Uncertainty in the Development and Use of Equation of State Models

V. Gregory Weirs*, Nathan Fabian*, Kristin Potter†, Laura McNamara*, and Thomas Otahal*

*Sandia National Laboratories   †Scientific Computing and Imaging Institute

ABSTRACT

In this paper we present the results from a series of focus groups on the visualization of uncertainty in Equation-Of-State (EOS) models. The initial goal was to identify the most effective ways to present EOS uncertainty to analysts, code developers, and material modelers. Four prototype visualizations were developed to presented EOS surfaces in a three-dimensional, thermodynamic space. Focus group participants, primarily from Sandia National Laboratories, evaluated particular features of the various techniques for different use cases and discussed their individual workflow processes, experiences with other visualization tools, and the impact of uncertainty to their work. Related to our prototypes, we found the 3D presentations to be helpful for seeing a large amount of information at once and for a big-picture view; however, participants also desired relatively simple, two-dimensional graphics for better quantitative understanding, and because these plots are part of the existing visual language for material models. In addition to feedback on the prototypes, several themes and issues emerged that are as compelling as the original goal and will eventually serve as a starting point for further development of visualization and analysis tools. In particular, a distributed workflow centered around material models was identified. Material model stakeholders contribute and extract information at different points in this workflow depending on their role, but encounter various institutional and technical barriers which restrict the flow of information. An effective software tool for this community must be cognizant of this workflow and alleviate the bottlenecks and barriers within it. Uncertainty in EOS models is defined and interpreted differently at the various stages of the workflow. In this context, uncertainty propagation is difficult to reduce to the mathematical problem of estimating the uncertainty of an output from uncertain inputs.

Index Terms: H.5.1 [Information Systems]: Multimedia Information—Evaluation/methodology; J.2 [Computer Applications]: Physical Sciences and Engineering

1 INTRODUCTION

The research objective of the Material Model Uncertainty Visualization (MMUV) project is to develop effective techniques for displaying the uncertainty of material model data, with the eventual goal of providing a software tool to users in the in materials modeling community at Sandia National Laboratories. An example of a recently developed capability for three-dimensional equation-of-state surface visualization is shown in Fig. 1. This image was generated by Prism, a plugin that is distributed with ParaView [15]. The key convenience that Prism provides to users is that data from simulations can be simultaneously displayed in both the physical geometry space and the thermodynamic (EOS) space.

   The initial thrust of the research in this paper was to address how one should add uncertainty information to the display of an

---

*e-mail:{vgweirs, ndfabia, tjotaha, lamcnam}@sandia.gov
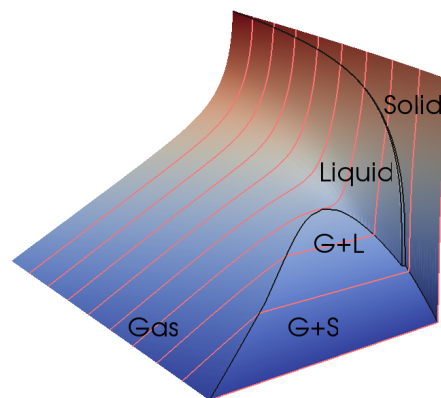
†e-mail:kpotter@sci.utah.edu

Figure 1: Pressure is shown as a function of density and temperature for a fictitious material. All three dimensions are logarithmic. Logarithmically distributed isotherms are shown as pink contours. Black lines identify discontinuous phase boundaries. Regions of different phases are labeled, including mixed regions ("G+L" for mixed gas and liquid, "G+S" for mixed gas and solid.)

EOS surface, such as in Fig. 1. Narrowly defined, this question would not be too difficult to answer. As we describe in this article, however, the focus group approach revealed that to address this particular community's needs, just presenting the uncertainty information would not be sufficient. The key challenges derive from the technical complexity of material modeling, and the related organizational complexity and diversity of roles across the material model stakeholder community. In this context, uncertainty quantification is not (just) a mathematical problem in which probability distributions of uncertain inputs are propagated through a model to compute output probability distributions. Instead, it is the communication and integration of various technical uncertainties about material behavior from material science theoreticians and experimentalists, through software developers and analysts, to decision makers. In this sequence, the material uncertainties are convolved with uncertainties from many other sources; the material model development, software development, numerical simulations, and decision may span decades; and rarely have all the people involved communicated with each other directly.

   In the next subsection, the nature of material models is described, with specific emphasis on EOS models, their associated uncertainties, broadly classified by their sources, and their role in engineering simulations. Following subsections describe the focus group methodology and our process in particular.

1.1 Material Modeling

A material model describes the behavior of a specific material or class of materials. A material model is often comprised of a number of submodels. The different submodels may describe different physical or chemical phenomena, for example equation-of-state models describe equilibrium thermodynamics, while solid mechanics models describe the relationship between stress and deformation. A solid mechanics model may be used in concert with other submodels, e.g., for the yield strength, fracture, or damage. Submodels may cover different regimes; an EOS model for nitrogen at lower temperatures might not include chemical reactions or ionization, but a high temperature model might not include liquid or solid phases. Material models range from the relatively simple to the highly complex, and the development of each submodel requires specialized expertise. Each model is based on theories and assumptions, but experimental data is usually required to determine or calibrate model parameters for a specific material, or for a particular regime.

In this work our interest is in EOS and solid mechanics material models. EOS models describe relationships between thermodynamic variables such as pressure, density, temperature, internal energy, and the speed of sound in a material. Solid mechanics models describe relationship between the stress state and the strain (or deformation) of a material in the solid state. In this paper we focus on EOS models because we have existing techniques for visualizing EOS surfaces, e.g., Fig. 1. Solid mechanics models are more complicated because stress and strain are 3×3 tensor quantities, and because solid mechanical behavior is often dependent on the deformation history of the material. Visualizing solid mechanics models, or at least the quantities of greatest interest to analysts, will be the subject of future work.

Material model domain complexity makes the development of useful and usable uncertainty visualizations quite challenging. Uncertainty is an abstract concept, even when the domain of interest and its associated uncertainties are relatively uncomplicated and/or well-studied. Material science is complicated and significant knowledge gaps and aleatory uncertainties exist for most materials; yet all technology development requires at least minimal characterization of material behavior and properties.

