LA-UR- 08-6599

Approved for public release;
distribution is unlimited.

Title:	Verification and Validation for Induction Heating
Author(s):	Kin Lam, LANL Trevor B. Tippetts, LANL David W. Allen, Knowledge Reef Systems, Inc.
Intended for:	Proceedings for International Modal Analysis Conference XXVII, February 9-12, 2009 Orlando, Florida

Los Alamos
NATIONAL LABORATORY
EST. 1943

Form 836 (7/06)

Verification and Validation for Induction Heating

Kin Lam, Trevor B. Tippetts
Los Alamos National Laboratory
Los Alamos, NM 87545
klam@lanl.gov, tippetts@lanl.gov

David W. Allen
Knowledge Reef Systems, Inc
dave@knowledgereefsystems.com

Abstract

Truchas is a software package being developed at LANL within the Telluride project for predicting the complex physical processes in metal alloy casting. The software was designed to be massively parallel, multi-material, multi-physics, and to run on 3D, fully unstructured meshes. This work describes a Verification and Validation assessment of Truchas for simulating the induction heating phase of a casting process. We used existing data from a simple experiment involving the induction heating of a graphite cylinder, as graphite is a common material used for mold assemblies. Because we do not have complete knowledge of all the conditions and properties in this experiment (as is the case in many other experiments), we performed a parameter sensitivity study, modeled the uncertainties of the most sensitive parameters, and quantified how these uncertainties propagate to the Truchas output response. A verification analysis produced estimates of the numerical error of the Truchas solution to our computational model. The outputs from Truchas runs with randomly sampled parameter values were used for the validation study.

1 Introduction

Truchas is a multi-physics software package being developed at LANL under the sponsorship of Department of Energy (DOE)'s Advanced Simulation and Computing (ASC) Program to deliver a computational capability for predicting the complex physical processes in key DOE manufacturing operations. An important application of Truchas is the simulation of metal alloy casting at LANL foundries, which typically involves electromagnetic induction heating of the mold assembly as well as melting the metal charge in the crucible. The thermal environment within the casting furnace determines the heating and cooling behaviors of both the part and mold, which in turn control defect development, microstructural evolution, internal stress build-up, distortion and structural integrity. Therefore, accurate and credible prediction of the thermal environment during induction heating is important for optimum model and process designs as well as for assurance of part quality and performance.

In this work, we performed a verification and validation (V&V) assessment of Truchas for simulating the induction heating phase of a casting process. (Here we follow the definition of V&V provided in the ASC Program Plan [3] which states that *verification* is the process of confirming that a computer code correctly implements the algorithms that were intended, and *validation* is the process of confirming that the predictions of a code adequately represent measured physical phenomena.) Virtually all casting processes in LANL's foundries are carried out in evacuated furnaces and use electromagnetic induction to heat the metal charge and mold assembly. The thermal environment within the casting furnace determines the heating and cooling behaviors of both the part and mold, which in turn control defect (e.g., porosity) development, microstructural evolution, internal stress build-up, distortion and structural integrity. Therefore, accurate and credible prediction of the thermal environment during the preheat stage is important for optimum mold and process designs as well as for assurance of part quality and performance.

For the V&V study, we used existing data from a simple experiment involving the induction heating of a graphite cylinder, as graphite is the most common material used for mold assemblies in LANL foundries. Because we do not have complete knowledge of all the conditions and properties in this experiment (as is the case in many other experiments), we performed a parameter sensitivity study, modeled the uncertainties of the most sensitive parameters, and quantified how these uncertainties propagate to the Truchas output response. The outputs from Truchas runs with randomly sampled parameter values were used for the validation study.

2 Methods in Truchas for Induction Heating

All the physical models and solution methods used in Truchas are described in the Physics and Algorithms Manual [4]. A User's Manual [5] and a Reference Manual [6] also available with document in detail how to use the code for various problems. Here we give a brief summary of the physical models and their solution algorithms in Truchas that are relevant to induction heating in casting simulation problems. These methods are grouped into the two major physical phenomena of electromagnetics (EM) and heat transfer (HT).

Electromagnetics The electromagnetics modeling in Truchas is targeted toward induction heating problems, and is based on a direct solution of the Maxwells equations in the time domain. Specifically, the equations are the Faradays law and the Amperes law, plus three constitutive equations involving material properties .

The boundary condition is in the form of a tangential component of the magnetic field on the boundary. However, the user does not specify this boundary condition directly. The boundary magnetic field is assumed to be driven solely by the induction coil, whose propertiesposition, radius, length, current, number of turns, and frequencyare defined via code input.

The Faraday's and Ampere's equations in the primary dependent variables, the magnetic flux density ($\vec{B}$) and the electric flux intensity ($\vec{E}$) are solved on a tetrahedral mesh using a mimetic spatial discretization procedure based on the Whitney family of finite element spaces [7]. The time derivatives are discretized using the trapezoid rule. This implicit treatment in time allows use of a time step size of the order of the temporal variation of the magnetic driving field, which is much larger than the explicit stability limit. The final time and space discretized equations can be combined into a single matrix equation (in the discretized unknown $\vec{E}$ field), which is symmetric, positive definite, and is solved by the conjugate gradient method using a relaxation scheme adapted from Hiptmair [8] as a preconditioner. After the electromagnetic solution is obtained, the joule power density is calculated as the product of the electrical conductivity and the square of the magnitude of $\vec{E}$ over a period, which is assumed to be much shorter than the time scale of heat transfer.

Heat Transfer The main governing equation for heat transfer is the conservation of energy, which is written in the enthalpy formulation, i.e., the primary dependent variable to be solved for is the material mixture enthalpy. Temperature is a derived solution quantity, which is determined from the total enthalpy and material mass fractions. The rate of change of the mixture enthalpy depends on the balance between advection, conduction, release/absorption of latent heat due to phase change, and volumetric source terms such as the electromagnetic joule power density.

For boundary conditions, both the Dirichlet and Neumann (flux) types are supported. A heat transfer coefficient can be specified at an external boundary to represent convective cooling to a bulk temperature, or at an internal material interface to model added resistance due to a gap or coating that is too thin to be represented by the computational grid. Radiation at an external surface can be modeled as a simple exchange with an ambient temperature (which can be time dependent) using the Stefan-Boltzmann law. Truchas also provides an enclosure radiation modeling capability (based on the net radiation method), in which selected material surfaces can be defined to form a single or multiple enclosures, which may be open or closed. Calculation of the view factors is based on the hemi-cube method as implemented in the Chaparral library developed at the Sandia National Laboratory [10].

The finite-volume discretized enthalpy equation is solved with an implicit treatment of the diffusion (conduction) and phase change terms, while the advection and other source terms are treated explicitly in time. The set of nonlinear discrete equations are solved with a method that achieves Newton-like super-linear convergence in the global nonlinear iteration, without the complexity of forming or inverting the Jacobian matrix.

3 Graphite Cylinder Heating Experiment

A series of graphite cylinder heating experiments were carried out at LANL's Sigma foundry. The purpose of these experiments was to provide data on the temperature response of graphite when it is heated inductively by several different electric coils at various power levels. Table 1 lists three of the experiments chosen for our present validation study, indicating the nominal power level, coil type, and the current sensor used.

In the experiments, an 8-in (20.32-cm) diameter, 24-in (60.96-cm) long cylinder of graphite was placed concentrically inside either a 12-in or 16-in diameter coil. Experiments 04K-450 and 04K-451 used the smaller coil while Experiment 04K-454 used the larger coil. The placement of the coil was adjusted so that its middle lined up with that of the cylinder. The 12-in coil has four turns, and an inner diameter of 28 cm and an outer diameter of 33 cm. The 16-in coil has three turns, and an inner diameter of 38 cm and an outer diameter of 43 cm. Both coils are 15

Table 1: Series of graphite cylinder heating experiments in K Furnace of Sigma Foundry. Experiments 04K-450, 04K-451, and 04K-454 are chosen for this validation study.

Experiment	Current Sensor	Nominal Coil Diameter	Nominal Power(kW)	Comments
04K-450	Person 1423	12"	20	20kW for 60 minutes
04K-451	Person 1423	12"	30	30kW for 60 minutes
04K-454	Person 1423	16"	20	20kW for 60 minutes

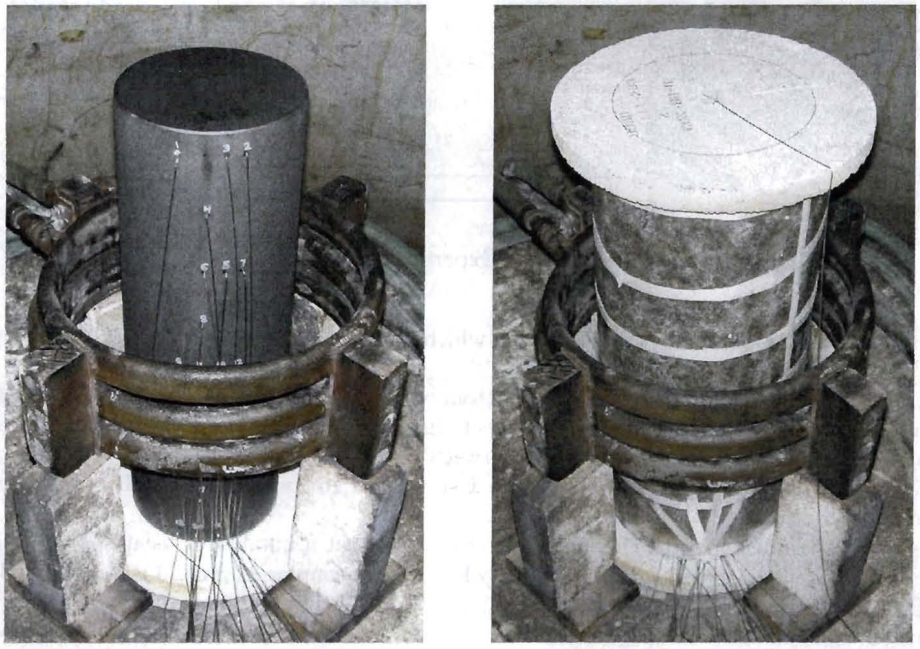

Figure 1: Setup for the 16-in coil experiment, showing the bare graphite cylinder instrumented with thermocouples prior to installation of insulation (left) and after the mica insulation and alumina top plate were installed.

