SANDIA REPORT

SAND2012-0107
Unlimited Release
Printed January 2012

Experiences Using DAKOTA Stochastic Expansion Methods in Computational Simulations

Joseph R. Ruthruff and Jeremy A. Templeton

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Approved for public release; further dissemination unlimited.

Sandia National Laboratories

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

NOTICE: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government, nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, make any warranty, express or implied, or assume any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represent that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government, any agency thereof, or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof, or any of their contractors.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from
U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-Mail: reports@adonis.osti.gov
Online ordering: http://www.osti.gov/bridge

Available to the public from
U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Rd
Springfield, VA 22161

Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-Mail: orders@ntis.fedworld.gov
Online ordering: http://www.ntis.gov/help/ordermethods.asp?loc=7-4-0#online

SAND2012-0107
Unlimited Release
Printed January 2012

Experiences Using DAKOTA Stochastic Expansion Methods in Computational Simulations

Joseph R. Ruthruff
Quantitative Modeling and Analysis
Sandia National Laboratories
P. O. Box 969
Livermore, CA 94551-0969
jruthru@sandia.gov

Jeremy A. Templeton
Thermal/Fluid Science and Engineering
Sandia National Laboratories
P. O. Box 969
Livermore, CA 94551-0969
jatempl@sandia.gov

Abstract

Uncertainty quantification (UQ) methods bring rigorous statistical connections to the analysis of computational and experiment data, and provide a basis for probabilistically assessing margins associated with safety and reliability. The DAKOTA toolkit developed at Sandia National Laboratories implements a number of UQ methods, which are being increasingly adopted by modeling and simulation teams to facilitate these analyses. This report disseminates results as to the performance of DAKOTA's stochastic expansion methods for UQ on a representative application. Our results provide a number of insights that may be of interest to future users of these methods, including the behavior of the methods in estimating responses at varying probability levels, and the expansion levels for the methodologies that may be needed to achieve convergence.

Acknowledgments

We thank the developers of the DAKOTA toolkit at Sandia National Laboratories who helped educate us as to the use of these methods and exercising them within DAKOTA. We are especially grateful to Mike Eldred for his assistance in using the methods, and Brian Adams and Bill Bohnhoff for their help in acquiring a developmental release of DAKOTA for use in these studies. We also thank Jerry McNeish for his comments on an earlier draft of this report. Bill Erickson developed the model that we make use of in this report. This work was supported by the Engineering V&V Assessments project within the Advanced Simulation and Computing (ASC) program.

Contents

Nomenclature **12**

1 Introduction **13**

2 Modeling Problem **15**

3 Stochastic Expansion UQ Methods in DAKOTA **17**

4 Case Study **21**

5 Results **27**

Polynomial Chaos Expansion: Tensor Product Quadrature Grid 27

Response Level 27

L^2 Norms 29

L^∞ Norms 32

Probability-level Convergence Plots 36

Stochastic Collocation: Tensor Product Quadrature Grids 39

Polynomial Chaos Expansion: Smolyak Sparse Grids 43

Response Level 43

L^2 Norms 45

L^∞ Norms 46

Probability-level Convergence Plots 48

Stochastic Collocation: Smolyak Sparse Grids 52

Response Level 53

L^2 Norms 53

L^∞ Norms 54

Probability-level Convergence Plots 54

6 Concluding Remarks **63**

Probability Levels 63

Convergence of Cumulative Distribution Functions 63

Tensor Product Quadrature Grids versus Smolyak Sparse Grids 64

Improvements to DAKOTA 64

Global Reliability Methods 65

References **66**

List of Figures

2.1 Computational simulation workflow of modeling problem. 15

3.1 Example growth of full tensor product quadrature and Smolyak sparse grids, as originally shown in the DAKOTA User's Manual. 19

4.1 Example DAKOTA input deck specifying the PCE methodology using a third-level expansion of Smolyak sparse grids. 22

4.2 A DAKOTA ensemble workflow. The process of computing CDFs for each response at the conclusions of the CALORE evaluations is not shown. 25

5.1 Polynomial Chaos Expansion: tensor product grid levels 2–7. 30

5.2 Polynomial Chaos Expansion: tensor product grid levels 2–4. 30

5.3 Polynomial chaos expansion, tensor product grid levels 2–7: L^2 norm for response level. 32

5.4 Polynomial chaos expansion, tensor product grid levels 2–4: L^2 norm for response level. 33

5.5 Polynomial chaos expansion, tensor product grid levels 5–7: L^2 norm for response level. 33

5.6 Polynomial chaos expansion, tensor product grid levels 2–7: L^∞ norm for response level. 35

5.7 Polynomial chaos expansion, tensor product grid levels 2–4: L^∞ norm for response level. 35

5.8 Polynomial chaos expansion, tensor product grid levels 5–7: L^∞ norm for response level. 36

5.9 Convergence of the calculated margins from PCE as the quadrature grid level increases. 37

5.10 Convergence of the calculated margins from PCE, based on quadrature grids, as the number of DAKOTA evaluations increases. 37

5.11 Polynomial chaos expansion, sparse grid expansion levels 1–7: Response level values. 45

5.12 Polynomial chaos expansion, sparse grid expansion levels 1–7: L^2 norm for response level values. 47

5.13 Polynomial chaos expansion, sparse grid expansion levels 1–4: L^2 norm for response level values. 47

5.14 Polynomial chaos expansion, sparse grid expansion levels 5–7: L^2 norm for response level values. 48

5.15 Polynomial chaos expansion, sparse grid expansion levels 1–7: L^∞ norm for response level values. 50

5.16 Polynomial chaos expansion, sparse grid expansion levels 1–4: L^∞ norm for response level values. 50

5.17 Polynomial chaos expansion, sparse grid expansion levels 5–7: L^∞ norm for response level values. 51

5.18 Convergence of the calculated margins from PCE as the sparse grid level increases. 51

5.19 Convergence of the calculated margins from PCE, based on sparse grids, as the number of DAKOTA evaluations increases. 52

5.20 Polynomial chaos expansion, sparse grid expansion levels 1–7: Response level values. 54

5.21 Stochastic collocation, sparse grid expansion levels 1–7: L^2 norm for response level values. 56

5.22 Stochastic collocation, sparse grid expansion levels 1–4: L^2 norm for response level values. 56

5.23 Stochastic collocation, sparse grid expansion levels 5–7: L^2 norm for response level values. 57

5.24 Stochastic collocation, sparse grid expansion levels 1–7: L^∞ norm for response level values. 57

5.25 Stochastic collocation, sparse grid expansion levels 1–4: L^∞ norm for response level values. 59

5.26 Stochastic collocation, sparse grid expansion levels 5–7: L^∞ norm for response level values. 59

5.27 Convergence of the calculated margins from SC as the sparse grid level increases. . 60

5.28 Convergence of the calculated margins from SC, based on sparse grids, as the number of DAKOTA evaluations increases. 61

List of Tables

5.1 The 28 DAKOTA ensembles conducted to evaluate the UQ methods. 27

5.2 Observed samples for PCE ensembles using tensor product quadrature grids. 28

5.3 Polynomial chaos expansion, tensor product grid expansion levels 2–7: Response level values. 28

5.4 Polynomial chaos expansion, tensor product grid expansion levels 2–7: L^2 norm for response level values. 31

5.5 Polynomial chaos expansion, tensor product grid expansion levels 2–7: L^∞ norm for response level values. 34

5.6 Convergence of estimated response to a "true" solution at the seventh-level quadrature grid expansion using PCE. 38

5.7 Observed samples for SC ensembles using tensor product quadrature grids. 39

5.8 Stochastic collocation, tensor product grid expansion levels 2–7: Response level values. 40

5.9 Stochastic collocation, tensor product grid levels 2–7: L^2 norm for response level. . 41

5.10 Stochastic collocation, tensor product grid expansion levels 2–7: L^∞ norm for response level values. 42

5.11 Observed samples for PCE ensembles using Smolyak sparse grids. 43

5.12 Polynomial chaos expansion, sparse grid expansion levels 1–7: Response level values. 44

5.13 Polynomial chaos expansion, sparse grid expansion levels 1–7: L^2 norm for response level values. 46

5.14 Polynomial chaos expansion, sparse grid expansion levels 1–7: L^∞ norm for response level values. 49

5.15 Convergence of estimated response to a "true" solution at the seventh-level sparse grid expansion using PCE. 49

5.16 Observed samples for SC ensembles using Smolyak sparse grids. 52

5.17 Stochastic collocation, sparse grid expansion levels 1–7: Response level values. . . . 53

5.18 Stochastic collocation, sparse grid expansion levels 1–7: L^2 norm for response level values. 55

5.19 Stochastic collocation, sparse grid expansion levels 1–7: L^∞ norm for response level values. 58

5.20 Convergence of estimated response to a “true” solution at the seventh-level sparse grid expansion using SC. 61

Nomenclature

ASC Advanced Simulation & Computing

CDF Cumulative Distribution Function

DAKOTA Design Analysis Kit for Optimization and Terascale Applications

FSM Full System Model

PCE Polynomial Chaos Expansion

QMU Quantification of Margins and Uncertainty

SC Stochastic Collocation

SL Strong-link

UQ Uncertainty Quantification

V&V Verification and Validation

WL Weak-link

Chapter 1

Introduction

As the United States has moved into a era of Science-Based Stockpile Stewardship after the moratorium on full system tests of nuclear weapons, Sandia National Laboratories has brought computational simulation efforts to bear on continuing assessments that seek to ensure the safety and reliability of the nation's nuclear deterrent. At Sandia, these efforts are primarily supported by the Advanced Simulation and Computing (ASC) program — formerly the Accelerated Strategic Computing Initiative — and make use of Full System Models that have been created for different environments of interest. The problem we describe here is concerned with a model designed for simulations concerning abnormal thermal environments, such as fuel fires.

Verification and validation (V&V) is an important methodology that supports quantification of confidence in the predictive capability of a continuum physics model, and computational simulations making use of that model [3, 5]. Verification is a process to confirm that the model is correctly computing results within the solution codes, while validation is a process to confirm that modeling results are consistent with real world phenomena. The Full System Model (FSM) teams in the Nuclear Weapons program at Sandia National Laboratories make extensive use of V&V to confirm the predictive capability of the modeling and computational simulations used as a technical basis for assessments regarding weapon safety and reliability.

An important process that must take place in the context of V&V is uncertainty quantification (UQ). Uncertainty quantification seeks to quantitatively assess the effect of input uncertainty on response metrics of interest [1]. In the context of modeling, these uncertainties can be aleatoric in nature, meaning they derive from inherent and irreducible variability that cannot be controlled in the phenomena being modeled, or epistemic in nature, meaning they derive from lack of knowledge about the variables of interest, but are therefore reducible if sufficient effort could be made to better characterize and understand the uncertainty. Uncertainty quantification is important to V&V from at least two standpoints: (1) it brings rigor to the statistical analysis of computational and experiment data, and the comparison between the two during the validation process of V&V; and (2) it provides a basis for quantifying the uncertainty in a margin that is computed in the course of a quantification of margins and uncertainty (QMU) study [3].

One toolkit that provides a variety of uncertainty quantification methods is DAKOTA (Design Analysis Kit for Optimization and Terascale Applications) [8, 1], which is developed at Sandia in conjunction with collaborators in academia. In addition to its many other capabilities, DAKOTA makes available a variety of uncertainty quantification methods, including sampling methods, relia-

bility methods, stochastic expansion methods, and epistemic non-deterministic methods [1]. Many of these methods, however, have not been empirically explored using problems other than the examples provided in the distribution of DAKOTA. This means that some facets of these methods, such as the number of samples needed to achieve convergence in results, are not as well understood as we would like. Because our modeling problems are very large, we would prefer to have guidelines for using these methods that are supported by a reputable technical basis, such as their performance on smaller yet non-trivial and representative problems.

This report explores the use of some of DAKOTA's UQ methods — specifically its stochastic expansion methods — on an unclassified model of a Sandia-developed weapon subsystem. This model, which will be described further in Chapter 2, is non-trivial in nature yet representative of other models seen in Full System Modeling work at Sandia. As such, the primary contributions of this report are results and insights regarding the performance of some DAKOTA stochastic expansion methods for a non-trivial problem instance that is derived from an important computational simulation domain.

The remainder of this report is organized as follows: Chapter 2 describes the model on which these UQ methods are explored, Chapter 3 provides a brief overview of the stochastic expansion methods explored in this report, Chapter 4 outlines the study conducted to exercise these methods, Chapter 5 presents and discusses the results of using the methods on the model, and Chapter 6 offers concluding remarks.

Chapter 2

Modeling Problem

In this report, we make use of a nuclear weapon subsystem model that contains limited but complex geometric fidelity to the subsystem of interest, and assigns a generic stainless steel material property to the entire subsystem.[1] The solid geometry of this model was meshed using Sandia's CUBIT toolkit [6] into over 150,000 tetrahedral finite elements (86 element blocks), making this is a non-trivial problem in terms of model size.

