

Invited Paper for Session B1

**Verification and Validation for Modeling and Simulation
in Computational Science and Engineering Applications**

Foundations for Verification and Validation in the 21st Century Workshop

October 22-23, 2002

Johns Hopkins University/Applied Physics Laboratory

Laurel, Maryland

Verification, Validation, and Predictive Capability in Computational Engineering and Physics

William L. Oberkampf

Validation and Uncertainty Estimation Department

wloberk@sandia.gov

Timothy G. Trucano

Optimization and Uncertainty Estimation Department

tgtruca@sandia.gov

Sandia National Laboratories

P. O. Box 5800

Albuquerque, New Mexico 87185

Charles Hirsch

Department of Fluid Mechanics

hirsch@stro.vub.ac.be

Vrije Universiteit Brussel

Brussels, Belgium

Summary

Computer simulations of physical processes are being relied on to an increasing degree for design, performance, reliability, and safety of engineered systems. Computational analyses have addressed the operation of systems at design conditions, off-design conditions, and accident scenarios. For example, the safety aspects of products or systems can represent an important, sometimes dominant, element of numerical simulations. The potential legal and liability costs of hardware failures can be staggering to a company, the environment, or the public. This consideration is especially crucial, given that we may be interested in high-consequence systems that cannot ever be physically tested, including the catastrophic failure of a full-scale containment building for a nuclear power plant, explosive damage to a high-rise office building, ballistic missile defense systems, and a nuclear weapon involved in a transportation accident.

Developers of computer codes, analysts who use the codes, and decision makers who rely on the results of the analyses face a critical question: How should confidence in modeling and simulation be critically assessed? Verification and validation (V&V) of computational simulations are the primary methods for building and quantifying this confidence. Briefly, verification is the assessment of the accuracy of the solution to a computational model. Validation is the assessment of the accuracy of a computational simulation by comparison with experimental data. In verification, the relationship of the simulation to the real world is not an issue. In validation, the relationship between computation and the real world, *i.e.*, experimental data, *is* the issue.

This paper presents our viewpoint of the state of the art in V&V in computational physics. (In this paper we refer to all fields of computational engineering and physics, *e.g.*, computational fluid dynamics, computational solid mechanics, structural dynamics, shock wave physics, computational chemistry, etc., as computational physics.) We do not provide a comprehensive review of the multitudinous contributions to V&V, although we do reference a large number of previous works from many fields. We have attempted to bring together many different perspectives on V&V, highlight those perspectives that are effective from a practical engineering viewpoint, suggest future research topics, and discuss key implementation issues that are necessary to improve the effectiveness of V&V. We describe our view of the framework in which predictive capability relies on V&V, as well as other factors that affect predictive capability. Our opinions about the research needs and management issues in V&V are very practical: What methods and techniques need to be developed and what changes in the views of management need to occur to increase the usefulness, reliability, and impact of computational physics for decision making about engineering systems?

We review the state of the art in V&V over a wide range of topics; for example, prioritization of V&V activities using the Phenomena Identification and Ranking Table (PIRT), code verification, software quality assurance (SQA), numerical error estimation, hierarchical experiments for validation, characteristics of validation experiments, the need to perform nondeterministic computational simulations in comparisons with experimental data, and validation metrics. We then provide an extensive discussion of V&V research and implementation issues that we believe must be addressed for V&V to be more effective in improving confidence in computational predictive capability. Some of the research topics addressed are development of improved procedures for the use of the PIRT for prioritizing V&V activities, the method of manufactured solutions for code verification, development and use of hierarchical validation diagrams, and the construction and use of validation metrics incorporating statistical measures.

Some of the implementation topics addressed are the needed management initiatives to better align and team computationalists and experimentalists in conducting validation activities, the perspective of commercial software companies, the key role of analysts and decision makers as code customers, obstacles to the improved effectiveness of V&V, effects of cost and schedule constraints on practical applications in industrial settings, and the role of engineering standards committees in documenting best practices for V&V.

Our top ten recommendations for needed research, development, and implementation issues in V&V are the following:

1. The application domain for the modeling and simulation capability should be well understood and carefully defined, and the accuracy requirements should be known for the applications of interest.
2. A validation tier hierarchy, including the validation pyramid, should be carefully constructed using a systems engineering approach.
3. A PIRT should be constructed to identify and rank the importance of physical processes and interactions of processes for all tiers and faces of the validation pyramid.
4. Code verification activities should be defined and prioritized, schedules set, and needed resources allocated.
5. SQA procedures should be defined, implemented, consistently adhered to, and documented.
6. Using the results of the PIRT, model validation activities, both computational and experimental, should be defined and prioritized, schedules set, and needed resources allocated.
7. Validation metrics should be carefully defined with requirements stated, and clearly connected to modeling and simulation requirements for the application domain.
8. Statistical data for both computational and experimental results should be used for the validation metrics to precisely quantify the accuracy of the model for all validation experiments.
9. Various fields of computational physics should begin to assemble, construct, and document V&V databases.
10. Analysts who use codes and decision makers who use results from codes should begin requiring extensive, well documented V&V activities from code developers.

Table of Contents

1.	Introduction.....	5
1.1	Background.....	5
1.2	Basic Terminology and Methodology.....	6
1.3	Outline of the Paper.....	12
2.	Primary Processes.....	13
2.1	Framework for Predictive Capability.....	13
2.2	Modeling and Simulation Requirements.....	17
2.3	Verification Activities.....	21
2.3.1	Fundamentals of verification.....	21
2.3.2	Numerical algorithm verification.....	24
2.3.3	Software quality assurance.....	27
2.3.4	Numerical error estimation.....	29
2.4	Validation Activities.....	31
2.4.1	Fundamentals of validation.....	31
2.4.2	Construction of validation experiment hierarchy.....	33
2.4.3	Characteristics of validation experiments.....	37
2.4.4	Uncertainty quantification in computations.....	40
2.4.5	Validation metrics.....	42
3.	Major Research Issues.....	46
3.1	Prioritization of Assessment Activities.....	46
3.2	Verification Activities.....	48
3.3	Validation Activities.....	49
4.	Major Implementation Issues.....	52
4.1	Management Issues.....	52
4.2	Practical applications in industrial settings.....	54
4.3	Commercial Software.....	56
4.4	Development of Standards.....	58
5.	Concluding Remarks.....	60
	Acknowledgements.....	63
	Acronyms and Abbreviations.....	63
	References.....	64

1. Introduction

1.1 Background

During the last three or four decades, computer simulations of physical processes have been used in scientific research and in the analysis and design of engineered systems. The systems of interest have been existing or proposed systems that operate, for example, at design conditions, off-design conditions, and failure-mode conditions in accident scenarios. The systems of interest have also been natural systems, for example, computer simulations for environmental impact, as in the analysis of surface-water quality and the risk assessment of underground nuclear-waste repositories. These kinds of predictions are beneficial in the development of public policy, the preparation of safety procedures, and the determination of legal liability. Thus, because of the impact that modeling and simulation predictions can have, the credibility of the computational results is of great concern to engineering designers and managers, public officials, and those who are affected by the decisions that are based on these predictions.

For engineered systems, terminology such as “virtual prototyping” and “virtual testing” is now being used in engineering development to describe numerical simulation for the design, evaluation, and “testing” of new hardware and even entire systems. This new trend of modeling and simulation-based design is primarily driven by increased competition in many markets, *e.g.*, aircraft, automobiles, propulsion systems, and consumer products, where the need to decrease the time and cost of bringing products to market is intense. This new trend is also driven by the high cost and time that are required for testing laboratory or field components as well as complete systems. Furthermore, the safety aspects of the product or system represent an important, sometimes dominant, element of testing or validating numerical simulations. The potential legal and liability costs of hardware failures can be staggering to a company, the environment, or the public. This consideration is especially critical, given that we may be interested in the reliability, robustness, or safety of high-consequence systems that cannot ever be physically tested. Examples are the catastrophic failure of a full-scale containment building for a nuclear power plant, a fire spreading through (or explosive damage to) a high-rise office building, ballistic missile defense systems, and a nuclear weapon involved in a transportation accident. In contrast, however, an inaccurate or misleading numerical simulation for a scientific research project has comparatively no impact.

Users and developers of computational simulations today face a critical question: How should confidence in modeling and simulation be critically assessed? Verification and validation (V&V) of computational simulations are the primary methods for building and quantifying this confidence. Briefly, verification is the assessment of the accuracy of the solution to a computational model by comparison with known solutions. Validation is the assessment of the accuracy of a computational simulation by comparison with experimental data. In verification, the relationship of the simulation to the real world is not an issue. In validation, the relationship between computation and the real world, *i.e.*, experimental data, *is* the issue.

In the United States, the Defense Modeling and Simulation Office (DMSO) of the Department of Defense (DoD) has been the leader in the development of fundamental concepts and terminology for V&V.[48,50] Recently, the Accelerated Strategic Computing Initiative (ASCI) of the Department of Energy (DOE) has also taken a strong interest in V&V. The ASCI program is focused on computational physics and computational mechanics, whereas the DMSO has traditionally emphasized high-level systems engineering, such as ballistic missile defense systems,

warfare modeling, and simulation-based system acquisition. Of the work conducted by DMSO, Cohen recently observed:[35] “Given the critical importance of model validation . . . , it is surprising that the constituent parts are not provided in the (DoD) directive concerning . . . validation. A statistical perspective is almost entirely missing in these directives.” We believe this observation properly reflects the state of the art in V&V, not just the directives of DMSO. That is, the state of the art has not developed to the place where one can clearly point out all of the actual methods, procedures, and process steps that *must* be undertaken for V&V.

It is fair to say that computationalists (code users and code developers) and experimentalists in the field of fluid dynamics have been pioneers in the development of methodology and procedures in validation. However, it is also fair to say that the field of computational fluid dynamics (CFD) has, in general, proceeded along a path that is largely independent of validation. There are diverse reasons why the CFD community has not perceived a strong need for code V&V, especially validation. One reason is that a competitive and frequently adversarial relationship (at least in the United States) has often existed between computationalists and experimentalists, resulting in a lack of cooperation between the two groups. We, on the other hand, view computational simulation and experimental investigations as complementary and synergistic. To those who might say, “Isn’t that obvious?” We would answer, “It should be, but they have not always been viewed as complementary.” In retrospect, the relationship between computationalists and experimentalists is probably understandable because it represents the classic case of a new technology (computational simulation) that is rapidly growing and attracting a great deal of visibility and funding support that had been the domain of the older technology (experimentation).

It is our view that the field of structural dynamics has enjoyed, in general, a more beneficial and synergistic relationship between computationalists and experimentalists. We believe this type of relationship has developed because of the strong dependence of structural dynamics models on experimental measurements. Most researchers in the field of structural dynamics have referred to this interaction as “model validation.” As discussed in Section 2.1, we believe a more precise term for this interaction is either “model updating” or “model calibration.” That is, the primary interaction between computation and experiment is to update or “tune” the unknown parameters in the computational model using the experimental results from modal testing. This approach in structural dynamics has proven to be very effective because it permits the estimation of specific constituents of poorly known physics in the computational models. In structural dynamics the problem primarily arises in poor understanding of the localized deformation of connectors and joints between structural elements in the computational models. A similar approach is used in fluid dynamics when dealing with turbulent reacting flows and two-phase flows.

From a historical perspective, the operations research (OR) and systems engineering communities have provided the philosophical foundations for verification and validation. With the recent interest in V&V from the CFD and computational physics communities, one recognizes significant differences in perspectives between the historical view and the view held by the computational physics community. (For simplicity, we will refer to all fields of computational engineering and physics, *e.g.*, CFD, computational solid mechanics, structural dynamics, shock wave physics, computational chemistry, etc., as *computational physics*.)

1.2 Basic Terminology and Methodology

There is a wide variety of different meanings used for V&V in the various technical disciplines. For example, the meanings used by the Institute of Electrical and Electronics Engineers (IEEE) and the software quality assurance community are different than the meanings used in the

DoD modeling and simulation community. And given that members of the different technical communities often work together on V&V activities, we expect there will be long-term ambiguity and confusion resulting from terminology differences. Although we have not reviewed all of the different meanings in this paper, we refer the reader to references that describe the varying usage.[3,7,49,50,96-98] For a review of the historical development of the terminology for verification, validation, and prediction, see, for example, Refs. [134,142,144].

The DMSO under the DoD has played a major role in attempting to standardize the definitions of V&V. In 1994 the DoD published definitions of V&V that are clear, concise, and directly useful by themselves.[48-50] From the perspective of the computational engineering and physics communities, however, the definition of verification by the DoD does not make it clear that the accuracy of the numerical solution to the partial differential equations (PDEs) should be included in the definition. To clarify this issue, the CFD Committee on Standards of the American Institute of Aeronautics and Astronautics (AIAA) proposed a slight modification to the DoD definition. This paper will use the DoD definitions, with the AIAA modification for verification:[3]

Verification: The process of determining that a model implementation accurately represents the developer's conceptual description of the model and the solution to the model.

Validation: The process of determining the degree to which a model is an accurate representation of the real world from the perspective of the intended uses of the model.

We think that the best graphical representation of these definitions was constructed by the Society for Computer Simulation (SCS) in 1979, as shown in Fig. 1. The figure identifies two types of models: a conceptual model and a computerized model. The conceptual model is composed of all mathematical modeling data and mathematical equations that describe the physical system or process of interest. The conceptual model is produced by analyzing and observing the physical system. In computational physics, the conceptual model is dominated by the PDEs for conservation of mass, momentum, and energy. The conceptual model also includes all of the auxiliary equations, such as turbulence models, constitutive models for materials, and electromagnetic cross-section models, and all of the initial conditions and boundary conditions of the PDEs. The computerized model is an operational computer program that implements a conceptual model. Modern terminology refers to the computerized model as the computer model or code. Figure 1 clearly shows that verification deals with the relationship between the conceptual model and the computerized model and that validation clearly deals with the relationship between the computerized model and experimental measurements. The SCS defined *qualification* as the “Determination of adequacy of the conceptual model to provide an acceptable level of agreement for the domain of intended application.” According to this definition, qualification would deal with issues such as, definition of the system of interest, effects of the environment on the system, and the choice of PDEs in computational physics. The topic of conceptual model qualification will not be addressed in this paper.

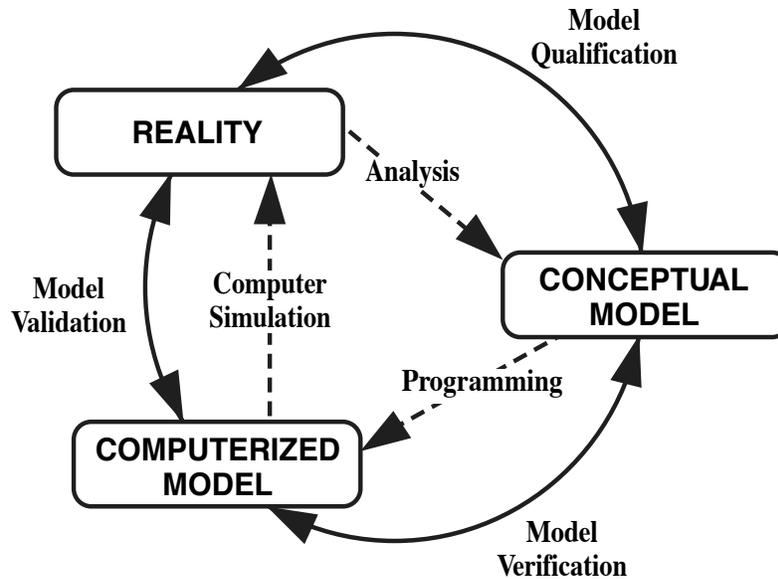


Figure 1
Phases of Modeling and Simulation and the Role of V&V[172]

Fundamentally, V&V are tools for assessing the accuracy of the conceptual and computerized models. For much of the OR work, the assessment was so difficult, if not impossible, that V&V became more associated with the issue of credibility, *i.e.*, the quality, capability, or power to elicit belief. In science and engineering, however, quantitative assessment of accuracy, at least for some important physical cases, is mandatory; it is a necessary condition for credibility. And in certain situations, assessment can only be conducted using subscale physical models or a subset of the active physical processes. Regardless of the difficulties and constraints, methods must be devised for measuring the accuracy of the model for as many conditions as the model is deemed appropriate. As the complexity of a model increases, its accuracy and range of applicability become more difficult to assess.

Some important implications and subtleties in the definitions of V&V need to be addressed at this point. The first significant feature is that both V&V are “process[es] of determining.” That is, they are ongoing activities that do not have a clearly defined completion point. Completion or sufficiency is usually determined by practical issues such as budgetary constraints and intended uses of the model. The definitions include the ongoing nature of the process because of an unavoidable but distressing fact: the veracity, correctness, and accuracy of a computerized, or computational, model cannot be demonstrated for all possible conditions and applications, except for trivial models. Trivial models are clearly not of interest. All-encompassing proofs of correctness, such as those developed in mathematical analysis and logic, do not exist in complex modeling and simulation. Indeed, one cannot prove that complex computer codes have no errors. Likewise, models of physics cannot be proven correct; they can only be disproved. Thus, V&V activities can only assess the correctness or accuracy of the specific cases tested.

The emphasis on “accuracy” is the second feature that is common in the definitions of V&V. This feature assumes that a measure of correctness can be determined. In verification activities, accuracy is generally measured in relation to benchmark solutions of simplified, or closely related, model problems. Benchmark solutions refer either to analytical solutions, *i.e.*, exact solutions to

the PDEs with the specified initial conditions and boundary conditions, or to highly accurate numerical solutions. However, we believe that in the solution of nonlinear PDEs or solutions with discontinuities or singularities the most reliable benchmark solutions are analytical solutions. In validation activities, accuracy is measured in relation to experimental data, *i.e.*, our best indication of reality. Since all experimental data have random (statistical) and bias (systematic) errors, the issue of “correctness,” in an absolute sense, becomes impossible. From an engineering perspective, however, we do not require “absolute truth”: we only expect a statistically meaningful comparison of computational results and experimental measurements. These issues are discussed in more detail in Section 2.4.

Effectively, verification provides evidence (substantiation) that the conceptual (continuum mathematics) model is solved correctly by the discrete-mathematics computer code. (Note: When we refer to “continuum mathematics” we are *not* referring to the physics being modeled by the mathematics. For example, the equations for noncontinuum fluid dynamics are commonly expressed with continuum mathematics.) Verification does not address whether the conceptual model has any relationship to the real world. Validation, on the other hand, provides evidence (substantiation) for how accurately the computational model simulates reality. This perspective implies that the model is solved accurately. However, multiple errors or inaccuracies can cancel one another and give the appearance of a validated solution. Verification, thus, is the first step of the validation process and, while not simple, is much less involved than the more complex statistical nature of validation. Validation addresses the question of the fidelity of the model to specific conditions of the real world. As Roache[161] succinctly states, “Verification deals with mathematics; validation deals with physics.”

As a final comment on terminology, it is our view that the DoD definition of validation does *not* include the concept of *adequacy of the computational result for the intended uses* of the model. Stated differently, we argue that validation is the process of determining the degree to which a computational simulation results agree with experimental data. We recognize that this interpretation of the meaning of validation is narrower than the interpretation that is widely accepted in the DoD community. It is our understanding that the DoD community considers validation to be the process of determining the degree to which the computational model results are *adequate* for the application of interest. This important topic of divergent interpretations of validation is briefly discussed in recommendation for future work, Section 4.4. Regardless of whether or not the reader agrees with our interpretation, we have chosen to clarify our view now to help avoid confusion throughout the paper. Stating our view succinctly: validation deals with comparisons between experimental data and computational data; *not* the adequacy of the comparisons.

In 1998, the Computational Fluid Dynamics Committee on Standards of the AIAA contributed to the basic methodology and procedures for V&V.[3] The *Guide for the Verification and Validation of Computational Fluid Dynamics Simulations*, referred to herein as the “AIAA Guide” was the first engineering standards document that addressed issues of particular concern to the computational physics community. In the following paragraphs we have briefly summarized the basic methodology for V&V from the AIAA Guide.

The fundamental strategy of verification is to identify, quantify, and reduce errors in the computational model and its numerical solution. Figure 2 depicts the verification process of comparing the numerical solution from the code in question with various types of highly accurate solutions.[3] Given a numerical procedure that is stable, consistent, and robust, the five primary sources of errors in computational physics solutions are (1) insufficient spatial discretization, (2) insufficient temporal discretization, (3) insufficient iterative convergence, (4) computer round-off, and (5) computer programming. The emphasis in verification is on identifying and quantifying

errors from these various sources, as well as on demonstrating the stability, consistency, and robustness of the numerical scheme. Stated differently, an analytical or formal error analysis is inadequate in the verification process; verification relies on *demonstration* and *quantification* of numerical accuracy. Note that the recommended methodology presented here applies to finite-difference, finite-volume, finite-element, and boundary-element discretization procedures.

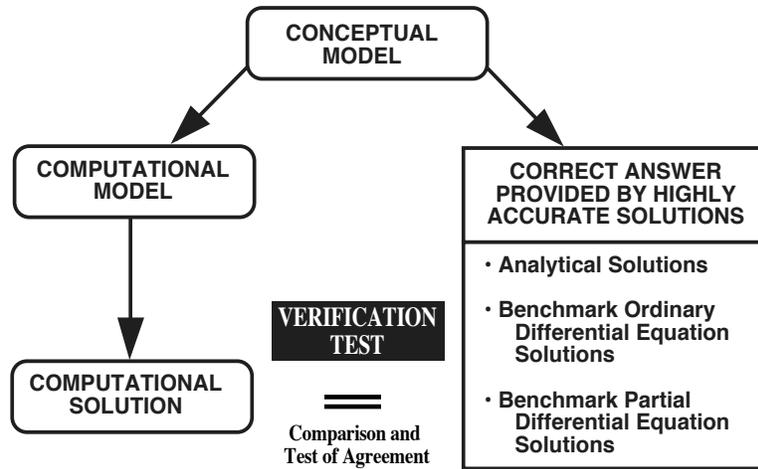


Figure 2
Verification Process[3]

The first three error sources listed in the previous paragraph (spatial discretization, temporal discretization, and iterative procedure) are considered to be within the traditional realm of computational physics, and there is extensive literature dealing with each of these topics. The fourth error source, computer round-off, is rarely dealt with in computational physics. Collectively, these four topics in verification could be referred to as solution verification or solution error assessment. The fifth error source, computer programming, is generally considered to be in the realm of computer science or software engineering. Programming errors, which can occur, for example, in input data files, source-code programming of the numerical algorithm, output data files, compilers, and operating systems, generally are addressed using methods and tools in software quality assurance (SQA), also referred to as software quality engineering (SQE).[110] The identification of programming errors is usually referred to as code verification, as opposed to solution verification.[161,162] In our opinion, the perspectives and emphases of SQA and computational physics are sufficiently different and thus should be addressed separately: code verification, which emphasizes the accuracy of the numerical algorithm that is used to solve the discrete form of the PDEs, and SQA, which emphasizes the issues in computer science. Each of these topics is discussed in Section 2.3.

The fundamental strategy of validation involves identifying and quantifying the error and uncertainty in the conceptual and computational models, quantifying the numerical error in the computational solution, estimating the experimental uncertainty, and then comparing the computational results with the experimental data. This strategy *does not* assume that the experimental measurements are more accurate than the computational results. The strategy only asserts that experimental measurements are the most faithful reflections of reality for the purposes of validation. Validation requires that the estimation process for error and uncertainty must occur

on both sides of the coin: mathematical physics and experiment. Figure 3 depicts the validation process of comparing the computational results of the modeling and simulation process with experimental data from various sources.[3]

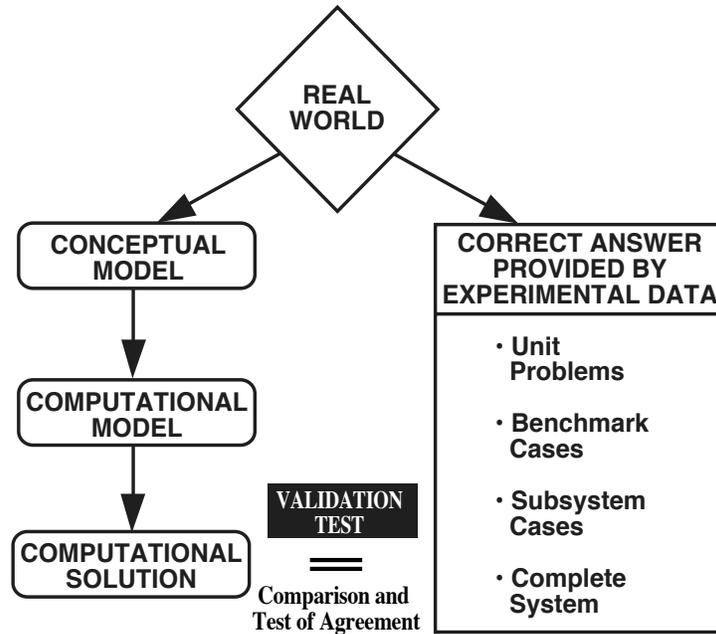


Figure 3
Validation Process[3]

Because of the infeasibility and impracticality of conducting true validation experiments on most complex or large scale systems, the recommended method is to use a building-block approach.[3,38,123,126,175,176] This approach divides the complex engineering system of interest into at least three progressively simpler tiers: subsystem cases, benchmark cases, and unit problems. The strategy in the tiered approach is to assess how accurately the computational results compare with the experimental data (with quantified uncertainty estimates) at multiple degrees of physics coupling and geometric complexity. The approach is clearly constructive in that it (1) recognizes that there is a hierarchy of complexity in systems and simulations and (2) recognizes that the quantity and accuracy of information that is obtained from experiments vary radically over the range of tiers. Furthermore, this approach demonstrates that validation experiments can be conducted at many different levels of physics and system complexity. Each comparison of computational results with experimental data allows an inference of validation concerning tiers both above and below the tier where the comparison is made. However, the quality of the inference depends greatly on the complexity of the tiers above and below the comparison tier. For simple physics, the inference may be very strong, *e.g.*, laminar, single phase, Newtonian, nonreacting flow, and rigid-body structural dynamics. However, for complex physics, the inference is commonly very weak, *e.g.*, turbulent reacting flow and fracture dynamics. This directly reflects the quality of our scientific knowledge about the experiments and calculations that are being compared for more complex tiers.

1.3 Outline of the Paper

This paper presents our viewpoint on the state of the art in V&V in computational physics. We have not attempted herein to provide a comprehensive review of the multitudinous contributions to V&V from many diverse fields. This literature represents distinctively different perspectives and approaches, ranging from engineering and physics to operations research. Recent reviews of the literature are given in Refs. [14,107,108,130,143,161,164,195]. Recent work providing wide-ranging procedures in V&V is described in Refs. [27,143,144,155,185]. We have attempted in this paper to bring together many different perspectives on V&V, highlight those perspectives that are effective from a practical engineering viewpoint, suggest future research topics, and discuss key implementation issues that are necessary to improve the effectiveness of V&V. Our views about the research needs and management issues in V&V are very practical: What methods and techniques need to be developed and what changes in the views of management need to occur to increase the usefulness, reliability, and impact of computational physics for decision making about engineering systems?

