

UNCERTAINTY QUANTIFICATION OF SIMULATION CODES BASED ON EXPERIMENTAL DATA

Kenneth M. Hanson¹ and François M. Hemez²

Los Alamos National Laboratory
Los Alamos, NM 87545

We present an approach for assessing the uncertainties in simulation code outputs in which one focuses on the physics submodels incorporated into the code. Through a Bayesian analysis of a hierarchy of experiments that explore various aspects of the physics submodels, one can infer the sources of uncertainty, and quantify them. As an example of this approach, we describe an effort to describe the plastic-flow characteristics of a high-strength steel by combining data from basic material tests with an analysis of Taylor impact experiments. A thorough analysis of the material-characterization experiments is described, which necessarily includes the systematic uncertainties that arise from sample-to-sample variations in the plastic behaviour of the specimens. The Taylor experiments can only be understood by means of a simulation code. We describe how this analysis can be done and how the results can be combined with the results of analyses of data from simpler material-characterization experiments.

INTRODUCTION

Our approach to understanding the uncertainties in simulation code predictions combines the principles of physics and Bayesian analysis. The focus is on understanding and quantifying the uncertainties in the simulation-code submodels and the numerical errors introduced in solving the dynamical equations. Bayesian analysis provides the underpinning for quantifying the uncertainties in models inferred from experimental results, which possess their own degree of uncertainty. The aim is to construct an uncertainty model that is based on inferences drawn from comparing the code's predictions to relevant experimental results. In the context of the proposed framework, it is possible to design new experiments that can best provide data for reducing prediction uncertainty [1, 2].

The sources of uncertainty in a simulation-code prediction of the outcome to a hypothesized physical situation include a) uncertainties in the dynamical equations, b) uncertainties in submodels that describe material properties, c) numerical-solution errors, and d) uncertainties in the initial and boundary conditions of the physical situation being simulated. All of these sources of uncertainty need to be taken into account in summarizing the uncertainties in simulation predictions. They must also be taken into account when analyzing the comparison between experimental data and

simulation code output in order to make inferences about the physics submodels.

We demonstrate our approach by analyzing a set of material-characterization experiments. We indicate how a similar data-fitting can be used to analyze the results of a Taylor impact test. In that test, a metal cylinder is propelled into a rigid wall, and the profile of the deformed cylinder is measured. A simulation code is needed to analyze the Taylor data. The goal is to be able to predict how well the code should be able to predict the results of the next Taylor test involving the same material. Fundamental to our approach is a thorough uncertainty assessment, performed as part of the analysis of each experiment, and the accumulation of information about the physics models gained by analyzing a hierarchy of experiments..

INFERENCE ABOUT PHYSICS MODELS FROM EXPERIMENTS

Physics and engineering simulation codes are based on our understanding of the relevant physics processes involved in the phenomenon being simulated. Such codes are usually constructed out of building blocks, each of which is supposed to account for a specific behavior, for example, the behavior of materials are described by constitutive models. General physical principles provide the glue. Our goal is to make the simulation code robust so as to be able to handle a variety of physical scenarios.

¹ Methods for Advanced Scientific Simulation Group (CCS-2), 505-667-1402 (Voice), kmh@lanl.gov (E-mail)

² Engineering Analysis Group (ESA-EA), 505-665-7955 (Voice), hemez@lanl.gov (E-mail)

Figure 1 schematically depicts the basic components of a simulation code that is meant to predict dynamic phenomena. To do a calculation, the simulation code requires full and complete specification of the physical situation. Thus, the code requires as input the properties of the materials involved over the range of conditions encountered in the experiment and the initial state of the system. Even though the diagram shows input from only one model of material behavior, more complicated situations could obviously require the input of models for several materials. The boundary conditions may be thought of as being included in the initial-state specification. The goal of the simulation code is to predict the behavior of the system at later times. To avoid extra complications, we consider only deterministic phenomena in which the state of the system at later times is uniquely determined by the initial state and material behavior.