Table 2: Summary of experimental conditions for 04K-450, 04K-451, and 04K-454 in terms of parameters required to define the Truchas validation problem.

	Coil		Average Current		Initial
Experiment	Radius (m)	No. Turns	Frequency (Hz)	Amplitude (A)	Temperature (K)
04K-450	0.1524	4	1885	1645	290
04K-451	0.1524	4	1964	2030	288
04K-454	0.2032	3	1975	2022	315

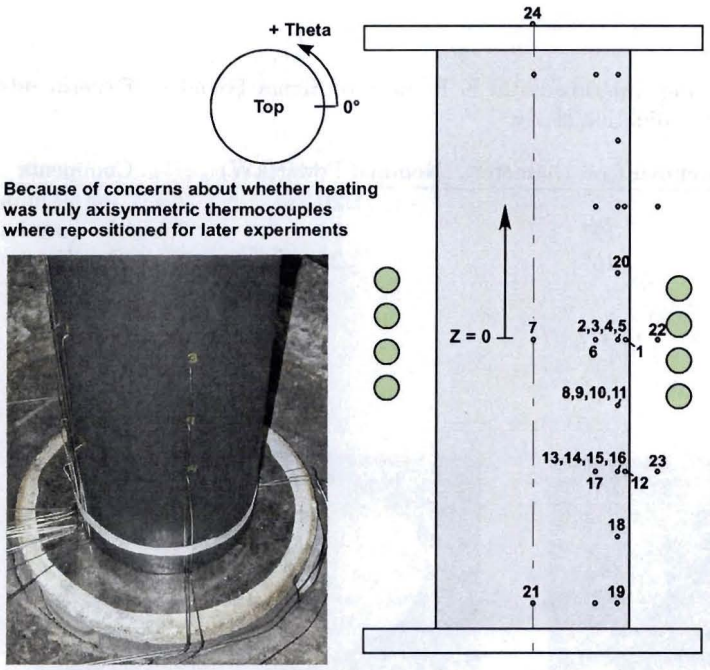

all dimensions in cm

TC	Depth	Z	Theta
1	0.47	0.0	23°
2	1.28	0.0	0°
3	1.26	0.0	90°
4	1.28	0.0	180°
5	1.28	0.0	270°
6	3.19	0.0	13°
7	10.16	0.0	346°
8	1.22	-7.0	0°
9	1.25	-7.0	90°
10	1.28	-7.0	180°
11	1.26	-7.0	270°
12	0.45	-14.0	23°
13	1.29	-14.0	0°
14	1.23	-14.0	90°
15	1.28	-14.0	180°
16	1.29	-14.0	270°
17	3.12	-14.0	13°
18	1.24	-21.0	0°
19	1.32	-28.0	0°
20	1.22	7.0	0°
21	10.16	-28.0	-13°
22	on insulation	0.0	
23		-14.0	

Figure 2: Thermocouple positions for Experiments 04K-450, 04K-451, and 04K-454.

cm long, and are made of 1-in diameter copper tubing which allowed cooling water to flow through while they were operating.

The cylinder was insulated on both the top and bottom with a 2.54-cm thick disk of firebrick (alumina). A 1-cm thick sleeve made of graphite felt wrapped in mica sheets insulated the cylinder radially. The mica insulation had an inner radius of 12.16 cm so there was a 2-cm gap between it and the graphite cylinder. Therefore, the major heat loss by the cylinder radially was via radiation. Figure 1 shows the 16-coil experimental setup before and after the insulating materials were installed.

The cylinder was instrumented with thermocouples extensively at a number of locations, as shown in Figure 2. Thermocouple 24 measured the firebrick temperature, while thermocouples 22 and 23 were intended to measure the temperature of the mica insulation. Several thermocouples were placed at the same axial and radial positions but at different azimuthal angles to confirm that the heating and temperature response were axi-symmetric.

In all cases, the graphite cylinder was heated with a uniform supply of power for an hour, and was then allowed to cool for another hour before the experiment ended. Table 2 lists the different experimental conditions for the three experiments used in this validation study. Averaged values of the frequency and current amplitude are given; these are used in Truchas for calculating the induction heating power.

4 Truchas Computational Model

General Problem Description In each experiment, a graphite cylinder was heated inductively by a concentric coil of current surrounding it. The current is alternating and is assumed to be sinusoidal, defined by an amplitude and frequency. The magnetic field created by the sinusoidal current couples with the electrically conductive graphite and induces an alternating current within the cylinder in the azimuthal direction. These eddy currents heat up the graphite via the joule power, which results from the material's electrical resistivity. The heat is then transported throughout the cylinder via conduction.

The cylinder loses heat to the firebrick plate it is in contact with at each end. It also loses heat via radiation across a vacuum annular gap to the mica sleeve. The external surfaces of the firebrick and mica sleeve then radiate to an ambient temperature, which serves as the ultimate heat sink. For heat transfer modeling, all structures beyond the firebrick and mica insulation are lumped into the ambient environment, which is assumed to be under a constant temperature.

The steady induction heating lasts for 3600 s, at which time the driving current is turned off instantly. The

problem then continues for another hour. Therefore, the entire duration of the problem is 7200 s, with the first hour being the heat-up phase, and the second hour being the cool-down phase.

In this Section, we discuss various features of the computational model. Only nominal model parameter values are given here. The effects of mesh and time resolutions adopted in the numerical model will be discussed in Section 5. The sensitivity and uncertainty of the physical model parameters will be discussed in Section 6.

Computational Domains The EM solver in Truchas requires a tetrahedral mesh. However such a mesh is not suitable for other physics within Truchas because many of the algorithms perform poorly on this type of mesh. A hexahedral mesh is recommended for all other physics packages such as heat transfer. Therefore, two computational meshes are necessary for this problem, a tet mesh for EM and a hex mesh for HT.

For the EM calculation, the material to be modeled is the graphite cylinder, which is the only material that couples with the EM field. The induction coil is not modeled explicitly. Instead, the driving current and associated magnetic field are modeled via boundary conditions. The other materials, such as the firebrick and mica sleeve, are ignored and treated as part of the free space surrounding the graphite cylinder. The amount of free space to include in the EM domain is problem dependent. Ideally the domain should be large enough so that the EM fields imposed as boundary conditions are the same the undisturbed fields that result from the driving current in the coil. Beyond the cylinder's ends, there is no restriction on the possible domain size, while radially the domain is limited by the radius of the coil. In our model, the free space extends beyond the end of the cylinder a distance that is the same as the radius of the cylinder. The radial extent of the free space is 70% of the available distance between the radius of the cylinder and that of the coil. This choice of the EM domain size is judged to be optimum as extending the free space radially to 80% of the way to the coil radius results in a joule power density that is about 0.5% higher in some preliminary test calculations using the conditions of Experiment 04K-450. For the same test problem, extending the free space beyond the cylinder's end to a distance twice the radius of the cylinder would result in a joule power density that is only 0.02% higher.

The EM domain as discretized by a tet mesh of relatively coarse resolution is shown in Figure 3. Only 1/8 of the entire geometry is included because this problem is assumed to be axisymmetric and symmetric about the mid-plane of the cylinder and coil. The z-axis representing the vertical is the axis of symmetry, while the x- and y-axes are the horizontal directions.

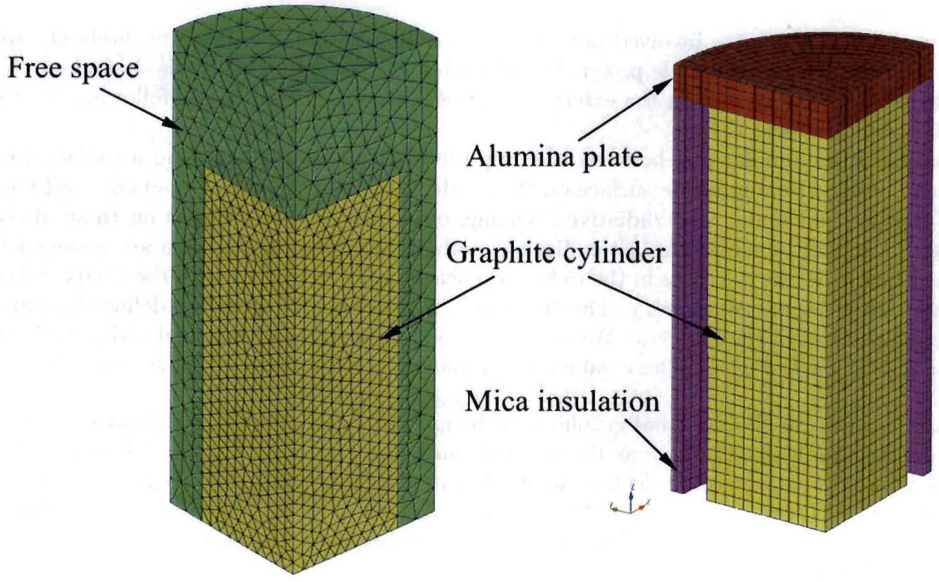

Figure 3: EM and HT domains as discretized by two coarse meshes: a 27,326-cell tetrahedral mesh (average cell size 1.0 cm, left) for the EM domain, and a 7,800-cell hexahedral mesh (average cell size 0.8 cm, right) for the HT domain. The domains are 1/8 of the total geometry, with the boundary defined by each of the three axes being a plane of symmetry.

Figure 3 also show the HT domain as meshed by a coarse hex grid. Again, only 1/8 of the real geometry is included because of symmetry. The materials of interest in the HT domain are the graphite cylinder, the firebrick end plate, and the mica sleeve. The physical presence of the coil is not considered for heat transfer because it is

water cooled. The coil is treated as part of the ambient environment which acts as the ultimate heat sink to which the mica and firebrick external surfaces radiate. Unlike for EM, vacuum or free space does not have to be included in the HT mesh.

Other EM and HT meshes, with different resolutions from those shown in Figure 3, have also been used. The effects of mesh resolution will be discussed in Section 5.