Section 2 describes the primary processes in V&V and the relationship of V&V to predictive capability. V&V are the key building blocks in assessing the confidence in the predictive capability of a computational physics code. The section begins with a description of the framework in which predictive capability relies on V&V, as well as other factors that affect predictive capability. We also briefly discuss how this framework is related to more traditional approaches of predictive capability, *e.g.*, Bayesian estimation. The importance of requirements for a computational-physics capability is stressed so that effective and efficient V&V activities can be conducted. Following the framework discussion, we present a summary of verification activities, emphasizing code verification, SQA, and numerical error estimation. A discussion of validation activities follows, highlighting methods for focusing on the validation experiments most important to the predictive capability required. The section concludes with a discussion of hierarchical experiments for validation, characteristics of validation experiments, the need to perform nondeterministic simulations in comparisons with experimental data, and validation metrics.

Section 3 discusses the research issues that we believe must be addressed for V&V to be more effective in improving confidence in computational predictive capability. We begin with a discussion of methods for prioritizing V&V assessment activities. Needed research in verification activities, such as statistical methods for code verification and the method of manufactured solutions, are discussed. Topics discussed with regard to validation research are: development and use of hierarchical validation diagrams; and the construction and use of validation metrics incorporating statistical measures. We close the section with the difficult research issue of how to quantify uncertainty in predictive capability, given an arbitrary set of validation experiments, *e.g.*, how to estimate the uncertainty of a computational prediction when, in some sense, we are extrapolating beyond the validation database.

Section 4 discusses issues related to improving the implementation of V&V activities in a realistic engineering environment and a commercial software environment. These issues include needed improvements in management of verification activities, the key role of analysts and decision makers as code customers, and obstacles to the improved effectiveness of V&V. Examples of obstacles are conflicting perspectives of code developers, analysts, hardware designers, and experimentalists, competition between organizations and nations, and the loss of focus on the needs of the customer for computational physics analyses. Also discussed are the effects of cost and schedule constraints on practical applications in industrial settings, the demands of customers on commercial software companies, and the large-scale validation database activity underway in

Europe.

Section 5 presents some closing remarks concerning the status of V&V and the need for improvement.

2. Primary Processes

2.1 Framework for Predictive Capability

The issues underlying the V&V of mathematical and computational models of physical systems, including those systems with strong human interaction, touch on the very foundations of mathematics, science, and human behavior. Verification is rooted in issues pertaining to continuum and discrete mathematics and to the accuracy and correctness of complex logical structures (computer codes). Validation is deeply rooted in the question of how formal constructs (mathematical models) of nature and human behavior can be tested by physical observation. In the OR field, the systems being analyzed can be extraordinarily complex, such as industrial planning models, marketing models, national and world economic models, monetary investment models, and military conflict models. For these types of situations, one must deal with statistical models where statistical calibration and parameter estimation are crucial elements in building the models. These complex models commonly involve a strong coupling of complex physical processes, human behavior, and computer-controlled systems. For such complex systems and processes, fundamental conceptual issues immediately arise about how to assess the accuracy of the model and the resulting simulations. Indeed, the predictive accuracy of most of these models cannot be assessed in any meaningful way, except for predictive cases that are very near, in some sense, the calibration database.

In the AIAA Guide, prediction is defined as “use of a computational model to foretell the state of a physical system under conditions for which the computational model has not been validated.” A prediction refers to the computational simulation of a specific case of interest that is *different* from cases that have been validated. This definition of prediction is more restrictive than the general scientific meaning of prediction because it eliminates past comparisons of computational results with experimental data. This definition segregates the general meaning of prediction and only refers to prediction, not retrodiction (replication of previously obtained results). If this restriction is not made, then one is only demonstrating previous agreement with experimental data in the validation database. The results of the process of validation should be viewed as historical statements. Thus, the validation database represents reproducible evidence that a model has achieved a given level of accuracy in the solution of specified problems. From this perspective, it becomes clear that validation comparisons do not directly allow one to make claims about the accuracy of predictions; they allow inferences to be made. The strength of the inferences depends on many factors. The suggested relationship between validation and prediction is shown in Fig. 4.

Figure 4 attempts to capture the distinction between validation and prediction. The bottom portion of the figure represents the validation process. Although it is not readily apparent, the validation process in Fig. 4 is fundamentally the same as that shown in Fig. 3. In Fig. 4, the block “Validation Experiments” produces one or more realizations of the “Real World.” The “Experimental Outcomes” are the physical realizations, *i.e.*, the experimental data from the experiment. The physical conditions from the actual validation experiments, *i.e.*, model input parameters, initial conditions, and boundary conditions, are input to the “Computational Model,” which produces the “Computational Results of Experimental Outcomes.” These results are then

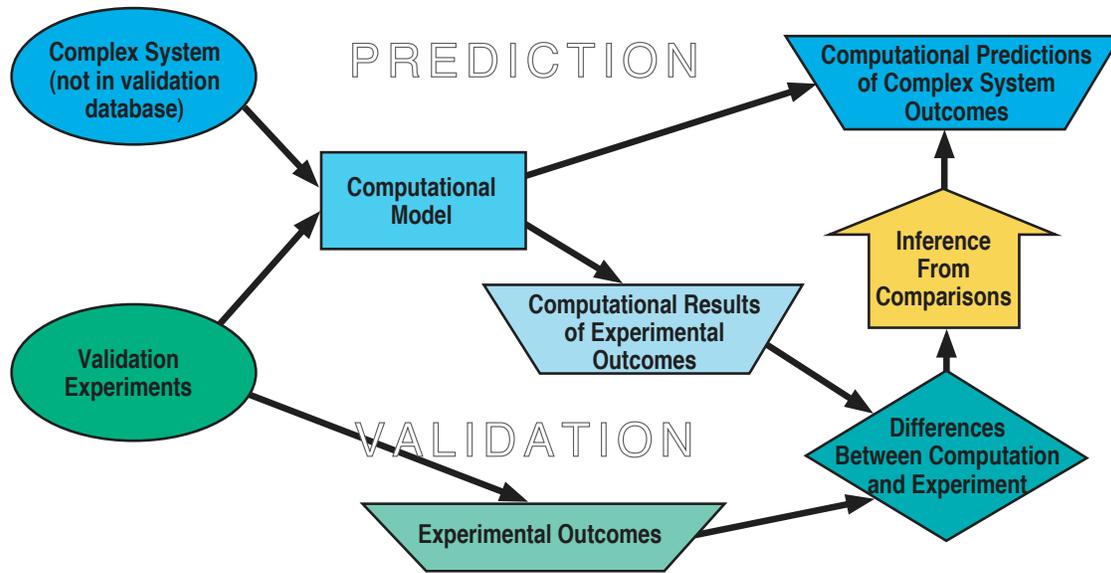


Figure 4
Relationship of Validation to Prediction[57,143,144]

compared with the experimentally determined outcomes in the block “Differences Between Computation and Experiment.” Based on the magnitude of these differences in quantities of interest in the simulation and on the understanding of the physical process, an “Inference from Comparisons” is made.

The upper portion of Fig. 4 represents the prediction process. The “Complex System” of interest should drive the entire modeling and simulation process, but most of the realizations of interest, *i.e.*, predictions, are not in the validation database. That is, when a physical realization is conducted as part of the validation database, regardless of the tier as discussed in Section 2.4.2, the realization becomes part of the “Validation Experiments.” Predictions for conditions of interest are made using the “Computational Model,” resulting in “Computational Predictions of Complex System Outcomes.” The confidence in these predictions is determined by the “Inference from Comparisons” and the level of understanding of the physical process.

The process of logical inference of accuracy of a computational model stemming from its associated validation database is analogous to similar processes and conclusions for classical scientific theories. However, we argue that the strength or confidence in the inference from computational simulation is, and should be, much weaker than traditional scientific theories. Computational simulation relies on the same logic as traditional science, but it also relies on many additional mathematical issues, *e.g.*, discretization algorithms and grid quality, and practical issues, *e.g.*, computer hardware, operating-system software, source-code reliability, and analyst skill, that are not present in classical science. (In this paper “analyst” refers to the individual, or group of individuals, who constructs the computational model, including all physical and numerical inputs to the model, and then uses the computer code to produce a simulation.) One of the key theoretical issues is the state of knowledge of the process being modeled. Zeigler et al[195] gives a detailed discussion of hierarchical levels of knowledge of a system. For physical processes that are well understood both physically and mathematically, the inference can be quite strong. For complex physical processes, the inference can be quite weak. A general mathematical method for

determining how the value and consequence of such an inference degrades as the physical process becomes more complex has not been formulated. For example, in a complex physical process how do you determine “how nearby” the prediction case is from cases in the validation database? Struggling with the strength or quantification of the inference in a prediction is presently an important topic of research.[29,195]

To properly deal with the issue of predictive capability requires a more careful distinction between error and uncertainty than has traditionally been made in the literature. The AIAA Guide[3] (see also Refs. [139-141]) defines error to be “A recognizable deficiency in any phase or activity of modeling and simulations that is not due to lack of knowledge.” This definition emphasizes the required feature that the deficiency is identifiable or knowable upon examination; that is, the deficiency is not caused by lack of knowledge. This definition leads to the so-called *acknowledged* and *unacknowledged* errors. An acknowledged error is characterized by knowledge of divergence from an approach or ideal condition that is considered to be a baseline for accuracy. Examples of acknowledged errors are finite precision arithmetic in a computer, approximations made to simplify the modeling of a physical process, and conversion of PDEs into discrete equations. An acknowledged error can therefore be measured in principle because its origins are fully identified. Solution verification, *i.e.*, solution error estimation, deals with acknowledged errors. For example, we know that the discretization of the PDE and the approximate solution of the discrete equations introduce acknowledged errors. Unacknowledged errors are blunders or mistakes such as programming errors, input data errors, and compiler errors. There are no straightforward methods for estimating, bounding, or ordering the contributions of unacknowledged errors. Code verification and SQA activities primarily deal with unacknowledged errors: another important reason these activities are essential.

In a technical sense, the term uncertainty seems to have two rather different meanings. The first meaning of uncertainty has its roots in probability and statistics: the estimated amount or percentage by which an observed or calculated value may differ from the true value. This meaning of uncertainty has proven its usefulness over many decades, particularly in the estimation of random uncertainty in experimental measurements.[36] The second meaning of uncertainty relates to lack of knowledge about physical systems, particularly in the prediction of future events and the estimation of system reliability. The probabilistic-risk and safety assessment communities, as well as the reliability engineering community, use the term *epistemic uncertainty* in this latter sense. The risk assessment community[40,63,65,84,94,151] refers to the former meaning, random uncertainty, as *aleatory uncertainty*, as does the information theory community.[6,54,109,111,112,179]

Aleatory uncertainty is used to describe the inherent variation associated with the physical system or environment being considered. Sources of aleatory uncertainty can commonly be singled out from other contributors to uncertainty by their representation as randomly distributed quantities that can take on values in an established or known range, but for which the exact value will vary by chance from unit to unit or from time to time. The mathematical representation most commonly used for aleatory uncertainty is a probability distribution. Aleatory uncertainty is also referred to in the literature as variability, irreducible uncertainty, inherent uncertainty, and stochastic uncertainty.

Epistemic uncertainty as a cause of nondeterministic behavior derives from some level of ignorance or lack of knowledge about the system or the environment. Thus an increase in knowledge or information can lead to a reduction in the predicted uncertainty of the system’s response—all things being equal. Epistemic uncertainty can be introduced from a variety of sources, such as limited or nonexistent experimental data for a fixed (but unknown) physical parameter, limited understanding of complex physical processes, and insufficient knowledge

concerning initial conditions and boundary condition in an experiment. Epistemic uncertainty is also referred to in the literature as reducible uncertainty, subjective uncertainty, and model form uncertainty.

To better understand the present perspective of the relationship of prediction to validation, consider Fig. 5. The “validation domain” in the figure suggests three features. First, in this region we have high confidence that the relevant physics is understood and modeled at a level that is commensurate with the needs of the application. Second, this confidence has been quantitatively demonstrated by satisfactory agreement between computations and experiments in the validation database for some range of applicable parameters in the model. And third, the boundary of the domain indicates that outside this region there is a degradation in confidence in the quantitative predictive capability of the model. Stated differently, outside the validation domain the model is credible, but its quantitative capability has not been demonstrated. The “application domain” indicates the region where predictive capability is needed from the model for the applications of interest.

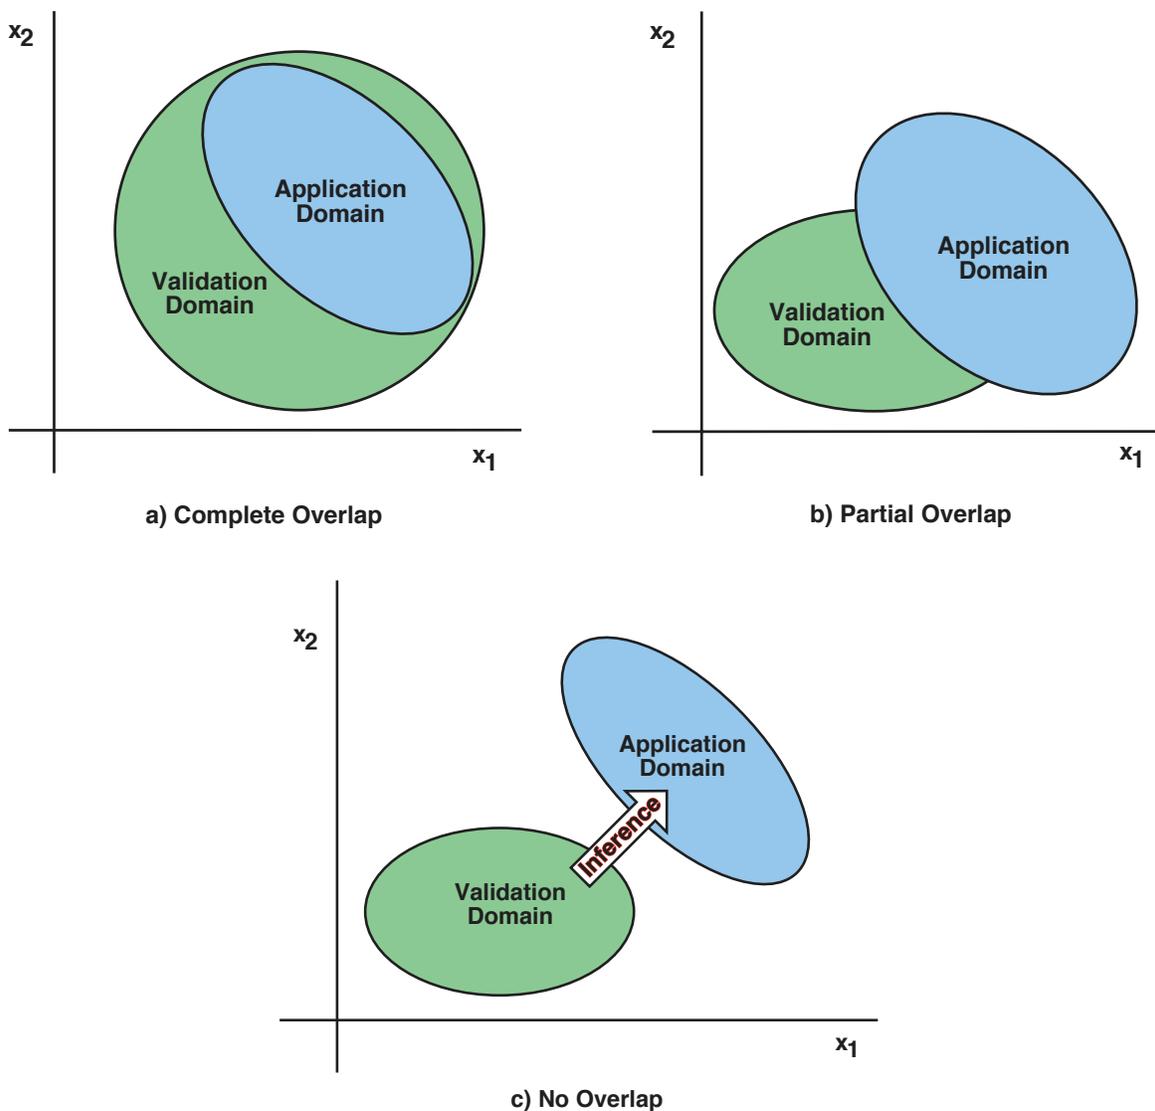


Figure 5
Possible Relationships of the Validation Domain to the Application Domain

Figure 5a depicts the prevalent situation in engineering which shows the complete overlap of the validation domain with the application domain. The vast majority of modern engineering system design is represented in Fig. 5a. Figure 5b represents the occasional engineering situation where there is significant overlap between the validation domain and the application domain. Some examples are: prediction of crash response of new automobile structures, entry of spacecraft probes into the atmosphere of another planet, and the structural response of new designs for deep-water offshore oil platforms. Figure 5c depicts the situation where there is no overlap between the validation domain and the application domain.[57] We believe that predictions for many high-consequence systems are in this realm because we are not able to perform experiments for closely related conditions. We believe that the inference from the validation domain can only be made using *both* physics-based models and statistical methods. The need to perform this extrapolation reinforces our need for models to be judged to achieve the right answer for the right reasons in the validation regime. Model calibration, which employs explicit tuning or updating of model parameters to achieve some degree of agreement with existing validation experiments, does not fully assess uncertainty in the predictive use of the model.

The requirement for predictive code use “far from” the validation database necessitates extrapolation beyond the understanding gained strictly from experimental validation data. This type of uncertainty in our statistical inference is primarily epistemic uncertainty. For example, the modeling of specific types of interactions of physical processes may not have been validated together in the given validation database. There is then uncertainty in the accuracy of model predictions describing such interactions. Traditional applications of computational physics have typically not had to deal with consequential predictive extrapolation. However, some exceptions to this statement are computational modeling performed as part of the safety assessment of nuclear-reactors during failure scenarios[102,113,127,132] and the assessment of nuclear-waste-repository performance.[86,87,128,152]

2.2 Modeling and Simulation Requirements

To improve the efficiency and confidence-building impact of V&V, we believe it is necessary to improve the coupling of V&V activities with the requirements of the intended application of the computational physics code. Figure 6 depicts a recent high-level view of the role of experimental validation in evaluating the predictive capability of a computational physics code.[185] Validation and the associated constraints on validation experiments cannot be separated from the requirements for application of the code, nor can validation be separated from the ultimate predictions that are demanded of the code. In Fig. 6 we suggest that the validation-prediction process relies upon the following activities:

1. Identification and specification of the application driver that focuses the use of the code under discussion
2. Careful planning of V&V activities, especially the use of the Phenomena Identification and Ranking Table (PIRT) to prioritize V&V activities for the application driver
3. Development, implementation, and documentation of code verification and SQA activities for the code, as well as solution error assessment of validation calculations
4. Design and execution of validation experiment activities in accordance with the PIRT

5. Development and definition of appropriate metrics for comparing computational results with experimental results to measure confidence in the intended application of the code
6. Assessing the success or failure of the code using the results of the validation metrics
7. Assessment of the required predictive accuracy of the code using the validation metrics and the application driver of the code
8. Accurate and full documentation of the planning, results, and consequences of the validation activities, especially their implications for predictive confidence for the application driver of the code

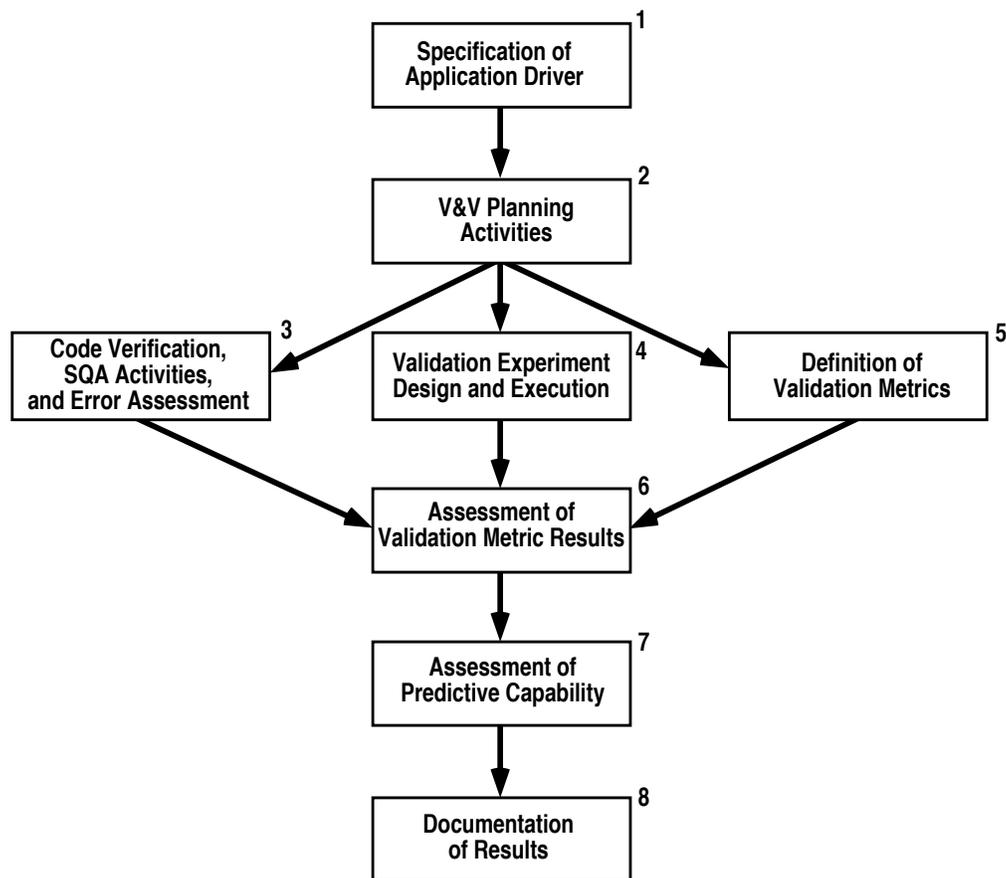


Figure 6
V&V Activities and Predictive Capability
 (Adapted from Ref. [185])

Although Ref. [185] discusses in detail each of these eight activities, we only comment on the two key activities that directly address the focus of V&V on the requirements of the application: block 1 “Intended Application” and block 2 “Planning.” The intended application is the application for which the modeling and simulation capability is being developed. Concerning application requirements, validation activities must assess confidence in the use of the code for a specified

application. The application requirement at which a particular validation activity is directed is a critical planning element and must be defined before the performance of any specific validation work. One of the methods in risk assessment for identifying specific applications deals with identifying system-application scenarios. In the nuclear weapons area, event scenarios are segregated into three wide categories to which the engineering system might be exposed: normal, abnormal, and hostile environments. Although these categories may not be appropriate for all engineering systems, we believe they are a helpful framework for many types of systems, *e.g.*, systems ranging from military and commercial aircraft to spacecraft to power generation facilities and public buildings.

Regarding planning, we believe the PIRT is the most important tool in the planning process for translating technical requirements for the relevant code application into prioritized requirements for V&V activities.[155] The PIRT was developed to assess the safety of nuclear reactors and has primarily been used for that purpose.[22,191-193,197] The PIRT is critical for planning validation experiments because it helps establish both sufficiency and efficiency of the validation activities. To demonstrate sufficiency requires a careful response to the question, What has to be done to establish a necessary level of confidence in the application of the code? To demonstrate efficiency requires evidence that limited resources (people, money, time) are balanced as a result of planning, not simply as a reaction to circumstances. We presume in this paper that dedicated validation experiments supporting the validation assessment of a particular code application are directed at the most important elements in the associated PIRT. If this is not true, there is already a revealed weakness in the planned validation activities. The planning for the dedicated validation experiments should make this direction explicit and clear.

The PIRT addresses five aspects of V&V that relate to the application requirements. Depicted in Fig. 7, these aspects are as follows:

- 1) Physical Phenomena Importance. The PIRT systematically identifies the physical phenomena that are needed for the modeling and simulation application requirements. Each phenomenon must be prioritized, and the criteria applied to accomplish this prioritization must be described. Occasionally, these criteria must span the scenarios of normal, abnormal, and hostile environments.
- 2) Conceptual Model Adequacy. The PIRT identifies the adequacy of the existing set of needed physical phenomena for the application requirements.
- 3) Code Verification Adequacy. The PIRT identifies the adequacy of the set of code-verification test cases.
- 4) Experimental Adequacy. The PIRT identifies the adequacy of existing experimental data required for model validation. If existing data are inadequate for the validation requirements, resources must be allocated and planning begun for the needed additional validation tasks.
- 5) Validation Adequacy. The PIRT identifies the validation metrics needed for the requirements and the accuracy requirements needed for these metrics. Two types of validation adequacy must therefore be addressed: (1) identification of the validation metrics and (2) specification of the magnitude of the metrics to satisfy the application requirements.

Stated differently, both the metrics and the success and failure criteria should be determined for assessing the expected adequacy of the outcomes of experimental-computational comparisons. However, as stated previously, the requirement for the adequacy of all of the specified metrics is coupled to the overall application requirements; hence, the process of specifying the validation metrics will interact with the process of specifying the application requirements.

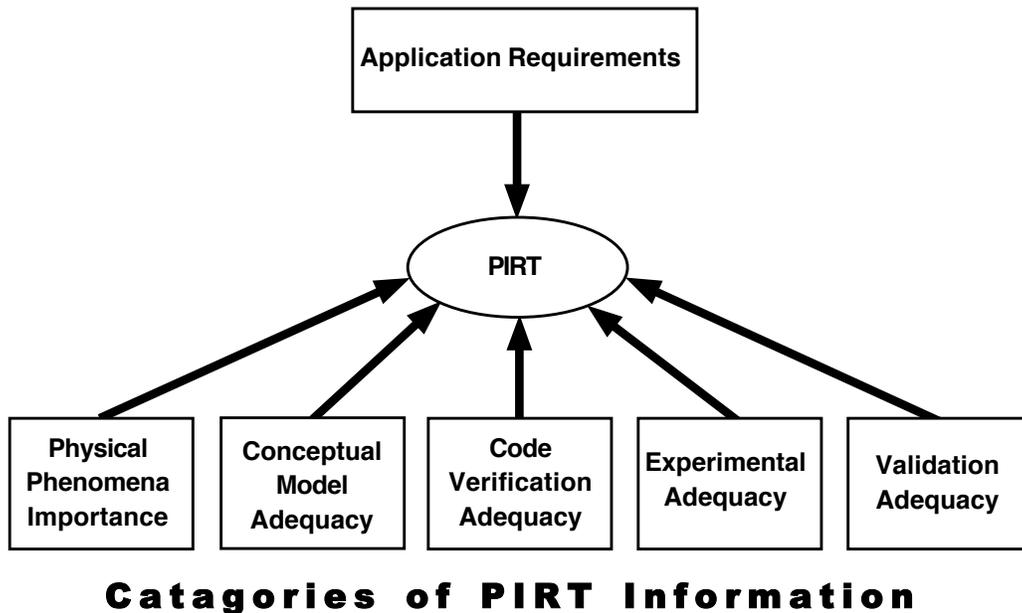


Figure 7
Aspects of PIRT Related to Application Requirements
 (adapted from Ref. [185])

As stressed by Boyack,[22] the PIRT is certainly not set in stone once it is formulated and documented. While a given formulation of the PIRT guides V&V activities, the PIRT must also be adaptable to reflect information gathered during the conduct of these activities. Importantly, to gain the greatest advantage of its value during planning, the PIRT can and probably should be adapted as experimental validation activities are conducted. Several different outcomes illustrate how experimental validation activities can interact with and influence changes to the PIRT:

- A validation experiment may be planned and conducted under the assumption that a specific PIRT element has high importance. After the results of the experiment are analyzed, the importance of that PIRT element is found to change from high to medium or low in response to these results. This does not argue that the underlying application requirements could or would change as a result of experiments, only that the technical importance of an element for validation may change.
- An experiment is conducted that reveals a shift of a PIRT element from low to high importance. This may require, for example, that a subsequent exploratory experiment be performed that was not identified in the existing PIRT.

