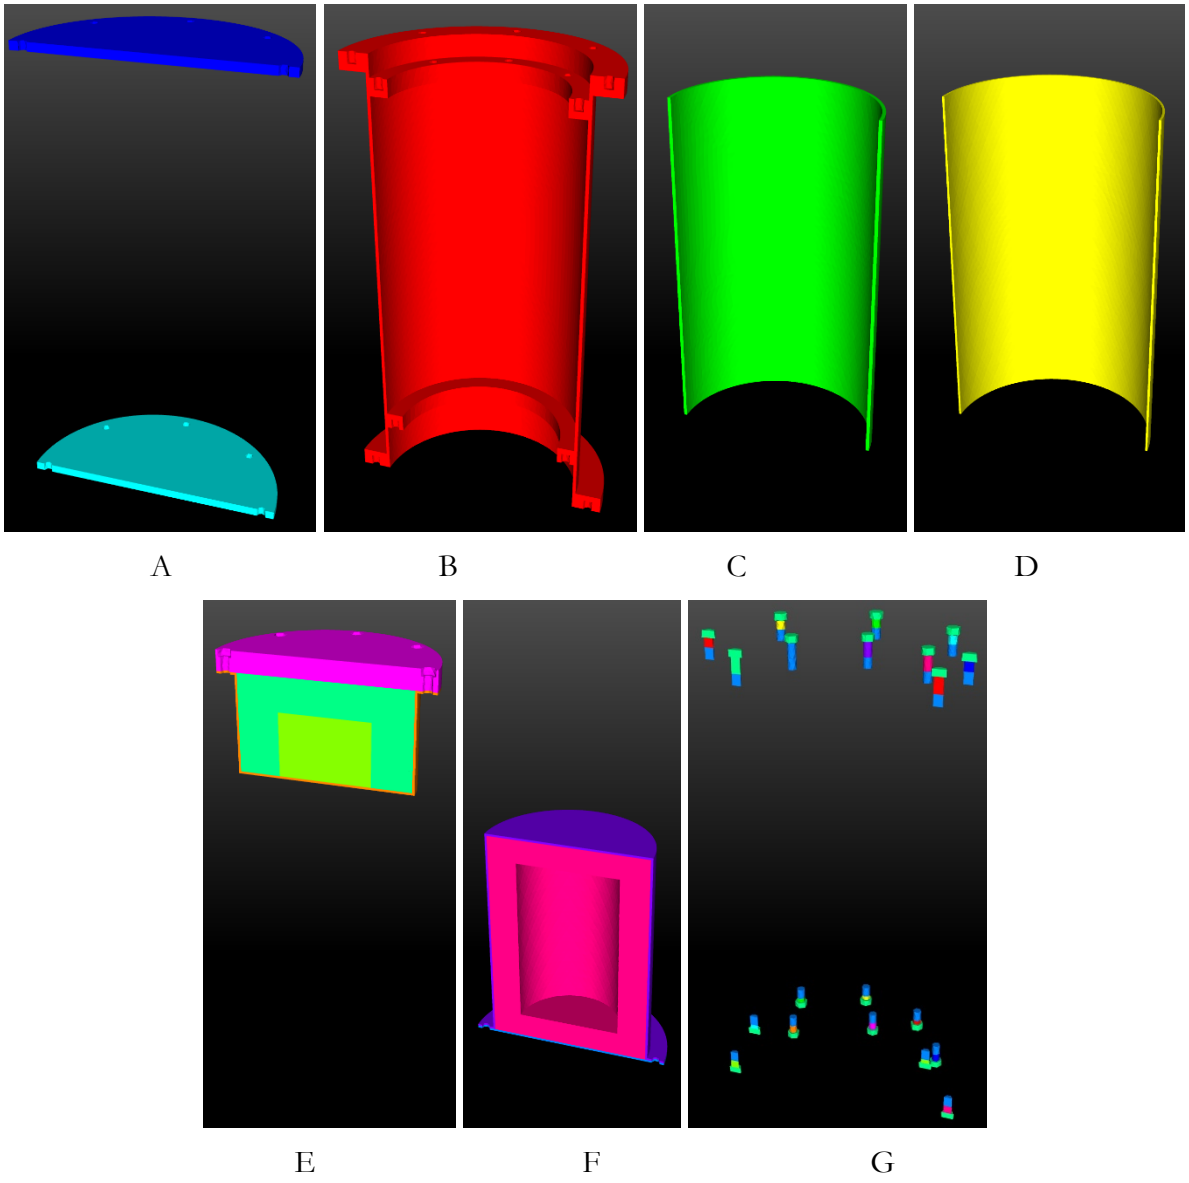

Figure 2-2. Crash-Burn geometry of components in cross-section (upper left to lower right): A) upper and lower case lids (blue/turquoise), B) outer case (red), C) structural composite layer (green), D) epoxy layer (yellow), E) mass mock (yellow-green) with its case(red), lid (purple), and potting(green), F) plastic (pink) with its case and lid (dark purple), and G) assembly screws (various)

2.1. Finite Element Model

Two different tetrahedral-based finite element model (FEM) meshes were used for the current investigation, one coarser and one finer. The coarser mesh is comprised of 248K nodes and 1.32M tetrahedral elements. This mesh was used for all NTE simulations. A more refined mesh was also created for additional investigation and is composed of 1.04M nodes and 5.79M tetrahedral elements. This second, more refined mesh was used for all ATE simulations due to the expectation

of higher thermal gradients and the need to spatially resolve the more active thermochemical mechanisms. Each finite element model contains 12 primary element blocks consistent with the larger components shown in Fig. 2-2. The screws and their heads and tails, seen in the last panel of Fig. 2-2, were placed in multiple additional element blocks for simplicity and ease of postprocessing. Element counts for each mesh block of the unrefined (NTE) and refined meshes are given in Appendix A, Table A-1.

Four unique sidesets were employed in this model. The first is composed of all exterior surfaces of the geometry and serves as the locus for application of the external thermal environment boundary condition. The remaining three sidesets define radiation enclosure cavities (voids) present on the internal surfaces of the model. Exterior and interior sidesets are shown in Figure 2-3 below.

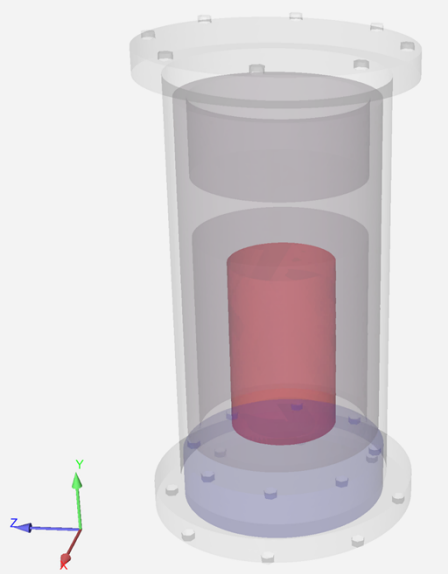

Figure 2-3. Exterior (white) and interior radiation enclosure (grey, blue, and red) sidesets for pristine Crash-Burn geometry.

The element and node counts for these surfaces (refined mesh) are shown in Table A-2 in Appendix A.

2.2. Materials and Chemistry Model (ChemEq)

2.2.1. Aluminum 6061

The outer case for this geometry, and both upper and lower case lids are constructed of Aluminum 6061 (Al6061) [Al6061-Conductivity/Emissivity, Phinney 2015].

The thermal material model for Al6061 is consistent with that typically used at SNL for many NW components and systems, including temperature-dependent specific heat, thermal conductivity, and emissivity. A constant density (2712.6 kg/m^3) and latent heat (391.25 kJ/kg) model are also employed. The latent heat model is used to properly account for the thermal energy needed to phase transition the Al6061 past its melting point for the ATE scenario.

2.2.2. Stainless Steel 304

The mass mock employed here is made of Stainless Steel 304 (SS304) [AWE 2018, Shurtz 2018, Phinney 2015, Tan-Torres 2020]. The mock is used in lieu of an actual nuclear material that would

normally be present in such NW system. The current design is meant to be representative of a standard system geometry but does not represent the actual design of any particular weapon in the current NW arsenal. The case enclosing this mass in addition to its lids are also composed of SS304. The case for the plastic as well as all screws used in this model are also SS304.

Similar to Al6061, SS304 utilizes a standard thermal model including temperature-dependent specific heat, thermal conductivity, and emissivity, with constant density (7916 kg/m^3). No latent heat model is needed for SS304 since temperatures never approach the melting point of steel, even in the abnormal thermal case. A constant heating rate of about 2800 W/m^3 representative of a constantly radiating mass that one might see in an actual NW system, is applied to the mock.

2.2.3. Epoxy

The epoxy layer is composed of epoxy material [Tan-Torres 2020]. The epoxy material properties are specified using a standard SNL thermal material data set. For epoxy, the specific heat varies with temperature. Thermal conductivity is a constant value of 0.228 W/m.K. Likewise, the emissivity is constant at 0.2.

2.2.4. Composite

The composite layer is composed of composite material [Tan-Torres 2020]. The composite has a constant density of ~1630 kg/m^3 and constant specific heat of 980 J/kg.K. It uses a directional (tensor) thermal conductivity with a nominal value of 0.4 W/m.K in a direction normal to the surface and 4.0 W/m.K in directions parallel to the surface. This is consistent with the layered structure of such composite materials. The radiative emissivity of the composite is a constant value of 0.5.

2.2.5. PMDI Foam

The potting for the mass as well as the plastic block is made of reacting foam (10_lb_foam_Reacting_PMDIpress) [Tan-Torres 2020, Scott 2014, Phinney 2015]. This foam material has an initial isotropic density of ~160 kg/m^3 and constant emissivity of 0.8. The specific heat follows a temperature-dependent form. Chemical reactions that govern species concentrations are managed by the ChemEq capability included in *Aria*. The foam is composed of several material constituents termed FOAMA, FOAMB, FOAMC, CHAR, CO2 (carbon dioxide), LMWO (low molecular weight oxides), HMWO (high molecular weight oxides), and N2 (nitrogen). The first four of these species are condensed phase, and the final four are gas phase products. In this model, FOAMA, FOAMB, and FOAMC, and CHAR have a constant solid density of 1500 kg/m^3. Foam components FOAMA, FOAMB, and FOAMC initially comprise 45%, 15%, and 40% of the foam by mass, respectively. Molecular weights for CO2, LMWO, HMWO, and N2 are 44, 80, 120, and 28, respectively. A three-reaction model, converting the foam components A, B, and C into gas and condensed phase products, is utilized:

$$[\text{FOAMA}] \rightarrow 0.56[\text{CO2}] + 0.44[\text{LMWO}]$$

$$[\text{FOAMB}] \rightarrow [\text{HMWO}]$$

$$[\text{FOAMC}] \rightarrow 0.5[\text{CHAR}] + 0.5[\text{HMWO}]$$

Activation energies for each reaction are held constant at a value of 179 MJ/kg. The foam thermal conductivity is the sum of both a radiative and bulk thermal component, the bulk component being defined by a polynomial form and the radiative component following a specific foam model which is adjusted for the solid mass fraction. Other specifics of the ChemEq chemistry model can be found in [Scott 2014].

2.3. Initial and Boundary Conditions

For the current study, both NTE and ATE conditions have been simulated. Both NTE and ATE simulations begin with the entire model at an isotropic temperature of 300K (26.85 °C).

2.3.1. Normal Thermal Environment

Under NTE conditions, environment (external) temperatures are set to a sinusoidal form that varies over a 24 hr time span as seen in Figure 2-4 below. Peak temperatures (daytime) are approximately 29.4 °C (84.9 °F) with nighttime temperature lows approximately 19.4 °C (66.9 °F), representing a modestly warm summer day.

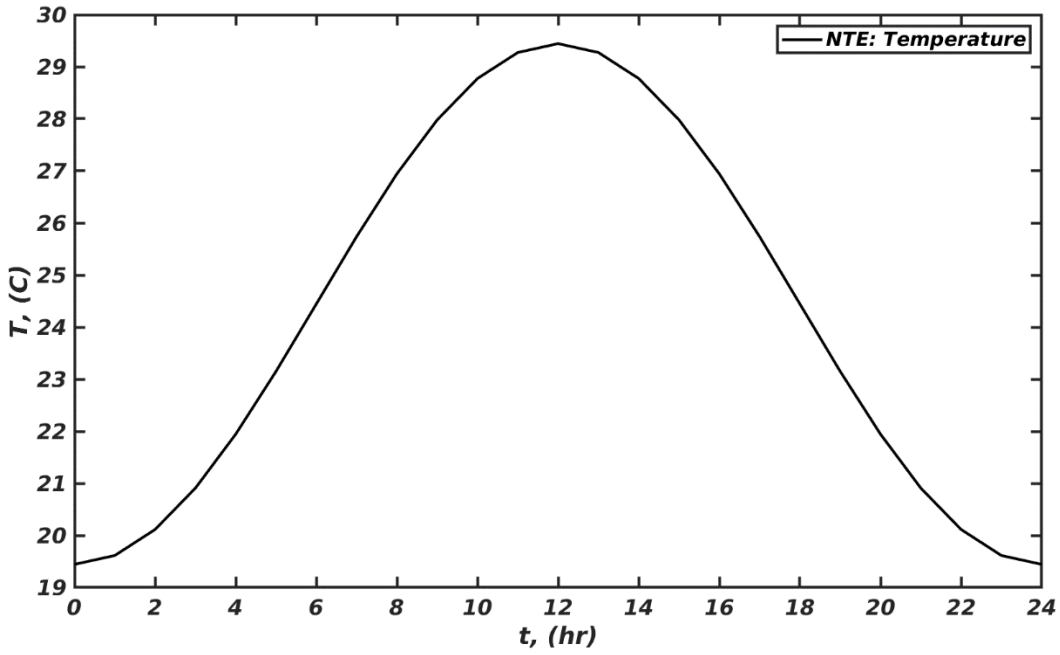

Figure 2-4. Environmental (external) temperature variation over the full normal thermal environment (NTE) model simulation duration (24hr)

2.3.2. Abnormal Thermal Environment

Under ATE conditions, as might be experienced by the object in a fire, the exterior of the device is exposed to an environmental temperature of 1010°C (1283.15 K) for the full hour of the abnormal thermal simulation. This temperature is commonly referred to as the "fully engulfing fire" boundary condition for NW qualification at SNL.

2.3.3. Boundary Condition Application

Both NTE and ATE exterior boundary conditions are imposed as a *radiative heat flux boundary condition*, indicating that the net heat flux on the exterior of the device normal to the surface is modeled as shown in Equation 1 below:

$$Q_{normal} = F\varepsilon\sigma_{SB}(T^4 - T_{BC}{}^4) \quad \text{Equation 2-1}$$

Where Q_{normal} is the incident heat flux and F is the radiation form factor, set to 1 for both NTE and ATE simulations. A form factor of 1.0 indicates that the full flux from the environment temperature is incident on the exterior surface, a limiting case. ε is the emissivity of the exterior surface of the device, which is composed of Al6061 (properties described above), and σ_{SB} is the Stefan-Boltzmann constant of approximately $5.67\text{x}10^{-8}$ W/m^{2}K^4. T is the instantaneous, spatially resolved temperature of the exterior surface of the device, and T_{BC} is the environment temperature, as specified above for the NTE and ATE scenarios.

While a convective term could be applied to the exterior, we have not included this effect in the current set of simulations. In typical NW qualification studies, a convective term would be included with either a constant convection coefficient or a more sophisticated convection correlation model. We have chosen to ignore convective heating/cooling for simplicity.

2.4. Enclosure Radiation

Both NTE and ATE simulations employ an enclosure radiation transport model, where empty cavities in the device are idealized as sets of diffuse-grey surface facets that isotropically emit/absorb/reflect thermal radiation based on their temperatures, emissivities, sizes, and relative positions and orientations [Brunini 2022]. To this end, pairwise geometric viewfactors are calculated between all facets within an enclosure, leading to the construction of a viewfactor matrix. In Sierra *Aria* (and by extension *Pressio_Aria*), enclosure radiation viewfactor and radiosity calculations are performed within the sub-module Chaparral. For this scenario, the hemicube approximation for viewfactor calculation is used with a geometric tolerance of 10^{-8}[Cohen 1985]. The radiosity solver is of a GMRES-type using a mason linearization of temperatures across each surface facet with a convergence tolerance of 10^{-6}. While viewfactor matrix smoothing is available in *Aria*, none is utilized for this scenario. The simulation geometry includes three separate radiation enclosures. For each surface, full viewfactor matrices are determined and blocking surface calculations are performed to determine which facets impede direct line of sight between pairs of facets so that row sums of the matrix converge to 1 as required by energy conservation. Radiosity terms from these enclosures contribute to the volumetric energy equation (thermal conduction) via surface heat flux contributions on enclosure surfaces. In typical thermal analyses, enclosure cavities are filled with a gas phase of known density and thermal capacity. In these cases, we define a convective flux between the gas in the cavity and the walls of the enclosure. This capability relies on a bulk element approach which is currently being implemented for *Pressio_Aria*. Since this is in the development stage, we do not include this convective term in the current study either for FOM or ROM simulations. Therefore, the gases contained in the three enclosure cavities in this model are assumed to be transparent to thermal radiation and non-interacting with the enclosure walls.

2.5. Thermal Conduction and Energy Sources

Standard thermal conduction is modeled in this simulation for all materials given their specific heat, density, and thermal conductivity as specified in the materials section above, including the effects of anisotropic thermal conductivity, as discussed for the composite material. Latent heat terms, to

represent the melting transition in Al6061, are also employed. Additionally, species dependent radiative and bulk thermal conductivities are used in the foam material which also affect thermal transport in the device. The chemistry model employed in the Crash-Burn scenario can add additional energy (heat) source terms in the foam material, though for this case the foam is non-energetic. The mass mock includes the constant energy (heating) source of 2799.16 W/m^3 as described in the materials section above.

2.6. Solver Parameters and Timestepping

The energy equation uses a conjugate gradient (CG) solver with a Jacobi preconditioner. The convergence tolerance for the solver is 10^{-8}. *Aria* is set to use an adaptive second order backward differentiation formula (BDF2) time-stepping scheme with an initial solution timestep set to 0.05 sec. BDF2 follows the form:

$$\frac{dy}{dt} = f(t,y) \qquad y(t_0) = y_0 \qquad \text{Equation 2-2}$$

$$y(t_{n+2}) - \frac{4}{3}y(t_{n+1}) + \frac{1}{3}y(t_n) = \frac{2}{3}\Delta t f[t_{n+2}, y(t_{n+2})] \qquad \text{Equation 2-3}$$

At each timestep, a maximum timestep size ratio of 1.25 is enforced so that the current timestep maintains a reasonable and stable rate of increase or decrease from step to step. The maximum timestep is set to 60 sec, with a small minimum timestep of 10^{-13} sec. A failed timestep size ratio of 0.5 is used to reduce the current timestep size by a factor of 0.5 when convergence isn't achieved with the current timestep size. A predictor-corrector tolerance of $5\text{x}10^{-5}$ is also employed. Additionally, the maximum temperature change is restricted to be less than 100K (100°C) in a single timestep. When exceeded, the timestep is reduced to avoid large temperature swings which may indicate solution instability.

The ChemEq solver that is used to solve the ordinary differential equation (ODE) chemical reaction terms is a CVODE (C-Language Variable ODE) type that employs an 4th order Adams functional iteration method with a relative tolerance of 10^{-3} and absolute tolerance of 10^{-6}. [Cohen 1996]

A newton nonlinear solution strategy is used for the energy equation with a relaxation factor of 1 and an iteration tolerance of 10^{-7}. A maximum of 10 nonlinear iterations is allowed at each timestep.

Each simulation performed in the NTE/ATE LHC set is propagated for a period of either 24 hrs (NTE) or 1 hr (ATE) of simulation time. Temperatures are recorded every 60 sec of simulation time for comparison between the ROM and FOM approaches, for a total of 1441 (NTE) or 61(ATE) recorded outputs.

3. PRESSIO ROM APPROACH AND METHODOLOGY

Pressio is a collection of C++ ROM libraries being continually developed in an ongoing open-source project that enables projection-based approaches for ROMs applicable to very large linear and nonlinear dynamical systems of a variety of types [Rizzi 2020]. Pressio has capabilities that enable a user to reduce the spatial degrees of freedom, a method termed hyper-reduction. Temporal degrees of freedom, that is the number of required computational timesteps for a transient run of fixed duration, can also be reduced in comparison to FOM solutions in some cases as ROMs have been observed to provide enhanced stability [Rizzi 2020, Blonigan 2020, Blonigan 2021c]. The utility of Pressio's design lies in its straightforward Application Programming Interface (API) which provides the potential to enable the use of ROM features in nearly every application, with the constraint that the dynamics of the system are expressible as a system of ordinary differential equations (ODEs). Other coding requirements, including API, are discussed by Rizzi [Rizzi 2020].

Projection based ROMs have notable advantages over other surrogate techniques for solution finding including data-fit and reduced-fidelity models. While data-fit methods circumvent the specific physics of the problem involved, reduced-fidelity approaches either simplify the physics fidelity of the FOM or otherwise reduce the model details via techniques like mesh coarsening. Projection based ROM methods like Pressio solve the same computational model as the corresponding FOM, and they can do this for the same set of physics and even the same spatial model (mesh) albeit using a projection-based scheme. In this context, the high order approximation of each solution state is given as:

$$\boldsymbol{x}_{approx}(t,\boldsymbol{\mu}) = \boldsymbol{x}_{ref}(\boldsymbol{\mu}) + g[\boldsymbol{x}_g(t,\boldsymbol{\mu})] \quad \text{Equation 3-1}$$

where $\boldsymbol{x}_{approx}$ is the approximate state of the system at time t with system parameters $\boldsymbol{\mu}$. $\boldsymbol{x}_{ref}$ is a reference state of the system, normally the state of the system at a specific simulation time, often the initial state at time $t = 0$. $\boldsymbol{g}$ is a mapping function (decoder) that maps a set of generalized coordinates $\boldsymbol{x}_g$ to the approximate current transient state of the system [Rizzi 2020]. Note that in this formulation, the function $\boldsymbol{g}$ maps from a smaller dimensional space to a higher dimensional one, the very purpose of a ROM. When the mapping function is linear, we can express the approximate solution as:

$$\boldsymbol{x}_{approx}(t,\boldsymbol{\mu}) = \boldsymbol{x}_0(\boldsymbol{\mu}) + \Phi\boldsymbol{x}_g(t,\boldsymbol{\mu}) \quad \text{Equation 3-2}$$

Where Φ is the constant Jacobian of the decoder function (commonly referred to as the basis), and $\boldsymbol{x}_{ref}(\boldsymbol{\mu}) = \boldsymbol{x}_0(\boldsymbol{\mu})$ is the initial state of the system.

