LA-UR- 02-6646

Approved for public release; distribution is unlimited.

Title: Estimating the Prediction Accuracy of High-rate, High-temperature Constitutive Models for Taylor Impact Applications

Author(s): François M. Hemez, LANL

Submitted to: COMPLAS VII, the 7th International Conference on Computational Plasticity, Barcelona, Spain, April 7-10, 2003.

VII International Conference on Computational Plasticity
COMPLAS 2003
E. Oñate and D. R. J. Owen (Eds)

ESTIMATING THE PREDICTION ACCURACY OF HIGH-RATE, HIGH-TEMPERATURE CONSTITUTIVE MODELS FOR TAYLOR IMPACT APPLICATIONS

François M. Hemez*

*Los Alamos National Laboratory
Engineering Sciences and Applications (ESA-WR)
Mail Stop T006, Los Alamos, New Mexico 87545, U.S.A.
E-mail: hemez@lanl.gov — Web: http://www.esa.lanl.gov

Key words: Model validation, predictive accuracy, test-analysis correlation, Taylor impact.

Abstract. *Material scientists commonly use Taylor anvil impact experiments to calibrate constitutive and equation-of-state models that include plasticity, high strain-rate and temperature dependency. The Taylor anvil experiment consists of impacting a sample of material against a rigid wall and measuring its deformed profile. The measured profiles are compared to numerical predictions and parameters of the material model are calibrated to improve the predictive accuracy. The application for which constitutive models, among other things, must be validated is the containment of hazardous residues. Such residues result from hydrodynamic experiments during which materials are compressed and radio-graphed with multiple-axis, high-energy X-rays. The discussion presented here overviews the successive steps of the validation process for a Zerilli-Amstrong plasticity model. A suite of validation experiments is designed, starting with static testing and progressing to Hopkinson bar tests and Taylor impact tests. It is shown how uncertainty is propagated through the forward and inverse calculations. Based on test-analysis correlation, statistical metamodels are created to estimate the predictive accuracy of the plasticity model in regions of the operational space where testing is not possible. The operational space is defined in terms of varying strain-rates and temperatures. The result of the predictive accuracy assessment is that the fidelity of a numerical simulation can be estimated, even before performing the calculation itself. The suite of validation tools helps analysts decide whether their models meet the accuracy requirement for a particular application.*

Approved for unlimited, public release. LA-UR-02-6646. ***Unclassified****.*

1. INTRODUCTION

Material scientists commonly use an experiment known as the Taylor anvil impact to develop constitutive and equation-of-state models. Because of the regimes for which they are developed, such models generally include plasticity, high strain-rate and temperature dependency. Examples used in engineering mechanics include the Johnson-Cook[1] and Zerilli-Amstrong[2] models. The Taylor anvil experiment consists of impacting a sample of material against a rigid wall and measuring its deformed profile. The measured profiles are compared to numerical predictions and parameters of the constitutive equations can be calibrated to improve the model's predictive accuracy.

The application we are interested in involves the containment of hazardous residues inside pressure vessels. Such residues are the byproduct of hydrodynamic tests during which materials are radio-graphed with multiple-axis, high-energy X-rays. During the tests, it may happen that small fragments impact and penetrate the vessel's wall. Pre-shot qualification requires that each test be numerically simulated to assess the potentially hazardous situations and the margin of safety. A pre-requisite to simulation-based qualification is the assessment of predictive accuracy of the numerical models. This publication does not address the development of constitutive models per say. Instead, the discussion illustrates the successive steps of the validation process. A suite of validation experiments is designed, starting with Hopkinson bar testing and proceeding with the higher strain-rate Taylor tests, to assess the predictive accuracy of material models over different regions of the operational space. The discussion overviews the technology developed and applied for assessing the predictive accuracy of material models. It includes sampling, design of experiments, metamodeling, Bayesian calibration and non-probabilistic uncertainty quantification.

Based on test-analysis correlation and statistical metamodeling, the predictive accuracy of a constitutive model is assessed in regions of the operational space where testing is not possible. The significance of this assessment is that the fidelity of a numerical simulation can be estimated, even before performing the calculation itself. The suite of tools developed for model validation helps analysts decide whether their models meet the accuracy requirement for a particular application.

2. MOTIVATION

The application we are interested in is the pre-testing qualification of containment vessels for the radiography of hydrodynamic experiments. A hydrodynamic experiment consists of compressing materials to the point where they “flow” and must be characterized by a distribution of density and pressure instead of strength. The containment vessels are also designed to facilitate the radiography of the experiment with multiple-axis, high-energy X-rays. The vessels, pictured in Figure 1, are 6 feet (1.83 m) in diameter with 2-inch (50.8 mm) thick walls and manufactured out of high-strength steel known as HSLA-100.

Because of the potentially toxic nature of the materials involved, each shot must be qualified before a review board will allow the testing team to proceed with the physical experiment. The difficulty, of course, is that experimental measurements are available only after the physical experiment has been performed. Measurements cannot be used for assessing the safety margin and qualification; they can only be used after the fact to verify that the simulations were correct. Qualification must therefore be assessed on the basis of simulations only, which implies that the predictive accuracy of the models going into the code must be established before the experiment can be performed.

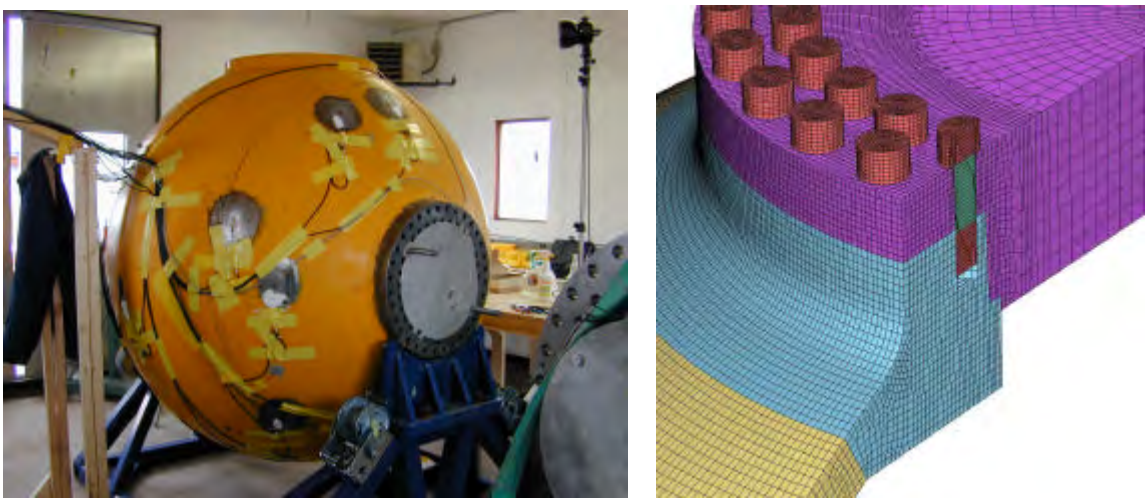

Figure 1: Photograph of a containment vessel (left) and illustration of the finite element mesh (right).