- An experiment addressing a high-importance PIRT element is performed. The current code implementation addressing that phenomenon is believed to be adequate. However, it is discovered unexpectedly that the code cannot even function properly in defining the proposed experiments, thereby changing the ranking of the implementation to inadequate.
- An experiment designed to probe fully coupled phenomena reveals the presence of a completely unexpected and unknown phenomenon that is of high importance for the application driver. Not only must the PIRT be changed to reflect this event, but also the overall V&V effort for the code application may require significant revision. For example, a previously low-ranked phenomenon may now be ranked high or a planned validation experiment may have to be redefined as a phenomenon-exploration experiment.
- A validation experiment for a single phenomenon reveals that certain models implemented in the code must be recalibrated. This changes the code implementation from adequate to incomplete and may require additional planning for calibration experiments to improve the pre-calibration model capabilities.

2.3 Verification Activities

2.3.1 Fundamentals of Verification

Two types of verification are generally recognized in computational modeling: code verification and solution verification.[161,162] Because of recent work by several investigators,[142,143] we now believe that code verification should be segregated into two parts: numerical algorithm verification and SQA (software quality assurance). Numerical algorithm verification addresses the software reliability of the implementation of all of the numerical algorithms that affect the numerical accuracy and efficiency of the code. In other words, this verification process focuses on how correctly the numerical algorithms are programmed (implemented) in the code. This issue is of paramount importance in computational physics codes, whereas in conventional areas of application of SQA, such as real-time control systems, this issue receives less emphasis. The major goal of numerical algorithm verification is to accumulate sufficient evidence to demonstrate that the numerical algorithms in the code are implemented correctly and functioning as intended. SQA emphasizes determining whether or not the code as a software system is reliable (implemented correctly) and produces repeatable results on specified computer hardware and a specified system with a specified software environment, including compilers, libraries, etc. SQA focuses on the code as a software product that is sufficiently reliable and robust from the perspective of computer science and software engineering. SQA procedures are needed during software development and modification, as well as during production-computing operations.

Solution verification deals with the quantitative estimation of the numerical accuracy of a given solution to the PDEs. Because, in our opinion, the primary emphasis in solution verification is significantly different from that in numerical algorithm verification and SQA, we believe solution verification should be referred to as *numerical error estimation*. That is, the primary goal is attempting to estimate the numerical accuracy of a given solution, typically for a nonlinear PDE with singularities and discontinuities. Assessment of numerical accuracy is the key issue in computations used for validation activities, as well as in application of the code for the intended

application.

The study of numerical algorithm verification and numerical error estimation is fundamentally empirical. Numerical algorithm verification deals with careful investigations of topics such as spatial and temporal convergence rates, iterative convergence, independence of solutions to coordinate transformations, and symmetry tests related to various types of boundary conditions. Analytical or formal error analysis is inadequate in numerical algorithm verification: the code must *demonstrate* the analytical and formal features of the numerical analysis. Numerical algorithm verification is conducted by comparing computational solutions with highly accurate solutions. We believe Roache's description of this, "error evaluation," clearly distinguishes it from numerical error estimation.[163] Numerical error estimation deals with approximating the numerical error for particular applications of the code when the correct solution is not known. In this sense, numerical error estimation is similar to validation assessment. We aim to understand what the computational error is for given calculations or types of calculations, not for code in general.

In our view, to rigorously verify a code requires rigorous proof that the computational implementation accurately represents the conceptual model and its solution. This, in turn, requires proof that the algorithms implemented in the code correctly approximate the underlying PDEs, along with the stated initial conditions and boundary conditions. In addition, it must also be proven that the algorithms converge to the correct solutions of these equations in all circumstances under which the code will be applied. It is unlikely that such proofs will ever exist for computational physics codes. The inability to provide proof of code verification is quite similar to the problems posed by validation. Verification, in an operational sense, then becomes the absence of proof that the code is incorrect. While it is possible to prove that a code is functioning incorrectly, it is effectively impossible to prove that a code is functioning correctly. Single examples suffice to demonstrate incorrect functioning, which is also a reason why testing occupies such a large part of the validation assessment effort.

Defining verification as the absence of proof that the code is wrong is unappealing from several perspectives. For example, that state of affairs could result from complete inaction on the part of the code developers or their user community. An activist definition that still captures the philosophical gist of the above discussion is preferable and has been stressed by Peercy.[154] In this definition, verification of a code is equivalent to the development of a legal case. Thus, numerical algorithm verification and SQA activities consist of accumulating evidence substantiating that the code does not have any apparent algorithmic or programming errors and that the code functions properly on the chosen hardware and system software. This evidence needs to be documented, accessible, repeatable, and capable of being referenced. The accumulation of such evidence also serves to reduce the regimes of operation of the code where one might possibly find such errors.

The view of code verification in this paper as an ongoing process, analogous to accumulating evidence for a legal case, is not universally accepted. In an alternative view,[161] code verification is not considered an ongoing process but one that reaches termination, analogous to proving a theorem. Obviously, the termination can only be applied to a fixed code; if the code is modified, it is a new code (even if the name of the code remains the same) and the new code must be verified again. In addition, all plausible non-independent combinations of input options must be exercised so that every line of code is executed before one can claim that the entire code is verified; otherwise, the verification can be claimed only for the subset of options tested. The ongoing code usage by multiple users still is useful, in an evidentiary sense (and in user training), but is referred to as confirmation rather than code verification. In this alternative view of verification, it is argued that contractual and regulatory requirements for delivery or use of a "verified code" can more easily

be met and that superficial practices are less likely to be claimed as partial verification. Ongoing usage of the code can possibly uncover mistakes missed in the code verification process, just as a theorem might turn out to have a faulty proof or to have been misinterpreted; however, in this view, code verification can be completed, at least in principle. Verification of individual calculations, as well as validation activities, are still viewed as ongoing processes.

Our view of integrating a set of code verification activities into a verification process for computational physics is conceptually summarized in Fig. 8. In this figure we have depicted a top-down process with two main branches: numerical algorithm verification and SQA practices. Numerical algorithm verification, which is the topic discussed in Section 2.3.2, focuses on the accumulation of evidence to demonstrate that the numerical algorithms in the code are implemented correctly and functioning properly. The main technique used in numerical algorithm verification is testing, which is alternately referred to in this paper as numerical algorithm testing or algorithm testing. The branch of SQA practices, discussed in Section 2.3.3, includes practices and procedures associated with SQA. SQA emphasizes programming correctness in the source program, system software, and compiler software. As shown in Fig. 8, software quality analysis and testing can be divided into static analysis, dynamic testing, and formal testing. Dynamic testing further divides into such elements of common practice as regression testing, black-box testing, and glass-box testing.

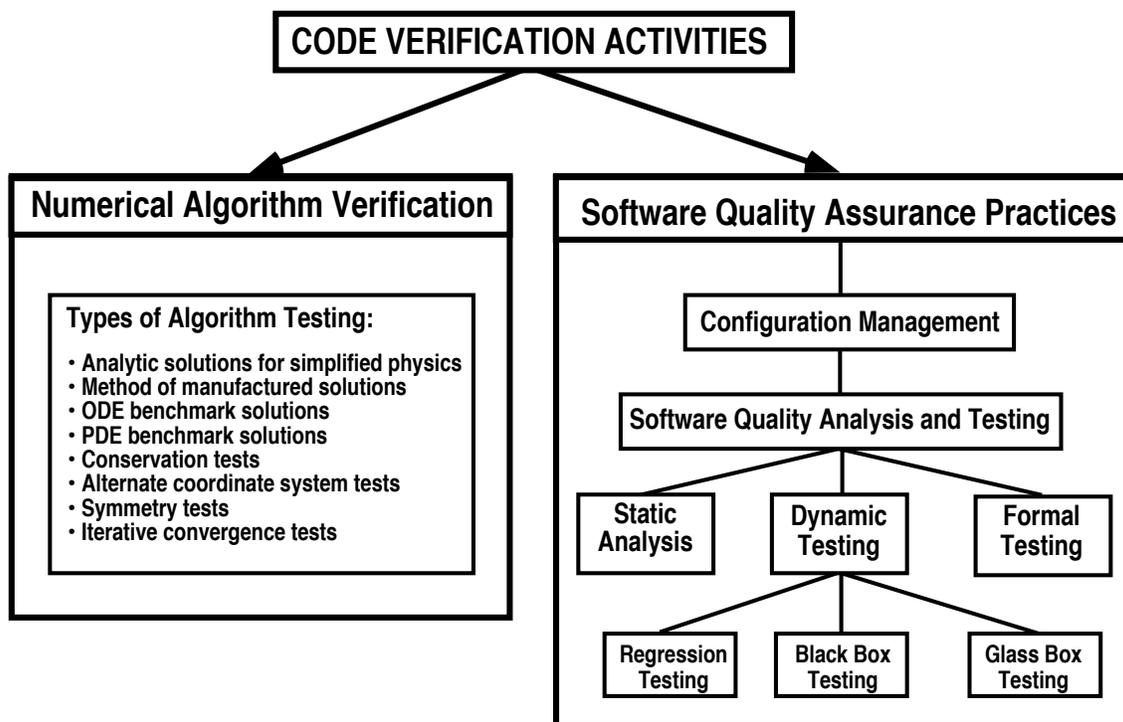


Figure 8. Integrated View of Verification Assessment for Computational Physics (Adapted from Refs. [143,144])

To the disbelief of many, a recent comprehensive analysis of the quality of scientific software by Hatton documented a dismal picture.[81] Hatton studied more than 100 scientific codes over a period of seven years using both static analysis and dynamic testing. The codes were submitted primarily by companies, but also by government agencies and universities from around the world.

These codes covered 40 application areas, including graphics, nuclear engineering, mechanical engineering, chemical engineering, civil engineering, communications, databases, medical systems, and aerospace. Both safety-critical and non-safety-critical codes were comprehensively represented. All codes were “mature” in the sense that the codes were regularly used by their intended users, *i.e.*, the codes had been approved for production use. The total number of lines of code analyzed in Fortran 66 and 77 was 1.7 million, and the total number of lines analyzed in C was 1.4 million. As the major conclusion in his study, Hatton stated, “The T experiments suggest that the results of scientific calculations carried out by many software packages should be treated with the same measure of disbelief researchers have traditionally attached to the results of unconfirmed physical experiments.” Hatton’s conclusion is disappointing, but not at all surprising in our view. We also observe that both Stevenson[182] and Gustafson[75] strongly agree with Hatton’s view that the problems uncovered by Hatton’s static analysis experiments are basically independent of the programming language used.

We believe there are several reasons why disciplined testing of computational physics software is not a common practice. The first reason is that people working on computational physics codes do not like to be thought of as software engineers. A second reason is that the perceived risk associated with the incorrect functioning of nearly all computational physics codes is less than the perceived risk associated with the incorrect functioning of other kinds of software, such as safety-critical systems. For example, Johnson[100] discusses rigorous approaches to developing fault-minimal and fault-tolerant avionics software. There are regulatory requirements governing software verification for avionics applications[61] that are driven by the safety-critical consequences of software failure. We have never encountered a similar paper in computational physics that addresses the underlying consequences of software failure as Johnson addresses them.

2.3.2 Numerical Algorithm Verification

Numerical algorithm testing focuses on numerical correctness and performance of the algorithms. The major components of this activity include the definition of appropriate test problems for evaluating solution accuracy and the determination of satisfactory performance of the algorithms on the test problems. Numerical algorithm verification rests upon comparing computational solutions to the “correct answer,” which is provided by highly accurate solutions for a set of well-chosen test problems. The correct answer can only be known in a relatively small number of isolated cases. These cases therefore assume a very important role in verification and should be carefully formalized in test plans for verification assessment of the code.

There are two pressing issues that need to be addressed in the design and execution of numerical algorithm testing. The first issue is to recognize that there is a hierarchy of confidence in highly accurate solutions. The AIAA Guide,[3] for example, suggests the following hierarchical organization of confidence for the testing of computational physics codes: (1) exact analytical solutions, (2) semianalytic benchmark solutions (reduction to numerical integration of ordinary differential equations [ODEs], etc.), and (3) highly accurate benchmark solutions to PDEs.

The second pressing issue in the design and execution of algorithm testing is to choose application-relevant test problems that will be used. There are two possible approaches for making this selection. One approach is to pick test problems that people have a great deal of experience with. And it would be very advantageous if these problems developed into industry standards that could be used in verification activities for specific engineering or physics disciplines.

Unfortunately, no industry-standard test problems currently exist. A second approach is to construct specialized test problems that address specific needs that arise in the structure of the test plan. These test problems are specifically constructed to exercise the portions of the software that one requires.

Both existing test problems and specially constructed test problems can be either analytical solutions or benchmark solutions. Analytical solutions are closed-form solutions to special cases of the PDEs that are represented in the conceptual model. These closed-form solutions are commonly represented by infinite series, complex integrals, and asymptotic expansions. Numerical methods are usually used to compute the infinite series, complex integrals, and asymptotic expansions in order to obtain the solutions of interest. However, the accuracy of these solutions can be quantified much more rigorously than can the accuracy of the numerical solutions of the conceptual model. When computational solutions are compared with highly accurate solutions, either the comparisons should be examined along boundaries of interest or the error norms should be computed over the entire solution domain. The accuracy of each of the dependent variables or functionals of interest should be determined as part of the comparisons. Note that the most significant practical shortcoming of analytical solutions is that they exist only for very simplified physics and geometries.

A technique for developing a special type of analytical solution to be used for testing numerical algorithms is the “Method of Manufactured Solutions” (MMS).[161,170,173,181] The MMS is a method of custom-designing verification test problems of wide applicability, where a specific form of the solution function is assumed to satisfy the PDE of interest. This function is inserted into the PDE, and all the derivatives are analytically derived. Typically these derivatives are derived by using symbolic manipulation software such as MACSYMA[©] or Mathematica[©]. The equation is rearranged such that all remaining terms in excess of the terms in the original PDE are grouped into a forcing-function or source term. This source term is then considered to be simply added to the original PDE so that the assumed solution function satisfies the new PDE exactly. For example, in the Navier-Stokes equations this term can be considered to be a source term, *i.e.*, a new term on the right-hand side of the PDE. The boundary conditions for the new PDE can be chosen to be the value of the solution function on the boundary (Dirichlet condition), a Neumann condition that can be analytically derived from the solution function, or a boundary condition of the third kind. This approach could be described as finding the problem, *i.e.*, the PDE, for which we have assumed a solution.

Using the MMS in code verification requires that the computed source term and boundary conditions are programmed into the code and that a numerical solution is computed. This technique verifies, although for a narrow range of physical modeling, a large number of numerical aspects in the code, such as the numerical method, differencing technique, spatial-transformation technique for grid generation, grid-spacing technique, and correctness of algorithm coding. Shih and his colleagues[174] have applied the MMS to the incompressible Navier-Stokes equations for laminar two-dimensional flow and have obtained an impressive exact solution to a variant of the classical lid-driven cavity problem for an arbitrary Reynolds number. It is highly recommended that incompressible Navier-Stokes codes be verified with this exact solution.

Salari and Knupp[170] have systematically applied the MMS to the compressible and incompressible Navier-Stokes equations. These researchers have also presented an interesting study of the “bug-finding” capability of the method to better understand what errors the MMS is capable of resolving. In the study, the performance of the MMS on a set of 21 different coding mistakes was examined. The MMS correctly diagnosed every error in the set that prevented the governing equations from being solved correctly. Salari and Knupp (see also Ref. [161]) noted,

however, that the MMS is incapable of detecting errors such as algorithm-efficiency mistakes, which occur in situations where the equations are still solved correctly by the algorithms but the coding is less efficient than is optimal. These results suggest the desirability of coupling studies like Hatton's T1 experiments with a dynamic testing diagnosis like the MMS to better correlate potential structural flaws with actual incorrect numerical performance. Such a study has yet to be performed.

In numerical algorithm testing the key feature to determine is the observed, or demonstrated, order of accuracy from particular numerical solutions. This should then be compared with the formal order of accuracy of the numerical method. Researchers have found a number of reasons why the observed order of accuracy can be less than the formal accuracy. Some of the reasons are: (1) a programming error exists in the numerical algorithm, (2) insufficient grid resolution so that the grid is not in the asymptotic convergence region of the Taylor series expansion for the particular solution variable of interest, (3) the formal accuracy for interior grid points is different than the formal accuracy for boundary conditions with derivatives resulting in a mixed order of accuracy, (4) singularities, discontinuities, and contact surfaces interior to the domain of the PDE, (5) grid clustering, (6) singularities and discontinuities in the boundary conditions, (7) inadequate convergence of an iterative procedure in the numerical algorithm, and (8) over-specified boundary conditions. It is beyond the scope of this paper to discuss these in detail, however some of the representative references in these topics are [8,17,19,26,47,64,104,161,167,168,180,186]. It is our view that most of these reasons for degradation in formal order of accuracy are poorly understood and significant research is needed in this area.

Returning to the topic of types of highly accurate solutions, benchmark ODE solutions are very accurate numerical solutions to special cases of the general PDEs. These solutions commonly result from simplifying assumptions, such as simplified geometries that allow solutions in the formation of similarity variables. Benchmark PDE solutions are also very accurate numerical solutions to special cases of the PDEs or the boundary conditions. The accuracy of the benchmark solutions, whether ODE or PDE solutions, clearly becomes more of an issue as one moves away from analytical solutions. In the literature, for example, one can find descriptions of computational simulations that are considered to be of high accuracy by the author but are later found to be lacking. It is strongly recommended that no published solution be considered a numerical benchmark solution until: (1) the code used in producing the solution has been thoroughly verified and documented, (2) very comprehensive numerical error estimation is reported with the solution, and (3) the solution has been accurately calculated by independent investigators, preferably those who use different numerical approaches and computer codes.

Consistency tests can also be usefully employed as numerical algorithm tests. Global as well as regional tests can be made for the conservation of mass, momentum, and energy.[82] A number of tests can also be made that are related to the effect of numerical boundary conditions on the solution. One group of tests is used to determine if the same numerical solution is obtained in different coordinate systems.[177] Another group of tests is used to evaluate whether certain symmetry features are preserved in the solution. For example, if a plane of symmetry exists in the conceptual model, then the normal gradient of appropriate variables can be set to zero and a solution can be obtained. The same solution should also be obtained if this plane-of-symmetry condition is not imposed and the entire domain is solved. For unbounded domains, the boundaries of the computational domain are conceptually considered to be at infinity, *i.e.*, the boundaries are infinitely far from the spatial region of interest. Typically, a user-defined parameter specifies how "far out" these boundaries are. If the boundaries are too close, the asymptotic conditions applied there may not be accurate. The usual method of determining the size of the computational domain is

to systematically increase the domain until the solution is no longer dependent on the size of the domain that is compatible with the objectives of the computation. Consistency tests should be considered as complementary to the other types of algorithm tests described in this section.

We believe that the primary responsibility for numerical algorithm verification should be placed upon the code development team. Computational analysts who are not directly a part of this team may conduct the algorithm testing, but we believe the code development team needs to have in-depth knowledge of the algorithms used and their expected performance. This information should be documented either as part of the code development process or during testing. The documentation on algorithm testing may be combined with that for SQA testing, as discussed in Section 2.3.3. It is important that analysts and code customers have access to this information.

2.3.3 Software Quality Assurance

Software quality analysis and testing, as highlighted in Fig. 8 presented previously, rests in three techniques: static analysis, dynamic testing, and formal testing.[190] Static analysis techniques analyze the form, structure, and consistency of the code without executing the code. Examples of static analysis techniques are software reviews, complexity analysis, inspections, audits, and analyses of data flows. Dynamic testing techniques involve execution of the code. The results of executing the code are analyzed to detect coding errors or weaknesses in design that can cause coding errors. Regression testing, which re-evaluates the accuracy of computed results, is one example of a dynamic testing technique. Formal testing, also referred to as formal methods, is directed toward rigorously demonstrating that the code exactly represents the underlying conceptual model. Further remarks on static, dynamic, and formal testing are included below. Two particularly useful books for readers who wish to acquaint themselves with the vast subject of SQA testing are by Beizer[16] and by Kaner et al.[103] The differing viewpoints in these books add interesting nuances to a subject as seemingly dry (to practitioners of computational physics) as SQA testing. Beizer's book is somewhat more formal in its approach, while Kaner and his colleagues are less structured in their approaches and recommendations.

It is probably safe to claim that static analysis has traditionally been dominated by techniques to detect compilation errors on the specific code of interest. One doesn't tend to see published studies of, say, applying architectural analysis, such as the complexity metrics used in Hatton's analysis,[81] to computational physics codes. We strongly advocate the expansion of static analysis as part of verification assessment for computational physics codes. The complexity of modern programming languages (C++ especially) and the use of massively parallel computing platforms for computational physics increase the need for, and the impact of, static analysis.

When considering software testing, it is important to understand the distinction between glass-box testing and black-box testing. Glass-box testing refers to testing when all of the source-code of interest is available. Black-box testing refers to testing when only the code inputs and code outputs are observable. Glass-box testing is practiced primarily by code developers because it assumes sufficient knowledge of the design and architecture of the code to design, populate, and assess test plans based on this knowledge. Glass-box testing (also called structural testing) is heavily weighted to what we have called SQA testing, although that is not mandatory. When we discuss static analysis, test problems generated by coverage analysis, and regression testing, we will be discussing examples of glass-box testing. We note that the MMS, discussed in Section 2.3.2, is actually code invasive and thus related to glass-box testing. The MMS cannot be applied exclusively in black-box mode because of the need for specific coding intervention for each test that it creates, at least in the examples of which we are aware. Once that coding is accomplished,

however, the gathering of results proceeds in a fully black-box manner. There does not appear to be any analog of the MMS testing methodology in the work of Beizer[16] and Kaner et al.,[103] which is rather surprising.

Black-box testing (also called functional testing) is primarily the paradigm for the algorithm testing discussed above. Black-box testing can be performed by anyone involved with a given computational physics code, but it tends to be associated with independent testing entities. Certainly in the computational physics world, much black-box testing is performed by users of the code. (If a developer of the code is also a user, for purposes of this discussion we emphasize their role as a user.) Black-box testing requires no detailed knowledge of the code software, although such knowledge can certainly help. Even when those who designed the algorithms and implemented the code execute a test problem to investigate the empirical performance of the code, they are performing black-box testing. When a computational physics algorithm is described in a publication, for example, and compared with some highly accurate solution, a black-box test has been performed. The goal of the authors in this case is to assess the functionality and accuracy of the output results, not to test specific software elements. When users test codes by running their favorite test problems prior to application, they are performing black-box testing. Users are an enormous resource for black-box testing. In our discussion of dynamic testing, aspects of regression testing, and statistical testing, we are addressing issues that are particularly relevant to black-box testing. The goal of a coordinated, optimal test strategy for a computational physics code is to blend black-box testing and glass-box testing.

Coverage analysis and code sensitivity analysis are methods that can be used to assess the complexity of the computational physics code in order to design collections of test problems. Coverage analysis enables one to determine what components of the code enter into various options exercised in the code. The relative importance of the selected component contributions to that calculation is then determined via a sensitivity analysis. Ideally, one would start coverage analysis with a tool of sufficient capability to assess code lines and units exercised by the code in performing the desired calculation. For example, PureCoverage™ [156] is a commercial tool that detects lines of code executed during the operation of C++ software and reports this information in a useful form. A major challenge of coverage analysis for algorithm testing is to understand execution paths as well as lines touched in the code in order to execute the intended application calculation. One then designs tests that specifically target these modules and paths. This is clearly a glass-box-testing strategy and is discussed extensively in both Beizer[16] and Kaner et al.[103] While a discussion of sensitivity analysis is beyond the scope of this paper, the topic is brought up in the context of software quality testing in Ref. [133].

A natural partner of coverage analysis for software quality testing is regression testing. A widely used commercial testing technique, regression testing is defined by Beizer[16] as “any repetition of tests (usually after software or data change) intended to show that the software’s behavior is unchanged except insofar as required by the change to the software or data.” What this means in practice for computational physics is that a compendium of tests is assembled, a baseline for code performance against these tests is established, and the tests are run periodically. When the current performance is determined to have deviated from the baseline, either the decision is made that a bug has been introduced or the regression test suite is baselined again. There are two requirements that make this approach feasible. First and foremost, the test suite must run sufficiently fast so that the testing can be performed frequently. Second, the regression testing should be performed for every software modification. For multiple developers of a computational physics code, for example, this means that no modified software is accepted before regression

testing is performed. This, of course, can impede the software development process if not managed properly.

Regression tests that are glass-box tests are often correlated with coverage analysis and designed to provide testing of as many lines of code as possible. Systematic, cyclic-regression testing, say once a suitable set of test problems has been defined, is known to be important in code development and maintenance. Unfortunately, this requires a trade-off between available resources that influences the fidelity of the testing. For example, a high-fidelity test suite that takes longer than 5 to 10 hours to execute is probably not practical for execution every day. Thus, such a suite cannot be used to demonstrate the stability of daily code builds.

The main use of regression testing is to minimize the effort devoted to fixing incorrect new software that is introduced into existing code in a software development project. There is a balance that must be determined between the effort devoted to regression testing and the effort devoted to finding and fixing incorrect modifications of the existing code at a later date. At the same time, regression testing aims to cover as much of the code as possible. In the computational physics projects that we are familiar with, the anticipated coverage target for regression testing is typically on the order of 80% of the lines of code. Accepting 80% line coverage is commonly based on the developers willingness to define specific code elements as not requiring coverage. More importantly, these coverage thresholds are established in the attempt to prevent the effort devoted to regression testing from becoming larger than the effort required to fix introduced errors.

A great deal of effort has recently been devoted to the development and application of formal methods.[21,150,169] However, the actual goal of these methods—rigorous “proof” that a system of software is correctly implemented—remains controversial in our opinion. The application of formal methods is also complicated by disputes over cost, appropriateness, utility, and impact.[20] Formal methods have certainly not been applied to software systems like those of interest in computational physics, namely, those systems in which floating-point arithmetic is dominant and an effectively infinite variability in software applications is the norm. The utility of these methods to computational physics codes, even if resource constraints were not issues, has not been established. This fact has led to interest in the application of formal methods to more restricted problems that still may be of significant interest in computational physics software. For example, there is current interest in applying formal methods to aid in verification of the mathematical formalism of the conceptual model that underlies a computational physics code rather than in the full details of the software implementation.[34,114]

We conclude this section on SQA by providing a few brief remarks and references for some of the key literature that is relevant to the overall problem of verification assessment. A number of modern texts are available that describe current practice in the SQA field.[9,16,41,43,55,62,101,122,153,165] Much of the research and development in SQA procedures has been fueled by computer-controlled systems that require extremely reliable and secure software, as in Johnson’s avionics-systems application referenced previously. Such systems are commonly referred to as “high-integrity systems.” Examples of other high-integrity systems are control systems for nuclear power reactors and software for nuclear-weapon security and safety. The scientific software community has much to learn from SQA procedures that have been developed for these systems as high-consequence scientific computing becomes more prevalent.