1.1.1 Equation-of-State Models and Sources of Uncertainty

An equation of state describes relationships between thermodynamic variables for a given material. Given any two variables, all other variables can be computed through the EOS under the assumption of thermodynamic equilibrium. The theory of thermodynamics provides some of these relationships [5, 13], such as the Maxwell relations. EOS models can cover a very wide range of conditions, and different physical phenomena dominate material behavior in different regimes [7]. The relevant physics for a particular regime provide additional equations and relationships. At this level of description, uncertainties arising from the assumptions or approximations in the theory, as reflected in the equations, are called *model form uncertainties*.

The various thermodynamics theories often result in models that have parameters that must be provided to fully specify the EOS. Some parameters may be independent of the specific material, but in many cases, the particular parameter values distinguish the model for one material from models for other materials. When a theory does not provide the parameter values, they must be determined by experimental measurements or by more sophisticated physics theories or models. These experiments and more sophisticated models have their own sources of error, which appear in the material model as uncertainty in the specific values of the model parameters, or *parameter uncertainty*.

Sometimes the boundaries between regimes are sharp, such as the phase transition between ice and water, but in other cases there is a gradual transition, for example between liquid and gas phases above the critical point. Again, see Fig. 1 for an illustration. For a multiphase material model, EOS models for different regimes are blended or combined to describe the behavior across a number of regimes. Materials science does provide some theory to determine where discontinuous phase transitions occur, and constraints on thermodynamic variables across the phase transitions; however, for various reasons EOS models may not locate these transitions accurately. Depending on the particular material model, the uncertainty in the location of phase transitions may be classified as model form or parameter uncertainty, but in either case it can have a large effect on the accuracy of the material model.

For the work considered here, material models are input data for engineering simulation codes; the material model is not a result of the engineering simulation. In these codes, EOS models are often stored in tabular form, e.g., pressure values are stored for a number of discrete density-temperature points. For an arbitrary density and temperature, the pressure is interpolated from nearby points in the table. Over the range of densities and temperatures, the pressure can be shown as a surface. Most tabular EOS models contain several thermodynamic variables; the most common are density, temperature, internal energy, and pressure. This format for representing the EOS introduces interpolation error, which may be reduced by sampling the underlying material model at more closely spaced points, but cannot be eliminated in practice. A second potential problem is that while the consistency between the different variables can be maintained at the points in the table, consistency is difficult to enforce for interpolated values [13, 6]. Finally, the most common tabular format, "SESAME" [11], uses a rectangular grid of density and temperature values, which was originally driven by limited choices for data structures; tables often provide data outside of the regimes for which a model has been developed and tested simply because of the constraints of the SESAME format. While the uncertainties associated with the tabular format would seem to be easy to address, they remain a source of significant concern for practitioners.

The results of engineering simulations depend on the EOS models used, but the sensitivity of particular responses of interest (extracted from the simulation results) to these models varies widely. One reason is that the response can be chosen from an essentially infinite number of possibilities; the analyst may be interested in a response that depends very strongly on a material model, or one that is practically insensitive to the material model. If the response is insensitive to a material model, then the uncertainties in that material model will be irrelevant. Assuming the response does depend on the material model, a second factor is the thermodynamic conditions the material experiences during the simulation; a material model may have regions of large uncertainty, but the material in the simulation might not enter those regions.

The uncertainty originating from the material models is convolved with uncertainties from other sources during the simulation. For example, other simulation inputs will include the geometry of the device or control volume, detailed specification of the boundary conditions and the initial conditions, and various numerical algorithm control parameters, and all of these have their own model form and parameter uncertainties. During the simulation, these uncertainties interact, and the uncertainty in the response of interest is a function of the individual uncertainties and their generally nonlinear interactions.

1.1.2 The Mie-Grüneisen Equation-of-State Model

In this work we have used a Mie-Grüneisen (MG) model for aluminum to provide test data for prototype visualizations. The MG EOS describes a single phase, and has a number of parameters that engineering codes allow the user to specify. The model is defined

by

$$P(\rho,E) = P_R(\rho) + \Gamma_0\rho_0(E - E_R(\rho)) \quad (1)$$
$$E(\rho,T) = E_R(\rho) + C_V(T - T_R(\rho)) \quad (2)$$

where ρ is the density, P is the pressure, E is the internal energy per unit mass, T is the temperature, Γ_0 is the Grüneisen parameter, ρ_0 is the reference density, and C_V is the heat capacity. The last three are material-specific constants. The MG model relates the pressure and energy to a reference Hugoniot, a special curve comprised of the locus of shock states for a given initial state; the subscript R denotes the reference Hugoniot. For many metals in the solid phase, a linear Hugoniot relation expressed in the shock velocity, U_s, and the particle velocity, u_p provides an excellent fit to experimental data:

$$U_s = C_s + S_1 u_p \quad (3)$$

where C_s is the speed of sound at reference conditions and S_1 is the slope of the linear relationship (both material-specific constants.) The Rankine-Hugoniot equations relate the states on each side of a shock wave, and can be used to derive

$$P_R(\rho) = P_0 + \rho_0 U_s u_p \quad (4)$$
$$E_R(\rho) = E_0 + (P_R(\rho) + P_0)\mu/2\rho_0 \quad (5)$$
$$T_R(\rho) = e^{\Gamma_0\mu}[T_0 + C_V^{-1}\int_0^\mu e^{-\Gamma_0\mu}\mu^2 U_s \frac{dU_s}{d\mu}d\mu], \quad (6)$$

where $\mu = 1 - \rho_0/\rho$ and P_0, E_0, and T_0 define the initial state.

As we have noted, uncertainties are abundant throughout the materials models workflow. For clarification, some of the different types of uncertainties are now identified in the MG EOS model. As an example of model form uncertainty, many materials are not are well-described by the linear $U_s - u_p$ relationship, and no calibration of the parameters of the MG model will be able to compensate for the error or uncertainty introduced by the linear relationship over a significant range of densities. Parameter uncertainties are associated with the material-specific model parameters Γ_0, C_v, ρ_0, C_s, S_1, P_0, E_0, and T_0. For this single-phase model, the range of validity is not provided.

1.2 Research Method

A focus group is a structured group interview, facilitated by a moderator, in which participants explore an issue or set of issues of research importance. Because so many disciplines use focus groups, approaches to designing, deploying, and analyzing focus groups vary tremendously. However, all focus groups begin with the same basic principle: that exchanges among participants facilitate the expression of ideas, knowledge, behaviors, and opinions that may be invisible to individualized methods (such as a questionnaire or a one-on-one interview). Groups enable researchers to access a broader range of skills and experiences than a single respondent; listening to others express ideas and opinions can spur participants to remember and share information that might not have emerged in a one-on-one setting. Not surprisingly, focus groups are an excellent way of eliciting the kinds of information that people naturally express in group settings; or for documenting how knowledge emerges in the context of group interactions.