Two families of models are developed to simulate the response of the pressure vessels to the impulse and estimate the safety margin. They are referred to as the "Low Fidelity" (LoFi) and "High Fidelity" (HiFi) models. The engineering LoFi models count about 91,000 finite elements and 600 contact surfaces, which results into more than one million equations. The analysis software is HKS/Abaqus^{TM}, a general-purpose finite element package that can handle arbitrary non-linearity, impact and contact dynamics through its explicit integration schemes.[3]

The engineering models are complemented with HiFi models analyzed with ParaDyn, a parallel finite element package developed by the Lawrence Livermore National Laboratory.[4] High-fidelity models make it possible to study the leak path likely to develop through the bolted connection between the lids that close openings and the vessel's wall. The computational mesh developed to study this critical area is illustrated in Figure 1. The HiFi models result into problems that are one order of magnitude larger than the LoFi models. They are analyzed on hundreds of processors of massively parallel machines such as the Los Alamos BlueMountain and ASCI-Q super-computers, both capable of achieving several TERAOPS of computational power. (One TERAOPS is defined as 10^{+12} floating-point multiplications per second.) Figure 2 shows four snapshots of the deformation profile of a quarter-symmetry vessel simulated with HKS/Abaqus^{TM}.

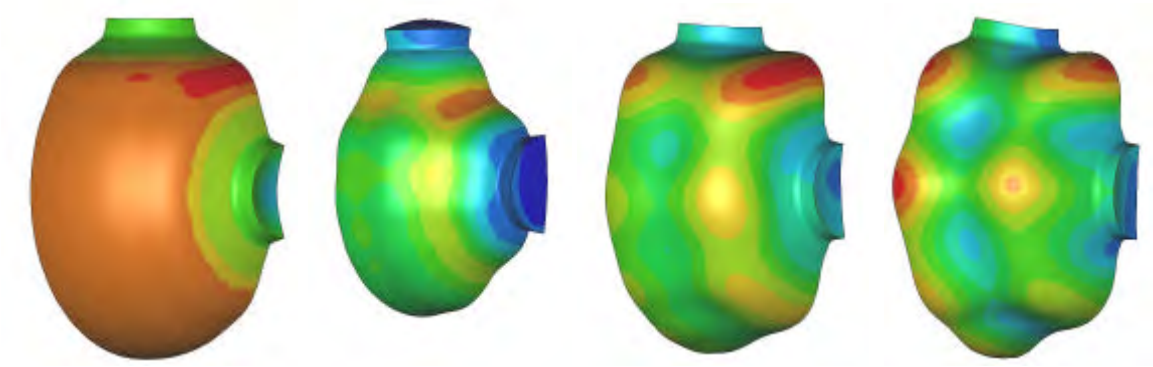

Figure 2: Simulated deformation profiles at four instants after impulse (0.5 ms, 0.9 ms, 1.3 ms, 2.0 ms).

Figure 3 illustrates the agreement obtained between a series of numerical simulations and measurements performed during a hydrodynamic experiment. The visual agreement between the curves is a clear indication of the excellent predictive accuracy of the simulation. Note that predicted and measured responses start to show significant disagreement after 4 ms, when only the first two milliseconds of response in fact matter to estimate the safety margin.

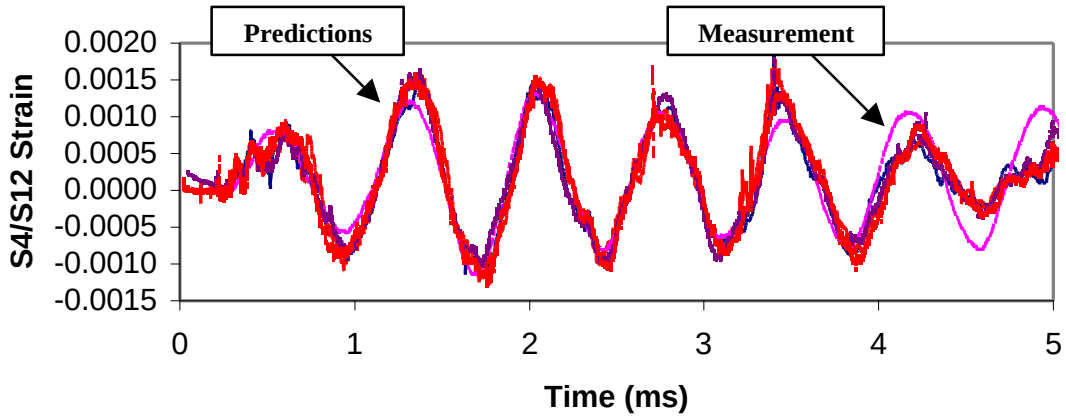

Figure 3: Predicted and measured signals. (Red and blue lines: Simulations. Pink line: Measurement.)

Obviously, reaching such good agreement is only possible to the extent where validated models have been implemented. This publication focuses on the validation of the material's constitutive model. The other components of the Verification and Validation (V&V)—such as verifying the contact algorithm; verifying the numerical solvers; making sure that the grid size provides a converged solution; and selecting an appropriate time resolution—are addressed separately. Because of the nature of the transient applied and its effect on the structure, a material model must be implemented that includes plasticity, high strain-rate and temperature dependency. In the remainder, the validation of the Zerilli-Amstrong constitutive model of plasticity is discussed even though other models, such as the Johnson-Cook plasticity and equation-of-state models,[1,5] have been documented as well.

3. THE VALIDATION DOMAIN

Before proceeding with the description of some of the validation's steps, the notion of validation domain must be introduced. Generally speaking, a numerical simulation is always developed to analyze a given operational domain because the interest of a point-prediction, that is, the prediction of a model that could not—even slightly—be modified, is limited. For example, a simulation of aero-elasticity is parameterized with design variables such as flow speed and angle-of-attack. The usefulness of a simulation that could be analyzed only for a single combination of these variables would be limited.

Likewise, the plasticity model is developed to run numerical simulations at different combinations of strain-rates and temperatures. For our application, these two input variables define the operational space of interest. The validation domain is here simply defined as the region of the operational space where the mathematical or numerical model provides acceptable accuracy for the application of interest. This concept is illustrated in Figure 4. Simply speaking, validation is achieved when the predictive accuracy of the model has been assessed within the operational domain, a consequence of which is the identification of the region—or validation domain—that provides sufficient accuracy.