2.3.4 Numerical Error Estimation

The two basic approaches for estimating the error in a numerical solution to a PDE are *a priori* and *a posteriori* approaches. An *a priori* approach uses only information about the

numerical algorithm that approximates the partial differential operators and the given initial and boundary conditions. *A priori* error estimation is a significant element of classical numerical analysis for PDEs, especially those underlying the finite element and finite volume methods.[64,92,93,115,129,146,161] An *a posteriori* approach uses all of the *a priori* information, plus computational results from a previous numerical solution using the same numerical algorithm on the same PDE and initial and boundary data. In this paper we also refer to *a posteriori* error estimates as “empirical” data or, as Roache[161] call it, “observed” data. Empirical data could be, for example, a solution on a single coarse grid or a sequence of solutions on consecutively finer grids. We believe the only quantitative assessment of numerical error that can be achieved in practical cases of nonlinear PDEs is through *a posteriori* error estimates. As a result, only *a posteriori* will be discussed here.

A posteriori error estimation has primarily been approached through the use of either Richardson extrapolation[161] or estimation techniques based on finite element approximations.[4,12] Richardson’s method can be applied to both finite difference methods and finite element methods. As pointed out by Roache,[161] Richardson’s method produces different estimates of error and uses different norms than the traditional *a posteriori* error methods used in finite elements.[93,158] A Grid Convergence Index (GCI) based on Richardson’s extrapolation has been developed to assist in the estimation of grid convergence error.[159-161] The GCI converts error estimates that are obtained from any grid-refinement ratio into an equivalent grid-doubling estimate. Recent studies have shown that the GCI method is fairly reliable, even for solutions that are not in the asymptotic-convergence region.[24,28,59,166]

Richardson’s extrapolation can be applied to compute error estimates of dependent variables at all grid points, as well as error estimates for solution functionals. Solution functionals are integrated and differentiated quantities such as body lift and surface heat flux, respectively. Venditti and Darmofal[188] discuss the coupling of error estimation and adaptive grid generation to reduce numerical errors in such computed functionals. Their particular example is the quasi-one-dimensional flow in a variable-area duct. It is important to emphasize that different dependent variables and functionals converge at different rates as a function of grid size. For example, the grid and time step that are required to show second-order convergence in heat flux on a body surface are extremely finer than for total lift on a body. This fact was also pointed out by earlier researchers,[99,161] as well as in the AIAA Guide.[3]

A posteriori error estimates are also important for finite element adaptivity, where both the spatial grid density (h-adaptivity) and the order of the finite element scheme (p-adaptivity) can be adapted. The important role of *a posteriori* estimates in adaptivity for finite elements is discussed in many papers, as well as in the recent books by Ainsworth and Oden[4] and Babuska and Strouboulis.[12] For numerical error estimation, however, finite element adaptivity is not the issue. Although traditional *a posteriori* methods used in finite elements are very useful and computationally efficient, these methods do not demonstrate asymptotic convergence or directly address useful engineering error measures. Recent work in *a posteriori* error estimation addresses error in local quantities of engineering interest, as well as error estimation for certain types of nonlinearities in elliptic problems.[4,12] Regardless of the *a posteriori* error estimation approach chosen, the estimation should always be made for the system response outputs that are relevant to the application of interest.

The assumption of smoothness in solutions, *i.e.*, no singularities and discontinuities, is quite demanding in estimating local errors in the solution domain. Singularities and discontinuities commonly occur in fluid dynamics, solid mechanics, and structural dynamics. The “pollution” of particular regions of a calculation by the presence of singularities such as shock waves, geometrical

singularities, or crack propagation is a subject of grave concern in error estimation. Often, the only clear sense of this pollution available to us is through careful empirical assessment. Since the technical aspects of this issue are beyond the scope of this paper, the reader should consult a series of papers by Babuska and his colleagues[10,11,13] as well as the paper by Oden[145] for a discussion of this problem from a finite-element point of view. Roache[161] has a wide-ranging discussion of this topic, and the recent work of Zhang and his colleagues[196] discusses how the presence of a shock wave structure affected *a posteriori* error estimates for the Euler equations. A recent paper of Botella and Peyret[19] discusses similar problems associated with computing singular solutions of the Navier-Stokes equations.

An additional numerical error estimation complexity should be mentioned that can occur in certain types of computational physics simulations. This complexity can be conceptually described as a coupling between numerical error and the appearance of new spatial and temporal scales in certain types of physical models. In fluid dynamics, for example, when grids are sufficiently refined then completely new processes or characteristics on the smaller spatial scales can develop. It is not uncommon that a steady flow field can become unsteady as the grid is refined and very small-scale phenomena can develop that did not exist on the coarser grids. This occurs in the case of the trailing edge region of a compressor or turbine blade, where vortex shedding is known to exist and will appear in the simulation if the grid around this region is sufficiently refined. Refining grids does not ensure that the physics modeled will remain unchanged as the grid is resolved. This observation directly influences the accuracy and reliability of any type of *a posteriori* error estimation method, especially extrapolation methods.

Extracting an estimate of numerical accuracy on underresolved grids is a topic of current research. We previously mentioned one aspect of the problem, generically referred to as pollution error—the propagation of discretization error from less accurately resolved regions of a calculation into more accurately resolved regions. Another issue is simply to understand what information can rationally be synthesized from calculations with more- or less-known resolution accuracy. Though this topic is beyond the scope of this paper, Chorin[31-33] and Glimm[70,73] demonstrate that techniques for understanding the numerical accuracy of underresolved calculations are challenging and specialized. A critical theme in this literature, and one that might be generalized to other important problems, is the role of statistical estimation in performing such assessment. We conclude by commenting that studying V&V for underresolved calculations directly addresses the challenges posed by Gustafson in his 1998 paper.[75]

2.4 Validation Activities

2.4.1 Fundamentals of Validation

We have emphasized that validation experiments must be designed purposefully with specific goals linked to application objectives and to specific elements of the PIRT. The primary goal of directed validation experiments is to ensure that experimental data are sufficiently related to the application driver to provide stringent confidence assessment when code calculations are compared with these data. These data must be precisely located within the structure of the application-specific PIRT and must be unambiguous in defining confidence in the application of the code to the phenomena exemplified in the experiment. Because this task must be accomplished by quantitatively comparing code calculations with experimental data, important requirements are placed on validation experiments to create the greatest opportunities for performing these comparisons. It is critically important to design and execute validation experiments that allow

precise and conclusive comparisons of calculations with experimental data for the purpose of assessing code fidelity and credibility. In Section 2.4.5 we further discuss code-experiment comparisons. It will hopefully be made clear that the design and execution of experiments must allow us to quantify meaningful and useful metrics.

Particular validation experiments may achieve these goals to a greater or lesser degree, but any such attempt rests on a foundation of rational design concepts that are based on the needs of the application of the code. The more purposefully we design an experimental validation activity, the greater the probability that we can optimize the results of that activity in the face of these complex constraints. It is imperative that validation experiments balance resource constraints, including time, level of effort, available expertise, and desired fidelity. The approach for achieving this balance should be defined in the experimental plan. How to achieve this balance will depend strongly on which tier of physics complexity the experiment is conducted. As briefly mentioned in Section 1.2, validation experiments can be conducted at the unit-problem, benchmark, subsystem, or system tier. The resources needed vary greatly from one tier to the next. Section 2.4.2 discusses some of these topics in more detail.

The experimental, computational, and comparison activities should expedite the credibility assessment of computational models for system applications that are within the application domain. It is important to attempt to quantify the boundary separating the region of acceptability of the model from the region where the model is not acceptable for the application. Designing experiments that test a code in regions where the code is believed to be insufficiently accurate for the intended application helps locate this boundary and provides a means for quantitatively assessing the degree of the expected inaccuracy. Because such experiments are performed purposefully rather than accidentally, these experiments also further test our grasp of the conceptual models underlying the code that are probed by the validation experiments. Obviously, this goal only makes sense when experiments that probe code inaccuracy lie close enough to the boundary of applicability to be relevant. It is desirable to have experimental validation tasks that have the explicit goal of defining those application domains where use of the code is questionably adequate to better quantify the boundary of applicability of the code. To perform these validation tasks in a conscious and scientifically directed manner is not necessarily any easier than to achieve desirable levels of confidence in other application domains.

Figure 9 illustrate a typical relationship between the application domain and the validation domain, *i.e.*, the apparent region of acceptable computational capability. These domains are defined in a multidimensional parameter space where only two of these dimensions, X_1 and X_2 are shown. Two kinds of points are also shown. An “A” denotes a location in the application domain where a specific application of a code has previously been performed; a “V” denotes a location in the domain where validation has been performed. The boundary of the validation domain represents the apparent limit where the physical models in the code would no longer be reliable. The boundary of the validation domain could be estimated, for example, by using alternate plausible computational models or by using expert opinion. We draw the reader’s attention to two pairs of intended applications, each pair denoted “A_I” and “A_O”. The applications “A_I” each lie inside of the validation domain, while the applications “A_O” each lie outside of the apparent validation domain. Our point is that it is essentially as important to know that the code is questionable for the applications “A_O” as it is to know that the code is acceptable for the applications “A_I”. A complete validation effort could quantify this situation, but such an effort is commonly not possible, for example, because of programmatic schedules, budgets, safety and environmental concerns, or international treaties.

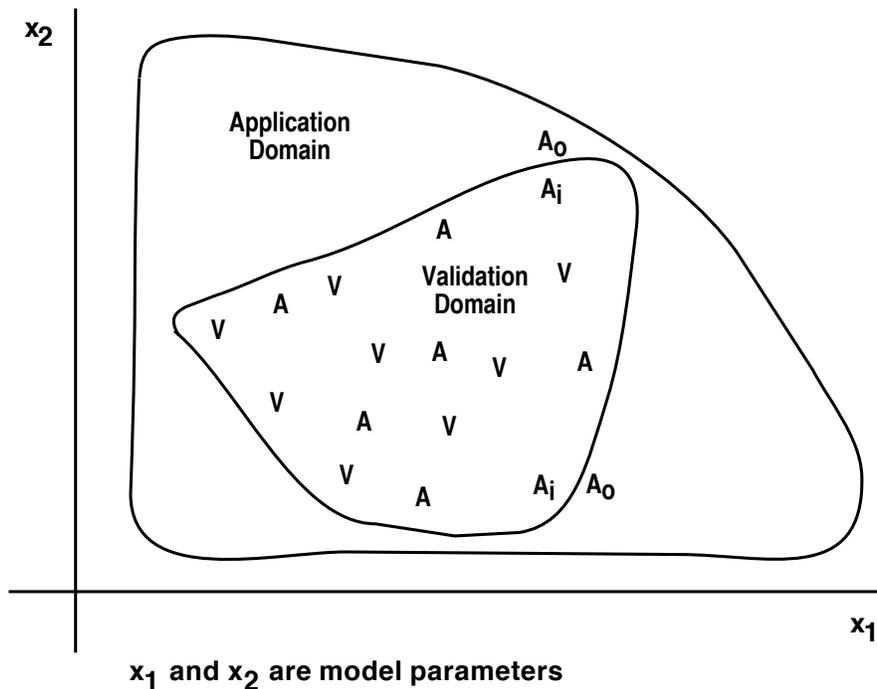


Figure 9
Typical Relationship Between the Application Domain and the Validation Domain
(from Ref. [185])

There should be a well-reasoned basis for choosing between two potential validation experiments. The basis should depend on the potential benefits of defining the acceptable domain boundary and on the use of available time and resources. The code-application V&V plan should be helpful in choosing between the two potential experiments, especially through utilization of PIRT prioritizations. In partnership with the PIRT, two approaches are recommended: competing alternate plausible models and statistical experimental designs. Competing alternate plausible models can be of two different types: alternate models that are generally at the same level of modeling fidelity, and hierarchical models. Using hierarchical models, one formulates conceptual models that are clearly at different levels of fidelity or physical detail. Then results from the models at different levels of complexity are compared to attempt to assess where the lower fidelity model could begin to fail. Statistical experimental design (for example, see Cox[39] and Dean and Voss[42]) is an attractive basis for attacking the resource optimization problem for experiments with a focus on uncertainty quantification. Gunter[74] has argued in favor of using statistical experimental design in planning physical-science experiments. The specifics of how the statistical design of experiments may be applied in particular validation activities are dependent upon the subject-matter focus of these experiments.

2.4.2 Construction of Validation Experiment Hierarchy

Any appropriate methodology for the purposeful definition, design, and analysis of experiments for validation of a code must consider the intended application. As discussed in Section 2.2, the PIRT is the key methodology for identifying the predictive-capability needs and

setting the validation priorities. To implement the needs and priorities identified in the PIRT, the recommended methodology is the use of a validation hierarchy, as discussed in Section 1.2. The hierarchical view of validation is distinctly an engineering perspective, as opposed to a scientific or model-building perspective. The purpose of the validation hierarchy is to help identify a range of experiments, possible separation of coupled physics, and levels of complexity—all of which are related to an engineering system—so that computer codes from different disciplines can be evaluated. Stated differently, the validation hierarchy must be *application* driven to be of engineering value, *not* code driven.

As one constructs each lower tier of the hierarchy, the emphasis moves from multiple coupled-physics codes to single codes (or single-physics options in multi-physics codes) simulating a particular type of physics. In parallel with the simplification process, the focus on the actual operating conditions of the complete system should not be lost. Constructing these hierarchical tiers and identifying the types of experiments that should be conducted at each tier are formidable challenges. There are many ways of constructing the tiers; no single construction is best for all cases. We would draw the analogy of constructing validation hierarchies to the construction of control volumes in fluid-dynamic analyses. Many varieties of control volumes can be drawn; some lead nowhere, and some are very useful for the task at hand. The construction should emphasize the modeling and simulation capability that is desired to be validated, whether that capability be computational physics or other computational disciplines. Analogous tier structures can be developed for structural dynamics and electrodynamics, for example, when the engineering system of interest involves these disciplines.

A good hierarchical tier construction is one that accomplishes two tasks. First, the construction carefully disassembles the complete system into tiers in which each lower-level tier has one less level of physical complexity. For complex engineered systems, this may require more than the three building-block tiers shown in Fig. 2, presented previously. The types of physical complexity that could be uncoupled from one tier to the next are spatial dimensionality, temporal nature, geometric complexity, and physical-process coupling. The most important type of physical complexity is physical-process coupling, which commonly contains the highest nonlinearity of the various contributors. The nonlinear coupling is, of course, eliminated by conducting experiments with uncoupled and separate physical effects, *e.g.*, separating turbulence and combustion. It is important to recognize the potential nonlinear nature of all of the contributors in the construction of the tiers because the philosophy of the tier construction rests heavily on linear-system thinking. That is, the belief that confidence in the computational capability for the complete system can be built from an assessment of the computational capability of each of its parts. The complete systems of interest clearly do not have to be linear, but the philosophy of the hierarchical validation approach loses some of its utility and strength for strong nonlinear coupling from one tier to the next.

The second task accomplished by a good hierarchical tier construction is the selection of individual experiments in a tier that are practically attainable and able to produce validation-quality data. In other words, the individual experiments should be physically achievable given the experimental test facilities, budget, and schedule, and they should be capable of producing quantitative experimental measurements of multiple system-response measures that can test the code. For complex systems, the ability to conduct a true validation experiment at the complete system tier is extremely difficult, if not impossible. At the subsystem tier, it is usually feasible to conduct validation experiments, but it is still quite difficult and expensive. One usually chooses a single hardware subsystem or group of subsystems that are closely related in terms of physical processes or functionality. For complex subsystems, one might want to add a new tier called

“components” below the subsystem tier. As with the subsystem tier, the components tier would consist of actual operational hardware components. When one defines the individual experiments at the benchmark-tier level, special hardware should be fabricated. By “special hardware” we mean that the hardware need not be constructed to perform its system tier or subsystem tier functions.

The benchmark tier is probably the most difficult to construct because it represents the transition from a hardware focus in the two top tiers to a physics-based focus in the bottom tiers of the hierarchy. At the bottom tier, unit problems, one should identify simple geometry experiments that have a single element of physical-process complexity. As with the subsystem tier, an additional tier may need to be added to attain only one element of physics at the bottom tier. Also, the experiment must be highly quantitatively characterized so that it can provide the necessary data to the computational code, and the experiment must be conducted so that experimental uncertainty can be estimated precisely. High-quality validation experiments are practically attainable at the benchmark and unit-problem tiers, but usually not at the system or subsystem tiers for complex systems.

Recently, an example of a hierarchical tier structure for a complex, multidisciplinary system was presented in Ref. [142]. The example features an air-launched, air-breathing, hypersonic cruise missile. The missile is referred to as the complete system, and the following are referred to as systems: propulsion, airframe, guidance, navigation, and control (GNC), and warhead. The example assumes that the missile has an autonomous GNC system, an on-board optical target seeker, and a warhead. Figure 10 shows the system-level hierarchical validation structure for the hypersonic cruise missile. The structure shown is not unique, nor is it necessarily optimum for every computational-simulation perspective of the missile system. In addition, the structure shown in Fig. 10 focuses on the airframe system and the aero/thermal protection subsystem for the purpose of analyzing the aero/thermal performance of the missile.

To better demonstrate how the validation hierarchy of the aero/thermal protection subsystem is related to the validation hierarchy of the propulsion, GNC, and warhead systems, the example was discussed further in Ref. [142]. Figure 11 shows how the validation hierarchy of each of these four systems could be viewed as the primary facets of a four-sided pyramid. The airframe facet was divided into three additional facets, each representing the three subsystems: aero/thermal protection, structural, and electrodynamics. The propulsion system was divided into four additional facets to represent its subsystems: compressor, combustor, turbine, and thermal signature. Similarly, the GNC and the warhead systems could be divided into subsystems appropriate to each. On the surface of this multifaceted pyramid, one can more clearly and easily indicate the coupling from one facet to another. For example, we discussed the coupling of laminar and hypersonic flow with ablation to the optical seeker of the GNC system. This coupling would be shown by an arrow connecting these hypersonic flow elements to appropriate elements on the GNC facet of the pyramid. (Suggested references stressing a systems engineering perspective are Refs. [18,76,117,139,195].)

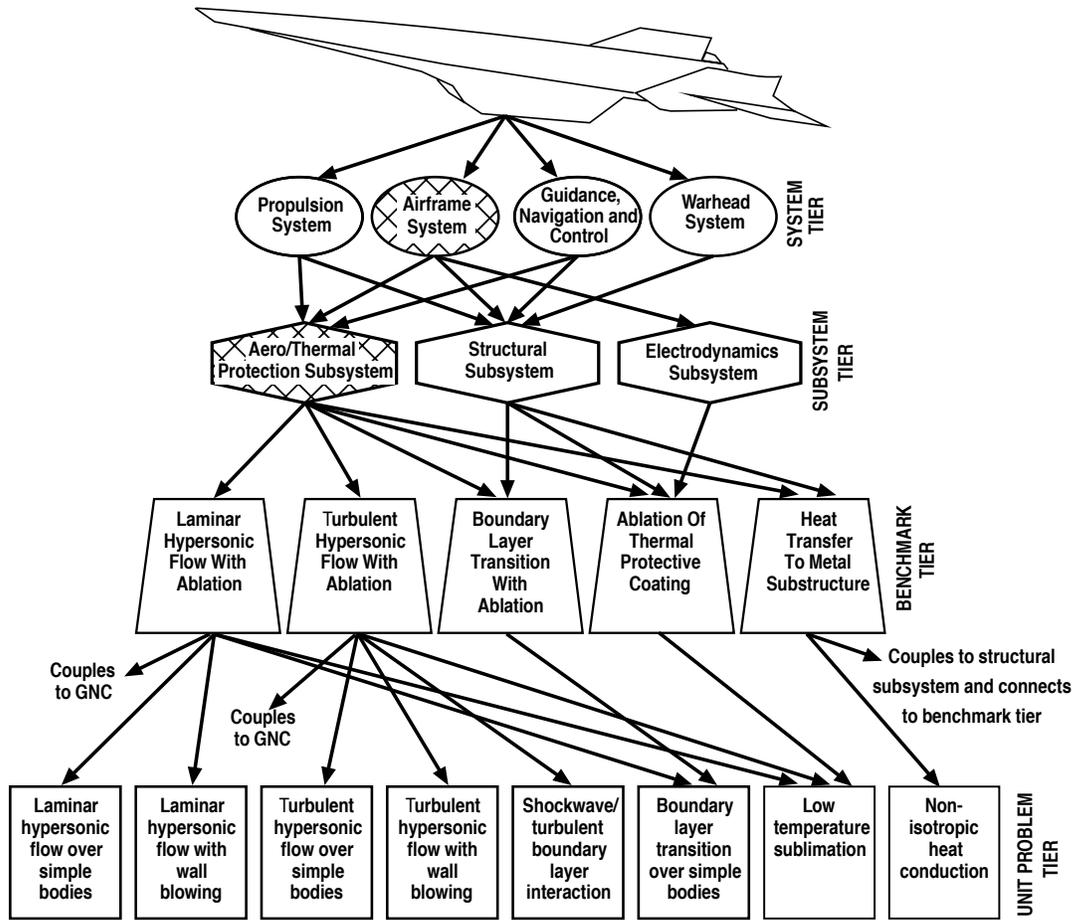


Figure 10
Validation Hierarchy for a Hypersonic Cruise Missile[142]

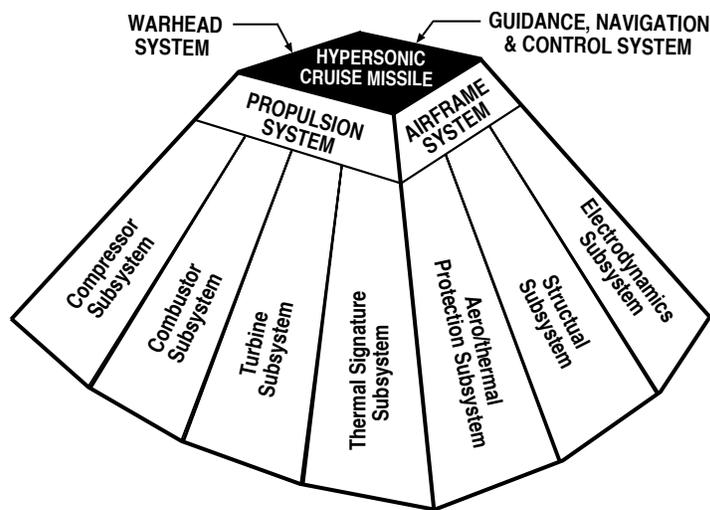


Figure 11
Validation Pyramid for a Hypersonic Cruise Missile[142]

The validation pyramid stresses the systems-engineering viewpoint in modeling and simulation-based design, as opposed to the viewpoint of a specific discipline. Each facet of the pyramid can then be devoted to identifying validation experiments for each computational code responsible for part of the design of the system. As one traverses around the top of the pyramid, the number of facets is equal to the number of systems that are identified. As one traverses around the bottom of the pyramid, the number of facets is equal to the total number of major computer codes used in the analysis of the engineering system, *i.e.*, the number of codes that require validation activities for the intended application. For the example of the hypersonic cruise missile, if the code that simulates surface ablation is a separate code from the aerodynamics code, then an additional facet on the pyramid is added on the aero/thermal subsystem facet. We strongly believe this type of system-level thinking is necessary to increase confidence in complex systems that are designed, manufactured, and deployed with reduced levels of testing.

2.4.3 Characteristics of Validation Experiments

Many researchers, analysts, and managers ask, What is a validation experiment? or how is a validation experiment different from other experiments? These are appropriate questions. We suggest that traditional experiments could generally be grouped into three categories. The first category comprises experiments that are conducted primarily to improve the fundamental understanding of some physical process. Sometimes these are referred to as physical-discovery experiments. The second category of traditional experiments consists of those conducted primarily for constructing or improving mathematical models of fairly well-understood flows. The third category of traditional experiments includes those that determine or improve the reliability, performance, or safety of components, subsystems, or complete systems. These experiments are commonly called “tests” of engineered components or systems.

We argue that validation experiments constitute a new type of experiment. A validation experiment is conducted for the primary purpose of determining the validity, or predictive accuracy, of a computational modeling and simulation capability. In other words, a validation experiment is designed, executed, and analyzed for the purpose of quantitatively determining the ability of a mathematical model and its embodiment in a computer code to simulate a well-characterized physical process. Thus, in a validation experiment “the code is the customer” or, if you like, “the computational analyst is the customer.” And only during the last 10 to 20 years has computational simulation matured to the point where it could even be considered as a customer. As modern technology increasingly moves toward engineering systems that are designed, and possibly even fielded, based on modeling and simulation, then modeling and simulation itself will increasingly become the customer of experiments.

In our view, there are three aspects that should be used to optimize the effectiveness and value of validation experiments. The first aspect is to define the expected results of the experimental validation activity using the code itself. The second aspect is to design specific validation experiments by using the code in a predictive sense. The third aspect is to develop a well-thought-out plan for analyzing the computational and experimental results. These aspects emphasize tight coupling of the subject code to the experimental activity and ensure we are in the best position to learn from the comparison of computational and experimental results, regardless of whether these comparisons are good or bad, as explained below.

Consider the second aspect above related to the design of a validation experiment. Suppose through a series of exploratory calculations for a particular application of the code that an

unexpectedly high sensitivity to certain physical parameters appears. Specific PIRT elements for the application may have postulated little sensitivity to the parameters, or the PIRT may have stated that the sensitivity was unknown. If this unexpected sensitivity has an important impact on the application of interest, a validation experiment assessing the correctness of this sensitivity should be pursued. The code that is the object of the validation exercise should participate in defining the principles and goals of a planned validation experiment. An alternative phrase to keep in mind when the word “define” appears is “define the purpose of the experiment.” The role of the code in the definition of validation experiments may be purely exploratory until some progress has been made on validation activities.

One way of designing a validation experiment is to use the code calculations as specific guidance on where to locate the instrumentation and what kind of data to acquire for assessing the anticipated sensitivity. In the proposed validation experiments, “design” means specifying to the greatest degree possible the initial and boundary conditions, material properties, diagnostic locations and characteristics (Strain gauge? Stress gauge?), and data fidelity. The probability that these conditions “defined” through code calculations will be met precisely in any complex validation experiment is probably zero, but the probability that they will be met in simpler experiments is relatively high. In most cases, the success of a validation experiment will often be determined by the degree to which the experiment matches these specifications. Deviations can be acceptable—but if the intent of the experiment was to measure x at location y for material z and geometry g and in all cases the experiment is significantly different from most or all of these factors, it is unlikely that the experiment would be a successful validation experiment.