Our goal is to compare the output of a simulation code to experimental measurements. If we think about the experimental counterpart to the simulation depicted in Fig. 1, we realize that each aspect of the experimental situation that is uncertain is related to an uncertainty in each input to the code. In our approach, uncertainties are quantitatively described in terms of probability density functions (pdfs). We should, therefore, associate with each input variable a pdf describes our uncertainty in its value. The properties of the material may be uncertain for a number of reasons. For example, the exact composition and preparation of the material used in the experiment may not be known. Uncertainties can also arise because of variability in the material characteristics, such as variations in grain structure that result from specimen preparation, which may not be adequately enough controlled or measured to be neglected. Braces are used in Fig. 1 to indicate that uncertain quantities.

In specifying the input for a simulation meant to match an experimental situation, it is important to specify the degree of uncertainty in the initial state, described in terms of its geometry, initial velocities, materials, etc. To obtain the most useful results, it is important in designing the experiment to minimize the significant sources of uncertainty, a concept that Oberkampf includes under his definition of “validation experiments.” [3]

It is relatively easy to propagate uncertainties in inputs through a simulation code, a process that is called forward propagation of uncertainties. One way to think about propagating uncertainties through a simulation code is in terms of the Monte Carlo (MC) technique. The parameters that describe the initial conditions and the material behavior are drawn from the pdfs that

describe our uncertainties in them. The variability in the output states of the simulation obtained for the set of random input values represents the uncertainty in simulation predictions.

In the process of comparing experimental results with simulations, we wish to conduct inference about the simulation code, namely draw conclusions about the submodels. The usually approach for matching a simple model to experimental data is to employ least-squares or chi-squared fitting of the model to the data. We will argue later that this also feasible in the context of simulation codes.




Figure 1. The simulation code calculates the dynamic state of a system from a specified initial state and a submodel that defines the behavior of the material involved, represented here by the parameter vector α . Uncertainties in the initial state and material models, which are designated by braces, propagate through the simulation code to yield uncertainties in the predictions made by the simulation.

The Bayesian approach to analysis is perfectly suited for conducting inference about models [1, 3]. Fundamental to the Bayesian approach is the use of pdfs to quantify uncertainties. Bayes theorem gives the probability for a parameter vector, designated by \mathbf{a} , as

$$p(\mathbf{a} | \mathbf{d}, I) \propto p(\mathbf{d} | \mathbf{a}, I) p(\mathbf{a} | I), \quad (1)$$

where the vector \mathbf{d} stands for the data. The first factor on the right is called the likelihood, which is the probability of the data, given the parameters and whatever relevant background information I is available, for example, which physical model is appropriate for the physical situation. The second factor is called the prior. It expresses the uncertainty we have in the parameters \mathbf{a} before we consider the new data \mathbf{d} . Bayes theorem provides the appropriate means for updating our knowledge about the parameters. The expression (1) is called the posterior (probability

distribution) for the parameters because it refers to our state of knowledge after the analysis is done.

Prior information is not incorporated in many analyses, partly because of a lack of prior knowledge, and partly to keep the analysis objective. In these analyses, the prior in Bayes theorem (1) is effectively taken to be uniform. Consequently, the posterior distribution is proportional to the likelihood, and Bayesian analysis just amounts to the familiar likelihood analysis. We stress the importance of including a full uncertainty analysis in any complete analysis. Experimental uncertainties result in parameter uncertainties, which must be quantified for the analysis results to be useful.

HIERARCHY OF EXPERIMENTS

To best understand the submodels contained in a simulation code, one should employ a hierarchy of experiments [1, 3], as schematically shown in Fig. 2. The level of integration is based on the complexity of the physical phenomenon involved, or roughly the number of physics models needed to describe each experiment. Ideally, one learns about the individual physics models used in a simulation code through basic experiments, which are designed to isolate and characterize each physics model. Then by doing more and more complicated experiments, one can learn more about the physics models by extending the range of physical conditions probed and also about possible interactions between different individual models, e.g., involving different materials. The ultimate goal is to combine the results from many (or all) experiments to reduce as much as possible the uncertainties in the models and their parameters. In the process, one should check that the models are consistent with all experiments. Any deviations from what is expected on the basis of the models may indicate that something is wrong with them.