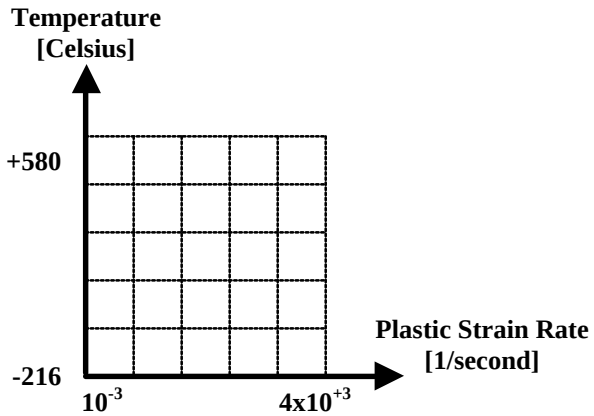

Figure 4: Definition of the validation domain.

As mentioned previously, model validation cannot rely on measurements collected during the hydrodynamic experiment because the models must be validated before the experiment is performed. A suite of validation experiments is therefore designed. A validation experiment is a somewhat simpler procedure that isolates the phenomenon of interest. The suite of validation experiments hence provides increasing levels of understanding of the fundamental physics for the application. Static material testing comes first, which allows the identification of bulk mechanical properties such as the modulus of elasticity. Because such tests are static in nature, they must be augmented with Hopkinson bar tests. Hopkinson bar tests however do not provide sufficient resolution in the regime of interest, that is, at high strain-rates and varying temperatures. To provide more insight into the behavior of the plasticity throughout

the validation domain, Taylor anvil tests are performed next. The validation experiments explore different regions of the validation domain while providing successive material models that are hopefully consistent with each other. The discussion presented in this publication focuses on the definition of an error metric between inferences made from the Hopkinson bar tests and inferences made from the Taylor anvil impact tests (see section 6).

The Zerilli-Amstrong model estimates the stress s resulting from a plastic deformation as:

$$s = C_o + C_1 e^{-C_3 T + C_4 T \log\left(\frac{de_P}{dt}\right)} + C_5 e_P^N \tag{1}$$

where the symbol T represents temperature and e_P denotes plastic strain. The six parameters C_0, C_1, C_3, C_4, C_5 and N are material-dependent constants that can be calibrated to improve the predictive accuracy of the model. Because of the large spread of strain-rates for which a validated model is sought (from the quasi-static rate of 10^{-3}/second to $4\text{x}10^{+3}$/second), another symbol S_R is introduced that defines the logarithm of the plastic strain rate:

$$S_R = \log_{10}\left(\frac{de_P}{dt}\right) \tag{2}$$

It is emphasized that the Zerilli-Amstrong model is chosen for simplicity, not necessarily because it represents the "best" model for this application. Also, it is important not to confuse the six calibration variables $(C_0;C_1;C_3;C_4;C_5;N)$ with the two input parameters $(T;S_R)$ that define the operational space or validation domain. The main difference between the two is that calibration variables are introduced by our particular choice of plasticity model. Should another physical model be adopted, the number of calibration variables would likely change. The dimensionality of the operational space, however, never changes and the plasticity models—whatever they are—must still be validated at various combinations of $(T;S_R)$.

4. FORWARD PROPAGATION OF UNCERTAINTY

Two of the key technologies that support V&V are the propagation and analysis of uncertainty. This is because model validation is essentially an exercise in the quantification of uncertainty, whether it originates from the physical experiments, computations, or modeling assumptions. In this section and the next, some of the tools employed to propagate uncertainty are briefly illustrated. They include Monte Carlo sampling for propagation through forward calculations (section 4) and Bayesian calibration for backwards propagation (section 5).

Uncertainty quantification tools that are not discussed, although they have been used in this and other studies, include the design of experiments and analysis-of-variance. Our application of these techniques, that we have found to be of paramount importance when dealing with large-size problems, are documented in several other publications.[6,7,8]

Figure 5 shows several deformed profiles simulated with the finite element package HKS/Abaqus™. In this numerical simulation, an axi-symmetric mesh is impacted against a

perfectly rigid surface, which produces large deformations (over 260%) and significant plastic strain at the crushed end of the cylinder. The profiles compared in Figure 5 are obtained by varying the calibration variables $(C_0;C_1;C_3;C_4;C_5;N)$ according to a probability law.

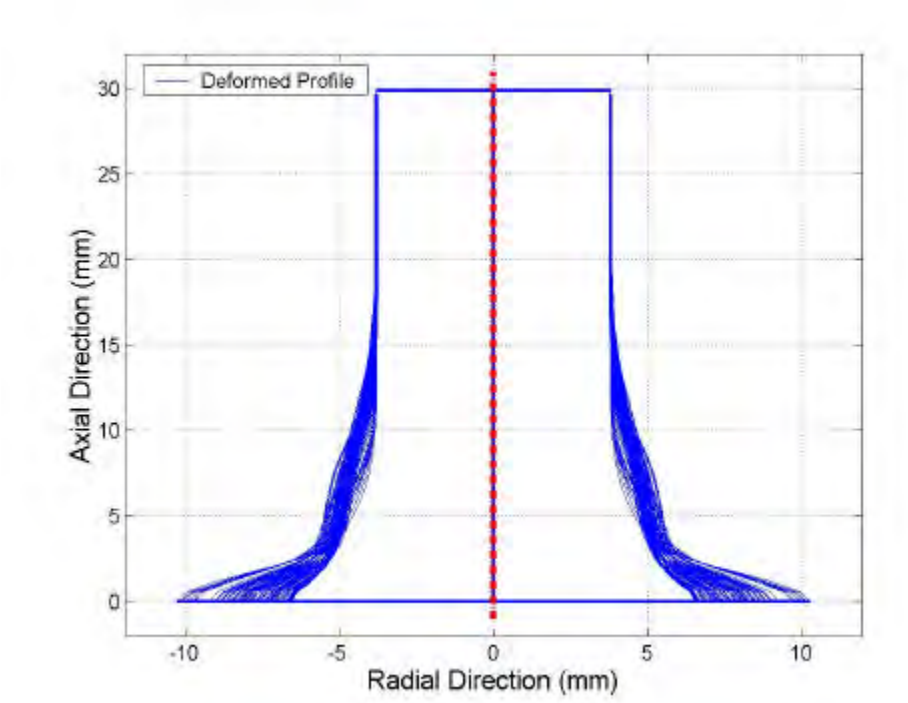

Figure 5: Several examples of randomly selected simulated deformation profiles.

