Journal of the Mechanics and Physics of Solids 60 (2012) 983–1001

Contents lists available at SciVerse ScienceDirect

Journal of the Mechanics and Physics of Solids

journal homepage: www.elsevier.com/locate/jmps

Rigorous model-based uncertainty quantification with application to terminal ballistics, part I: Systems with controllable inputs and small scatter

A. Kidane, A. Lashgari, B. Li, M. McKerns, M. Ortiz *, H. Owhadi, G. Ravichandran, M. Stalzer, T.J. Sullivan

Division of Engineering and Applied Science, California Institute of Technology, Pasadena, CA 91125, USA

ARTICLE INFO

Article history:
Received 23 August 2011
Received in revised form
1 December 2011
Accepted 3 December 2011
Available online 9 December 2011

Keywords:
Certification
Uncertainty quantification
Concentration of measure
Terminal ballistics

ABSTRACT

This work is concerned with establishing the feasibility of a *data-on-demand* (DoD) uncertainty quantification (UQ) protocol based on concentration-of-measure inequalities. Specific aims are to establish the feasibility of the protocol and its basic properties, including the tightness of the predictions afforded by the protocol. The assessment is based on an application to terminal ballistics and a specific system configuration consisting of 6061-T6 aluminum plates struck by spherical S-2 tool steel projectiles at ballistic impact speeds. The system's inputs are the plate thickness and impact velocity and the perforation area is chosen as the sole performance measure of the system. The objective of the UQ analysis is to certify the lethality of the projectile, i.e., that the projectile perforates the plate with high probability over a prespecified range of impact velocities and plate thicknesses. The net outcome of the UQ analysis is an M/U ratio, or confidence factor, of 2.93, indicative of a small probability of no perforation of the plate over its entire operating range. The high-confidence ($>99.9\%$) in the successful operation of the system afforded the analysis and the small number of tests (40) required for the determination of the modeling-error diameter, establishes the feasibility of the DoD UQ protocol as a rigorous yet practical approach for model-based certification of complex systems.

1. Introduction

Steady advances in computational power are presently forcing a reevaluation of computational paradigms and of the very notion of *scientific prediction*. Many applications of interest are concerned with systems that exhibit complex behavior and that are stochastic, be it as a result of uncertainties in the inputs of the system during normal operating conditions, uncertainties in the properties of the system, or the intrinsic stochasticity of the system itself. The performance of systems is characterized by outputs, or performance measures, whose value must remain within acceptable limits in order for the systems to operate safely and according to specifications. A fundamental question in systems engineering is whether the safe operation of a system can be certified with sufficient confidence to warrant deployment or continued operation. The ability to certify a system with high confidence is particularly important where high-consequence decisions are

* Corresponding author. Tel.: +1 626 395 4529.
E-mail address: ortiz@aero.caltech.edu (M. Ortiz).

0022-5096/$ - see front matter © 2011 Elsevier Ltd. All rights reserved.
doi:10.1016/j.jmps.2011.12.001

concerned regarding systems whose failure may result in high economic or societal costs (for instance, for background on the question of certification from a national security perspective see, e.g., Sharp and Wood-Schultz, 2003; Leader, 2005; Pilch et al., 2006; Committee on the Evaluation of Quantification of Margins and Uncertainties Methodology for Assessing and Certifying the Reliability of the Nuclear Stockpile, 2008; Helton, 2009). In these cases, it is often not enough to simply probe the performance of the system by performing a limited number of *hero* calculations, or even computing the mean performance and design margin of the system by means of extensive sampling. Instead, it is imperative to be able to predict the performance of the system with *rigorously quantified uncertainties*, so that design margins can be carefully weighed against the attendant uncertainties. This requirement for uncertainty quantification (UQ) represents a paradigm shift in predictive science in that it requires a careful assessment of the extent to which the performance of the system is likely to deviate from its mean. In particular, uncertainty quantification necessitates a systematic exploration of the performance of the system over its entire range of operation. A current grand challenge in predictive science thus concerns how to precisely effect such an exploration through a judicious combination of experiment and physics-based modeling and simulation, leading to a rigorous quantification of margins and uncertainties in the performance of complex systems.

This paper is concerned with the formulation and critical assessment of a rigorous protocol, which we term *data-on-demand* (DoD) protocol, for uncertainty quantification analysis of the performance of complex systems. In this approach, the precise sequence of tests and simulations required for uncertainty quantification of the system diameter is determined by an optimization algorithm and it is not known *a priori*. Thus, the feasibility of the protocol depends critically on the ability to execute tests *on demand* over the entire operating range of input parameters, hence the name of *data-on-demand* given to the protocol.

Following Lucas et al. (2008), we specifically take a *certification* point of view of uncertainty quantification. For definiteness, we consider systems whose operation can be described in terms of N scalar performance measures $(Y_1,\ldots,Y_N)\equiv Y\in\mathbb{R}^N$. The response of the systems of interest is *stochastic* due to the intrinsic randomness of the system, or randomness in the input parameters defining the operation of the system, or both. Suppose that the safe operation of the system requires that $Y\in A$ for some measurable *admissible set* $A\subset\mathbb{R}^N$. Ideally, we would then like Y to be always contained within A, i.e.,

$$\mathbb{P}[Y\in A]=1 \tag{1}$$

where $\mathbb{P}$ denotes probability in the sense of some appropriate probability space.[1] Systems satisfying this condition can be certified with complete certainty. However, this absolute guarantee of safe performance may be unattainable, e.g., if $\mathbb{P}$ lacks compact support. In these cases, we may relax condition (1) to

$$\mathbb{P}[Y\in A^c]\leq\epsilon \tag{2}$$

for some appropriate *certification tolerance* ϵ, where $A^c=\mathbb{R}^N\backslash A$ is the *inadmissible set*. Inequality (2) expresses the requirement that the probability of failure of the system be acceptably small.

However, in many cases the probability of failure $\mathbb{P}[Y\in A^c]$ is not readily available, e.g., if its determination is prohibitively expensive or unfeasible due to restrictions on testing or other constraints. In these cases, A *conservative* certification criterion can still be obtained if the probability of failure $\mathbb{P}[Y\in A^c]$ can be bounded from above and the upper bound is verified to be below the certification tolerance ϵ. Evidently, for an upper bound to be useful it must be *tight*, i.e., it must be close to the actual probability of failure $\mathbb{P}[Y\in A^c]$ and accessible in the laboratory, by computational means, or by a combination of both. In many areas of application, considerable effort is invested in the development of physics-based, high-fidelity computational models of the systems of interest (cf., e.g., Blue Ribbon Panel on Simulation-based Engineering Science, 2006). In addition, considerable resources are allocated to the development and deployment of computational platforms of ever increasing capacity, and to the development and deployment of new laboratory facilities with unique observational capabilities. *Therefore, within the present certification approach to UQ the essential mathematical, experimental and computational challenge is to devise a methodology leading to the determination of tight upper bounds on the probability of failure of complex systems.*

Following Lucas et al. (2008), we specifically consider probability-of-failure upper bounds of the *concentration-of-measure* type (Boucheron et al., 2004; McDiarmid, 1989; Lugosi, 2006; Ledoux, 2001). In the simplest form of the approach, the system is characterized by a deterministic *response function* $Y=G(X)$ that maps system inputs into outputs. When McDiarmid's (1989) inequality is used to bound probabilities of failure, the system may be certified on the sole knowledge of its mean performance $\mathbb{E}[Y]$ and a certain measure D_G of the spread of the response, known as *system diameter*, which provides a rigorous quantitative measure of the uncertainty in the response of the system. The quantification of system uncertainties, as measured by the system diameter D_G, entails the solution of N global optimization problems for the evaluation of subdiameters for each of the input parameters of the system. Every objective function evaluation in each of those N optimization problems requires two evaluations of the response function of the system for two different sets of input parameters. Since global optimization algorithms typically converge slowly, the evaluation of uncertainties directly from laboratory testing may be prohibitively expensive.

[1] Here and subsequently throughout this work we use standard notation of probability theory (cf, e.g., Chapter 2 of Evans, 2004 for an introduction).

A common response to such challenges is to devise physics-based models, often requiring extensive use of computational resources, capable of predicting the response of the system with high-fidelity in lieu of testing. Within the present context, a model is a function $F : E \subset \mathbb{R}^N \to \mathbb{R}$ that maps input parameters into performance measures. Evidently, a simple replacement of the test G by a model F does not result in rigorous certification in general, since there is no *a priori* guarantee that the predictions of the model are of sufficient fidelity. The assessment of model fidelity inevitably requires testing, though the tests need not be necessarily integral, cf. Topcu et al. (2011). Thus, the central challenge of *model-based UQ* is to *rigorously certify complex systems with a maximum of computation and a minimum of testing*. Specifically, the objective of model-based certification is the determination of rigorous upper bounds on the probability of failure of the system by a judicious combination of simulation and testing.

The main goal of the present work is to demonstrate the feasibility of the DoD UQ protocol for uncertainty quantification in its simplest form by means of an actual example of application, namely, the terminal ballistics of 6061-T6 aluminum plates impacted by spherical S-2 tool-steel projectiles. All tests of the system were conducted at Caltech's GALCIT Powder-Gun Plate-Impact Facility. The system's inputs are the plate thickness and impact velocity and, for simplicity, we choose the perforation area as the sole performance measure of the system. The objective of the uncertainty quantification analysis is to certify the *lethality* of the projectile, i.e., that the projectile perforates the plate with high probability over a prespecified range of impact velocities and plate thicknesses.

Terminal ballistics sets in motion complex deformation and material failure mechanisms and it therefore poses a challenging and exacting test of uncertainty quantification protocols. The response function of the system exhibits a number of features that renders certification particularly challenging. For instance, the dependence of the perforation area on projectile velocity exhibits a threshold, or *ballistic limit*, below which the perforation area vanishes. The extent of perforation rises rapidly for projectile velocities in excess of the ballistic limit and eventually saturates at sufficiently high projectile velocities, a type of response sometimes known as *cliff behavior*. The terminal ballistics application considered in this work thus provides an effective basis for ascertaining a number of theoretical and practical questions raised by uncertainty-quantification protocols, including: matters of feasibility, such as the number of tests required for the determination of system uncertainties; and issues of robustness, such as the tightness of the probability-of-failure upper bounds supplied by the protocol.

The paper is structured as follows. A summary of the theoretical basis of the DoD UQ protocol is provided in Section 2, with particular focus on McDiarmid's inequality (McDiarmid, 1989) and model-based certification. The application to terminal ballistics is presented in Section 3. This section begins with a brief description of the facility and of the experimental set-up, Section 3.1. The computational model used to simulate the terminal ballistics system, namely, the Optimal-Transportation MeshFree (OTM) method of Li et al. (2010), is described in Section 3.2. The optimization framework used for the computation of the system diameter and the modeling-error diameter is presented in Section 3.3. Finally, the results of the DoD UQ analysis are reported in Section 3.4. Concluding remarks and a *segue* to further extensions of this study are collected in Section 4.