In Fig. 2, experiments 1, 3, and 4 are considered to be basic experiments because they involve only the individual material models, represented by the parameter vectors α , γ , and δ , respectively. Experiment 2 is partially integrated because it involves two models, represented by α and β . With prior information about α coming from experiment 1, the analysis of the data from experiment 2 will provide information about β , as well as further refining what is known about α . Experiment 5 is fully integrated in this scenario because it involves all four models [1]. The nomenclature used here is slightly different from that used by Oberkamp [3], but the intent is similar.

Figure 2 depicts how knowledge about the physics models flows from the analysis of experiments at one

level to the next. Reference [1] suggests that this knowledge may be quantitatively specified in terms of uncertainty distributions in the parameter vectors α , β , γ , and δ . Bayes law may be used in each bubble to update our knowledge of the physics models. Thus, Fig. 2 is a probabilistic network in which probabilities flow between nodes and are updated in each node on the basis of new experimental information. From this perspective, it appears feasible to systematically incorporate the information from all the experiments into a consistent set of parameters for the models involved. Because this approach is grounded in quantitative uncertainty assessment, the uncertainties in simulation predictions may be calculated through the Monte Carlo procedure described above.




Figure 2. Diagram of the conceptual process for combining information from several experiments conducted for a hierarchy of experiments involving various levels of integration. Four underlying physics models, represented here by the parameter vectors α , β , γ , and δ , are required to account for the behavior of the fully integrated system.

AN EXAMPLE RELATED TO CHARACTERIZING THE PLASTIC BEHAVIOR OF STEEL

As a demonstration of how the above notions can be employed to update our knowledge of a physics model, we consider the following problem. The engineering task is to design a vessel that can adequately contain an experiment involving chemical explosives. For this project, the vessel is to be constructed from a high-strength steel, HSLA 100. The first requirement is that the vessel must be able to handle a specified impulse in internal pressure without rupturing or leaking. For this discussion, we will focus on a secondary requirement, that shrapnel fragments produced in the explosion will not penetrate the vessel-wall. Our ultimate goal is to be able to predict the depth of penetration of a shrapnel fragment with a specific mass and impact velocity. Naturally the prediction must be accompanied by

estimates of its uncertainties. The estimated uncertainty in this prediction should be linked to the safety factor that needs to be built into the vessel design.

Consistent with the preceding discussion, our approach to solving this problem is to determine what experiments are needed to characterize the stress-strain relationship for plastic flow of the metal. In the analysis, one would track the uncertainty through the analysis of a series of experiments to obtain an uncertainty estimate for the final prediction. The variables to consider include temperature, strain rate, variability in material behavior caused by composition, and processing.

To be more specific, appropriate basic experiments might consist of quasi-static tests and Hopkinson-bar experiments. In the quasi-static experiments, specimens of the material are squeezed (or pulled), and the stress required to produce a measured strain is measured. The measurements in these experiments are typically very accurate (1%) and quite reproducible when the specimens are carefully selected to make sure their properties are uniform from sample to sample.

In Hopkinson-bar experiments, a shock wave is passed through a small disc-shaped specimen, and the length of the specimen is measured as a function of time. Knowing the properties of the apparatus, it is possible to determine the stress-vs.-strain curve of the material at moderately high strain rates (for example, up to around 10^3 s^{-1}). Hopkinson-bar data typically possess larger errors than quasi-static data.

Because the stress-strain behavior of metals like HSLA 100 are known to depend on the rate at which the strain develops, the Hopkinson-bar experiments are needed to provide data that are more pertinent to the target application. Even so, the strain rates of the target application are higher than occur in Hopkinson-bar tests. This consideration argues for conducting an experiment of intermediate complexity that provides information about very high strain rates. The well-known Taylor impact test serves the purpose well; strain rates of 10^5 s^{-1} are routinely achieved. Thus, the Taylor test bridges the gap between moderate and high strain rates.

To complete the hierarchy of experiments and to check the ability of the simulation code to accurately predict shrapnel penetration, it clearly would be desirable to conduct experiments that are as close to the intended application as possible. Other basic- and intermediate-level experiments may potentially be required to describe the penetration of steel plate by a fragment, such as material fracture and the equation of state of the

material as it heats and melts. Furthermore, a different type of simulation code, for example, a Eulerian code, will most likely be required to simulate the fragment penetration phenomenon. The common thread is the material model, which evolves throughout the process of analyzing the hierarchy of experiments.