In the offline stage of the Pressio ROM process, $\boldsymbol{g}$ (or Φ) is determined via a singular value decomposition (SVD) approach. That is, a lower dimensional "trial subspace" is determined for approximating the solution [Tencer 2021]. FOM solutions are used to build the SVD basis, where solution data is provided via a set of transient solution exodus files. The size or number of basis vectors can be chosen independently by the analyst.

Galerkin and least-squares Petrov-Galerkin (LSPG) formulations are both available via Pressio [Rizzi 2020]. We leave a discussion of both approaches to other sources, except to indicate that we have chosen to employ a LSPG methodology for the current study as that is the approach currently available within *Pressio_Aria*.

Over the last three years, Pressio's ROM capabilities have been incorporated into the Sierra multiphysics module *Aria* via the executable module *Pressio_Aria*. *Pressio_Aria* utilizes the same

input deck style as Sierra *Aria* and makes use of *Aria*'s infrastructure for calculating residuals and Jacobians, necessary inputs for Pressio solvers. This similarity simplifies the (FOM→ROM basis→ ROM simulation) paradigm. In this paradigm, *Aria* FOM simulations (one or several) are run with the full system geometry and description. FOM results from exodus output files, are used as the input to scripts that manage creation of SVD bases in the form of auxiliary exodus files using tools contained in the *pressio_basis* repository. These auxiliary files are then used to perform the ROM simulation.

4. CRASH-BURN FOM RESULTS TO SVD BASIS

In the current study, we have followed a multi-step approach for generating *Pressio_Aria* ROM results. The first step of this process includes running one or more FOM (*Aria*) Crash-Burn scenario simulations, varying parameters as one would in a standard LHC UQ study. In our case, we have followed the same parameter variational approach used during the original Crash-Burn L2 milestone performed at SNL [Tan-Torres 2020]. Here, 16 material parameters were varied over expected distribution limits. These included the thermal conductivity, specific heat, and emissivity of the system materials described above. 96 NTE and 256 ATE independent simulations were executed using a Latin Hyper-Cube (LHC) approach for parameter sampling using SNL's Dakota package. [Dakota 2020]. Additionally, single NTE and ATE simulation sets with nominal parameters were also executed.

In a UQ study of this type, each parameter is varied within a range specified by the analyst. Here, a gaussian distribution was assumed for each of the 16 varied parameters, as shown in Table 4-1 below. Additionally, a lower bound and upper bound for each parameter was enforced to keep material properties within physical limits. All 16 parameters were defined as multiplicative factors to the nominal properties of the various materials including thermal conductivity (k), specific heat (Cp), and emissivity (emis). In the case of the PMDI foam, two conductivity terms are present, a radiative conductivity (k_rad) and bulk conductivity (k_bulk). Using this list of parameters and distribution information, which the analyst defines in an input file, Dakota generates and executes a set of independent simulations based on LHC sampling. For the NTE case, 96 independent simulation sets were generated and executed. For the ATE case, 256 independent simulations were employed. In both cases, FOM (*Aria*) and ROM (*Pressio_Aria*) simulation sets with the same parameter definitions were run.

Table 4-1. Dakota Crash-Burn Parameter Variation

Parameter	Mean	Standard Deviation	Lower Bound	Upper Bound
f_al_k	1.0	0.05	0.9	1.1
f_al_Cp	1.0	0.05	0.9	1.1
f_al_emiss	1.0	0.1	0.8	1.17647
f_ss304_k	1.0	0.025	0.95	1.05
f_ss304_Cp	1.0	0.1	0.8	1.2
f_ss304_emis	1.0	0.1	0.8	1.2
f_epoxy_k	1.0	0.15	0.7	1.3
f_epoxy_Cp	1.0	0.15	0.7	1.3
f_epoxy_emis	1.0	0.15	0.8	1.3
f_PMDI_k_bulk	1.0	0.175	0.65	1.35
f_PMDI_k_rad	1.0	0.175	0.65	1.35
f_PMDI_Cp	1.0	0.1	0.8	1.2
f_PMDI_emis	1.0	0.1	0.8	1.2
f_composite_k	1.0	0.15	0.7	1.3
f_composite_Cp	1.0	0.15	0.7	1.3

Parameter	Mean	Standard Deviation	Lower Bound	Upper Bound
f_composite_emis	1.0	0.1	0.8	1.2

Several approaches were explored to generate SVD bases for ROM simulations. This includes using some of the LHC FOM results as the input to our SVD basis creation. Clearly, it was not feasible to create an SVD basis set for every one of the 96/256 FOM LHC parameter sets in the UQ studies. We have explored the effects of using only simulation 1 from the FOM LHC set as well as LHC simulations numbered 1 to 2, 1 to 4, 1 to 8, and 1 to 16. We additionally attempted using *only* the FOM simulation with the nominal set of parameters. Our best and most representative ROM results *overall* came from the use of only the nominal parameter FOM simulations. This is somewhat unexpected as we initially considered it likely that adding more simulation data from a varied set of parameters would allow for a more robust basis set. Without exception, when more than one simulation was used as the FOM input to the SVD creation, the ROM results were less accurate than using a single set and at significantly increased cost due to the inclusion of additional data into the FOM results file from the multiple simulations. As the nominal set of parameters is the "central" parameter set, it is not surprising that this single set provided the most ideal *single* set of FOM results to be used as input to the SVD basis creation scripts.

The time to create a basis also varies with the number of basis vectors, with this number being user selected. Additionally, Pressio enables the user to run a ROM simulation that corresponds to the FOM version while on a spatially reduced (sample) mesh, so called hyper-reduction. Hyper-reduction reduces the overall ROM simulation time significantly, and we have employed that technique here. For the hyper-reduction as utilized in this study, a fraction of the nodes of the original mesh are chosen by a method described below. All mesh elements that contain those nodes as well as their additional accompanying nodes are included in the hyper-reduced sample mesh. As previously indicated, for NTE simulations, a sample(reduced) mesh at a mesh fraction $\varphi = 0.03$ was used for *Pressio_Aria* ROM simulations, while for ATE we used both $\varphi = 0.03$ and $\varphi = 0.003$ for comparison of results. We found that for both NTE and ATE scenarios, a basis set consisting of 64 basis vectors resulted in well-resolved results while still being feasible to generate within a time period of a few hours. We have used a psample method for basis creation. Psampling makes use of QR decomposition to decompose the matrix of solution snapshots S, that is the matrix of solutions at all nodal locations (rows) and output times (columns), into a product of an orthogonal matrix Q and an upper diagonal matrix R.

$$SP = QR \qquad \text{Equation 4-1}$$

Here P is a matrix that resorts S such that R has a diagonal that is non-increasing. The effect of this QR decomposition is that specific nodes from the mesh can be chosen based on their importance in the QR decomposition. An alternative "random" method randomly selects nodes from the mesh without regard to importance in such a decomposition.

5. FOM/ROM RESULTS AND COMPARISON

Evaluation of results for this FOM/ROM comparison consists of two basic themes. The first is that of computational efficiency, that is the speedup offered by ROM calculations over FOM. The second focuses on solution accuracy.

5.1. Comparison of FOM and ROM Simulation Timings

Total processor timings (CPU hours) for NTE and ATE thermal environments are displayed graphically in Figures 5-1 and 5-2 and in tabular form in Tables 5-1 and 5-2. Timing values are the sum of all processor times for these multi-core simulations. Figures 5-1 and 5-2, shown in log-log for clarity, display the distributions of these timings. Tables indicate averages, standard deviations, minima, and maxima for NTE and ATE, for both FOM and ROM timings. For NTE, a total of 96 Dakota LHC simulations, both FOM and ROM were performed, and for the ROM simulations, a mesh fraction of $\varphi = 0.03$ was utilized. For ATE, a total of 256 simulations Dakota simulations were performed as part of the LHC set for both FOM and ROM, with ROM using mesh fractions of $\varphi = 0.003$ and $\varphi = 0.03$.

All FOM and ROM simulations for both NTE and ATE were run on SNL's Eclipse HPC architecture. Eclipse has 1488 computational nodes with 36 processor cores per computational node for a total of 53,568 cores. The processor type used on all Eclipse nodes is the 2.1 GHz Intel Broadwell E5-2695 v4 2Sx18C with 128 GB of RAM per node. All FOM simulations were run using 128 processor cores. ROM simulations used only 4 processor cores.

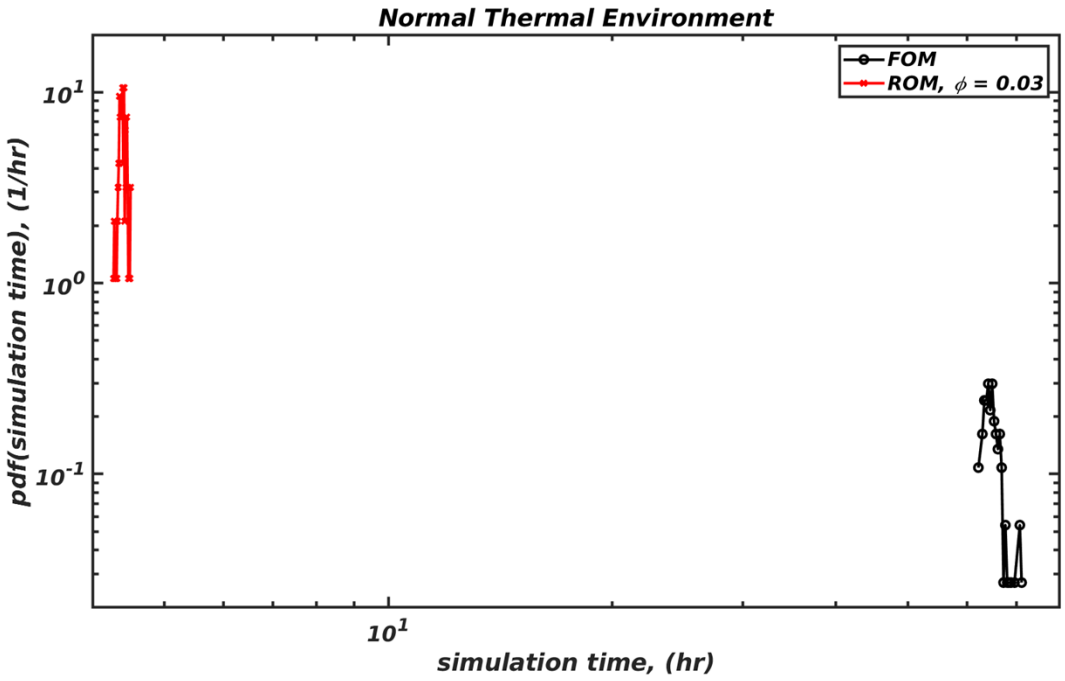

Figure 5-1. PDFs of NTE simulation times for FOM(black) and ROM(red)

Table 5-1. Normal thermal environment (NTE) CPU timings including average, standard deviation (σ), minimum and maximum of total CPU times (all cores) in hrs over 96 Dakota LHC parameter variations for FOM and ROM (mesh fraction = 0.03).

Measure	FOM (hr)	ROM (φ = 0.03) (hr)
Average	64.8	4.39 (6.81% FOM)
σ	1.84	0.050 (2.72% FOM)
Minimum	61.8	4.26 (6.89% FOM)
Maximum	71.0	4.450 (6.33% FOM)

For the NTE set, ROM average, minimum, and maximum times for NTE are at 6-7% of FOM times, indicating a significant speedup in simulation times. Standard deviations for ROM timings are also relatively smaller than FOM timings, at less than 3% of the FOM value.

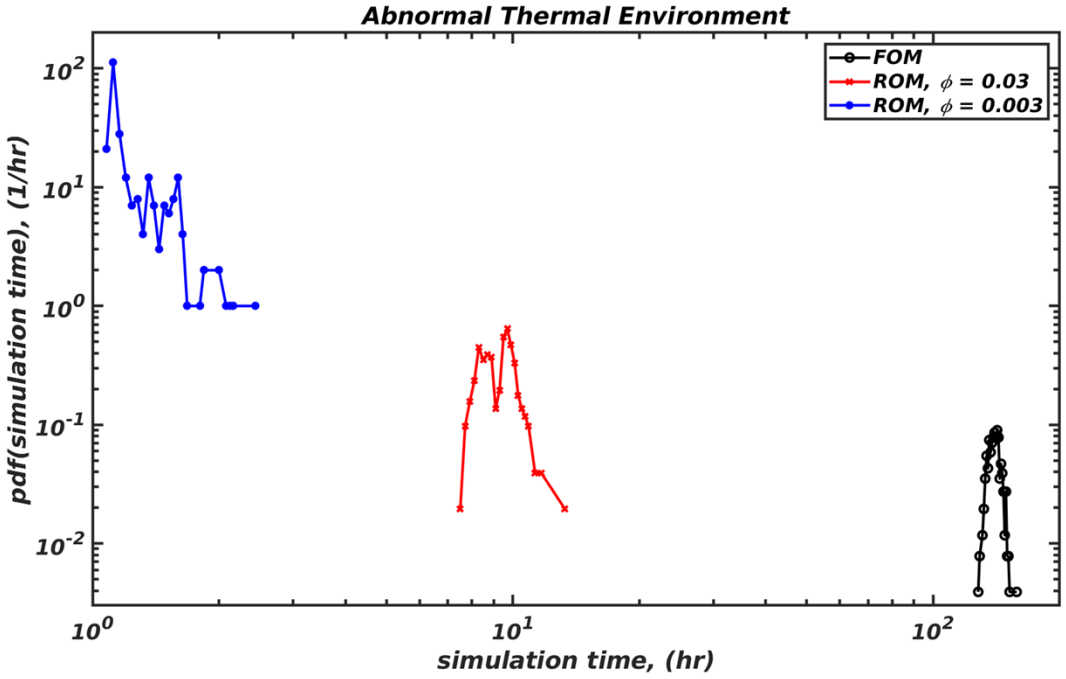

Figure 5-2. PDFs of ATE simulation times for FOM (black) and ROM [φ = 0.03(red), 0.003(blue)]

Table 5-2. Abnormal thermal environment (ATE) CPU timings including average, standard deviation (σ), minimum and maximum of total CPU times (all cores) in hrs over 256 Dakota LHC parameter variations for FOM and ROM (mesh fractions = 0.003, 0.03).

Measure	FOM (hr)	ROM (φ = 0.03) (hr)	ROM (φ = 0.003) (hr)
Average	140	9.20 (6.59% FOM)	1.23 (0.884% FOM, 13.4% φ = 0.03)

Measure	FOM (hr)	ROM (φ = 0.03) (hr)	ROM (φ = 0.003) (hr)
σ	4.80	0.881 (18.4% FOM)	0.221 (4.61% FOM, 25.1% φ = 0.03)
Minimum	128	7.41 (5.79% FOM)	1.05 (0.824% FOM, 14.2% φ = 0.03)
Maximum	157	13.1 (8.36% FOM)	2.412 (1.54% FOM, 18.4% φ = 0.03)

For the ATE set, ROM average, minimum, and maximum times are in the range of approximately 6-8% of FOM times for φ = 0.03 and 1-2% for the lower mesh fraction of φ = 0.003. The ratio φ = 0.03 ROM to FOM timings is consistent with what is observed for NTE simulations as discussed above. The timings for the φ = 0.003 ROM simulations are in the range of 13-18% of those for φ = 0.03 despite the 10x reduction in mesh size. Thus, we do see significant but not linear speedup in simulation times with ROM mesh fraction. Standard deviations for φ = 0.03 ROM timings relative to average timing values are larger here than for FOM simulations, at greater than 18% of the FOM value. When we move to φ = 0.003, the standard deviation is also relatively higher at nearly 5% of the FOM value and greater than 25% of the φ = 0.03 value, despite the 10x reduction in mesh fraction.

In summary, *Pressio_Aria* ROM simulations provide a path for significant speedup in performance in comparison to *Aria* FOM calculations, with speedup being a strong function of mesh fraction. Speedups are slower than linear, but the aggregated time savings over large LHC sets can be significant. Additionally, the use of reduced mesh fractions can enable analysts to run such investigations on desktop resources. As demonstrated here, 4 processor cores were sufficient to run ROM LHC sets, while comparatively large resources (128 cores) were used for FOM calculations.

5.1.1. FOM ROM Timing Comparison varying CPU core count

Below we present timing data for a slight variation of the ATE ROM configuration defined above. In this case we employed SNL's HPC Platform Attaway, a computational cluster with 1488 nodes housing 53,568 2.3 GHz Intel Skylake Gold processors (36 cores/node). Each computational node has 192 GB of available RAM. This timing investigation focuses on FOM and ROM results for the Crash-Burn scenario with the nominal parameter set only, varying the number of CPU cores per simulation. A higher ROM hyper-reduction mesh fraction of 0.1 and 128 basis vectors were used. Figure 5-3 below displays the total CPU time spent for FOM(Aria) and ROM(*Pressio_Aria*) simulations run to completion. It further decomposes these timings into the time spent performing distinct physics calculations, specifically those associated with radiosity (thermal radiation enclosure) and chemistry (PMDI Foam decomposition). A final set provides the sum of the times for all non-radiosity and non-chemistry calculations, labeled here as "Other".

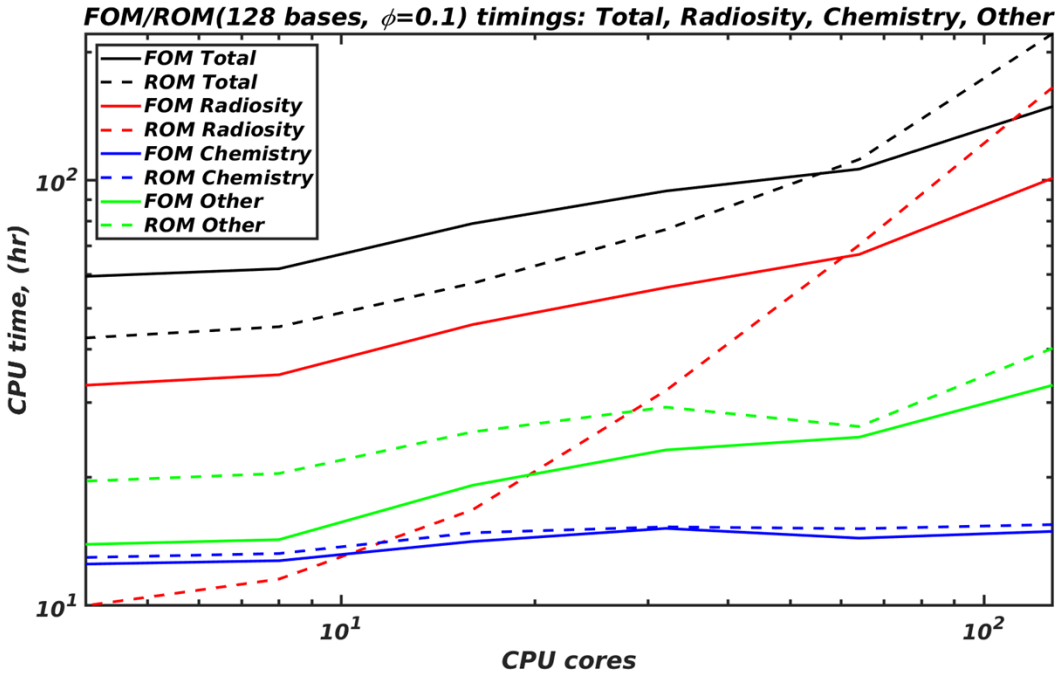

Figure 5-3. Total(black), Radiosity(red), Chemistry(blue), and Other(green) ATE CPU timings for FOM (solid line) and ROM (dashed line, φ = 0.1, 128 bases) calculations vs. CPU core count

Ideal strong scaling for FOM and ROM simulations would be evidenced by horizontal lines in the plots of total CPU time as a function of core count, that is, independence of core count. While we observe relatively flat curves for both FOM and ROM at low core counts (4 to 8 cores), upward trends are observed at higher core counts, indicating greater than power-law increases on this log-log scale. ROM CPU times are less than FOM for less than 64 processor cores. At 64 cores, the FOM is operating at ~16K mesh nodes per core, the ROM at fewer mesh nodes per core (~12.6K). Aria is known to demonstrate less than ideal scaling at mesh node-per-core counts $O(10^4)$ or less. *Pressio_Aria* relies on the same underlying computational framework and is expected to exhibit similar behavior. The deviation from ideal strong scaling is due, in large part, to the growing additional cost of communication of data between neighboring processor cores at processor boundaries as the core count increases. Alternatively, one could say that as the number of processor cores increases for a mesh of fixed size, the number of mesh nodes that hold data that must be communicated to other processors becomes an increasing fraction of the total number of mesh nodes, thus increasing the relative cost of communication versus computation. Notably, the strong scaling behavior of the ROM starts to trend upward at a lower core count than the FOM, most likely due to the relatively smaller mesh size leading to fewer mesh nodes per core than for the FOM. Similar scaling for both ROM and FOM is observed in the radiosity calculations, with ROM speedups of more than 3x at the lowest core count (4) and crossover between FOM and ROM timings occurring at 64 cores. Interestingly, chemistry timings remain nearly flat, approaching ideal strong scaling limits with both FOM and ROM from 4 to 128 cores, timings increasing only slightly over this range. Also interesting is the fact that FOM timings are faster by a modest amount than ROM timings throughout the entire range of processor core count. Timings for all "other" calculations also exhibit FOM times lower than ROM times, again over the entire range of core

counts. These other calculations do not demonstrate the same quality of strong scaling as the chemistry calculations, but retain better scaling than radiosity calculations. At smaller mesh fractions, including at 0.003 and 0.03, we have shown above that CPU timings are significantly reduced when utilizing ROM calculations over FOM. A mesh fraction of 0.1 as used here, while still able to provide superior timings at low core counts, is much less competitive with FOM timings than smaller mesh fractions of 0.003 and 0.03 are. A subtle point should be made regarding these results. A user-specified mesh fraction of 0.1 does not mean that 10% of the nodes from the FOM mesh will be present in the ROM mesh. Rather, 10% of the FOM mesh nodes will be chosen via p-sampling or random selection and included *along with* all mesh nodes in elements that contain this selection, as described for hyper-reduction above. In this case, a 0.1 mesh fraction resulted in 77% of the nodes and 34% of the elements from the FOM mesh being included in the ROM hyper-reduced sample mesh. Thus, instead of a 10x speedup, we may reasonably expect a 1.5x-3x increase in efficiency based on node/element count alone. We have not compared FOM and ROM timings without hyper-reduction in this study and so cannot authoritatively distinguish here how much of this speedup is a result of ROM methods in lieu of FOM methods separately from hyper-reduction.