2. Concentration-of-measure approach to UQ

The use of concentration-of-measure inequalities and, specifically, McDiarmid's inequality, as a basis for rigorous UQ was proposed by Lucas et al. (2008). We proceed to summarize the salient aspects of the resulting UQ protocol, which we term *data-on-demand* (DoD) protocol. We specifically consider systems characterized by N uncorrelated real random inputs $X = (X_1, \ldots, X_N) \in E \subseteq \mathbb{R}^N$ and a single real performance measure $Y \in \mathbb{R}$. We also begin by supposing that the function $G : \mathbb{R}^N \to \mathbb{R}$ that describes the response function of the system is deterministic and is known exactly, either experimentally or through the availability of an exact model. We additionally suppose that the system fails when $Y \leq a$, where a is a threshold for the safe operation of the system, and that the mean performance $\mathbb{E}[Y]$ is known exactly. These assumptions represent ideal conditions in which all uncertainty regarding the response of the system is *aleatoric uncertainty*, i.e., stems from the randomness of the system inputs, and there is no *epistemic uncertainty*, i.e., the behavior of the system is known exactly, including its mean response.

2.1. McDiarmid's inequality

McDiarmid's (1989) inequality is a result in probability theory that provides an upper bound on the probability that the value of a function depending on multiple independent random variables deviate from its expected value. It is an example of a general class of inequalities in probability known as *concentration-of-measure* inequalities (Boucheron et al., 2004; Lugosi, 2006; Ledoux, 2001).

A central device in McDiarmid's inequality is the *diameter* of a function. We begin by recalling that the *oscillation* $\mathrm{osc}(f,E)$ of a real function $f : E \to \mathbb{R}$ over a set $E \in \mathbb{R}$ is

$$\mathrm{osc}(f,E) = \sup\{|f(y) - f(x)| : x,y \in E\} \tag{3}$$

Thus, $\mathrm{osc}(f,E)$ measures the spread of values of f that may be attained by allowing the independent variable to range over its entire domain of definition. For functions $f : E \subset \mathbb{R}^n \to \mathbb{R}$ of several real values, component-wise *suboscillations* can be

defined as

$$\mathrm{osc}_i(f,E) = \sup\{|f(x)-f(y)| : x,y \in E,\ x_j = y_j \text{ for } j \neq i\} \tag{4}$$

Thus, $\mathrm{osc}_i(f,E)$ measures the maximum oscillation among all one-dimensional fibers in the direction of the ith coordinate. The *diameter* $D(f,E)$ of the function $f : E \to \mathbb{R}$ is obtained as the root-mean square of its component-wise suboscillations, i.e.,

$$D(f,E) = \left(\sum_{i=1}^{n} \mathrm{osc}_i^2(f,E)\right)^{1/2} \tag{5}$$

and it provides a measure of the spread of the range of the function. In particular, $D(C,E) = 0$ for all constant functions C.

An important property of the diameter is that it defines a *seminorm*.

Proposition 2.1. *For each subset $E \subset \mathbb{R}^N$, the diameter is a seminorm on the space of bounded real-valued functions on E.*

Proof. Evidently, $D(f,E) \geq 0$ for any function f and set E. In addition, the homogeneity property:

$$D(\lambda f,E) = |\lambda| D(f,E) \tag{6}$$

of the diameter follows directly from the definition. Finally, for any pair of functions $f,g : E \subset \mathbb{R}^n \to \mathbb{R}$ we have

$$\begin{aligned}
\mathrm{osc}_i(f+g,E) &= \sup\{|f(x)-f(y)+g(x)-g(y)| : x,y \in E,\ x_j = y_j \text{ for } j \neq i\} \\
&\leq \sup\{|f(x)-f(y)| + |g(x)-g(y)| : x,y \in E,\ x_j = y_j \text{ for } j \neq i\} \\
&\leq \sup\{|f(x)-f(y)| : x,y \in E,\ x_j = y_j \text{ for } j \neq i\} + \sup\{|g(x)-g(y)| : x,y \in E,\ x_j = y_j \text{ for } j \neq i\} \\
&= \mathrm{osc}_i(f,E) + \mathrm{osc}_i(f,E)
\end{aligned} \tag{7}$$

whereupon

$$\begin{aligned}
D(f+g,E) = \left(\sum_{i=1}^{n} \mathrm{osc}_i^2(f+g,E)\right)^{1/2} &\leq \left(\sum_{i=1}^{n} (\mathrm{osc}_i(f,E) + \mathrm{osc}_i(g,E))^2\right)^{1/2} \\
&\leq \left(\sum_{i=1}^{n} \mathrm{osc}_i^2(f,E)\right)^{1/2} + \left(\sum_{i=1}^{n} \mathrm{osc}_i^2(g,E)\right)^{1/2} = D(f,E) + D(g,E)
\end{aligned} \tag{8}$$

and the diameter satisfies the triangular inequality. □

The diameter provides a measure of the variability of a function value given the variability of its inputs. The diameter fails to be a norm since, as noted earlier, it vanishes for constant functions. However, the triangular-inequality property of the diameter will prove invaluable in subsequent developments as a means of deriving useful upper bounds for the diameters of functions.

The diameter may also be thought of as providing a quantitative measure of the uncertainty in the value of a function whose inputs are themselves uncertain. Evidently, it is possible to contrive many other measures of uncertainty, and a wide range of such *ad hoc* measures have been proposed, often without much justification other than convenience. Within the present framework, what sets diameters apart—and confers them a privileged status as an uncertainty measure—is that they supply just the information that is required to formulate a rigorous upper bound on the probability of failure of the system. Specifically, we have the following.

Theorem 2.2 (*McDiarmid (1989) inequality*). *Assume that the random variables $X_1, \ldots, X_N$ are independent random variables, let $f : E \subset \mathbb{R}^N \to \mathbb{R}$ be integrable and let $D(f,E)$ be its diameter. Then, for every $r \geq 0$:*

$$\mathbb{P}[f(X) \geq \mathbb{E}[f] + r] \leq \exp\left(-2\frac{r^2}{D^2(f,E)}\right) \tag{9}$$

McDiarmid's inequality may be taken as a basis for rigorous uncertainty quantification, as discussed next.

2.2. UQ based on McDiarmid's inequality

Suppose, that the mean performance $\mathbb{E}[G]$ and the diameter $D_G = D(G,E)$ of the system are known exactly. Then, McDiarmid's inequality (9) gives the following upper bound, not necessarily optimal (Owhadi et al., submitted for publication), on the probability of failure of the system (Lucas et al., 2008), namely,

$$\mathbb{P}[G(X) \leq a] \leq \exp\left(-2\frac{(\mathbb{E}[G]-a)_+^2}{D_G^2}\right) \tag{10}$$

where we write $x_+ \coloneqq \max(0,x)$ and $D_G \equiv D(G,E)$ for the diameter of the response function. From (10) it in turn follows that the inequality

$$\frac{M}{U} \geq \sqrt{\log\sqrt{\frac{1}{\epsilon}}} \tag{11}$$

supplies a conservative certification criterion, where

$$M = (\mathbb{E}[G]-a)_+ \tag{12a}$$

$$U = D_G \tag{12b}$$

are the *system margin* and *system uncertainty*, respectively. With these identifications, the certification criterion can be expressed in the form:

$$\mathrm{CF} = \frac{M}{U} \geq \sqrt{\log\sqrt{\frac{1}{\epsilon}}} \tag{13}$$

where CF is the *confidence factor*.

Remarkably, if follows from the preceding analysis that systems can be rigorously certified solely on the basis of their the mean performance and their diameter. It also bears emphasis that, in formulating the McDiarmid bound, only ranges of input parameters, and not their detailed probability distribution functions, are required. McDiarmid's inequality supplies rigorous quantitative definitions of system margin and system uncertainty. In particular, the latter is measured by the *system diameter* D_G. We note from definition (5) that

$$D_G = \left\{\sum_{i=1}^{N} \mathrm{osc}_i(G,E)^2\right\}^{1/2} \tag{14}$$

where the oscillation $\mathrm{osc}_i(G,E)$ may now be regarded as the *subdiameter* corresponding to variable X_i. That subdiameter may in turn be regarded as a measure of the uncertainty contributed by the variable X_i to the total uncertainty of the system.

We note that the quantification of system uncertainty, as measured by the system diameter D_G, entails the solution of N global optimization problems for the evaluation of the subdiameters $\mathrm{osc}_i(G,E)$. Every objective function evaluation in each of those N optimization problems requires the execution of two tests for the evaluation of $G(X_1,\ldots,X_i,\ldots,X_N)$ and $G(X_1,\ldots,X_i',\ldots,X_N)$. The precise sequence of tests required for the evaluation of the system diameter is determined by the optimization algorithm and it is not known *a priori*. Thus, the feasibility of the protocol depends critically on the ability to execute tests *on demand* over the entire operating range of input parameters, hence the name of *data-on-demand* given to the protocol.

2.3. Empirical mean estimation

In the foregoing we have assumed that the mean performance $\mathbb{E}[G]$ of the system is known *a priori*. However, in many situations of practical interest such information is not available and the mean performance must instead be estimated. Suppose that, to this end, we perform m tests based on an unbiased sampling of the input parameters, resulting in predicted performances $Y^1,Y^2,\ldots,Y^m$. We may then define an *empirical mean performance* as

$$\mathbb{E}_m[G] = \frac{1}{m}\sum_{i=1}^{m} Y^i \tag{15}$$

Under these conditions, the probability of failure $\mathbb{P}[Y \in A^c]$ of the system can now only be determined to within confidence intervals reflecting the randomness of the empirical mean $\mathbb{E}_m[G]$, namely (Lucas et al., 2008),

$$\mathbb{P}\left[\mathbb{P}[Y \in A^c] \geq \exp\left(-2\frac{[\mathbb{E}_m[G]-a-\alpha]_+^2}{D_G^2}\right)\right] \leq \epsilon' \tag{16}$$

where ϵ' is a prespecified *estimation tolerance* and

$$\alpha = D_G m^{-1/2}(-\log\epsilon')^{1/2} \tag{17}$$

Thus, with probability $1-\epsilon'$ we have

$$\mathbb{P}[Y \in A^c] \leq \exp\left(-2\frac{[\mathbb{E}_m[G]-a-\alpha]_+^2}{D_G^2}\right) \tag{18}$$

and the certification criterion (2) in turn becomes

$$\frac{[\mathbb{E}_m[G]-a-\alpha]_+}{D_G} \geq \sqrt{\log\sqrt{\frac{1}{\epsilon}}} \tag{19}$$

This certification criterion is again of the form (13) with margin and uncertainty given by

$$M = [\mathbb{E}_m[G] - a - \alpha]_+ \tag{20a}$$

$$U = D_G \tag{20b}$$

Comparison of (12) and (20) shows that the estimation of the mean performance of the system by means of an empirical mean effectively reduces the margin in the amount α. Evidently, this margin hit can be reduced to an arbitrarily small value by carrying out a sufficiently large number of tests. The certification criterion (20) again shows that, in the absence of epistemic uncertainty, certification can be rigorously achieved from the sole knowledge of the system diameter and an empirical mean performance.