MATERIAL-CHARACTERIZATION EXPERIMENTS

The first phase of the project is to conduct the basic material-characterization experiments consisting of quasi-static tests and Hopkinson-bar experiments, as described above.

Figure 3 shows the results of seven material characterization experiments conducted at LANL on HSLA 100 samples. Because large plastic deformations and strain rates in excess of 10^5 per second are expected in the target application, plasticity and strain-rate dependence must be included in the stress-strain model.




Figure 3. Comparison between data from material-characterization experiments for a variety of temperatures and strain rates and the ZA model fit to the data. The vertical bar to the right of each curve indicates the size of the estimated systematic offset for that curve.

We use the Zerilli-Armstrong (ZA) model for rate-dependent plasticity:

$$\sigma_p = \alpha_1 + \alpha_3 \varepsilon_p^{\alpha_6} + \alpha_2 \exp \left[\left(-\alpha_3 + \alpha_4 \log \frac{\partial \varepsilon_p}{\partial t} \right) T \right], \quad (2)$$

where ε_p and σ denote the equivalent plastic strain and resulting stress, respectively, and T is the temperature at the start of the experiment. The parameters α_i are material specific, and once more, they may depend on

how the material has been processed. The ZA formula is based on a dislocation-mechanics model, which may not hold for all materials or experimental conditions.

From Eq. (2), we see that only two parameters (α_5 & α_6) specify the basic dependence of the stress as a function of strain. The four remaining parameters provide an additive offset as a function of temperature and strain rate. While the Zerilli-Armstrong model may not capture all aspects of the behavior of the data, we will assume that it is one that has been specified for this project. Our job is to determine the best parameters to use for HSLA 100 and their uncertainties, and then predict the uncertainties in simulations based on that model.

LIKELIHOOD ANALYSIS

Before describing the details of the analysis that has been done using the data displayed in Fig. 3, let us briefly review a standard approach to fitting a nonlinear model to data by the minimum chi-squared method [5].

The likelihood is the probability of the measured data for a specified parameter vector \mathbf{a} . If the errors in the measured data \mathbf{d} are Gaussian distributed and independent, the likelihood is given by

$$p(\mathbf{d} | \mathbf{a}) \propto \exp \left\{ -\frac{1}{2} \sum_i \frac{[d_i - y_i(\mathbf{a})]^2}{\sigma_i^2} \right\}, \quad (3)$$

where σ_i is the expected rms deviation of the measurement d_i . The corresponding value given by the model for a specified parameter set \mathbf{a} is designated by y_i .

We recognize the sum in the exponential in Eq. (3) as the familiar chi squared, χ^2 , which is often used to quantify the discrepancy between measurements and a model. The parameters that best fit the data are typically taken as those that maximize the likelihood, or, equivalently, minimize χ^2 . A standard tactic is to expand the model value y_i at the value of the independent variable x_i in terms of a Taylor series,

$$y_i = y_i(x_i, \mathbf{a}) = y_i^0 + \sum_j \left. \frac{\partial y_i}{\partial a_j} \right|_{\mathbf{a}^0} (a_j - a_j^0) + \dots \quad (4)$$

The complete set of derivatives make up the so-called the Jacobian matrix, \mathbf{J} . By dropping higher-order terms, chi-squared can be approximated as a quadratic function around its minimum,

$$\chi^2(\mathbf{a}) = \frac{1}{2} (\mathbf{a} - \hat{\mathbf{a}})^T \mathbf{K} (\mathbf{a} - \hat{\mathbf{a}}) + \chi^2(\hat{\mathbf{a}}), \quad (5)$$

where $\hat{\mathbf{a}}$ is the parameter vector at the minimum in chi-squared, and \mathbf{K} is the curvature matrix of chi squared at $\hat{\mathbf{a}}$, which is commonly called the Hessian. It can be written in terms of the Jacobian, evaluated at $\hat{\mathbf{a}}$, as

$$[\mathbf{K}]_{jk} = \left. \frac{\partial^2 \chi^2}{\partial a_j \partial a_k} \right|_{\hat{\mathbf{a}}} = \mathbf{J} \mathbf{J}^T. \quad (6)$$