We exercise this weapon subsystem model using the CALORE code [2] in the Sierra Mechanics tool suite [7], which is also developed and maintained by Sandia. CALORE performs thermal simulations as specified by input decks describing the conditions and parameters of a simulation. These input decks are provided to CALORE along with the meshed finite element model. The present CALORE simulation with these artifacts takes less than 30 minutes to complete, with the time step control parameters we have specified, using 16 processors on Sandia's capacity computing clusters. The resulting output of this simulation provides temperature data for each node in the model at various time steps, which is written using the Exodus finite element data model [9]. This workflow is outlined in Figure 2.1.

For a FSM, we can use this temperature data over time to compute various measures of the state of this subsystem in the abnormal thermal environment of interest. The response we measure

[1] Because this model is not approved for unlimited release, we do not provide further details about its design in this report.

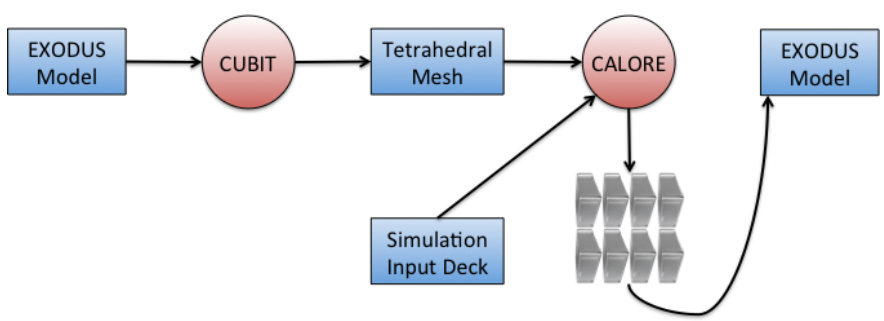

Figure 2.1. Computational simulation workflow of modeling problem.

in this report is calculated by examining the temperature of two components in the model at specific periods of time in order to create a weak-link/strong-link thermal race. Weak-link/strong-link systems are a common and important safety characteristic of high-consequence systems, including nuclear weapons. Strong-links (SLs) act as a barrier to signals or energy that might trigger a response from a high-consequence system. These stimuli originate from components designed to fail before the strong-links in an unintended scenario where, from a safety perspective, the system must remain non-operational (e.g., an accident condition); the components providing these signals or energy are therefore termed weak-links (WLs). In the context of the thermal race, failures of both weak-links and strong-links are induced by their exposure to extreme heat in an abnormal thermal scenario, such as if a system were exposed to a fire, and can be characterized by an expected failure temperature for each component.

The failure of the weak-links ensures that the signals or energy required to generate a system are not present, thereby diffusing the system with respect to being triggered by that component. (In normal scenarios, the strong-links and weak-links must unambiguously agree about the instructions for the system before the strong-links allow these signals or energy to pass.) A thermal race margin is the time-based safety margin in the "race" to ensure the weak-links fail before the strong-links in abnormal thermal conditions. High-consequence systems are designed to maximize this margin in the direction of the weak-links failing before the strong-links. This thermal race is "lost" if the strong-links fail before the weak-links (which would generate a negative margin less than 0.0). Note that this does not mean that a system response will be triggered — only that the design features preventing such a response have been disabled in this safety scenario.

The goal of this report is not to evaluate a representative safety question or metric, but rather to explore the use of DAKOTA for performing uncertainty quantification on a non-trivial and representative problem space. To generate this scenario in this investigation, we treat one component in our subsystem model as a single strong-link and another component as a single weak-link, creating a thermal race scenario that can be measured in the context of our computational CALORE simulations. We have specifically chosen components (and failure temperatures for those components) according to the following criteria:

1. The components will create an arbitrary thermal race that is not representative of an evaluation of the system from which this model was derived.

2. The arbitrary WL/SL thermal race is expected to ensure that the weak-link always fails well before the strong-link (meaning the quantitative margin will always be greater than 0.0). This does not consider any probabilities of events occurring that may be associated with such margin calculations.

3. The selected failure criterion for each component will be temperatures that are not representative of the system, and attempt to ensure that the arbitrary WL/SL thermal race will finish in the time allotted to the CALORE simulations (meaning both the weak-link and strong-link reach their failing temperatures so that a margin can be calculated). Otherwise, an inordinately long simulation time might be needed to ensure we are able to generate a response for each simulation in a stochastic expansion ensemble driven by DAKOTA.

Chapter 3

Stochastic Expansion UQ Methods in DAKOTA

A common approach to uncertainty quantification is to characterize and quantify the role of uncertain variables contributing to one or more calculated responses of interest. DAKOTA provides a number of methods to perform UQ, including sampling methods, reliability methods, stochastic expansion methods, and epistemic non-deterministic methods [1].

Each of these methods have their own advantages, making them an improvement on the state-of-the-art for certain circumstances. For example, reliability methods are often better at computing statistics in the tails of response distributions (i.e., events with low probability) than traditional sampling-based approaches. The stochastic expansion methods we explore in this report express their final solution as a random process, rather than a set of statistics as do many other nondeterministic methodologies [1], making them particularly attractive for the multi-physics applications in our problem domains. However, there remains a need to exercise all these methods on real-world problems to both gauge their effectiveness in general and to better understand how to optimally employ these methods to ensure a characterization of sufficient pedigree.

In this report, we explore the use of stochastic expansion methods for performing UQ on the modeling problem presented in Chapter 2 because, as discussed earlier, they are well-suited for multi-physics applications. (Time and resource constraints have thus far prevented us from considering other methods.) In particular, we explore the use of Polynomial Chaos Expansion (PCE) and Stochastic Collocation (SC) for characterizing responses of interest to this problem, such as our arbitrary margin. Each of these two methods are described in detail in [1], and we refer readers to that source for a more complete account of these methods. In these chapter, we briefly overview both PCE and SC to provide the necessary context for this work.

Stochastic expansion methods are useful when the response(s) of the system(s) of interest can be characterized by a governing set of equations with stochastic coefficients [1].[1] Because their solutions are expressed as a random process, the methods are well-suited for the multi-physics problem domain of interest in this work [1]. This is due to inherent, aleatoric uncertainties in aspects of multi-physics applications such as material properties (i.e., the precise properties of each material may slightly vary from system to system). As we later discuss for future work, we hope to

[1] Because not all our uncertainties fall into this category, it is possible that other UQ methods may prove more appropriate over time. We plan to investigate such methods in the future.

explore in detail other UQ methods implemented by DAKOTA, especially the toolkit's reliability methods, as low-probability events are of particular importance to designers of high-consequence systems.

Polynomial Chaos Expansion (PCE) is the first stochastic expansion method we explore. This method models the effect of continuous random response variables in terms of truncated finite-dimensional series expansions taking place in standardized random variables [1] for the purpose of estimating the stochastic coefficients. To do this, an expansion method (of which there are several for PCE) is used to cast the problem into a mathematical space based on known orthogonal polynomial basis functions. From these functions, one can generate sets of denumerable (i.e., countable) random values for the uncertain variables of interest and measure the response of interest accordingly. A problem solution is then generated from these observations by estimating coefficients for the basis functions alongside the rest of the expansion. These coefficients represent an approximated mapping between your selected response function set and the underlying random variables, which provides a surrogate for the true simulation-based input-output mapping.

Stochastic collocation (SC) is closely related to PCE. The primary difference between the two methods is that SC forms interpolation functions for known coefficients, rather than estimating coefficients for known basis functions as PCE does [1]. In the context of SC, the expansion coefficients are simply the response values at each collocation point within the interpolation function. Thus, while PCE must define an expansion formulation and estimate coefficients accordingly, SC formulates an interpolation for known coefficients [1].

In the case of both approaches, a spectral projection approach is required to project the response against the basis functions to generate the coefficients. We explore two different approaches that are supported by DAKOTA for both polynomial chaos expansion and stochastic collocation. The first approach is to utilize a full tensor product quadrature grid. A nested abscissas approach is used within this grid to identify evaluations (in terms of values for the uncertain variables) that are needed to generate the coefficients based on the response values obtained for those evaluations [1]. The expansion level of this approach can rise to increase the number of collocation points in the tensor grid selected for the evaluations. Because the same expansion level is used for all random dimensions in our studies, this approach generates m^n function evaluations, where m is the expansion level and n is the ensemble's dimensionality (i.e., number of uncertain variables) [1]. Furthermore, the number of function evaluations is identical when this method is used in conjunction with both PCE and SC.

The second approach we evaluate is Smolyak sparse grids, which attempt to reduce the number of collocation points while preserving accuracy in the generated coefficients [1]. Stated simply, Smolyak sparse grids use linear combinations of the product formulas used in tensor product quadratures to generate collocation points. However, to reduce the number of needed evaluations, they purposefully use isotropic products with only a small number of points. This offers isotropic sparse grids a scaling rate of $m^{\log n}$, where m is again the expansion level and n is the ensemble dimensionality, rather than the m^n growth seen in tensor product quadrature grids [1]. An example depicting the growth of both full tensor product quadrature and Smolyak sparse grids are depicted in Figure 3.1; both plots are originally found in [1].

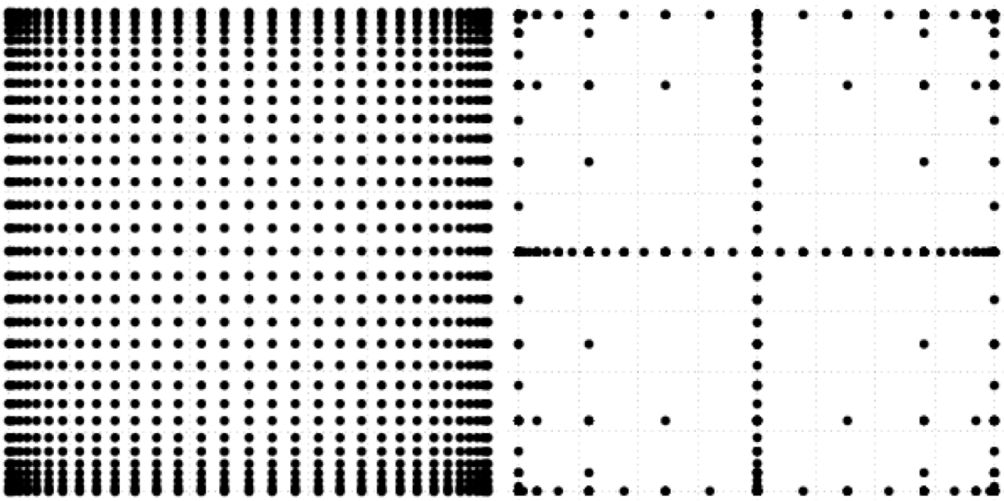

Figure 3.1. Example growth of full tensor product quadrature and Smolyak sparse grids, as originally shown in the DAKOTA User's Manual.

This page intentionally left blank.

Chapter 4

Case Study

For our investigation, we designed a case study to evaluate the PCE and SC stochastic expansion by examining the role of the generic stainless steel material property assigned across the model described in Chapter 2. In particular, in the context of an abnormal thermal environment simulated using CALORE, there are three material properties of stainless steel that may impact the temperature response in the model: thermal conductivity (κ), heat capacity (c_p), and emissivity (ε). We treat a scaling of these temperature dependent properties of the stainless steel as three uncertain variables whose effect on a thermal response will be separately examined by ensembles of evaluations driven by DAKOTA's PCE (DAKOTA syntax: `nond_polynomial_chaos`) and SC (`nond_stoch_collocation`) methods. We used a developmental release of DAKOTA 5.0+, built on June 23, 2010 from Subversion revision 6835 to drive these ensembles.

For a spectral projection method for DAKOTA, we used both the tensor product grid (DAKOTA syntax: `quadrature_order`) and Smolyak sparse grid (`sparse_grid_level`) methods overviewed in Section 3. For each projection method, we explore expansion levels 1–7 for both PCE and SC, resulting in a total of 28 DAKOTA ensembles. Each DAKOTA ensemble invokes multiple directed CALORE simulations to explore the specified uncertainty space for κ, c_p, and ε, with ensembles driven by higher-level expansions of the grids resulting in additional CALORE simulations (i.e., increased cost).[1] We expect that the highest expansion levels (e.g., a seventh-level expansion) are far more than is needed to adequately capture the distribution of the response; in fact, experienced DAKOTA developers suggested to us that third- to fourth-level expansions are generally sufficient in most cases. However, the DAKOTA-based UQ ensembles described in this report are far less expensive than the ensembles we plan to ultimately conduct. Thus, we consider the higher-level expansions in this report to help evaluate whether such high expansion levels provide additional value that is warranted their exorbitant cost, and to demonstrate convergence of quantities of interest.

An example of an input deck for one of the 28 DAKOTA ensembles is presented in Figure 4.1. This input deck specifies the polynomial chaos expansion method (line 7) with a third-level expansion of Smolyak sparse grids (line 8). A seed of 12345 for the random number generator (line 10) was chosen for all DAKOTA ensembles in this case study.