2.4. Model-based UQ

As noted in the foregoing, the direct determination of system diameters by means of laboratory testing may require a large number of tests and hence be prohibitively expensive. However, suppose that a model of the system is available, i.e., a function $F : E \subset \mathbb{R}^N \to \mathbb{R}$ that maps input parameters into performance measures and that approximates the actual response function G of the system. Evidently, a simple replacement of G by a model F does not result in rigorous certification in general, since there is no *a priori* guarantee that the predictions of the model are of sufficient fidelity. The aim of *model-based certification* is to achieve rigorous certification with a maximum use of modeling and simulation and a minimum use of testing.

Models can be worked into the certification protocol described in the foregoing as follows. Begin by noting the monotonicity property

$$\exp\left(-2\frac{M^2}{D_G^2}\right) \le \exp\left(-2\frac{M^2}{U^2}\right) \tag{21}$$

provided that

$$U \ge D_G \tag{22}$$

Hence, (18) remains a conservative upper bound with probability $1-\epsilon'$ provided that the margin M is computed according to (20a) and the uncertainty U satisfies the inequality (22). A convenient upper bound on the system diameter follows from the triangular-inequality property of the diameter, cf. Proposition 2.1, as

$$D_G = D_{G-F+F} \le D_{G-F} + D_F \equiv U \tag{23}$$

where again we write $D_{G-F} = D(G-F,E)$ and $D_F = D(F,E)$ for short. We note that U indeed supplies an upper bound on the diameter D_G of the system. The diameter D_F may be regarded as a *predicted model diameter*, i.e., the system diameter predicted by the model. From the definition of the diameter it follows that D_F is computed by exercising the model in the absence of any testing. The requisite quantitative measure of model fidelity is supplied by the *model-error diameter* D_{G-F}. This diameter measures the discrepancy between model prediction and experimental observation. Evidently, many such measures may be contrived in an *ad hoc* manner. However, what confers the diameters D_F and D_{G-F} their unique value is that they furnish a rigorous upper bound on the probability of failure of the system and, therefore, may be taken as a basis for conservative certification.

We note that the quantification of the model-error diameter D_{G-F}, entails the solution of N global optimization problems for the evaluation of the subdiameters $\mathrm{osc}_i(G-F,E)$. Every objective function evaluation in each of those N optimization problems requires the execution of two integral tests for the evaluation of $G(X_1,\ldots,X_i,\ldots,X_N)$ and $G(X_1,\ldots,X_i',\ldots,X_N)$, respectively and, simultaneously, two nominally identical simulations $F(X_1,\ldots,X_i,\ldots,X_N)$ and $F(X_1,\ldots,X_i',\ldots,X_N)$. Thus, the approach is practical if the error function $G-F$ is better behaved than either G or F separately and, hence, the computation of the diameter D_{G-F} may be effected by means of rapidly converging iterative schemes, with the result that the number of tests is minimized. This expectation is indeed born out in the terminal ballistics example of Section 3.1.

2.5. UQ as an optimization problem

The McDiarmid concentration-of-measure approach to UQ is attractive because it requires tractable information on response functions (subdiameters) and measures (independence and mean response function). In the preceding sections we have described how to insert this information into McDiarmid's concentration inequality to obtain an upper bound on probabilities of deviation. A question of theoretical and practical importance concerns whether it is possible to obtain an optimal concentration inequality from the same information. A related question concerns the possible use of information other than subdiameters and mean values. These questions have been addressed in Owhadi et al. (submitted for publication) and we proceed to summarize their main results for completeness. Assume, for definiteness, that we want to certify that

$$\mathbb{P}[G(X) \ge a] \le \epsilon \tag{24}$$

based on the information that $\mathrm{osc}_i(G,E) \le D_i$, $X=(X_1,\ldots,X_N)$, $\mathbb{E}[G] \le 0$ and that the inputs $X_1,\ldots,X_N$ are independent under $\mathbb{P}$. In this situation, the optimal bound $\mathcal{U}(\mathcal{A}_{McD})$ on the probability of failure $\mathbb{P}[G(X) \ge a]$ is the solution of the following optimization problem:

$$\mathcal{U}(\mathcal{A}_{MD}) = \sup_{(f,\mu)\in\mathcal{A}_{MD}} \mu[f(X) \ge a] \tag{25}$$

where

$$\mathcal{A}_{MD} = \left\{ (f,\mu) \;\middle|\; \begin{array}{c} f: E_1 \times \cdots \times E_N \to \mathbb{R} \\ \mu \in \mathcal{M}(E_1) \otimes \cdots \otimes \mathcal{M}(E_m) \\ \mathbb{E}_\mu[f] \le 0 \\ \mathrm{osc}_i(f,E) \le D_i \end{array} \right\} \tag{26}$$

and $\mathcal{M}(E_k)$ denotes the set of probability measures on E_k. McDiarmid's concentration-of-measure inequality provides the following upper bound:

$$\mathcal{U}(\mathcal{A}_{MD}) \le \exp\left(-2 \frac{a^2}{\sum_{i=1}^N D_i^2} \right) \tag{27}$$

Similarly, an information set $\mathcal{A}$ other than (26) results in an optimal probability of deviation:

$$\mathcal{U}(\mathcal{A}) = \sup_{(f,\mu)\in\mathcal{A}} \mu[f(X) \ge a] \tag{28}$$

In practical applications, the available information does not determine $(G,\mathbb{P})$ uniquely but an information set $\mathcal{A}$ consisting of possible values of $(G,\mathbb{P})$. Optimal Uncertainty Quantification (OUQ) (Owhadi et al., submitted for publication), is based on the observation that, given a set of assumptions and information about the system, there exist optimal bounds on uncertainties; these are obtained as extreme values of well-defined optimization problems corresponding to extremizing probabilities of deviation subject to the constraints imposed by assumptions and information.

Although the optimization problems (25) and (28) are extremely large (with optimization variable in infinite-dimensional spaces of measures and functions), under general moment and independence conditions, Owhadi et al. (submitted for publication) have shown that they have finite-dimensional reductions. The reduction theorems are a generalization of classical Chebyshev inequalities and of the Markov–Krein theorem (Karlin and Studden, 1966; Marshall and Olkin, 1979). The reduction theorems rely on an extension of Choquet theory (von Weizsäcker and Winkler, 1979; Winkler, 1988) and on characterization of simplices and vector lattices (Kendall, 1962).

An application of OUQ that is relevant to the present work concerns the development of explicit and optimal concentration inequalities of the McDiarmid type. Namely, it is shown in Owhadi et al. (submitted for publication) that

$$\mathcal{U}(\mathcal{A}_{MD}) = \mathcal{U}(\mathcal{A}_C) \tag{29}$$

where

$$\mathcal{U}(\mathcal{A}_C) := \sup_{(C,\alpha)\in\mathcal{A}_C} \alpha[h^C \ge a] \tag{30}$$

is a finite-dimensional (therefore computationally tractable) optimization problem defined by

$$\mathcal{A}_C = \left\{ (C,\alpha) \;\middle|\; \begin{array}{c} C \subset \{0,1\}^N, \\ \alpha \in \otimes_{i=1}^N \mathcal{M}(\{0,1\}), \\ \mathbb{E}_\alpha[h^C] \le 0 \end{array} \right\}$$

h^C is a real function mapping $\{0,1\}^N$ onto $\mathbb{R}$, parameterized by C (a coloring of the nodes of the N-dimensional hypercube), and defined as follows:

$$h^C(t) = a - \min_{s\in C} \sum_{i:s_i\ne t_i} D_i$$

Explicit solutions of problem (30) can be computed for $N=2$ and $N=3$. For $N=2$, we have

$$\mathcal{U}(\mathcal{A}_{MD}) = \begin{cases} 0 & \text{if } D_1+D_2 \le a \\ \frac{(D_1+D_2-a)^2}{4D_1D_2} & \text{if } |D_1-D_2| \le a \le D_1+D_2 \\ 1-\frac{a}{\max(D_1,D_2)} & \text{if } 0 \le a \le |D_1-D_2| \end{cases} \tag{31}$$

We also refer to Owhadi et al. (submitted for publication) for an explicit expression for $N=3$ and for the tail of the distribution (in a) for $N \ge 4$. The preceding optimal bounds provide a means of improving on the simple McDiarmid bounds that are taken as the basis for the present work. Since the bounds are optimal, further improvements inevitably require information other than—or in addition to—system diameters and mean performance.

3. Application to ballistic penetration

We proceed to demonstrate the feasibility and performance of the DoD UQ protocol introduced in the foregoing by means of an application to terminal ballistics. The particular system under consideration consists of 6061-T6 aluminum plates struck by spherical S-2 tool steel projectiles. The system's inputs are the plate thickness and impact velocity. For simplicity, we choose the perforation area as the sole performance measure of the system. The objective of the uncertainty quantification analysis is to certify the *lethality* of the projectile, i.e., that the projectile perforates the plate with high probability over a prespecified range of impact velocities and plate thicknesses.

3.1. Experimental facility

All tests required for the determination of the modeling error diameter D_{F-G} were conducted at Caltech's GALCIT Powder-Gun Plate-Impact Facility. This facility thus plays several crucial roles as part of the uncertainty quantification analysis: it defines the system to be certified and provides a specific and concrete realization of that system; and it supplies the *data-on-demand* required for the computation of the modeling error diameter D_{F-G}, cf. Section 2.4. A brief description of the facility follows for completeness.

3.1.1. Materials

An aluminum alloy (6061-T6) was used as a target material. Selected material properties of this alloy are listed in Table 1. Four different plate thicknesses, 0.81 mm (0.032 in), 1.02 mm (0.040 in), 1.27 mm (0.050 in) and 1.62 mm (0.063 in) were used in the tests. The plates were cut from plate stock to nominally 152.4 mm (6 in) square in size. The projectile was a 7.94 mm (5/16 in) diameter S-2 tool-steel grade 200 steel sphere with high hardness and strength (R_c 55–58, yield strength 2 GPa). This aluminum/steel target/projectile system is greatly overmatched and, at impact velocities exceeding the ballistic limit, the projectile perforates the target without much plastic deformation.