Physical Models and Boundary Conditions The problem begins with EM computations to obtain the joule power density. It is assumed that the time scale of the alternating EM fields is much shorter than the time scale of HT. This allows us to solve the Maxwell's equations to a periodic steady state while assuming other physics remain unchanged. The rapid temporal variation in joule heat is then averaged over a cycle to give a steady heating rate. This volumetric energy rate within the graphite is then interpolated from the EM tet mesh onto the HT hex mesh via a conservative mapping algorithm to drive the HT computations.

As mentioned previously, the EM domain consists of the graphite cylinder and free space. The coil is not modeled explicitly. Its effects are imposed as EM fields on the boundary of the domain. The physical coil is defined by the length, radius, center position, and number of turns in the idealized stack of current loops. The EM fields due to these loops carrying a sinusoidal current, defined by the specified amplitude and frequency, are computed and imposed on the boundary of the domain. Truchas then solves the Maxwell's equations for the EM fields within the graphite, from which solution the joule power density is derived and passed to the heat transfer calculations as a volumetric heat source.

The relevant EM model parameters are the coil radius, length, number of turns, and the amplitude and frequency of the current. The coil length is 0.13335 m and values for the other parameters as used in the three experiments are taken from Table 2.

As seen in Figure 3, the heat transfer model consists of three materials, the graphite cylinder, firebrick plate, and mica sleeve. The joule heat generated by EM induction is received by the HT package as a volumetric energy source. The heat produced within the graphite diffuses throughout the cylinder. This thermal energy is transported within the cylinder and leaves at the top (and bottom) via contact with the firebrick plate and on the side via radiation to the mica sleeve across the vacuum gap. In both the firebrick and mica sleeve, heat is conducted from the inside surfaces to the outside boundaries, where it is rejected to the ambient environment via radiation as the entire furnace is evacuated.

In summary, the major physics involved here heat conduction within the graphite, firebrick, and mica sleeve. The heat flow is determined by the joule power density within the graphite and by the boundary conditions at the interfaces between the materials and on the external surfaces of the materials. In the following, we will discuss each of these boundary conditions.

The boundaries of the annular gap between the graphite and the mica sleeve form a radiation enclosure. This enclosure is then bounded by a graphite surface on the inside, a mica surface on the outside, and a narrow firebrick surface at the top (and bottom). The radiative exchange of energy at the cell faces on these discretized surfaces is calculated with a model based on the net radiation method, in which the surfaces are assumed to be grey and diffuse. The view factors of the cell faces in the radiation enclosure are computed with the Chaparral code, which has been implemented into Truchas as a library. The only physical parameters required to define this radiation enclosure are the surface emissivities of the surfaces. We specify 0.8 as the nominal surface emissivity value for the graphite and firebrick, and 0.75 for the mica. These values are estimates based on consulting several internet and textbook resources [13][14].

The interface between the firebrick and graphite can be ignored if both materials are smooth and perfect contact is assumed. However this is not the case as the firebrick surface is rather rough. Therefore a contact resistance is added, which is modeled as an internal heat transfer coefficient (HTC) between the two materials across the interface. We specify a value of 10 W/(m^2-K) for this HTC, which is equivalent to the conductance across a 1-mm material with a conductivity of 0.01 W/(m-K). Likewise the small contact surface between the firebrick and mica is also modeled with an HTC of 10 W/(m^2-K).

Both the external surfaces of the firebrick and mica are modeled with an ambient radiation boundary condition. Two physical parameters required for this type of boundary condition, namely the surface emissivity and the reference or ambient temperature. For these external surfaces radiating to the ambient, we specify the effective emissivity to be 0.5, a value lower than the emissivities defined for the radiation enclosure. This is an approximation to account for the fact that the external surfaces do not radiate completely to the cold ambient and the fact that the ambient is not purely a black body. The ambient environment is assumed to remain at 300 K.

Material Properties Three materials are present in this problem, graphite, firebrick, and mica sleeve. The graphite participates in both the EM and HT calculations, while the firebrick and mica sleeve are involved only in

HT physics.

Material properties for graphite (for 2020 Stackpole graphite by Carbone of America, Inc.) are based on manufacturer data sheets. The density, $\rho_{graphite}$, is 1750 kg/m^3. The thermal conductivity, $k_{graphite}$, in W/(m-K), has been fitted with a third-order polynomial in temperature (T in K):

$$k_{graphite} = 240.6 + 2.5017\,T - 0.001123\,T^2 + 1.7 \times 10^{-7}\,T^3 \tag{1}$$

Similarly, the temperature-dependent specific heat capacity, $(C_p)_{graphite}$, in J/(kg-K), is modeled as

$$(C_p)_{graphite} = 112.35 - 0.13876\,T + 9.042 \times 10^{-5}\,T^2 - 2.0 \times 10^{-8}\,T^3 \tag{2}$$

The electrical conductivity, σ, in $(\Omega\text{-m})^{-1}$ can also be fitted with a polynomial in T:

$$\sigma = 2.3298 \times 10^4 + 136.247\,T - 0.09722\,T^2 + 2.139 \times 10^{-5}\,T^3 \tag{3}$$

Modeling the electrical conductivity as a function of temperature necessitates recomputing a new EM solution with the changed property as the temperature increases during the heat-up. As this adds significant computational time to the simulation, the approximation of a constant σ is usually assumed in practical applications. In this problem, we have also assumed that σ is constant, at a value of $7.0{\times}10^4$, which is the average of $5.6{\times}10^4$ at 300 K and $8.4{\times}10^4$ at 1100 K.

For the firebrick (alumina) properties, we use the values (all in SI units) given in Lienhard [14] and assume they are independent of temperature:

$$\rho_{firebrick} = 2000, \quad k_{firebrick} = 0.1, \quad (C_p)_{firebrick} = 1000 \tag{4}$$

For the mica sleeve, which is made of graphite felt wrapped in mica sheets (with imperfect contact and gaps between them), we assume the composite material has the following constant properties:

$$\rho_{mica} = 2000, \quad k_{mica} = 0.02, \quad (C_p)_{mica} = 1000 \tag{5}$$

Note that the above properties are used in nominal simulations runs. Discussion of the sensitivity and uncertainty of these material property parameters will be given in Section 6.

5 Verification

5.1 Introduction

When finite-volume discretizations are used to solve partial differential equations, discretization error is introduced into the computed solutions. This error may be a significant component of the uncertainty in simulation results. Therefore, discretization error quantification via calculation verification, also known as solution verification, is vitally important.

Many discretization error estimation methods use an ansatz, *i.e.*, an assumed form for the discretization error. The assumption is justified by Richardson extrapolation,[15, 16] based on a series expansion of the error in a grid scale parameter, Δx. Only the first term of the series is believed to be significant, and the truncated series is used to model the error. Extrapolation-based error estimators depend on the assumption that all of the grids used to solve the PDEs are in this asymptotic region.

Previous work by Kamm, Rider, and Brock[21] and Hemez[22] addressed the problem of combined space and time convergence with an ansatz that accounted for both space and time discretization. Their approach was applied to code verification analyses, using exact solutions to evaluate L_1 norms of the error. The method was demonstrated with a variety of problems, including both smooth and discontinuous solution fields. Section 5.2 describes a similar combined ansatz that was used in this work.

In this verification analysis, the absolute value of the error is assumed to converge exponentially in the cell size, as in some previous works.[20, 23]. Following Tippetts, Timmes, Brock, and Kamm [26], we obtain the solution of the ansatz parameters by solving an optimization problem. This approach allows more than the minimum number of three grids to fit a single ansatz. Analytical expressions for the ansatz parameters are given for the special case of solutions from three grids, for both monotonic and oscillatory convergence.

Table 3: The reference discretizations, Δt_r and Δx_r, normalize the ansatz discretizations. The values are equivalent to the discretizations used in the validation simulations.

	Simulation Time (s)	Δt_r (s)	Δx_r (cm)
Electromagnetics	Throughout heat up	1/37700	0.7
Heat Up	3600	10	0.5
Cool Down	7200	10	0.5

5.2 Combined Space and Time Ansatz

In all cases considered in Sections 5.3 and 5.4, the response quantities were observed to converge with Δt monotonically. This suggests the use of a combined space and time ansatz that allows for oscillatory convergence in space but assumes monotonic convergence in time.

$$\xi_i - \hat{\xi} = s_i A \left(\frac{\Delta x_i}{\Delta x_r} \right)^q + B \left(\frac{\Delta t_i}{\Delta t_r} \right)^p \tag{6}$$

Any product term between Δx and Δt in the ansatz is assumed to be small relative to the individual terms. This makes the ansatz separable, so that the space and time convergence behaviors may be analyzed individually. The same code used to fit single-term ansatzes can fit each of the separated equations,

$$\xi_i - \left(\hat{\xi}_{\Delta t} + s_i A \left(\frac{\Delta x_i}{\Delta x_r} \right)^q \right) = B \left(\frac{\Delta t_i}{\Delta t_r} \right)^p \tag{7}$$

and

$$\xi_i - \left(\hat{\xi}_{\Delta x} + B \left(\frac{\Delta t_i}{\Delta t_r} \right)^p \right) = s_i A \left(\frac{\Delta x_i}{\Delta x_r} \right)^q . \tag{8}$$

This produces two estimates of $\hat{\xi}$, $\hat{\xi}_{\Delta t}$ and $\hat{\xi}_h$, which are not necessarily equal due to fit error in both ansatzes. The magnitude of the difference between them, relative to the total error as expressed in Equation 6, can be an indication of the validity of the zero product term assumption.

Truchas solves an electromagnetics model to determine a heat source distribution. This simulation is carried out on a different mesh from the one used later for heat transfer. Therefore, verification analyses in Section 5.3 are used to verify convergence and estimate errors in the induction heat source. Section 5.4 covers verification for temperatures of probes 1, 7, 19, and 21 during the heat transfer simulation. Temperature snapshots at two times are analyzed: 3600 s and 7200 s. The first time was chosen because, as the end of the heat up period, it will see the greatest temperatures and likely some of the largest gradients in temperature. The second time is at the end of the simulation, so errors that accumulate over time will be at their maximum. In all cases, calculation verification is studied with respect to both space and time discretization. All phases of the simulations use constant, fixed time steps and unstructured meshes. The discretizations are normalized by the reference values, Δt_r and Δx_r, shown in Table 3. These values are equivalent to the discretizations used in the verification simulations.