Engaging a code in the definition and design of validation experiments can easily be distinguished from more common experimental practice. In the conventional case of nonvalidation experiments, the experimenters state the following to the computational analyst in the typical approach: “We are performing the following experiment. Tell us what the expected response of gauge a will be, given this experimental design and location for the gauge.” For a validation experiment, on the other hand, the computational analyst makes the following statement to the experimenter: “Please perform the following experiment. We recommend locating gauge a at location x based on our computational predictions. It will be very useful to see whether our predictions agree with your experiment. If your experiment has deviated from our requested design, we still anticipate achieving useful validation consequences through further postexperimental analysis, as long as the design deviation is not great.”

A number of researchers, particularly experimentalists, have slowly been developing the concepts of a validation experiment. During the past several years, a group of researchers at Sandia National Laboratories has been developing philosophical guidelines and procedures for designing and conducting a validation experiment. Although the following six guidelines and procedures were developed in a joint computational and experimental program conducted in a wind tunnel, they apply over the entire range of computational physics:[1,2,135-138,189]

Guideline 1: A validation experiment should be jointly designed by experimentalists, model developers, code developers, and code users working closely together throughout the program, from inception to documentation, with complete candor about the strengths and weaknesses of each approach.

Guideline 2: A validation experiment should be designed to capture the essential physics of interest, including all relevant physical modeling data and initial and boundary conditions required by the code.

Guideline 3: A validation experiment should strive to emphasize the inherent synergism between computational and experimental approaches.

Guideline 4: Although the experimental design should be developed cooperatively, independence must be maintained in obtaining both the computational and experimental results.

Guideline 5: A hierarchy of experimental measurements of increasing computational difficulty and specificity should be made, for example, from globally integrated quantities to local measurements.

Guideline 6: The experimental design should be constructed to analyze and estimate the components of random (precision) and bias (systematic) experimental errors.

Validation experiments should not produce data that fundamentally depend on code calculations for critical data-reduction tasks. Such data do not correctly address our need for independent comparability of experimental data with calculations and violate our need for robustness of the experimental data. Experimental data that require code calculations for evaluation can never be a desirable outcome for a validation experiment, although exactly this problem may arise in other kinds of experiments. An example of what we mean is the problem of inference of material temperature from experimentally acquired shock hydrodynamic data (density, pressure, and velocity fields) using code calculations of shock-wave physics rather than some type of diagnostic to directly measure the temperature. The only possible validation data that will emerge from shock hydrodynamics experiments without temperature diagnostics are the shock hydrodynamic data. This problem is relevant, since it arises in investigations of temperature dependence in high-pressure shock-driven material response. Such experiments often need code calculations to estimate the associated thermal conditions under shock loading. For purposes of scientific discovery, this is permissible though speculative. Such experiments, however, cannot be claimed to provide validation data for the high-pressure thermomechanical response of the material because of the lack of independence of calculations and experiment.

The purposeful design of validation experiments is enabled and accomplished through a diverse team of people that participate in the experimental validation activity. Most obviously, one or more experimenters are participants in this team. However, code users (analysts, designers) must also participate, given our proposed primary role of the code in the definition, design, and analysis of validation experiments. Finally, one or more code developers should be on the validation experiment team. Their presence provides valuable expert knowledge about the perceived *a priori* capability of the code in all three aspects of experimental validation: define, design, and analyze. In fact, code developers, including experts in the physical models in the code, are the most knowledgeable about the boundary of the acceptable application domain discussed in Fig. 9.

Attempting to sharply define the boundary of applicability of the code for a given application domain through a deliberate choice of experiments close to the boundary, but on the invalid side, has greater leverage when code predictions are used to design such experiments. We stress once again that achieving a satisfactory level of performance for the code in comparison with validation experimental data in a case of true prediction has far greater power to raise our level of confidence in the application of the code. We may not be able to quantitatively measure this fundamental

observation at this time, but it is obviously true. Predictive accuracy is gold, while posttest consistency is merely brass at best and possibly fool's gold at worst.

2.4.4 Uncertainty Quantification in Computations

As mentioned in Section 2.4.1, it is common when simulating validation experiments that one encounters physical parameters, *e.g.*, in the PDEs or in the initial or boundary conditions, that are not precisely known or measured for an experiment. This situation is much more common at the system and subsystem tiers than at the benchmark and unit-problem tiers. Examples of such parameters are thermochemical transport properties, flow rates in complex systems, and inflow nonuniformities. Another common situation occurs during a series of experiments, *e.g.*, weather-driven experiments and flight tests, when there are certain parameters that are poorly controlled or not controllable at all. For any parameter, we make the assumption that a value of that parameter is required to perform the computational simulation. The situations just mentioned are all characterized by stochastic parametric uncertainty. In the following discussion, we simply refer to “uncertainty,” but we will always be referring to parametric uncertainty. This type of uncertainty requires nondeterministic simulations in validation of computational physics.

One standard approach is to estimate, by one means or another, a single value of such a parameter and compute a solution with that selected value. This might be a fully adequate way of dealing with this uncertainty, especially if experience suggests that the range of potential parameter values is very small and that the calculation is known not to be extremely sensitive to the parameter in question. The resulting calculation intrinsically is interpreted as “representative,” or “best estimate” or “conditional” for that parameter.

The shortcomings with the standard approach referred to above begin to be noticeable when the range of variation of the parameter is large or when the calculation is known to be sensitive to the parameter values. If multiple required parameters in a computation of a validation experiment are uncertain, it is entirely possible that the interaction of these parameters in the calculation may magnify the influence of their uncertainty on the final results of the calculation. In our opinion, this statement is especially important when performing calculations that are intended for direct quantitative comparison with validation experiments. We believe that the uncertainty of the parameters should be incorporated directly into the computational analysis.

The simplest strategy for incorporating uncertainty of this kind directly into the computation is performed in three steps. The first step, called *characterizing the source* of uncertainty, is based on the assumption that the uncertainty in the parameters of interest is characterized by probability distributions. Sometimes such distributions can be directly estimated if a large quantity of experimental data is available for the parameters. Sometimes such distributions simply must be assumed. At any rate, the first step involves specifying the probability distributions and understanding the credibility of these assumptions.

In the second step, *ensemble computing*, values from the input probability distributions specified in the previous step are selected using statistical sampling procedures, such as Monte Carlo or Latin Hypercube sampling methods (see, for example, Refs. [40,67]). These sampled values are then used in a set of computations. Because this latter statement is so important, we will restate it. The assumed prior probability distributions for the uncertain parameters are used to generate a set of calculations. This set is sometimes called an ensemble of calculations.

The key issue is that a single calculation is no longer sufficient; a set of calculations must be performed. Obviously, this need is disturbing—where once one might have performed a single

calculation, now one must perform a potentially large number of calculations. We have not raised nor answered the question of whether sufficient computational resources are available to execute more than the one calculation. However, the constraints enforced by the availability of computing may be formidable.

After the set of calculations has been generated, the third step, *uncertainty quantification of the output*, is performed. This step involves analysis of the set of calculations, typically using statistical inference, to estimate a probability distribution for the output variables of interest that results from the given input parameter distributions. In general, we cannot deduce the exact output distribution that results from the assumed form of the parameter distributions used to generate the computational input associated with those parameters. Instead, the common practice is to use statistical procedures to determine estimates of important parameters associated with that output distribution.

Such statistically determined estimates are useful for comparing computational results with experimental measurements. For example, the mean of the output calculations provides us with an estimate of the expected value of the output quantity of interest, given the uncertainty structure specified by the input distributions. This mean-value estimate is of primary importance when comparing computational results with the mean value of multiple experimental realizations. Another statistic of interest is the estimated variance of the output distribution, which can be interpreted as a measure of computational output uncertainty, or scatter, given the input uncertainty.

For readers unfamiliar with this methodology, we stress that it is rarely true that the mean of the output given the input uncertainty can be determined by performing a calculation for a single set of inputs that is chosen to be the mean of each of the input distributions. Stated in another way, the mean value of the output cannot be computed by performing a simulation using the mean value of all input parameters except when the mapping of inputs to outputs is linear in the parameters. Linearity in the parameters essentially never occurs when the mapping of inputs to outputs is given by a differential equation: *even a linear* differential equation. Instead, we must perform the ensemble of calculations to develop a statistically rational estimator for the mean. The previous statement is also true for estimating other output statistics. Kleijnen[105] provides a broad summary of methodologies that go beyond Monte Carlo for assessing output distributions statistically.

In summary, the general methodology we are thus advocating for incorporating parameter uncertainty into computational physics computations is to execute all three steps in the manner suggested above. Performing the three-step strategy will clearly be nontrivial for hard computational problems simply because of the computational burden imposed. There are also certain subtleties that we have failed to mention, such as whether complex structure in the resulting output probability distribution can actually be discovered using such a crude approach. We simply state that extracting an intricate output-distribution structure will require either a large number of sample input values or considerably more sophisticated approaches for performing the methodology. The only modeling and simulation applications we are aware of within computational physics that pursue this methodology are the fields of underground transport of toxic waste materials, or pollutants,[15,83,85,120,183] and climatology.[5,30,53,77,149,171]

We now return to our mention of Bayesian inference. In step 2 of the three-step strategy, the problem of propagating input uncertainty through a computational model to understand the resultant output as described above is sometimes referred to as the *forward uncertainty problem*. There is an associated *inverse uncertainty problem*, or *backward problem*, which is conceptually and mathematically much more difficult.[72] The backward problem asks whether we can reduce the output uncertainty by updating the statistical model using comparisons between computational

and experimental results. The inverse uncertainty problem is a calibration problem, except in the inverse uncertainty problem the parameters being calibrated are given by probability distributions. For example, might we be able to improve our original prior distributions that characterize the parameter uncertainty? This problem can be cast as a problem in Bayesian statistical inference. (See Refs. [56,66,69] for an introduction to Bayesian inference.)

Part of the difficulty alluded to above is related to providing a better understanding of computational accuracy when it is known that we are in an underresolved grid or time-step situation. This is an important research topic because of its practical consequences. Recent work attacks this problem and illustrates the formidable difficulties in two distinctly different areas: fluid dynamics.[46,70] and dynamical systems[31-33] While the applications and specific technical attacks of researchers in these two areas are distinctly different, we find it fascinating that a common deep thread in their work is the treatment of insufficient information using statistical methods. We should stress, however, that these authors do not discuss one problem of interest to us—the problem of validation of underresolved computational models. If reliable estimates of error due to underresolved grids can be made, then we believe model validation can be conducted. The disadvantage, however, is that if large error bounds are estimated, there can be a great deal of room for the computational results to agree or disagree with the experimental data. Stated differently, only *weak* validation conclusions can be drawn if either the computational or experimental error estimates are large. If *no* estimate of grid convergence error can be made for the computational solution, we strongly believe *no* validation conclusion can be made. The same argument can be made if only one experimental measurement is made. However, it is traditional in engineering and science that relying on only one experimental measurement is more acceptable than not estimating grid convergence error, regardless of whether or not this position is actually defensible. The issue of how experimental uncertainty should be treated in validation is addressed in Section 2.4.5.

2.4.5 Validation Metrics

The specification and use of validation metrics comprises the *most important* practice in validation activities. Validation metrics are used to quantitatively compare the results of code calculations with the results of validation experiments. The straightforward interpretation of this word is that a “metric” is simply a “measure.” Thus, the choice of one or more metrics defines the means used to measure the differences between computational results and experimental data. Because we emphasize that the overarching goal of validation experiments is to develop sufficient confidence so that the code can be used for its intended application, we do not view the use of metrics as simply a passive means of measuring differences between computational results and experimental data. Metrics must be devised that actively resolve assessment of confidence for relevant system response measures for the intended application of the code. Referring back to Fig. 5, it is clear that validation metrics are important for situations portrayed in Figs. 5a and 5b, but they are *crucial* for quantitatively estimating confidence in modeling and simulation capability for Fig. 5c. The definition, design and application of these metrics should directly flow from the PIRT activities described in Section 2.2.

A metric should quantify both errors and uncertainties in the comparison of computational results and experimental data. Specific examples of using metrics can be found in a number of references.[37,51,52,57,58,68,78,90,91,106,116,118,119,142,147,148,184,187] Although there is a large amount of recent work in this topic, it is our view that this area of validation quantification techniques is in its infancy, despite its great importance.

Figure 12 helps to clarify where we believe the present state of the art is in computational-

experimental comparisons and to identify the direction in which we believe the topic needs to progress. This figure illustrates the conceptual increase of quantification in performing validation comparisons as increased attention is paid to both experimental and computational uncertainty, as well as numerical error. The figure depicts a desirable progression from (a) qualitative comparisons in which experimental data and calculations are presented side by side on a viewgraph without any information about uncertainties in either one, to (f) analysis of quantitative differences between experiments and calculations accounting for uncertainty and error. Let us consider this figure in some detail.

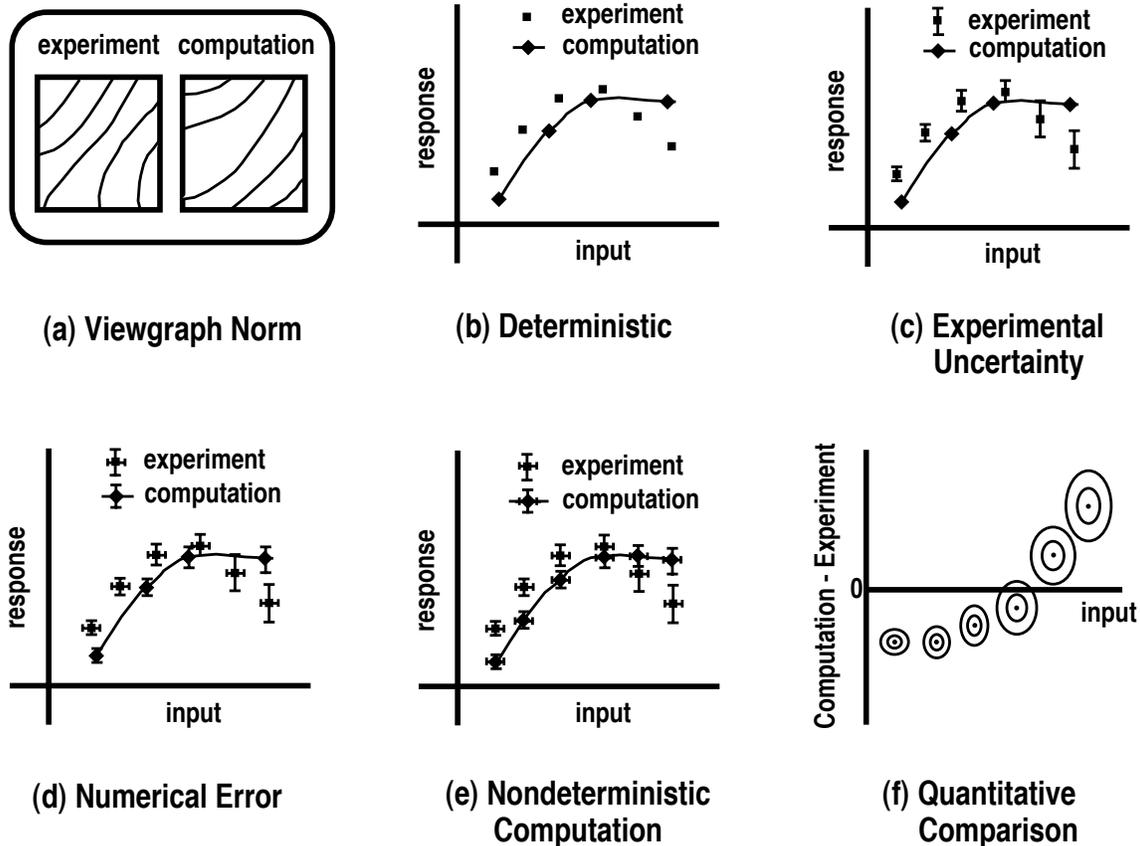


Figure 12
Increasing Quality of Validation Metrics
 (adapted from Ref. [185])

Figure 12a depicts a qualitative comparison we refer to as a “viewgraph-norm” comparison of experimental and computational data that is often seen in practice. This comparison is marked by the qualitative nature of the chosen metric, typically comparison of color contours over a domain. It is also distinguished by no information on experimental and computational uncertainty. No quantitative statement about validation can seriously be made based on such a comparison, although the statement may provide some level of confidence in the application of the code at an intuitive level. Intuition and degree of agreement between computational results and experimental data, of course, is in the eyes of the beholder.

The plot in Fig. 12b portrays the most common type of comparison between computational results and experimental data. Generic axis labels of system “input” and “response” are shown,

although both axes can also be system response quantities over a range of an input quantity. While discrete experimental and computational points are shown in this plot, the concept also encompasses curve overlays without any points shown. The key problem with metrics implemented at the level of Figure 12b is that there is no recognition of uncertainty in the comparison. Yet, conclusions drawn from this type of comparison are really only qualitative, such as, “fair agreement” or “generally good agreement.”

Figure 12c suggests that the next step for improving the method of comparison is to place qualitatively estimated error bars around the experimental data. By “qualitatively estimated error bars” we mean the common practice of the experimentalist either (a) quoting the manufacturer’s stated accuracy of the diagnostic sensor, or (b) making a statement such as “Essentially all of the experimental data fell within the error bars shown.” At this stage of comparison a validation metric could be constructed that accounts for two characteristics: the difference between the computation and experiment integrated over the range of the input parameter, and some estimate of the experimental uncertainty.

Figure 12d represents the case where there is a more quantitative estimate of experimental uncertainty and there is an estimate of computational error. The information concerning experimental uncertainty could be improved, for example, in two ways. First, multiple experimental realizations could have been obtained so that the experimental data point shown would represent the mean of all of the samples, and the error bar would represent two standard deviations of an assumed Gaussian probability distribution. Second, an estimate of the experimental uncertainty in the measured input quantity is also obtained and is shown as an error bar. Concerning the computational error, an *a posteriori* numerical error estimate from the computation would be given for the specific system response quantity that was measured. The computational error estimate may have come from, for example, a single-grid solution or an estimate based on solutions from multiple grids.

Figure 12e suggests a further improvement in estimation of experimental uncertainty and also computation of nondeterministic simulations. For example, the experimental uncertainty could include an estimate of correlated bias errors using the procedure discussed earlier, or an estimate of uncertainty based on the same experiment conducted at separate facilities. Concerning the nondeterministic computational simulations, we are referring to an ensemble of computations at each of the experimental conditions, as discussed earlier. For example, multiple computations would be made using the experimentally estimated probability distribution for the input quantity. As a result, the computational data point would be the mean of the nondeterministic computations for both the input and the system response. The “error bars” for the input quantity and the system response would represent two standard deviations for each of the computational probability distributions. Note that to use the experimental uncertainty distributions for the input quantity, computations would need to be made at the measured input conditions.

Figure 12f shows a true quantitative measure of the comparison between the computations and the experimental measurements over the range of the input quantity. In terms of information content, one could have the same data as contained in Fig. 12e. However, we believe a validation metric should be based on the difference between computations and experiments over the range of measurements of the input parameter. Assuming that probability distributions for both computational and experimental data are known, as discussed in Fig. 12e, comparing computations and experiments will require a difference, or more properly, a convolution of pairs of probability distributions. (See, for example, Refs. [23,44] concerning convolutions of distributions.) The elliptical symbols in Fig. 12f are meant to signify one and two standard deviation contours of the

convolutions of the computational and experimental probability distributions. The “dot” in the center of each contour is the difference in the mean, or expected value, of the computational and experimental distributions.

As proposed in Refs. [142,144] we believe that all reasonable validation metrics should include a number of useful properties, although validation metrics need not be restricted to satisfying only these properties. The following is a list of properties that we believe a validation metric should have:

- 1) A metric should incorporate an estimate of the numerical error in the computational simulation. This estimate may only be an estimate from solutions on underresolved grids. However, we believe that representing the uncertain estimate of error as a probability, *e.g.*, 1 out of 100 chance, is not defensible, since convergence of numerical error is more closely analogous to bias errors in experimental measurements.
- 2) A metric should not exclude any modeling assumptions or approximations used in the computation of the simulation result. A metric should reflect all uncertainties and errors incurred in the modeling process.
- 3) A metric should incorporate estimates of the random errors in the experimental data that are the basis of comparison. In addition, we believe a metric should also include an estimate of the correlated bias errors in the experimental data, if possible.
- 4) A metric should depend on the number of experimental replications of a given measurement quantity. That is, a metric should reflect the level of confidence in the experimental mean that has been estimated, not just the variance or scatter in the data.
- 5) A metric should be able to incorporate computational uncertainty that is due to both random uncertainty in experimental parameters required for defining the calculations and any uncertainty that is due to lack of experimental measurement of needed computational quantities. Thus, a metric should use nondeterministic methods to propagate uncertainty through the subject computational model.

The importance of quantifying the experimental uncertainty has been emphasized in the discussion of experimental design above. Quantifying experimental uncertainty is crucial for validation metrics. The inference from validation experiments to the application domain as discussed in Section 2.1 is greatly weakened in the absence of quantifying the experimental uncertainty. Similarly, it is critically important to estimate numerical errors in the computational simulation that are compared with the experimental results. Otherwise, in computational-experimental comparisons, the numerical errors can be indistinguishable from the modeling uncertainties.

It is also important to properly represent the uncertainty observed in the experiment. For example, observed experimental variations can be propagated through ensembles of calculations using techniques for propagating uncertainty, and the computational results can then be analyzed to provide statistical data about the resulting computational uncertainty. In practice, however, difficulties arise for several reasons. First, the raw effort involved in obtaining ensembles of calculations to propagate uncertainty becomes highly taxing for complex computations. Second, it is very difficult to accurately assemble the experimental uncertainty data into precise statistical

information for the routine propagation of computational uncertainty. The inability to accurately characterize experimental uncertainty as probability distributions may be due to, for example, inadequate diagnostic information on physical parameters or to limited time and financial resources. When information is lacking, referred to previously as epistemic uncertainty, some of the newer methods in information theory could be used. However, these methods are not well developed for practical engineering problems. Even with these difficulties, we believe the definition and use of validation metrics can still proceed even if information about computational uncertainty is lacking or poor. In effect, when a poor job is done in quantifying the computational uncertainty, the quality of the conclusions that can be drawn from the validation metrics is adversely affected. However, when a poor job is done in quantifying the experimental uncertainty, the goal of the validation experiment can be defeated.

Comparing an ensemble of calculations with an ensemble of experimental data is a statistical problem of quantifying uncertainty. Statistical analysis of the differences between computational and experimental results must play a key role in the analysis and presentation of metric data. A fundamental goal of a validation metric is then to present at least the impact of the uncertainty of the experimental data upon the inferences that are drawn from the validation exercise. It is important to recognize that uncertainty in the experimental data (and of the computations) affects the credibility of the results of a validation metric. This fact, in turn, will impact the confidence in the proposed application domain of the code.

3. Major Research Issues

3.1 Prioritization of assessment activities

In this section we advocate widespread use and further research into the use of the PIRT for prioritizing V&V activities and explain our reasons for doing so. The use of the PIRT for this purpose is fairly well developed in the field of reactor-safety risk assessment; however, its use in other fields has been very limited. We believe the PIRT is the only method presently available to account for the needs of the intended application, the believed present capability of a code, and areas where the code is weak. When these factors are combined, with prioritization of multiple intended applications, one is able to identify and prioritize the necessary V&V activities. In our use of the PIRT in the ASCI program at Sandia National Laboratories we have learned some important lessons that can inform future practice in V&V. The first lesson learned is that when the intended applications of a code are being identified and prioritized, the process requires close interaction between the computational analysts and the “customers” of the computational analyses. The customers may be design engineers, safety engineers, reliability engineers, and project engineers who would all typically serve as decision makers about engineering-related issues. Without such interaction, computational analysts can lose touch with the perspectives and priorities of their analysis customers. The second lesson learned is that the PIRT can uncover requirements for computational simulations that had not been recognized as important or that had not been recognized at all. For example, the PIRT can be particularly helpful in identifying important unrecognized areas of coupled physics, either on one face of the validation pyramid or coupling between different faces of the pyramid (see Fig. 11).

We also learned that the PIRT can be helpful in identifying a spectrum of different types of experiments and tests that are likely to be needed in the future. The following types of experiments and tests have been identified:[185]

- Phenomena-exploration experiments
- Mathematical-model-development experiments
- Calibration (model-updating) experiments
- Validation experiments
- System and certification tests

The key difference between this spectrum and the hierarchical validation experiments discussed earlier is the goal of each: validation experiments focus on the *assessment* of modeling capability, regardless of the level of physical complexity.

This spectrum of experiments influences the development and growth of confidence in computational models and their applications in an evolutionary sense, as depicted in Fig. 13. In this figure a serial view of the role of the above experiments is emphasized, as well as the interaction of experiments and the code. In the left column we cast our view of experimental activities in an ordered sequence consisting of the four types of experiments: phenomena exploration, mathematical model development, calibration, and validation. In reality, of course, there are possibly local loops that have not been emphasized in Fig. 13 that further connect each of these experimental elements. For example, it is apparent that the conduct of experiments that are designed to be validation experiments could reveal the need to perform experiments from any other element in the column.

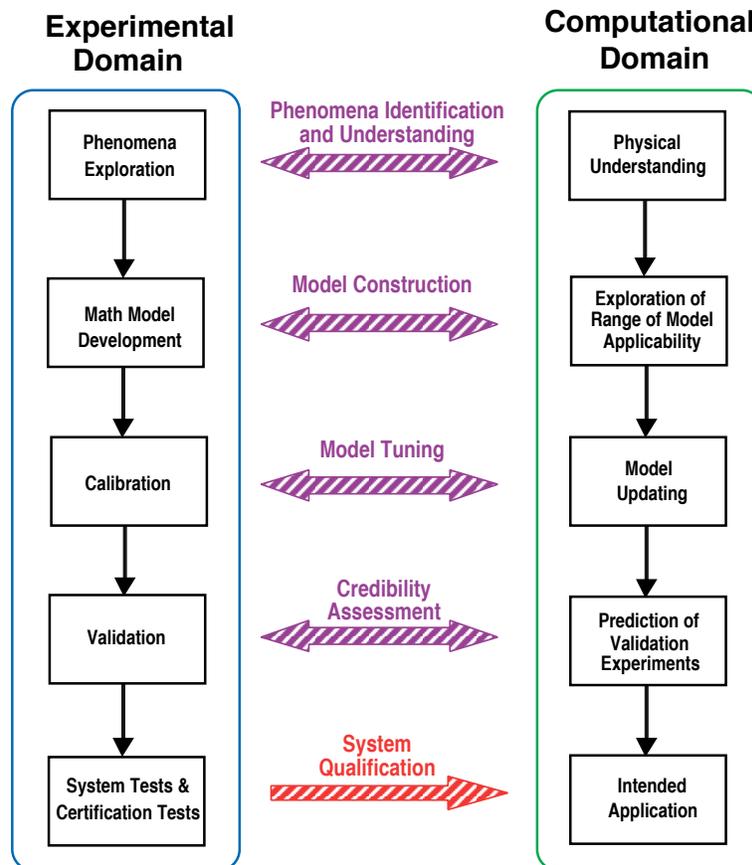


Figure 13
Interaction of Various Experimental and Computational Activities
 (adapted from Ref. [185])

Each element in the experimental domain directly interacts with an element in the computational domain. The motivating concepts generating these computational elements are suggested by the arrows in the middle of the figure that depict the flow of experimental information directly to the computational domain and vice versa. We have simplistically characterized the computational domain through the types of activities expected to be most associated with the directly correlated experimental element. We believe that our depiction in Fig. 13 highlights the right connections between the experimental and computational domains if the experimental domain is not restricted only to validation experiments as they are defined in this report. In particular, the ultimate goal of both experimental and computational activities is to provide a basis of confidence and decision making for the intended application of the code. We have highlighted the great weight that system tests and certification tests evoke by emphasizing a more unidirectional flow of information between the experimental and computational domains at this level. The technical weakness to their genuine physical reality, if the system can be tested, is that the test is only one physical realization of a complex, nondeterministic system.