Focus groups play an important role in computer science research and software engineering, in ways that are germane to the goals of the MMUV project. Studies of software engineering processes have used focus groups to gather data about workflow patterns in engineering teams [9]. User- and interaction-oriented design paradigms suggest the use of focus groups to gather qualitative data on user expectations and system requirements, and to evaluate prototype interfaces [17, 4, 10]. For technology developers, focus groups also afford the opportunity to demonstrate sincere interest in user concerns. In that sense, focus groups can enhance relationships

between the user community and the technology developers by establishing a foundation for ongoing communication and exchange of information. Over time, the user community perceives itself as a dedicated stakeholder in the work of the technology developers, instead of recipients of a product tossed over the proverbial fence.

The focus group approach does have significant drawbacks: while participants often generate excellent contextual insights, focus group data is less useful for analyzing long-term trends or generalizing about large populations. Moreover, focus groups are prone to groupthink bias and social dominance bias (i.e., when one of the group members exerts undue influence on the interactions or content of the group, either consciously or unconsciously [12]. A structured script, pilot runs, careful moderation, and a sound qualitative sampling strategy can enhance the quality and dependability of focus group findings. Even so, software developers should treat focus groups as a starting point for technology design and evaluation, if only because because focus groups only capture information on what users "say they do – not how they actually do it" [14]. Other approaches, including observation, user participation in design teams, multiple prototyping and iterative re-designs, are necessary to develop technologies that people perceive as truly adoptable.

1.3 Focus Group Process

As noted above, one of the major goals of the MMUV project was the design and development of visualizations that would be usable and useful to experts who generate, interact with, or rely on material models in their work – what we have described above as the "stakeholder community." The domain complexity of material modeling and the organizational complexity of the stakeholder community are precisely what makes material modeling an interesting computational science, information visualization, and technology design problem; but they also make it difficult to understand what "usability" and "utility" mean in the many contexts of work where material model visualization might be useful. Accordingly, we decided to use focus groups as a way of gathering expert knowledge about material modeling challenges at Sandia National Laboratories, and as a way of opening a dialogue with the user community so that we could better understand the current state of practice. For specific feedback on representing uncertainty in EOS models, four visualization prototypes were developed by the MMUV project. In May, June, and July of 2011, we conducted four focus groups with participants representing the various material modeling stakeholder communities described above. Participants were technical staff at, or in some way affiliated with, Sandia.

The diversity of the material modeling stakeholder community made composition of the focus groups a bit challenging, since members of the stakeholder subcommunities tend not to interact with each other on a regular basis (more on this issue below). The success of a focus group depends on the composition of the participant pool; people need to have enough common ground that they can communicate productively about the topic of discussion, but diversity of perspective can spark insights that might not emerge in a homogeneous group. All four focus groups included representatives of each of the three primary subcommunities – material modelers, code developers, and analysts. To recruit participants, we drew on contacts from our own Sandia networks. We scheduled the focus groups and invited participants but offered no additional incentives (neither snacks nor money), and all participation was completely voluntary.

In moderating the groups, we decided to use a team facilitation approach, in which a technical leader and a process leader managed the group logistics and flow of conversation. To ensure a smooth process, we developed a script with timing notations to ensure adequate and balanced coverage of the topics of interest. In this script, we split the focus groups into four phases of discussion: an intro-

duction, a general discussion about material modeling and uncertainty; presentation and discussion of the prototypes; and a wrap-up discussion. Rather than ask participants to dive right into assessing the MMUV prototypes, we decided to prime the discussion by asking the participants to talk about the role of material modeling in their work. In doing so, participants exchanged observations about the importance of material models for engineering analyses; identified key sources of material model uncertainties; discussed the impact of material uncertainty on their work; and described strategies for representing and managing uncertainty. These exchanges set the stage for the second half of the focus group, during which the visualization prototype developers took turns presenting and discussing their prototype designs with the group participants. As expected, putting prototype designs in front of the experts generated intense discussion about the problem of understanding and managing uncertainty in material models.

In the following sections, we describe the focus groups and summarize key themes. The prototypes are described in Section 3. The participants provided detailed comments on these prototypes and concrete suggestions for enhancing visualization utility and interactivity, as discussed in Section 4. However, the participants' discussions also illuminated the complicated organizational and technical relationships through which information about material properties and performance is exchanged and incorporated into Sandia's engineering research and development. Although there was an awareness that the stakeholder community was only loosely connected, a much clearer picture emerged during our focus groups, and since it provides the context for all of the participants' comments, we begin with an overview of these relationships in the next section.

2   THE DISTRIBUTED WORKFLOW OF MATERIAL MODELING

The focus group sessions revealed a great deal about the state of practice in material modeling and the use of material models in Sandia's research and engineering domains. Material model development and use exist in a distributed information workflow: a particular material model is developed by material modelers, incorporated into a continuum engineering simulation code by code developers, and used by analysts when they run simulations for specific applications. Ultimately, a decision maker chooses actions that are informed by those simulations. Each of these different groups is a stakeholder in the material model, but their knowledge about the material behavior and their use of the material model vary widely. Naturally, the meaning of material model "uncertainty" also varies widely across these stakeholder roles.

2.1   Material Modelers

As described in Section 1, material model development begins with theories that may be incomplete, contain acknowledged gaps in applicability, or have other known deficiencies. A typical model has a number of parameters that must be determined for each particular material; these parameters are calibrated to match available experimental data, or lacking that, to data from simulations of more fundamental models such as density functional theory (DFT) or molecular dynamics (MD) simulations. These parameters might be exposed to end users of the model, or might be internal to the model and relevant only to the material model developer. Finally, while a material model may have a functional interface that accepts input values and returns output values, in many cases the model is incorporated into the engineering code in precomputed, tabular form and output values are interpolated. For a material modeler, each of these steps is a source of a different kind of error that contributes to the overall uncertainty of the model.