5.3 Electromagnetics Simulation of Induction Heating

5.3.1 Time Discretization

The induction heating is simulated by an approximate solution of the electromagnetics PDEs on a tetrahedral mesh. Finite time steps are required to model the sinusoidally varying current amplitude that drives the heating. Convergence of the total induction heating power with respect to this discretization in time is apparently asymptotic over a very wide range. A calculation verification was performed by dividing the current period of 1/1885 s into 10, 20, 30, 40, 50, 60, 70, 80, 90, and 100 equal time steps.

The high quality of fit over a broad discretization range is clear from Figure 4. The fact that the observed exponent of 2.028 is close to the theoretical value of 2 also gives confidence in the validity of the error ansatz. The reference time step, Δt_r, is 1/37700 s; *i.e.*, 1/20 of the current period. This is the electromagnetics time step that was used for the verification and validation simulations to follow. The extrapolated value, $Q_{extrapolated}$, is 992.439 W. Using a reference time step makes the prefactor, A, equal to the difference between computed and extrapolated solutions at a Δt of Δt_r. In this case, the ansatz predicts a 7.135 W difference.

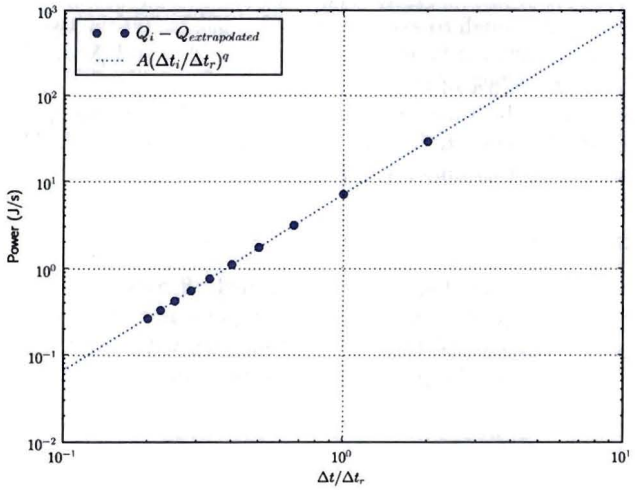

Figure 4: Convergence of induction heating rate with respect to the time step.

5.3.2 Space Discretization

Spatial convergence of the induction heating rate is more difficult to characterize. Unstructured meshes complicate the very definition of a discretization parameter. It is not surprising that the extent of the asymptotic region and the comparison between observed and theoretical convergence rates are both worse than for time convergence.

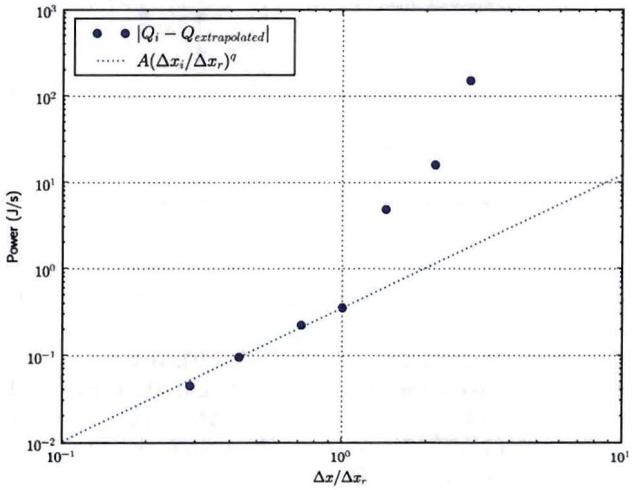

Figure 5: Convergence of the induction heating rate with respect to mesh refinement. Only the finest four meshes were used to fit the error ansatz.

The plot in Figure 5 makes it clear that meshes coarser than Δx_r are not in the same asymptotic region as those finer than Δx_r. An ansatz fit to the four finest meshes gives an exponent of 1.540, lower than the theoretical value of 2. The ansatz predicts a 0.3628 W error at the reference resolution of 0.7 cm, relative to an extrapolated power of 999.9218 W.

5.3.3 Combined Space and Time Error

The extrapolated values from the decoupled time and space verifications are now compared to obtain a combined extrapolation. Subtracting the spatial error from the time extrapolation gives a $\hat{\xi}_{\Delta t}$ of 992.802 W. Subtracting time error from the spatial extrapolation gives a $\hat{\xi}_{\Delta x}$ of 992.786 W. The 0.0156 W difference between these is small compared to both the time and space components of the error. This confirms the assumption that any cross term

with a product of Δt and Δx is small enough to exclude from the ansatz. Of the two extrapolated values, 992.786 W yields the greater difference with the 999.563 W power computed at Δt_r and Δx_r. The estimated induction heating power error is therefore 6.777 W, or 0.678% of the computed value.

The combined error is lower than the time discretization error because the spatial discretization error has an opposite sign. It is not always safe, however, to assume that the errors will cancel. The root-mean-square of the two components is a more conservative error estimate at 7.144 W.

5.4 Heat Transfer Simulation

The coil current remains constant until 3600 s, when it is turned off and the setup is allowed to cool until 7200 s. Separate verification analyses attempted to verify probe temperatures at these two times because of the jump discontinuity in the energy source. Truchas used a constant time step with implicit time integration over both time periods. As with the induction heat rate, the probe temperature errors are assumed to be separable in space and time.

5.4.1 Time Discretization

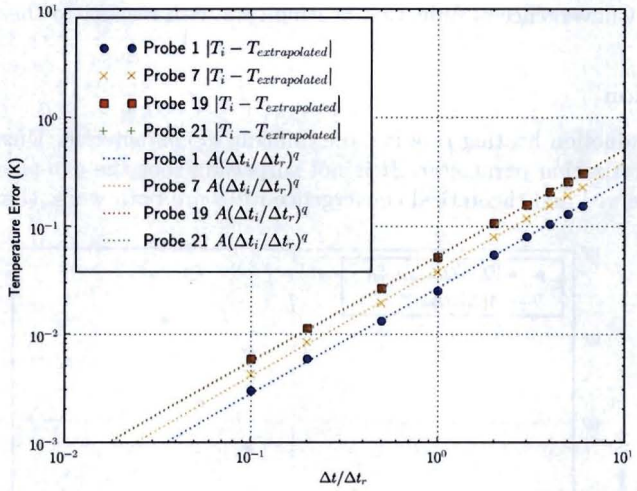

Figure 6: Time convergence of probe temperatures at the end of heat up.

Heat Up At 3600 s, all four probes exhibit linear convergence with respect to the time step, with exponents of 0.9867, 0.9907, 0.9808, and 0.9807 respectively for probes 1, 7, 19, and 21. Figure 6 shows that ansatz fits the data well, even at time steps as large as 60 s. The estimated errors at $\Delta t_r = 10$ s are 0.0268, 0.0393, 0.0532, and 0.0522 K for the four probes. The error due to time discretization is therefore quite small relative to the extrapolated temperatures of 995.041, 964.383, 777.312, and 778.988 K.

Cool Down It is evident from Figure 7 that none of the temperatures are in the asymptotic region at 60 s, and probes 19 and 20 are not asymptotic at 50 s. Only time steps smaller than 50 s were used for the ansatz fits for all of the probes. At the end of the cool down, probes 1 and 7 show sublinear convergence, with exponents of 0.8696 and 0.8640. This might be due to the discontinuity in the derivative of temperatures with respect to time and high cool rate along the midplane of the graphite. Probes 19 and 21 are still very close to linear at 0.9754 and 0.9791. All of the probes continue to exhibit small errors: 0.0181, 0.0179, 0.0291, and 0.0288 K are the predicted errors at $\Delta t_r = 10$ s. The extrapolated temperatures from the ansatz fit are 782.140, 785.036, 780.948, and 783.620 K.

5.4.2 Space Discretization

Just as with the induction heating, interpreting verification results in space is more difficult than in time. Unstructured meshes, nonlinearities, complex boundary conditions, and other complications can all cause the convergence of a multiphysics code in space discretization to deviate from a simple Richardson extrapolation. Multiple regions with different material properties also create nonuniform dependence on space discretization.

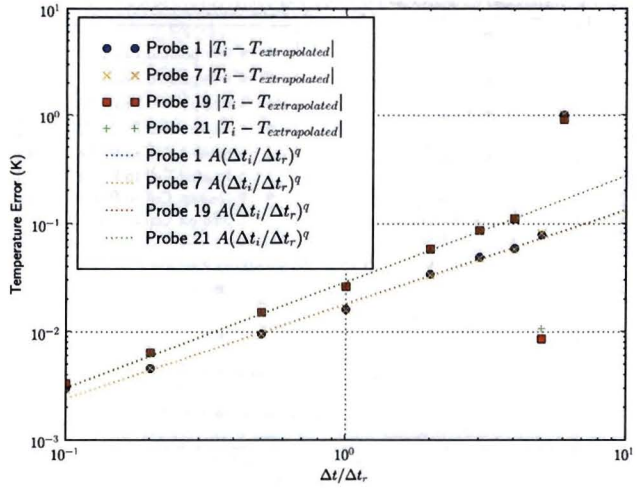

Figure 7: Time convergence of probe temperatures at the end of cool down. Only $\Delta t/\Delta t_r < 5$ were used to fit the error ansatz.

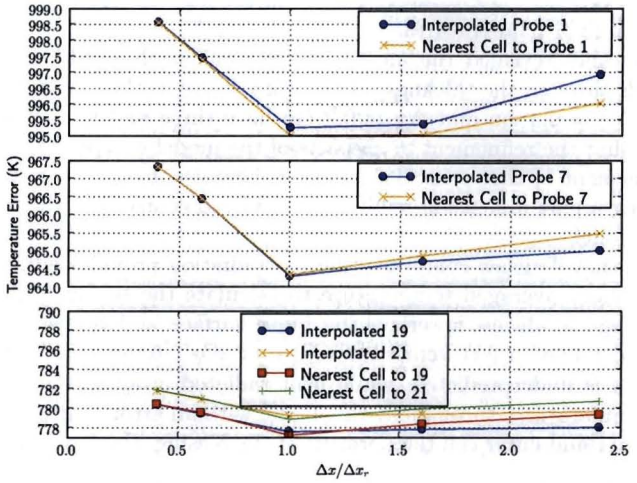

Figure 8: Probe temperatures at 3600 s. Mesh discretizations are relative to $\Delta x_r = 0.5$ cm.