In this illustration, each calibration variable is assumed to vary according to a Gaussian Probability Density Function (PDF). The mean values are nominal values obtained by best fitting the Zerilli-Amstrong model of plasticity to Hopkinson bar test data. The standard deviations are arbitrarily initialized to 20% of the mean values. Because the calibration variables are assumed to be independent and uncorrelated, sampling the six individual PDF laws is straightforward. Random samples of coefficient values are drawn from the six normal distributions independently. A combination of variables $(C_0;C_1;C_3;C_4;C_5;N)$ then defines a specific material model and the impact simulation is repeated for each model.

This procedure defines a Monte Carlo simulation and a total of 1,000 finite element calculations are performed. This number of simulations is selected somewhat arbitrarily based on the time necessary to perform a single analysis and the available computing resource. Results of the Monte Carlo simulation are illustrated in Figures 6 and 7.

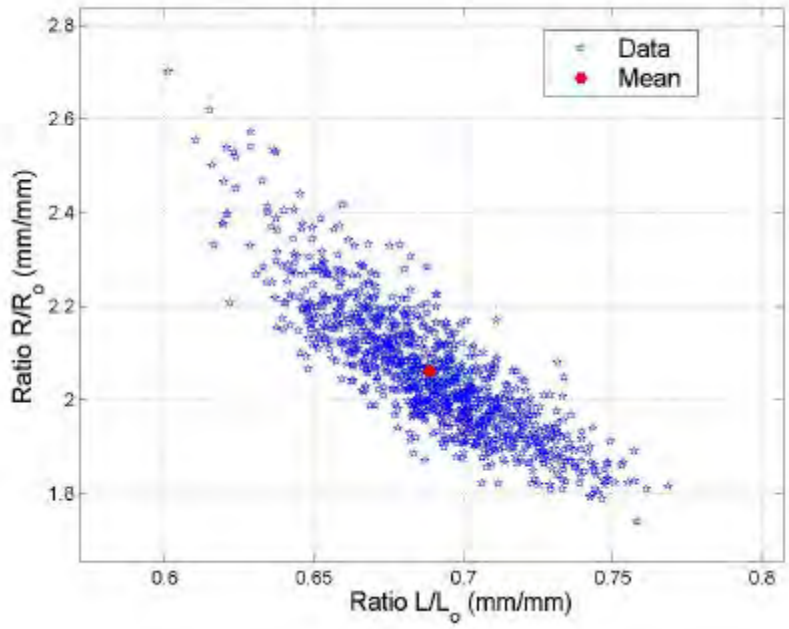

Figure 6: Distribution of features (L/L_o) and (R/R_o) obtained from a 1,000 Monte Carlo simulation.

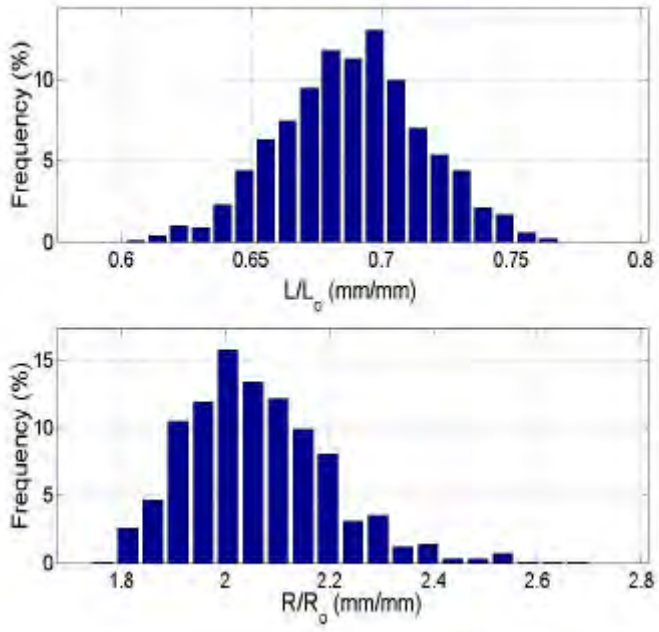

Figure 7: Histograms of simulated features (L/L_o) and (R/R_o).

In Figures 6 and 7, two features of the response are defined to characterize the deformed profiles, the ratios of final-to-initial lengths (L/L_o) and footprints (R/R_o). In Figure 6, each

point corresponds to the result of an impact simulation for a particular material model. It can be observed that there is a significant correlation between the two output features, as one would expect because the shorter the cylinder, the larger its footprint. The histograms shown in Figure 7 are obtained by projecting the distribution of output features on the horizontal and vertical axes. Each axis is then discretized in twenty bins and the histograms simply show how many features are counted within each bin. The histograms approximate the output PDF. It can be observed from their asymmetries and long tails that the probability laws of response features (L/L_o) and (R/R_o) are not normal. It is well known that a Gaussian PDF propagated through a non-linear system such as this finite element simulation does not stay Gaussian.

Monte Carlo simulations are popular because of their simplicity and well-established convergence properties. Our simple illustration demonstrates the propagation of uncertainty from inputs to outputs and the estimation of the response's probability structure, from which statistics can be calculated. To guarantee convergence, however, large numbers of samples may be required in which case other sampling strategies—for example, stratified sampling such as the Latin Hypercube Sampling[9] or orthogonal arrays[10]—offer attractive alternatives. Screening experiments and analysis of variance can also be performed to understand which input variables or combinations of variables are most responsible for explaining an observed or simulated spread of responses such as the one pictured in Figure 6.[6,8]

5. INVERSE PROPAGATION OF UNCERTAINTY

The calibration of a model's parameters is a technique often employed to improve the accuracy of numerical simulations when compared to measurements. Parameter calibration is generally formulated as a deterministic inverse problem. A cost function is defined as the "distance" in some sense between measurements and predictions. The parameters of the model are then optimized to minimize the cost function. In the context of statistics where it is recognized that both calibration parameters and response features are random variables, a mechanism must be found to propagate uncertainty from the measurements back to the inputs. This is here referred to as the inverse propagation of uncertainty.