3.1.2. Experimental setup

A general view of the powder gun used in the tests is shown in Fig. 1a. The gun consists of a pressure chamber fitted with a smooth-bore steel barrel of 25.4 mm (1 in) in inner diameter and 152 cm (5 feet) in length. The target plate is positioned at the end of the barrel and attached to a 25.4 mm (1 in) wide steel frame. The two vertical edges of the plate are simply supported and the other two edges (top and bottom) are clamped to the steel frame as shown in Fig. 1b, which results in a 101.6 mm (4 in) $\times$ 101.6 mm (4 in) target area. The spherical projectile is glued on to a light-weight styrofoam sabot 25.4 mm (1 in) in diameter and 50 mm (2 in) in length, Fig. 1c. In tests, the projectile is first inserted into the end of the barrel. Subsequently, the gun chamber is filled with helium and the pressure is set to the desired value of 138–552 kPa (20–80 psi), depending on the required velocity. The gas in the chamber is then released by rapid opening of a valve, resulting in the projectile being propelled in the barrel and eventually impacting the center of the target. Near the impact end of the barrel, light emitting diodes are placed in two orifices on the periphery of the barrel and two fast-response photo detectors are placed in orifices diametrically opposite to the diodes. The two detectors are 108 mm (4.25 in) apart. When the projectile passes through these orifices it generates a ramp pulse whose duration is recorded using a high-speed WaveSurfer 24Xs oscilloscope (LeCroy, Chestnut Ridge, NY), which has a bandwidth of 200 MHz and a sampling rate of 2.5 Giga samples/s. The projectile velocity is then calculated from the time of travel of the projectile across the two detectors. The calculated projectile velocity is verified by calculating the pulse duration corresponding to the projectile crossing a single detector. The velocity of the projectile is calibrated against gas pressure and can be controlled within $+2$–5 m/s in the impact velocity range of 100–400 m/s. The perforation area is measured using high-precision gage pins and verified using a surface profile scanning technique. The surface-profile scanning technique employs a ConoScan 3000 (Optimet, North Andover, MA), which uses a conoscopic holography technology and has a precision of 10 m over a scan area of up to 120×120 mm^2.

3.1.3. Experimental results

Targets were removed from the clamping frame fixture post-mortem and the perforated area was measured. Typical images of the aft and rear faces of a perforated plate are shown in Fig. 2a and b. The perforated area was measured using

Table 1
Physical and mechanical properties of 6061-T6 aluminum alloy.

Density (kg/m^3)	Melting point (°C)	Young's modulus (GPa)	Shear modulus (GPa)	Brinell hardness
2700	652	68.9	26	95
Tensile yield strength (MPa)	Ultimate tensile strength (MPa)	Shear strength (MPa)	Fracture toughness (MPa$\sqrt{\text{m}}$)	
276	310	207	29	

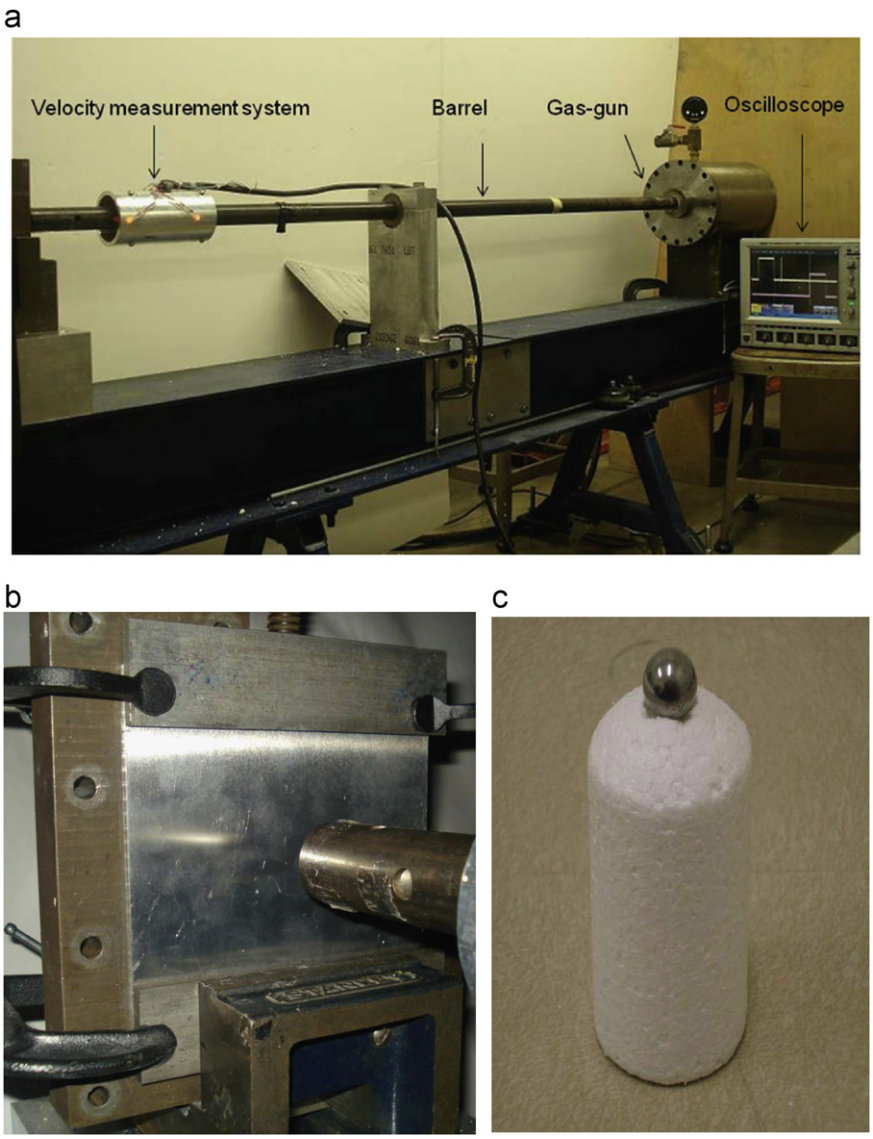

Fig. 1. Experimental set up. (a) General view of Caltech's GALCIT Powder-Gun Plate-Impact Facility. (b) 6061-T6 aluminum alloy target held in the fixture at the end of the barrel. (c) S-2 tool steel (7.94 mm diameter) projectile glued to a styrofoam sabot.

high precision gage-pins and verified using the surface profile scanning system. A typical image of a scanned surface of a perforated plate is shown in Fig. 2c. Beyond the ballistic limit, the perforated area is fairly independent of the projectile velocity and plate thickness. For the material system under consideration, perforation happens by the classical mechanism of plugging. The projectile velocity at which the 6061-T6 aluminum alloy plate is fully perforated is plotted in Fig. 3a as a function of plate thickness. Using the data from the gage-pins, the perforated area is plotted in Fig. 3b as a function of projectile velocity for different plate thicknesses. This information was used to establish the ballistic limit of the aluminum alloy for a given plate thickness. For the range of plate thicknesses investigated here, this ballistic limit is in the range of 120–150 m/s. As expected, the ballistic limit increases with increasing plate thickness.

3.2. Computational model

The accurate simulation of terminal ballistics is a grand challenge in scientific computing that places exacting demands on physics models, solvers and computing resources. In order to meet these challenges, in calculations we use the Optimal-Transportation MeshFree (OTM) method of Li et al. (2010) extended to account for contact and fracture (cf. Schmidt et al., 2009). The OTM method combines: (i) optimal transportation concepts such as the Wasserstein distance between successive mass densities in order to discretize the action integral in time; (ii) Maximum-entropy (max-ent) meshfree interpolation (Arroyo and Ortiz, 2006) from a nodal-point set in order to avoid mesh entanglement and the need for

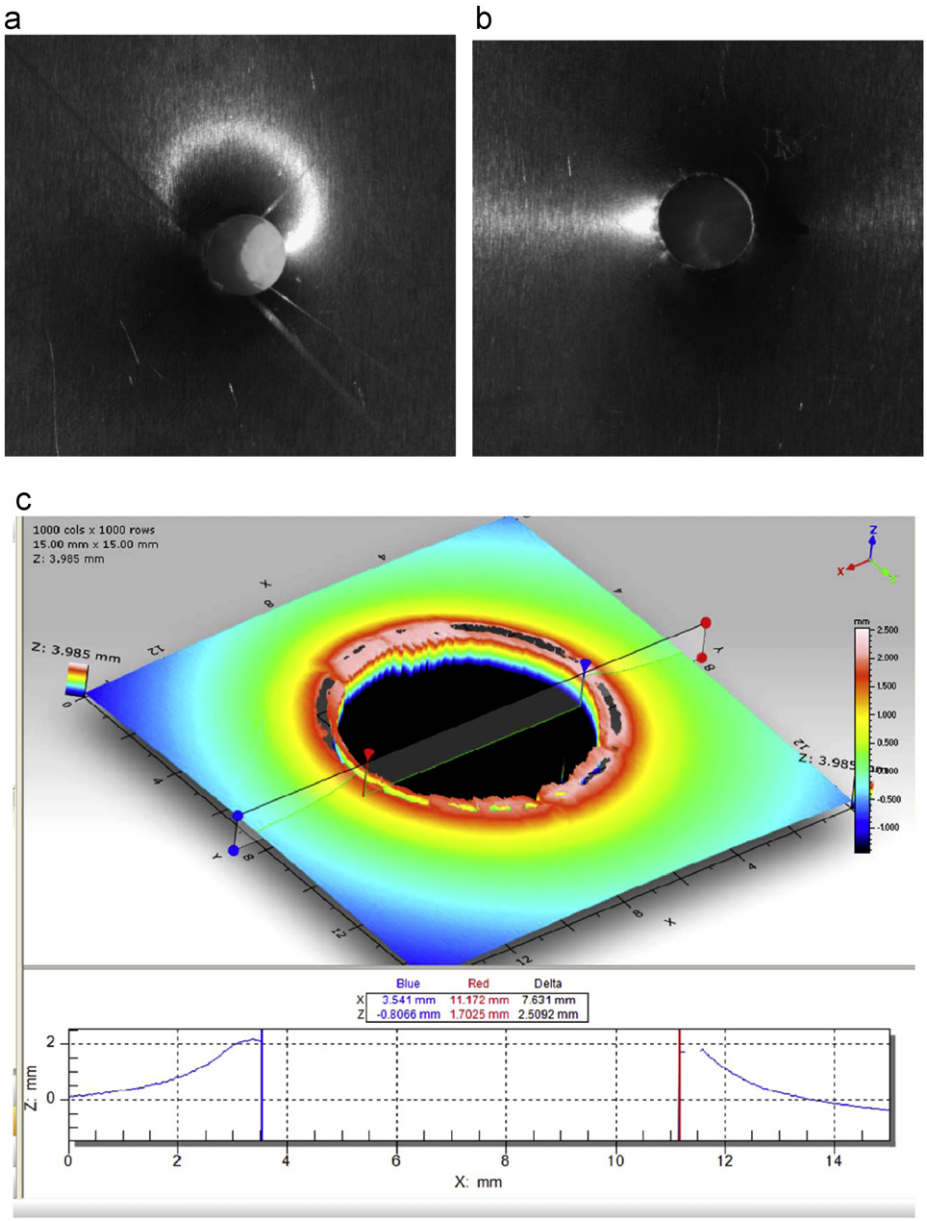

Fig. 2. Typical images of the perforated 6061-T6 aluminum plate: (a) Impact face; (b) Rear face and (c) Scanned profile.

continuous remeshing in simulations of unconstrained flows; and (iii) material-point sampling in order to track the local state of material points and carry out constitutive updates. A detailed description of the OTM terminal ballistic model employed in the present work may be found in Li et al. (2012).