Heat Up Figure 8 shows interpolated and nearest cell temperatures for the four probes. The figure is a cut axis plot to show small variations relative to the spread between the various probe temperatures. The mesh discretizations are relative to $\Delta x_r = 0.5$ cm, the cell size later used for the validation runs.

None of the probes appears to exhibit behavior that can be well-represented with an exponential ansatz. Attempts were made to run simulations finer than 0.2 cm, which is the finest mesh with results plotted in Figure 8. However, the simulations failed to complete due to problems encountered in calculating the very large view factor matrix of cell surfaces participating in the enclosure radiation. It would be possible to fit the ansatz to the finest three meshes, which show a monotonic temperature increase in all probes. No information would be available to validate the asymptotic assumption, however, and the next finest mesh at 0.8 cm is certainly not asymptotic.

It can also be seen that interpolation correction is noticeable at the coarser meshes. At $\Delta x = \Delta x_r$, it is very small relative to the uncertainty due to mesh discretization because the cells are already small compared to temperature gradients near the probes. For this reason, all of the results in the validation analysis were generated with nearest cell temperatures. In the rest of this section, however, results from both interpolated and nearest cell temperatures will continue to be shown to facilitate comparison.

Cool Down A similar non-asymptotic situation is encountered at 7200 s, as shown in Figure 9. In this case, however, a change in phenomena is evident at Δx_r. Coarser meshes have very little difference in probe temperatures,

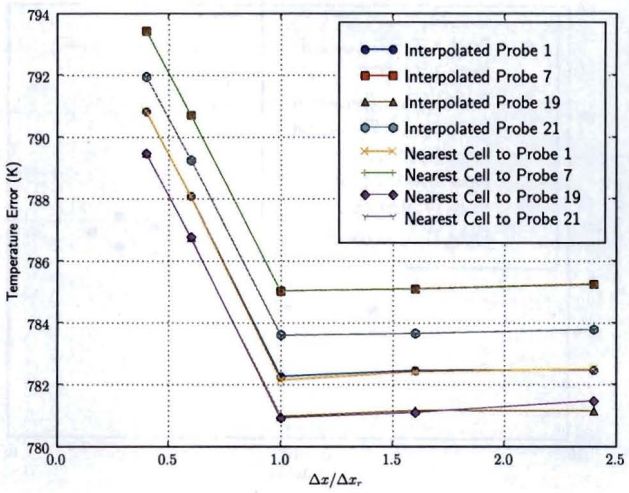

Figure 9: Probe temperatures at 7200 s. Mesh discretizations are relative to $\Delta x_r = 0.5$ cm.

yet finer meshes rapidly rise in final temperatures.

An examination of the meshes revealed the most likely reason for this sudden change in behavior. All heat transfer meshes have two cells through the thickness of the mica sleeve as they are refined from $\Delta x = 1.2$ to 0.5 cm. The meshes produced with $\Delta x = 0.3$ cm and $\Delta x = 0.2$ cm have three and four cells in the thickness dimension, respectively. Figure 9 shows that the refinement in all parts of the model except the mica thickness have little effect on final temperatures. Refinement of that critical dimension, however, has a much more dramatic effect. In some models, two to four cells along a part dimension might be considered coarse, but the temperature gradient is nearly constant in this part of the model.

Consideration of the way that Truchas computes thermal radiation provides additional insight. Truchas' finite volume formulation uses a volume-averaged temperature to calculate the surface radiation. The thermal gradient means that the volume average is always lower on the inner surface and higher on the outer surface than the continuous temperature. These two surface temperature errors affect temperatures throughout the model. Heat radiated back to the graphite is underpredicted, while heat radiated away to the ambient is overpredicted. The higher net heat loss lowers temperatures throughout the model with an error that increases over time. Even though the difference between the inner and outer cell temperatures is increasing, the gradient is slightly decreasing because the cell centroids are farther apart.

The verification analysis that follows uses the mica thickness discretization in recognition of the mica's dominant contribution to the error. The non-asymptotic nature of $\Delta x = 0.8$ and 1.2 cm no longer contraindicates an error ansatz, but there is no way to validate the asymptotic assumption for the remaining three points. Regardless of the difficulty this estimation, it is apparent that the spatial discretization error greatly outweighs the time discretization error from Section 5.4.1.

The convergence exponent for probes 1, 7, 19, and 21 are respectively 1.194, 1.135, 1.279, and 1.134. The estimated errors are 15.49 K, 15.43 K, 14.53 K, and 15.36 K. As the graphite is nearly isothermal at this time, the probes all show nearly identical convergence behavior. All the exponents are between 1.1 and 1.3, and errors are within 0.5 K of 15 K.

Total Discretization Error Estimate The verification of total induction heat rate was carried out because its calculation uses a different mesh from the heat transfer simulation. The temperatures at the probe locations are the response quantities of interest, so the error estimate should be translated to its effect on temperature. If heat loss rate to the surroundings were independent of temperature, a percent error in the heat source rate would have a proportional effect on average temperature for a constant specific heat. There is some conservatism in such an assumption because erroneously higher temperatures would produce a higher heat loss. The effect on the geometric distribution is, of course, neglected by simply scaling to an equivalent relative temperature error at each probe. Even such a simple transfer of heat source error might still be profitable, however, to compare magnitudes with other errors. A 7.144 W error is 0.7196% of the 992.786 W extrapolated total power. At a final temperature of 800 K, this would correspond to a 5.757 K error.

The two error estimates (15 K and 5.757 K), while each less than 2% of the absolute temperature, are big enough to be comparable to the validation uncertainty. An attempt to "correct" a solution by subtracting the estimated error would eliminate information about discretization uncertainty. Furthermore, it is not always possible to predict how multiple error estimates interact. In this case, however, the electromagnetics discretization was observed to contribute a positive temperature error. The mica discretization in the heat transfer simulation negative temperature error It seems safe to assume that, whatever their magnitudes, the two errors will have opposite sign and will partially cancel each other. This leads to a total discretization error estimate of approximately 10 K. This error is probably an underprediction of temperature, given the magnitude of the heat transfer mesh discretization error. Once again, caution is required in applying this estimate because of the untested assumption that the three mica discretizations are in the asymptotic region.

6 Uncertainty Quantification

This project takes a probabilistic approach to the model uncertainty quantification (UQ) and validation. In the past validation of the Truchas code has been performed in a purely deterministic fashion with a model being run using a single set of parameters. This single realization is then compared qualitatively with experimental evidence. Underneath this deterministic approach is either the assumption that the input values are known with 100% confidence, or more likely, that both aleatoric and epistemic uncertainty of the problem is being ignored. The probabilistic approach to validation first requires that the uncertainty be acknowledged, and secondly that it be quantified.

Parameters of the model are no longer a single number, but are instead treated as distributions of probable values. This fundamental shift allows a model to perform more like a natural system. In experiments it is expected that some stochastic process, variation of properties, or operations will yield similar, but distinctly different outcomes, a distribution of responses. UQ attempts to capture these mechanisms, as well as the inherent lack of knowledge about the systems that are modeled. Once propagated through the model, these uncertainties yield distributions of responses, similar to an experiment. These distributions represent the possible range of outcomes based on the combined knowledge and uncertainty about the modeled system.

To perform the UQ, first the parameters are identified, the sensitivity of the system is analyzed, the important parameters are refined, and finally the uncertainty is propagated through the model.

6.1 Phenomena Identification and Ranking Table (PIRT)

To begin the UQ process, 20 model parameters and six first order interactions between the parameters were quantified by experts using previous experience, general material values, and rough measurements of the experimental system. The parameters were ranked in order of their perceived importance on the graphite based temperature probes using expert opinion. This ranking resulted in Table 4 which shows the input parameters, their estimated mean value , and an estimate on the possible spread.

Typically engineering judgement could be used in conjunction with the PIRT to eliminate several of the bottom factors from the analysis. Because this problem has short runtimes, and there is so much epistemic uncertainty in the parameter values, it was determined to continue by considering all the parameters.

6.2 Sensitivity Analysis

In order to down select from the 20 model input parameters and focus on the important information, a sensitivity analysis was performed. The output quantities of interest are temperatures (K) at five locations: TC1, TC7, TC19, TC21 in the graphite and TC24 on top of the firebrick (Figure 2). TC24 is included as this output should be most affected by firebrick properties and boundary conditions, while having little or no effect from graphite and mica material properties.

As shown in Figure 2, TC1 is on the outer radius of the cylinder at mid-height and is most affected by the heating coils. TC7 is also mid-height, but is located on the center line of the cylinder and will be most affected by the penetration of the heating coils via conduction. TC19 and TC21 are located near the lower boundary condition, and TC24 is on top of the firebrick. These locations are chosen to provide coverage of the different physics and interactions of interest contained within the model.

The parameter upper and lower bounds in the PIRT (see Table 4) have been estimated at a variety of bounding standard deviations. For accurate comparison of parameter sensitivities, these bounds must be normalized to a common spread. In dynamics applications responses typically have quickly varying derivatives, therefore a couple of sensitivity analyses are performed on small variations about the mean, typically 0.1σ and 0.5σ. In this application, the responses of interest are very smooth, so a single analysis was deemed to be sufficient.

Table 4: Phenomena Identification and Ranking Table for the graphite cylinder experiment with relevant inputs into the Truchas code ranked by expert opinion. Shaded regions show areas of measured, or reference values.