Although many formulations are possible, the discussion focuses on the Bayesian formulation implemented for the Taylor impact application.[11] Like in the deterministic case, a procedure for inverse propagation of uncertainty starts with the definition of a cost function. The main difference is to take into account that input parameters p and output features y are random variables, which generally implies that the cost function becomes a statistical test. In the case of Bayesian inference, the cost function is defined as the posterior probability that the parameters $p=(C_0;C_1;C_3;C_4;C_5;N)$ are correct given the available measurements y^{Test}. The measurements provide evidence against which predictions are tested. The posterior PDF is defined as the conditional probability law of the calibration variables p:

$$e^2 = -2\,\mathbf{log}\,(\mathbf{Prob}\ (\mathrm{p}\,|\,\mathrm{y}^{\mathrm{Test}})) \tag{3}$$

The Bayes Theorem states that the posterior probability *e* is the product of the likelihood function—likelihood to predict the measurements based on a given model—multiplied by the prior probability of *p*. Under the assumption of Gaussian probabilities, the likelihood function is defined as the mean square error between measurements and predictions. One advantage of Bayesian calibration is that the cost function obtains a closed-form expression:

$$e^2 = \sum_{k=1...N} \left(y_k^{Test} - y_k(p)\right)^T \left(\Sigma_{y_k}^{Test}\right)^{-1} \left(y_k^{Test} - y_k(p)\right) + (p - p_o)^T \left(\Sigma_p\right)^{-1} (p - p_o) \quad (4)$$

where the inverted matrices represent covariance matrices and the symbol p_o denotes the prior or nominal values of the calibration variables. The quantify *y* collects the output features, in our case $y=\{(L/L_o); (R/R_o)\}$. The summation represents a summation over potentially different experimental configurations, which is addressed in section 6 (see Figure 9). For simplicity, the covariance matrices of measurements and prior PDF are defined as diagonal matrices, meaning that the corresponding random variables are assumed to be uncorrelated.

In fact, the cost function defined in equation (4) based on the assumption of Gaussian probability distributions becomes the well-known chi-square statistical test that attempts to reject the null hypothesis that measurements y^{Test} and model predictions *y* are sampled from the same parent population.

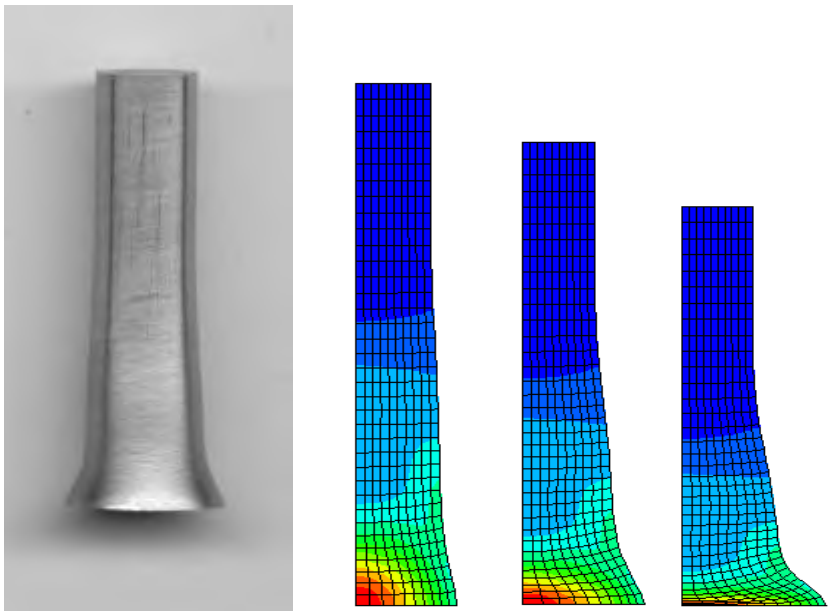

Figure 8: Test-analysis comparison for the Taylor impact experiment.
(Left: Measured profile. Right: Simulated profiles at 17 ms, 33 ms, and 50 ms after impact.)

The general procedure for calculating the chi-square statistics goes as follows. First, the impact simulation is analyzed for a given selection of calibration variables *p* and experimental configuration defined by the temperature and strain-rate parameters $(T;S_R)$. Next, the output

features (L/L_o) and (R/R_o) are calculated from the final deformed profile. These two values are then compared to the measurements. Figure 8 illustrates the test-analysis comparison. The whole procedure is repeated for each combination of temperature and strain-rate $(T;S_R)$. The chi-square (4) is calculated by accumulating the experiment's prediction error metrics.

Once a numerical procedure has been defined to compute the cost function, the calibration variables $p=(C_0;C_1;C_3;C_4;C_5;N)$ are optimized to search for the lowest possible chi-square. Because the prior and posterior PDF laws have been assumed, a deterministic optimization solver can be used. In the case where no evidence is available to suggest a particular distribution, the main difficulty becomes the estimation of a posterior PDF whose functional form is unknown. This can be resolved with a Markov Chain Monte Carlo optimizer that exhibits the attractive property of being able to sample an unknown probability law.[8,12]

Calibration Variable	Prior Mean	Posterior Mean	Prior Standard Deviation		Posterior Standard Deviation	
C_0	175.0	102.5	35.0	20.0% of mean	32.9	32.1% of mean
C_1	950.0	954.3	190.0	20.0% of mean	62.7	9.6% of mean
C_3	$3.0\text{x}10^{-3}$	$4.1\text{x}10^{-3}$	$0.6\text{x}10^{-3}$	20.0% of mean	$0.6\text{x}10^{-3}$	14.6% of mean
C_4	$8.5\text{x}10^{-5}$	$11.7\text{x}10^{-5}$	$1.7\text{x}10^{-5}$	20.0% of mean	$2.9\text{x}10^{-5}$	24.8% of mean
C_5	675.0	996.2	135.0	20.0% of mean	22.4	2.2% of mean
N	0.275	0.247	0.055	20.0% of mean	0.021	8.5% of mean

Table 1: Posterior mean and standard deviation values of the calibration variables.

	C_0	C_1	C_3	C_4	C_5	N
C_0	100.0%	-8.3%	37.2%	20.7%	48.8%	26.7%
C_1	-8.3%	100.0%	34.4%	31.1%	8.2%	13.0%
C_3	37.2%	34.4%	100.0%	80.2%	45.3%	-62.1%
C_4	20.7%	31.1%	80.2%	100.0%	27.1%	-46.6%
C_5	48.8%	8.2%	45.3%	27.1%	100.0%	86.0%
N	26.7%	13.0%	-62.1%	-46.6%	86.0%	100.0%

Table 2: Posterior correlation matrix obtained through Bayesian calibration.

The results presented in Table 1 and 2 are obtained by minimizing the cost function (4) with a conventional, gradient-based optimization solver. Table 1 lists the mean and standard deviation values of the calibration variable's prior and posterior distributions. Because the variables are assumed to be uncorrelated before calibration, the knowledge of their individual standard deviations suffices to define their joint PDF. Calibration, however, updates the entire covariance matrix and the resulting correlation coefficients are listed in Table 2.