For a given material model, material modelers struggle to integrate the uncertainties from these different sources into a single uncertainty field for a given thermodynamic variable. Some material modelers in our focus groups were hesitant to even attempt to put a number on the sum uncertainty for any model, because they saw no constructive value in, essentially, guessing. This emphasizes the technical complexity of material modeling, and indicates that rigorous quantification of material model uncertainty is at an early stage. With that said, serious efforts are underway for quantifying model form uncertainty, parameter uncertainty, uncertainty in experimental data, and errors in discretization, and providing this information in a form that is accessible to the engineering simulation codes.

Material modelers implement their models in software libraries. A typical library life cycle begins with a minimal set of models for a particular engineering code, but over time more models are added. At some point the library is refactored so that it can be used in another code, which may have different algorithms or be intended for a different class of engineering applications. Material model libraries are reused in different engineering codes because they are expensive to develop and test. Most libraries have many contributors over a number of years or decades, sometimes from several labs and companies. The level of documentation on how a model was developed can vary significantly across models. Historically, uncertainty information has not been included in material mode libraries beyond a few very general text comments.

2.2   Code Developers

The code developers in our focus groups are programmers, but their educational degrees were PhDs in various engineering disciplines, physics, or applied math. Development and maintenance of engineering simulation codes require a solid foundation in numerical discretization techniques as well as domain-specific expertise. Regarding material modeling, the essential concern of code developers is the interaction between the material models and the discretization algorithms. The discretization algorithms assume the EOS possesses certain properties, such as convexity, or a positive speed of sound. Likewise, the EOS library often assumes it will only be given valid input data, e.g., density and temperature points for which the EOS has been validated. In practice both sets of assumptions are sometimes violated, and occasionally the code crashes and does not produce a solution. In engineering codes, developers address these cases with a number of techniques that vary widely in their theoretical credibility.

Ideally, code developers could treat a material model library as a black box and focus on the interface between the models and discretization techniques. Unfortunately, this separation of concerns is rarely possible for a number of reasons. The primary reason is that it is difficult to completely separate the material model implementation from the host code (the code that calls the model from the library.) Complicated material behavior, described by a number of different submodels, often requires specialized treatment by the algorithms in the engineering code. A particularly difficult example is the treatment of mixed-material elements, or computational cells that contain more than one material. To properly treat such an element, the engineering code will need to know what materials are included, and the phases of each material; different algorithms apply for each combination of phases. The relative positions of the materials are also important when at least one material is in the solid phase, requiring complicated interface reconstruction algorithms.

Because some models in a library might have been intended for a different engineering code, or with a different application in mind, code developers spend a significant amount of time updating material models to be consistent with other models and the host code. Consequently, the maintenance of material model libraries is often shared between code developers and material modelers. Code developers work with material modelers, but rarely do they know the authors of all the material models included in the libraries that their engineering code uses.

2.3 Analysts

Analysts are quite interested in how the uncertainty of the material model affects their simulation results, and less concerned about the material model uncertainty itself. It is analysts that view material model uncertainty as just one type of contribution among many to the overall uncertainty in their simulations. As a group, the analysts that participated in our focus groups had the most widely varying backgrounds. Some had been code developers or worked on specific aspects of material modeling, while others had only an introductory understanding of material modeling. Analysts have minimal access to information about material model uncertainties; if material modelers have this information, it is rarely in an easily accessible, published document, and analysts have minimal contact with material modelers. Analysts are more likely to contact code developers with their questions and concerns. A disturbing theme, heard repeatedly, was that analysts have little guidance on choosing among the EOS models available for the same material – they do not have a way to obtain and compare the uncertainties from two EOS models for the conditions that matter in their application.

2.4 Uncertainty Propagation Through the Workflow

Uncertainty quantification is most often presented as a mathematical problem, in which a number of uncertain inputs produce an uncertain output. If probability distributions of the inputs are known, and the output can be computed for specific values of the inputs, the probability distribution of the output can be determined. In this mathematical problem, the process that produces an output value from a set of input values is arbitrary: the outputs can be determined by controlled physical experiments, or observations for measured inputs in uncontrolled experiments (e.g., weather), or produced by a model or numerical simulation. The main impact of the process on the mathematical problem is that in most cases, it limits the amount of output values available; a model evaluation or controlled experiment may be expensive or time consuming to perform, or may provide output values for only a small sampling of input values. As a result of limited output data, estimates of output probabilities are less accurate.

In this context, the uncertainty of simulation outputs on some material model uncertainties can be examined. The uncertainty of an input (e.g., a material model parameter) is described by a probability distribution, and a variety of techniques are available to propagate a number of input uncertainties to the uncertainty of a simulation output. Nonintrusive methods rely on running a number of simulations for different input values and examining the distribution of the output quantity. Software for propagating uncertain inputs through an engineering simulation code to determine output uncertainties is available, such as DAKOTA [1]. Intrusive techniques are also available but are very difficult to retrofit to existing engineering codes. Considering uncertain material models, this approach is effective for examining parameter uncertainties and to a lesser degree, the uncertainties that can be represented by discrete inputs.

While this mathematical formulation of the uncertainty quantification is necessary and useful, it does not capture all of the complexity of the the material modeling distributed workflow. It is daunting to consider how many parameters would be required to describe the many different material model uncertainties individually, the level of effort required to accurately characterize the probability distribution of each, and the number of simulations required to adequately sample the input hypercube. And yet, the effects of model form uncertainty and interpolation error for tabular EOSs on simulation results are not accessible in this approach. As noted earlier, material modelers struggle to determine the total uncertainty in a material model, and are only beginning to think about how to express that uncertainty in a material model library; however, this is the most pragmatic next step. Ultimately, code developers would incorporate this uncertainty into discretization and solution algorithms to provide output uncertainties directly to analysts, but this a research area in a very early stage.

In the near future, the propagation of material model uncertainties through the distributed workflow is more likely to be through people, more qualitative than quantitative, and more interactive than automated. Software tools for visualizing material model uncertainty will aid communication between the various stakeholders, as well as help each group with their specific tasks. This state of affairs is simply a consequence of the technical complexity and diversity of material science.

3 PROTOTYPES

To facilitate discussion within our focus groups, we developed four visualization prototypes, each of which present uncertainty within a material model in a unique way. Participants evaluated specific features of each prototype and described scenarios in which different elements could prove helpful. We present each of the prototypes used in the focus groups, using the Mie-Grüneisen EOS model. Because it models just the solid phase, the surfaces in the prototypes will not exhibit the discontinuities observed in the pressure surface in Fig. 1.