	Parameter		**Parameter Uncertainty**			
Rank	Name	Description	Mean	Low	High	Spread
Important	micaEps	mica surface emissivity	0.5	0.3	0.7	$\pm 2\sigma$
	micaK	mica thermal conductivity	0.1	0.02	0.5	$\pm 2\sigma$
	mica_eVF	mica view factor	0.5	0.2	0.8	$\pm 2\sigma$
	graphiteHTC	graphite/firebrick heat transfer coef.	10	1	100	$\pm 3\sigma$
	firebrickK	firebrick thermal conductivity	0.1	0.02	0.5	$\pm 2\sigma$
	micaTamb	mica ambient radiation temp.	300	290	350	$\pm 3\sigma$
	firebrick_eVF	firebrick viewfactor	0.5	0.2	0.8	$\pm 2\sigma$
	firebrickTamb	firebrick ambient radiation temp.	300	290	350	$\pm 3\sigma$
Moderate	graphiteSig	graphite electrical conductivity	7.0E+04	5.6E+04	8.3E+04	$\pm 2\sigma$
	coilRad	coil radius	0.1524	0.146	0.1588	$\pm 3\sigma$
	coilCurrent	current amplitude	1645	1625	1665	$\pm 2\sigma$
	coilLength	coil length	0.13335	0.13017	0.13653	$\pm 2\sigma$
	graphiteEps	graphite surface emissivity	0.8	0.7	0.9	$\pm 2\sigma$
	graphiteK	graphite thermal conductivity	poly. fit	poly. fit - 3	poly. fit + 3	$\pm\sigma$
	graphiteCp	graphite specific heat	poly. fit	poly. fit - 50	poly. fit + 50	$\pm\sigma$
Low	frequency	current frequency	1885	1875	1895	$\pm 2\sigma$
	firebrickEps	firebrick emissivity	0.2	0.1	0.4	$\pm 2\sigma$
	micaHTC	mica/firebrick heat transfer coef.	10	1	100	$\pm 3\sigma$
	T_initial	initial material temperature	290	288	292	$\pm 3\sigma$

Interactions

graphiteSig x coilCurrent	graphiteSig x frequency
coilCurrent x frequency	graphiteK x graphiteCp
firebrickK x graphiteHTC	mica_eVF x micaK

Table 5: Parameter distributions and normalized upper and lower bounds.

Parameter	Distribution	σ_{lower}	σ_{upper}
micaEps	Normal	0.4	0.6
micaK	LogNormal	0.044	0.224
mica_eVF	Normal	0.35	0.65
graphiteHTC	LogNormal	4.64	21.5
firebrickK	LogNormal	0.044	0.224
micaTamb	LogNormal	309	329
firebrick_eVF	Normal	0.35	0.65
firebrickTamb	LogNormal	309	329
graphiteSig	Normal	63500	76500
coilRad	Normal	0.150	0.155
coilCurrent	Normal	1635	1655
coilLength	Normal	0.132	0.135
graphiteEps	Normal	0.75	0.85
graphiteK	Normal	poly. fit - 3	poly. fit + 3
graphiteCp	Normal	poly. fit - 50	poly. fit + 50
frequency	Normal	1880	1890
firebrickEps	LogNormal	0.141	0.282
micaHTC	LogNormal	4.64	21.5
T_initial	Normal	289	291

Once Table 5 was completed, a two-level ($\pm 1\sigma$) fractional-factorial design of experiments (DoE) [27] was created with the 20 input parameters and the six first-order interactions of interest. Because higher order interactions are able to be ignored, a full factorial DoE is not required and interactions can be aliased. This assumption greatly reduces the number of computational runs from $2^{20} = 1,048,576$ to the 28 sensitivity runs.

The results of the 28 sensitivity runs were analyzed using the analysis of variance (ANOVA) [28] technique generalized for N factors. An N-way ANOVA [29] allows the analysis of whether or not the mean of the observations change with respect to the changes in the input parameters or specified higher order interactions. This ANOVA analysis assumes a fixed-effects model.

Two results of the ANOVA are typically used for the sensitivity analysis: The F-value and the p-value. The F-value is a measurement of distance between two distributions. In this case, the F-value reflects how much the output distribution changes with changes to the input parameter of interest. Larger F-values indicate greater sensitivity to the parameter. Conversely, as the F-value increases, the p-value decreases. The p-value is an measure of confidence in a change being due to the parameter. Typically a $p < 0.05$ is considered statistically significant.

Heat-up Portion The first 3600 seconds of the test are dominated by the injection of energy into the system through the electro-magnetic coils placed around the cylinder. At 3600 seconds, the maximum total energy has been applied to the system and the thermal response should be predominantly affected by parameters closely tied to the EM energy deposition and conduction.

As is seen in figure 10 the parameters which affect the various temperature responses make physical sense. TC24 is most affected by the firebrick material properties and the graphite/firebrick interface. The other seven significant parameters affect the other 4 probes similarly, and are related to the graphite HT properties or coil characteristics. Another indication that the analysis makes physical sense is that the parameters affect the TC1/TC7 and TC19/TC21 pairings in the same way. TC1/TC7 are located at the center height of the log and are more affected by the coil parameters, while TC19/TC21 are close to the end of the log and are more dominated by the mechanisms of heat transfer and radiation. For example, the *graphiteSig* (electrical conductivity) and *coilRad* affect the responses at TC1

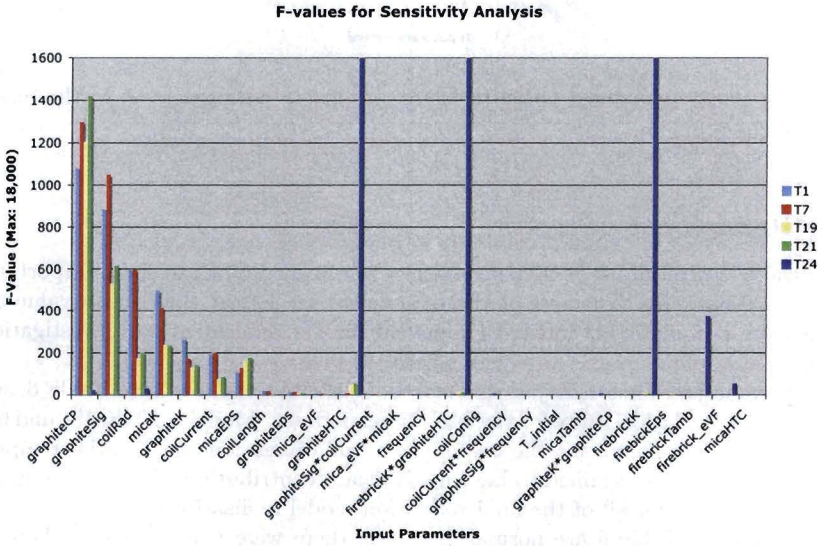

Figure 10: The F-Values at the heat-up phase termination. These parameters show the most sensitivity during the EM field application.

and TC7 far more than the other responses. These responses intuitively would have the highest coupling with the EM field which both depends on the *graphiteSig* and the *coilRad*. Also, the graphite responses are most affected by the *graphiteCp*, with TC7, TC19, and TC21 showing the highest sensitivity. This result also makes physical sense, as the responses are dominated primarily by heat conduction within the graphite.

The first seven most sensitive parameters also showed significant ($p < 0.05$) effect on the output responses of interest, and hence are retained for the further refinement and propagation of uncertainty.

Cool-down Portion The cool down portion of the experiment is dominated by radiation from the graphite cylinder, through the mica sleeve, and out to the ambient furnace walls. The experiment is performed in a vacuum, so

there is no convection. The various materials also quickly become isothermal minimizing conduction effects. As can be seen in figure 11, the final temperature of the cylinder is predominantly affected by the mica insulating layer. TC24, on the firebrick, is dominated by the firebrick's thermal conductivity, and the graphite heat transfer coefficient, which affects the graphite/firebrick interface.

Again, because of the typical casting application, only the properties affecting the graphite are retained for the uncertainty quantification. Only the mica thermal conductivity (*micaK*) is statistically significant ($p < 0.05$) and should be retained. This parameter is also important in the heat-up portion of the problem; no new parameters are added to the initial seven sensitive parameters.

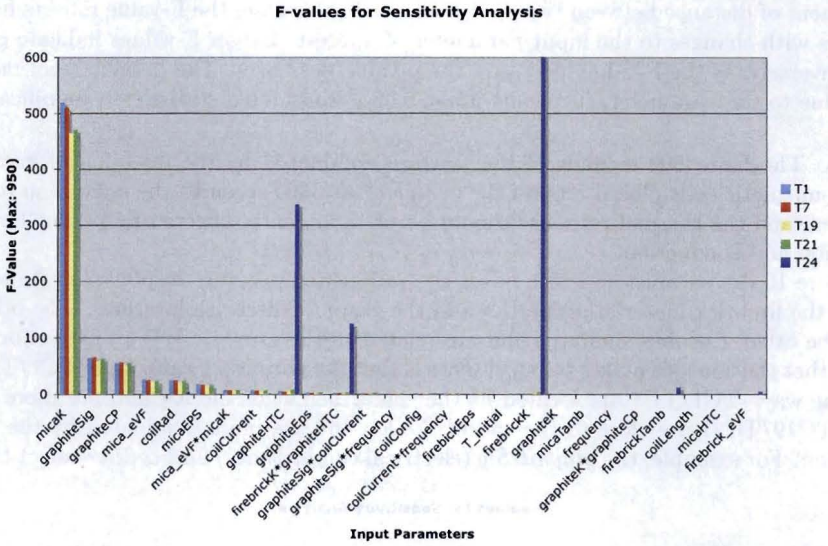

Figure 11: The F-Values at the cool-down phase termination. These parameters show the most sensitivity during the radiation portion of the problem.

6.3 Parameter Refinement

After the sensitivity study, the initial list of 20 parameters was screened down to seven important parameters which affect the four responses of interest. The rest of the parameters are left at their mean value as designated in the PIRT (Table 4). The seven parameters of interest underwent further refinement and investigation from the initially assigned distributions.

After refining the parameters, uncertainty estimates, and distributions, we have Table 6, which represents the parameters used for the final validation study for the 04K-450 experimet. Models of 04K-451 and 04K-454 experiments are similar, but have appropriate values for the coil current, coil dimensions, and initial temperature. Parameters listed without a distribution were determined to have insignificant contributions to the responses of interest, and are therefore left at their mean value for all of the final validation model realizations.