It can be first observed from Table 1 that calibration tends to reduce the standard deviation values, except for the variables C_3 and C_4. It tends to indicate that information is learned from the test-analysis correlation exercise. Second, Table 2 shows no strong posterior correlation except for the pairs $(C_3;C_4)$ and $(C_5;N)$. (A significant correlation is arbitrarily defined here as a coefficient greater than 70% in Table 2.) It tends to indicate that the initial assumption of

uncorrelated variables p was correct. What causes the strong correlation between variables $(C_3;C_4)$ and $(C_5;N)$ is unknown but believed to be an artifact of the modeling assumptions. Further investigation, for example with a Markov Chain Monte Carlo optimizer, is warranted to study if the assumption of Gaussian priors and posteriors is judicious. In conclusion, the example illustrates how an optimized set of calibration variables $(C_0;C_1;C_3;C_4;C_5;N)$ that best predicts the outcome of physical experiments performed at various settings of temperatures and strain-rates $(T;S_R)$ can be obtained in the context of uncertainty propagation.

6. ASSESSMENT OF PREDICTIVE ACCURACY

In section 5, the problem of calibrating the model's parameters under uncertainty has been illustrated. Calibration, however, is only a tool in support of predictive accuracy assessment. A calibrated model is likely to provide small prediction errors in the neighborhood of the points used for calibrating its parameters but the question of adequacy in other regions of the operational space remains. In the remainder, an assessment of predictive accuracy for the plasticity model is illustrated. The illustration is purposely simplified for the sake of clarity.

The concept of operational space (or validation domain) introduced in section 3 is essential to the discussion. The operational domain, that is, the set of conditions for which a validated model of plasticity is sought, is defined by the combination of temperature T and strain-rate S_R, as shown in Figure 9. The red dots symbolize the settings $(T;S_R)$ at which physical experiments have been performed. (These are the same seven experiments used previously to calibrate the variables $(C_0;C_1;C_3;C_4;C_5;N)$, see section 5.) The question we would like to answer is the following one: "Given the uncertainty quantification work performed previously—that includes calibration under uncertainty—can the predictive accuracy of the plasticity model be estimated throughout the operational domain?"

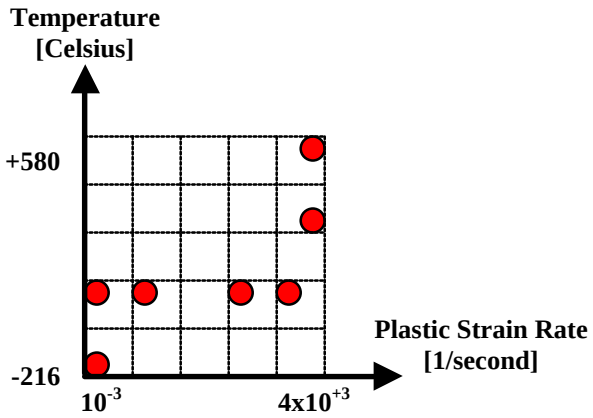

Figure 9: Location of the seven Taylor anvil impact tests in the operational space.

Answering this question is important for three main reasons. First, simulation-based decision-making requires validated models otherwise decisions might be based on erroneous information. To validate a model, its predictive accuracy must be estimated for all potential

operating environments or configurations, which is precisely the object of this study. Second, a formal protocol is required to make comparisons between competing models. The procedure described here provides a formalism to decide, for example, which one of the Zerilli-Amstrong or Johnson-Cook models is more suitable for the application of interest. Based only on test-analysis comparison at locations where experiments have been performed, it may happen that one model provides better accuracy than the other one. But a different model might be more suitable (less overall prediction error), or more robust to the uncertainty introduced by the modeling assumptions. An assessment of predictive accuracy throughout the entire design space captures such behaviors.

Third, being able to quantify the prediction error throughout the operational space provides a practical tool to decide where additional testing should take place. In environments where budget restrictions and time constraints prohibit an exhaustive testing campaign, it is critical to know which experimental configuration will provide the most useful information about the model's predictive accuracy. When budget becomes available for one more experiment, the configuration that could potentially improve the most the overall predictive accuracy of the model can be tested, instead of designing a confirmatory test. These three aspects of predictive accuracy assessment are briefly illustrated at the end of this section.

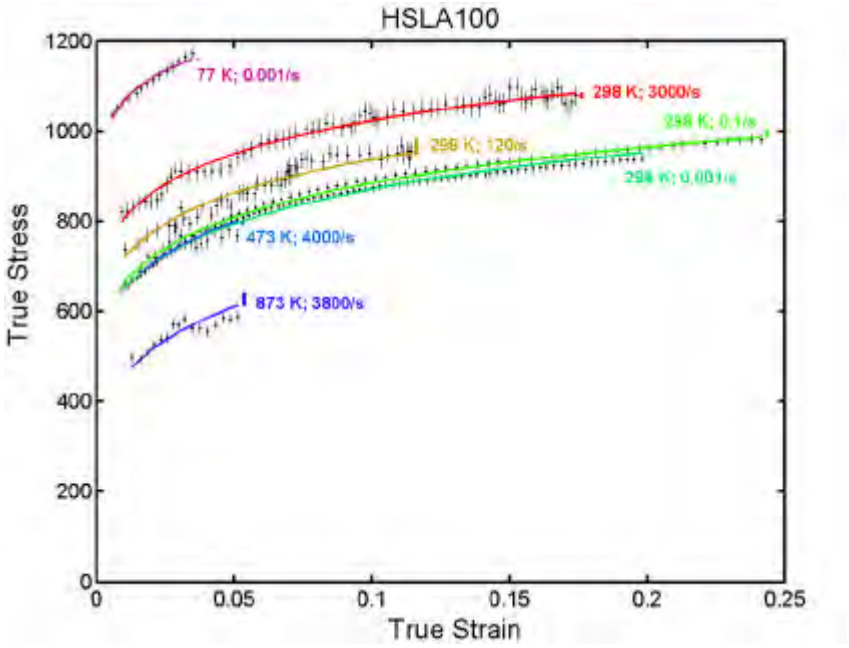

Figure 10: Comparison between Zerilli-Amstrong models and Hopkinson bar test data sets.

Figure 10 pictures the results of the inverse propagation of uncertainty from section 5. The solid lines represent the strain-stress curves predicted by the calibrated Zerilli-Amstrong model (1) for the seven configurations $(T;S_R)$ shown in Figure 9. Predictions are compared to experimental measurements obtained with the Hopkinson bar tests. The vertical bars represent

the experimental uncertainty. Simply speaking, Figure 10 compares the consistency obtained when material models are inferred from two different experimental procedures (Taylor anvil impacts and Hopkinson bar tests). In the remainder, the model's predictive accuracy is assessed based on this consistency.