The optimal-transportation approach to time discretization leads to geometrically exact updates of the local volumes and mass densities, thus bypassing the need for solving a costly Poisson equation for the pressure and entirely eliminating the mass conservation errors that afflict Eulerian formulations. In addition, by adopting a discrete Hamilton principle based on the time-discrete action furnished by optimal transportation, the discrete trajectories have exact conservation properties including symplecticity, linear and angular momentum. The total energy can also be exactly conserved, subject to solvability constraints, if time is taken as an independent generalized coordinate and the successive time steps are computed so as to render the action stationary (Kane et al., 1999). Max-ent interpolation (Arroyo and Ortiz, 2006) offers the advantage of being mesh-free and entirely defined—essentially explicitly—by the current nodal-set positions. In addition, max-ent interpolation satisfies a Kronecker-delta property at the boundary, which greatly facilitates the enforcement of essential boundary conditions, and has good accuracy, convergence and monotonicity properties that render it well-suited to applications involving shocks.

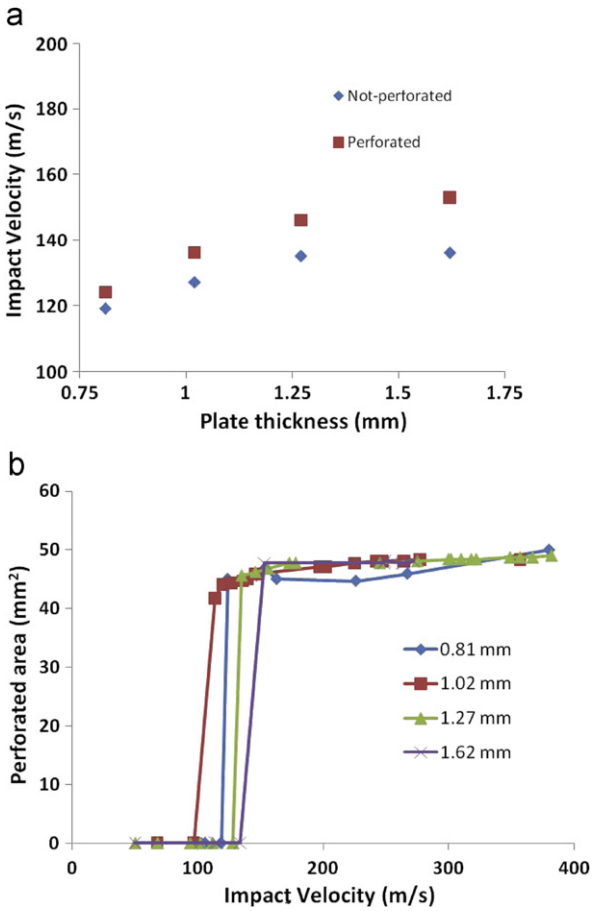

Fig. 3. (a) Ballistic limit for 6061-T6 aluminum alloy as a function of plate thickness. (b) Perforated area versus impact velocity for different thicknesses of 6061-T6 aluminum plates.

All local state data is stored—and constitutive calculations are performed—at an evolving material point set. As their name indicates, material points designate fix material points of the body and, therefore, are convected by the deformation. Material points also carry a fixed volume and mass and serve the purpose of integration points for the calculation of the effective nodal forces and masses. In calculations, the max-ent shape functions are recalculated at every time step, which effectively results in a dynamic or *on-the-fly* reconnection of nodes and material points. Since max-ent functions have essentially local support, the calculation of the max-ent functions and derivatives at a material point involves a *local neighborhood* of nodal points only. Such local neighborhoods are updated dynamically using range searches. Conveniently, the resulting reconnection between material points and nodes leaves the material points invariant, thus entirely eliminating the need for state-variable remapping.

In addition to transporting mass efficiently, a terminal ballistics solver must account for complex contact, fracture and fragmentation phenomena. A convenient feature of OTM, which is common to other material point methods (Sulsky et al., 1994) is that *seizing contact* is accounted for automatically, essentially for free, by simply allowing nodal points from different bodies to belong to the local neighborhoods of material points. The ensuing cancelation of linear momenta then automatically accounts for dynamic contact interactions of the seizing type. A final component of the material model concerns the simulation of fracture and fragmentation. Specifically, we simulate fracture by a variational material-point failure scheme known as *eigenfracture* (Schmidt et al., 2009). In this scheme, the energy-release rate attendant to the failure of a material point is estimated by a local energy-averaging procedure, and material points are failed when the attendant energy-release rate exceeds the specific fracture energy of the material. The eigenfracture scheme is known to properly converge to Griffith fracture in the limit of vanishingly small mesh sizes (Schmidt et al., 2009).

We conclude this section with some typical calculations for purposes of illustration. All materials are described by means of engineering J_2-viscoplasticity models with power-law hardening, rate-sensitivity and thermal softening. Specifically, the sensitivity law is assumed to be of the form:

$$\dot{\epsilon}^p = \dot{\epsilon}_0^p \left(\frac{\sigma - \sigma_c(\epsilon^p, T)}{\sigma_0} \right)^m \tag{32}$$

Table 2
Mechanical constants.

Material	ρ (kg/m^3)	E (GPa)	v	σ_{y0} (MPa)	ϵ_0^p	n	$\dot{\epsilon}_0^p$	m
Al6061-T6	2700	69	0.33	276	0.001	13.5	1000	11.5
S2 tool steel	12 695	193	0.3	2000	0.001	22	1	340

Table 3
Thermal constants.

Material	c (J/kgK)	T_0 (K)	T_m (K)	l	β
Al6061-T6	896	298	853	0.5	0.9
S2 Tool steel	477	298	1777	1.17	0.9

where ϵ^p is the Mises effective plastic strain, σ is a Mises effective stress, T is the absolute temperature, $\dot{\epsilon}_0^p$ is a reference effective plastic-strain rate, σ_0 is a reference stress and m is the rate-sensitivity exponent. The critical stress for plastic yielding is assumed of the form:

$$\sigma_c(\epsilon^p,T) = \sigma_y(T)\left[1+\left(\frac{\epsilon^p}{\epsilon_0^p}\right)^{1/n}\right] \tag{33}$$

where $\sigma_y(T)$ is the yield stress, ϵ_0^p is a reference effective plastic strain, and n is the hardening exponent. Finally, the yield stress is assumed of the form:

$$\sigma_y(T) = \sigma_{y0}\left(1-\frac{T}{T_0}\right)^l \tag{34}$$

where σ_{y0} is the yield stress at zero absolute temperature, T_0 is the melting temperature and l is the thermal softening exponent. The elastic response is assumed to be quadratic in the elastic logarithmic strains, with isotropic elastic coefficients depending linearly on temperature and vanishing at the melting temperature. The equation of state, which governs the volumetric response of the material, is assumed to be of the Mie–Gruneisen type. The material constants used in calculations are collected in Tables 2 and 3.

Fig. 4 shows a snapshot of a calculation concerned with an aluminum plate struck by a spherical steel impactor at 370 m/s. The thickness of the plate is 1.6 mm and the diameter of the projectile is 7.94 mm. Fig. 4a shows the nodal set color coded according to von-Mises stress, and Fig. 4b shows the material point set colored according to energy-release rate. Fig. 5a and b shows level contours of Mises effective plastic stress and temperature, respectively, at the same moment in time. As may be seen from the figures, the projectile easily defeats the plate, which is perforated by a conventional plugging mechanism. As expected, the energy-release rate is concentrated around the shear rupture that separates the plug from the rest of the plate. The temperate field also peaks at the plug boundary, resulting in thermal softening of the plate. This thermal softening in turn facilitates and promotes localization of deformation, eventually resulting in plug formation. The large plastic deformations undergone by the plate are also noteworthy. By contrast, the projectile remains comparatively undeformed.

Typical comparisons between experimental data and model predictions are shown in Fig. 6. As may be seen from the figure, the OTM model qualitatively predicts the ballistic limit, the subsequent steep rise in perforation area, and the eventual plateau at impact velocities greatly in excess of the ballistic limit. It also predicts qualitatively the boundary between perforation and non-perforation in the thickness-impact velocity plane. Whereas the predictions of the simulations may thus be regarded to compare favorably with experiment, whether the accuracy of the predictions suffices or not depends critically on their intended application and needs to be quantified accordingly. In the specific framework of the DoD UQ protocol under consideration here, the appropriate quantitative measure of model fidelity is supplied by the modeling-error diameter D_{F-G}, cf. Sections 2.4 and 3.4.

3.3. Optimization framework

Our implementation of the DoD UQ protocol utilizes an approximate global optimization algorithm to compute the model diameter D_F, where the accuracy depends primarily on the optimum being found within a fixed tolerance. The computation of the model-error diameter D_{G-F} is also done by approximate global optimization. Because each evaluation of G requires an expensive experimental measurement, the optimization of D_{G-F} has the additional requirement of minimizing the number of evaluations of G. The details of the computations of D_F and D_{G-F} are presented below.