The goal of each UQ study conducted in DAKOTA is to measure the effect of the uncertain

[1] For the remainder of this chapter, we differentiate between a DAKOTA "ensemble," which consists of individual evaluations of CALORE simulations and their associated post-processing, and the case study that is the subject of this report, which consists of 28 DAKOTA ensembles.

```
strategy
  tabular_graphics_data
    tabular_graphics_file = 'PCEWizard.dat'
  single_method

method
  nond_polynomial_chaos
    sparse_grid_level 3
    samples 1000000
    seed 12345
    probability_levels .000001 .00001 .0001 .001 .01 .05 .1 .2 .3 .4 .5
  .6 .7 .8 .9 .95 .99 .000001 .00001 .0001 .001 .01 .05 .1 .2 .3 .4 .5 .6
  .7 .8 .9 .95 .99 .000001 .00001 .0001 .001 .01 .05 .1 .2 .3 .4 .5 .6 .7
  .8 .9 .95 .99 .000001 .00001 .0001 .001 .01 .05 .1 .2 .3 .4 .5 .6 .7 .8
  .9 .95 .99 .000001 .00001 .0001 .001 .01 .05 .1 .2 .3 .4 .5 .6 .7 .8 .9
  .95 .99 .000001 .00001 .0001 .001 .01 .05 .1 .2 .3 .4 .5 .6 .7 .8 .9 .95
  .99

variables
  uniform_uncertain 3
    lower_bounds 0.5 0.5 0.75
    upper_bounds 1.5 1.5 0.95
    descriptors 'k' 'cp' 'e'

interface
  fork
    analysis_driver = 'calore_driver.sh'
    parameters_file 'calore.in'
    results_file 'calore.out'
    file_tag
    file_save

responses
  num_response_functions 3
  descriptors 'L2_norm' 'Linf_norm' 'Margin'
  no_gradients
  no_hessians
```

Figure 4.1. Example DAKOTA input deck specifying the PCE methodology using a third-level expansion of Smolyak sparse grids.

variables κ, c_p, and ε on our arbitrary margin. To do this, we use the `probability_levels` method parameter to instruct DAKOTA to generate Cumulative Distribution Functions (CDFs) for the response level for each ensemble. These CDFs help us gauge the range of possible values for the response, and the probability of those values, given perturbations in κ, c_p, and e. To formulate these CDFs, we began by instructing DAKOTA to compute the response level at 10% increments for the probability levels (i.e., 10%, 20%, ..., 90%). We then included the 1%, 5%, 95%, and 99% probability levels, as these are levels that are commonly recommended by the statistical community. Finally, for the actual margins we are ultimately interested in evaluating, we will be interested in the behavior of our response at the lower end of the CDFs, where the calculated response might approach zero (indicative of a smaller margin of safety). Therefore, we also included probability levels of 10^{-3}, 10^{-4}, 10^{-5}, and 10^{-6}. The response levels at each probability level are computed by DAKOTA by sampling the expansion after the stochastic coefficients have been computed; in these DAKOTA ensembles 1,000,000 samples were performed.[2]

We treat κ, c_p, and ε as uniformly-distributed uncertain variables (DAKOTA syntax: `uniform_uncertain`). By this we mean that we instruct DAKOTA to uniformly explore the ranges for each uncertain variable, as specified by the lower and upper bounds for each variable in the input deck (lines 15–16 in Figure 4.1). Although thermal conductivity and heat capacity can be represented as constants, in this model these material properties are defined by temperature-dependent piecewise linear functions to make them more representative to actual problems of interest. The specified lower and upper bounds for these variables scale each function by 50% in each direction to provide a range for κ and c_p when treated as uniformly-distributed uncertain variables. The emissivity of the steel in the original CALORE input deck that was provided to us by the model designer was a constant value: 0.85. Scaling emissivity is different than for heat capacity or thermal conductivity because, in addition to being a constant value, ε must be between 0.0 and 1.0. Therefore, we choose to scale it between 0.75 and 0.95 to bring it reasonably close to its upper bound, while scaling it no more severely in the opposite direction.

We post-process the Exodus data models that are created by CALORE using the ENCORE application in the Sierra Mechanics suite [7]. Using ENCORE, we extract the temperatures of our arbitrarily weak-link and strong-link at various time steps. We use customized Python scripts to calculate the time of failure for each component by interpolating the temperatures between the two time steps surrounding each component's failure temperature. These scripts then calculate the thermal race margin by subtracting the weak-link failure time from the strong-link failure time ("Margin" in line 29 of Figure 4.1).

In addition the margin calculations, for every CALORE simulation, we compute the L^2 and L^∞ norms ("L2_norm" and "Linf_norm" in line 29) of the temperatures across the entire model as compared to a "nominal" evaluation where κ, c_p, and ε are not perturbed from their original values. In general terms, an L^2 is a global metric representing what can be thought of as a standardized average of the values in an arbitrary vector, while the L^∞ norm is the maximum value in that vector.

[2]One drawback of this approach is that performing 1,000,000 samples is actually less than desired for the 10^{-6} probability level. This sampling is chosen due to limitations in DAKOTA for storing data in memory from larger sample sizes. However, 1,000,000 samples has the effect of basing the response level for the 10^{-6} probability level on only one data point.

Specifically, for a vector:

$$\mathbf{x} = \begin{vmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{vmatrix}$$

the L^2 norm is

$$| \mathbf{x} | = \sqrt{\sum_{i=1}^{n} |x_i|^2}$$

while the L^∞ norm is

$$| \mathbf{x} |_\infty = max(x_1, x_2, \ldots, x_n)$$

In this case, the L^2 norm is a measure of the average difference in observed temperaturess between each evaluation and the nominal case, while the L^∞ norm is the maximum difference. Some areas of the nominal model may have a higher temperature than a given evaluation, while other areas may have a lower temperature, and we do not want such directions of discrepency to counteract each other and indicate that the magnitude of temperature differences is less than it really is. Therefore, the norms we compute do not take negative values: a norm of 0.0 indicates no temperature differences between an evaluation and the nominal case, while norms increasingly greater than zero indicate increasingly greater temperature differences, irrespective of whether the nominal model has higher or lower temperatures, or a mix in each direction.

Both the L^2 and L^∞ norms are computed using ENCORE. While extracting the weak-link and strong-link temperatures for the margin calculation, ENCORE computes these norms by post-processing comparing the temperatures across the Exodus model resulting from each CALORE evaluation to the nominal case.

This workflow is depicted in Figure 4.2. DAKOTA selects input parameters for each uncertain variable, which are substituted into the CALORE simulation input deck. After an approximately 30-minute CALORE simulation using 16 processors, ENCORE begins the post-processing step by processing the resulting Exodus data model, and the Python completes post-processing by calculating the thermal race margin from the resulting temperature data. This margin, along with the L^2 and L^∞ norms calculated by ENCORE, are then provided back to DAKOTA so they can be used to calculate the CDFs of each response at the conclusion of each DAKOTA ensemble.

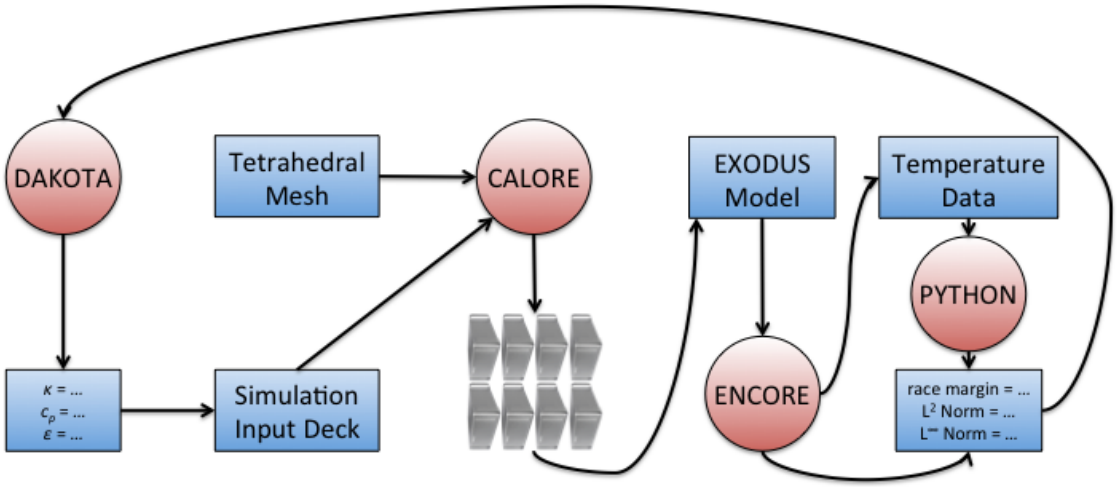

Figure 4.2. A DAKOTA ensemble workflow. The process of computing CDFs for each response at the conclusions of the CALORE evaluations is not shown.

This page intentionally left blank.

Chapter 5

Results

In this chapter, we present the results for the 28 DAKOTA ensembles for the purposes of uncertainty quantification. We divide these ensembles into four sets for each UQ methodology and spectral projection method combination (see Table 5.1).

Polynomial Chaos Expansion: Tensor Product Quadrature Grid

The first set of DAKOTA ensembles we present are those driven by polynomial chaos expansion using tensor product quadrature grids. The number of evaluations required by DAKOTA to generate the basis functions for each grid expansion level, which are then sampled in order to obtain the cumulative distribution functions, are provided in Table 5.2.

Response Level

Table 5.3 presents the results for the polynomial chaos expansion ensembles using tensor product quadrature grids, as well as providing the number of evaluations needed for those ensembles from Table 5.2. We present the results for only the ensembles with expansion levels 2–7, as the ensemble with a first-level expansion contains only one evaluation. (Recall that this methodology requires m^n evaluations, where m is the expansion level and n is the number of uncertain variables. As seen in Table 5.2, with an expansion level of one, $m^n = 1^3 = 1$.)

It is worth mentioning that there is some fluctuations in the estimated response level values

	Spectral Projection	
UQ Methodology	*Quadrature Grid*	*Sparse Grid*
Polynomial Chaos Expansion	Expansion Levels 1–7	Expansion Levels 1–7
Stochastic Collocation	Expansion Levels 1–7	Expansion Levels 1–7

Table 5.1. The 28 DAKOTA ensembles conducted to evaluate the UQ methods.

— particularly for the probability levels at the tails of the CDFs — at the lower expansion levels. This is because the method is still converging to the actual estimated distribution of the response. In general, nothing should be read into these fluctuations other than the methods do not yet have enough data upon which to base their estimations, and thus a DAKOTA ensemble based on a

Expansion Level	**Observed Evaluations**
1	1
2	8
3	27
4	64
5	125
6	216
7	343

Table 5.2. Observed samples for PCE ensembles using tensor product quadrature grids.

Probability Level	**PCE-QUAD2**	**PCE-QUAD3**	**PCE-QUAD4**	**PCE-QUAD5**	**PCE-QUAD6**	**PCE-QUAD7**
0.000001	256.0886	306.5850	290.6769	295.6303	294.3255	294.5847
0.00001	257.8407	306.8070	291.9730	296.6318	295.4277	295.6837
0.0001	261.4189	307.7672	295.2509	298.9111	298.0695	298.1935
0.001	274.4286	313.0581	305.7019	307.3761	307.2024	307.2466
0.01	316.3085	337.8036	337.9673	337.3026	337.4439	337.4783
0.05	398.5474	401.1710	403.2651	403.3936	403.3421	403.2890
0.1	464.4763	460.9952	461.5682	461.9706	462.0855	462.0222
0.2	563.3894	559.9280	559.0895	559.0692	559.0980	559.0946
0.3	645.3881	646.2299	645.8927	645.7309	645.6514	645.6443
0.4	725.8265	727.3683	727.5404	727.5061	727.4863	727.5075
0.5	807.9255	808.7758	809.0145	809.0269	809.0020	809.0106
0.6	900.6417	889.6241	890.2811	889.8611	889.9059	889.9160
0.7	1008.9456	983.6282	985.4548	985.8844	985.6563	985.7250
0.8	1142.4200	1127.4547	1120.2338	1121.0794	1121.3338	1121.2942
0.9	1324.3481	1351.5760	1345.0337	1342.4166	1342.3705	1342.4607
0.95	1459.1637	1531.8185	1537.3787	1534.9826	1534.1477	1534.0130
0.99	1644.9036	1800.1327	1838.4151	1845.9102	1846.8578	1846.7348
Evaluations	8	27	64	125	216	343

Table 5.3. Polynomial chaos expansion, tensor product grid expansion levels 2–7: Response level values.

higher-level expansion is needed.