As noted above, when a flat prior is assumed, the posterior is proportional to the likelihood. In the quadratic approximation for χ^2 given in Eq. (5), the posterior is a Gaussian

$$p(\mathbf{a} | \mathbf{d}) = \frac{1}{\det[\mathbf{C}] (2\pi)^{n/2}} \times \exp \left[-\frac{1}{2} (\mathbf{a} - \hat{\mathbf{a}})^T \mathbf{C}^{-1} (\mathbf{a} - \hat{\mathbf{a}}) \right] \quad (7)$$

which is written in a way to explicitly display the covariance matrix, \mathbf{C} :

$$\text{cov}(\mathbf{a}) = \langle (\mathbf{a} - \hat{\mathbf{a}})(\mathbf{a} - \hat{\mathbf{a}})^T \rangle \equiv \mathbf{C} = 2\mathbf{K}^{-1}. \quad (8)$$

The covariance matrix describes the degree of correlation among the uncertainties in the parameters. The analysis of nearly every type of experiment leads to off-diagonal terms in the covariance matrix, which must be stated for a full specification of the uncertainties in the parameters. The off-diagonal elements of the covariance matrix are often expressed in terms of the correlation coefficients, $\rho_{ij} = C_{ij} / \sqrt{C_{ii} C_{jj}}$.

ANALYSIS OF MATERIAL-CHARACTERIZATION EXPERIMENTS

We now outline the analysis of the material-characterization experiments to estimate the ZA parameters and their uncertainty. The approach is, in a sense, a straightforward chi-squared analysis, but it has some noteworthy aspects, for example, the inclusion of systematic uncertainties.

In looking at the unedited data sets from the seven test results shown in Fig. 3, the first conclusion is that the data are inconsistent with the simple dependences contained in the ZA model. Specifically, at high and low temperatures, the data show a different stress-strain behavior than at room temperature. As stated above, one of the requirements of the project is to employ the

ZA model. To make do with the ZA model, the wise thing is to avoid trying to incorporate inconsistent data in fitting the ZA model. Thus, from the high- and low- T experiments, we include only the low strain data (less than about 0.05) in determining the ZA parameters. At temperatures that differ significantly from the target operating conditions of room temperature and high strain rate, a different material-damage model seems to be in effect at large strains, so these are avoided.

A well-known feature of Hopkinson-bar experiments is the presence of wiggles in the measured curve at low strains, which is particularly evident at strains of 0.02 and below (not included in Fig. 3), but also observable at higher strains. These oscillations, caused by reflections of shocks within the apparatus, make portions of the Hopkinson data less useful. Taking the variability of the measurements into account and considering the inherent difficulties in the experiments, we assign standard (rms) error bars to the Hopkinson data of 2%. Because of their better reproducibility for carefully selected specimens, we assign standard error of 1% to the quasi-static data.

Difficulties in matching the strain-rate dependence of the data sets at fixed (room) temperature have prompted us to consider sources of systematic uncertainties. Auxiliary experiments on specimens taken from various locations in a 5.08-cm-thick plate of HSLA 100 indicate that the stress-strain curves can systematically differ from each other by an additive constant. From a limited number of tests, we estimate the rms variation in the offset between different samples is approximately 20 MPa.

We include this sample-to-sample variability in the analysis by treating it as a systematic uncertainty. Because the observed differences between different samples amount to a small vertical shift of the curves, it is appropriate to incorporate them in terms of an additive offset parameter for each curve. This effectively adds seven more parameters to the model, which need to be estimated as well as the six α 's in the ZA model. To include this systematic effect in a full analysis of the seven data sets, the appropriate expression for the minus-log-likelihood is

$$-\log p(\mathbf{a} | \mathbf{d}, I) = \varphi(\mathbf{a}) = \frac{1}{2} \sum_k \chi_k^2 + \frac{1}{2} \sum_k \frac{\Delta_k^2}{\sigma_k^2}, \quad (9)$$

where the index k identifies the data set. The first term is a sum of χ^2 values for each data set at a specific temperature and strain rate. The second sum represents the prior probability for the offset parameters Δ_k , and

σ_k is the rms deviation of the prior distribution on Δ_k , taken to be 20 MPa in the current analysis. This term is needed to constrain the vertical offset of the curves. Without it, parameters α_1 through α_4 would be indeterminate.