As noted earlier, few, if any, material model libraries contain EOS uncertainty data, and in fact, it is not clear how best to express such information so that engineering codes could use it. To obtain data for our prototype visualizations, we have treated two parameters of the MG EOS as uncertain. From several sets of values for these uncertain parameters, different realizations of the model were generated. Each model realization was evaluated at a number of density and temperature values to produce a pressure surface, i.e., a pressure surface is the pressure as a function of density and temperature, and a different surface was obtained for each pair of uncertain parameter values. The two uncertain parameters are C_s and S_1 in the linear Hugoniot (Eq. 3), which were simultaneously varied within about 1% of the nominal values for aluminum. The differences in the surfaces produced are quite small, but are smooth functions of density and temperature. From these surfaces, several of the prototypes calculated the mean and standard deviation, also as functions of density and temperature.

One could argue the merits of this approach to generating a test dataset, but the key feature is that the uncertainty is a field variable, not a constant value. Whether the standard deviation is the best measure to represent uncertainty, or if the mean pressure surface is a better reference than the nominal surface, are important questions for stakeholders using the visualizations; for analysts interpreting the visualizations it is critical to understanding what they are seeing. However, these issues do not affect the comparisons of the visualization prototypes.

For many metals, the linear Hugoniot relationship describes experimental data very well, so the uncertainties in the parameters are small. In most of the prototypes the uncertainty has been rescaled or exaggerated to improve the display. (In a production tool this rescaling would be controlled by the user.) There are many other use cases (more complex models and surfaces, uncertain phase boundaries, or just larger parameter uncertainties) that might raise different utility issues for the visualization prototypes, but the dataset generated represents a realistic situation.

Finally, recall that the goal of the MMUV prototypes is to visualize the uncertainty in the material model itself, and not the uncertainty in engineering simulation results. We expect that simulation data (with or without uncertainty information) could be displayed in addition to the material model uncertainty.

3.1 Point Cloud

The first prototype, shown in Fig. 2, implements a technique presented in [8] that represents uncertainty by a cloud of points. The

emphasis of this prototype is to show uncertainty in in direct relation to the surface. The distance between each point and the surface is random, but within a range defined by the local uncertainty associated with a point on the surface. The algorithm creates a cloud of points that extends further away from the surface in regions of high uncertainty, and remains closer to the surface in regions of lower uncertainty. Additionally, the transparency of each point can be varied with the uncertainty, so points of higher uncertainty become more transparent. This creates a visual effect that feeds the expectation of the human visual system, where regions of low uncertainty appear crisp and solid, and regions of higher uncertainty appear hazy and indistinct. Finally, the points can be colored by another scalar value, such as temperature or internal energy, and thus simultaneously convey scalar data in addition to uncertainty information.
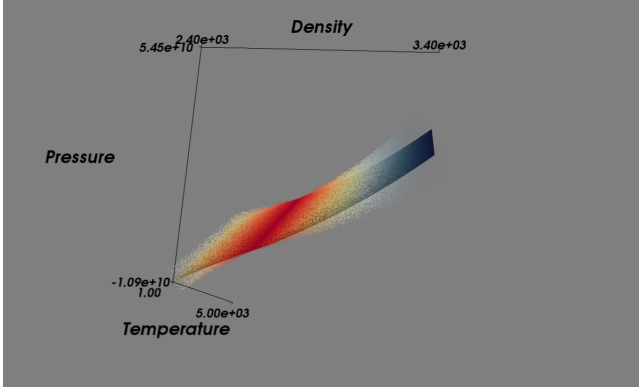

Figure 2: View of the point cloud prototype. The mean pressure surface is viewed nearly from the edge, and is colored in red for the lowest uncertainty, and dark blue for the highest. The point cloud is colored in the same way, but in addition, the points are opaque for the lowest uncertainty and transparent for the highest.

3.2 Surface Animation

The second prototype is based on a technique described by [3] that uses animated visual vibrations of the points defining a surface to show uncertainty in the surface location. A fixed, semi-transparent mean surface provides a reference. In the animation, another solid surface sweeps through one standard deviation above and below the mean surface, with the animation transition defined by the sinusoid equation:

$$V = \frac{c\sin(2\pi pt + \frac{\pi}{2}) + 1}{2} + f \qquad (7)$$

where V is the location of the vertex along the surface normal, c is the amplitude of the oscillation, p is the period, f is the floor of the oscillation, and t is time. The sinusoid defines a smooth transition between the floor and amplitude over time for each vertex in the surface mesh. If the floor and amplitude for each vertex corresponds to the uncertainty at that point on the surface, then the viewers eye will naturally be drawn to areas of high uncertainty as the surface animates. Other oscillation functions could be used that cause more rapid transitions between states, such as step and sawtooth functions. Figure 3 shows three frames of the animation.

3.3 Bounding Statistics

The third prototype uses statistics similar to the traditional box-plot [16] to bound the valid regions of the simulation. The minimum, maximum, and mean surfaces are calculated point-wise (i.e., for each density-temperature point), as is the standard deviation across all surfaces. The user controls what is displayed through a graphical interface, which provides options to show each of the original pressure surfaces and contextual surfaces such as the mean +/- standard deviation. Data values can be colormapped onto the mean surface and the user may choose which data values are displayed. Figure 4 shows a screenshot of the prototype. The mean surface is shown colored by the pressure, and is flanked by the minimum and maximum surfaces which are partially transparent to reduce visual clutter. Two of the original surface realizations are shown (pink and purple, mostly obscured) below the mean surface. The main goal of this prototype is to show the range of possible pressures (as a function of density and temperature), as well as indicate where the data is most likely to reside.
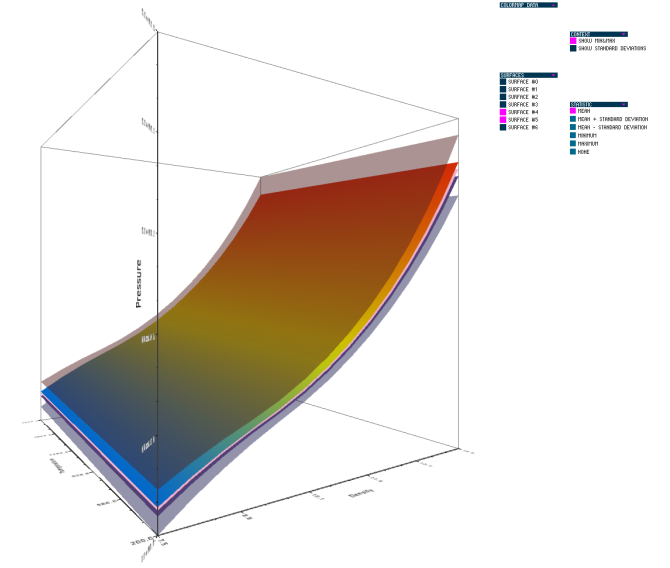

Figure 4: Prototype using bounding statistics and a graphical user interface to explore the dataset. The mean surface is shown along with the minimum and maximum surfaces (with partially transparency) and two surface realizations (in pink and purple). Through a series of buttons, the user can control which surfaces to display.