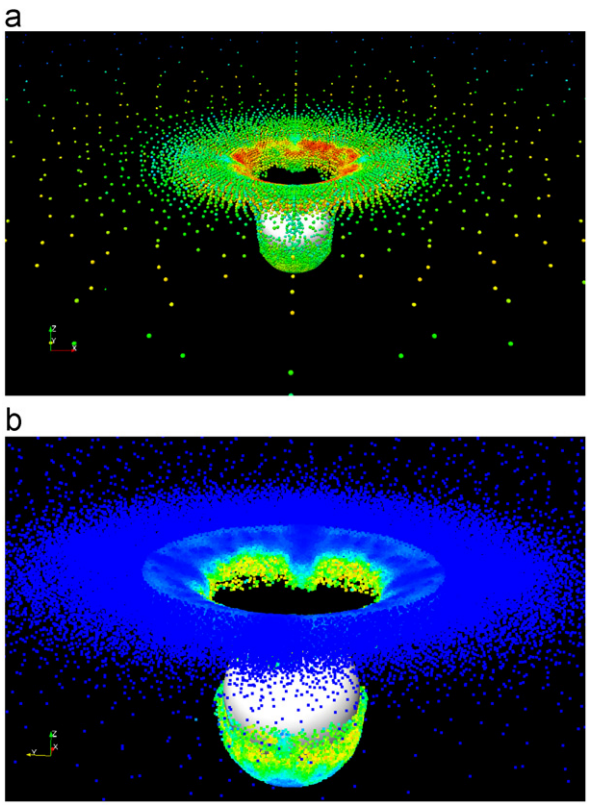

Fig. 4. Snapshot of aluminum plate after perforation by steel sphere. (a) Nodal point set. (b) Material point set. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

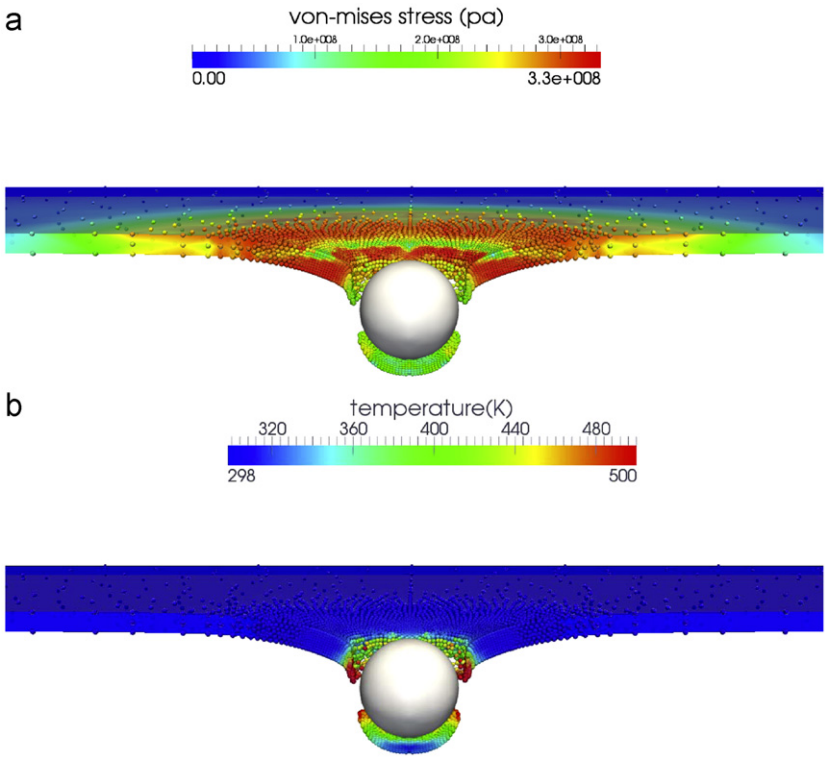

Fig. 5. Snapshot of aluminum plate after perforation by steel sphere. (a) Mises effective plastic stress. (b) Temperature.

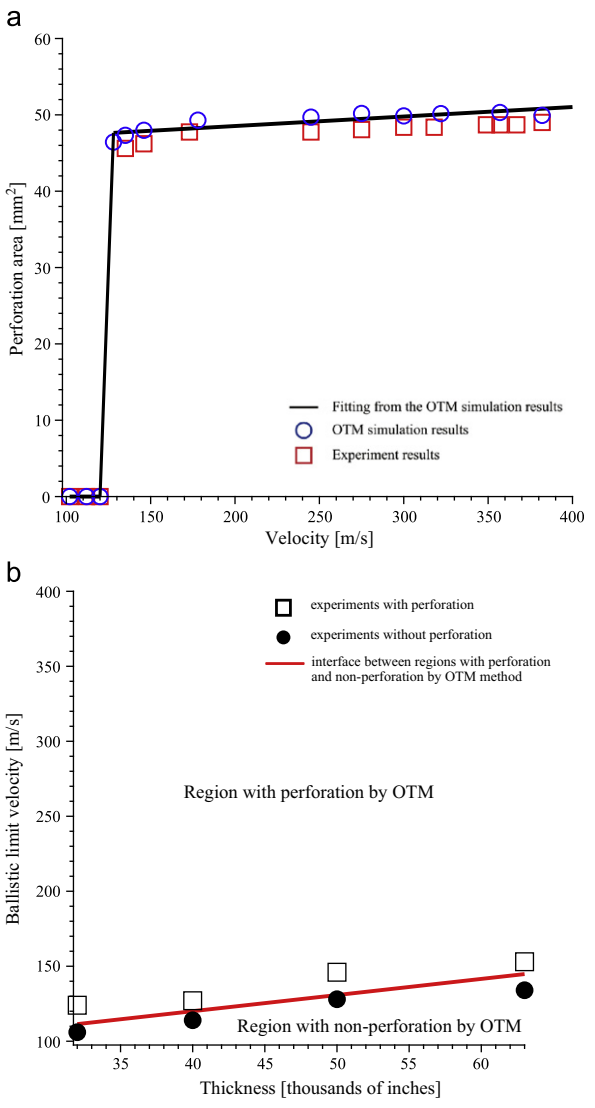

Fig. 6. Typical comparison between experimental data and model predictions. (a) Ballistic curve for 50 mil plate thickness. (b) Boundary between perforation and non-perforation in the thickness-impact velocity plane.

3.3.1. *The model diameter D_F as an optimization problem*

D_F^2 is the sum of the N subdiameters $\mathrm{osc}_i^2(F,E)$, Eq. (14), and computing each subdiameter requires solving a global optimization problem. During the computation of the i th subdiameter, the optimization algorithm takes an $N+1$-dimensional set of input parameters $\mathcal{P}=(X_1,\ldots,X_i,X_i',\ldots,X_N)$ and the current measure of oscillation $C(\mathcal{P})=-[F(X_1,\ldots,X_i',\ldots,X_N)-F(X_1,\ldots,X_i,\ldots,X_N)]^2$ as inputs, and produces an updated instance of $\mathcal{P}$. The optimization loop iterates until the condition $C(\mathcal{P})<-(\mathrm{osc}_i^2-\epsilon)$ is satisfied. The final set $\mathcal{P}$ contains the optimal (worst) input parameter set $(X_1,\ldots,X_i,X_i',\ldots,X_N)$ resulting in the largest oscillation in the response function.

3.3.2. *Computing the model diameter D_F*

Due their stochastic nature, global search algorithms often require hundreds of iterations and thousands of function evaluations to find a global optimum. Local methods, such as Powell's (1989,1994) method, may require orders of magnitude fewer iterations and evaluations, but do not generally converge to a global optimum for complex, highly non-convex objective functions. To compute diameters, we use a *lattice*-Powell method that carries out local optimizations at the centers of lattice cubes in parameter space. The local optimizations are quadratically convergent, based on a multivariate gradient descent algorithm and do not require implicit first derivatives. Specifically, we partition the range of each parameter i to be optimized into n_i equal pieces giving a total of $M=\prod_i n_i$ cubes. If we let $n_i\to\infty$, then we are guaranteed a

global optimum albeit at the expense of long execution times. However, in our experience modest values of n_i tend to produce good results for the class of problems under consideration. In addition, the algorithm allows for the execution of the M local optimizations in parallel. After all M optimizations are completed, the optimal $\mathcal{P}$ for D_F is returned.

A computation of D_F for the computational model of ballistic impact described in Section 3.2 was performed using the lattice-Powell method. The physical system is described in Section 3.1 and consists of a steel sphere impacting an aluminum plate of thickness h at a speed v and obliquity $\theta=0$ from plate normal. We assume that the thickness and speed are independent random variables and calculate subdiameters for h and v. Each subdiameter calculation utilized M local searches constrained to the input parameter range chosen for the UQ analysis, namely, $h\in[32,63]$ mils and $v\in[200,400]$ km s^{-1}. These ranges were partitioned into $n_h=5$ and $n_v=3$ intervals for an $M=15$ lattice. While this partition is somewhat coarse, it was selected due to the rough planarity of the response function F in the region of interest. Optimizations terminated when the relative error between the tenth previous generation and the current generation was less than a prespecified tolerance. The maximum number of iterations was set to 20 in order to terminate optimizations meandering in a shallow well. Computations were performed in parallel on sixteen cores: two evaluations of F at a time utilizing eight cores each. This scheme trivially scales to much larger computational resources. The computation of D_F was expedited by storing previously computed values of $F(X)$ as pairs $\langle X,F(X)\rangle$ in a table. Given an X', a match is the smallest $|X-X'|<r$, if any, for some fixed search radius r. This radius must be chosen small enough so that the optimization converges to the true value. If there was no match, then $F(X')$ was computed from scratch and stored in the table. For h and v, the radii were $r_h=2.5$ and $r_v=20$, respectively.

The results of the computation of D_F are presented in Section 3.4. Each of the M optimizations reached numerical convergence in roughly five iterations. The total number of evaluations of F required over all M optimizations was 23,000 for D_F^h and 36,000 for D_F^v. All but 1000 of the F evaluations were already in the table. Increasing the n_i to produce a finer lattice showed no significant improvement. As argued in Section 2.4, the large number of function evaluations required for the computation of D_F clearly points to the infeasibility of performing the uncertainty quantification analysis solely on the basis of laboratory testing.

3.3.3. Computing the modeling error D_{G-F}

Computing the model error diameter D_{G-F} also requires solving N global optimization problems for the evaluation of the subdiameters $\mathrm{osc}_i^2(G-F,E)$, and is formulated similarly to the computation of D_F. However, each evaluation of G requires an experiment and, therefore, we need to modify our optimization strategy in order to minimize these evaluations. One approach is to use expert judgement to choose an appropriate starting point for a local optimization algorithm with a fast convergence rate. Alternatively, if an analytical surrogate function G_s is available that provides a good approximation of the experimental function G, we can perform a calculation of D_{G_s-F} to determine a good set of starting points for the evaluation of D_{G-F}. In our calculations, we use a surrogate fitted to the experimental data presented in Section 3.1 for guidance on initial starting points for the D_{G-F} computation. The surrogate is fitted to the experimental data by the least squares method. The least square residual is 6.14 mm^2 and the maximum difference between the experimental measurements and surrogate is 1.86 mm^2. This difference is indeed small compared to the empirical mean of 48.08 mm^2.