However, we are not particularly interested in the actual response levels shown in Table 5.3, as the response is an arbitrary margin for a non-representative model. However, we did note that in the left tail of the CDF approaching the 10^{-6} probability level, there were notable changes in the response level for the ensembles with lower expansion levels. For example, the ensemble with a fourth-level grid expansion (`PCE-QUAD4`) saw the response level drop from 337.9673 at probability level 0.01 (10^{-2}) to 290.6769 at probability level 10^{-6}. We take this to mean that there is possible utility in actually calculating these ever-smaller response levels in our actual models of interest, if the consideration of such small levels is warranted.

We also observed that there is little change in the response around the middle of the CDF curve as the expansion level of the ensemble increases and more observations are considered, while there are notable changes in the response at both tails of the CDF. For example, for probability level 0.5, the response level only increases from 807.9255 to 809.0106 as the expansion level rises from two to seven — this despite the evaluations of the ensemble increasing from $2^3 = 8$ to $7^3 = 343$. However, in the left tail at probability level 0.01, the response level increases from 316.3085 to 337.4783; the change is even more pronounced in the right tail at probability level 0.99, where the response level increases from 1644.9036 to 1846.7348. As we discussed earlier, this is a strong indication that higher-level expansions are needed for these estimated responses to converge.

But perhaps our greatest interest for this work is how these UQ techniques converge as the expansion level is increased. That is, at what point does the increase in the expansion level, and the rise in the evaluations (and therefore cost) required for the UQ method as a result of that increase, result in little to no change in the estimated response? To gauge this, we plot all the CDFs for the expansion level 2–7 ensembles in Figure 5.1. As can be seen, there appears to be little change to the response overall from a third-level expansion onward; note, for example, how the right-tail of the CDF for the second-level expansion (blue) is to the left of the CDF tail for the third-level expansion (red), which is still slightly to the left of the remaining CDFs, which are stacked on top of each other. This is perhaps better seen in Figure 5.2, which plots the CDFs for only the second-, third-, and fourth-level expansions. There are small changes in the tails of the CDFs between the third- and fourth-level expansions, but generally speaking the CDFs are quite similar, as can be observed in Table 5.3. As such, there is a reasonable question as to whether the small changes in the CDF of fourth-level expansions and onward are worth the non-linear increase in evaluations that are required to generate those estimations.[1]

L^2 Norms

Our arbitrary margin metric provides a localized measure to study in the context of uncertainty quantification; it is highly relevant to a particular modeling scenario (i.e., a safety question for a high-consequence system), but does not directly consider how the uncertain variables impact other

[1] This judgment is merely qualitative based on a visual inspection of the CDFs, which was suitable for our purposes of better understanding these UQ methods. These differences could be quantified, but such rigor is beyond the scope of what we sought out of these results.

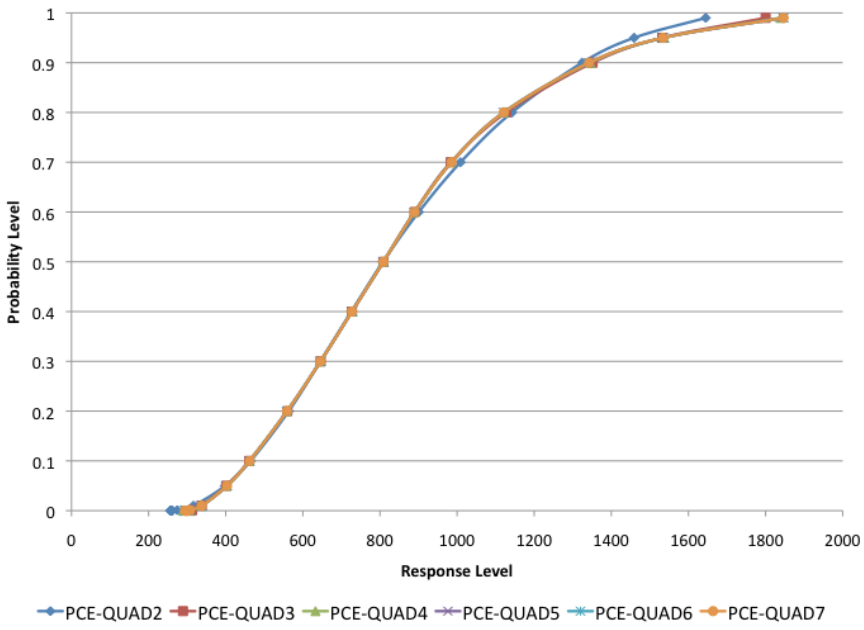

Figure 5.1. Polynomial Chaos Expansion: tensor product grid levels 2–7.

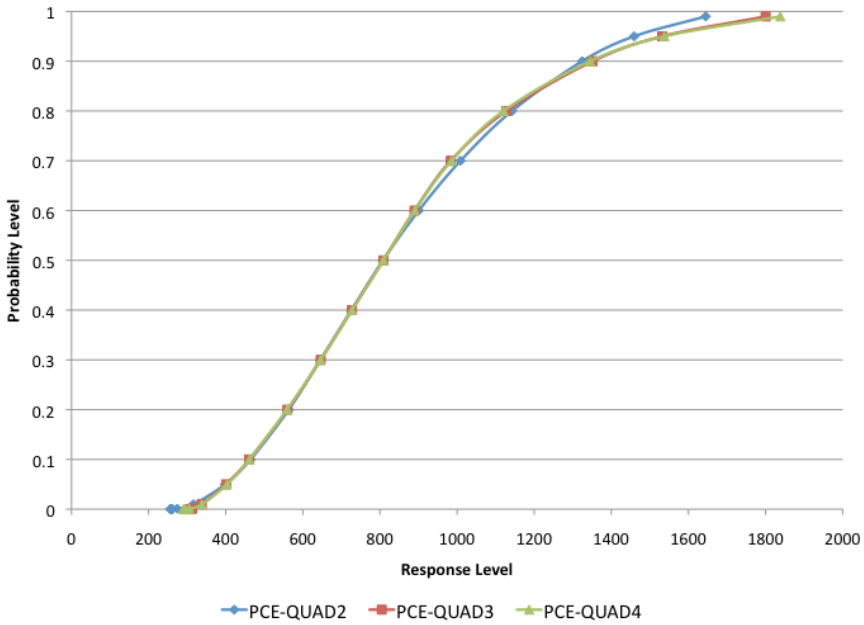

Figure 5.2. Polynomial Chaos Expansion: tensor product grid levels 2–4.

Probability Level	PCE-QUAD2	PCE-QUAD3	PCE-QUAD4	PCE-QUAD5	PCE-QUAD6	PCE-QUAD7
0.000001	-0.1125	-0.1041	-0.0820	-0.0622	-0.0465	-0.0343
0.00001	-0.1110	-0.1013	-0.0779	-0.0569	-0.0403	-0.0274
0.0001	-0.1078	-0.0943	-0.0677	-0.0442	-0.0259	-0.0123
0.001	-0.0971	-0.0721	-0.0375	-0.0111	0.0038	0.0048
0.01	-0.0726	-0.0177	0.0062	0.0072	0.0074	0.0076
0.05	-0.0339	0.0075	0.0107	0.0113	0.0115	0.0113
0.1	-0.0078	0.0134	0.0151	0.0148	0.0142	0.0141
0.2	0.0247	0.0227	0.0219	0.0214	0.0217	0.0224
0.3	0.0457	0.0336	0.0312	0.0326	0.0340	0.0351
0.4	0.0600	0.0437	0.0446	0.0472	0.0490	0.0503
0.5	0.0688	0.0579	0.0613	0.0643	0.0661	0.0665
0.6	0.0775	0.0764	0.0809	0.0837	0.0840	0.0836
0.7	0.0920	0.0996	0.1033	0.1043	0.1030	0.1017
0.8	0.1131	0.1282	0.1293	0.1263	0.1246	0.1241
0.9	0.1458	0.1623	0.1579	0.1554	0.1558	0.1564
0.95	0.1723	0.1823	0.1807	0.1798	0.1796	0.1802
0.99	0.2121	0.2251	0.2207	0.2273	0.2273	0.2251
Evaluations	8	27	64	125	216	343

Table 5.4. Polynomial chaos expansion, tensor product grid expansion levels 2–7: L^2 norm for response level values.

areas of the model. Our L^2 and L^∞ norms seek to evaluate the degree of differences explored by the evaluations in these UQ methodologies across the entire model, based on a nominal case where the κ, c_p, and ε material properties are not perturbed. The distribution of these norms is also quite different from a thermal race margin, which unlike the norms does not have a minimum value of 1.0 (or any minimum, for that matter). As such, studying these variables provides possible insight as to the behavior of these UQ methods on a different class of variables.

Table 5.4 provides the estimations of the relative L^2 norm, which measures the average difference between a given evaluation for PCE ensembles using tensor product quadrature grids. From this table, it seems that these ensembles are slower to converge for higher levels of grid expansion — even in the middle probability levels where we previously found the response level to converge quite quickly. This can be graphically seen quite clearly; for example, the three CDF curves plotted in Figure 5.2 line up much more closely than do the curves in the upcoming figures. This indicates that uncertainty quantification efforts based on relative L^2 and L^∞ norms may require additional effort in terms of grid expansions in order for the calculated norms to converge toward each other.

One potential issue that can be drawn from Table 5.4 is that some of the estimated response levels at the edge of the left tail are less than zero. These negative estimates occur even for seventh-level expansions of the PCE ensemble when over 300 observations are gathered for only three uncertain variables. These negative numbers for the L^2 norms are notable because the estimates

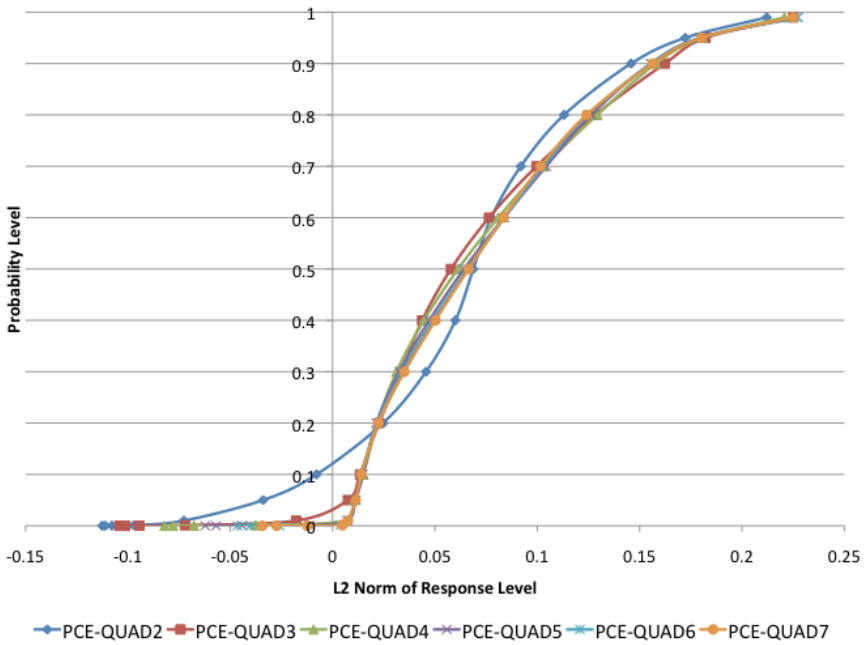

Figure 5.3. Polynomial chaos expansion, tensor product grid levels 2–7: L^2 norm for response level.

should fall within the 0.0–1.0 range that is valid for a standardized L^2 norm metric. This suggests that careful consideration should be given to results so far out in the left tail of the CDF, especially for the lower-level PCE ensembles where smaller numbers of evaluations are available upon which to base estimations. These negative estimates may also be reflective of the limitations of the stochastic expansion methods themselves, and may be explained by Gibbs' phenomenon, which would account for the oscillating of the negative L^2 norm predictions.

Figure 5.3 illustrates how results according to this metric are slower to converge than with the raw response level. In particular, Figure 5.4 reveals that the ensembles for expansion levels 2–4 show notable differences among each other, while Figure 5.5 suggests that levels 5–7 are generally consistent with each other with respect to differences found from the nominal case. The fact that differences still exist in the CDFs plotted in Figure 5.4

L^∞ Norms

Table 5.5 provides the estimations of the relative L^∞ norm, which measures the maximum difference between a given evaluation and the nominal case for PCE ensembles using tensor product quadrature grids. This table suggests that the L^∞ norm metric for the response is another case where the PCE ensembles are slower to converge, at least when based on tensor product quadrature grids, than for the response itself.