The above model is fit to the data shown in Fig. 3 using the general approach described earlier to minimize the function given in Eq. (9). The Jacobian (sensitivity) matrix is estimated at each optimization iteration by the straightforward method of finite differences.

The final fit to the data is shown in Fig. 3. The vertical bars at the end of each curve display the fitted value for the offset of that curve. The seven offset values range from -37 MPa to 11 MPa, with a mean of -9 MPa and an rms value of 19 MPa, which is reasonably consistent with their assumed rms deviation, 20 MPa. The ZA parameters obtained from the fit, and their rms uncertainties, are given in Table 1. As important as the uncertainties in the individual parameters are their correlation coefficients, which are presented in Table 2. Use of the rms errors without consideration of the correlation coefficients would seriously misrepresent the results of this analysis, as will be demonstrated shortly.

Standard Monte Carlo techniques can be used to draw random ZA parameter vectors from their uncertainty distribution specified in Tables 1 and 2. Figure 4 shows a plot of the curves that result from 20 such random draws for the conditions of room temperature and high strain rate. From the comparison of these curves to the Hopkinson-bar measurements, we conclude from this plot that the parameters and their uncertainties plausibly represent the data for the conditions of the target application.

Table 1. The ZA model parameters and their standard errors, estimated from the data shown in Fig. 3. The units correspond to those when the variables in Eq. (2) are measured in MPa, deg K, and seconds.

α_1	103 ± 33
α_2	954 ± 63
α_3	0.00408 ± 0.00059
α_4	0.000117 ± 0.000029
α_5	996 ± 22
α_6	0.247 ± 0.021

Table 2. Correlation coefficient matrix for the ZA model parameters obtained with the fit to the data shown in Fig. 3. The covariance matrix is estimated using Eq. (6).

	α_1	α_2	α_3	α_4	α_5	α_6
α_1	1	-0.083	0.372	0.207	-0.488	0.267
α_2	-0.083	1	0.344	0.311	0.082	0.130
α_3	0.372	0.344	1	0.802	0.453	-0.621
α_4	0.207	0.311	0.802	1	0.271	-0.466
α_5	-0.488	0.082	0.453	0.271	1	-0.860
α_6	0.267	0.130	-0.621	-0.466	-0.860	1

The above procedure is highly recommended whenever one does model fitting. In effect, it closes the loop between the data and their uncertainties and inferences made in terms of parameters.

A major contribution to the uncertainty in the stress-versus-strain relation arises from the sample-to-sample variability, taken into account in the above analysis as a systematic error. The importance of specifying and including correlations in the use or treatment of uncertainties is graphically demonstrated in Fig. 5. The stress-versus-strain curves are derived by randomly drawing 20 ZA parameter vectors from the uncertainties in the individual parameters stated in Table 1, but neglecting the correlation coefficients specified in Table 2. The variation among the curves is now unreasonably large compared to the error bars in the data.




Figure 4. A set of plausible stress-strain curves for room temperature and high strain rate obtained by drawing Monte Carlo samples from the uncertainty distribution in the ZA parameters as derived from the data shown in Fig. 3.

The above description of the analysis of the HSLA 100 is meant to emphasize the importance expertise plays in what seems to be a straightforward analysis. The selection of which data to include in the fit is based on the conditions of the target application, and an understanding of the plastic behavior of metals. Furthermore, the inclusion of systematic effects involves considerations that go beyond the data themselves.

The systematic uncertainties can in principle be reduced through careful selection or preparation of the samples. However, in the end, the steel that is used to manufacture the containment vessel will most likely not be so carefully selected. The sample-to-sample variation in material properties would have to be taken into account at that stage. The uncertainties must be included in the overall process in a careful and systematic way, and need to be documented throughout




Figure 5. A set of stress-strain curves for room temperature and high strain rate obtained by drawing Monte Carlo samples from the same uncertainty distribution used to generate Fig. 4, except for **neglecting** the correlations given in Table 2.

the process to ensure that all uncertainties are included in their proper place, but not any more than necessary.