Most of the distributions in Table 6 are normal, however, there were two instances where a truncated normal was applied to remove the possibility of sampling non-physical results. For instance, the *coilRad* could not physically be smaller than the radius of the graphite cylinder. The *micaK* parameter is modeled as a lognormal distribution because it was believed that the value could not be much smaller than 0.01 W/(m*K) and most of the probability mass would be located around 0.02, however, it could range up to 0.1 W/(m*K).

6.4 Uncertainty Propagation

After refinement of the parameter probability distributions, the Latin Hypercube Sampling (LHS) approach was chosen for uncertainty propagation through the model. The LHS design is a space filling technique that forces sampling in the tails of the distribution with a lower number of samples than Monte-Carlo sampling. Even though model realizations completed in a reasonable amount of time (approximately 1 hour), hundreds or thousands of model realizations required for other sampling techniques could not be afforded.

Table 6: Parameter values for uncertainty propagation in Experiment 04K-450. Parameters without distributions are held at their mean value because the model was determined to be insensitive to them.

Parameter	Distribution	$\bar{\mu}_x$	σ_x
graphiteSig	Normal	70,000	7000
coilRad	Truncated Normal	0.1524	3.175e-3
coilCurrent	Normal	1646	9.8
graphiteK	Normal	poly. fit	2
graphiteCp	Normal	poly. fit	25
micaK	LogNormal	0.04	0.02
micaEps	Truncated Normal	.75	0.05
frequency		1885.0	
coilLength		0.13335	
graphiteEps		0.8	
firebrickEps		0.8	
graphiteHTC		10	
firebrickK		0.02	
mica_eVF		0.5	
micaTamb		300	
firebrick_eVF		0.5	
firebrickTamb		300	
micaHTC		10	
T_initial		290	

Given the amount of time and man power available to the team, a balance between defining the output distributions well and run time of the simulations was needed. To estimate how well the output distributions would be defined, the standard error (S_E) was calculated:

$$S_E(\sigma, n) = \frac{\sigma}{\sqrt{n}} \tag{9}$$

where the output distributions from the sensitivity analysis were used as an approximation of the final σ. We decided that 50 realizations for each of the three experiments would be a good compromise between fidelity and resources. There is not a large difference in the resolved temperature between any of choices of n, but 50 puts the distribution in the area of $\pm 3K$ which was determined as acceptable. Southwest Research Institute's NESSUS program was used to generate the seven dimensional LHS design for the 150 validation runs.

7 Validation

Validation of the project is approached in a *statistical process control* fashion. The process is the heat up and cool down of the graphite cylinder, with the uncertain model realizations providing the statistical bounds at various time points. The validation is therefore determined on how well the experiment falls within the process bounds dictated by the model.

On top of the process control analogy, the various test points in time are considered Bernoulli trials. There are 12 such trials for the full process, and at each of the 12 points the process can either pass (fall within the model uncertainty) or fail (fall outside of the model uncertainty. A hypothesis test is performed to determine whether or not the model is a good representation of the data.

7.1 Metric Description & Decisions

In their paper "Validation of a Mathematical Model Using Weighted Response Measures", Paez, et. al. [30] demonstrate the method of validating the model using a Bernoulli process. The validation process therefore is as follows:

1. Given a set of model realizations for a point in time (n), use a Kernel Density Estimator (KDE) to estimate the upper and lower confidence values (CI_U, CI_L).

2. Determine if the experimental response (X_n) passes (1) or fails (0) using $CI_L < X_n < CI_U$.
3. Repeat steps 1 and 2 for all n and calculate the success rate: $S = \frac{successes}{n}$.
4. From a Binomial distribution [28], $B(p, n)$, determine the required success rate: $S_{val} = \frac{P_{success} \geq (1-\alpha_{fail})}{n}$.
5. The model is determined to be valid if $S \geq S_{val}$.

It was suggested that decisions on thresholds for accuracy, adequacy, and possible weighting of individual trials based on importance, should be made before comparing the UQ runs with the experiment. Based on these ideas, we decided on the following:

- The model will be evaluated at $n = 12$ points in time.
- The experiment should fall within a two-sided CI of 90% determined by the UQ model runs at each comparison point.
- There will be no expansion of these probability intervals for the concept of adequacy. Because this problem does not posses the intricacies of real mold shapes, the concept of an adequacy interval on the temperatures does not translate well.
- The null hypothesis (H_o) confidence level for the Bernoulli trials is set to 95%. This level is interpreted as having a 5% chance of rejecting a good model ($\alpha_{fail} = 0.05$). The level is drawn from $B(12, 0.90)$.
- Weighting is not applied to any of the points. Because this is an example problem it is determined that every point is as important as the others. In some applications the way the mold heats up, or the peak temperature, or the cool down may be more important and weighting can be applied in those instances.

An assumption of the binomial distribution, from which the 95% pass rate for the Bernoulli trials are drawn, is that of (Identical and Independent) IID trials. The probability of success in all cases ($p = 0.90$) and are therefore identical. The chance of passing one trial and then the next is not truly independent as the process shows correlation from one point in time to the next. Looking into working with dependent trials, however, showed a complicated and niche research area of statistics. Therefore, the dependance is ignored and the validation proceeded using the IID assumption. Future work may incorporate the dependence of success or failure based on previous results.

7.2 Results

The validation results for the entire heat up and cool down process are plotted in Figure 12 for the TC7 sensor and summarized in Table 7. While the plots follow the TC7 sensor, as a representative sample, for a graphical it is neither the best nor worst sensor out of the group; rather it illustrates well the visual shapes and biases in the model and represents most of the physics in the system.

As was pointed out in the sensitivity study, the parameters most affecting the heat up portion of the process are: *graphiteCp, graphiteSig, coilRad, micaK, graphiteK, coilCurrent,* and *micaEps*, with *graphiteK*, and *coilCurrent* having the most effect on TC1 and TC7. As can be seen in Table 7 TC1 and TC7 are valid for the 12" coil experiments, TC19 and TC21, however, do not meet the validation requirements for these experiments. For 04K-454, with a 16" coil, none of the probes meet the validation requirements during heat up. As seen in Figure 12, the model consistently under predicts the heat up process.

Cool down is completely dominated by the *micaK* parameter. Quantitatively the model passes the cool down validation for all sensors in the 04K-450 and 04K-451 simulations. The 04K-454 realizations, however, show a failure for the TC1/TC7 probes. This result is interesting as it is the opposite from what was observed during the heat up phase, when these two probes typically outperform the TC19/TC21 probes. Looking at Figure 12, it can be seen that the model realizations do not lose energy as fast as the experiment.

The failures in the 04K-454 experiment are related to the heat up portion which, it appears, failed to deposit enough energy into the cylinder. This lower final temperature means that cool down portion is biased to producing too low of temperatures. Combining both the heat up and cool down trials into a single Bernoulli process shows that overall for experiment 04K-451 the process is validated. Experiment 04K-450 is valid near the coil with the probes toward the end showing invalid results. Experiment 04K-454 shows that the model is rejected and that more work needs to be done to understand the 16" coil configuration.

Looking specifically at Experiment 04K-454 there are a few areas into which further investigation could be performed to determine an accurate model. One possibility is that the experimental data itself is an outlier. Because

Table 7: Validation metric results for the full duration of the three experiments 04K-450, 04K-451, and 04K-454.

	TC1	TC7	TC19	TC21
S_{val}	0.6667	0.6667	0.6667	0.6667
S_{450}	1.0000	0.9167	0.5000	0.5000
S_{451}	1.0000	1.0000	0.6667	0.6667
S_{454}	0.1667	0.1667	0.3333	0.2500

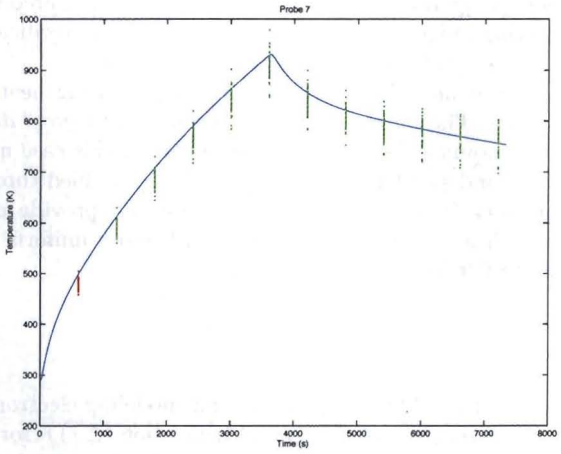

(a) 04K-450

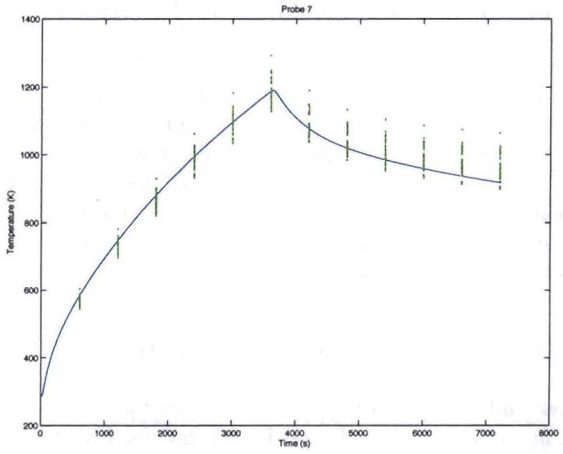

(b) 04K-451

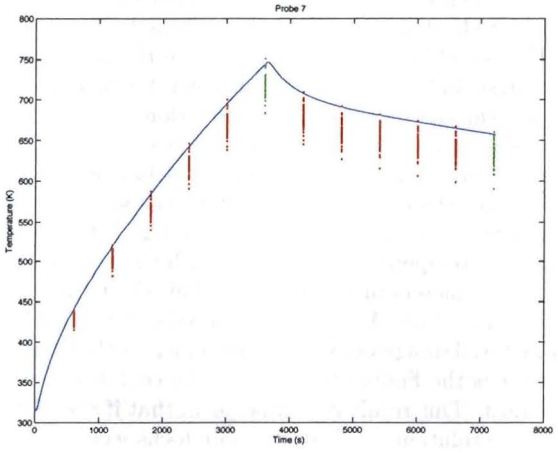

(c) 04K-454

Figure 12: Validation results for the heat up and cool down processes for probe 7. Experimental data are represented by the blue lines. The colored dots represent model realizations from the simulations of temperature at the given time points. Red dots indicate that the experiment is outside of the $p = 90\%$ interval (failure), while green indicates the experiment is within the interval (success).

there is only a single run it is inconclusive whether this is a mean or merely a 99th percentile realization of the system. More experiments would confirm the true mean behavior of the system with a 16" coil.