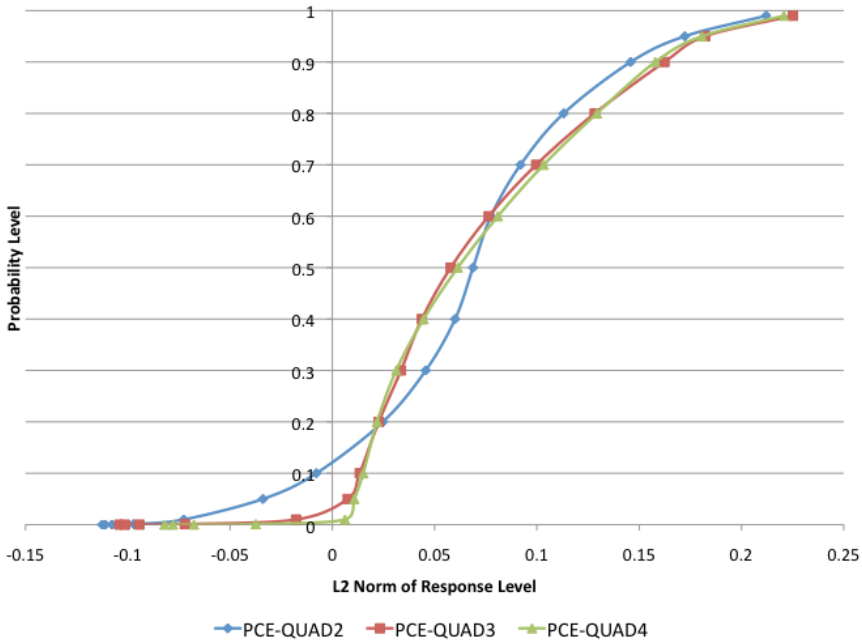

Figure 5.4. Polynomial chaos expansion, tensor product grid levels 2–4: L^2 norm for response level.

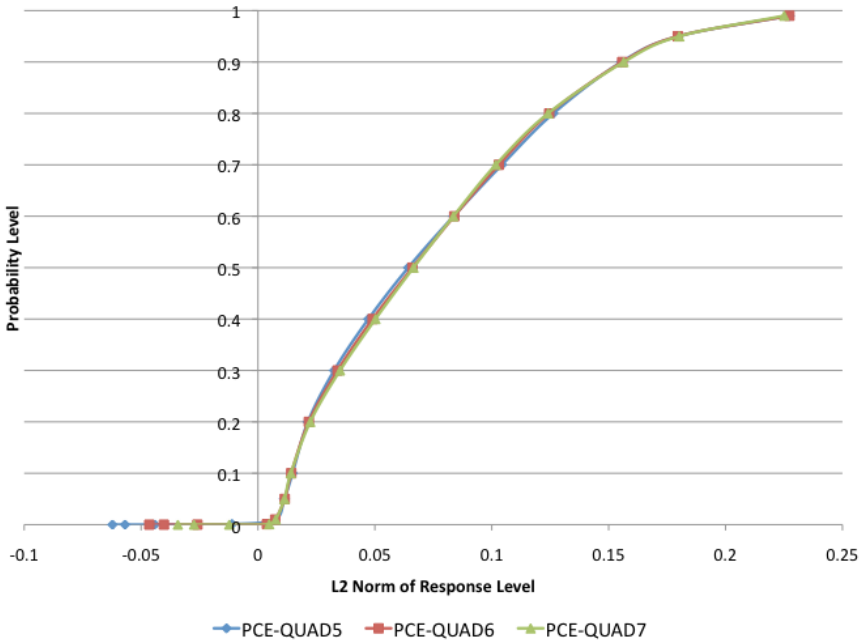

Figure 5.5. Polynomial chaos expansion, tensor product grid levels 5–7: L^2 norm for response level.

Probability Level	PCE-QUAD2	PCE-QUAD3	PCE-QUAD4	PCE-QUAD5	PCE-QUAD6	PCE-QUAD7
0.000001	-0.0623	-0.0482	-0.0160	0.0111	0.0214	0.0358
0.00001	-0.0601	-0.0439	-0.0112	0.0154	0.0269	0.0397
0.0001	-0.0533	-0.0338	0.0003	0.0254	0.0389	0.0504
0.001	-0.0397	-0.0123	0.0240	0.0465	0.0558	0.0579
0.01	-0.0056	0.0429	0.0584	0.0663	0.0595	0.0624
0.05	0.0381	0.0804	0.0751	0.0742	0.0705	0.0683
0.1	0.0661	0.0861	0.0823	0.0782	0.0750	0.0726
0.2	0.1006	0.0955	0.0903	0.0870	0.0850	0.0842
0.3	0.1227	0.1061	0.1007	0.1004	0.1009	0.1015
0.4	0.1379	0.1164	0.1160	0.1178	0.1189	0.1198
0.5	0.1472	0.1314	0.1348	0.1367	0.1385	0.1391
0.6	0.1565	0.1514	0.1560	0.1578	0.1586	0.1594
0.7	0.1714	0.1772	0.1795	0.1802	0.1809	0.1809
0.8	0.1933	0.2101	0.2107	0.2077	0.2091	0.2076
0.9	0.2274	0.2533	0.2512	0.2501	0.2501	0.2488
0.95	0.2548	0.2812	0.2820	0.2819	0.2799	0.2815
0.99	0.2951	0.3296	0.3225	0.3208	0.3221	0.3225
Evaluations	8	27	64	125	216	343

Table 5.5. Polynomial chaos expansion, tensor product grid expansion levels 2–7: L^∞ norm for response level values.

It is notable that from a fifth-level grid expansion onward we do not have any "erroneous" estimates of a L^∞ norm that are less than 0.0. This may be one indication regarding the expansion level needed for an ensemble to ensure accurate and reliable results. It is not surprising that L^∞ norms were estimated to be greater than 0.0, for the lowest probability levels, for expansion levels whose same levels were consistently less than 0.0 for L^2 norms. This is because L^∞ norms take a maximum and will naturally be greater than L^2 norms, so we would expect L^∞ norm estimations to more rapidly gravitate above 0.0 (perhaps due to this norm being smoother near 0.0).

As with the previously presented L^2 norms, we plot the CDFs for the L^∞ norms for these PCE ensemble using tensor product quadrature grids in Figure 5.6, and then divide those CDFs up between expansion levels 2–4 and 5–7 in Figures 5.7 and 5.8.

Once again, in Figure 5.7 we see notable differences in the CDFs resulting from the lower grid expansion levels. These differences are much smaller in Figure 5.8, although we can see from this figure that differences can be seen in the left tails of the CDFs (which, as stated in Section 4, we are especially interested in). Because the L^∞ norms aim to illustrate the maximum difference from the nominal case, Figure 5.8 suggests these differences may be greatest in magnitude within the left tails of the CDFs. This is of key significance because left tails (i.e., low probability events) are often of special interest in safety-related questions involving high-consequence systems.

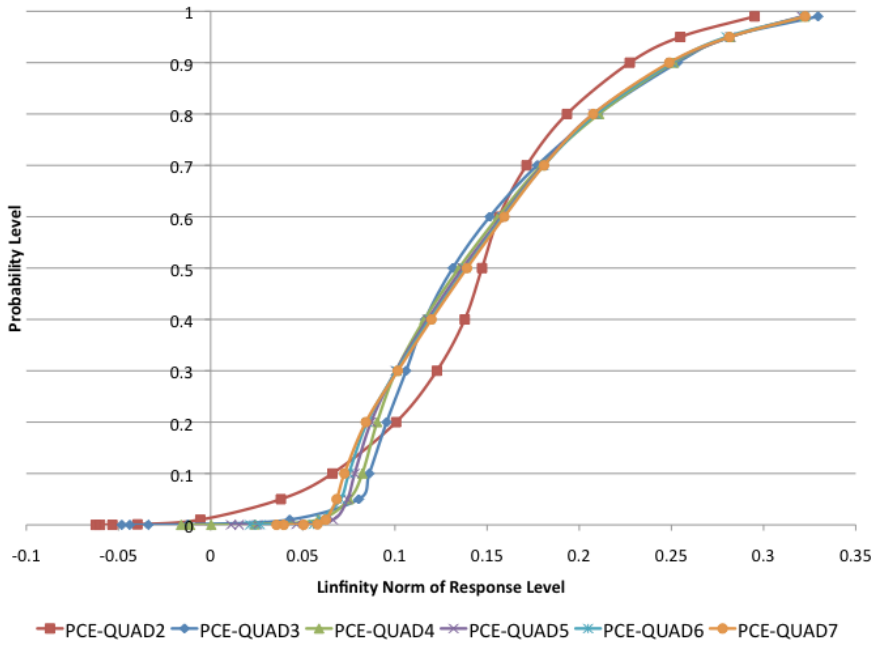

Figure 5.6. Polynomial chaos expansion, tensor product grid levels 2–7: L^∞ norm for response level.

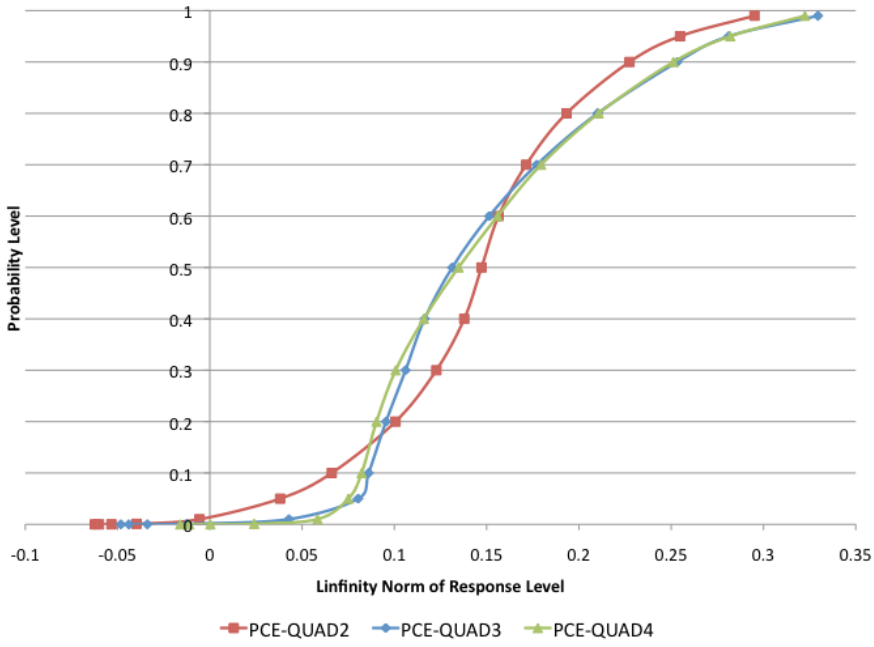

Figure 5.7. Polynomial chaos expansion, tensor product grid levels 2–4: L^∞ norm for response level.

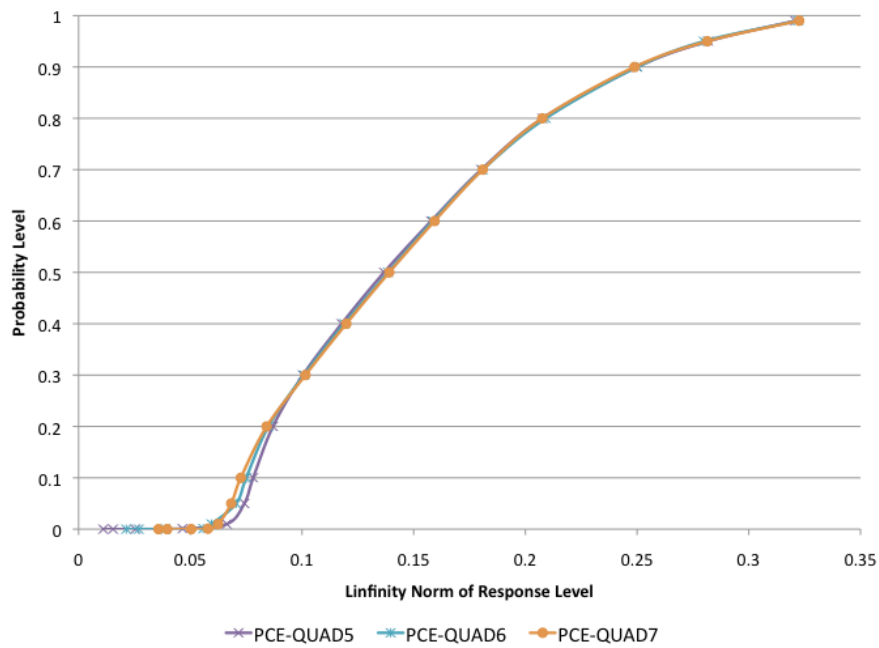

Figure 5.8. Polynomial chaos expansion, tensor product grid levels 5–7: L^∞ norm for response level.

Probability-level Convergence Plots

The final analysis we perform for each stochastic expansion method explored is a verification of whether increasing the expansion levels of the grids lead to a quantifiable convergence of the method. In this case, convergence is measured in terms of the differences between each lower-level evaluation and the highest-level evaluation performed (in this case, the seventh-level expansion of the tensor product quadrature grid), which is taken to represent the "true" solution to assess convergence in the Cauchy sense [4].