We can also gain insight about CPU timings via the minimum and maximum times spent by CPU cores for radiosity and chemistry calculations as a function of processor count. Figures 5-4 and 5-5 below show FOM and ROM minimum and maximum CPU core timings for radiosity and chemistry calculations, respectively, as a function of core count. Dashed lines in these plots indicate perfect strong scaling predictions assuming the lowest core count times (4 cores) scales exactly to higher core counts up to 128 cores.

For radiosity calculations, ROM min and max core timings are both less than their FOM counterparts for all core counts less than 64. At 64 cores, FOM calculations are faster for both minimum and maximum timings. Interestingly, min and max FOM radiosity timings, while not perfect in strong scaling, are at least monotonically decreasing from 4 to 128 cores. This is not the case for ROM calculations, where minimum timing actually become larger for core counts greater than 8, and maximum radiosity timings increase above 64 cores.

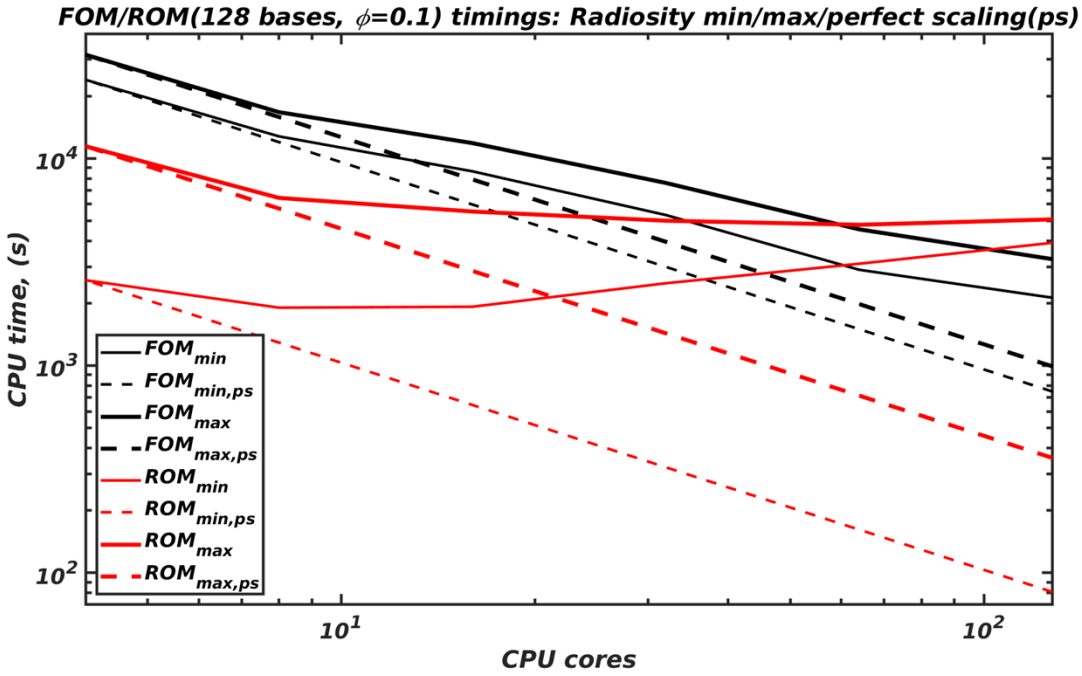

Figure 5-4. Minimum(thin), and Maximum(thick) CPU timings for FOM(black) and ROM (red, φ = 0.1, 128 bases) radiosity calculations vs. CPU core count. Perfect strong scaling predictions (based on 4 core results) for each shown as dashed lines.

Chemistry calculations, shown in Figure 5-5, also exhibit interesting timing behavior. Here, minimum and maximum timings are barely distinguishable between FOM and ROM. For minimum timings, better than perfect strong scaling is observed for both FOM and ROM, with timings falling off at greater than power law rates from 4 to 128 cores. Maximum timings for FOM and ROM are both nearly ideal with timings tracking perfect strong scaling predictions closely from 4 to 128 cores with a slight increase as the core count approaches 128 cores. Chemistry timings present challenges to a clear understanding about the effect of using ROM versus FOM methods. We propose that total chemistry timings could be largely unaffected by the use of ROM techniques since a reduction of spatial resolution in hyper-reduction ROM scenarios such as this could make chemistry solver convergence more difficult so that total solution times remain largely unaffected. Further study is certainly warranted, and investigation with a variety of chemical mechanisms, both fast and slow reactions, with and without exo-/endothermic energy terms, could help illuminate effects on ROM versus FOM timings in a clearer way than this single scenario.

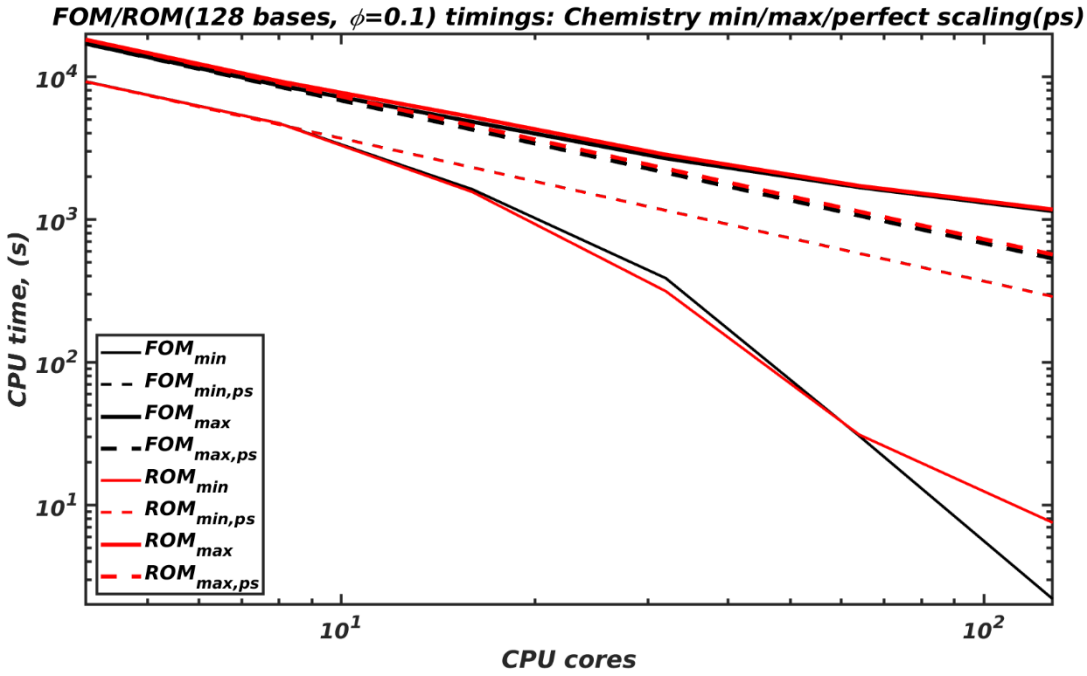

Figure 5-5. Minimum (thin), and Maximum (thick) CPU timings for FOM (black) and ROM (red, φ = 0.1, 128 bases) chemistry calculations vs. CPU core count. Perfect strong scaling predictions (based on 4 core results) for each shown as dashed lines.

5.2. Comparison of FOM and ROM Transient Thermal Responses

Here we discuss the accuracy of *Pressio_Aria* ROM calculations when compared to Aria FOM results for both NTE and ATE LHC simulation sets.

5.2.1. Best and Worst Case QOI across Dakota LHC

Figure 5-6 displays transient data for NTE FOM and ROM simulations at φ = 0.03 for the QOI that was selected for this LHC study, that is the temperature at the center of mass mock in the Crash-Burn geometry and mesh described and indicated above. Displayed in the top pane of Figure 5-6 are both FOM and ROM results for two of the simulations from the LHC set. Runs #84 and #40 are shown, with run #84 (black) displaying the best agreement between FOM and ROM results and run #40 (red) showing the worst agreement between data sets. Here "agreement" indicates consistency between FOM and ROM results for the QOI only, not for temperatures over the entire simulation domain. Consistency is measured by the mean-squared error between QOI FOM and ROM values over the entire simulation time range of 24 hrs with data taken for the QOI every 60 sec. One can easily observe the excellent agreement for run #84 where FOM and ROM results are nearly indistinguishable over the entire simulation time of 24 hrs (86,400 sec). Even in the worst case of run #84, agreement is quite good, with temperature variations of less than 0.25 K over nearly the entire duration of the simulation. FOM/ROM differences are displayed for both sets in the bottom pane of Figure 5-6, further highlighting the difference between best and worst case agreement.

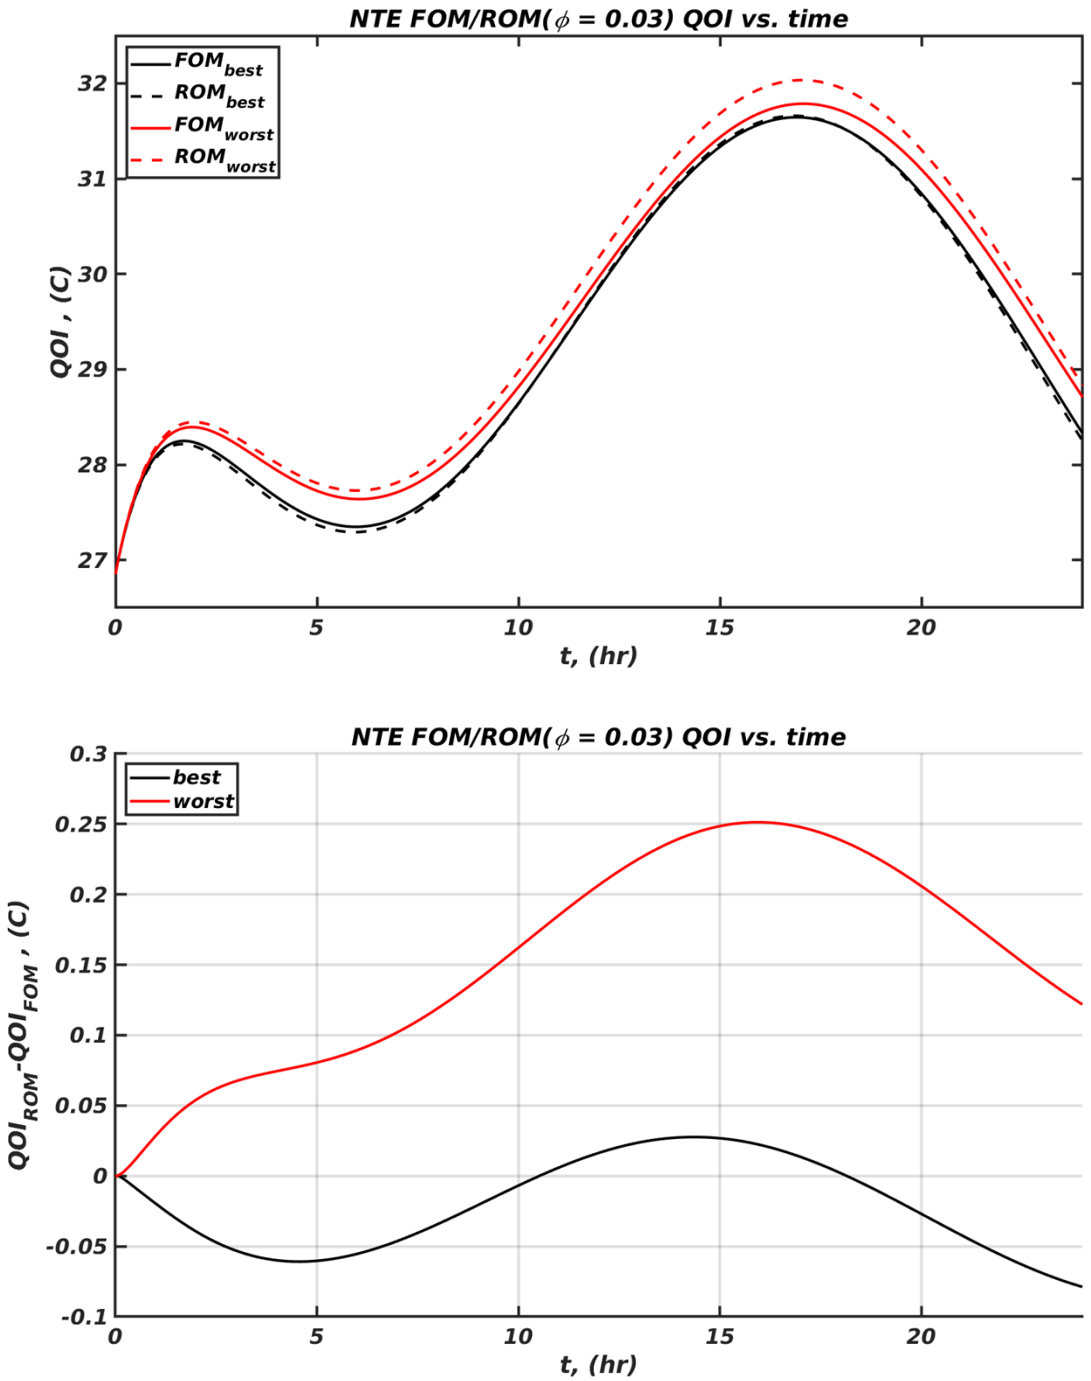

Figure 5-6. FOM/ROM QOI vs. time for best (run 84) and worst (run 40) case agreement across NTE LHC simulations. Upper pane displays QOI vs. time for FOM and ROM simulations. Lower pane shows difference between FOM and ROM simulations

Figure 5-7 similarly displays transient QOI results for ATE FOM and ROM simulations at φ = 0.003 and φ = 0.03. For both mesh fractions, we again display the best and worst cases from the LHC parameter study. Runs #158 (best) and #173 (worst) are shown for φ = 0.003, and runs #76

(best) and run #168 (worst) are displayed for φ = 0.03. For φ = 0.003 (Fig. 5-7, upper figure), the best case (run #158), shows deviations from the FOM as high as 30K while the worst case (run #178) has deviations as high as 100K towards the end of the simulation. For φ = 0.03 (Fig. 5-7, lower figure), the best case (run #76) has almost no deviation during intermediate times in the simulation and small variations, on the order of 10K towards the beginning and end of the simulation. The worst case (run #168) has the excellent agreement with FOM results at beginning and end of the simulation, with some deviation, on the order of 50K-75K at intermediate times. As expected, the higher mesh fraction ROM results are a significant improvement over the lower mesh fraction results, albeit at a significant increase in computational cost, as described in the previous section. An increase of mesh fraction above 0.03, to perhaps 0.1, would likely result in even closer agreement, but at significantly increased computational cost, and depending on the processor core count, may lead to ROM timings that exceed FOM timings, as was demonstrated for high processor core counts above.

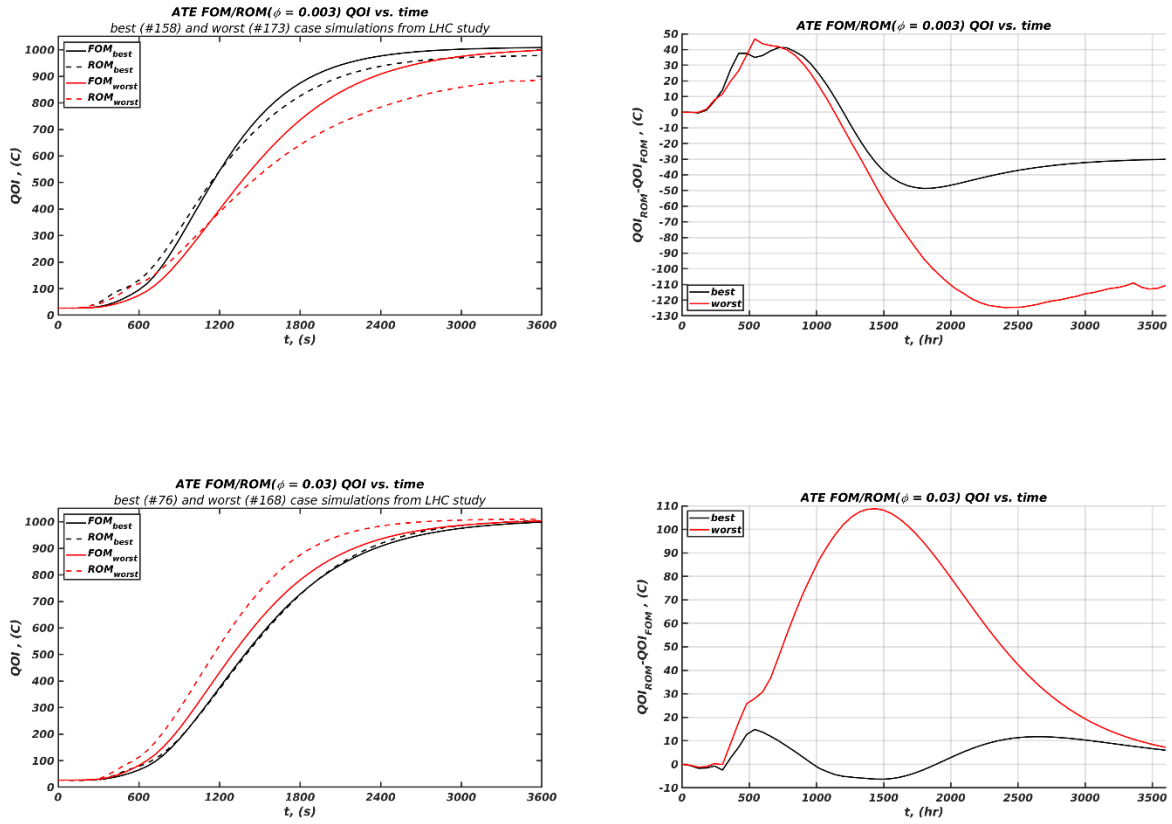

Figure 5-7. FOM/ROM QOI vs. time for best and worst agreement across ATE LHC simulations for ϕ = 0.003 (upper two panes) and 0.03 (lower two panes). Left panes displays QOI for FOM and ROM. Right panes show difference between ROM and FOM simulations.

Comparing NTE to ATE results, we see that NTE simulations at φ = 0.03 are more accurate in comparison to FOM results than ATE simulations. NTE simulations explore a much narrower range of temperatures than ATE do, varying the external temperature over a narrow ~10K (10°C) range, while ATE simulations, initialized at 300K, experience an external temperature almost 1000K (1000°C) higher, at 1283.15 K. In the materials section we discussed the temperature dependence of material properties including specific heat and thermal conductivity as well as latent heat and

temperature-dependent chemistry terms. With the relatively small variation in temperature seen in NTE simulations, material properties are close to constant. Variation from one LHC simulation to the next is simply a matter of the scaling factors used for material properties in the LHC parameter variation described above. In addition, the aluminum components do not reach a melting point, thus eliminating the effects of a latent heat. Chemical terms are also nearly inactive at these relatively low temperatures. For ATE, all these temperature dependencies become important, and sensitivity of material models to variations in temperature are much more significant. This is convolved with the multi-factor parameter variation of the LHC. Latent heat and chemistry terms are also important. The combined effect is that we see a larger variation between FOM and ROM when we move from NTE to ATE, and we find a greater importance in using sufficiently high mesh fractions. Acceptable deviation between ROM and FOM results is a decision to be made by the analyst in the context of specific program requirements.

5.2.2. QOI Statistics

In Figure 5-8 below, we display transient statistical metrics of FOM and ROM (φ = 0.003/0.03) QOI temperatures over the full 256 parameter variations in the LHC set for the ATE condition, including the ensemble average (<T>) and standard deviation (σ(T)). These metrics give a different view of the agreement between QOI values obtained from FOM and ROM LHC simulations than the best and worst cases shown above.