TAYLOR IMPACT TEST

The Taylor impact test represents an experiment of an intermediate level of complexity in the hierarchy of experiments chosen to elucidate the material model for HSLA 100. In the Taylor impact test, a cylindrical sample of material is propelled into a fixed, rigid surface, as depicted in Fig. 6. Taylor tests are often performed to confirm the plastic behavior of a material under severe strain conditions. Extremely high plastic strains and strain rates occur at the crushed end of the rod, resulting in severe local deformation [6, 7]. The experimental data usually consist of measurements of its final deformed profile.

Simulations of the Taylor impact test have been performed using Abaqus, a general-purpose finite-element modeling and analysis package, which employs explicit time integration [8]. The modulus of elasticity is taken to be $E = 310$ GPa and the material density, $\rho = 7,750$ kg/m³. Verification tests were performed to make sure that solution errors are sufficiently small.

We use the Monte Carlo technique to illustrate the forward propagation of uncertainty through the simulation code. We consider a situation in which the only parameters that are uncertain are those in the above ZA stress-strain model, stated in Eq. (2). We employ the full uncertainty distributions derived in the previously-described analysis of the material characterization experiments.




Figure 6. In the Taylor impact test, a cylinder is driven with high velocity into a fixed, rigid plate (left) producing a significant permanent deformation in the cylinder after rebound (right).

As demonstrated in Fig. 4 for the ZA model, the Monte Carlo process consists of sampling parameter values from a specific uncertainty distribution and running the simulation code for each set of values. Based on the ZA parameters and their uncertainties obtained in the above analysis of the material-characterization data, Fig. 7 shows the results of following this process to generate seven final deformed profiles of a cylinder simulated for an impact velocity of 260 m/s. The observed variation in shape indicates the uncertainty in the profile that is produced by the uncertainty in the plastic stress-strain relation derived in the preceding analysis. A quantitatively accurate characterization of the uncertainties in the simulation output would require many more realizations.

In this exercise, we have ignored the uncertainties in the experimental set up. These could readily be included in the Monte Carlo technique. For example, to include the uncertainty in the impact velocity, it could be included as an additional random variable, with its uncertainty described by its own pdf. By drawing random samples for the velocity from its pdf and including these random values in the simulation process, the effects of uncertain




Figure 7. Seven representative Taylor test profiles predicted for an impact velocity of 260 m/s, obtained with a Monte Carlo technique employing uncertainties come from the analysis described above. The variability in these profiles indicates their uncertainty. The data points represent hypothetical measurements of the deformed-cylinder profile from a corresponding experiment.

velocity are easily included in the Monte Carlo uncertainty calculation. Other uncertainties in the This approach to assessing the consequences of uncertainties in the experimental conditions provides a good basis for designing validation experiments, for example, as discussed in Ref. [9] for structural dynamics applications.

ANALYSIS OF TAYLOR DATA USING A SIMULATION CODE

The data from a Taylor experiment may be analyzed in much the same way as was done above for the material-characterization experiments. Systematic experimental uncertainties, for example, in the impact velocity, may be included in a way similar to that used above to include the sample-to-sample variations. A contribution similar to the second term in Eq. (9) is necessary to account for the systematic offset for the specific sample used in the Taylor test. One viable approach to chi-squared minimization is to use the same methods as described above. When the simulation code is treated as a black box, the Jacobian matrix may be estimated by the method of finite differences. When the simulation code is available, however, the more sophisticated method of automatic differentiation may prove useful [10, 11].

Figure 2 depicts the general procedure that we have in mind to update our knowledge of model parameters with the analysis of each new experiment. Having analyzed the material-characteristic experiments, which are considered to be basic experiments in Fig. 2, we may use those results as priors to the analysis of the more complex Taylor experiments. Bayes theorem, given in Eq. (1), provides the proper means for combining the prior pdf from the first analysis with the likelihood of the subsequent Taylor analysis [1]. The uncertainties from the above analysis may be included by adding to the expression given in Eq. (9) a term to represent the prior, $\frac{1}{2}(\mathbf{a} - \hat{\mathbf{a}}_1)^T \mathbf{C}_1^{-1} (\mathbf{a} - \hat{\mathbf{a}}_1)$, where $\hat{\mathbf{a}}_1$ is the ZA parameter vector estimated in that analysis and \mathbf{C}_1 is the estimated covariance matrix.