Another source of error in our model, as described in the Verification section, comes from spatial discretization of the mica insulating sleeve. Estimations from this error indicate the modeling results could be up to 15 K too low. This level of error has been ben shown to affect the outcome of the validation results. Other sources of modeling error and assumptions certainly exist in the problem, but more than likely lead to smaller errors than the bias seen in the 04K-454 validation.

As mentioned previously there is a large uncertainty in the *micaK* parameter from the physical construction. Reducing this error could come in the form of casting an insulation sleeve. Another large source of uncertainty is in the *graphiteCp* parameter. This parameter was determined to be the most important during heat up, yet is based on the manufacturers data for a generic sample and the uncertainty bounds are from expert judgement. Material property tests could reduce this uncertainty on the mean and take the spread from epistemic to aleatoric.

As with many of the high fidelity engineering models created under the ASC program, the qualitative evidence points to the code physics being correct, but the quantitative results show portions of the model are invalid. This result appears to be because the codes capture the physics well, but there is a tremendous amount that is unknown about the parameters going into the models, and assumptions made by the modelers. In this case material properties are drawn and modified from manufacturer data sheets at room temperature, other properties are unmeasured such as those for the mica sleeve, and still others are purely engineering judgement. It was also found during verification that the spacial discretization of the mica sleeve may contribute a significant error to the final results.

In general the results suggest that heat is not being transfered into the graphite fast enough during heat up, and also not enough heat is leaving the graphite during cool down. These phenomena are dominated by *graphiteCp*, *graphiteK*, and the*micaK* properties, all of which present large amounts of epistemic uncertainty. In this case many of the important graphite material properties could be measured independently and statistically quantified through experimentation. The experiments would change the uncertainty from epistemic to aleatoric and provide more accurate model results. Also, the experiment could be run with a cast insulation sleeve with more uniform and measurable properties rather than the layered mica and graphite felt sleeve that was used.

8 Summary & Conclusions

In this study, we performed an V&V assessment of the capability of the ASC code Truchas for modeling electromagnetic induction heating. It was found that for EM, the order of convergence of the spatial truncation (Δx) error was 1.5, while the order of time-step (Δt) convergence was 2.0 which agrees with the theoretical value for the second-order time advancement scheme for the Maxwell equations. The magnitude of the combined spatial and time truncation error in the total joule power calculated, at the nominal tet-mesh and EM time resolution levels, was about 7 W (0.7% relative error), which was dominated by the contribution from the time discretization.

For HT, the Δt convergence analysis indicates a fitted order of 1.0, which is consistent with the backward Euler time stepping method used in solving the energy conservation equation. The truncation error due to time discretization alone, at the nominal HT Δt value of 10 s, was negligible, i.e., less than 0.1 K in the graphite temperature. The Δx convergence analysis, however, did not yield well defined error behavior. Asymptotic convergence of the spatial discretization error was assumed (rather than demonstrated) with the solutions from the finest three grids (Δx = 0.5, 0.3, and 0.2 cm), resulting in a fitted order of convergence of 1.2 and an estimated error magnitude of up to 15 K (under prediction) in the graphite temperature computed with the nominal (0.5-cm) hex mesh.

There are twenty physical input parameters in this problem that affect the output quantity of interest, which is the graphite temperature as a function of time. A sensitivity analysis with the ANOVA technique of the run output results from a two-level fractional-factorial design of experiments indicates that the twenty parameters can be screened down to seven. Of particular interest is the finding that during the cool-down, the single most sensitive parameter is the thermal conductivity of the mica. This result is significant in that if the high importance of conduction across the mica sleeve were known before the solution verification, more focus would be placed on resolving the mica layer which would reduce discretization errors.

Probability distributions were then developed for the seven selected parameters in order to propagate the uncertainties in these input parameters to the output quantities. The uncertainty quantification was performed, for each of the three experiments, with an ensemble of 50 realizations each consisting of a set of randomly sampled parameter values. The parameter values were chosen with the orthogonal Latin Hypercube Sampling method, which forces the sampling in the tails of the distribution with a low number of samples. The resulting distribution of outputs were then used for validation against the experimental data.

Our validation metric is based on using a Bernoulli process in which a hypothesis test is formed by asking how

many successes are required so the model is not rejected with 95% confidence. We decided that the experimental data should fall within a two-sided confidence interval of 90% as determined by the UQ model runs at each comparison point (test) represented by the temperature at each of four locations within the graphite cylinder and at each of 12 time instants during the experiment. The concepts of adequacy (to expand the probability interval) and possible weighting (of individual trials) were not applied; and the identical and independent assumption was made of the trials.

According to the metric described above, the model is valid for Experiment 04K-451 at all four tested thermocouple locations. For Experiment 04K-450, the model is valid for the locations at mid-height of the cylinder, which are close to the coil, while it is not valid towards the ends of the cylinder. For Experiment 04K-454, the validation test does not pass at any of the thermocouple locations. The primary reason for this is the large amount of uncertainty in some of the model parameters, especially the thermal conductivity of the mica sleeve, whose geometry and internal material arrangement were not well defined in the experiments. We note, however, that the validation metric chosen is a stringent one, with pure accuracy as the primary criterion, i.e., with no considerations for adequacy, and for weighting and inter-dependency of the individual tests.

Finally, most of the uncertainties in the model parameters of this problem are of epistemic nature. In particular, the thermal properties and surface emissivities of the mica sleeve and firebrick insulation are not known. The experiments would be more valuable for validation assessment if all important materials were well characterized.

Acknowledgment

We would like to thank Rob Aikin of MST-6, LANL, for providing all the data used in this project and the photographs of the experimental setup presented in the report. We appreciate his patience and enthusiasm when explaining the experiments and related matters to us. We would also like to thank Neil Carlson of CCS-2, LANL, for the numerous discussions on the electromagnetics model which he developed and implemented into Truchas. Neil also performed some initial calculations of the graphite cylinder heating experiment, upon which our current computational model is built. Finally the support of NNSA's ASC Program is greatly appreciated.

References

[1] Office of Advanced Simulations & Computing. Advanced Simulation and Computing Program Plan FY05, September 2004.

[2] The Telluride Team. Truchas Physics and Algorithms, Version 2.3.0. Technical Report LA-UR-07-4517, Los Alamos National Laboratory, 2007.

[3] The Telluride Team. Truchas User's Manual, Version 2.3.0. Technical Report LA-UR-07-4516, Los Alamos National Laboratory, 2007.

[4] The Telluride Team. Truchas Reference Manual, Version 2.3.0. Technical Report LA-UR-07-4519, Los Alamos National Laboratory, 2007.

[5] A. Bossavit. *Computational Electromagnetism: Variational formulations, Complementarity, Edge Elements.* Academic Press, San Diego, CA, 1998.

[6] R. Hiptmair. Multigrid method for Maxwell's equations. 36(1):204–225, 1998.

[7] M.W. Glass. Chaparral v2.x: A library for solving enclosure radiation heat transfer problems. Technical Report Distribution UC-700, Sandia National Laboratory Report, May 2007.

[8] Internet sites containing miscellaneous surface emissivity values:
http://www.omega.com/literature/transactions/volume1/emissivityb.html
http://www.engineeringtoolbox.com/emissivity-coefficients-d_447.html
http://www.coleparmer.com/techinfo/techinfo.asp?htmlfile=emissivity.htm&id=254.

[9] John H. Lienhard. *A Heat Transfer Textbook.* Prentice-Hall, Inc., Englewood Cliffs, N.J., 1981.

[10] L. F. Richardson. The approximate arithmetical solution by finite differences of physical problems involving differential equations with an application to the stresses in a masonry dam. *Transactions of the Royal Society of London*, Ser. A(210):307 – 357, 1910.

[11] L. F. Richardson and J. A. Gaunt. The deferred approach to the limit. *Transactions of the Royal Society of London*, Ser. A(226):299 – 361, 1927.

[12] J. R. Kamm, W. J. Rider, and J. S. Brock. Combined space and time convergence analyses of a compresible flow algorithm. In *16th AIAA Computational Fluid Dynamics Conference*, number AIAA-2003-4241, June 2003. LA-UR-03-2628.

[13] F. M. Hemez. Non-linear Error Ansatz Models for Solution Verification in Computational Physics. Technical Report LA-UR-05-8228, Los Alamos National Laboratory, 2005.

[14] J. S. Brock. Bounded numerical error estimates for oscillatory convergence of simulation data. In *18th AIAA Computational Fluid Dynamics Conference*, number AIAA-2007-4091, June 2007.

[15] D. P. Smitherman, J. R. Kamm, and J. S. Brock. Calculation verification: pointwise estimation of solutions and their method-associated numerical error. *Journal of Aerospace Computing, Information, and Communication*, 4:676 – 692, March 2007. LA-UR-06-2006.

[16] T. Tippetts, F. X. Timmes, J. S. Brock, and J. R. Kamm. Field-wide Calculation Verification for Finite Volume Hydrodynamics Simulations. In *18th AIAA Computational Fluid Dynamics Conference*, number AIAA-2007-4090, June 2007. LA-UR-07-4009.

[17] SAS Institute, Inc., SAS Campus Drive, Cary, NC 27513. *JMP Design of Experiments Guide*, 7 edition, 2007.

[18] George G. Roussas. *A First Course in Mathematical Statistics.* Addison-Wesley, 1973.

[19] The Mathworks, 3 Apple Hill Drive, Natick, MA 01760-2098. *Statistics Toolbox User's Guide*, 6 edition, 2007.

[20] T. L. Paez, J. Massad, T. Hinnerichs, C. C. O'Gorman, and P. Hunter. Validation of a Mathematical Model Using Weighted Response Measures. In *Proceedings of the 2008 International Modal Analysis Conference.* SEM, 2008.