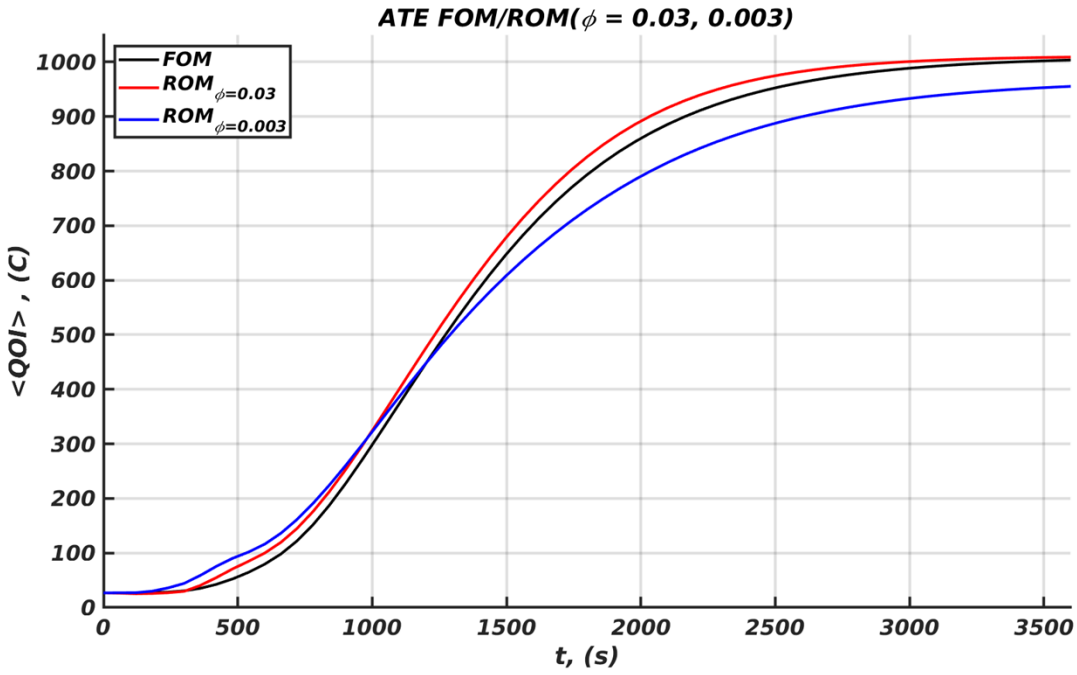

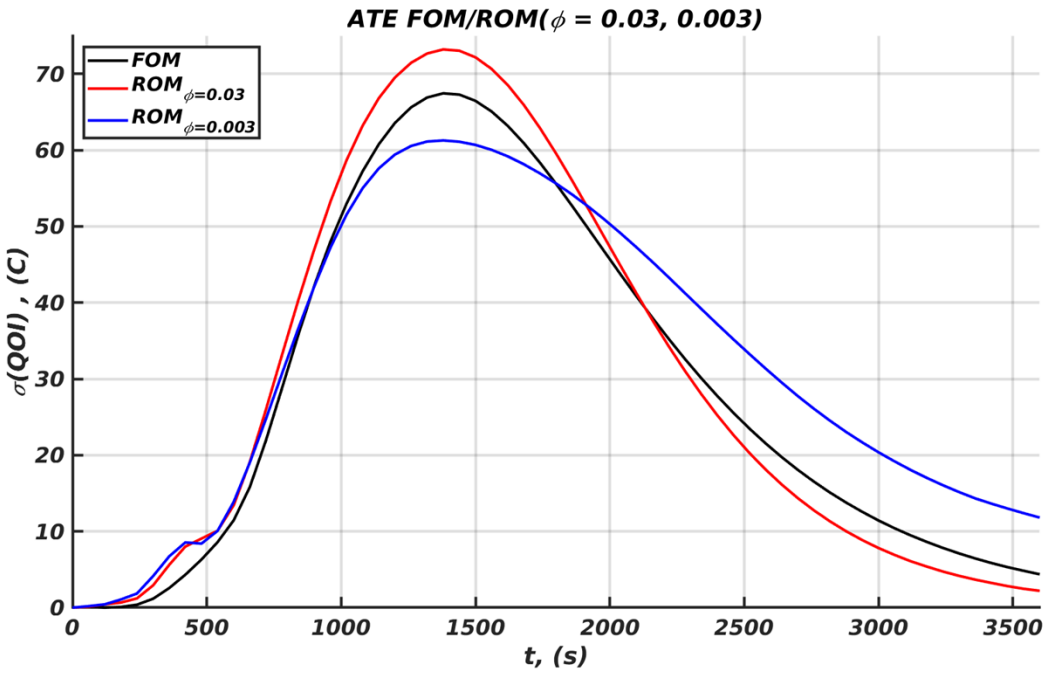

Figure 5-8. ATE FOM/ROM (φ = 0.003, 0.03) LHC QOI statistics vs. time: <T> (upper) and σ(T) (lower)

From these metrics, we see that φ = 0.03 results in significantly better agreement with FOM results than φ = 0.003. Somewhat unexpectedly, at intermediate to late times, where φ = 0.003 results show under-prediction relative to FOM results, φ = 0.03 tends to show slight over-prediction. Temporal trends tend to be consistent for both metrics between FOM and ROM results, indicating further that this ROM paradigm does result in distributional data that is representative of FOM studies at significantly reduced computational cost. That is, the computationally cheaper ROM LHC study not only reproduces QOI metric averages that are similar to their FOM counterparts, but also variances.

Figure 5-9 below focuses on additional statistical metrics, in this case for NTE simulations. The figure displays distributional results for several measures of the difference between ROM and FOM QOI temperatures over the LHC study. These include the minimum (min), maximum(max), and average (< >)temperature difference (over time) between ROM and FOM QOI temperatures for each run as well as the average absolute value(<| |>) and the standard deviation (σ) of the ROM and FOM QOI temperature difference across each run. Histograms here indicate how these metrics are distributed over the full LHC set.

$\phi = 0.03$

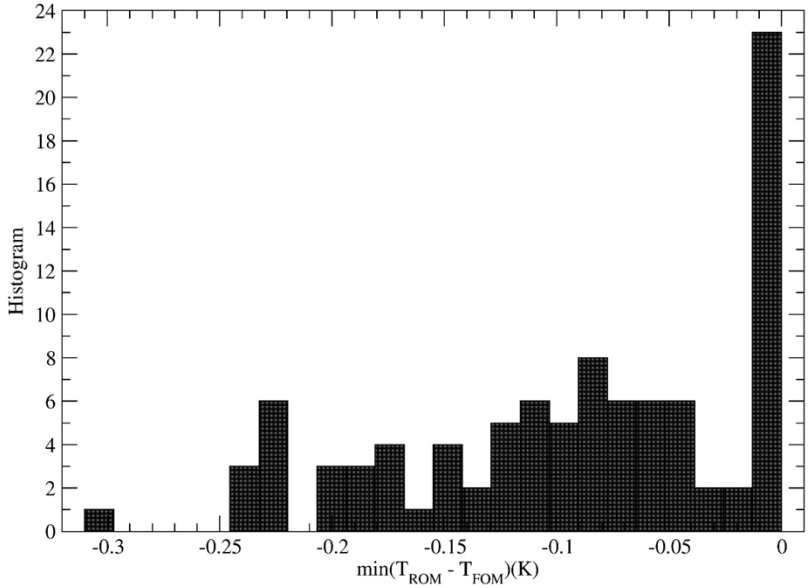

$\phi = 0.03$

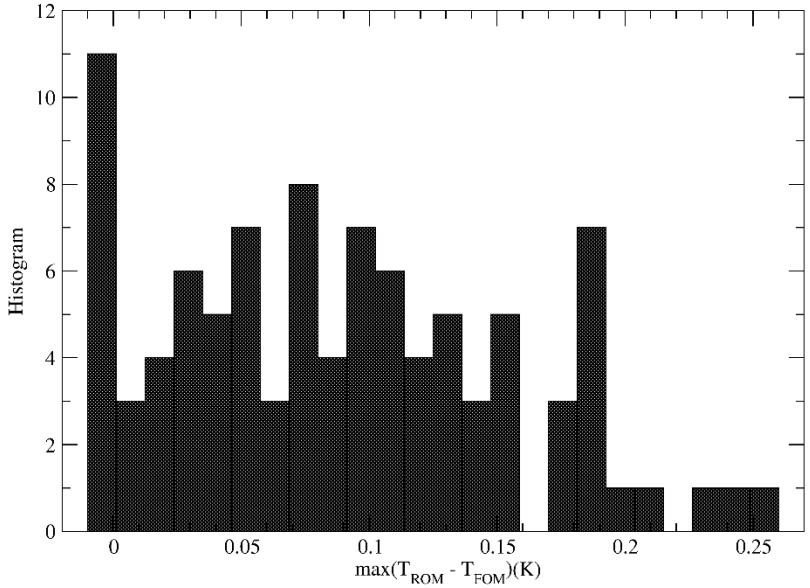

$\phi = 0.03$

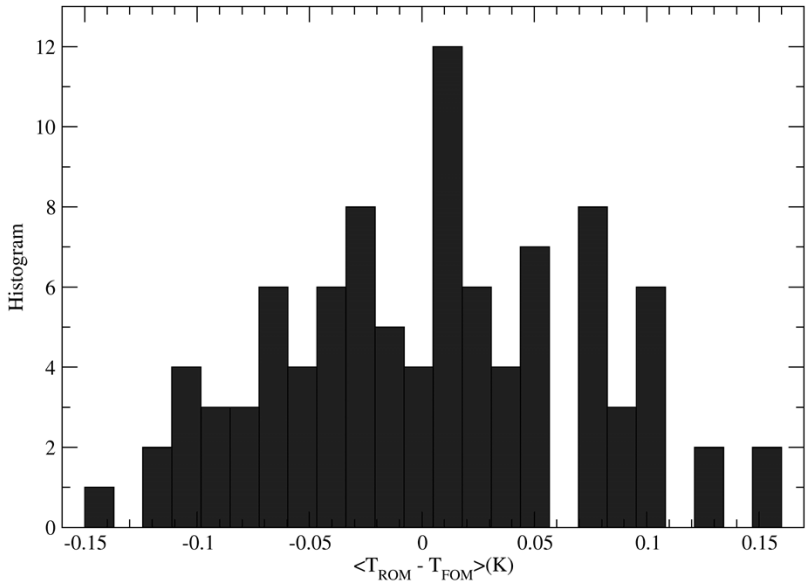

$\phi = 0.03$

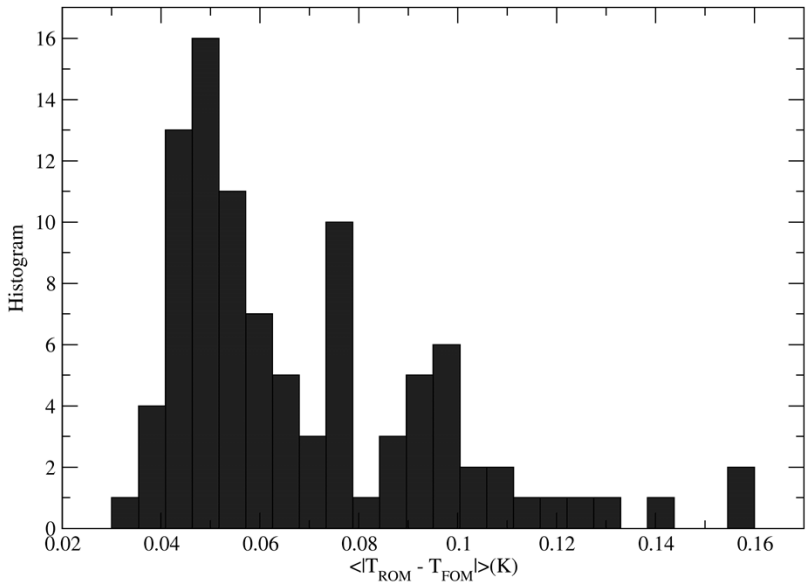

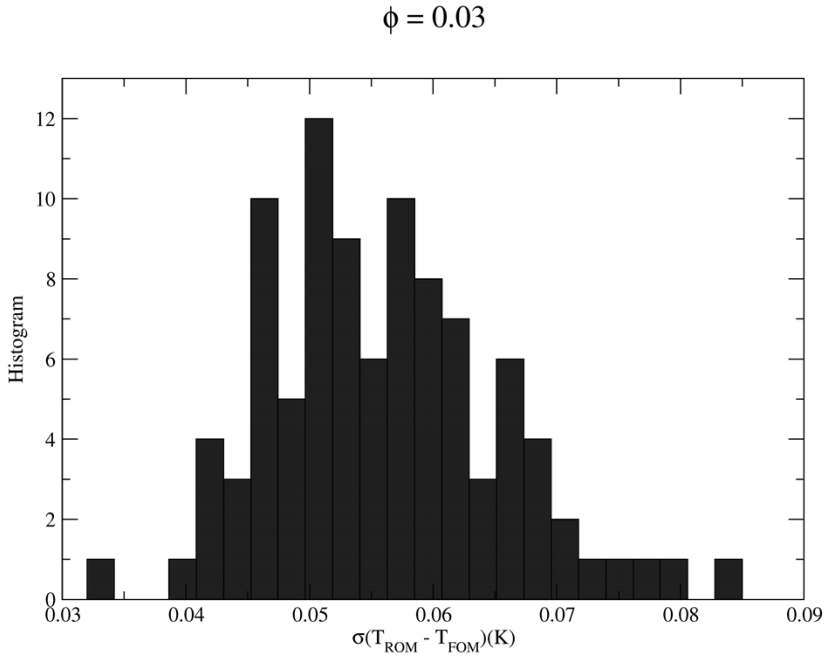

Figure 5-9. NTE LHC QOI statistics for difference (over full simulation time) between ROM and FOM QOI including (top to bottom): minimum (min), maximum (max), average(< >), average absolute value(<| |>, and standard deviation(σ).

For minimum difference, the distribution is centered around -0.1K with a peak near 0K and minimum at ~-0.3K. Maximum differences again center around 0.1K with a peak at ~0K and maximum of ~0.25K. This is consistent with our results for best and worst NTE cases as described above. Average temperature differences are centered and peaked at 0K with extremes at -0.15K and 0.15K, respectively. Average absolute differences are peaked at 0.05K and also extend to ~0.15K. Standard deviations peak at ~0.05K and extend to ~0.85K. These distributions further confirm the excellent agreement between FOM and ROM results for NTE conditions at $\varphi = 0.03$. Such findings point us toward a general application of ROM techniques in normal thermal analysis scenarios. To remind the reader, the SVD bases used to run ROM simulations came from a single FOM scenario run with the nominal set of variational parameters identified for the LHC study. The excellent agreement between FOM and ROM QOI results across the entire ensemble of the LHC study in which those parameters were varied within expected norms would indicate that a thermal analyst could reasonably utilize a ROM approach for the LHC in lieu of a full FOM approach to obtain sensitivities. As demonstrated above, doing so would require significantly less computational expense than a FOM LHC study. However, the agreement/disagreement between FOM and ROM results is a function of many factors including the FEM geometry/mesh, the initial and boundary conditions, material properties, specific physics in the model, and simulation duration. Thus, one must proceed with caution in applying ROM techniques to a new scenario.

Figure 5-10 below displays the same statistical metrics, this time for ATE conditions. We display each metric with $\varphi = 0.003$ and $\varphi = 0.03$ results together for ease of comparison.