3.2 Verification activities

This section identifies areas of improvement that can help advance the effectiveness of the verification process in computational physics. We recommend the development of standardized test problems, the further development of the Method of Manufactured Solutions, and the use of statistical methods for the quantification of SQA. The issue of verification in chaotic systems is also briefly addressed.

Previously we have argued for the special importance that highly accurate solutions play in numerical algorithm verification. For brevity, we will simply refer to highly accurate solutions as “benchmarks.” These benchmarks can be analytical or numerical solutions. There are two senses in which we mean benchmarks for use in algorithm testing. Weak-sense benchmarks are test problems that are in common *ad hoc* use by various algorithm developers or code projects for the purpose of assessing numerical accuracy. We define strong-sense benchmarks fundamentally to be *engineering standards*. We will discuss strong-sense benchmarks in Section 4.4.

We believe using benchmarks constructed by the Method of Manufactured Solutions (MMS) are the most effective technique for numerical algorithm verification. Although many code developers question the effectiveness of MMS, we have found that everyone who has used the technique is now a strong believer in its power. Users of the method have demonstrated that it is extraordinarily sensitive to coding errors of the numerical algorithms: as Roache puts it “annoyingly so.” In fluid dynamics, MMS benchmarks should be developed, for example, in supersonic flows with shock waves, turbulent flows, reacting flows, multiphase flows, non-Newtonian flows, free-surface flows, and large-eddy-simulation models. We recommend that MMS benchmarks be developed and applied to other areas of physics, for example, radiation transport, large plastic deformation, fracture dynamics, electromagnetics, and quantum mechanics. Two additional topics of MMS that require further research, for all areas of physics, are (1) correct imposition of boundary conditions for a mixture of elliptic, parabolic, and hyperbolic systems and (2) determination of types of code errors that are *not* detectable by individual applications of the MMS. Stated differently, the limits of the MMS should be investigated more thoroughly so that any redundancy with traditional analytical-solution verification testing could be better understood.

Statistical testing methods for the quantification of SQA are becoming more widely used. We

recommend that these methods should be developed more thoroughly, primarily by research teams composed of computer scientists and statisticians. We believe that statistical testing methods can be effectively used in the computational engineering and physics community, however, we are not aware of any studies to date. The recent text by Singpurwalla and Wilson[178] is an excellent introduction to the topic of statistical testing and detailed methods for statistical quantification. To contrast the view of statistical testing of software with the view expressed in the present paper, we note that the book by Singpurwalla and Wilson does not use the term “verification” in anywhere the Table of Contents or in the Index.

The final research topic in verification activities that we discuss in this paper goes beyond the topic of error estimation. We can ask fundamental questions of how well the computational world matches the “true” dynamics of the underlying equations. For example, how well are the dynamical system features of the underlying differential equations matched by the computational dynamics? In the case of steady flows, this question is concerned with the stability and bifurcation behavior of the computed solutions. Attempting to measure the computational “error” in this case is very difficult. For example, it is only in relatively simple cases that the “true” stability and bifurcation behavior may be well understood for the conceptual model. An entire taxonomy of potential threats to computational accuracy arises in this point of view. Potential dangers are summarized in Ref. [194] and the work cited there.

3.3 Validation activities

In this section we identify areas of improvement that can help advance the effectiveness of the validation process in computational physics. We recommend the construction and use of a validation hierarchy and the specification and use of quantitative assessment criteria for validation metrics at the different tiers of the hierarchy. Several suggestions are also made regarding the need for developing additional validation metrics based on statistical measures, as well as those that are applicable to a wider range of physical processes.

The construction and use of a validation hierarchy for complex engineering systems is relatively new, but we believe this approach will be fundamentally useful in the future. In fact, we believe it is the *only* constructive approach for decomposing complex systems that can build demonstrable confidence in individual computational models. Consequently, we recommend that organizations that use computational modeling in the design, certification, or production of engineered systems should begin constructing and using the validation hierarchical structure. Admittedly, effective validation hierarchies—those that emphasize single-physics phenomena at the lower levels and engineered system-performance priorities at the higher levels—are difficult to construct. These hierarchies can only be constructed in an iterative fashion with input from a wide range of professionals—from engineering analysts to design engineers to product marketing staff. In our limited use of the approach, we have found that a validation hierarchy quickly points out areas where little or no validation evidence exists. The ability of this approach to elucidate insufficiency or weakness in the validation evidence may not be appealing to some advocates of computational simulation, but knowledge of such deficiencies is essential from an engineering perspective.

Validation experiments conducted at the higher tiers of the hierarchy (subsystem and system) will invariably have poorly characterized experimental data for input to the computational simulation. As a result, validation will necessitate probabilistic treatment of uncertain parameters in the computational physics submodels or in the initial conditions or boundary conditions for the PDEs. Propagation of these parameters or conditions through the computational physics model will

likely rely on probabilistic sampling methods like Monte Carlo or Latin Hypercube sampling. Importantly, major increases in computational resources will be required for the tens or hundreds of solutions needed for the sampling techniques. Additionally, improved training and understanding for probabilistic sampling techniques will be required for the computational physics analysts involved in statistical validation activities. On this topic, we believe that the computational physics community can learn a great deal from the probabilistic structural dynamics community and the probabilistic risk assessment community and can use computational tools developed by these communities.

We believe the validation team must be able to assess the results of the validation metrics. This activity is referred to as “Metrics Assessment” (block 6) in Fig. 6 presented previously. The primary goal of such an assessment is to credibly establish whether the agreement of calculations with experimental data satisfies the application requirements. This is the key purpose of dedicated validation experiments, and it is what distinguishes validation from a straightforward computational simulation of an experiment. Failure to assess the validation metrics with respect to the intended application makes it virtually impossible to understand what information about model confidence has been generated as a result of the validation activity. In conventional practice, assessment criteria for validation metrics are typically formulated, if at all, after the experimental data have been collected. We believe it is very useful to define some of these assessment criteria, assuming this can be done, during the “Intended Application” activity and the “Planning” activity (see Fig. 6). More detailed specification of the assessment criteria will be possible after the validation experiments are complete, but it will still be difficult.

Specifying the assessment criteria for validation metrics will be difficult for two reasons. First, application requirements for the accuracy of computational analyses are commonly very vague. When application requirements are given, they are typically given for relatively high-level measures of system response. For example, a system-level requirement for debris containment during the catastrophic failure of a gas turbine engine may simply be that all engine-part debris must be contained within the engine shroud. Second, if the application requirement gives little or no guidance for detailed system-response measures, how should these assessment criteria be determined? For example, in the case of debris containment described above, how should the application requirement be determined for models for metal tearing or large plastic-deformation metal folding?

Two approaches are recommended for guidance in determining the assessment criteria at the different tiers that are shown in Fig. 10 presented previously. First, the metrics assessment activity (identified in Fig. 6) is not intended to completely answer the question of whether the code is suitable for the intended application. Clearly, the intended application should dominate the criteria for assessing the results of validation metrics. The fact remains that the quality of comparisons of computational and experimental results can also be judged somewhat independently of their ultimate application, purely as a problem in a scientific area. People who historically engage in computational modeling of experiments make these judgments as a matter of course, but rarely are these judgments quantified or formalized to the degree being advocated for validation purposes in this paper. Second, guidance for assessment criteria for validation experiments done at the unit-problem tier and the benchmark tier should also be derived from sensitivity and uncertainty analyses done at both the subsystem tier and the system tier. Stated differently, the impact on the system tier of given levels of accuracy at the lower tiers should be determined with sensitivity and uncertainty analyses. Although these analyses are hard to accomplish early in the development of a code, they will need to be done during the intended application activity and the planning activity.

While the questions and concerns listed above for specifying the assessment criteria are

significant, it is our view that the quantification of requirements at multiple tiers is the correct approach. When this approach is compared to the traditional test-based engineering approach, it is seen that the approaches are actually analogous. The test-based engineering approach, however, is more *ad hoc* at the lower tiers, and there is much more emphasis on testing of the full-scale system.

Making validation-metric success criteria as quantitative as possible is important. It is possible that success criteria may be qualitatively or indirectly defined, perhaps through a process of assimilated experience during the course of the validation activity. In our opinion, however, qualitative criteria, or criteria that are developed after the activity of comparing computational and experimental results has begun, make the validation job more difficult and undermine the credibility of the assessment. In addition, qualitative criteria that are developed after the execution of experimental validation activities lead to the danger of focusing the definition of success on discovering any feature in the comparison that is “good” at the expense of answering the harder question: Is there an important reason that this feature in the computational-experimental comparison should be good?

Attaching quantitative success criteria to the definitions of validation metrics eliminates ambiguity from the meaning of computational-experimental comparisons. The net effect of this statement is that a validation metric could also be viewed as a pair of concepts—one being the definition of the metric, the other being a definition of what it means for the resulting comparison to be “good” or “successful.” Defining the metric and its associated success criteria may be quite difficult in various circumstances, but this goal is very desirable, nonetheless, and should be attempted in all validation activities.

Additional research is needed to define validation metrics for a wide range of system-response measures. For example, in fluid and solid dynamics we note that validation metrics are needed that weight important regions of the domain more heavily than less important regions. Metrics are also needed for unsteady continuum mechanics. If the computational solutions are periodic or quasi-periodic, we believe that validation metrics should be formulated in the frequency domain, not in the time domain. As is well known from structural dynamics, the frequency domain eliminates the unimportant feature of phase shifts in the periodic response and concentrates on natural modes in the system. For example, Urbina and Paez[187] have focused development of validation metrics for structural dynamics models using frequency domain metrics, in particular a windowed spectral response function metric. Apparently random, possibly chaotic, temporal acceleration histories at various locations in complex structural systems result from close, but not solid contact, between structural elements. We observe that the application of frequency-domain validation metrics in complex fluid flows has some relationship to certain techniques that have been developed for assessing prediction uncertainty in climate modeling using singular-value decompositions.[149] For other kinds of flows, of course, elaboration of complicated time-domain metrics may still be important. For example in an oil reservoir flow problem, Glimm et al.[71] devise a complex time-domain metric that is used to predict confidence intervals for simulation predictions of reservoir production given a stochastic geology.

Another research problem that needs vigorous attention is to formalize the propagation of validation metric information from the unit problem/benchmark tiers in the validation hierarchy to the subsystem/system tiers. We mentioned above that this problem is related to sensitivity analysis and propagation of uncertainty in the validation hierarchy. What makes the problem even harder is that second order uncertainty is present in the validation hierarchy. That is, quantitative specification of information-flow within the hierarchy is itself highly uncertain, at least for strongly-coupled-physics problems. While it may be the case that non-traditional uncertainty

quantification methods are required to address this problem, researching and applying probabilistic methodologies is certainly important to begin with. Recently, Hills et al.[89] have begun a study of the propagation of validation metrics and their associated uncertainty within a simple validation hierarchy using first-order-sensitivity methods and a maximum likelihood based validation metric. The problem of propagation of validation metric information in a validation hierarchy also has substantial formal similarity to the problem of system reliability methods. Recent work applying Bayes net methods originally developed for system reliability methods to understanding system validation metrics has just begun.[125]

Finally, we believe that improvements in validation metrics can be gained by adoption of some of the methods used in model calibration. Validation puts a premium on obtaining a measure to compare computational and experimental results, and relating the magnitude of this measure to the intended application. Calibration emphasizes methods to improve the agreement between the model and the experiment by optimization of model parameters. It is clear that metrics used in the first step of calibration, *i.e.*, before optimization is introduced, are directly pertinent to validation metrics discussed here. After a model is calibrated or updated, the same question remains - "How do we quantify the confidence in our calibrated model?" A recent review of some of these issues with a focus on calibration and Bayesian updating is found in Ref. [25]. Bayesian model updating requires a probabilistic validation metric for improvement of probabilistic parameters in the model. Hasselman et al,[79,80] has developed a different approach to compute probabilistic validation metrics for structural dynamics analyses. He uses principal components analysis and singular-value decomposition in the formation of the metric.

4. Major Implementation Issues

Many of the research issues just discussed will require mathematical methods development and careful thought and debate over a number of years to resolve. We believe there are also a number of implementation issues that are just as important for progress of V&V, but they are issues related to management and implementation. In this section we discuss: requirements for improved understanding of differing perspectives, new management initiatives, new "customer" initiatives for using commercial software, and the development of engineering standards for code verification. Without attention, several of these implementation issues may actually be more detrimental to the maturation of trust and confidence in computational simulation than the lack of research progress.

4.1 Management issues

In this section we address implementation issues related to the need for significantly improved code verification, the need for improved cooperation between computational analysts and systems engineers, the need for synergism between analysts and experimentalists, and the need for more quantitative V&V indicators. All of these issues present difficult management challenges to transform the status quo. We note that the issues discussed in this section predominantly relate to organizations that develop their own computational physics software or organizations that use commercial software in their activities.

The growing impact associated with computational physics software failure, the increased size and cost of code projects, and the need for interaction of large numbers of software developers greatly leverage the involvement of, and reliance upon, formal SQA practices in the development of

computational physics software. Resistance to the application of formal SQA practices in computational physics is created by the degree of formality, constraints, and seeming costs of SQA, as well as by psychological reasons (science versus product development, for example). As mentioned by Roache,[161] SQA issues are downright unattractive to small computational physics projects whose primary purpose is to develop software for exploration and scientific insight, rather than to minimize the risk of software failure at seemingly great cost.

Significantly, as the goals and expectations of computational physics evolve to have impact that goes far beyond scientific exploration and insight, the consequence of software failure greatly magnifies. In particular, inaccurate answers rather than obvious code failures are especially dreaded in high-consequence computing because it may be very difficult to determine that a calculation is inaccurate. SQA methodologies provide some additional means of addressing such a problem. But does the cost justify intensive application of such methodologies for a given computational-physics-code project? Part of the confidence-optimization problem for verification is the difficult question of how to measure the consequence of failure in a computational physics code. Consequence measures could involve, for example, potential loss of future business, possible liability costs (both human costs and property costs), magnitude of environmental damage, and national security impact. Such a measure is important in clearly understanding when the application of formal SQA methodologies provides unquestioned value in code verification activities.

Two deeply held and widespread difficulties are encountered when applying computational physics predictions to real systems-engineering decision making. The first difficulty concerns the radically different time scales for making engineering decisions and producing results from computational simulations; that is, engineering decisions need to be made rather quickly on design issues, and there is not much time to run many computational simulations. For example, it is not uncommon for the design engineer to tell the analyst “If you get the analysis to me by next week, I will use the results to make a design decision. If you don’t, I will make the decision without your help.” The second difficulty is related to the radically different value systems that are in play between analysts and systems engineers. In our opinion, computational analysts typically have an infatuation with physical detail, while systems engineers are concerned with high-level issues such as design optimization, manufacturing produceability, system robustness, and assessment of system safety and reliability. We believe that systems engineers are the ultimate customers of the code and should be the dominant player in their interaction with computational analysts. Systems engineers are responsible for designing and producing products, and the primary role of the codes is to help them make better decisions.

As we have discussed previously, a competitive or adversarial relationship between computationalists and experimentalists exists in many organizations. This type of relationship can be found at the individual level or between computational and experimental groups. It could be due to competition over organizational resources or could be the result of recognition of one group over the other. Management often doesn’t recognize the problem, or if they do, tend to subconsciously ignore it. And even if there are not competitive or adversarial pressures, there is commonly a strong difference in cultures between computationalists and experimentalists. For V&V to be successful, it is imperative that management assess the state of the relationship between computationalists and experimentalists in their organizations and create opportunities for bringing these different individuals and groups together in cooperative and mutually beneficial efforts. Management must make it clear that the success of the validation effort will benefit both groups equally, and that failure will be the responsibility of both groups.

The issue of how to succinctly communicate quantitative information concerning the status of

V&V activities on a code or computational result is an important issue. There is a practical need for quickly transmitting this information, for example, to potential customers for commercial software, customers of computational analyses, decision makers, and policy makers. This is particularly important when this information must be transmitted at a meeting or seminar-type presentation. Logan and Nitta[124] have suggested an intuitive and effective method for transmitting this information using graphical symbols of a dial-face meter for code verification, model validation, and even numerical error estimation. They have suggested using a zero to ten scale on each V&V meter to quantitatively indicate the level of completion, maturity, or accuracy of the activity. The mapping of the numerical scale to various levels of completion, by necessity, must be qualitative. However, the same qualitative mapping would be used for each computational result in an ensemble of results, or computational capability exercised in the code. In that way the observer could consistently see the relative maturity of different computations or capabilities. Achieving consistency, even on a relative and qualitative scale, is a significant implementation goal for V&V. It would be expected that there could be significant differences of opinion on the mapping between numerical and qualitative descriptors between, say, commercial software competitors. However, we believe Logan and Nitta's VER- and VAL-meters are an effective way to move V&V information toward a more quantitative basis and conspicuous stance when brevity is of the essence.

4.2 Practical applications in industrial settings

A critical issue in improving the implementation of V&V activities, especially validation, is to understand the particular needs of the customers of the computational models. From the perspective of a code developer, these customers may be analysts who use the models to produce results or they may be system engineers who use the results for decision making. Understandably, an analyst could also be a system engineer. In this section we point out some of the practical concerns that are important to customers in industry and discuss how several of these concerns are dealt with in the process of attempting to validate and calibrate the results of the codes.

As the use of codes, particularly commercial codes, in industry becomes more widespread, industrial customers are requiring that these codes be more responsive to their needs and demonstrate more reliability. The primary concern of the industrial code user is the level of confidence he has in the daily production use of the code, particularly in situations where the validation domain is poorly known (see Fig. 9). The accuracy and reliability demands of industrial customers vary greatly, however, based primarily on the level of physical detail in which the customers are interested. This spectrum of physical complexity can range, for example, from output quantities restricted to a limited number of global response measures, such as engine or compressor efficiency, drag and lift coefficients, global forces and moments, up to detailed physical characteristics, such as local temperature or plastic strain in selected regions, or to gradients of properties such as local heat flux or fluid dynamic turbulence quantities. Given this spectrum of possible physical complexity, the validation requirements of industrial customers will vary greatly.

An illustration of how the validation requirements for industrial customers can vary is taken from the ASME benchmark exercise of the transonic compressor rotor R37 from NASA.[45] This exercise was handled as a blind validation computation by a large number of codes. When the results were compiled at the conclusion of the exercise, the scatter of results on global compressor efficiency was on the order of $\pm 2\%$. This small level of scatter was welcomed in the CFD community, but the important observation from the exercise was that different computations

showing the same level of accuracy for computed efficiency were due to very different combinations of errors in numerics and physical modeling. Since the global compressor efficiency is obtained by a ratio involving average stagnation pressure and temperature, the same value of efficiency can be obtained by under predicting both pressure and temperature or by exactly the opposite—over predicting both quantities. In essence, different numerators and different denominators can produce the same acceptable global efficiency result. This kind of example raises serious and difficult questions for code developers regarding the importance of understanding and responding to a wide range of customer requirements. To some customers, the accuracy of the ratio would be sufficient. To other customers, the accuracy requirements would dictate that both numerator and denominator be correct, not just the ratio of the two quantities. We believe it is just as important to understand *why* a comparison between computation and experiment is good or bad, as it is to obtain good agreement between computation and experiment.

All predictions will inevitably include a wide range of uncertainties, given the complexity of the systems being modeled. Industrial customers must find ways to cope with these uncertainties, as opposed to the initial expectation of accurately quantifying them. Generally, industrial customers tend to use computational models in more pragmatic ways than do code developers and researchers.[95] The rigorous application of computational models by code developers and researchers is often a luxury not available to industrial customers, who are constrained by the time and costs of industry design procedures. These constraints have consequences on the level of requirements that industry can put on a simulation. For example, industrial simulations tend to limit grid density, grid quality, and iterative convergence bounds; and sometimes the industrial simulations do not estimate the modeling uncertainty by comparing the results from different physical models. Consequently, the goals of numerical error estimation and the validation recommendations discussed previously are clearly out of reach in many industrial environments. Instead, the dialog between code developers and their customers focuses on the critical issue of *managing the uncertainties*. This is a new concept in computational simulation, but it is analogous to risk management in engineering design projects. This concept requires significant research to define rational methodologies, quantitative evaluations, and techniques for integrating the management of uncertainties in the decision-making process.

One way industry is attempting to deal with the severe constraints of cost and schedule affecting numerical error estimation and model uncertainty estimation is to concentrate on the “deltas” from computations. That is, the analyst considers only the changes of output quantities from run to run that are due to, for example, changes in geometry or boundary conditions. The basic assumption in this approach is that, for a given level of combined errors, the balance of errors will not vary significantly when relatively small changes in geometrical or input parameters are performed during a design-cycle process. This assumption is reasonable and pragmatic, but there is no assurance that the deltas will be accurate.

To improve the reliability of this approach, industrial customers calibrate their computational results with experimental data that are available for their existing hardware products. For example, industrial customers typically have large quantities of proprietary experimental data for different response measures of their products. As part of their calibration procedure, industrial customers will choose the option for the mathematical model in the code that best reproduces a wide range of their experimental data. When this calibration procedure is used, one of the important requirements is that there be “no change” in computational results between old and new versions of the code for the same set of geometrical conditions, boundary conditions, grids, and physical models chosen. This no-change requirement allows an industrial code user to keep the same calibration of parameters for the way in which the code was used. This calibration procedure in no way should

be viewed as validation. It is a method of ensuring that all of the possible numerical error sources and all of the physical model uncertainties consistently sum to zero for a given set of response measures. The issue of summation and cancellation of errors and uncertainties was addressed in Ref. [143].

A recent benchmark exercise organized by AIAA and NASA to validate lift and drag predictions of a wing-body combination illustrates the concerns of industrial customers for absolute predictive capability in a most striking way.[88,121] This exercise, conducted as a workshop, was set up to assess the level of confidence of CFD codes for force and moment predictions for the DLR-F4 wing-body data set. Thirty-five different CFD calculations were presented using various grids, both structured and unstructured, including some grids recommended to minimize grid dependency on some of the test cases. Although most of the codes were among the most advanced research codes available, providing better results than did some of the commercial codes, the outcome clearly showed the current limitations of CFD. Statistical analysis indicated that the standard deviation of the CFD results, even after exclusion of the computational “outliers,” was still 21 drag counts. (Note that 1 drag count is equal to a drag coefficient of 0.0001.) In comparison, wind tunnel experimental data is believed to be accurate to 4 drag counts. The industry objective, however, is 1 drag count: on the order of 0.5% of the total drag of a modern transport aircraft. Citing from Ref. [121]:

More experience needs to be gained where CFD is used in conjunction with wind tunnel data on development projects that culminate in a flight vehicle. Then the methods can be “calibrated” to a known outcome. Note that experimental methods went through a similar process long ago. Wind tunnel testing is not regarded as “perfect”, but it is useful as an engineering tool because its advantages and limitations are well known. CFD needs to go through the same process.

4.3 Commercial software

This section addresses implementation issues in V&V that are relevant to commercial software companies, the commercial software packages they produce, and the customers with whom these companies interact.

Commercial software companies are subjected to severe quality and reliability requirements from their customers. In addition, many of the recommended activities discussed in this paper, for example, which focused attention on modeling and simulation requirements, cannot be practically carried out by the commercial code company because of the extremely wide range of applications required by its customers. In light of such demands, the following factors are important for these companies to consider as part of their business activities:

- 1) Both the software company and the customer must understand that it is impossible to cover all possible combinations of physical, boundary, and geometrical conditions during the V&V process. As pointed out repeatedly in this paper, a code cannot be validated for all possible options available in the code. Hence, both parties, software company and customer, must learn to cope with management of the uncertainties.
- 2) The software company must put in place and practice a high level of SQA for every stage of the software life cycle. These stages range from a new software module to intermediate releases, to simple bug-fixing patches. Each of these types of code releases will require

separate SQA procedures for its quality assessment.

- 3) A commercial software system is much more than a numerical solver of the physical models. The system also contains the global communication and interaction system between the user, the various modules of the software system, and the computer as accessed through the graphical user interface (GUI). The GUI must act in the following ways: (a) as the task manager, allowing the user to direct computer runs to different computer systems on the network and/or launching a series of runs with varying parameters; (b) as a preprocessing unit for input geometries provided by computer-aided design (CAD) software; (c) as an interface for grid generation tools; (d) as input data to establish all of the needed boundary conditions and initial conditions; and (e) as the mechanism to transfer data to the postprocessing software.
- 4) The complete software system has to be compiled and ported to a wide variety of computer platforms (both serial and parallel processor machines), which run under different operating systems, and must be verified for a range of compiler options.

The V&V procedures within a software company environment constitute an ongoing process. With every software release, the company must have automatic SQA test procedures that are set up under the supervision of a software quality group. These procedures must include automatic reporting features for the results of the testing. In particular, the activities related to fixing, tracking, and managing software bugs are a critical part of the global SQA procedures.

Commercial software companies have long recognized the importance of documentation for their products. This documentation has mainly described the use of the GUI, the capabilities of the code, grid generation capabilities, numerical algorithm techniques, and postprocessing capabilities. However, it is our view that V&V activities have been poorly documented, if at all. This claim, of course, is also valid for most software development projects in computational physics. It is necessary to document the content and results of all major V&V activities, even though this practice could present a risk to the software company in relation to its competitors. The old adage of quality engineering applies here as well: If it isn't documented, it didn't happen.

Documentation can be in traditional paper form, or in the form of a computer-based records management system. It is our belief that a computer-based system will provide the best vehicle for two reasons. First, a computer-based system can more easily be updated and accessed, and it can include much more detail. For example, documentation can be provided on how the code was used in the definition, design, and analysis of each validation experiment. Enough information should be included to allow reproduction of the described calculations by others. Additionally, a complete description of each experiment should be given, and all of the experimental data should be easily available for visual and quantitative presentation. Second, a properly designed computer-based system can easily be searched from a wide variety of perspectives, e.g., computational, experimental, and engineering project, and can point, or link, to other documented information associated with the V&V activities that intersects a wide range of perspectives. Internet access greatly facilitates these desirable search and link capabilities, both within an organization's computer network and between organizations that are widely spread geographically but that work together as part of a project team.