3.4 View Dependent Opacity

The final prototype uses an approach similar to the Blinn lighting model [2], in that the view angle is compared with the normal of the surface at each individual point. Instead of using this to modulate the lighting, it is used to modulate the opacity of the surface at each point. In this prototype the technique is extended to display a collection of surfaces, each rendered individually. For this prototype the internal energy was used for the test data, rather than the pressure, because the internal energy surfaces have a greater variation in the surface normal than the pressure surfaces. For each density (ρ) and temperature (T) pair, the transparency $\sigma_{\rho,T}$ is computed as the standard deviation of the internal energy across the surface realizations; in regions where the differences between the surfaces are large, the transparency will be large. At each density-temperature point, the transparency is evenly distributed to each surface. That is, the transparency of each of the N surfaces is $N * \sigma_{\rho,T}$, so that when looking through all N surfaces the transparency at ρ, T, is $\sigma_{\rho,T}$. The last step of the algorithm is to apply the dependence of the viewpoint to the transparency. If the vector from the viewpoint to a point on the surface is $v_{\rho,T}$ and the surface normal at that point is $n_{\rho,T}$, then the transparency is multiplied by $v \cdot n$. When the viewpoint is normal to the surface, the transparency of each surface is unchanged, but as the viewpoint becomes more oblique, the transparency is reduced. When the viewer is at an oblique angle, the
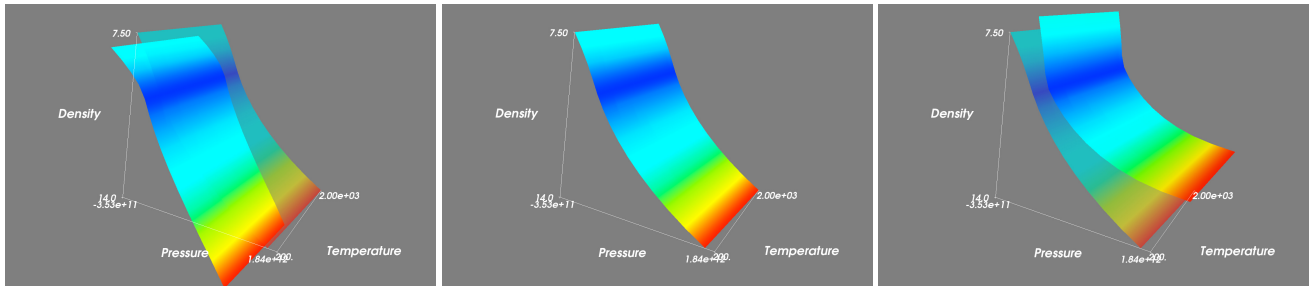
Figure 3: Three frames of the surface animation. The transparent surface represents the fixed, mean pressure surface and is visible in every frame for reference. The opaque surface sweeps through a region defined by one standard deviation distance above and below the mean surface. Here, we show the minimal, mean, and maximal position of the animated surface. Both surfaces are colored by standard deviation, with the largest point of deflection away from the mean surface shown in red.

surface is essentially opaque, regardless of the underlying surface's transparency.

The motivation behind this approach can be understood by considering objects embedded within the collection of surfaces. When the viewpoint is normal to the surfaces, the surfaces are transparent and objects inside (such as point glyphs representing simulation data) are clearly visible. This allows the viewer to determine where the object is positioned in depth because an object farther from the viewer will be obscured by more partially transparent surfaces, while a closer object will be less obscured, see Fig. 5. When the surfaces are viewed at oblique angles, or from the side, the surfaces become opaque. The embedded object's position with respect to the surfaces is apparent because the individual surfaces can be more easily distinguished from this viewpoint, as shown in Fig. 6. Of the different prototypes, this technique is the most complex and the least mature, and significant effort would be required to develop a production capability.
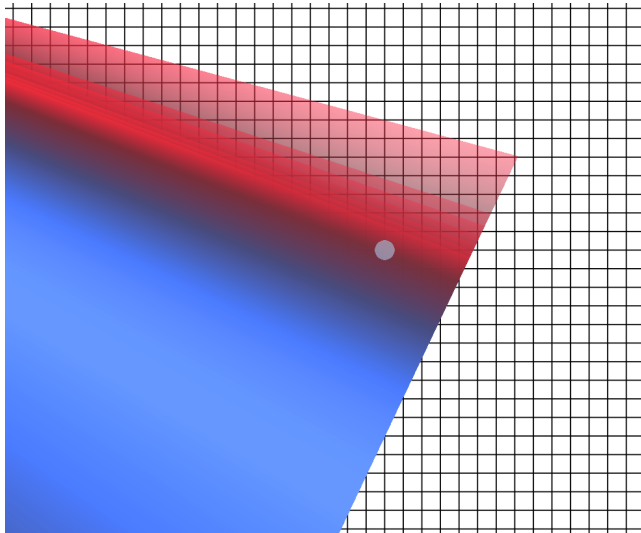