When this process is employed to set the values of the model parameters, it is often called calibration, which is different from what we are proposing. By basing this parameter-updating process on Bayes theorem and quantitative uncertainty estimates, the process becomes one of inference [10]. In a sense, the Monte Carlo technique for estimating uncertainties in simulation-code output described above is reversed; the uncertainties in the parameters are determined from the combined uncertainties in the measurements and the

effects on the simulation of uncertainties in experimental set up. Since the inference procedure involves determining the uncertainties in model parameters, it provides the means for predicting the uncertainty in simulation output in other physical scenarios. Further explanation of the details of the process are presented in Ref. [1].

CONCLUSION

We have presented an approach to understanding the uncertainties in predictions made by simulation codes. The first place to look for uncertainties in simulation code output is the uncertainties in the physics models that are incorporated in the simulation code. All sources of output uncertainties must be considered, including the numerical implementation of the physics models, especially the finite size of the finite elements or mesh, and aspects of the physics that are not accounted for. This framework highlights the importance of conducting experiments that are thoughtfully designed to provide results that can be quantitatively compared to simulation codes. A hierarchy of experiments should be included in the overall scheme to provide validation of the physics models in a variety of circumstances and over a wide range of experimental conditions.

The focus here has been on the parameters in the ZA model. The Bayesian methodology can address other questions, for example, comparison of two or more competing models to decide between them. It is generally applicable to answering all questions that one might pose about models.

ACKNOWLEDGMENTS

We gratefully acknowledge Shuh-Rong Chen's assistance, both by supplying the experimental data presented here and by providing considerable guidance in regard to material-science issues. The following people have kindly shared their insights with us: Rusty Gray, Paul Maudlin, Larry Hull, Mark Anderson, Bill Oberkampf, Tim Trucano, Scott Doebing, Mike McKay, Kathy Campbell, C. Shane Reese, and Tom Duffey. This work was supported under US-DOE Contract W-7405-ENG-36.

REFERENCES

- [1] K. M. Hanson, "A framework for assessing uncertainties in simulation predictions," *Physica D* **133**, pp.179-188 (1999)
- [2] K. M. Hanson et al., "A framework for assessing confidence in computational predictions," *Exp. Tech.* **25**, pp. 50-55 (2001)
- [3] W. L. Oberkampf, "What are validation experiments?," *Exp. Tech.* **25**, pp. 35-40 (2001)
- [4] D. S. Sivia, *Data Analysis – A Bayesian Tutorial* (Clarendon, Oxford, 1996)
- [5] P. R. Bevington and D. K. Robinson, *Data Reduction and Error Analysis for the Physical Sciences* (McGraw-Hill, Boston, 1992)
- [6] G. T. Gray III et al., "Bridging the experimental-modeling gap - the Taylor test: an integrated test used to validate constitutive models," to appear in *AMPTIAC Newsletter*
- [7] P. J. Maudlin et al., "On the modeling of the Taylor cylinder impact test for orthotropic textured materials: experiments and simulations," *Int. J. Plasticity* **15**, pp. 139-166 (1999)
- [8] *Abaqus™/Explicit User's Manual* (Hibbitt, Karlsson & Sorensen, Pawtucket, RI, 1998)
- [9] F. M. Hemez et al., "Design of computer experiments for improving an impact test simulation," *Proc. 19th SEM Int. Conf. on Modal Analysis*, pp. 977-985 (2001)
- [10] K. M. Hanson et al., "Inversion based on computational simulations," in *Maximum Entropy and Bayesian Methods*, G. J. Erickson et al., eds., pp. 121-135 (Kluwer Academic, Dordrecht, 1998)
- [11] A. Griewank, *Evaluating Derivatives: Principles and Techniques of Automatic Differentiation* (SIAM, Philadelphia, 2000)