One important way that software companies can improve the reliability of their codes is to provide sufficient training and education for their customers. For new users of a commercial code,

the training should focus on basic operation of the code and use of the GUI interface. For more experienced users, the training should focus on techniques for sophisticated grid generation, on understanding the limitations of models in the code, and on approaches for dealing with the wide range of uncertainties in real engineering systems. It is our view that large organizations that develop their own codes for internal use sometimes neglect to adequately address user support issues such as training, bug fixing, and consulting on the use of their codes.

4.4 Development of Standards

The maturation of V&V depends on concerted efforts by individuals and groups throughout the world to develop standardized terminology, procedures and tools that can be used by the computational physics community. In this section we call for a further commitment to such efforts, recommend the need to clarify the present terminology on validation, address the contributions that standards organizations can make and are making in this process, and underscore the critical need for benchmark databases for both verification and validation. Current European efforts to construct validation databases are highlighted.

Concerning terminology, we have used the DoD/AIAA definition of validation to mean that validation does *not* include the concept of *adequacy of the computational result for the intended uses* of the model. We argue that this is what the words mean in the definition of validation given in Section 1.2. As we noted however, our interpretation is narrower than the interpretation that is widely accepted in the DoD community. The DoD community considers validation to be the process of determining the degree to which the model is adequate for the application of interest. We argue that there are two significant reasons not to use the meaning ascribed by the DoD community. First, if the DoD definition was used, then individual researchers or analysts who are comparing computational and experimental results *cannot* claim they are validating a model. The *only* individuals, or group of individuals, that can do this are those who understand the intended application for the modeling and simulation capability and can state that the level of agreement between computation and experiment is adequate for the intended application. Second, the adequacy requirements for the intended application are typically ill defined. Adequacy requirements are ill defined for practical reasons; for example, adequacy may be defined fairly well at the complete system tier but very poorly for lower tiers of the validation hierarchy. Also, adequacy requirements are ill defined for technical reasons; for example, the analysis of complex systems involves the interaction of many different types of computational models, each with its own adequacy. The computational adequacy for some high-level system response measure is actually a combination of many adequacy measures from each of the contributing models. As a result, one must deal with adequacy requirements at a system level in an iterative and coupled manner by propagating uncertainty through many different types of physics models and multiple tiers of hierarchical complexity. Whether or not one accepts the arguments given here for our interpretation, the issue of interpretation of the meaning of validation needs further discussion and debate by the DoD community, the DOE community, engineering standards-writing committees, and other stakeholders relying on modeling and simulation.

We believe there is an urgent need to develop industry-wide standards for V&V procedures and tools that are applicable to the computational physics community. In our view, the most appropriate organizations for defining these standards are professional engineering societies that have officially recognized standards-writing committees. A complementary approach, one that is appropriate at the worldwide level, is the involvement of the International Organization for Standardization (ISO). We believe there is also an important contribution to be made by national

standards organizations, such as the U.S. National Institute of Standards and Technology (NIST), and by similar organizations in the Europe Union, such as the European Research Community on Flow, Turbulence And Combustion (ERCOFTAC)[60] and the National Agency for Finite Element Methods and Standards (NAFEMS).[131]

We believe standards organizations are best able to contribute in two areas of V&V: codification of recommended procedures and tools. As discussed earlier in this paper, the AIAA CFD Committee on Standards contributed the first standards document on V&V in 1998.[3] The document provided a foundation on which to build, but it was primarily focused on background, terminology, and fundamental concepts. New standards documents are needed that go beyond the AIAA Guide in terms of detailed procedures and tools, and that address technical disciplines other than fluid dynamics. For example, a new standards committee was formed in the American Society of Mechanical Engineers (ASME): the Committee on Verification and Validation in Computational Solid Mechanics (Performance Test Codes No. 60). In Europe, a particularly effective activity for the development of V&V procedures has been conducted by ERCOFTAC. ERCOFTAC privately funded an initiative on “Quality and Trust in Industrial CFD,” which led to the publication of the comprehensive report by Casey and Wintergerste.[27] This document, although not a standards document, gives general as well as specific practical advice on modeling and numerical techniques for CFD.

Concerning V&V tools, we contend that the key tools that need to be constructed are separate benchmark databases for verification and validation. Verification databases should be composed of all three types of highly accurate solutions, which we will simply refer to as benchmark solutions. Strong-sense benchmarks are test problems that are precisely defined and formally documented, typically by professional societies, academic institutions, or nonprofit organizations. It is our view that a strong-sense benchmark is described by the following four factors:

- 1) An exact, standardized, frozen, and promulgated definition of the benchmark.
- 2) An exact, standardized, and promulgated statement of the purpose of the benchmark. This statement addresses the benchmark’s role and application in a comprehensive test plan for a code, for example.
- 3) Exact, standardized, frozen, and promulgated requirements for comparison of codes with the benchmark’s results.
- 4) An exact, standardized, frozen, and promulgated definition of acceptance criteria for comparison of codes with the benchmark’s results. The criteria can be phrased either in terms of success or failure.

Do strong-sense benchmarks exist? We believe that the answer is “certainly not in computational physics.” We believe that standards organizations are the most appropriate organizations to compile and codify strong-sense benchmarks. We note that none of these groups have dealt with the most sensitive issue in our definition of strong-sense benchmarks, namely, the issue of establishing stringent standards for comparison with benchmarks and measuring success or failure. It is very clear to us that establishing standards will be a long and difficult process. The realities of competition among commercially available computational physics codes, competition among companies that rely on computational physics codes, and competition among countries may stifle this endeavor. However, we believe that the maturation of computational physics will suffer

if strong-sense benchmarks are not developed.

Validation benchmark databases are also critically needed. In the near term, these databases should be populated using existing high-quality experimental data that meet most, if not all, of the requirements for validation experiments specified earlier in this paper. The construction of experimental databases is in development both in the United States and in Europe. In the United States, for example, the National Project for Applications-oriented Research in CFD (NPARC) initiative has concentrated on validation of internal reacting flows. In Europe, the FLOWNET project is constructing databases for a wide range of industrial applications in support of validation. Closely related to FLOWNET is the QNET-CFD Thematic Network.[157] This network has more than 40 participants from several countries who represent research establishments and many sectors of the industry, including commercial CFD software companies. The QNET-CFD Thematic Network is funded by the European Commission. The objective of the network is to establish a list of so-called Application Challenges, or ACs, and Underlying Flow Regime Challenges. The Application Challenges are specific applications that represent simulations at the subsystem tier shown in Fig. 10. The Underlying Flow Regime Challenges represent simulations similar to the benchmark tier shown in Fig. 10.

The QNET-CFD Thematic Network has completed its second year of activity, with the participating organizations contributing 53 Application Challenges divided over six Thematic Areas aligned with the following industrial sectors: External Aerodynamics (10 ACs); Combustion & Heat Transfer (7 ACs); Chemical & Process, Thermal Hydraulics, & Nuclear Safety (11 ACs); Civil Construction (6 ACs); Environment (7 ACs); and Turbomachinery Internal Flows (12 ACs). The main requirement for each Application Challenge is that the experimental results become available so that CFD validation can be conducted. Even though the QNET-CFD Thematic Network is relatively new, it is our view that it is an exemplary initiative in the construction of validation databases. We believe that industry, government, and academia in the United States should promptly begin to formulate a similar activity.

In the long term, new validation experiments should be funded, and these results should then be entered in the validation database. We believe that identification of the validation experiments should be the responsibility of the application community. Of course, there would probably be little agreement in the application community on the specific experiments to be conducted. Funding for high-priority experiments could be obtained from research establishments, governmental institutions, and even joint ventures between private industrial organizations. The organizational role and facilitation of discussions regarding which experiments should be conducted would be best served by standards organizations. The validation databases, constructed in both the near term and the long term, could either be completely open to the public or accessible only to member organizations.

5. Closing Remarks

Implementation of most of the approaches and procedures recommended here, for V&V computations and experiments, will be neither inexpensive nor easy. Furthermore, some of these approaches may even be technically or economically impractical in particular situations. In addition, some of the approaches and procedures have not been developed satisfactorily for implementation in an engineering environment. With each included step, however, the quality of the code V&V processes will be improved. We firmly believe that V&V is a process, *not* a product. We also believe that an effective path forward to improving the V&V processes is to contrast processes that are ideal, possible, and ineffective. The following gives examples of these

as they relate to various V&V activities, as viewed from the perspective of developing modeling and simulation capability:

1. **Ideal:** The application domain for the modeling and simulation capability is well understood and carefully defined, and the accuracy requirements are known.
Possible: Portions of the application domain are understood and some of the accuracy requirements are known.
Ineffective: The modeling and simulation capability is expected to accurately compute results for every advertised option in the code.
2. **Ideal:** The validation tier hierarchy, analogous to the validation pyramid discussed, has been carefully constructed using a systems engineering approach.
Possible: The validation hierarchy has been constructed for individual subsystems or components, but the interactions between subsystems have not been studied.
Ineffective: No validation hierarchy has been constructed and only expedient validation experiments have been identified.
3. **Ideal:** The PIRT has been used to identify and rank the importance of physical processes and interactions of processes for all tiers and faces of the validation pyramid.
Possible: The PIRT has been used to identify and rank processes on one face of the validation pyramid.
Ineffective: No PIRT was constructed and only well exercised physical-process options in the code are considered for V&V activities.
4. **Ideal:** Using the results of the PIRT, code verification activities have been defined and prioritized, schedules set, and needed resources allocated.
Possible: Using the results of the PIRT, some of the code verification activities are defined, but the code development resources or schedule dictate which activities will be conducted.
Ineffective: No PIRT was constructed and code verification activities are hap-hazard and only conducted if resources and time become available.
5. **Ideal:** SQA procedures have been defined, implemented, consistently adhered to, and documented.
Possible: SQA procedures are poorly defined and partially implemented, only adhered to just before a new code release, and informally documented.
Ineffective: SQA procedures have not been considered or implemented, and are only mentioned if the code customer asks about them.
6. **Ideal:** Using the results of the PIRT, model validation activities, both computational and experimental, have been defined and prioritized, schedules set, and needed resources allocated.
Possible: The PIRT is used to define and prioritize some model validation activities, but code-application schedules and limited resources require elimination of many validation activities.
Ineffective: No PIRT was constructed and model validation will only be conducted using existing experimental data.

7. **Ideal:** The code is used to define, design, and analyze all of the validation experiments.
Possible: The code is applied to define and design some of the validation experiments, but many of the code results are only available after the validation experiment is completed.
Ineffective: The validation experiments are conducted completely separate from any input from the code development team and code results.
8. **Ideal:** Highly converged numerical solutions, in both space and time, have been obtained for every validation experiment conducted.
Possible: Quantitative numerical error estimates are computed for most validation experiments, but some complex-system experiments only have qualitative error estimates.
Ineffective: No quantitative numerical error estimates are available and only single-grid, single-time-step solutions are available.
9. **Ideal:** Validation-experiment data are well characterized, and a large number of experimental realizations are available for estimating random and bias errors.
Possible: The experimental data are moderately characterized, and only two experimental realizations are obtained.
Ineffective: Only one experimental realization is available and no estimates are available for random or bias errors in needed computational quantities or in measured experimental results.
10. **Ideal:** Validation metrics are carefully defined and requirements stated, and clearly connected to modeling and simulation requirements for the application domain.
Possible: Validation metrics are only defined and requirements stated for high-level system responses, and not defined or requirements stated for low-level (*i.e.*, unit tier) physical response.
Ineffective: Validation metrics are not defined and only qualitative assessment is to be used for low-level physical responses.
11. **Ideal:** Using statistical data for both computational and experimental results for the validation metrics, one is able to precisely quantify the accuracy of the model for all validation experiments.
Possible: Statistical data are available from the experimental results but computational results are limited to deterministic metrics because of highly-constrained budgets and schedules.
Ineffective: Validation metrics are not defined and the “view-graph norm” is used for assessment.
12. **Ideal:** Lessons learned from all of the validation activities are clearly formulated, documented, and broadly communicated.
Possible: Lessons learned from some validation experiments are documented, but only those where the validation experiment and the code accuracy was considered a success.
Ineffective: No time or resources are available for determining or documenting lessons learned, and lessons that are learned, are verbally and randomly communicated.

Some will argue that the costs of V&V activities can exceed their value added. In some instances they might. However, we remind these skeptics that the costs of V&V must be weighed against the costs of incorrect or improper decisions made based on computational modeling and simulation. Analogous to probabilistic risk assessment activities, the risk is typically defined as the product of the probability of the occurrence of the event and the consequence of the event. For example, suppose an erroneous conclusion is made on the physics of some process described in an article in a research journal. The erroneous result would rightly be viewed as a low consequence risk. On the other hand, if erroneous conclusions based on modeling and simulation are made on high-consequence events, decision makers could place their constituency at extreme risk. This is exceedingly more true for systems that cannot be tested. For these situations, the only path forward is to *drastically improve* our confidence and understanding in computational simulations.

Acknowledgements

The authors sincerely thank Dean Dobranich, Robert Paulsen, and Marty Pilch of Sandia National Laboratories for reviewing the manuscript and providing many helpful suggestions for improvement of the manuscript. The first author thanks Robert Thomas of Sandia National Labs. for his generous support to complete this work. We also thank Rhonda Reinert of Technically Write Inc. for providing extensive editorial assistance during the writing of the manuscript. This work was performed at Sandia National Laboratories, which is operated by Lockheed Martin Corp. for the U. S. Department of Energy under contract No. DE-AC04-94AL85000.

Acronyms and Abbreviations

AC	Application Challenges
AIAA	American Institute of Aeronautics and Astronautics
ASME	American Society of Mechanical Engineers
ASCI	Accelerated Strategic Computing Initiative
CFD	computational fluid dynamics
DMSO	Defense Modeling and Simulation Office
DoD	Department of Defense
DOE	Department of Energy
ERCOTAC	European Research Community On Flow, Turbulence, And Combustion
GCI	Grid Convergence Index
GNC	guidance, navigation and control
IEEE	Institute of Electrical and Electronics Engineers
MMS	Method of Manufactured Solutions
NAFEMS	National Agency for Finite Element Methods and Standards
NIST	National Institute of Standards and Technology
NPARC	National Project for Applications-oriented Research in CFD
PDE	partial differential equation
PIRT	Phenomena Identification and Ranking Table
Q-NET-CFD	Thematic Network on Quality and Trust for the Industrial Applications of CFD
SCS	Society for Computer Simulation
SQA	software quality assurance

SQE	software quality engineering
ODE	ordinary differential equation
OR	operations research
V&V	verification and validation

References

- [1] Aeschliman, D. P., and Oberkampf, W. L. Experimental Methodology for Computational Fluid Dynamics Code Validation, *AIAA Journal*, Vol. 36, No. 5, 1998; 733-741.
- [2] Aeschliman, D. P., Oberkampf, W. L., and Blottner, F. G. A Proposed Methodology for CFD Code Verification, Calibration, and Validation, Paper 95-CH3482-7, *16th International Congress on Instrumentation for Aerospace Simulation Facilities*, Dayton, OH, 1995.
- [3] AIAA. Guide for the Verification and Validation of Computational Fluid Dynamics Simulations, American Institute of Aeronautics and Astronautics, AIAA-G-077-1998, Reston, VA, 1998.
- [4] Ainsworth, M., and Oden, J. T. *A Posteriori Error Estimation in Finite Element Analysis*, John Wiley, New York, 2000.
- [5] Alapaty, K., Raman, S., and Niyogi, D. S. Uncertainty in the Specification of Surface Characteristics: A Study of Prediction Errors in the Boundary Layer, *Boundary-Layer Meteorology*, Vol. 82, No. 3, 1997; 473-500.
- [6] Almond, R. G. *Graphical Belief Modeling*, 1st ed., Chapman & Hall, London, 1995.
- [7] ANS. Guidelines for the Verification and Validation of Scientific and Engineering Computer Programs for the Nuclear Industry, American Nuclear Society, ANSI/ANS-10.4-1987, 1987.
- [8] Axelsson, O. *Iterative Solution Methods*, Cambridge University Press, Cambridge, U.K., 1996.
- [9] Baber, R. *The Spine of Software; Designing Provably Correct Software: Theory and Practice*, John Wiley, New York, 1987.
- [10] Babuska, I., Ihlenburg, F., Strouboulis, T., and Gangaraj, S. K. A Posteriori Error Estimation for Finite Element Solutions of Helmholtz' Equation - Part II: Estimation of the Pollution Error, *International Journal of Numerical Methods in Engineering*, Vol. 40, 1997; 3883-3900.
- [11] Babuska, I., and Oh, H.-S. Pollution Problem of the p- and h-p Versions of the Finite Element Method, *Communications in Applied Numerical Methods*, Vol. 3, 1987; 553-561.
- [12] Babuska, I., and Strouboulis, T. *The Finite Element Method and its Reliability*, Oxford University Press, Oxford, UK, 2001.
- [13] Babuska, I., Strouboulis, T., Upadhyay, C. S., and Gangaraj, S. K. A Posteriori Estimation and Adaptive Control of the Pollution Error in the h-Version of the Finite Element Method, *International Journal of Numerical Methods in Engineering*, Vol. 38, 1995; 4207-4235.
- [14] Balci, O. Principles of Simulation of Model Validation, Verification, and Testing, *Transactions of the Society for Computer Simulation International*, Vol. 14, 1997; 3-12.
- [15] Beck, M. B. Water Quality Modeling: A Review of the Analysis of Uncertainty, *Water Resources Research*, Vol. 23, No. 8, 1987; 1393-1442.
- [16] Beizer, B. *Software Testing Techniques*, Van Nostrand Reinhold, New York, 1990.
- [17] Blottner, F. G. Influence of Boundary Approximations and Conditions on Finite-Difference Solutions, *Journal of Computational Physics*, Vol. 48, No. 2, 1982; 246-269.

- [18] Bossel, H. *Modeling and Simulation*, 1st ed., A. K. Peters, Ltd., Wellesley, MA, 1994.
- [19] Botella, O., and Peyret, R. Computing Singular Solutions of the Navier-Stokes Equations with the Chebyshev-Collocation Method, *International Journal for Numerical Methods in Fluids*, Vol. 36, No. 2, 2001; 125-163.
- [20] Bowen, J. P., and Hinchey, M. G. 10-Commandments of Formal Methods, *Computer*, Vol. 28, No. 4, 1995; 56-63.
- [21] Bowen, J. P., and Hinchey, M. G. *Applications of Formal Methods*, Prentice-Hall, Englewood Cliffs, NJ, 1995.
- [22] Boyack, B. E., Catton, I., Duffey, R. B., Griffith, P., Katsma, K. R., Lellouche, G. S., Levy, S., Rohatgi, U. S., Wilson, G. E., Wulff, W., and Zuber, N. Quantifying Reactor Safety Margins Part 1: An Overview of the Code Scaling, Applicability, and Uncertainty Evaluation Methodology, *Nuclear Engineering and Design*, Vol. 119, 1990; 1-15.
- [23] Brandt, S. *Data Analysis: Statistical and Computational Methods for Scientists and Engineers*, Springer, New York, 1998.
- [24] Cadafalch, J., Perez-Segarra, C. C., Consul, R., and Oliva, A. Verification of Finite Volume Computations on Steady State Fluid Flow and Heat Transfer, *Journal of Fluids Engineering*, Vol. 124, No. 1, 2002; 11-21.
- [25] Campbell, K. A Brief Survey of Statistical Model Calibration Ideas, Los Alamos National Laboratory, LA-UR-02-3157, Los Alamos, NM, 2002.
- [26] Carpenter, M. H., and Casper, J. H. Accuracy of Shock Capturing in Two Spatial Dimensions, *AIAA Journal*, Vol. 37, No. 9, 1999; 1072-1079.
- [27] Casey, M., and Wintergerste, T. eds. *ERCRAFTAC Special Interest Group on Quality and Trust in Industrial CFD: Best Practices Guidelines*, European Research Community on Flow, Turbulence and Combustion, 2000.
- [28] Chen, C.-F. A., Lotz, R. D., and Thompson, B. E. Assessment of Numerical Uncertainty Around Shocks and Corners on Blunt Trailing-Edge Supercritical Airfoils, *Computers and Fluids*, Vol. 31, No. 1, 2002; 25-40.
- [29] Chiles, J.-P., and Delfiner, P. *Geostatistics: Modeling Spatial Uncertainty*, John Wiley, New York, 1999.
- [30] Chlond, A., and Wolkau, A. Large-Eddy Simulation of a Nocturnal Stratocumulus-Topped Marine Atmospheric Boundary Layer: An Uncertainty Analysis, *Boundary-Layer Meteorology*, Vol. 95, No. 1, 2000; 31-55.
- [31] Chorin, A. J., Kast, A. P., and Kupferman, R. On the Prediction of Large-Scale Dynamics Using Unresolved Computations, Lawrence Berkeley National Laboratory, LBNL-42283, Berkeley, CA, 1998.
- [32] Chorin, A. J., Kast, A. P., and Kupferman, R. Optimal Prediction of Underresolved Dynamics, *Proceedings of the National Academy of Sciences*, Vol. 95, 1998; 4094-4098.
- [33] Chorin, A. J., Kast, A. P., and Kupferman, R. Unresolved Computation and Optimal Prediction, *Communications in Pure and Applied Mathematics*, Vol. 52, 1999; 1231-1254.
- [34] Clark, E. M., Jr., Grumberg, O., and Peled, D. *Model Checking*, MIT Press, Cambridge, MA, 1999.
- [35] Cohen, M. L., Rolph, J. E., and Steffey, D. L. eds. *Statistics, Testing, and Defense Acquisition: New Approaches and Methodological Improvements*, National Academy Press, Washington, DC, 1998.
- [36] Coleman, H. W., and Steele, W. G., Jr. *Experimentation and Uncertainty Analysis for Engineers*, 2nd ed., John Wiley, New York, 1999.
- [37] Coleman, H. W., and Stern, F. Uncertainties and CFD Code Validation, *Journal of Fluids*

- Engineering*, Vol. 119, 1997; 795-803.
- [38] Cosner, R. R. CFD Validation Requirements for Technology Transition, AIAA Paper No. 95-2227, *26th AIAA Fluid Dynamics Conference*, San Diego, CA, 1995.
 - [39] Cox, D. R. *Planning of Experiments*, (republished in 1992) ed., John Wiley, New York, 1958.
 - [40] Cullen, A. C., and Frey, H. C. *Probabilistic Techniques in Exposure Assessment: A Handbook for Dealing with Variability and Uncertainty in Models and Inputs*, Plenum Press, New York, 1999.
 - [41] Dahl, O. *Verifiable Programming*, Prentice-Hall, Englewood Cliffs, NJ, 1992.
 - [42] Dean, A., and Voss, D. *Design and Analysis of Experiments*, Springer-Verlag, New York, 1999.
 - [43] DeMillo, R. A., McCracken, W. M., Martin, R. J., and Passafiume, J. F. *Software Testing and Evaluation*, Benjamin/Cummings, Menlo Park, CA, 1987.
 - [44] Denker, M., and Woyczynski, W. A. *Introductory Statistics and Random Phenomena*, Birkhauser, Boston, 1998.
 - [45] Denton, J. D. Lessons from Rotor 37, *Proceedings of the 3rd International Symposium on Experimental and Computational Aerothermodynamics of Internal Flows*, Beijing, China, 1996.
 - [46] DeVolder, B., Glimm, J., Grove, J. W., Kang, Y., Lee, Y., Pao, K., Sharp, D. H., and Ye, K. Uncertainty Quantification for Multiscale Simulations, *Journal of Fluids Engineering*, Vol. 124, No. 1, 2002; 29-41.
 - [47] Diskin, B., and Thomas, J. L. Analysis of Boundary Conditions for Factorizable Discretizations of the Euler Equations, NASA/ICASE, NASA/CR-2002-211648, Hampton, VA, 2002.
 - [48] DoD. *DoD Directive No. 5000.59: Modeling and Simulation (M&S) Management*, Defense Modeling and Simulation Office, Office of the Director of Defense Research and Engineering, www.dmsomil/docslib.
 - [49] DoD. *DoD Instruction 5000.61: Modeling and Simulation (M&S) Verification, Validation, and Accreditation (VV&A)*, Defense Modeling and Simulation Office, Office of the Director of Defense Research and Engineering, www.dmsomil/docslib.
 - [50] DoD. *Verification, Validation, and Accreditation (VV&A) Recommended Practices Guide*, Defense Modeling and Simulation Office, Office of the Director of Defense Research and Engineering, www.dmsomil/docslib.
 - [51] Dowding, K. Quantitative Validation of Mathematical Models, *ASME International Mechanical Engineering Congress Exposition*, New York, 2001.
 - [52] Draper, D. Assessment and Propagation of Model Uncertainty, *Journal of the Royal Statistical Society B*, Vol. 57, No. 1, 1995; 45-97.
 - [53] Du, J., Mullen, S. L., and Sanders, F. Short-Range Ensemble Forecasting of Quantitative Precipitation, *Monthly Weather Review*, Vol. 125, No. 10, 1997; 2427-2459.
 - [54] Dubois, D., and Prade, H. eds. *Fundamentals of Fuzzy Sets*, Kluwer Academic Publishers, Boston, MA, 2000.
 - [55] Dyer, M. *The Cleanroom Approach to Quality Software Development*, John Wiley, New York, 1992.
 - [56] Earman, J. *Bayes or Bust?*, MIT Press, Cambridge, MA, 1992.
 - [57] Easterling, R. G. Measuring the Predictive Capability of Computational Models: Principles and Methods, Issues and Illustrations, Sandia National Laboratories, SAND2001-0243,

- Albuquerque, NM, 2001.
- [58] Easterling, R. G. Quantifying the Uncertainty of Computational Predictions, Sandia National Laboratories, SAND2001-0919C, Albuquerque, NM, 2001.
- [59] Eca, L., and Hoekstra, M. An Evaluation of Verification Procedures for CFD Applications, *Proceedings of the 24th Symposium on Naval Hydrodynamics*, Fukuoka, Japan, 2002.
- [60] ERCOFTAC. *Portal to Fluid Dynamics Database Resources*, European Research Community on Flow, Turbulence and Combustion, <http://ercoftac.mech.surrey.ac.uk>.
- [61] FAA. System Design Analysis, Federal Aviation Administration, Advisory Circular 25.1309-1A, Washington, DC, 1988.
- [62] Fairley, R. E. *Software Engineering Concepts*, McGraw-Hill, New York, 1985.
- [63] Ferson, S., and Ginzburg, L. R. Different Methods are Needed to Propagate Ignorance and Variability, *Reliability Engineering and System Safety*, Vol. 54, 1996; 133-144.
- [64] Ferziger, J. H., and Peric, M. *Computational Methods for Fluid Dynamics*, Springer-Verlag, New York, 1996.
- [65] Frank, M. V. Treatment of Uncertainties in Space Nuclear Risk Assessment with Examples from Cassini Mission Applications, *Reliability Engineering and System Safety*, Vol. 66, 1999; 203-221.
- [66] French, S., and Smith, J. Q. eds. *The Practice of Bayesian Analysis*, Hodder Arnold, London, 1997.
- [67] Gamerman, D. *Markov Chain Monte Carlo*, Chapman & Hall, London, 1997.
- [68] Gass, S. I. Model Accreditation: A Rationale and Process for Determining a Numerical Rating, *European Journal of Operational Research*, Vol. 66, 1993; 250-258.
- [69] Gelman, A. B., Carlin, J. S., Stern, H. S., and Rubin, D. B. *Bayesian Data Analysis*, Chapman & Hall, London, 1995.
- [70] Glimm, J., Hou, S., Kim, H., Sharp, D. H., and Ye, K. A Probability Model for Errors in the Numerical Solutions of a Partial Differential Equation, Los Alamos National Laboratory, LAUR-99-5352, Los Alamos, NM, 1999.
- [71] Glimm, J., Hou, S. L., Kim, H. J., Lee, Y., Sharp, D. H., Ye, K., and Zou, W. S. Risk Management for Petroleum Reservoir Production: A Simulation-Based Study of Prediction, *Computational Geosciences*, Vol. 5, No. 3, 2001; 173-197.
- [72] Glimm, J., and Sharp, D. Stochastic Methods for the Prediction of Complex Multiscale Phenomena, *Quarterly of Applied Mathematics*, Vol. 56, No. 4, 1998; 741-765.
- [73] Glimm, J., and Sharp, D. H. Stochastic Methods for the Prediction of Complex Multiscale Phenomena, Los Alamos National Laboratory, LAUR-97-3748, Los Alamos, NM, 1997.
- [74] Gunter, B. H. How Statistical Design Concepts Can Improve Experimentation in the Physical Sciences, *Computers in Physics*, Vol. 7, No. 3, 1993; 262-272.
- [75] Gustafson, J. Computational Verifiability and Feasibility of the ASCI Program, *IEEE Computational Science & Engineering*, Vol. 5, No. 1, 1998; 36-45.
- [76] Haimes, Y. Y. *Risk Modeling, Assessment, and Management*, John Wiley, New York, 1998.
- [77] Hamill, T. M., and Wilks, D. S. A Probabilistic Forecast Contest and the Difficulty in Assessing Short-Range Forecast Uncertainty, *Weather and Forecasting*, Vol. 10, No. 3, 1995; 620-631.
- [78] Hanson, K. M. A Framework for Assessing Uncertainties in Simulation Predictions, *Physica D*, Vol. 133, 1999; 179-188.
- [79] Hasselman, T. K. Quantification of Uncertainty in Structural Dynamic Models, *Journal of Aerospace Engineering*, Vol. 14, No. 4, 2001; 158-165.