Figure 5.9 presents the convergence of the differences in margin calculations between the seventh-level expansion of the tensor product quadrature grids, taken to be the true solution, with each of the first- to sixth-level grid expansions. These differences are plotted on a logarithmic scale due to the wide numerical range of data. Figure 5.10 plots the same differences against the number of DAKOTA evaluations performed for each expansion level (also plotted logarithmically). The calculations plotted in this figures are offered in Table 5.6. We considered the predicted response levels at the 10^{-5} probability level because those were the lowest probability level where DAKOTA was able to conduct enough samples off the basis functions to provide a reasonable estimate of the response. The square of these differences between each estimate and the seventh-level estimate is plotted to expand the space between estimates very close to the solution and those that are not, as well as to ensure the plotted differences are positive and not sensitive to the direction of the difference from the seventh-level response.

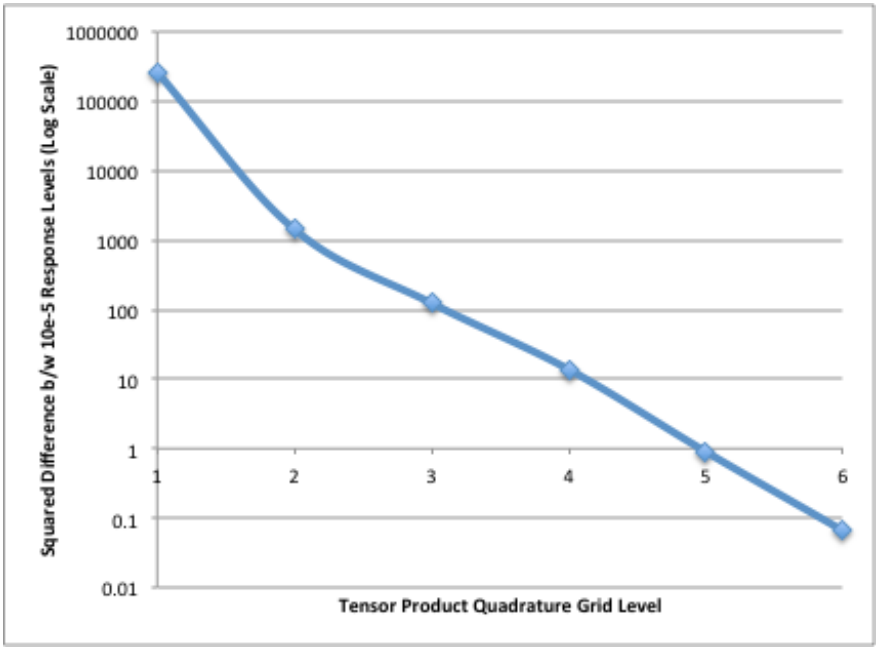

Figure 5.9. Convergence of the calculated margins from PCE as the quadrature grid level increases.

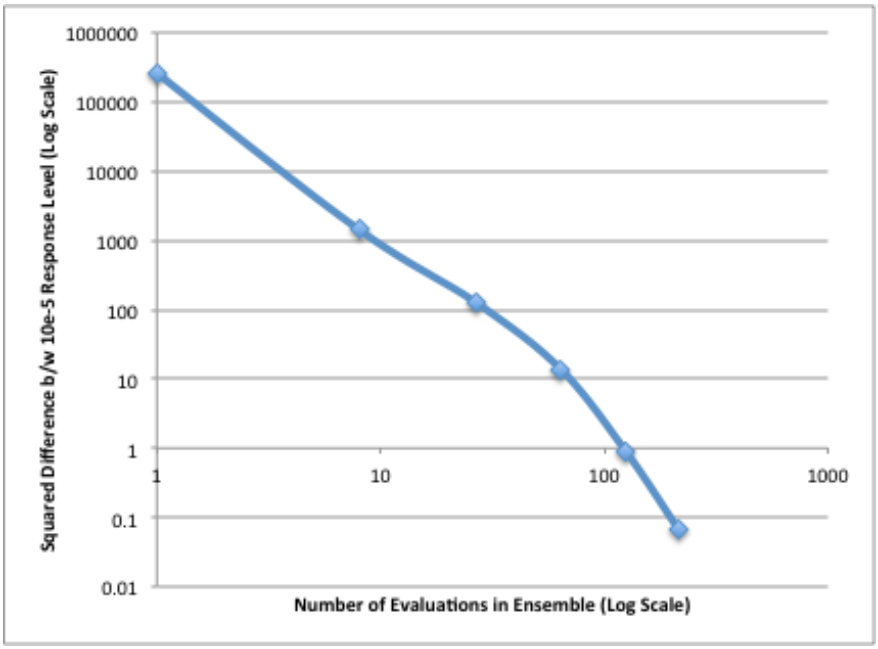

Figure 5.10. Convergence of the calculated margins from PCE, based on quadrature grids, as the number of DAKOTA evaluations increases.

Expansion Level	DAKOTA Evaluations	Response at 10^{-5} Probability Level	Difference from 7th-level Response	Square of the Difference
1	1	796.6431073	-500.9593939	250960.3144
2	8	257.8406706	37.84304285	1432.095892
3	27	306.8069935	-11.12328005	123.7273591
4	64	291.9729886	3.71072476	13.76947824
5	125	296.6318	-0.94808662	0.898868239
6	216	295.4276891	0.25602433	0.065548458
7	343	295.6837134	0.0	0.0

Table 5.6. Convergence of estimated response to a "true" solution at the seventh-level quadrature grid expansion using PCE.

As can be seen in this figures, there is still demonstrable progress toward converging to this "solution" even up to the sixth-level expansion. This actually runs contrary to some of the visual trends we observed in Figures 5.1 and 5.2, where it appeared the CDFs were virtually identical from the fourth-level expansion onward. However, this progress that is being made is very small — in some cases less than a second — and may not be justified by the expense of these higher-level ensembles.

Stochastic Collocation: Tensor Product Quadrature Grids

The second set of ensembles we present are those driven by SC using tensor product quadrature grids. The observed number of evaluations to generate the basis functions in these ensembles are provided in Table 5.7. Note that these numbers are the same as those previously presented in Table 5.2 for the polynomial chaos expansion results based on tensor product quadrature grids.

In fact, when either the polynomial chaos or stochastic collocation methods are based on tensor product quadrature grids, the expansion levels result in the same 3^n samples being selected for the three uncertain variables being studied. This means that the results for each method are the same. Tables 5.8–5.10 are presented below for comparison against the results using the polynomial chaos expansion based on tensor product quadrature expansion levels, previously presented in Tables 5.3–5.5. Because these results are the same, we do not present the duplicate figures or convergence plots as well.

Expansion Level	Observed Evaluations
1	1
2	8
3	27
4	64
5	125
6	216
7	343

Table 5.7. Observed samples for SC ensembles using tensor product quadrature grids.

Probability Level	**SC-QUAD2**	**SC-QUAD3**	**SC-QUAD4**	**SC-QUAD5**	**SC-QUAD6**	**SC-QUAD7**
0.000001	256.0886	306.5850	290.6769	295.6303	294.3255	294.5847
0.00001	257.8407	306.8070	291.9730	296.6318	295.4277	295.6837
0.0001	261.4189	307.7672	295.2509	298.9111	298.0695	298.1935
0.001	274.4286	313.0581	305.7019	307.3761	307.2024	307.2466
0.01	316.3085	337.8036	337.9673	337.3026	337.4439	337.4783
0.05	398.5474	401.1710	403.2651	403.3936	403.3421	403.2890
0.1	464.4763	460.9952	461.5682	461.9706	462.0855	462.0222
0.2	563.3894	559.9280	559.0895	559.0692	559.0980	559.0946
0.3	645.3881	646.2299	645.8927	645.7309	645.6514	645.6443
0.4	725.8265	727.3683	727.5404	727.5061	727.4863	727.5075
0.5	807.9255	808.7758	809.0145	809.0269	809.0020	809.0106
0.6	900.6417	889.6241	890.2811	889.8611	889.9059	889.9160
0.7	1008.9456	983.6282	985.4548	985.8844	985.6563	985.7250
0.8	1142.4200	1127.4547	1120.2338	1121.0794	1121.3338	1121.2942
0.9	1324.3481	1351.5760	1345.0337	1342.4166	1342.3705	1342.4607
0.95	1459.1637	1531.8185	1537.3787	1534.9826	1534.1477	1534.0130
0.99	1644.9036	1800.1327	1838.4151	1845.9102	1846.8578	1846.7348
Evaluations	8	27	64	125	216	343

Table 5.8. Stochastic collocation, tensor product grid expansion levels 2–7: Response level values.

Probability Level	**SC-QUAD2**	**SC-QUAD3**	**SC-QUAD4**	**SC-QUAD5**	**SC-QUAD6**	**SC-QUAD7**
0.000001	-0.1125	-0.1041	-0.0820	-0.0622	-0.0465	-0.0343
0.00001	-0.1110	-0.1013	-0.0779	-0.0569	-0.0403	-0.0274
0.0001	-0.1078	-0.0943	-0.0677	-0.0442	-0.0259	-0.0123
0.001	-0.0971	-0.0721	-0.0375	-0.0111	0.0038	0.0048
0.01	-0.0726	-0.0177	0.0062	0.0072	0.0074	0.0076
0.05	-0.0339	0.0075	0.0107	0.0113	0.0115	0.0113
0.1	-0.0078	0.0134	0.0151	0.0148	0.0142	0.0141
0.2	0.0247	0.0227	0.0219	0.0214	0.0217	0.0224
0.3	0.0457	0.0336	0.0312	0.0326	0.0340	0.0351
0.4	0.0600	0.0437	0.0446	0.0472	0.0490	0.0503
0.5	0.0688	0.0579	0.0613	0.0643	0.0661	0.0665
0.6	0.0775	0.0764	0.0809	0.0837	0.0840	0.0836
0.7	0.0920	0.0996	0.1033	0.1043	0.1030	0.1017
0.8	0.1131	0.1282	0.1293	0.1263	0.1246	0.1241
0.9	0.1458	0.1623	0.1579	0.1554	0.1558	0.1564
0.95	0.1723	0.1823	0.1807	0.1798	0.1796	0.1802
0.99	0.2121	0.2251	0.2207	0.2273	0.2273	0.2251
Evaluations	8	27	64	125	216	343

Table 5.9. Stochastic collocation, tensor product grid levels 2–7: L^2 norm for response level.

Probability Level	**SC-QUAD2**	**SC-QUAD3**	**SC-QUAD4**	**SC-QUAD5**	**SC-QUAD6**	**SC-QUAD7**
0.000001	-0.0623	-0.0482	-0.0160	0.0111	0.0214	0.0358
0.00001	-0.0601	-0.0439	-0.0112	0.0154	0.0269	0.0397
0.0001	-0.0533	-0.0338	0.0003	0.0254	0.0389	0.0504
0.001	-0.0397	-0.0123	0.0240	0.0465	0.0558	0.0579
0.01	-0.0056	0.0429	0.0584	0.0663	0.0595	0.0624
0.05	0.0381	0.0804	0.0751	0.0742	0.0705	0.0683
0.1	0.0661	0.0861	0.0823	0.0782	0.0750	0.0726
0.2	0.1006	0.0955	0.0903	0.0870	0.0850	0.0842
0.3	0.1227	0.1061	0.1007	0.1004	0.1009	0.1015
0.4	0.1379	0.1164	0.1160	0.1178	0.1189	0.1198
0.5	0.1472	0.1314	0.1348	0.1367	0.1385	0.1391
0.6	0.1565	0.1514	0.1560	0.1578	0.1586	0.1594
0.7	0.1714	0.1772	0.1795	0.1802	0.1809	0.1809
0.8	0.1933	0.2101	0.2107	0.2077	0.2091	0.2076
0.9	0.2274	0.2533	0.2512	0.2501	0.2501	0.2488
0.95	0.2548	0.2812	0.2820	0.2819	0.2799	0.2815
0.99	0.2951	0.3296	0.3225	0.3208	0.3221	0.3225
Evaluations	8	27	64	125	216	343

Table 5.10. Stochastic collocation, tensor product grid expansion levels 2–7: L^∞ norm for response level values.

Polynomial Chaos Expansion: Smolyak Sparse Grids

The third set of ensembles we present are those driven by polynomial chaos expansion, but this time using Smolyak sparse grids. The observed number of evaluations to generate the basis functions in these ensembles are provided in Table 5.11. Note that these numbers are much larger than those previously presented in Table 5.2, despite the fact that sparse grids are generally supposed to result in slower growth rates than do tensor product quadrature grids. This is because having only three uncertain variables in this formative case study slows the m^n growth of the quadrature grid methodology. If the number of uncertain variables were much larger (such as 10 or 20), the quadrature grid growth would rapidly eclipse that of the sparse grids.

Response Level

Table 5.12 presents the results for the PCE ensembles using Smolyak sparse grids. For this set of ensembles, we also present the results from the first-level sparse grid expansion, which for stochastic collocation contains seven evaluations instead of one, in addition to the second- through seventh-level ensembles. We also provide the observed number of evaluations to generate the basis functions in this table.