$\phi = 0.003$

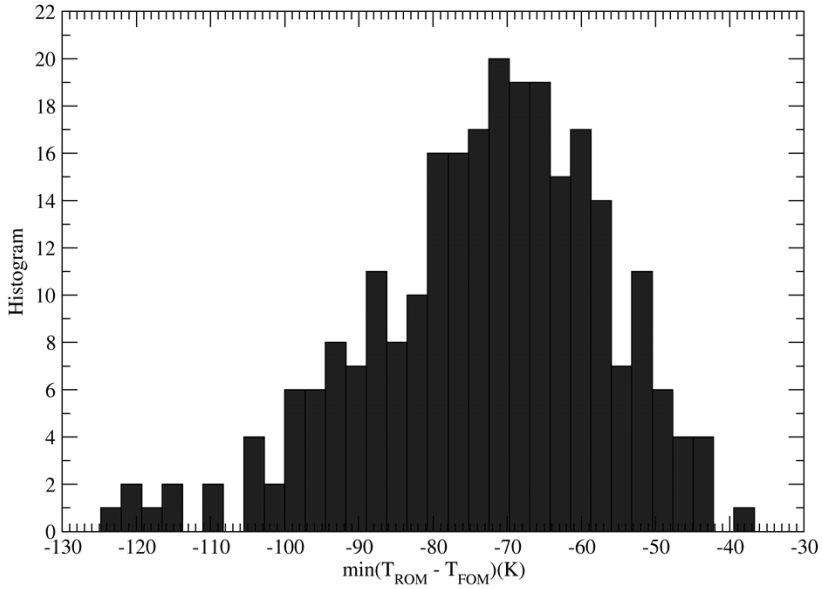

$\phi = 0.03$

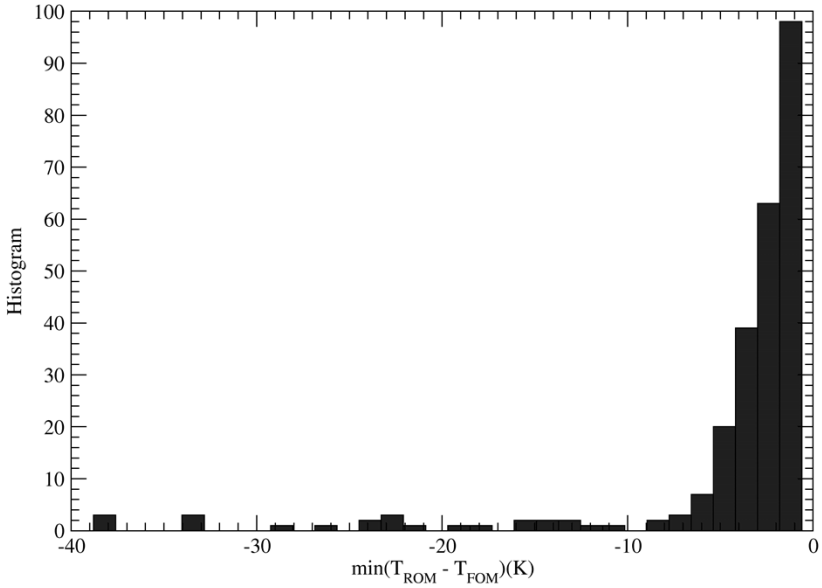

$\phi = 0.003$

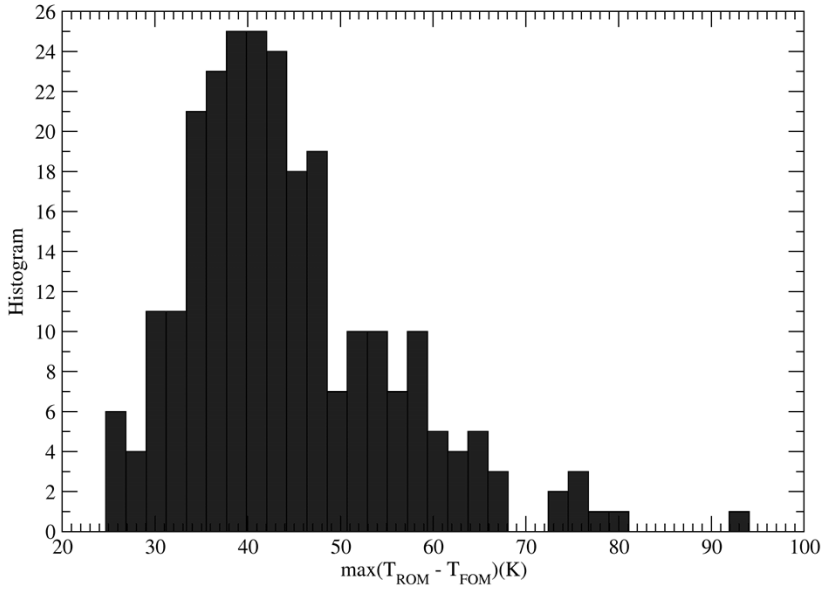

$\phi = 0.03$

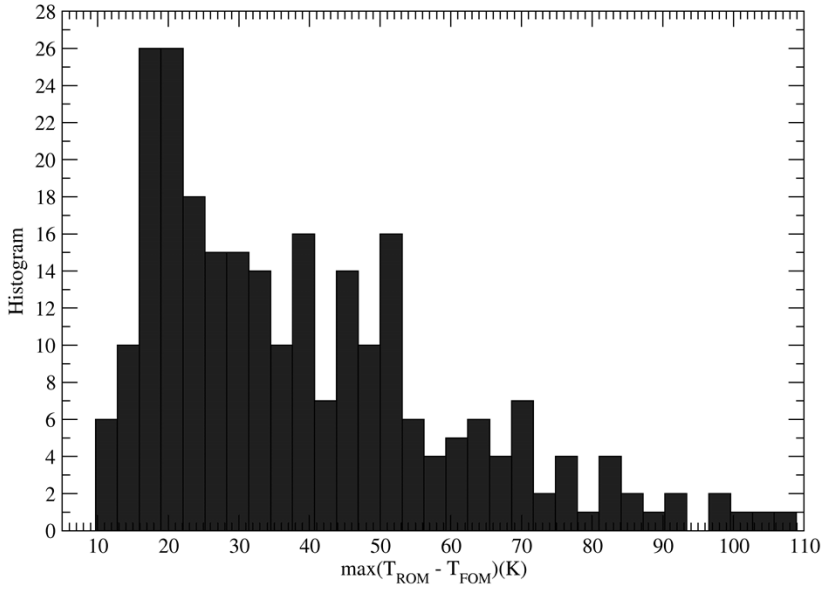

$\phi = 0.003$

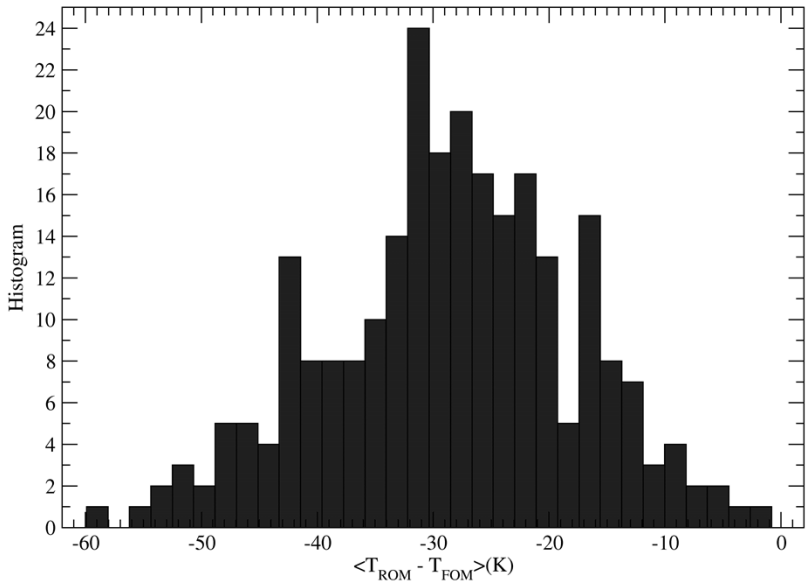

$\phi = 0.03$

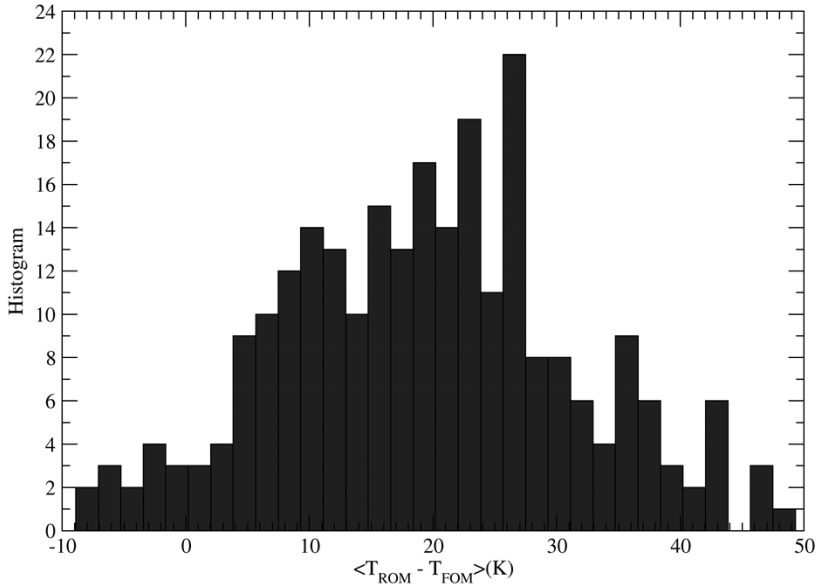

$\phi = 0.003$

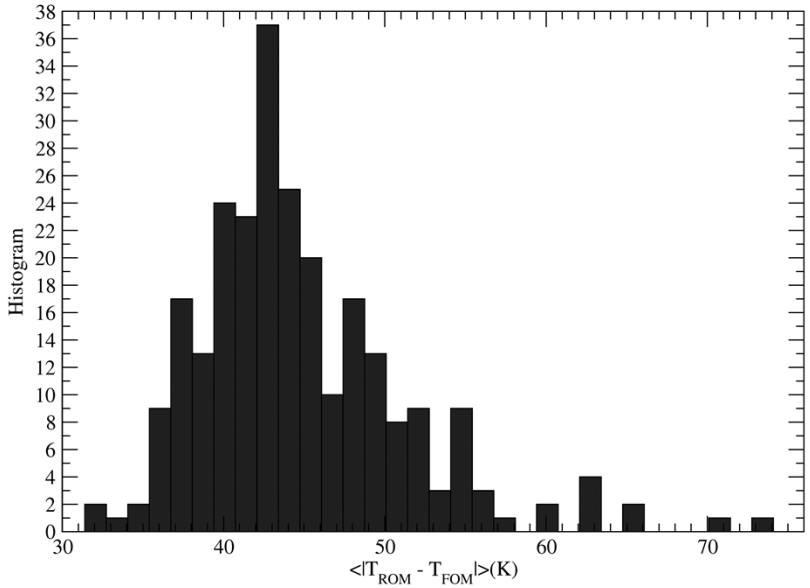

$\phi = 0.03$

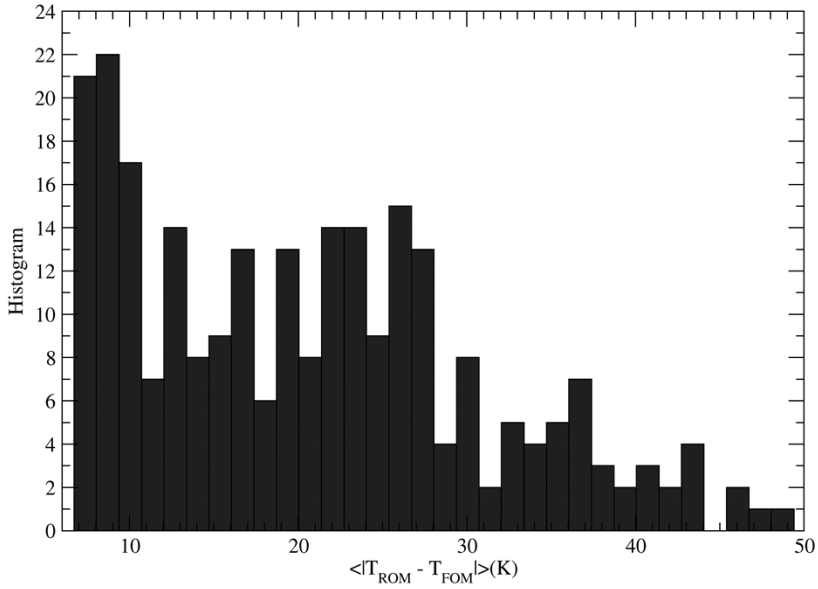

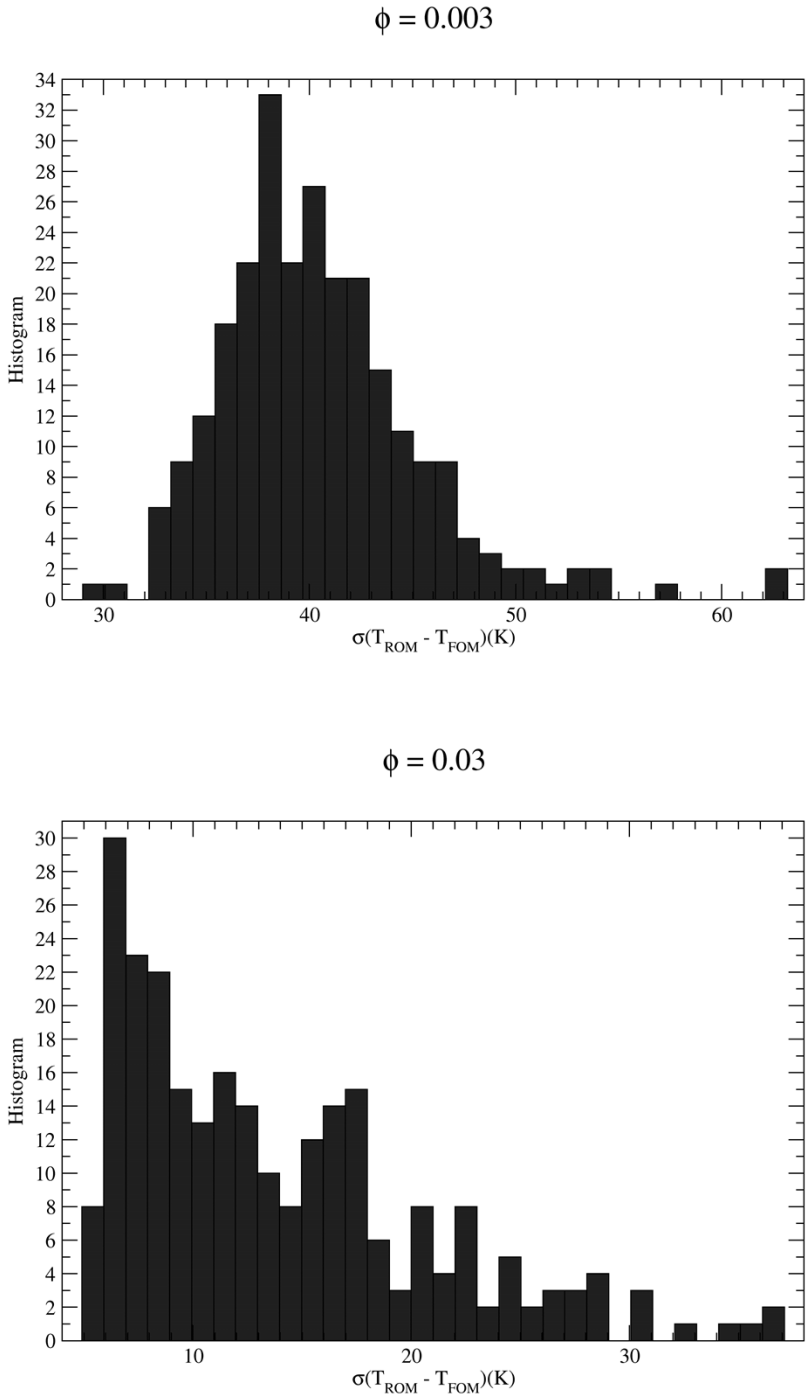

Figure 5-10. ATE LHC QOI statistics for difference (over full simulation time) between ROM and FOM QOI including minimum, maximum, average, average absolute value, and standard deviation. φ = 0.003 above and φ = 0.03 below for each metric.

Minimum QOI temperature differences for $\varphi = 0.003$ vary from ~-125K to ~-35K, peaking at ~-70K. Moving to $\varphi = 0.003$ drastically improves the agreement, shifting the distribution to a sharp peak at ~0K with a minimum at ~-38K. The vast majority of the distribution lies within the -10K to 0K range. Maximum QOI temperature differences for $\varphi = 0.003$ vary from ~25K to ~95K, peaking at ~40K. For $\varphi = 0.03$, the peak shifts to ~20K, but the range remains relatively large at 10K to nearly 110K. Average QOI differences for $\varphi = 0.003$ vary from -60K to 0K, peaked at nearly 30K. Moving to $\varphi = 0.03$, the average difference range shifts to the -10K to 50K range, with a peak at ~25K. It is somewhat surprising that we don't see a significant improvement in the average difference between FOM and ROM when we move from a mesh fraction of 0.003 to 0.03. For average absolute differences, $\varphi = 0.003$ shows variation from ~30K to ~75K with a peak at ~43K. $\varphi = 0.03$ has a range of 0K to ~50K with a slow decay from a peak at ~10K. Standard deviations for $\varphi = 0.003$ peak at ~40K varying from ~30K to ~65K. For $\varphi = 0.03$, the range extends from ~3k to ~35K with a slow decay from a peak at ~5K.

For ATE, this view of the statistics of QOI temperature differences between ROM and FOM highlights some important points. For one, the magnitude of the average difference over time between ROM and FOM and the width of the distribution is not a strong function of the mesh fraction, though it is slightly better in the $\varphi = 0.03$ case. This is surprising, and not what we might have expected from the best- and worst- case transient data shown above. While the minimum QOI difference distribution is much better for $\varphi = 0.03$ than $\varphi = 0.003$, the maximum differences are slightly worse. Additionally, the actual average differences between FOM and ROM QOI (peaks at 30K and 25K) are not large relative to the actual system temperatures which reach ~1300K, though the need to be precise within a specific temperature range is a function of the analyst's accuracy requirements. We do not speak to that here but leave this to the user.

5.2.3. Comparison Across the Larger Model Geometry

Also important to an investigation of the applicability of ROM based solutions to problems similar to the Crash-Burn scenario is an understanding of the accuracy of ROM solutions across the larger simulation domain. Up to this point, we have been discussing ROM accuracy with respect to a single QOI in the simulation domain, that is, the temperature at the center of mass for the mass mock described above. To understand the impact of ROM simulations on the full simulation domain for NTE scenarios, we display both FOM temperatures and the relative difference between FOM and ROM ($\phi = 0.03$) temperatures for a selection of simulation times on a cross-section through the simulation domain in Figure 5-11. We show results for both the best agreement (run #84) and worst agreement (run #40) cases from the LHC set as described above.

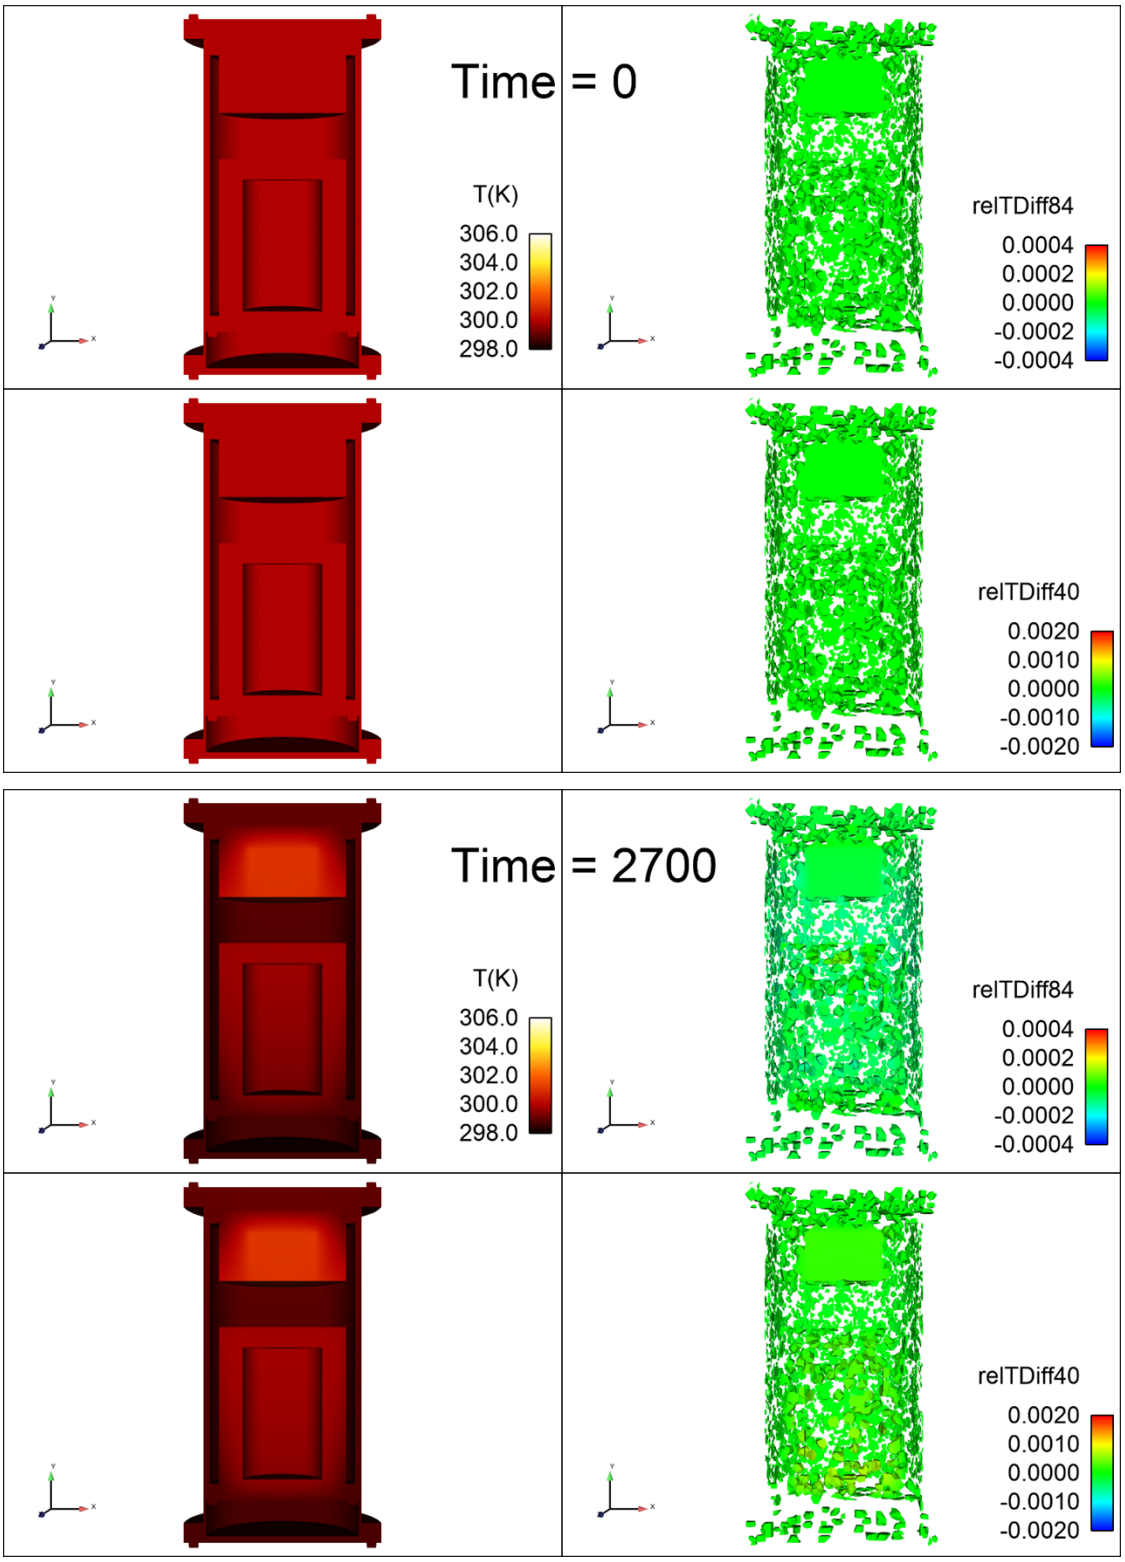
Time = 0
T(K)
306.0
304.0
302.0
300.0
298.0
relTDiff84
0.0004
0.0002
0.0000
-0.0002
-0.0004
relTDiff40
0.0020
0.0010
0.0000
-0.0010
-0.0020
Time = 2700
T(K)
306.0
304.0
302.0
300.0
298.0
relTDiff84
0.0004
0.0002
0.0000
-0.0002
-0.0004
relTDiff40
0.0020
0.0010
0.0000
-0.0010
-0.0020

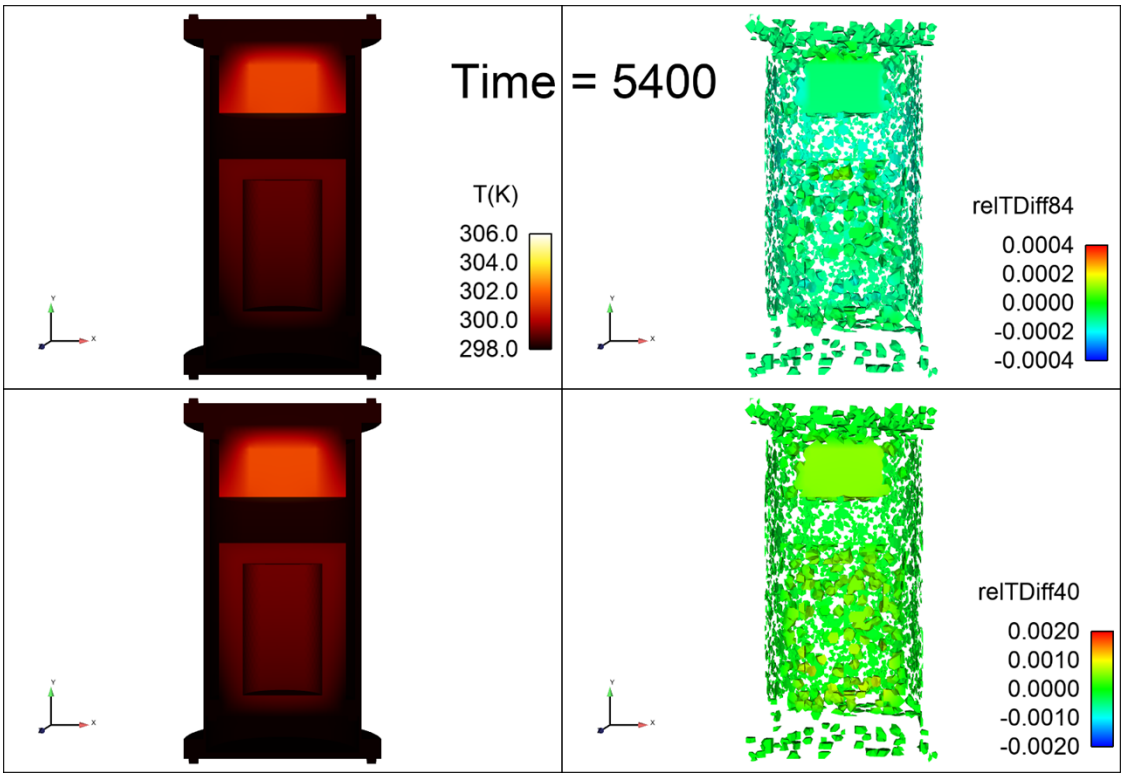
Time = 5400
T(K)
306.0
304.0
302.0
300.0
298.0
relTDiff84
0.0004
0.0002
0.0000
-0.0002
-0.0004
relTDiff40
0.0020
0.0010
0.0000
-0.0010
-0.0020