- [80] Hasselman, T. K., Wathugala, G. W., and Crawford, J. A Hierarchical Approach for Model Validation and Uncertainty Quantification, <http://wccm.tuwien.ac.at>, *Fifth World Congress on Computational Mechanics*, Vienna, Austria, 2002.
- [81] Hatton, L. The T Experiments: Errors in Scientific Software, *IEEE Computational Science & Engineering*, Vol. 4, No. 2, 1997; 27-38.
- [82] Haworth, D. C., Tahry, S. H. E., and Huebler, M. S. A Global Approach to Error Estimation and Physical Diagnostics in Multidimensional Computational Fluid Dynamics, *International Journal for Numerical Methods in Fluids*, Vol. 17, No. 1, 1993; 75-97.
- [83] Helton, J. C. Uncertainty and Sensitivity Analysis Techniques for use in Performance Assessment for Radioactive Waste Disposal, *Reliability Engineering and System Safety*, Vol. 42, No. 2-3, 1993; 327-367.
- [84] Helton, J. C. Uncertainty and Sensitivity Analysis in the Presence of Stochastic and Subjective Uncertainty, *Journal of Statistical Computation and Simulation*, Vol. 57, 1997; 3-76.
- [85] Helton, J. C. Uncertainty and Sensitivity Analysis in Performance Assessment for the Waste Isolation Pilot Plant, *Computer Physics Communications*, Vol. 117, No. 1-2, 1999; 156-180.
- [86] Helton, J. C., Anderson, D. R., Baker, B. L., Bean, J. E., Berglund, J. W., Beyeler, W., Economy, K., Garner, J. W., Hora, S. C., Iuzzolino, H. J., Knupp, P., Marietta, M. G., Rath, J., Rechar, R. P., Roache, P. J., *et al.* Uncertainty and Sensitivity Analysis Results Obtained in the 1992 Performance Assessment for the Waste Isolation Pilot Plant, *Reliability Engineering and System Safety*, Vol. 51, No. 1, 1996; 53-100.
- [87] Helton, J. C., Anderson, D. R., Basabilvazo, G., Jow, H.-N., and Marietta, M. G. Conceptual Structure of the 1996 Performance Assessment for the Waste Isolation Pilot Plant, *Reliability Engineering and System Safety*, Vol. 69, No. 1-3, 2000; 151-165.
- [88] Hensch, M. Statistical Analysis of CFD Solutions from the Drag Prediction Workshop, AIAA-2002-0842, *40th AIAA Aerospace Sciences Meeting and Exhibit*, Reno, NV, 2002.
- [89] Hills, R. G., Leslie, I. H., and Trucano, T. G. Statistical Validation of Engineering and Scientific Models: Validation Experiments to Application, Sandia National Laboratories, (to appear), Albuquerque, NM, 2002.
- [90] Hills, R. G., and Trucano, T. G. Statistical Validation of Engineering and Scientific Models: Background, Sandia National Laboratories, SAND99-1256, Albuquerque, NM, 1999.
- [91] Hills, R. G., and Trucano, T. G. Statistical Validation of Engineering and Scientific Models: A Maximum Likelihood Based Metric, Sandia National Laboratories, SAND2001-1783, Albuquerque, NM, 2001.
- [92] Hirsch, C. *Numerical Computation of Internal and External Flows: Vol. 1: Fundamentals of Numerical Discretization*, Vol. 1, John Wiley, New York, 1988.
- [93] Hirsch, C. *Numerical Computation of Internal and External Flows: Vol. 2: Computational Methods for Inviscid and Viscous Flows*, Vol. 2, John Wiley, New York, 1990.
- [94] Hora, S. C. Aleatory and epistemic uncertainty in probability elicitation with an example from hazardous waste management, *Reliability Engineering and System Safety*, Vol. 54, 1996; 217-223.
- [95] Hutton, A. G., and Casey, M. V. Quality and Trust in Industrial CFD - A European Perspective, AIAA2001-0656, *39th AIAA Aerospace Sciences Meeting*, Reno, NV, 2001.
- [96] IEEE. IEEE Standard Dictionary of Electrical and Electronics Terms, ANSI/IEEE Std 100-1984, New York, 1984.
- [97] IEEE. IEEE Standard Glossary of Software Engineering Terminology, IEEE Std 610.12-

- 1990, New York, 1991.
- [98] ISO. ISO 9000-3: Quality Management and Quality Assurance Standards - Part 3: Guidelines for the Application of ISO 9001 to the Development, Supply and Maintenance of Software, International Standards Organization, Geneva, Switzerland, 1991.
 - [99] Jameson, A., and Martinelli, L. Mesh Refinement and Modeling Errors in Flow Simulation, *AIAA Journal*, Vol. 36, No. 5, 1998; 676-686.
 - [100] Johnson, D. M. A Review of Fault Management Techniques Used in Safety-Critical Avionic Systems, *Progress in Aerospace Sciences*, Vol. 32, No. 5, 1996; 415-431.
 - [101] Jones, C. *Applied Software Measurement*, 2nd ed., McGraw-Hill, New York, 1997.
 - [102] Kafka, P. Important Issues Using PSA Technology for Design of New Systems and Plants, *Reliability Engineering and System Safety*, Vol. 45, No. 1-2, 1994; 205-213.
 - [103] Kaner, C., Falk, J., and Nguyen, H. Q. *Testing Computer Software*, 2nd ed., John Wiley, New York, 1999.
 - [104] Keller, H. B. Accurate Difference Methods for Linear Ordinary Differential Systems Subject to Linear Constraints, *SIAM Journal on Numerical Analysis*, Vol. 6, 1969; 8-30.
 - [105] Kleijnen, J. P. C. *Statistical Tools for Simulation Practitioners*, 1st ed., Marcel Dekker, Inc., New York, 1987.
 - [106] Kleijnen, J. P. C. Case-Study: Statistical Validation of Simulation Models, *European Journal of Operational Research*, Vol. 87, 1995; 21-34.
 - [107] Kleijnen, J. P. C. Verification and Validation of Simulation Models, *European Journal of Operational Research*, Vol. 82, 1995; 145-162.
 - [108] Kleindorfer, G. B., O'Neill, L., and Ganeshan, R. Validation in Simulation: Various Positions in the Philosophy of Science, *Management Science*, Vol. 44, No. 8, 1998; 1087-1099.
 - [109] Klir, G. J., and Wierman, M. J. *Uncertainty-Based Information: Elements of Generalized Information Theory*, Vol. 15, Physica-Verlag, Heidelberg, 1998.
 - [110] Knepell, P. L., and Arangno, D. C. *Simulation Validation: A Confidence Assessment Methodology*, 1st ed., IEEE Computer Society Press, Washington, DC, 1993.
 - [111] Kohlas, J., and Monney, P.-A. *A Mathematical Theory of Hints - An Approach to the Dempster-Shafer Theory of Evidence*, Springer, Berlin, 1995.
 - [112] Kramosil, I. *Probabilistic Analysis of Belief Functions*, Kluwer, New York, 2001.
 - [113] Kumamoto, H., and Henley, E. J. *Probabilistic Risk Assessment and Management for Engineers and Scientists*, 2nd ed., IEEE Press, New York, 1996.
 - [114] Kurshan, R. P. Program Verification, *Notices of the American Mathematical Society*, Vol. 47, No. 5, 2000; 534-545.
 - [115] Laney, C. B. *Computational Gasdynamics*, Cambridge University Press, New York, 1998.
 - [116] Laskey, K. B. Model Uncertainty: Theory and Practical Implications, *IEEE Transactions on Systems, Man and Cybernetics-Part A: Systems and Humans*, Vol. 26, No. 3, 1996; 340-348.
 - [117] Law, A. M., and Kelton, W. D. *Simulation Modeling and Analysis*, 2nd ed., McGraw-Hill, New York, 1991.
 - [118] Lee, L. H., and Poolla, K. Statistical Validation for Uncertainty Models, Lecture Notes in Control and Information Sciences, Vol. 202, *Feedback Control, Complexity, and Identification: A festschrift for Professor George Zames*, Montreal, Canada, 1994; 131-149.
 - [119] Lee, L. H., and Poolla, K. On Statistical Model Validation, *Journal of Dynamic Systems*,

- Measurement and Control*, Vol. 118, 1996; 226-236.
- [120] LeGore, T. Predictive Software Validation Methodology for Use with Experiments Having Limited Replicability, in *Benchmark Test Cases for Computational Fluid Dynamics*, FED-Vol. 93, I. Celik and C.J. Freitas eds., American Society of Mechanical Engineers, New York, 1990; 21-27.
- [121] Levy, D. W., Zickuhr, T., Vassberg, J., Agrawal, S., Wahls, R. W., Pirzadeh, S., and Hensch, M. J. Summary of Data from the First AIAA CFD Drag Prediction Workshop, AIAA-2002-0841, *40th AIAA Aerospace Sciences Meeting and Exhibit*, Reno, NV, 2002.
- [122] Lewis, R. O. *Independent Verification and Validation*, 1st ed., John Wiley, New York, 1992.
- [123] Lin, S. J., Barson, S. L., and Sindir, M. M. Development of Evaluation Criteria and a Procedure for Assessing Predictive Capability and Code Performance, *Advanced Earth-to-Orbit Propulsion Technology Conference*, Marshall Space Flight Center, Huntsville, AL, 1992.
- [124] Logan, R. W., and Nitta, C. K. Verification & Validation Guidelines and Quantitative Reliability at Confidence (QRC): Basis for an Investment Strategy, Lawrence Livermore National Laboratory, in press, Livermore, CA, 2002.
- [125] Mahadevan, S. Personal Communication, 2002.
- [126] Marvin, J. G. Perspective on Computational Fluid Dynamics Validation, *AIAA Journal*, Vol. 33, No. 10, 1995; 1778-1787.
- [127] Modarres, M. *What Every Engineer Should Know about Reliability and Risk Analysis*, Marcel Dekker, Inc., New York, 1993.
- [128] Mohanty, S., Codell, R. B., Ahn, T. M., and Cragolino, G. A. An Approach to the Assessment of High-Level Radioactive Waste Containment II: Radionuclide Releases from an Engineered Barrier System, *Nuclear Engineering and Design*, Vol. 201, 2000; 307-325.
- [129] Morton, K. W. *Numerical Solution of Convection-Diffusion Problems*, CRC Press, Boca Raton, FL, 1996.
- [130] Murray-Smith, D. J. Methods for the External Validation of Continuous Systems Simulation Models: A Review, *Mathematical and Computer Modelling of Dynamics Systems*, Vol. 4, 1998; 5-31.
- [131] NAFEMS. *CFD Working Group*, International Association for the Engineering Analysis Community, www.NAFEMS.org.
- [132] NRC. Severe Accident Risks: An Assessment for Five U.S. Nuclear Power Plants, U.S. Nuclear Regulatory Commission, Office of Nuclear Regulatory Research, Division of Systems Research, NUREG-1150, Washington, DC, 1990.
- [133] NRC. *Statistical Software Engineering*, National Research Council, National Academy Press, Washington, DC, 1996.
- [134] Oberkampf, W. L. A Proposed Framework for Computational Fluid Dynamics Code Calibration/Validation, AIAA Paper No. 94-2540, *18th AIAA Aerospace Ground Testing Conference*, Colorado Springs, CO, 1994.
- [135] Oberkampf, W. L., and Aeschliman, D. P. Joint Computational/Experimental Aerodynamics Research on a Hypersonic Vehicle: Part 1, Experimental Results, *AIAA Journal*, Vol. 30, No. 8, 1992; 2000-2009.
- [136] Oberkampf, W. L., Aeschliman, D. P., Henfling, J. F., and Larson, D. E. Surface Pressure Measurements for CFD Code Validation in Hypersonic Flow, AIAA Paper No. 95-2273, *26th AIAA Fluid Dynamics Conf.*, San Diego, CA, 1995.

- [137] Oberkampf, W. L., Aeschliman, D. P., Tate, R. E., and Henfling, J. F. Experimental Aerodynamics Research on a Hypersonic Vehicle, Sandia National Laboratories, SAND92-1411, Albuquerque, NM, 1993.
- [138] Oberkampf, W. L., and Blottner, F. G. Issues in Computational Fluid Dynamics Code Verification and Validation, *AIAA Journal*, Vol. 36, No. 5, 1998; 687-695.
- [139] Oberkampf, W. L., DeLand, S. M., Rutherford, B. M., Diegert, K. V., and Alvin, K. F. Estimation of Total Uncertainty in Computational Simulation, Sandia National Laboratories, SAND2000-0824, Albuquerque, NM, 2000.
- [140] Oberkampf, W. L., DeLand, S. M., Rutherford, B. M., Diegert, K. V., and Alvin, K. F. Error and Uncertainty in Modeling and Simulation, *Reliability Engineering and System Safety*, Vol. 75, No. 3, 2002; 333-357.
- [141] Oberkampf, W. L., Diegert, K. V., Alvin, K. F., and Rutherford, B. M. Variability, Uncertainty, and Error in Computational Simulations, ASME-HTD-Vol. 357-2, *AIAA/ASME Joint Thermophysics and Heat Transfer Conference*, Albuquerque, NM, 1998; 259-272.
- [142] Oberkampf, W. L., and Trucano, T. G. Validation Methodology in Computational Fluid Dynamics, AIAA 2000-2549, *Fluids 2000 Conference*, Denver, CO, 2000.
- [143] Oberkampf, W. L., and Trucano, T. G. Verification and Validation in Computational Fluid Dynamics, *Progress in Aerospace Sciences*, Vol. 38, No. 3, 2002; 209-272.
- [144] Oberkampf, W. L., and Trucano, T. G. Verification and Validation in Computational Fluid Dynamics, Sandia National Laboratories, SAND2002-0529, Albuquerque, NM, 2002.
- [145] Oden, J. T., Feng, Y., and Prudhomme, S. Local and Pollution Error Estimation For Stokesian Flow, *International Journal of Numerical Methods in Fluids*, Vol. 27, 1998; 33-39.
- [146] Oden, J. T. Error Estimation and Control in Computational Fluid Dynamics, in *The Mathematics of Finite Elements and Applications*, J.R. Whiteman ed. John Wiley, New York, 1993; 1-23.
- [147] Paez, T., and Urbina, A. Validation of Structural Dynamics Models Via Hypothesis Testing, *Society of Experimental Mechanics Annual Conference*, Portland, OR, 2001.
- [148] Paez, T. L., and Urbina, A. Validation of Mathematical Models of Complex Structural Dynamic Systems, *Proceedings of the Ninth International Congress on Sound and Vibration*, Orlando, FL, 2002.
- [149] Palmer, T. N. Predicting Uncertainty in Forecasts of Weather and Climate, *Reports on Progress in Physics*, Vol. 63, 2000; 71-116.
- [150] Partsch, H. A. *Specification and Transformation of Programs*, Springer-Verlag, New York, 1990.
- [151] Paté-Cornell, M. E. Uncertainties in Risk Analysis: Six Levels of Treatment, *Reliability Engineering and System Safety*, Vol. 54, 1996; 95-111.
- [152] Paté-Cornell, M. E. Conditional Uncertainty Analysis and Implications for Decision Making: The Case of WIPP, *Risk Analysis*, Vol. 19, No. 5, 1999; 1003-1016.
- [153] Paulk, M. C., Weber, C. V., Curtis, B., and Chrissis, M. B. eds. *The Capability Maturity Model: Guidelines for Improving the Software Process*, Addison-Wesley, Reading, MA, 1994.
- [154] Peercy, D. E. Personal Communication, 2000.
- [155] Pilch, M., Trucano, T. G., Moya, J. L., Froehlich, G. K., Hodges, A. L., and Peercy, D. E. Guidelines for Sandia ASCI Verification and Validation Plans - Content and Format:

- Version 2, Sandia National Laboratories, SAND2000-3101, Albuquerque, NM, 2001.
- [156] PSI. PureCoverage: User's Guide, Pure Software, Inc., Sunnyvale, CA, 1994.
 - [157] QNET-CFD. *Thematic Network on Quality and Trust for the Industrial Applications of CFD*, www.qnet-cfd.net.
 - [158] Roache, P. J. Need for Control of Numerical Accuracy, *Journal of Spacecraft and Rockets*, Vol. 27, No. 2, 1990; 98-102.
 - [159] Roache, P. J. Perspective: A Method for Uniform Reporting of Grid Refinement Studies, *Journal of Fluids Engineering*, Vol. 116, 1994; 405-413.
 - [160] Roache, P. J. Quantification of Uncertainty in Computational Fluid Dynamics, in *Annual Review of Fluid Mechanics*, 29, J.L. Lumley and M. Van Dyke eds., Annual Reviews, Inc., Palo Alto, CA, 1997; 126-160.
 - [161] Roache, P. J. *Verification and Validation in Computational Science and Engineering*, Hermosa Publishers, Albuquerque, NM, 1998.
 - [162] Roache, P. J. Verification of Codes and Calculations, *AIAA Journal*, Vol. 36, No. 5, 1998; 696-702.
 - [163] Roache, P. J. Code Verification by the Method of Manufactured Solutions, *Journal of Fluids Engineering*, Vol. 114, No. 1, 2002; 4-10.
 - [164] Robinson, S. Simulation Verification, Validation and Confidence: A Tutorial, *Transactions of the Society for Computer Simulation International*, Vol. 16, 1999; 63-69.
 - [165] Rook, P. *Software Reliability Handbook*, Elsevier Science Publishers, New York, 1990.
 - [166] Roy, C. J. Grid Convergence Error Analysis for Mixed-Order Numerical Schemes, AIAA-2001-2606, *AIAA Fluid Dynamics Conference*, Anaheim, CA, 2001.
 - [167] Roy, C. J., and Blottner, F. B. Assessment of One- and Two-equation Turbulence Models for Hypersonic Flows, *Journal of Spacecraft and Rockets*, Vol. 38, No. 5, 2001; 699-710.
 - [168] Roy, C. J., McWherter-Payne, M. A., and Oberkampf, W. L. Verification and Validation for Laminar Hypersonic Flowfields, AIAA2000-2550, *Fluids 2000 Conference*, Denver, CO, 2000.
 - [169] Rushby, J. Formal Methods and the Certification of Critical Systems, Computer Science Laboratory, SRI International, SRI-CSL-93-7, Menlo Park, CA, 1993.
 - [170] Salari, K., and Knupp, P. Code Verification by the Method of Manufactured Solutions, Sandia National Laboratories, SAND2000-1444, Albuquerque, NM, 2000.
 - [171] Sanders, F., Mullen, S. L., and Baumhefner, D. P. Ensemble Simulations of Explosive Cyclogenesis at Ranges of 2 - 5 Days, *Monthly Weather Review*, Vol. 128, No. 8/pt. 2, 2000; 2920-2934.
 - [172] Schlesinger, S. Terminology for Model Credibility, *Simulation*, Vol. 32, No. 3, 1979; 103-104.
 - [173] Shih, T. M. A Procedure to Debug Computer Programs, *International Journal for Numerical Methods in Engineering*, Vol. 21, No. 6, 1985; 1027-1037.
 - [174] Shih, T. M., Tan, C. H., and Hwang, B. C. Effects of Grid Staggering on Numerical Schemes, *International Journal for Numerical Methods in Fluids*, Vol. 9, 1989; 193-212.
 - [175] Sindir, M. M., Barson, S. L., Chan, D. C., and Lin, W. H. On the Development and Demonstration of a Code Validation Process for Industrial Applications, AIAA Paper No. 96-2032, *27th AIAA Fluid Dynamics Conf.*, New Orleans, LA, 1996.
 - [176] Sindir, M. M., and Lynch, E. D. Overview of the State-of-Practice of Computational Fluid Dynamics in Advanced Propulsion System Design, AIAA Paper No. 97-2124, *28th AIAA Fluid Dynamics Conference*, Snowmass, CO, 1997.

- [177] Singhal, A. K. Key Elements of Verification and Validation of CFD Software, AIAA 98-2639, *29th AIAA Fluid Dynamics Conference*, Albuquerque, NM, 1998.
- [178] Singpurwalla, N. D., and Wilson, S. P. *Statistical Methods in Software Engineering: Reliability and Risk*, Springer-Verlag, Berlin, 1999.
- [179] Smithson, M. *Ignorance and Uncertainty: Emerging Paradigms*, Springer-Verlag, New York, 1989.
- [180] Srivastava, B. N., Werle, M. J., and Davis, R. T. A Finite Difference Technique Involving Discontinuous Derivatives, *Computers and Fluids*, Vol. 7, No. 1, 1979; 69-74.
- [181] Steinberg, S., and Roache, P. J. Symbolic Manipulation and Computational Fluid Dynamics, *Journal of Computational Physics*, Vol. 57, No. 2, 1985; 251-284.
- [182] Stevenson, D. E. A Critical Look at Quality in Large-Scale Simulations, *Computing in Science and Engineering*, Vol. 1, No. 3, 1999; 53-63.
- [183] Stockman, C. T., Garner, J. W., Helton, J. C., Johnson, J. D., Shinta, A., and Smith, L. N. Radionuclide Transport in the Vicinity of the Repository and Associated Complementary Cumulative Distribution Functions in the 1996 Performance Assessment for the Waste Isolation Pilot Plant, *Reliability Engineering and System Safety*, Vol. 69, No. 1-3, 2000; 369-396.
- [184] Trucano, T. G., Easterling, R. G., Dowding, K. J., Paez, T. L., Urbina, A., Romero, V. J., Rutherford, R. M., and Hills, R. G. Description of the Sandia Validation Metrics Project, Sandia National Laboratories, SAND2001-1339, Albuquerque, NM, 2001.
- [185] Trucano, T. G., Pilch, M., and Oberkampf, W. L. General Concepts for Experimental Validation of ASCI Code Applications, Sandia National Laboratories, SAND2002-0341, Albuquerque, NM, 2002.
- [186] Turkel, E. Accuracy of Schemes with Nonuniform Meshes for Compressible Fluid-Flows, *Applied Numerical Mathematics*, Vol. 2, No. 6, 1986; 529-550.
- [187] Urbina, A., and Paez, T. L. Statistical Validation of Structural Dynamics Models, *Annual Technical Meeting & Exposition of the Institute of Environmental Sciences and Technology*, Phoenix, AZ, 2001.
- [188] Venditti, D. A., and Darmofal, D. L. Adjoint Error Estimation and Grid Adaptation for Functional Outputs: Application to Quasi-One-Dimensional Flow, *Journal of Computational Physics*, Vol. 164, No. 1, 2000; 204-227.
- [189] Walker, M. A., and Oberkampf, W. L. Joint Computational/Experimental Aerodynamics Research on a Hypersonic Vehicle: Part 2, Computational Results, *AIAA Journal*, Vol. 30, No. 8, 1992; 2010-2016.
- [190] Wallace, D. R., Ippolito, L. M., and Cuthill, B. B. Reference Information for the Software Verification and Validation Process, Rept. 500-234, 1996.
- [191] Wilson, G. E., and Boyack, B. E. The Role of the PIRT in Experiments, Code Development and Code Applications Associated With Reactor Safety Assessment, *Nuclear Engineering and Design*, Vol. 186, 1998; 23-37.
- [192] Wilson, G. E., Boyack, B. E., Catton, I., Duffey, R. B., Griffith, P., Katsma, K. R., Lellouche, G. S., Levy, S., Rohatgi, U. S., Wulff, W., and Zuber, N. Quantifying Reactor Safety Margins Part 2: Characterization of Important Contributors to Uncertainty, *Nuclear Engineering and Design*, Vol. 119, 1990; 17-31.
- [193] Wulff, W., Boyack, B. E., Catton, I., Duffey, R. B., Griffith, P., Katsma, K. R., Lellouche, G. S., Levy, S., Rohatgi, U. S., Wilson, G. E., and Zuber, N. Quantifying Reactor Safety Margins Part 3: Assessment and Ranging of Parameters, *Nuclear Engineering and Design*, Vol. 119, 1990; 33-65.

- [194] Yee, H. C., and Sweby, P. K. Aspects of Numerical Uncertainties in Time Marching to Steady-State Numerical Solutions, *AIAA Journal*, Vol. 36, No. 5, 1998; 712-724.
- [195] Zeigler, B. P., Praehofer, H., and Kim, T. G. *Theory of Modeling and Simulation: Integrating Discrete Event and Continuous Complex Dynamic Systems*, 2nd ed., Academic Press, San Diego, CA, 2000.
- [196] Zhang, X. D., Pelletier, D., Trepanier, J. Y., and Camarero, R. Verification of Error Estimators for the Euler Equations, AIAA-2000-1001, *38th AIAA Aerospace Sciences Meeting*, Reno, NV, 2000.
- [197] Zuber, N., Wilson, G. E., Ishii, M., Wulff, W., Boyack, B. E., Dukler, A. E., Griffith, P., Healzer, J. M., Henry, R. E., Lehner, J. R., Levy, S., and Moody, F. J. An Integrated Structure and Scaling Methodology for Severe Accident Technical Issue Resolution: Development of Methodology, *Nuclear Engineering and Design*, Vol. 186, No. 1-2, 1998; 1-21.