As before, we are not particularly concerned with the actual responses shown in Table 5.12, as they come from an arbitrary margin. We did notice for this set of ensembles that the convergence of those response values, however, occurred quite rapidly. In Figure 5.11, the CDF plots converged by the second-level expansion, even at the left-tail that was sometimes slower to converge in previous ensembles. This can also be seen from Table 5.12, where the values are already quite close to each other from the PCE-SGRID2 column and onward to the right side of the table. This more rapid convergence could be at least in part due to the increased number of evaluations that were required for the sparse grids (due to the small number of uncertain variables that were studied).

Expansion Level	Observed Evaluations
1	7
2	31
3	111
4	303
5	687
6	1375
7	2335

Table 5.11. Observed samples for PCE ensembles using Smolyak sparse grids.

Probability Level	PCE-SGRID1	PCE-SGRID2	PCE-SGRID3	PCE-SGRID4	PCE-SGRID5	PCE-SGRID6	PCE-SGRID7
0.000001	215.1997	282.4075	294.3085	294.8279	294.3912	294.5794	294.4411
0.00001	215.4527	283.9526	295.4558	295.8747	295.5438	295.6950	295.5937
0.0001	216.8676	287.5845	298.0398	298.3399	298.1524	298.2592	298.2137
0.001	224.7357	299.8211	307.2737	307.2833	307.2309	307.2553	307.2478
0.01	260.7595	337.1567	337.4329	337.4517	337.4691	337.4618	337.4817
0.05	351.2758	405.2467	403.2915	403.3420	403.2973	403.2892	403.2659
0.1	432.8379	462.1477	461.9964	462.0257	462.0323	462.0257	462.0102
0.2	562.5460	557.9411	559.0494	559.1181	559.0891	559.1130	559.0758
0.3	670.3572	644.7116	645.6414	645.6487	645.6729	645.6890	645.6827
0.4	767.1856	727.4888	727.5009	727.5024	727.5140	727.5232	727.5390
0.5	856.2118	809.0351	808.9883	809.0080	809.0088	808.9969	808.9794
0.6	939.3272	889.7840	889.9094	889.9101	889.9169	889.9230	889.9130
0.7	1019.2579	984.8884	985.7016	985.7002	985.6932	985.7191	985.7254
0.8	1130.3803	1123.1665	1121.2531	1121.2539	1121.2792	1121.2813	1121.2882
0.9	1297.6923	1347.1099	1342.4779	1342.4507	1342.4532	1342.4596	1342.5075
0.95	1424.8485	1535.7407	1533.8969	1533.9950	1534.0185	1533.9955	1534.0297
0.99	1604.9770	1833.5808	1846.6070	1846.6492	1846.6357	1846.6175	1846.6890
Evaluations	7	31	111	303	687	1375	2335

Table 5.12. Polynomial chaos expansion, sparse grid expansion levels 1–7: Response level values.

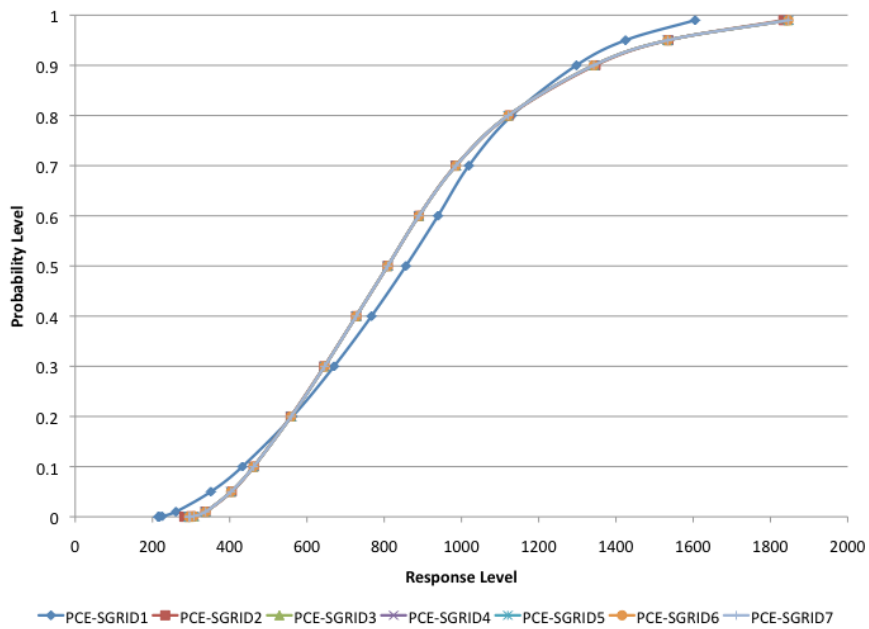

Figure 5.11. Polynomial chaos expansion, sparse grid expansion levels 1–7: Response level values.

L^2 Norms

The observed L^2 norms depicting the average differences seen in each evaluation a "nominal" simulation are shown in Table 5.13, with Figure 5.12 providing the plots of the CDFs for the seven expansion levels for the sparse grids, Figure 5.13 splitting out the first- through fourth-level plots, and Figure 5.14 showing the fifth- through seventh-level plots.

Figure 5.12 shows that the CDFs for the first- through fourth-level grid expansions vary quite noticeably. This is even more clear in Figure 5.13, which provides only these four cumulative distribution functions. However, as can be seen in Figure 5.14, the CDFs for these L^2 norms settle in relatively nicely by the fifth-level expansion, although there are still slight differences seen in the left tail of the functions. This behavior with Smolyak sparse grids mimics what we saw for the early sets of ensembles that utilized tensor product quadrature grids. It appears that a fifth-level expansion is needed for less smooth metrics such as L^2 and L^∞ norms, based on this result as well as the previous sets of ensembles.

In Table 5.13, we again observed estimated L^2 values that were less than zero, which should not be possible since the relative norms must have values greater than zero. In this set of ensembles, a sixth-level grid expansion was required before these negative values were not seen at our lowest probability levels. Also, recall that we only sampled the generated expansions 1,000,000 times due to limitations in DAKOTA with keeping so many samples in memory in order to predict responses at each probability level. This means that only one data point is used to generate the estimated

Probability Level	PCE-SGRID1	PCE-SGRID2	PCE-SGRID3	PCE-SGRID4	PCE-SGRID5	PCE-SGRID6	PCE-SGRID7
0.000001	0.0046	-0.1196	-0.0075	-0.0533	-0.0004	0.0018	0.0015
0.00001	0.0047	-0.1159	-0.0074	-0.0446	0.0004	0.0019	0.0016
0.0001	0.0048	-0.1100	-0.0067	-0.0225	0.0026	0.0023	0.0020
0.001	0.0055	-0.0921	-0.0040	-0.0020	0.0054	0.0033	0.0035
0.01	0.0076	-0.0446	0.0074	0.0043	0.0090	0.0065	0.0067
0.05	0.0144	0.0085	0.0163	0.0111	0.0133	0.0111	0.0111
0.1	0.0229	0.0211	0.0241	0.0156	0.0176	0.0145	0.0145
0.2	0.0398	0.0403	0.0348	0.0248	0.0261	0.0252	0.0252
0.3	0.0567	0.0524	0.0436	0.0370	0.0365	0.0377	0.0377
0.4	0.0736	0.0589	0.0535	0.0518	0.0501	0.0505	0.0505
0.5	0.0905	0.0673	0.0669	0.0678	0.0653	0.0654	0.0651
0.6	0.1075	0.0826	0.0842	0.0837	0.0818	0.0824	0.0823
0.7	0.1257	0.1061	0.1043	0.1002	0.1014	0.1014	0.1015
0.8	0.1483	0.1353	0.1247	0.1236	0.1256	0.1251	0.1249
0.9	0.1788	0.1621	0.1523	0.1572	0.1560	0.1564	0.1564
0.95	0.2056	0.1838	0.1801	0.1805	0.1802	0.1798	0.1801
0.99	0.2529	0.2191	0.2272	0.2239	0.2272	0.2259	0.2263
Evaluations	7	31	111	303	687	1375	2335

Table 5.13. Polynomial chaos expansion, sparse grid expansion levels 1–7: L^2 norm for response level values.

response value at the 10^{-6} probability level. Thus, it is probably more reasonable to note that a fifth-level expansion was required to eliminate these negative estimates, rather than a sixth-level expansion where the only negative estimated response was based on an insufficient number of samples.

L^∞ Norms

The results for the L^∞ norms, provided in Table 5.14, closely resemble those for the L^2 norms in that there are significant differences in the CDFs until we reach a fifth-level grid expansion; this can be seen in both Figures 5.15 and 5.16. At this point, as can be seen in Figure 5.17, the functions are quite close together. However, there are still slight differences in the left tails for the fifth-level expansion as compared to the rest of the higher-level expansions.

Similar to the previous polynomial chaos expansion ensembles using tensor product quadrature grids, we find few instances in Table 5.14 of negative estimated L^∞ norms, which in practice should not be seen. Another interesting observation from Table 5.14 is column PCE-SGRID4, where the estimated response values at the lower probability levels of the CDF's left tail are quite different from those in the third- and fifth-level expansions. It is not clear to us why this was the case, other

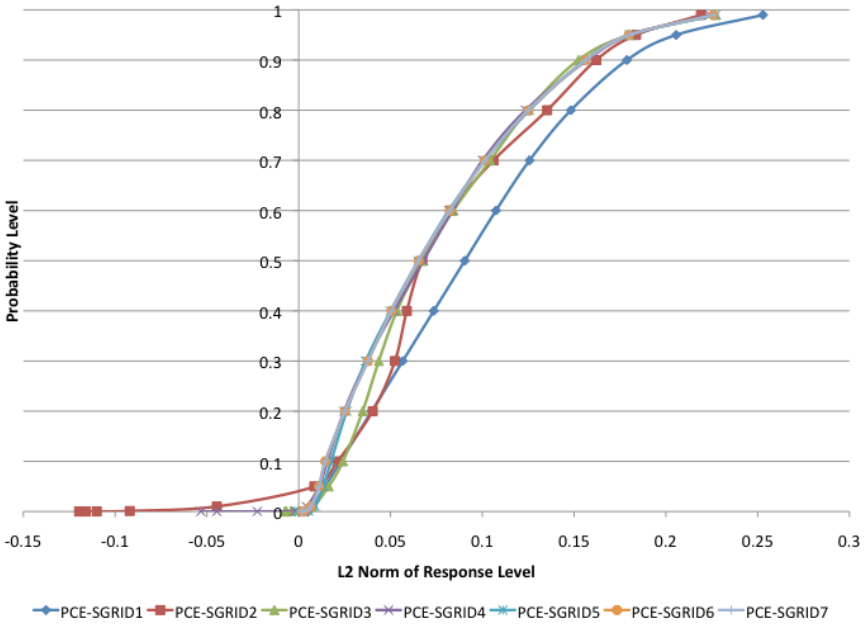

Figure 5.12. Polynomial chaos expansion, sparse grid expansion levels 1–7: L^2 norm for response level values.

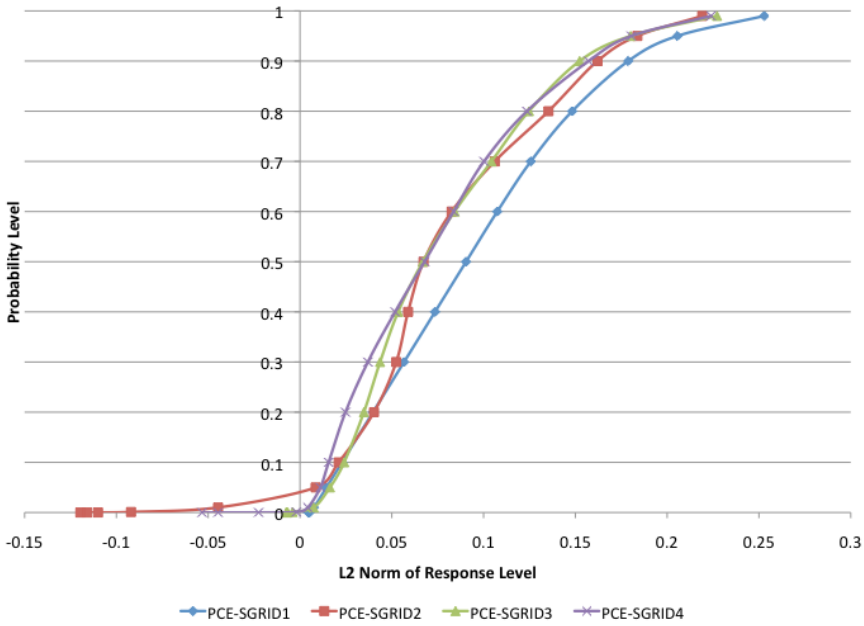

Figure 5.13. Polynomial chaos expansion, sparse grid expansion levels 1–4: L^2 norm for response level values.