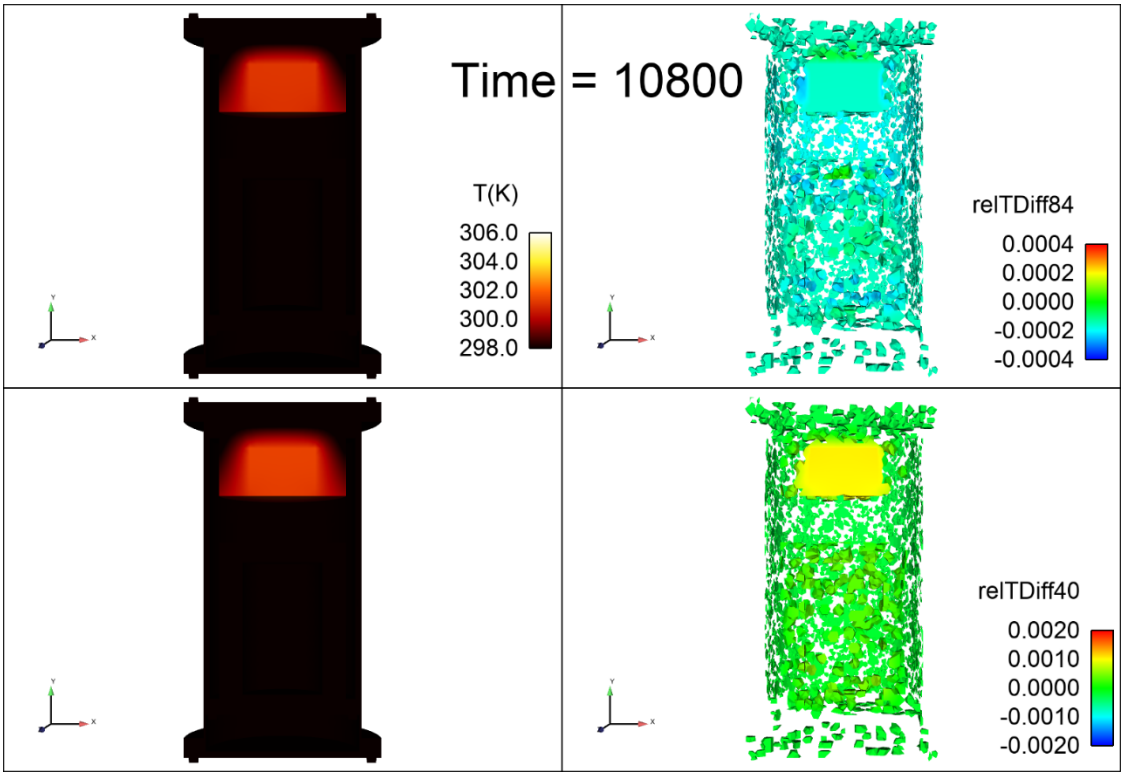
Time = 10800
T(K)
306.0
304.0
302.0
300.0
298.0
relTDiff84
0.0004
0.0002
0.0000
-0.0002
-0.0004
relTDiff40
0.0020
0.0010
0.0000
-0.0010
-0.0020

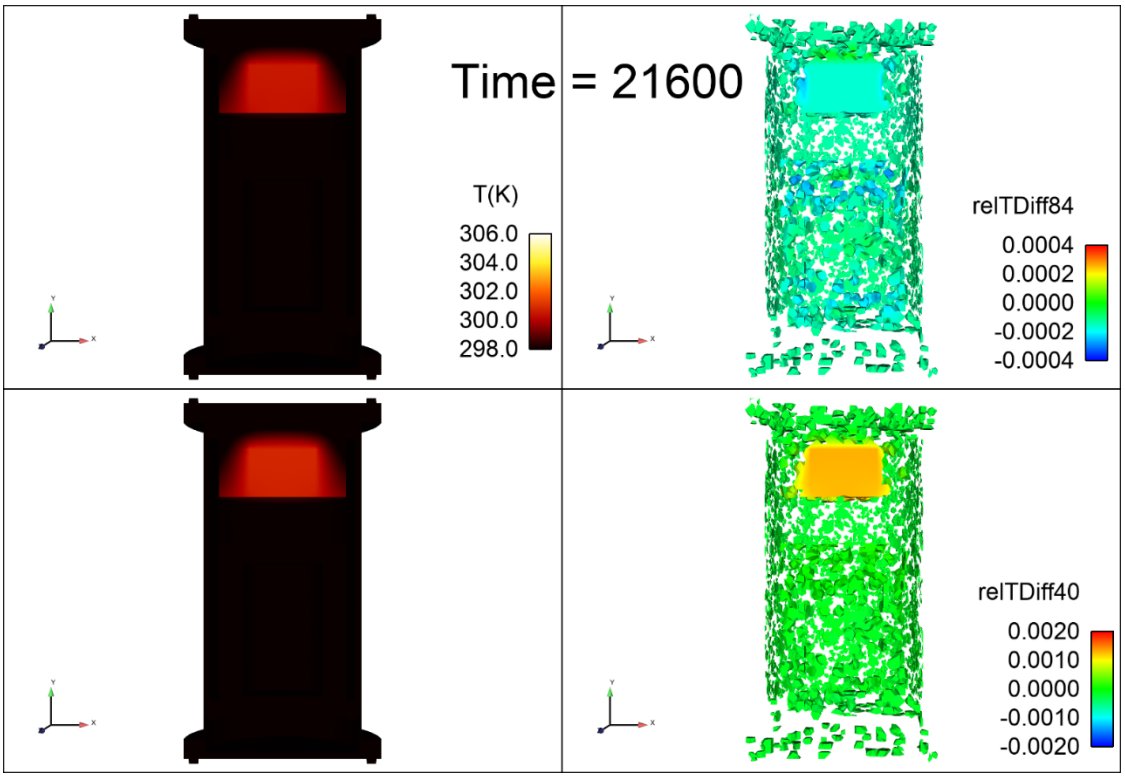
Time = 21600
T(K)
306.0
304.0
302.0
300.0
298.0
relTDiff84
0.0004
0.0002
0.0000
-0.0002
-0.0004
relTDiff40
0.0020
0.0010
0.0000
-0.0010
-0.0020

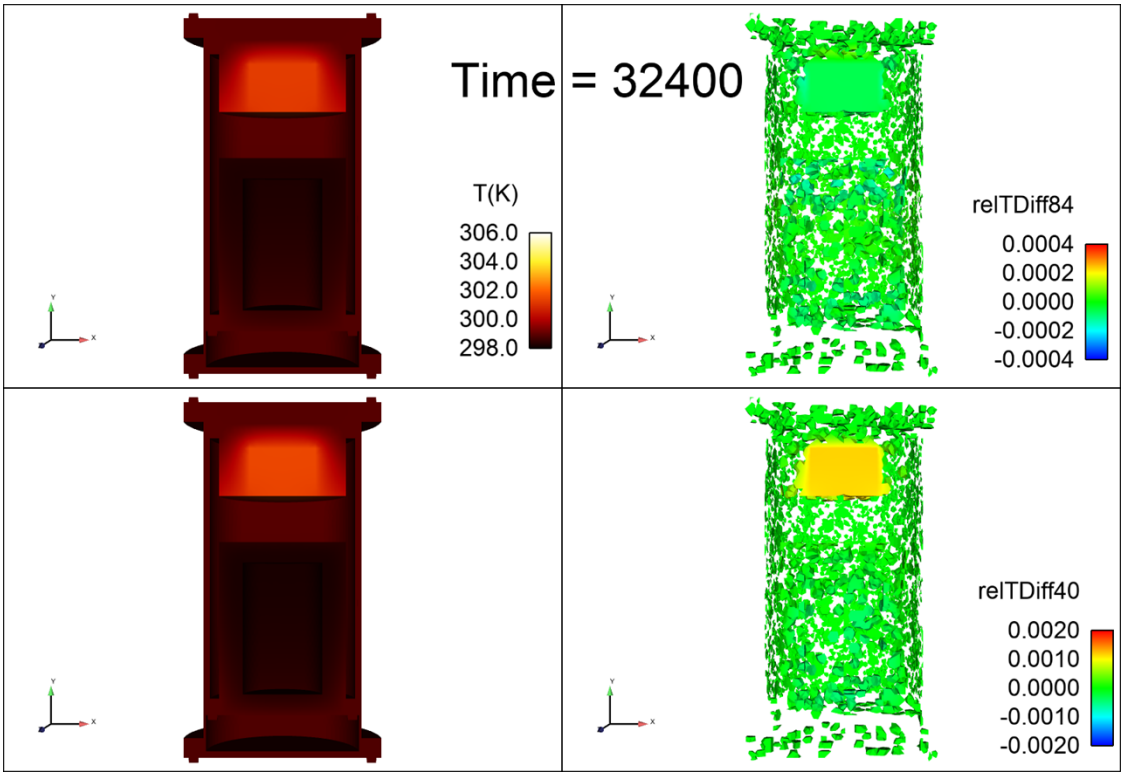
Time = 32400
T(K)
306.0
304.0
302.0
300.0
298.0
relTDiff84
0.0004
0.0002
0.0000
-0.0002
-0.0004
relTDiff40
0.0020
0.0010
0.0000
-0.0010
-0.0020

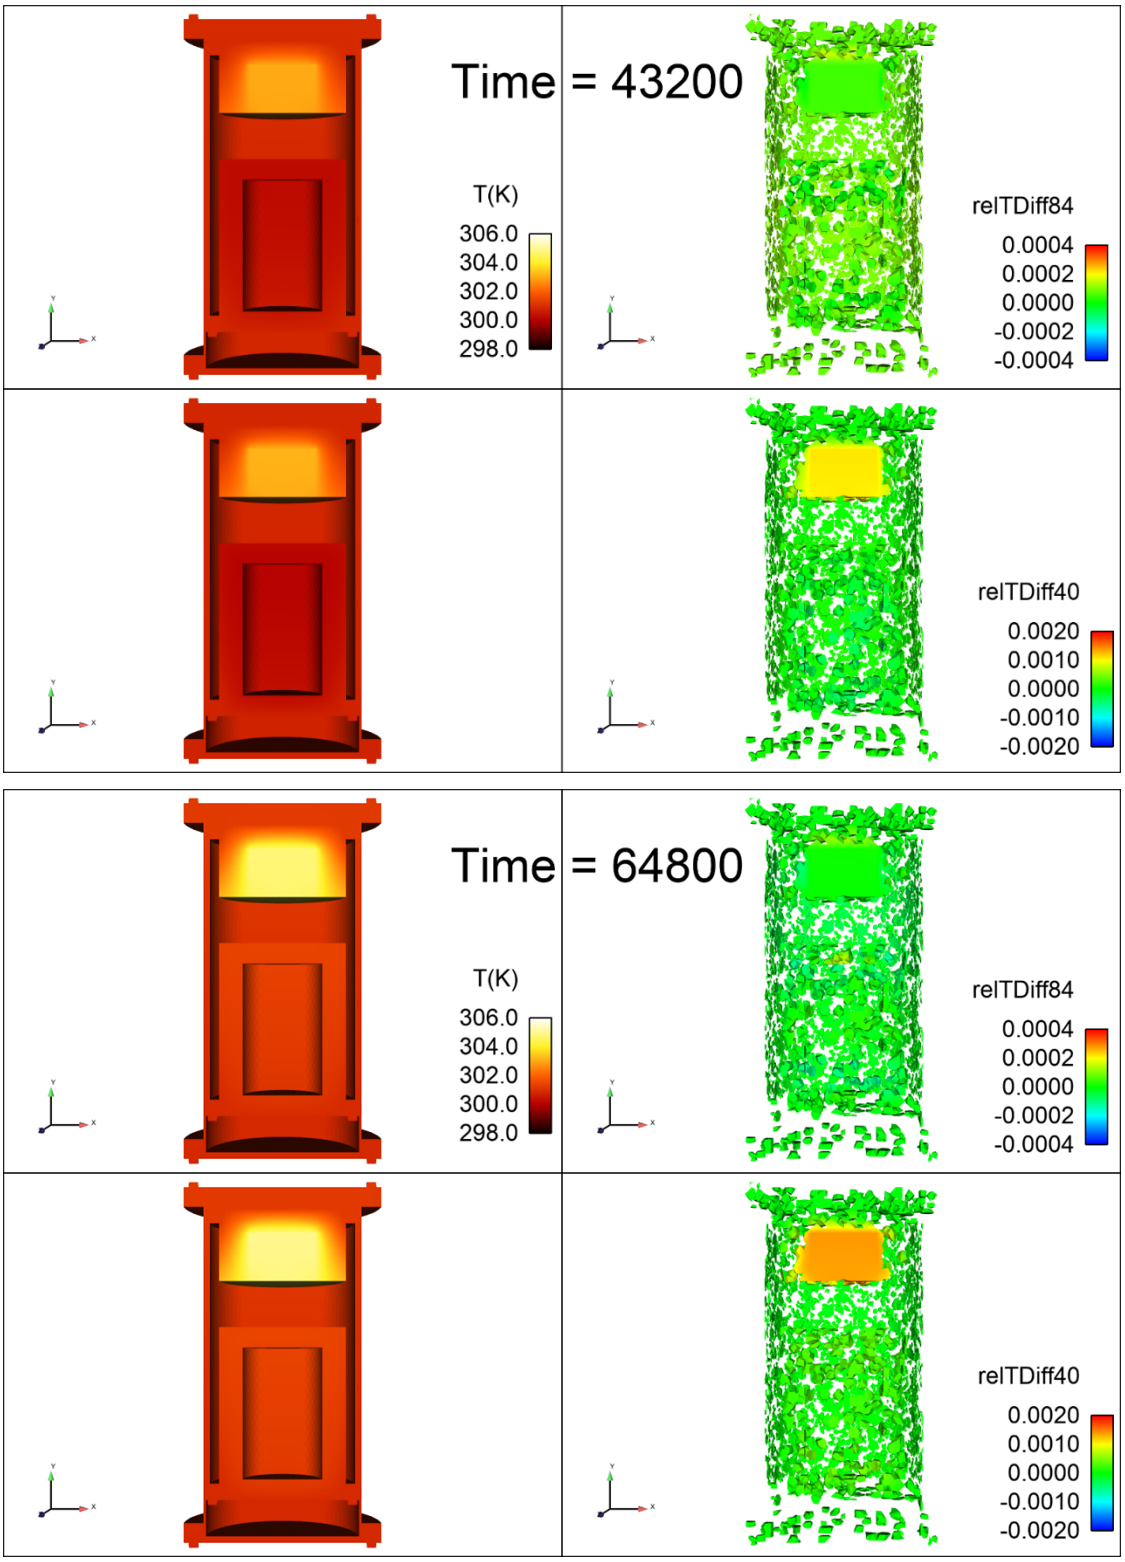
Time = 43200
T(K)
306.0
304.0
302.0
300.0
298.0
relTDiff84
0.0004
0.0002
0.0000
-0.0002
-0.0004
relTDiff40
0.0020
0.0010
0.0000
-0.0010
-0.0020
Time = 64800
T(K)
306.0
304.0
302.0
300.0
298.0
relTDiff84
0.0004
0.0002
0.0000
-0.0002
-0.0004
relTDiff40
0.0020
0.0010
0.0000
-0.0010
-0.0020

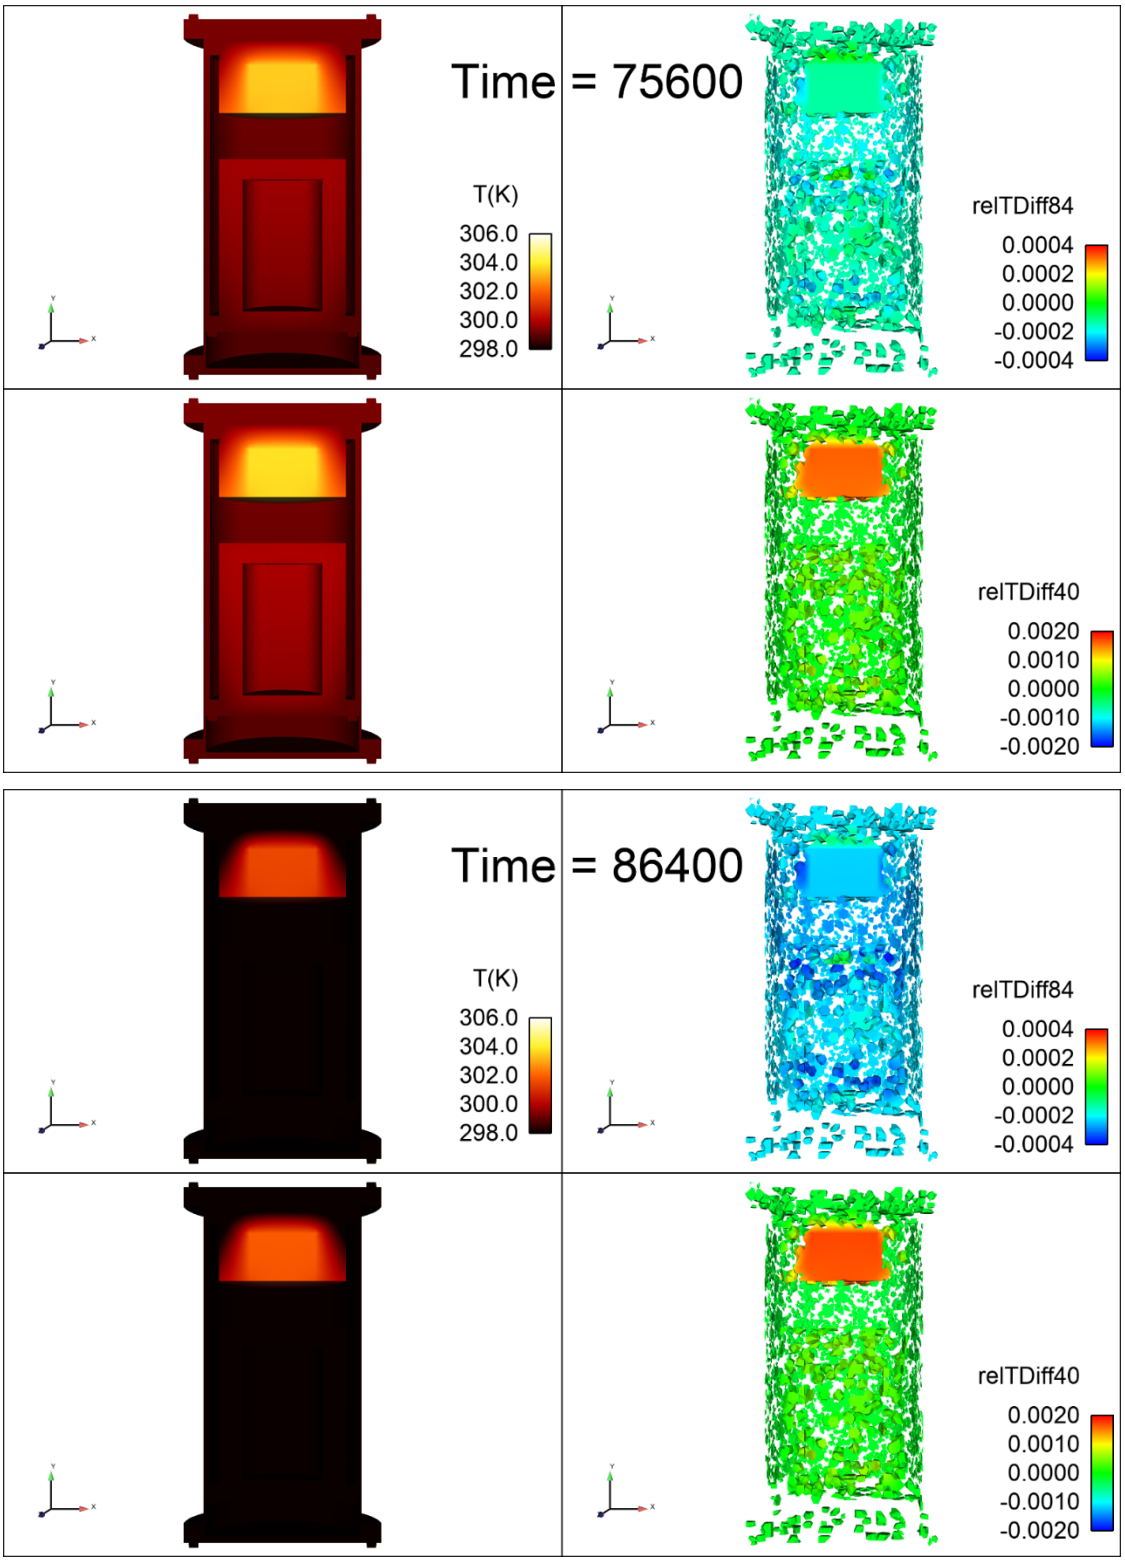

Figure 5-11. (left) NTE FOM temperatures shown in cross section at times throughout simulation for best (run 84) and worst (run 40) agreement with ROM. (right) Relative (fractional) differences between FOM and ROM temperatures shown on ROM sample mesh for these cases shown in cross section. ROM mesh fraction is $\phi = 0.03$.

Here the relative temperature difference is defined as:

$$\text{relTDiff} = T_{ROM}/T_{FOM} - 1$$

Note that the color scale for the best case is defined over the range ($-4x10^{-4}$, $4x10^{-4}$) whereas the scale for the worst case is defined over the 5x larger range of ($-2x10^{-3}$, $2x10^{-3}$). The largest relative difference in the best case occurs towards the end of the simulation (24 hrs = 86,400 sec) and is seen in two vertical bands towards the vertical middle of the domain, though relatively high values are also seen at 3 hrs (10,800 sec), 6 hrs (21,600 sec), and 21 hrs (75,600 sec) in the same regions. For the worst case, the largest values for the relative temperature difference are seen at the end of the simulation at 24 hrs (86,400 sec), located at the mass mock. The largest variations for the worst case at any of the displayed times is also at the mass mock, and the relative temperature difference shows some non-monotonic behavior with time. Relative differences for best and worst QOI cases over the entire domain and over the full simulation time are quite good, and it is likely that these are adequate for a standard UQ study.

Figure 5-12 shows similar results for FOM and ROM ATE scenarios at $\varphi = 0.03$, again for best (run #76) and worst (run #168) cases from the LHC study. Here, the range of relative temperature differences are limited to the (-0.25, 0.25) range. We have artificially limited the range to this, though relative temperature differences in both best and worst cases exceed this range at certain simulation times. This limited range allows us to more adequately view regions of the domain in which variations are relatively small. Any relative difference over 25% should be considered quite high.