For the computation of D_{G_s-F}, the lattice-Powell solver described in the previous section was used with $M=15$ cubes. Each of the optimizations reached numerical convergence in less than three iterations. The total number of function evaluations for the computation of $D^h_{G_s-F}$ was 6200, and 10,400 evaluations were required for $D^v_{G_s-F}$. The values for $D^v_{G_s-F}$ and $D^h_{G_s-F}$ were sorted by size, thus producing an ordered set of local optima based on the surrogate. The five largest corresponding input parameter sets $\mathcal{P}$ were selected and used as starting points for the computations of D^v_{G-F} and D^h_{G-F}. Again, each of the optimizations reached numerical convergence in less than three iterations, and less than 40 function evaluations of G (experiments) were required to obtain the final modeling-error diameter. It bears emphasis that each of these 40 experiments were carried out concurrently with the model simulations and according to input parameters prescribed by the optimization iteration. In particular, the precise sequence of experiments required was not known *a priori*, and the experimental data was supplied *on-demand* during the course of the optimization iteration.

3.4. UQ analysis

The results of the UQ analysis are shown in Table 4. The operating range of the gun assumed in the analysis corresponds to an impact velocity in the range of 200–400 m/s. This operating range excludes the cliff immediately above the ballistic limit in the curve of perforation area versus impact velocity for plates with a thickness in the range of 0.81–1.62 mm, cf. Fig. 3. The average perforation area is 48.08 mm ^{2}, computed by reusing the tests run for the calculation of D_{F-G}. The margin hit due to estimation of mean performance by an empirical mean is 6.68 mm^2. The total model diameter D_F, which measures the variability of the perforation area as predicted by the model, is computed to be 7.79 mm^2, which is a small fraction of the average perforation area. Thus, the model predicts a modest variation of the perforation area over the entire range of impact velocities and plate thicknesses. Likewise the total modeling-error diameter D_{F-G}, which measures the 'badness' of the model as compared with experiment, is computed to be 6.36 mm^2, which is also a small fraction of the average perforation area. In addition, the amount of scatter in the experimental measurements is small and is neglected in

Table 4
Summary of results from the UQ analysis of a spherical S-2 tool-steel projectiles/6061-T6 aluminum plate system. The lethality of the projectiles for plate thicknesses in the range 0.81–1.62 mm, and impact velocities in the range 200–400 m/s, can be certified with exceedingly large confidence.

Model diameter D_F	Thickness	5.92 mm^2
	Velocity	5.06 mm^2
	Total	7.79 mm^2
Modeling-error diameter D_{F-G}	Thickness	4.78 mm^2
	Velocity	4.19 mm^2
	Total	6.36 mm^2
Total uncertainty $U=D_F+D_{F-G}$	14.15 mm^2	
Empirical mean $\mathbb{E}_m[G]$	48.08 mm^2	
Margin hit α ($\epsilon'=0.1\%$)	6.68 mm^2	
Confidence factor M/U	2.93 mm^2	

the computation of uncertainties. With this approximation, the total perforation-area uncertainty is $U=D_F+D_{F-G}=14.15\ \text{mm}^2$, whereas the total margin is $M=[\mathbb{E}_m[G]-\alpha]_+=41.4\ \text{mm}^2$, which gives an M/U ratio, or confidence factor, of 2.93.

This is a large confidence factor indicative of an exceedingly small probability of no-perforation of the plate over its entire operating range. Indeed, McDiarmid's inequality gives an upper bound of 3.5×10^{-8} for the probability of failure, or no-perforation, of the system. The lethality of the system is thus certified with a large degree of confidence (i.e. 99.9% since $\epsilon'=0.1\%$ is larger than 3.5×10^{-8}). Moreover, the iterative computation of the modeling-error diameter D_{F-G} converges quickly and requires 40 function evaluations, a clear indication of the high fidelity of the model, whereas the computation of the model diameter D_F requires in the order of 59,000 function evaluations. It seems reasonable to assume that a direct determination of the system diameter D_G by means of laboratory testing would require a comparable number of tests ($\sim 59,000$). The net pay-off of model-based certification, and the feasibility of the DoD UQ protocol, may thus be identified with the striking reduction (from $\sim 59,000$ to ~ 40) in the number laboratory tests that it affords.

3.5. Optimal UQ analysis

McDiarmid's inequality is not the only certification method that can be applied to the diameter and mean performance data summarized in Table 4. Alternatively, the *Optimal Concentration Inequalities* of Owhadi et al. (submitted for publication) can be applied. Since the McDiarmid subdiameters are seminorms and satisfy the triangle inequality, Table 4 implies that the subdiameters of G with respect to plate thickness h and impact velocity v are bounded as follows:

$$D_G^h \le (5.92+4.78)\ \text{mm}^2 = 10.70\ \text{mm}^2 \tag{35}$$

$$D_G^v \le (5.06+4.19)\ \text{mm}^2 = 9.25\ \text{mm}^2 \tag{36}$$

We also conclude that, with probability at least 0.999 on the experimental samples, the mean of G over the specified parameter ranges satisfies

$$\mathbb{E}[G] \ge (48.08-6.68)\ \text{mm}^2 = 41.40\ \text{mm}^2 \tag{37}$$

It then follows from Owhadi et al. (submitted for publication, Theorem 4.2) that, with probability at least 0.999 on the experimental samples, the maximal probability of non-perforation is 0, and the perforation area G is in fact strictly positive throughout the domain of certification (this result can also be stated with probability 1 using the fact that G is almost surely bounded from below by its empirical mean minus its maximum variation $D_G^h+D_G^v$).

Observe that the information associated with the (bounds on) subdiameters D_G^h and D_G^v is deterministic in nature (i.e., these subdiameters are solely functionals of G and do not depend on $\mathbb{P}$) whereas the information on $\mathbb{E}[G]$ is a functional of both, G and $\mathbb{P}$ (the measure of probability on velocity and thickness). In situations where the mean is unknown, Owhadi et al. (submitted for publication, Theorem 4.2) provide a bound on the probability of non-perforation as a function of $\mathbb{E}[G]$ that is optimal, given the information contained in the (bounds on) subdiameters D_G^h and D_G^v and the knowledge that h and v are independent random variables. This optimal bound is

$$\mathbb{P}[G=0] \le \begin{cases} 0 & \text{if } D_G^h+D_G^v \le \mathbb{E}[G] \\ \frac{(D_G^h+D_G^v-\mathbb{E}[G])^2}{4D_G^hD_G^v} & \text{if } |D_G^h-D_G^v| \le \mathbb{E}[G] \le D_G^h+D_G^v \\ 1-\frac{\mathbb{E}[G]}{\max(D_G^h,D_G^v)} & \text{if } 0 \le \mathbb{E}[G] \le |D_G^h-D_G^v| \end{cases} \tag{38}$$

Using the numerical values given in (35) we obtain that

$$\mathbb{P}[G=0] \leq \begin{cases} 0 & \text{if } 19.95 \text{ mm}^2 \leq \mathbb{E}[G] \\ \left(1-\frac{\mathbb{E}[G]}{19.9 \text{ mm}^2}\right)^2 & \text{if } 1.45 \text{ mm}^2 \leq \mathbb{E}[G] \leq 19.95 \text{ mm}^2 \\ 1-\frac{\mathbb{E}[G]}{10.7 \text{ mm}^2} & \text{if } 0 \leq \mathbb{E}[G] \leq 1.45 \text{ mm}^2 \end{cases} \tag{39}$$

We observe from Eqs. (39) and (40) that if $\mathbb{E}[G] \leq D_G^h - D_G^v$ then the optimal bound on the probability of non-perforation does not depend on D_G^v. More precisely, if $\mathbb{E}[G] \leq 1.45$ mm^2, reducing the uncertainty bound on velocity corresponding to D_G^v down to zero does not change the optimal bound $1-\mathbb{E}[G]/10.7$ mm^2 on the probability of non-perforation. This is a non-trivial phenomenon having its origin in the fact that D_G^h and D_G^v act as constraints for an infinite-dimensional optimization problem and extremizers may not live on boundaries defined by these constraints. This phenomenon can be understood as a *screening effect* where uncertainties in velocity are completely screened by uncertainties in thickness. From a broader point of view, these results show that uncertainties in input parameters, which propagate to output uncertainties in the classical sensitivity analysis paradigm, may not do so when the transfer functions (or probability distributions) are imperfectly known. Observe that this is a stronger statement than saying that input uncertainties do not add up on output. We refer to Owhadi et al. (submitted for publication) for additional discussions of this phenomenon.

4. Summary and concluding remarks

The UQ case study presented in the foregoing is chiefly concerned with the assessment of the feasibility of a *data-on-demand* (DoD) UQ protocol based on concentration-of-measure probability-of-failure inequalities. The assessment is based on a particularly simple system configuration consisting of 6061-T6 aluminum plates struck spherical S-2 tool steel projectiles at ballistic impact speeds. The system's inputs are the plate thickness and impact velocity and, for simplicity, the perforation area is chosen as the sole performance measure of the system. The objective of the UQ analysis is to certify the lethality of the projectile, i.e., that the projectile perforates the plate with high probability over a prespecified range of impact velocities and plate thicknesses. The requisite experimental input for uncertainty quantification was supplied by Caltech's GALCIT Powder-Gun Plate-Impact Facility. A principal emphasis of the DoD UQ protocol under consideration concerns the achievement of rigorous certification through modeling and simulation with a minimum of testing. This emphasis places a high premium on high-fidelity physics-based models capable of effecting accurate predictions. In our study we have used the Optimal-Transportation MeshFree (OTM) method (Li et al., 2010) extended to account for contact and fracture in order to simulate terminal ballistics (cf. also Li et al., 2012). The materials were described by means of engineering viscoplasticity models with power-law hardening, rate-sensitivity and thermal softening combined with a Mie–Gruneisen equation of state. Given the low degree of scatter in the data, the total performance uncertainty reduces to the sum of the model diameter and the modeling-error diameter. The calculation of these diameters was carried out using the lattice-Powell method of McKerns et al. (2009, accepted).

The net outcome of the UQ analysis is an M/U ratio, or confidence factor, of 2.93. This is a large confidence factor indicative of an exceedingly small probability of no perforation of the plate over its entire the operating range. The lethality of the system can thus be certified with an exceedingly large degree of confidence. Indeed, McDiarmid's inequality gives an upper bound of 3.5×10^{-8} for the probability of failure, or no-perforation, of the system. This low probability of failure showcases the ability of concentration-of-measure inequalities to deal effectively with rare failure events. The high-confidence in the successful operation of the system afforded by the analysis, together with the small number of tests (40) required for the determination of the modeling-error diameter, establishes the feasibility of the DoD UQ protocol as a rigorous yet practical approach for model-based certification of complex systems. Moreover, the striking reduction in the number of tests required for rigorous certification (from $\sim 59,000$, as estimated by the number of function evaluations required for the computation of the system diameter D_F, to ~ 40, the number of tests required for the determination of the modeling-error diameter D_{F-G}) vividly showcases the high potential pay-off afforded by model-based certification.