Figure 5: View-dependent opacity prototype, normal view. Pressure surfaces are colored by uncertainty, with blue indicating low uncertainty and red indicating high. Individual surface opacities are also dependent on uncertainty. The background grid is provided so that different levels of opacity can be distinguished. In this image, the view direction is nearly aligned with the surface normals of the red edges of the pressure surface realizations, and the surfaces are at nearly maximum transparency. The spherical object is partially visible behind several of the surfaces.
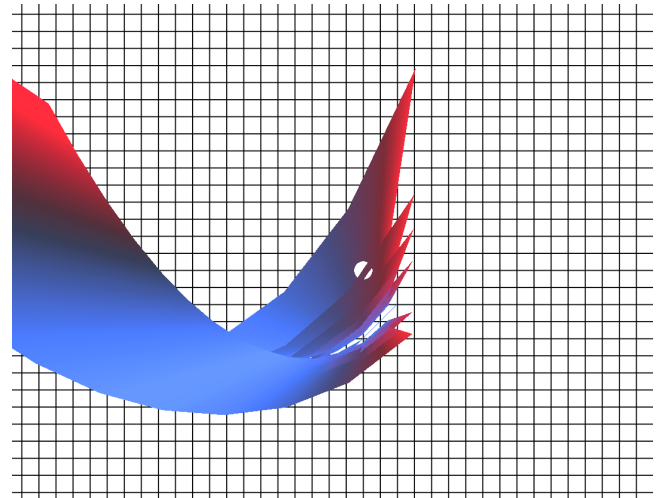

Figure 6: View-dependent opacity prototype, edge view. Pressure surfaces are colored by uncertainty, with blue indicating low uncertainty and red indicating high. Individual surface opacities are also dependent on uncertainty. Here the view direction is nearly orthogonal to the surface normals of the pressure surfaces, and the surfaces are essentially opaque. The location of the spherical object within the collection of surfaces is easy to determine because the surfaces are distinct from this viewpoint.

4 ROLES FOR MATERIAL MODEL UNCERTAINTY VISUAL-IZATION

After reviewing and analyzing the feedback from the focus groups, several themes emerged. In this section we begin with the themes more closely tied to the prototype visualizations, then move to the broader role visualization can play in addressing the needs of the material model stakeholders.

4.1 The different prototypes were suited to different uses

Each visualization prototype has different features. Participants found that the features could be be positive or negative, depending on the use case. The point cloud prototype was effective at showing how the uncertainty varied in different areas of the domain. As intended, the points immediately conveyed a sense that the surface was not known precisely. However, because points represented the uncertainty of the material model, participants thought another mechanism would be needed to show simulation data, which would, most naturally, also be represented by points. A second concern,

particularly for material modelers, was that important correlations were lost through the statistical processing of the EOS data – all the individual surfaces were averaged to compute the mean surface, and the point cloud was generated from the standard deviation with respect to the average surface.

The bounding statistics prototype also applied statistical processing, but maintained the original surfaces. While the emphasis of the prototypes was on visualization techniques, participants vocally supported the ability to display or hide the individual surfaces and statistically generated surfaces. The bounding surface technique worked well for this dataset, but participants were not sure how effective it would be for a larger number of surfaces or for multiphase surfaces, which have more geometric complexity.

The surface animation prototype was developed after the others, in response to participants' desire to see individual surfaces and variation within the set, but without overwhelming the viewer with all the information. The view-dependent opacity prototype was less intuitive to participants because initially, it wasn't clear why the opacity changed with the viewing angle. However, the representative point of simulation data allowed participants to grasp the value of seeing the EOS surface, the associated uncertainty, and simulation data at the same time – as the uncertainty increased the surfaces became more transparent, and one would see simulation data more clearly when it moved into an uncertain region, identifying a cause for concern.

4.2 Surface data is not enough

The EOS surface, whether for pressure, energy or another thermodynamic variable, provides an overview of the material behavior, but the surface in and of itself provides just the context. Participants expressed the need for various types of references to connect the surface to their understanding of the material behavior. Axes and contours labeled with numerical values would provide quantitative references. Domain-specific landmarks, such as phase boundaries, Hugoniots, isentropes and isotherms would provide a link back to traditional two-dimensional plots (see below.) For analysts, simulation results must be displayed on the surface to identify the region of interest; and the EOS uncertainty, while highly desired, is only meaningful to analysts in relation to the simulation data. Even when quantified uncertainty is not available, plotting the experimental data to which the model was calibrated would suggest a region of higher confidence to an analyst. For material modelers, simulation data is still informative, but curves with physical meaning (Hugoniots, isentropes, etc.) are more important. A key insight for visualization specialists is that making three-dimensional representations usable, useful, and adoptable across the community is likely to require careful interaction design.

4.3 Two-dimensional graphics remain the standard

The focus group participants were intrigued by the idea of having three-dimensional, interactive visualizations to enhance their understanding of material model dynamics. However, the participants also indicated that two-dimensional plots of, e.g., curves in the pressure-density plane are the existing lingua franca of material model representation; these plots are regularly shown in textbooks and the academic literature on the subject. The stakeholder community is already familiar with this visual form for material model information, so this view leverages previously developed mental models. Two-dimensional plots are relatively information impoverished, given the complexity and uncertainty that characterizes most material models; but they are easy to generate, read, and share. Three-dimensional, interactive visualizations can carry a great deal more information, and participants valued the "big-picture" view of the thermodynamic space. One participant remarked that he would start with a three-dimensional view to get a qualitative feel for the context, but wanted the ability to zoom in to a particular region

of interest and see traditional two-dimensional plots because they were "more quantitative". For this application, effective visualization may depend more on the timely delivery of a key existing technique than discovering or developing a new technique.

4.4 Visualization to compare and select material models

As mentioned in the Introduction, material models are inputs to the simulation code; that is, analysts specify the material models before the simulations can be run. Analysts noted that selecting one material model from several possible choices was hindered by a lack of information about each model, and often sought out a colleague or (if available) a material modeler for guidance. Minimally, one would like to know the range of validity of each model. In addition, the intended application, any articles or documents on the theory behind the model, and the experimental data used to calibrate the model would guide the user relative to their own application. Analysts usually have some idea of the densities, temperatures, and pressures a material is likely to experience in their application. Comparative visualization of the different surfaces for those conditions, particularly if uncertainty were displayed for each candidate model, would allow analysts to choose based on quantitative information.

4.5 Visualization to analyze results

The most readily apparent role for visualization is to aid the analysis of the simulation results. According to the focus groups, there are three primary use cases. The first is to "debug" a simulation that has crashed. For some engineering codes, a majority of simulation failures can be traced to the material model; sometimes the material model is a poor model, sometimes the EOS is being sampled outside (perhaps far outside) its range of validity, and sometimes the material model catches an unrelated error committed far upstream. In all these cases, visualization of simulation data on the EOS surface can provide insight and expose patterns in failures that are difficult to identify on a case by case basis.