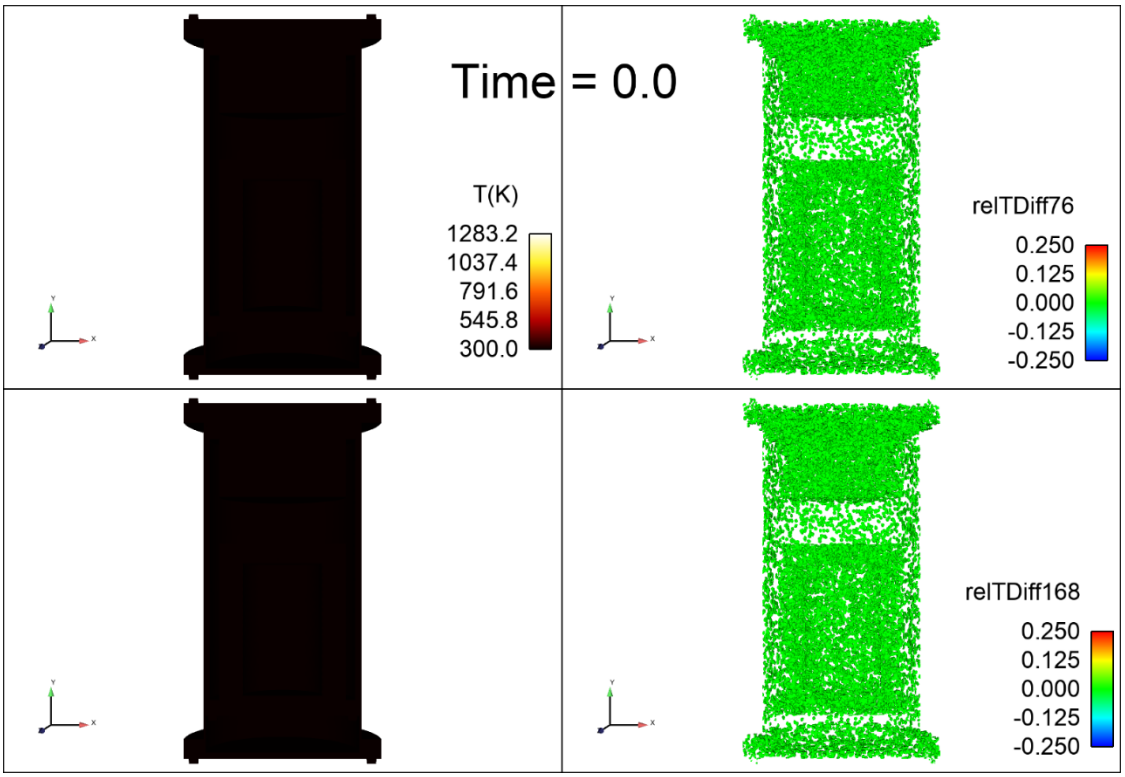

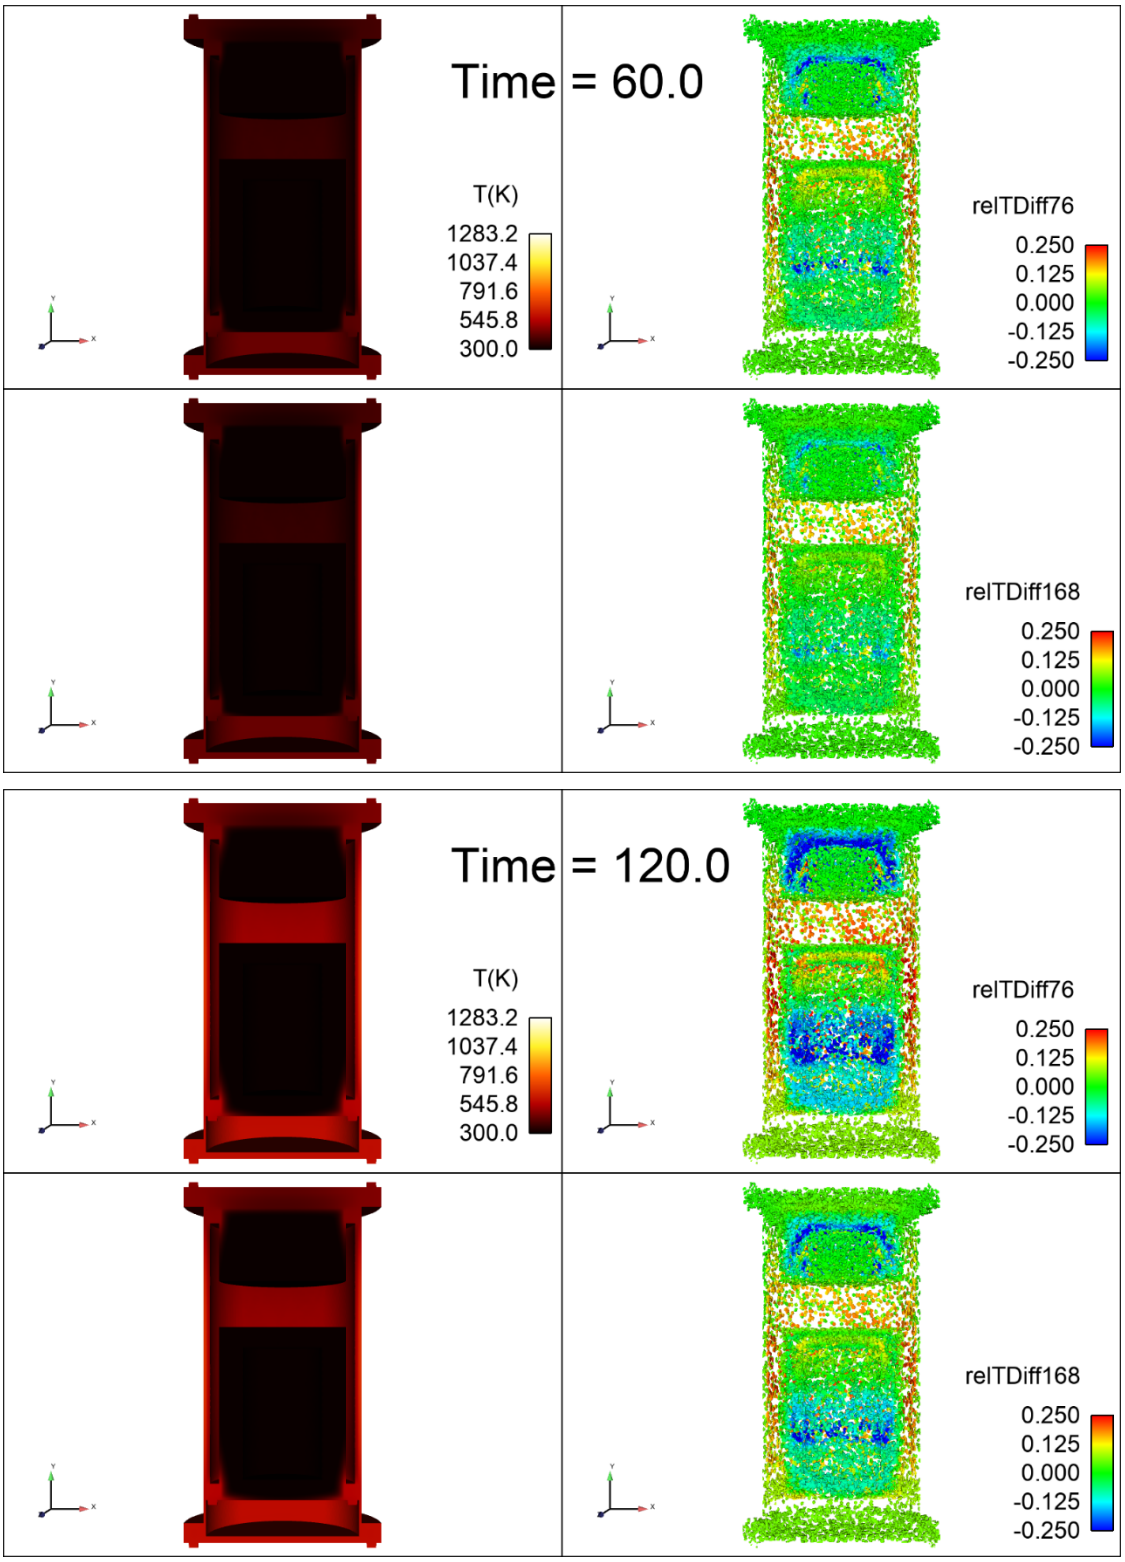
Time = 60.0
T(K)
1283.2
1037.4
791.6
545.8
300.0
relTDiff76
0.250
0.125
0.000
-0.125
-0.250
relTDiff168
0.250
0.125
0.000
-0.125
-0.250
Time = 120.0
T(K)
1283.2
1037.4
791.6
545.8
300.0
relTDiff76
0.250
0.125
0.000
-0.125
-0.250
relTDiff168
0.250
0.125
0.000
-0.125
-0.250

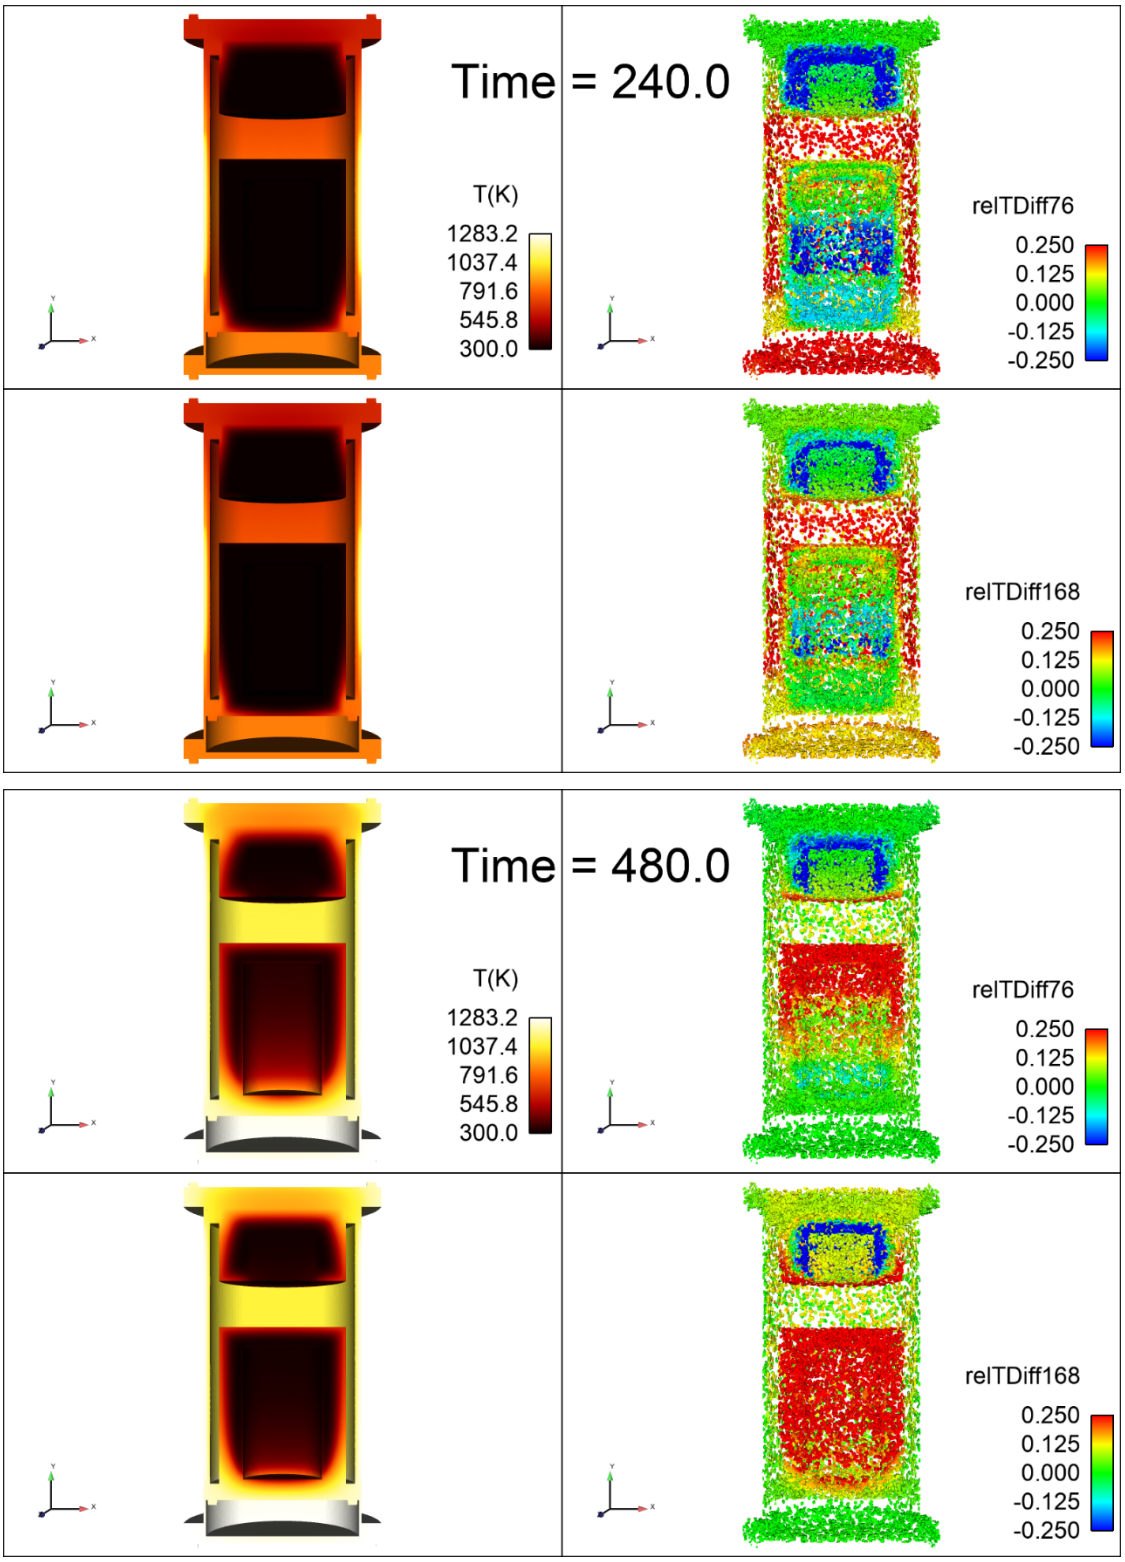
Time = 240.0
T(K)
1283.2
1037.4
791.6
545.8
300.0
relTDiff76
0.250
0.125
0.000
-0.125
-0.250
relTDiff168
0.250
0.125
0.000
-0.125
-0.250
Time = 480.0
T(K)
1283.2
1037.4
791.6
545.8
300.0
relTDiff76
0.250
0.125
0.000
-0.125
-0.250
relTDiff168
0.250
0.125
0.000
-0.125
-0.250

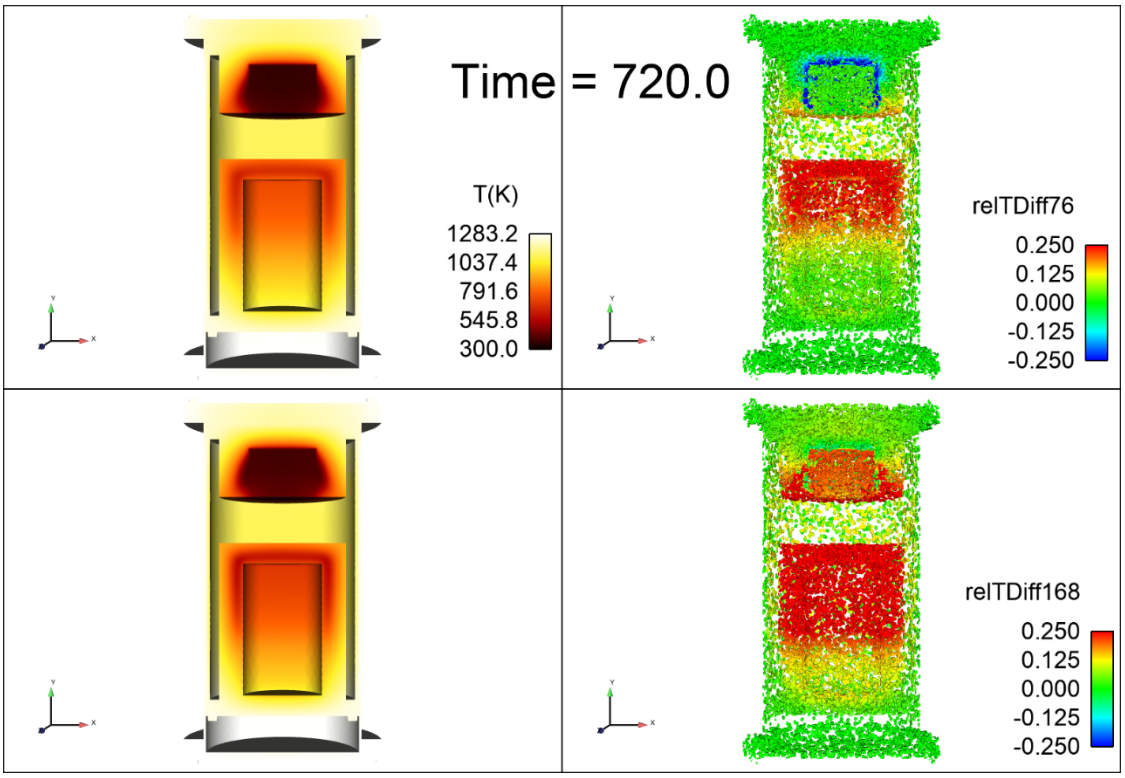
Time = 720.0
T(K)
1283.2
1037.4
791.6
545.8
300.0
relTDiff76
0.250
0.125
0.000
-0.125
-0.250
relTDiff168
0.250
0.125
0.000
-0.125
-0.250

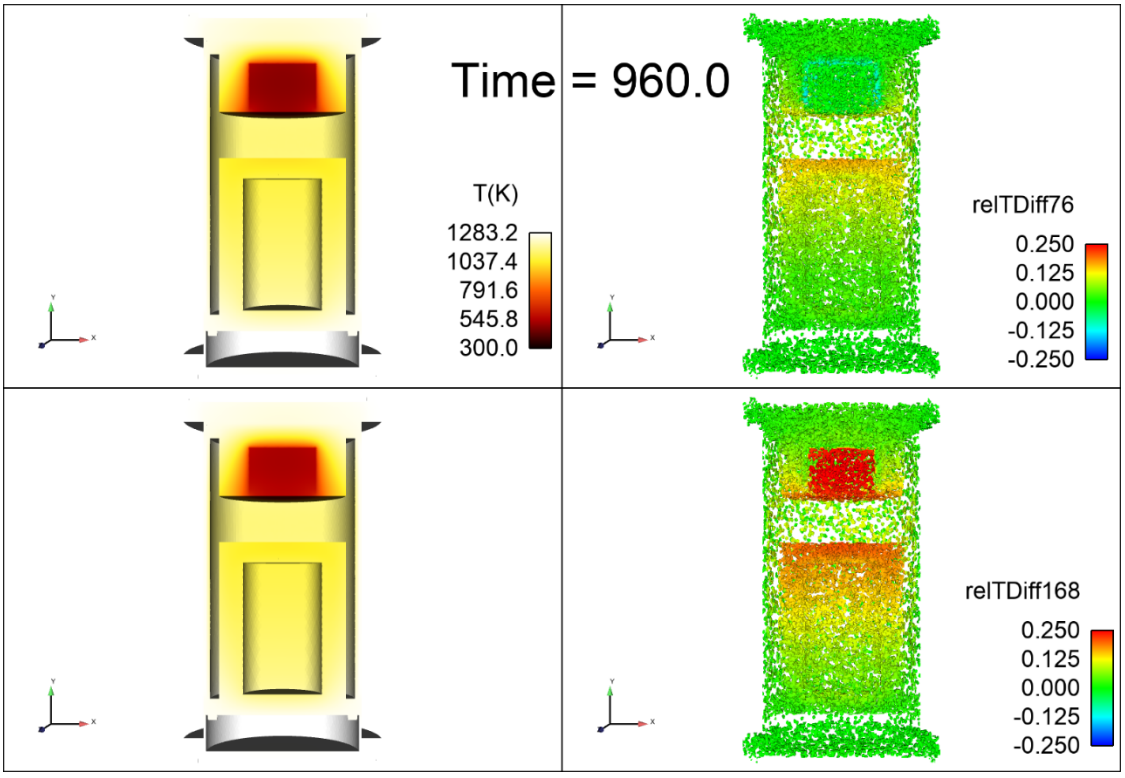
Time = 960.0
T(K)
1283.2
1037.4
791.6
545.8
300.0
relTDiff76
0.250
0.125
0.000
-0.125
-0.250
relTDiff168
0.250
0.125
0.000
-0.125
-0.250