A number of attributes of the DoD UQ protocol are worth emphasizing further. Firstly, the quantitative and reductionist qualities of the protocol are remarkable. Thus, the protocol supplies—and reduces the quantification of uncertainties to—two quantitative figures of merit: the system diameter D_F and the modeling-error diameter D_{F-G}. The former provides a measure of the oscillation of the response over the operating range of the system, i.e., it measures how well that response of the system can be *pinned down*, or how *predictable* the system is. The latter provides a measure of the *badness* of the model, i.e., of the extend of the divergence between model and reality. Together (in the absence of significant experimental scatter), those two measures aggregate to one single quantitative measure of system uncertainty. Evidently, many alternative uncertainty measures can be contrived, but in general there is no assurance that such *ad hoc* uncertainty measures can be taken as a basis for *rigorous* certification, which greatly detracts from their interest. What sets the system and modeling-error diameters apart from other *ad hoc* uncertainty measures is precisely their property that, together, they supply a rigorous upper bound on the probability of failure of the system and, therefore, can be taken as a basis for rigorous certification.

A common practice in uncertainty quantification is to compute sensitivities of performance measures with respect to input parameters. The present work shows that such sensitivities are not sufficient for rigorous certification in general.

Thus, the computation of the system diameter requires the careful quantification of variations in performance resulting from *arbitrary and finite* perturbations of the input parameters. Only such *nonlinear* sensitivity analysis is capable of quantifying uncertainties in system performance in a manner that enables rigorous certification.

It should also be carefully noted that uncertainty quantification based solely on mathematical or computational models is *a fortiori* incomplete as it does not account for the uncertainty due to the lack of model fidelity. In the DoD UQ protocol developed in this work, the modeling-error diameter supplies precisely the requisite quantitative measure of uncertainty. It is interesting to contrast such unambiguous quantification of model fidelity with traditional *validation*, which correctly emphasizes the need to assess models against experimental data but, in general, leaves unspecified the precise quantitative means by which such a comparison should be effected. This ambiguity is not surprising, since the specification of a precise validation measure must necessarily be *goal oriented*, i.e., whether a model is of sufficient fidelity or not depends critically on the intended application. Here again, certification provides a clear mathematical goal, namely, the bounding of probabilities of failure, that effectively disambiguates the choice model-fidelity measure. In particular, within this certification-oriented framework a model is of sufficient fidelity, and thus *validated* for the system and application at hand, if the resulting modeling-error diameter is sufficiently small, i.e., it constitutes a small fraction of the total uncertainty budget.

Finally, we point to some limitations of the approach and areas for further extensions. The system considered in this work, namely, the spherical S-2 tool steel projectile/6061-T6 aluminum target plate terminal ballistics system, is particularly clean in that all the input parameters are controllable and, by virtue of the large size of the projectile relative to the target plate thickness, the experimental data exhibits negligibly small scatter. Evidently, the controllability of the parameters, i.e., the ability to *dial in* specific values of the system for purposes of laboratory testing, is essential to the practical implementation of the DoD UQ protocol. In addition, the lack of significant experimental scatter greatly simplifies the quantification of system uncertainties. In a sequel to this work (Adams et al., in press), we extend the DoD UQ protocol to systems where those simplifying circumstances do not arise, i.e., to systems where some of the input systems cannot be dialed in at will and where the experimental data exhibits non-negligible scatter. It should also be pointed out that the McDiarmid's inequality employed in the present study tends to degrade, in the sense of loss of tightness, when the operating range of the system includes a cliff in the response function. In such cases, the determination of tight probability-of-failure bounds requires more elaborate bounding methods such as those based on the partitioning of the input domain (Sullivan et al., 2011).

Acknowledgments

The authors gratefully acknowledge the support of the Department of Energy National Nuclear Security Administration under Award Number DE-FC52-08NA28613 through Caltech's ASC/PSAAP Center for the Predictive Modeling and Simulation of High Energy Density Dynamic Response of Materials.

References

Adams, M., Lashgari, A., Li, B., McKerns, M., Mihaly, J., Ortiz, M., Owhadi, H., Rosakis, A.J., Stalzer, M., Sullivan, T.J. Rigorous model-based uncertainty quantification with application to terminal ballistics: systems with uncontrollable inputs and large scatter. Journal of the Mechanics and Physics of Solids, doi:10.1016/j.jmps.2011.12.002. In press.

Arroyo, M., Ortiz, M., 2006. Local maximum-entropy approximation schemes: a seamless bridge between finite elements and meshfree methods. Int. J. Numer. Meth. Eng. 65, 2167–2202.

Blue Ribbon Panel on Simulation-based Engineering Science, 2006. Simulation-Based Engineering Science: Revolutionizing Engineering Science Through Simulation. Technical Report, National Science Foundation, May, available at ⟨http://www.nsf.gov/pubs/reports/sbes_final_report.pdf⟩.

Boucheron, S., Bousquet, O., Lugosi, G., 2004. Concentration inequalities. in: Bousquet, O., von Luxburg, U., Rätsch, G. (Eds.), Advanced Lectures in Machine Learning, Springer, pp. 208–240.

Committee on the Evaluation of Quantification of Margins and Uncertainties Methodology for Assessing and Certifying the Reliability of the Nuclear Stockpile, 2008. Evaluation of quantification of margins and uncertainties for assessing and certifying the reliability of the nuclear stockpile. Technical Report, National Academy of Science/National Research Council (NAS/NRC). National Academy Press, Washington, DC (available at ⟨http://www.nap.edu/catalog.php?record_id=12531⟩).

Evans, L.C., 2004. An introduction to stochastic differential equations. Available at ⟨http://math.berkeley.edu/∼evans/SDE.course.pdf⟩.

Helton, J.C., 2009. Conceptual and computational basis for the quantification of margins and uncertainty. Technical Report SAND2009-3055.Albuquerque, Sandia National Laboratories (available at ⟨ http://infoserve.sandia.gov/sand_doc/2009/093055.pdf⟩).

Kane, C., Marsden, J.E., Ortiz, M., 1999. Symplectic-energy-momentum preserving variational integrators. J. Math. Phys. 40 (7), 3353–3371.

Karlin, S., Studden, W.J., 1966. Tchebycheff systems: with applications in analysis and statistics. Pure and Applied Mathematics, vol. XV. , Interscience Publishers John Wiley & Sons, New York, London, Sydney.

Kendall, D.G., 1962. Simplexes and vector lattices. J. London Math. Soc. 37 (1), 365–371.

Leader, D.E.S., 2005. Quantification of Margins and Uncertainties (QMU). Technical Report JSR-04-330, The MITRE Corporation, 7515 Colshire Drive McLean, Virginia, 22102.

Ledoux, M., 2001. The concentration of measure phenomenon. Mathematical Surveys and Monographs, vol. 89, American Mathematical Society.

Li, B., Habbal, F., Ortiz, M., 2010. Optimal transportation meshfree approximation schemes for fluid and plastic flows. Int. J. Numer. Meth. Eng. 83 (12), 1541–1579.

Li, B., Kidane, A., Ravichandran, G., Ortiz, M., 2012. Verification and validation of the Optimal Transportation Meshfree (OTM) simulation of terminal ballistics. Int. J. Impact Eng. 42, 25–36.

Lucas, L., Owhadi, H., Ortiz, M., 2008. Rigorous verification, validation, uncertainty quantification and certification through concentration-of-measure inequalities. Comput. Methods Appl. Mech. Eng. 197, 4591–4609.

Lugosi, G., 2006. Concentration-of-measure inequalities. Lecture Notes ⟨http://www.econ.upf.edu/∼lugosi/anu.pdf⟩.

Marshall, A.W., Olkin, I., 1979. Inequalities: theory of majorization and its applications. Mathematics in Science and Engineering, vol. 143, Academic Press Inc., Harcourt Brace Jovanovich Publishers, New York.

McDiarmid, C., 1989. On the method of bounded differences. Surveys in Combinatorics, London Mathematical Society Lecture Note Series, vol. 141, Cambridge University Press, pp. 148–188.
McKerns, M.M., Strand, L., Sullivan, T., Fang, A., Aivazis, M.A.G., Building a framework for predictive science. In: Proceedings of the 10th Python in Science Conference, accepted for publication.
Michael McKerns, Patrick Hung, Michael Aivazis, 2009. Mystic: a simple model-independent inversion framework. 〈http://dev.danse.us/trac/mystic〉.
Owhadi, H., Sullivan, T.J., McKerns, M., Ortiz, M., Scovel, C. Optimal uncertainty quantification, submitted for publication. 〈http://arxiv.org/abs/1009.0679〉.
Pilch, M., Trucano, T.G., Helton, J.C., 2006. Quantification of Margins and Uncertainties (QMU). Technical Report SAND 2006–5001, Sandia National Laboratories, Albuquerque, New Mexico 87185 and Livermore, California, 94550, September.
Powell, M.J.D., 1989. A tolerant algorithm for linearly constrained optimization calculations. Math. Program. 45, 547–566.
Powell, M.J.D., 1994. A direct search optimization method that models the objective and constraint functions by linear interpolation. in: Gomez, S., Hennart, J.P. (Eds.), Advances in Optimization and Numerical Analysis, Kluwer Academic, Dordrecht, pp. 51–67.
Schmidt, B., Fraternali, F., Ortiz, M., 2009. Eigendeformation: an eigendeformation approach to variational fracture. Multiscale Model Simul. 7, 1237–1266.
Sharp, D.H., Wood-Schultz, M.M., 2003. QMU and nuclear weapons certification: what's under the hood? Los Alamos Sci. 28, 47–53.
Sullivan, T.J., Topcu, U., McKerns, M., Owhadi, H. Uncertainty quantification via codimension-one partitioning, 2011. International Journal for Numerical Methods in Engineering 85, 1499–1521.
Sulsky, D., Chen, Z., Schreyer, H.L., 1994. A particle method for history-dependent materials. Comput. Meth. Appl. Mech. Eng. 118, 179–186.
Topcu, U., Lucas, L.J., Owhadi, H., Ortiz, M., 2011. Rigorous uncertainty quantification without integral testing. Reliab. Eng. Syst. Saf. 96, 1085–1091.
von Weizsäcker, H., Winkler, G., 1979. Integral representation in the set of solutions of a generalized moment problem. Math. Ann. 246 (1), 23–32.
Winkler, G., 1988. Extreme points of moment sets. Math. Oper. Res. 13 (4), 581–587.