It is emphasized that the choice of predictive accuracy metric is, to a great extent, application specific. Adequacy is here defined as an error between two material models. It could be argued that this is not a judicious choice because both material models can be biased, meaning that the resulting prediction metric could also be biased. This choice is made for simplicity. The reader will recognize that the successive steps of uncertainty quantification and predictive accuracy assessment proposed here can accommodate any error metric. The Mahalanobis distance is selected for its ability to account for experimental uncertainty:

$$e(T;S_R) = \left(y^{Test} - y(T;S_R)\right)^T \left(\Sigma_y^{Test}\right)^{-1} \left(y^{Test} - y(T;S_R)\right) \tag{5}$$

where y^{Test} represents the mean of the measured stress-strain curves and $y(T;S_R)$ is the model's prediction. As before, the covariance matrix is equal to a diagonal matrix—therefore neglecting the correlation between the data points—initialized with the variance coefficients associated to the measurements. The Mahalanobis distance extends to multivariate statistics the transformation $U = (X - \overline{X})/s_X$ that converts, for example, a normal random variable X with arbitrary mean and standard deviation into a zero-mean, unit-variance Gaussian U.

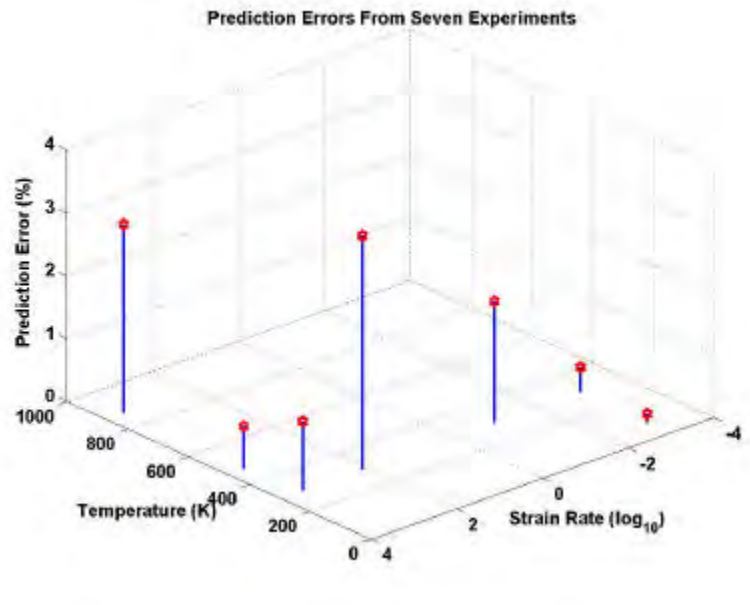

Figure 11: Seven Mahalanobis error values available for predictive accuracy assessment.

Figure 11 shows the Zerilli-Amstrong prediction errors computed at the seven settings in the two-dimensional space $(T;S_R)$ where Taylor anvil impacts have been performed. The

information pictured in Figure 11 is useful for local assessments of predictive accuracy because the prediction error is expected to vary smoothly as a function of coordinates $(T;S_R)$. Nevertheless, the question remains of knowing how well the model performs at locations where no physical experiment is available.

This question is addressed next by developing a metamodel of predictive accuracy. Essentially, an interpolation is performed by best fitting a statistical model to the seven data points. Such model is referred to as a metamodel because its development is not based on physical principles. Its only purpose is to capture an input-output relationship (between the inputs $(T;S_R)$ and the prediction error) and extrapolate beyond the available data with reasonable accuracy. For simplicity, polynomial models are adopted. Other choices might include fractions, Krigging models, exponential decays, neural networks, support vector machines, etc. Selecting the appropriate model form depends on factors such as the amount and quality of data available, prior knowledge of the input-output relationship, and the analyst's experience and preference. If the functional form of the metamodel is unknown, somewhat arbitrary choices are made and additional uncertainty is introduced. Naturally such uncertainty should be quantified and accounted for in the predictive accuracy assessment.

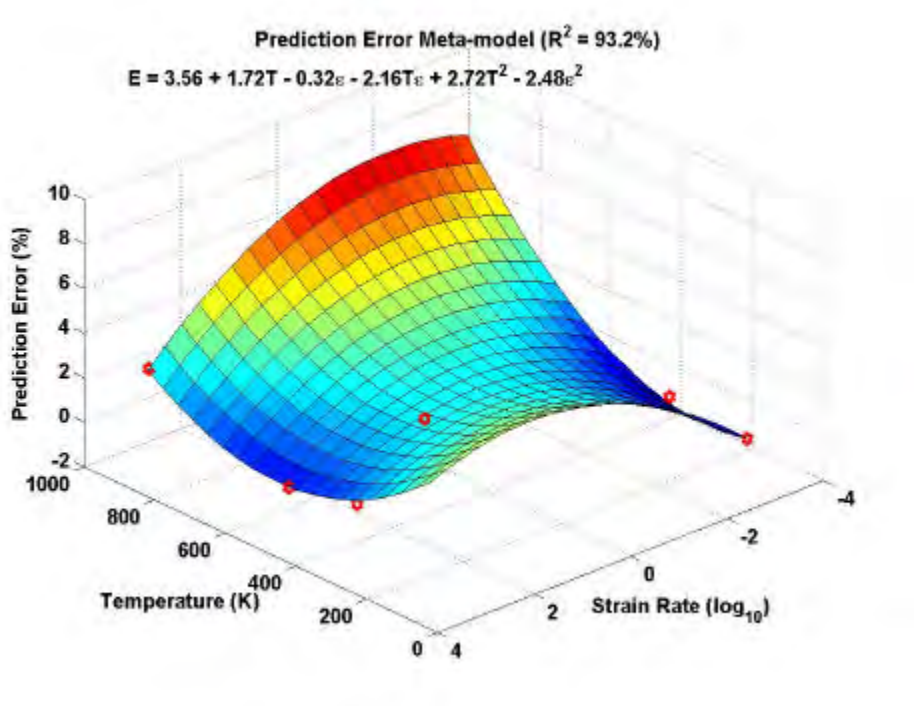

Figure 12: Metamodel that estimates the prediction error throughout the operational space.

Several polynomials are best fitted to the seven data points. They include linear, quadratic and cubic models as well as hybrid models that include some interaction effects between the parameters T and S_R but not others. Screening experiments and analysis of variance are useful

to provide guidance as to which effects should be included and which can be neglected. Metrics such as the mean square error and the R-square statistic—and its many variants—estimate the goodness-of-fit of a particular model. The model that best matches the prediction errors of Figure 11 without over-fitting the data is a quadratic polynomial:

$$e(T;S_R) = 3.56 + 1.72T - 0.32S_R - 2.16TS_R + 2.72T^2 - 2.48S_R^2 \quad (6)$$

The significance of the curve-fit to data is R^2 = 93.2% which is deemed acceptable. Obtaining this result includes the development of statistics for each one of the coefficients (not reported here). It leads to a statistical input-output relationship that can be sampled to propagate the uncertainty of the functional form itself. An estimation of predictive accuracy is illustrated in Figure 12 over the domain $(T;S_R)$. The response surface is simply obtained by calculating equation (6) over a 20-by-20 grid where T varies in the interval [-216; +580] degree-C and S_R varies in the interval [-3; +3.6] in units of $\log_{10}$(1/second).