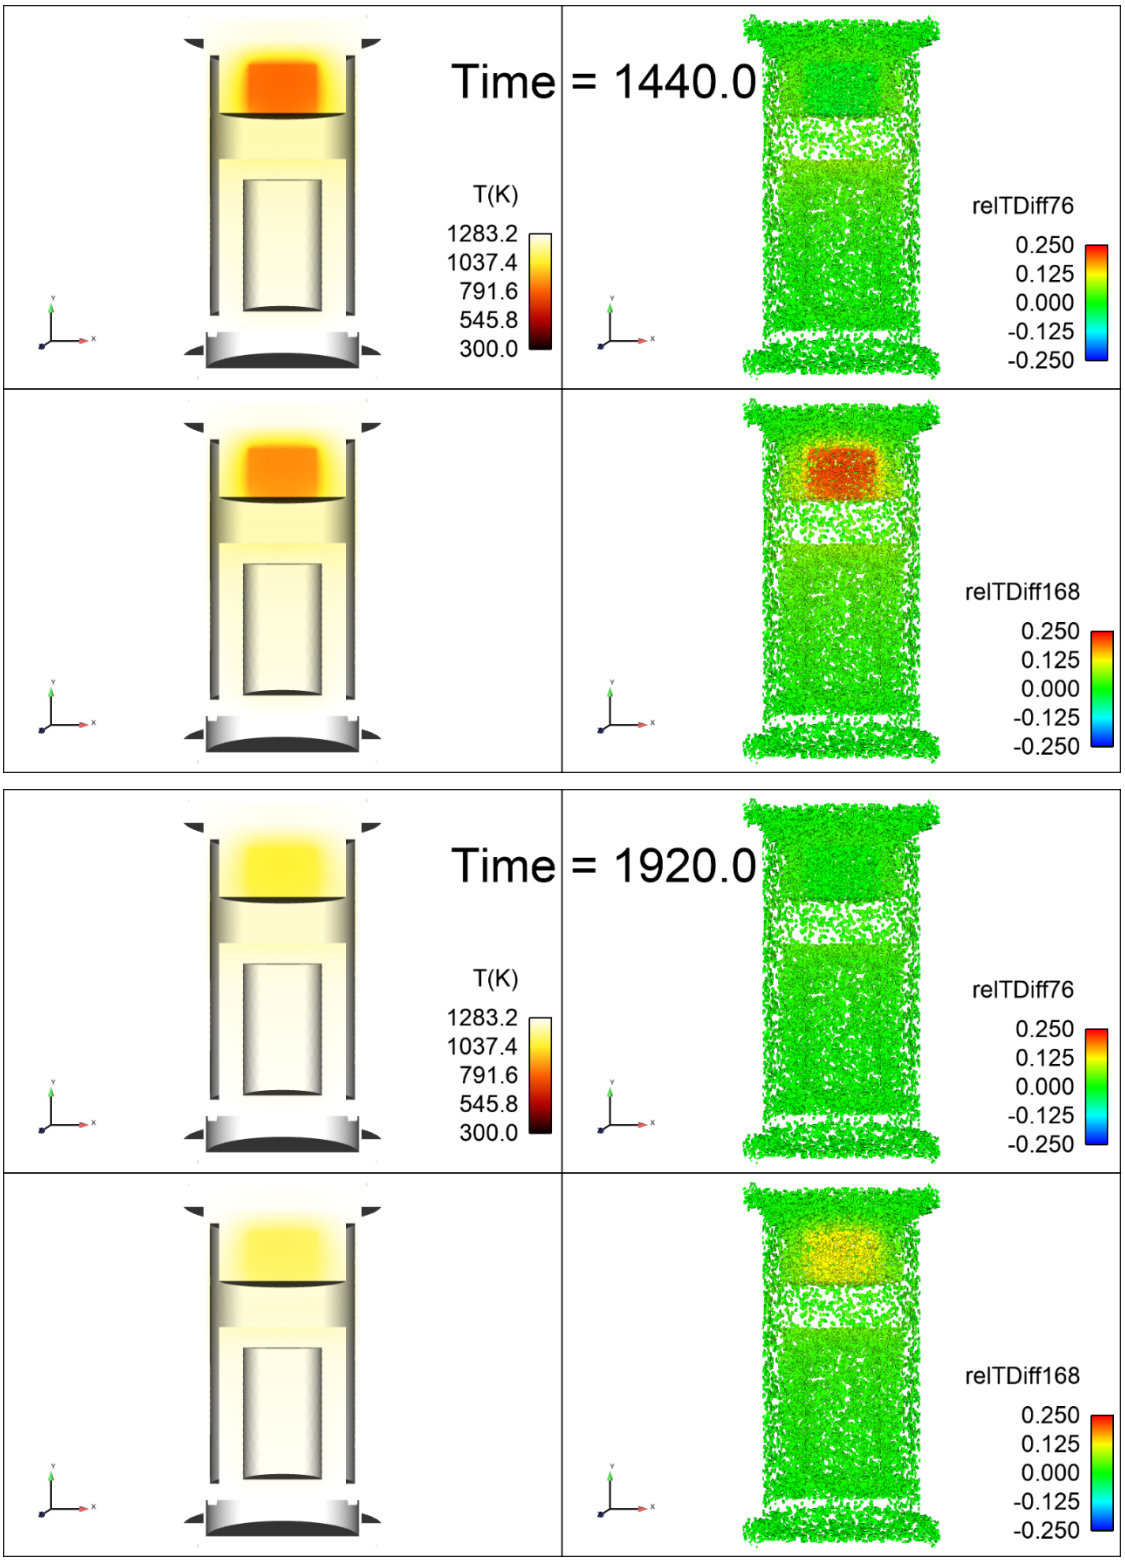
Time = 1440.0
T(K)
1283.2
1037.4
791.6
545.8
300.0
relTDiff76
0.250
0.125
0.000
-0.125
-0.250
relTDiff168
0.250
0.125
0.000
-0.125
-0.250
Time = 1920.0
T(K)
1283.2
1037.4
791.6
545.8
300.0
relTDiff76
0.250
0.125
0.000
-0.125
-0.250
relTDiff168
0.250
0.125
0.000
-0.125
-0.250

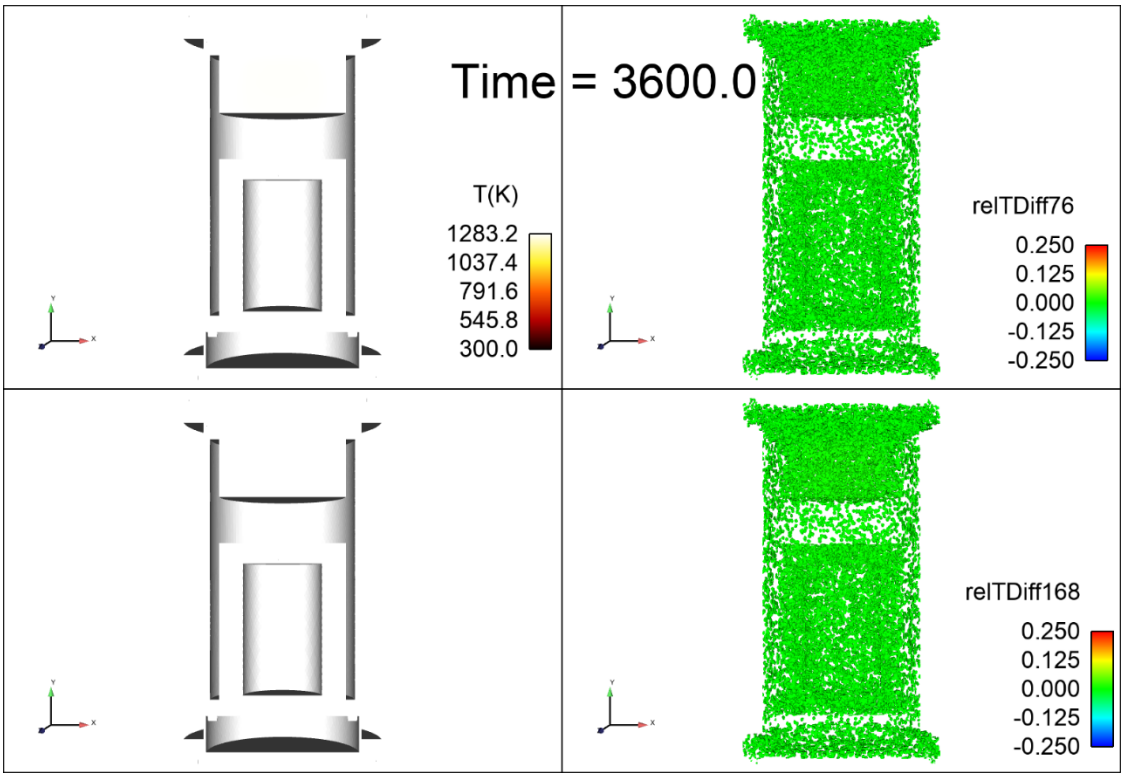

Figure 5-12. (left) ATE FOM temperatures shown in cross section at times throughout simulation for best (run 76) and worst (run 168) agreement with ROM at φ = 0.03. (right) Relative (fractional) differences between FOM and ROM temperatures shown on ROM sample mesh for these cases shown in cross section.

At 60 sec into the simulation, we observe the best case relative temperature differences actually exceed that of the worst case in the mass potting surrounding the mass mock and near the vertical center of the plastic component. The relative temperature difference grows for both best and worst cases up to 240 sec, with large variations seen in the regions just mentioned and in the case that surrounds the plastic component. By 480 sec, the best and worst cases have approximately equivalent relative temperature differences, with both cases showing decreasing values at later times. Relative to the best case, for times beyond 720 sec, the worst case has larger relative temperature differences, though for both at the end of the simulation (1 hr = 3600 sec), relative temperature differences are not visible on this scale.

We note an important finding here. While focusing on the QOI, that is the temperature at the center of the mass mocks, FOM and ROM simulations appear to produce similar results. Across the full simulation domain; however, significant temperature variations (greater than 25%) between FOM and ROM can be observed in specific regions at specific time intervals during the simulation, even at the relatively high mesh fraction of 0.03. This may be perfectly acceptable in the context of a UQ study focusing on this specific QOI, but some difficulty would be encountered if an analyst were interested in another location within the geometry for comparison.

6. SUMMARY

We have presented a comparison of the temperature response (QOI) at the center of a mass mock within a prototypical system (Crash-Burn model) based on full order model (FOM) and reduced order model (ROM) simulations using Sierra multimechanics module *Aria* and Sierra reduced order model module *Pressio_Aria*, applied to both NTE and ATE thermal environments. This model represents a mock weapons system reported on in a recent L2 milestone project at SNL. The current investigation has been performed in the context of a VV/UQ LHC approach using SNL's Dakota package. In the NTE portion of this study, both FOM and ROM utilized 96 simulation parameter variations, while the ATE portion explored 256 parameter variations. We have demonstrated the computational efficiency of the *Pressio_Aria* approach via FOM/ROM timings in all cases, showing both a significant speedup using *Pressio_Aria* and the ability to run ROM calculations on reduced computational resources. For NTE scenarios, we have demonstrated excellent agreement between FOM and ROM transient solutions at a ROM hyper-reduction mesh fraction of $\varphi = 0.03$. For ATE scenarios, agreement between FOM and ROM solutions are potentially satisfactory, but are dependent, to some degree, on the hyper-reduction mesh fraction. Likely, the NTE scenario would experience a similar dependency had hyper-reduction mesh fraction been varied, though it was not for this study. For ATE we have employed mesh fractions of $\varphi = 0.003$ and 0.03. In all cases, we have discussed not only best- and worst-case agreements between FOM and ROM for a particular quantity of interest (QOI), but also the distribution of transient temperature differences between FOM and ROM QOI solutions across the entire LHC set. We have also explored the broader agreement/disagreement between FOM and ROM solutions across the entire computational domain as a function of time.

This study demonstrates the possible utility of ROM approaches for thermal analysis problems. It also discusses the caveats of such methods. Normal thermal environments seem to be ideally suited for application of ROM methods. Abnormal thermal environments prove to be more challenging, and we have discussed here the nature of such difficulties, including the larger variation in temperature-dependent material properties and different physics excited that are seen in higher-temperature scenarios relative to scenarios involving lower temperatures. Additional challenges include producing agreement between FOM and ROM solutions across the entire simulation domain rather than just at specific locations as well as improving ROM timings for chemistry calculations. Further exploration of the limitations and constraints of these ROM methods is needed. Additional studies should consider variation of environments, geometries, and material properties, including chemically reacting materials. The conclusions drawn by the current study are both subject to variation in system geometry and the specific thermal scenario including initial and boundary conditions. We do not propose that specific conclusions can be extrapolated without care to every thermal analysis scenario. Rather, this prototypical scenario provides insight to those who wish to apply ROM techniques such as these to thermal analysis.

7. APPENDIX A: COARSE AND REFINED MESH ELEMENT COUNTS

Table A-1. Crash-Burn refined and coarse mesh element blocks by name with element count

Block Name	# Elements (refined)	# Elements (coarse)
case_outer	1170176	146272
epoxy	352776	44097
composite	339448	42431
case_lid1	243096	30387
case_lid2	248560	31070
mass_case_lid	348544	43568
mass_case	155032	19379
mass	189384	23673
mass_potting	770240	96280
plastic_case_lid	65304	8163
plastic_case	264056	33007
plastic	1458240	182280
scr1_s	2096	262
scr2_s	1480	185
scr3_s	1744	218
scr4_s	1856	232
scr5_s	1736	217
scr6_s	1400	175
scr7_s	1752	219
scr8_s	1784	223
scr9_s	3168	396
scr10_s	3608	451
scr11_s	3176	397
scr12_s	3808	476
scr13_s	3440	430
scr14_s	4048	506
scr15_s	2744	343
scr16_s	3552	444
scr17_s	1632	204
scr18_s	1424	178

Block Name	# Elements (refined)	# Elements (coarse)
scr19_s	2024	253
scr20_s	1528	191
scr21_s	1584	198
scr22_s	1496	187
scr23_s	1848	231
scr24_s	1664	208
scr25_s	408	51
scr26_s	952	119
scr27_s	1240	155
scr28_s	888	111
scr29_s	360	45
scr30_s	1040	130
scr31_s	1088	136
scr32_s	1200	150
screw_heads	51336	6417
screw_tails	69184	8648

Table A-2. Crash-Burn mesh sidesets with element and node count for the refined (ATE) mesh

Surface Name	# elements	# nodes
case_plastic_case_enc	32208	96624
plastic_enc	23376	70128
internal_void_enc	116104	348312
full_exterior	136720	410160

8. APPENDIX B: BEST AND WORST PARAMETER SETS FOR NTE AND ATE ROM SIMULATIONS

Table B-1. NTE Dakota parameter variation for hyper-reduction mesh fraction 0.03 (best #84, worst #40) agreement with FOM QOI for LHC study.

NTE: Parameter	Best (#84)	Worst (#40)
f_al_k	0.952136689	0.976584066
f_al_Cp	1.005155348	1.057632867
f_al_emis	1.024052861	1.009842347
f_ss304_k	1.028447484	0.9960958
f_ss304_Cp	1.136700033	1.128113881
f_ss304_emis	0.942631438	1.147335029
f_epoxy_k	0.926083427	1.075311507
f_epoxy_Cp	1.070196726	1.022855536
f_epoxy_emis	1.184964779	1.239550092
f_PMDI_k_bulk	1.090149606	0.671271722
f_PMDI_k_rad	1.175826648	1.041011056
f_PMDI_Cp	0.887188629	1.048658653
f_PMDI_emis	1.127971195	0.939927678
f_composite_k	1.241456905	0.997765725
f_composite_Cp	0.949106519	1.088701887
f_composite_emis	1.040388758	1.077065635

Table B-2. ATE Dakota parameter variation for ROM hyper-reduction mesh fraction 0.003 (best: #158, worst: #173) and 0.03 (best: #76, worst:#168) agreement with FOM QOI for LHC study.

ATE: Parameter	0.003: Best (#158)	0.003: Worst (#173)	0.03: Best (#76)	0.03: Worst (#168)
f_al_k	0.931993225	0.938388943	0.94424329	0.990513713
f_al_Cp	1.014626051	0.960198991	0.99951424	0.961025283
f_al_emis	1.022711028	1.105410435	1.036128132	0.968413939
f_ss304_k	0.992372014	1.001293132	0.983223243	1.024497983
f_ss304_Cp	0.905839458	1.088478125	1.160280892	0.914975073
f_ss304_emis	1.106897266	0.913985362	0.960798748	1.021652728

ATE: Parameter	0.003: Best (#158)	0.003: Worst (#173)	0.03: Best (#76)	0.03: Worst (#168)
f_epoxy_k	0.967033734	1.145715769	0.882207392	0.906247823
f_epoxy_Cp	1.071351756	1.219560624	1.099441442	0.970486377
f_epoxy_emis	1.053079747	0.835881453	1.205358139	1.006466752
f_PMDI_k_bulk	1.182018077	1.22561048	1.169955817	0.956917058
f_PMDI_k_rad	1.118600112	0.873271378	1.113192486	0.752632707
f_PMDI_Cp	1.124790304	1.021998002	1.018424508	1.082041312
f_PMDI_emis	0.875946957	1.131841864	1.04170172	1.030766588
f_composite_k	1.123700221	0.879332536	0.96995623	0.764985389
f_composite_Cp	0.879048413	1.012487634	1.11808767	1.172232986
f_composite_emis	1.121871891	0.866048114	1.08863111	0.98494731

REFERENCES

[Al6061-Conductivity] R. P. Tye, R. W. Hayden, and S. C. Spinney,"Thermal Conductivity of Selected Alloys at Low Temperatures,"in Advances in Cryogenic Engineering, K. D. Timmerhaus et al. (eds.), Springer Science, New York pp. 136-144, 1977.

[Al6061-Emissivity] Constant: Measured by Walt Gill, T-dependent: Dean Dobranich memo, Sandia National Laboratories, 5/21/2004

[Aria 2022] SIERRA Thermal/Fluid Development Team, 202. SIERRA Multimechanics Module: Aria User Manual – Version 4.40 (No. SAND2022-0256). Sandia National Lab. (SNL-NM), Albuquerque, NM (United States).

[Arienti 2021] Arienti, M., Blonigan, P.J., Rizzi, F., Tencer, J. and Howard, M., 2021. Projection-based model reduction for finite-element simulations of thermal protection systems. In AIAA Scitech 2021 Forum (p. 1717).

[AWE 2018] AWE materials database https://www.grantadesign.com/wp-content/uploads/2018/12/AWE-final.pdf

[Blonigan 2019a] Blonigan, P., Shimizu, Y., Carlberg, K., Hoang, C. and Rizzi, F., 2019. Pressio. Sandia National Lab. (SNL-NM), Albuquerque, NM (United States).

[Blonigan 2019b] Blonigan, P., Shimizu, Y., Carlberg, K., Hoang, C. and Rizzi, F., 2019. Pressio v. 1.0. Sandia National Lab. (SNL-NM), Albuquerque, NM (United States).

[Blonigan 2020] Blonigan, P.J., 2020. Projection-based ROMs at SNL: How? (No. SAND2020-8693PE). Sandia National Lab. (SNL-CA), Livermore, CA (United States).

[Blonigan 2021a] Blonigan, P.J., Rizzi, F.N., Parish, E.J. and Tencer, J., 2021. Pressio: A Computational Framework Enabling Projection-Based Model Reduction for Large-Scale Nonlinear Dynamical Systems (No. SAND2021-2283C). Sandia National Lab. (SNL-CA), Livermore, CA (United States); Sandia National Lab.(SNL-NM), Albuquerque, NM (United States)

[Blonigan 2021b] Blonigan, P.J., Rizzi, F., Howard, M., Fike, J.A. and Carlberg, K.T., 2021. Model reduction for steady hypersonic aerodynamics via conservative manifold least-squares Petrov–Galerkin projection. AIAA Journal, 59(4), pp.1296-1312.

[Blonigan 2021c] Blonigan, P.J., 2021. Projection-based ROMs at Sandia National Laboratories (No. SAND2021-1776PE). Sandia National Lab. (SNL-CA), Livermore, CA (United States).

[Brunini 2022] Brunini, V., Parish, E.J., Tencer, J. and Rizzi, F., 2022. Projection-Based Model Reduction for Coupled Conduction—Enclosure Radiation Systems. Journal of Heat Transfer, 144(6), p.062101.

[Ching 2020] Ching, D., Blonigan, P.J., Rizzi, F., Howard, M. and Fike, J.A., 2020. Quantifying the impact of the initial guess for projection-based model reduction of steady hypersonic aerodynamics (No. SAND2020-5692C). Sandia National Lab. (SNL-CA), Livermore, CA (United States); Sandia National Lab. (SNL-NM), Albuquerque, NM (United States).

[Dakota 2020] Adams, B., Bohnhoff, W., Dalbey, K., Ebeida, M., Eddy, J., Eldred, M., Hooper, R., Hough, P., Hu, K., Jakeman, J. and Khalil, M., 2020. Dakota, A Multilevel Parallel Object-Oriented Framework for Design Optimization, Parameter Estimation, Uncertainty Quantification, and Sensitivity Analysis: Version 6.13 User's Manual (No. SAND2020-12495). Sandia National Lab. (SNL-NM), Albuquerque, NM (United States).

[Parish 2020] Parish, E.J. and Rizzi, F., 2020. Windowed least-squares model reduction in Pressio (No. SAND2020-8806PE). Sandia National Lab. (SNL-CA), Livermore, CA (United States).

[Parish 2021] Reduced-order modeling for the net radiation method Eric Parish Sandia National Laboratories, Livermore, CA Unpublished

[Phinney 2015] Phinney, L., Thermal Conductivity and Specific Heat Recommendations with Uncertainties for Selected Materials for B61 Thermal Analysis, SAND2015-4242, Printed May 2015

[Rizzi 2020] Rizzi, F., Blonigan, P.J., Parish, E.J. and Carlberg, K.T., 2020. Pressio: Enabling projection-based model reduction for large-scale nonlinear dynamical systems. arXiv preprint arXiv:2003.07798.

[Scott 2014] Scott, S.N., Dodd, A.B., Larsen, M.E., Suo-Anttila, J.M. and Erickson, K.L., 2014. Validation of Heat Transfer Thermal Decomposition and Container Pressurization of Polyurethane Foam (No. SAND2014-17867). Sandia National Lab.(SNL-CA), Livermore, CA (United States); Sandia National Lab.(SNL-NM), Albuquerque, NM (United States).

[Shurtz 2018] Shurtz, R., 2018. Total Hemispherical Emissivity of Metals Applicable to Radiant Heat Testing (No. SAND2018-13271). Sandia National Lab. (SNL-NM), Albuquerque, NM (United States).

[Tan-Torres 2020] Tan-Torres, S., Beghini, L.L., Dillon, C.R., Francis, N.D., Hanson, A.A. and Murphy, A.W., 2020. Uncertainty Propagation in a Combined Crash then Burn Scenario (No. SAND2020-13215C). Sandia National Lab. (SNL-CA), Livermore, CA (United States); Sandia National Lab. (SNL-NM), Albuquerque, NM (United States).

[Tencer 2021] https://wiki.sandia.gov/display/1514KB/Reduced-Order+Modeling+with+ARIA

DISTRIBUTION

Email—Internal

Name	Org.	Sandia Email Address
Wyatt Hodges	1514	whodges@sandia.gov
Roy Hogan	1514	rehogan@sandia.gov
Leslie Phinney	1514	lmphinn@sandia.gov
Flint Pierce	1514	fpierce@sandia.gov
John Tencer	1514	jtencer@sandia.gov
Patrick Blonigan	8739	pblonig@sandia.gov
Francesco Rizzi	8739	fnrizzi@sandia.gov
Amanda Dodd	8750	ajbarra@sandia.gov
Victor Brunini	8751	vebruni@sandia.gov
Technical Library	01911	sanddocs@sandia.gov

Sandia National Laboratories

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia LLC, a wholly owned subsidiary of Honeywell International Inc. for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.