A second use case is to assess whether or not a simulation relies on EOS information from a region of high uncertainty. The prototypes display EOS models and their associated uncertainty; by adding the time-dependent paths of simulation data, such as element or cell values, or passive tracers used as diagnostics, an analyst could judge whether the EOS uncertainty was important for that particular simulation. For example, an analyst could conclude that while regions of high uncertainty exist in the model, none of the material in the simulation experienced conditions in those regions. Alternatively, they might observe that a critical part experienced a temperature near a phase transition, and the phase boundary is a region of high uncertainty; in this case, more simulations might be run to investigate the effect of crossing the phase boundary.

Finally, visualizing simulation data simultaneously with material model data can provide insight into simulation results. If in the previous example, the phase boundary was a region of low uncertainty, the analyst might modify their design so their part came closer to the phase transition to improve performance of the device. Of course, these are hypothetical examples, but participants had tried to answer similar questions about their own simulations and, if answers were obtainable, they often required a lot of time and effort. The material modelers who participated in the groups indicated that more detailed visualizations could support better understanding of the sources and degree of uncertainty in various regions of material behavior.

4.6 Visualization for Communication

Material modelers, code developers, and analysts all recognized that communication about material models was a weak point in their work. Analysts were concerned that they had little guidance in

choosing models for their simulations. Several code developers and material modelers recalled frustration that they could not explain a particular material model issue to an analyst in a way the analyst could understand it; they could not describe a complex structure in their own mental model to someone that did not have the same mental model and domain-specific vocabulary. The visualization of material models can alleviate these issues by providing a common view for the different stakeholders. Visualization can also be used as a training and learning tool, for more experienced practitioners as well as those new to the field. Visualization enables analysts to explain the reasoning and data behind their conclusions to sponsors and decision makers. Finally, material model visualization can be used to help researchers, managers and funding agencies identify gaps in knowledge and prioritize resources to close those gaps.

4.7 Provenance

Our assumption has been that uncertainty is a mathematical object that can be quantified, propagated, plotted, and visualized. But a number of participants identified a clarity about a model's origins and history, or provenance, as highly desirable. The lack of provenance is a subjective uncertainty that cannot be quantified. While quantified uncertainties for a particular model can, in principle, be propagated across the entire distributed workflow, stakeholders would invariably be more comfortable if the provenance of the model were known. Provenance provides subjective confidence not just in the model, but in the quantified uncertainty of the model.

5 CONCLUSION

The uncertainty of a material model can mean different things to material modelers, analysts, and the code developers that incorporate the models into the engineering simulation codes that the analysts use. The focus group approach revealed a distributed information workflow around the development and use of EOS models, and that information bottlenecks in this workflow had organizational as well as technical origins. Visualization prototypes anchored the discussions and better differentiated the perspectives of the different stakeholders. There is a clear demand for a visualization capability for EOS models and their uncertainties. This capability would improve communication across the workflow, as well as provide an analysis tool for material modelers and, if simulation data can be incorporated, for analysts and code developers. The complexity and diversity of material modeling and the abstract nature of uncertainty make the development of an effective tool challenging, but collaborative interaction between stakeholders and software developers will ensure the utility and usability of a such a tool.

REFERENCES

[1] B. M. Adams, W. J. Bohnhoff, K. R. Dalbey, J. P. Eddy, M. S. Eldred, D. M. Gay, K. Haskell, P. D. Hough, and L. Swiler. DAKOTA, a multilevel parallel object-oriented framework for design optimization, parameter estimation, uncertainty quantification, and sensitivity analysis: Version 5.1 user's manual. Technical Report SAND2010-2183, Sandia National Laboratories, 2010.

[2] J. F. Blinn. Models of light reflection for computer synthesized pictures. In *SIGGRAPH '77 Proceedings of the 4th annual conference on Computer graphics and interactive techniques*, New York, NY, 1977.

[3] R. Brown. Animated visual vibrations as an uncertainty visualisation technique. *Proceedings of the 2nd international conference on Computer graphics and interactive techniques in Austalasia and Southeast East Asia GRAPHITE 04*, 1(212):84, 2004.

[4] M. C. Brugnoli, J. Hamard, and E. Rukzio. User expectations for simple mobile ubiquitous computing environments. In *Proceedings of the Second IEEE International Workshop on Mobile Commerce and Services*, pages 2–10, Washington, DC, USA, 2005. IEEE Computer Society.

[5] H. B. Callen. *Thermodynamics*. John Wiley & Sons, 1960.

[6] G. A. Dilts. Consistent thermodynamic derivative estimates for tabular equations of state. *Phys. Rev. E*, 73(6):066704, Jun 2006.

[7] S. Eliezer, A. Ghatak, and H. Hora. *An Introduction to Equations of State: Theory and Applications*. Cambridge University Press, 1986.

[8] G. Grigoryan and P. Rheingans. Point-based probabilistic surfaces to show surface uncertainty. *IEEE Transactions on Visualization and Computer Graphics*, 10(5):546–573, 2004.

[9] T. Hall, S. Beecham, and A. Rainer. Requirements problems in twelve software companies: an empirical analysis. *IEE Proceedings - Software*, 149(5):153–160, 2002.

[10] Z. M. Kasirun and S. S. Salim. Focus group discussion model for requirements elicitation activity. In *Proceedings of the 2008 International Conference on Computer and Electrical Engineering*, pages 101–105, Washington, DC, USA, 2008. IEEE Computer Society.

[11] S. P. Lyon and J. D. Johnson, editors. *SESAME: The Los Alamos National Laboratory Equation-Of-State Database*, number LA-UR-92-3407, 1992.

[12] A. P. Massey and W. A. Wallace. Focus groups as a knowledge elicitation technique: An exploratory study. *IEEE Trans. on Knowl. and Data Eng.*, 3:193–200, June 1991.

[13] R. Menikoff and B. J. Plohr. The Riemann problem for fluid flow of real materials. *Reviews of Modern Physics*, 61(1):75–130, 1989.

[14] J. Nielsen. The use and misuse of focus groups. *IEEE Softw.*, 14:94–95, January 1997.

[15] ParaView. http://paraview.org.

[16] J. W. Tukey. *Exploratory Data Analysis*. Addison-Wesley, 1977.

[17] N. Van House, M. Davis, M. Ames, M. Finn, and V. Viswanathan. The uses of personal networked digital imaging: an empirical study of cameraphone photos and sharing. In *CHI '05 extended abstracts on Human factors in computing systems*, CHI EA '05, pages 1853–1856, New York, NY, USA, 2005. ACM.