The significance of Figure 12 is that it provides an estimation of the Zerilli-Amstrong model's adequacy everywhere within the operational domain, without having to perform any physics-based simulation. The validation domain can now be defined. For example, if a requirement of no more than 3% error must be satisfied, Figure 12 identifies the temperature and strain-rate settings where the plasticity model performs well. Once such assessment is available, competing models can easily be compared not just in terms of test-analysis accuracy but also throughout the entire operational domain. Response surfaces and contours of prediction accuracy provide a graphical way of assessing several models. The development of a model validation index is currently attempted at Los Alamos to further condense the information such as presented in Figure 12 into a single indicator that includes the entire uncertainty assessment. The third aspect previously mentioned is the selection of additional measurements. Figure 12 shows that model accuracy is the worst for combinations of high temperatures and low strain-rates. If it is deemed important that the model performs well in this region, we might want to perform more Taylor anvil impacts there.

7. CONCLUSION

An application of predictive accuracy assessment is presented for a non-linear numerical simulation. The application involves the development of temperature dependent and high-rate constitutive models of plasticity. Engineering plasticity models are developed to qualify the performance of pressure vessels subjected to impulse loading. Model validation involves a suite of Hopkinson bar and Taylor anvil impact tests, the quantification of uncertainty, and calibration of the model's parameters.

After having assessed the correlation between measurements and predictions, a statistical metamodel is developed to estimate the predictive accuracy when the temperature and strain-rate settings are varied. The end product is a polynomial that estimates the modeling error expected to be committed when the plasticity model is implemented to calculate the strain-stress curve at any combination of temperature and strain-rate. Obtaining such information

over an entire operating space is essential to answer questions such as: How appropriate is the model overall? Which one of several competing models is best? Which physical tests might be useful to improve the model's predictive accuracy?

Several key aspects are not discussed in this publication. They include the breakdown of total uncertainty, the quantification of modeling uncertainty and the validation index. The development of a validation index that captures the overall performance of a mathematical or numerical model—as well as the uncertainty associated to this assessment—is currently being pursued. The breakdown of total uncertainty refers to the necessity to make inferences of the modeling error from the total error obtained when measurements are compared to predictions. What makes this problem difficult is that numerous sources of uncertainty contribute to the total error. Examples include environmental variability, measurement error, grid convergence, time sampling, and algorithmic and solution convergence errors. Independent investigations of these potential contributions are necessary to eventually isolate the modeling error. Such studies are currently being performed for several applications.

Finally, modeling uncertainty refers to the uncertainty introduced by arbitrary choices such as functional forms and modeling assumptions. To be rigorous, an assessment of predictive accuracy must account for such uncertainty. Examples include the selection of data sets used to fit models, the form of the model itself, and variability of the statistical models. One roadblock is that many of these sources of uncertainty cannot be quantified probabilistically. Future work will attempt to quantify the variability of a predictive error assessment and study the robustness of model validation to the "unknown sources of uncertainty."

8. ACKNOWLEDGEMENTS

The author acknowledges the support of the U.S. Department of Energy's Advanced Scientific Computing (ASCI) program for engineering Verification and Validation. The guidance and assistance of Dr. Kenneth M. Hanson, from the Scientific Computing Group at Los Alamos, who provided the experimental data and performed some of the analyses, is greatly appreciated. Los Alamos National Laboratory is operated by the University of California for the United States Department of Energy under contract W-7405-ENG-36.

9. REFERENCES

[1] Adessio, F.L., Johnson, J.N., Maudlin, P.J., "The Effect of Void Growth on Taylor Cylinder Impact Experiments," *Journal of Applied Physics*, Vol. 73, No. 11, June 1993, pp. 7288-7297.

[2] Zerilli, F.J., Amstrong, R.W., "Dislocation Mechanics-based Constitutive Relations for Material Dynamics Calculations," *Journal of Applied Physics*, Vol. 61, No. 5, March 1987, pp. 1816-1825.

[3] **HKS/Abaqus™ Explicit User's Manual**, Version 5.8, Hibbitt, Karlsson & Sorensen, Pawtucket, RI, 1998.

[4] Hoover, C.G., De Groot, A.J., Sherwood, R.J., "ParaDyn: A Parallel Nonlinear, Explicit, Three-Dimensional Finite-Element Code for Solid and Structural Mechanics—User Manual," *LLNL UCRL-MA Report*, Lawrence Livermore National Laboratory, Livermore, CA, 2000.

[5] Maudlin, P.J., "On the Modeling of the Taylor Cylinder Impact Test for Orthotropic Textured Materials: Experiments and Simulations," *International Journal of Plasticity*, Vol. 15, 1999, pp. 139-166.

[6] Cundy, A.L., Schultze, J.F., Hemez, F.M., Doebling, S.W., Bingham, D., "Variable Screening in Metamodel Design for a Large Structural Dynamics Simulation," *Proceedings of the 20th SEM International Modal Analysis Conference*, Los Angeles, CA, Feb. 4-7, 2002.

[7] Hemez, F.M., Wilson, A.C., Doebling, S.W., "Design of Computer Experiments for Improving an Impact Test Simulation," *Proceedings of the 19th SEM International Modal Analysis Conference*, Kissimmee, FL, Feb. 5-8, 2001.

[8] Kerschen, G., Golinval, J.-C., Hemez, F.M., "Bayesian Model Screening for the Identification of Non-linear Mechanical Structures," *Journal of Vibration and Acoustics*, 2003, accepted for publication.

[9] McKay, M.D., Beckman, R.J., Conover, W.J., "A Comparison of Three Methods For Selecting Values of Input Variables in the Analysis of Output From a Computer Code," *Technometrics*, Vol. 21, No. 2, 1979, pp. 239-245.

[10] Hedayat, A.S., Sloane, N.J.A., Stufken, J., **Orthogonal Arrays: Theory and Applications**, Springer-Verlag Publishers, New York, NY, 1999.

[11] Hanson, K.M., "A Framework for Assessing Uncertainties in Simulation Predictions," *Physica D*, Vol. 133, 1999, pp.179-188.

[12] Carlin, B.P., Chib, S., "Bayesian Model Choice via Markov Chain Monte Carlo," *Journal of the Royal Statistical Society Series B*, Vol. 77, 1995, pp. 473-484.