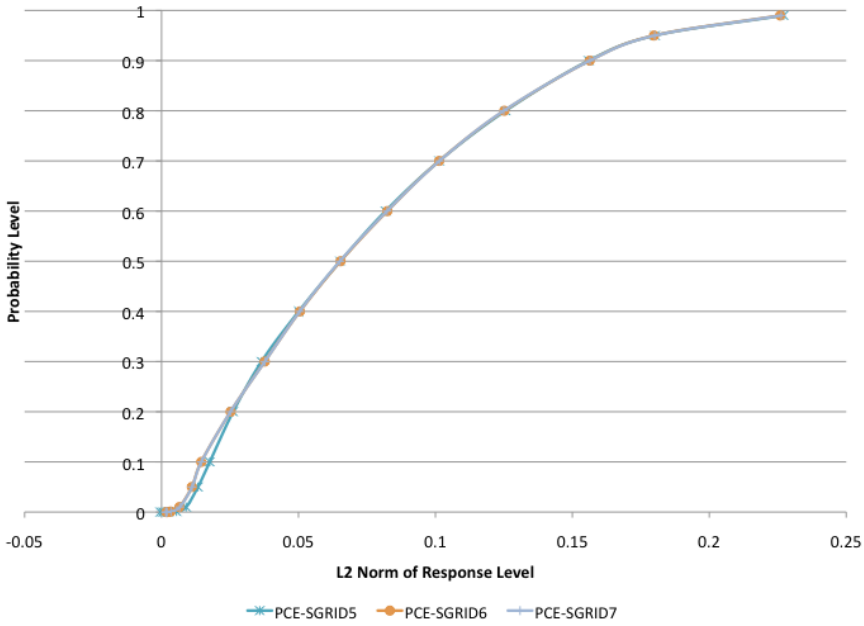

Figure 5.14. Polynomial chaos expansion, sparse grid expansion levels 5–7: L^2 norm for response level values.

than it being evidence that the fourth-level grid expansion was not sufficient.

Probability-level Convergence Plots

Figure 5.18 presents the logarithmic convergence plot showing differences in margin calculations between the seventh-level expansion of the Smolyak sparse grids — once again taken to represent a "true" solution — with each of the first- to sixth-level expansions, while Figure 5.19 plots the same information against the number of DAKOTA evaluations performed for each expansion level. These figures are based on calculations offered in Table 5.15.

Perhaps the most interesting characteristic of Figures 5.18 and 5.19 is the oscillation that occurs in the convergence plots beginning at the fourth-level expansion. This suggests that the DAKOTA ensembles have hit a point of diminishing returns where the estimates provided by the higher-level PCE ensembles are not improved in accuracy from the lower-level ensembles. In this particular case, given that the fourth-level ensemble's estimates are less accurate than the third-level ensemble's estimates (as measured by the seventh-level ensemble), a third-level ensemble is likely sufficient — especially from a cost-benefit criterion.

Probability Level	PCE-SGRID1	PCE-SGRID2	PCE-SGRID3	PCE-SGRID4	PCE-SGRID5	PCE-SGRID6	PCE-SGRID7
0.000001	0.0757	-0.0499	0.0411	0.0039	0.0384	0.0412	0.0411
0.00001	0.0758	-0.0441	0.0420	0.0108	0.0387	0.0415	0.0412
0.0001	0.0762	-0.0352	0.0433	0.0306	0.0396	0.0423	0.0418
0.001	0.0775	-0.0104	0.0441	0.0482	0.0424	0.0453	0.0454
0.01	0.0815	0.0444	0.0492	0.0551	0.0493	0.0525	0.0531
0.05	0.0891	0.0689	0.0628	0.0644	0.0608	0.0615	0.0608
0.1	0.0968	0.0787	0.0710	0.0718	0.0682	0.0693	0.0686
0.2	0.1122	0.0924	0.0843	0.0846	0.0834	0.0869	0.0866
0.3	0.1276	0.1010	0.0980	0.1010	0.1027	0.1037	0.1039
0.4	0.1431	0.1115	0.1140	0.1189	0.1225	0.1204	0.1206
0.5	0.1600	0.1264	0.1340	0.1385	0.1413	0.1388	0.1389
0.6	0.1788	0.1474	0.1561	0.1595	0.1590	0.1591	0.1589
0.7	0.1987	0.1776	0.1794	0.1822	0.1802	0.1809	0.1807
0.8	0.2195	0.2140	0.2065	0.2082	0.2055	0.2057	0.2061
0.9	0.2475	0.2529	0.2496	0.2480	0.2511	0.2509	0.2509
0.95	0.2747	0.2778	0.2837	0.2805	0.2820	0.2813	0.2820
0.99	0.3265	0.3349	0.3268	0.3225	0.3202	0.3213	0.3207
Evaluations	7	31	111	303	687	1375	2335

Table 5.14. Polynomial chaos expansion, sparse grid expansion levels 1–7: L^∞ norm for response level values.

Expansion Level	DAKOTA Evaluations	Response at 10^{-5} Probability Level	Difference from 7th-level Response	Square of the Difference
1	7	215.45274	80.14093221	6422.569015
2	31	283.9525515	11.64112077	135.5156928
3	111	295.4558495	0.13782275	0.01899511
4	303	295.8746728	-0.28100052	0.078961292
5	687	295.5437537	0.04991855	0.002491862
6	1375	295.6949642	-0.101292	0.010260069
7	2335	295.5936722	0.0	0.0

Table 5.15. Convergence of estimated response to a "true" solution at the seventh-level sparse grid expansion using PCE.

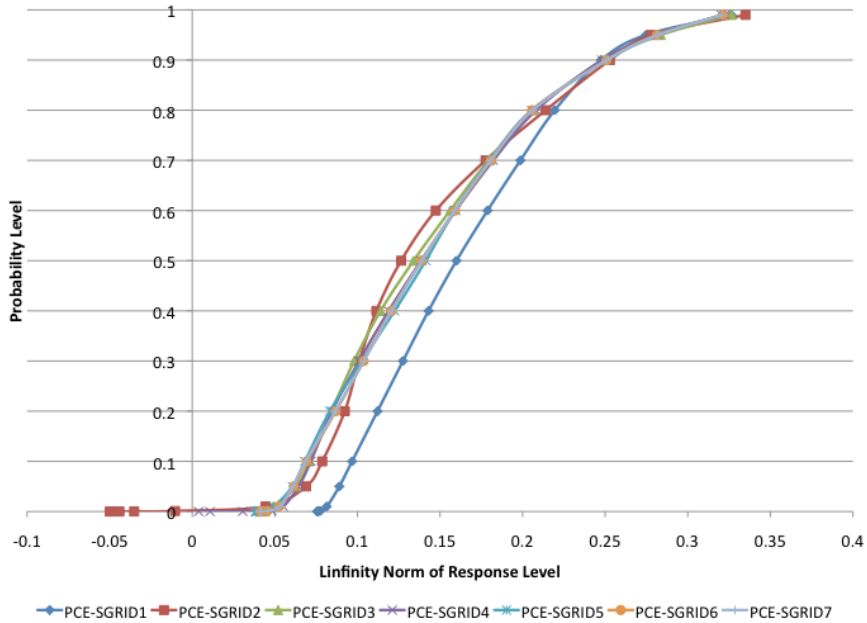

Figure 5.15. Polynomial chaos expansion, sparse grid expansion levels 1–7: L^{∞} norm for response level values.

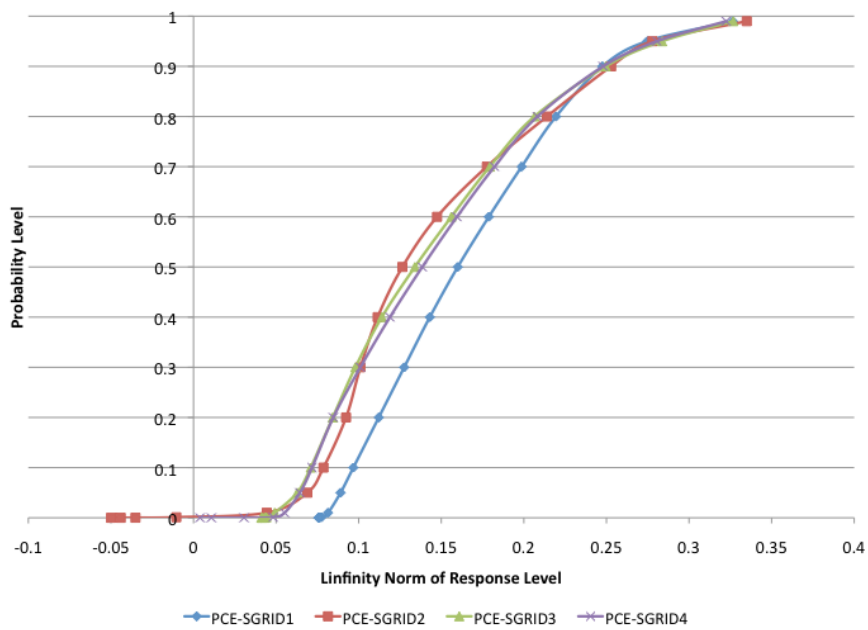

Figure 5.16. Polynomial chaos expansion, sparse grid expansion levels 1–4: L^{∞} norm for response level values.

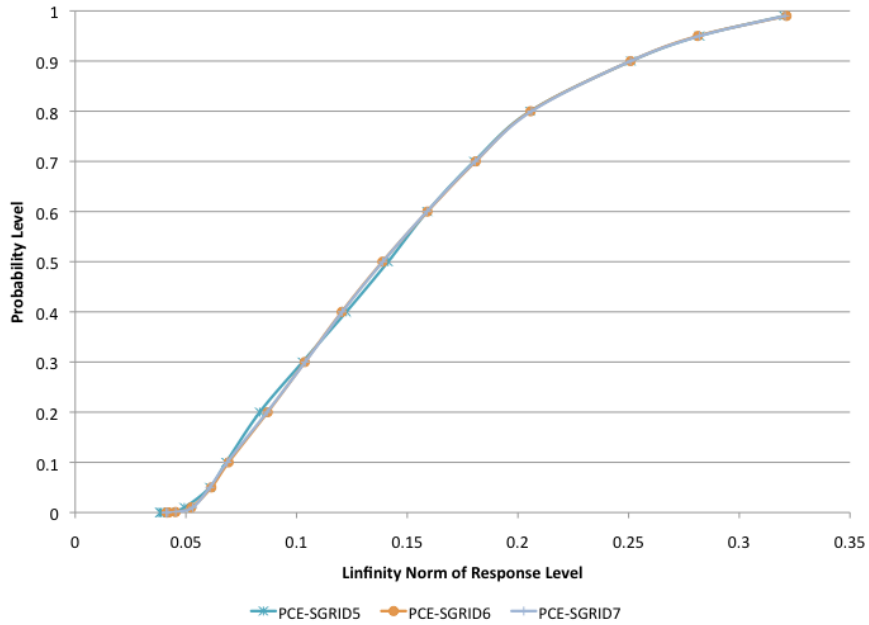

Figure 5.17. Polynomial chaos expansion, sparse grid expansion levels 5–7: L^{∞} norm for response level values.

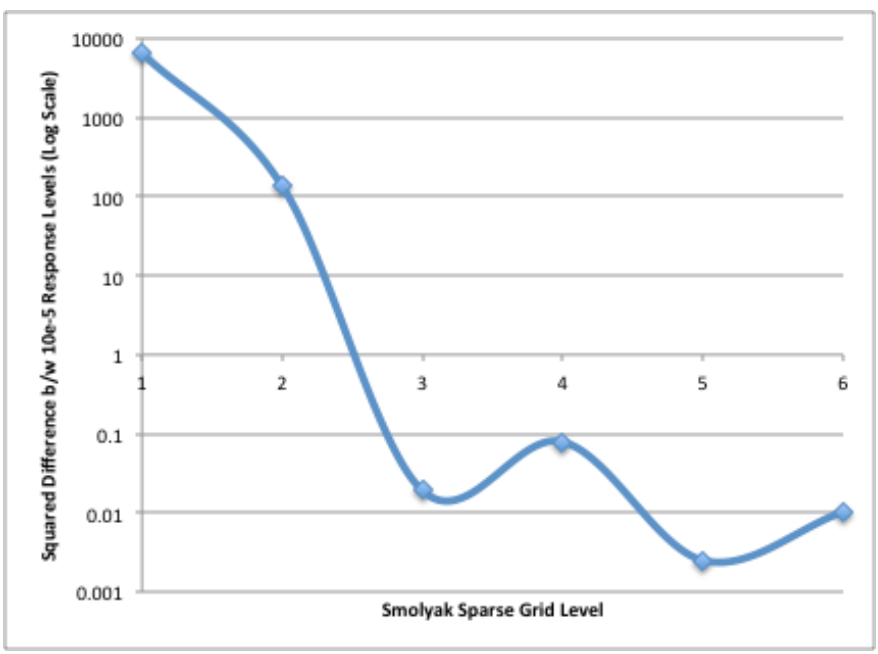

Figure 5.18. Convergence of the calculated margins from PCE as the sparse grid level increases.

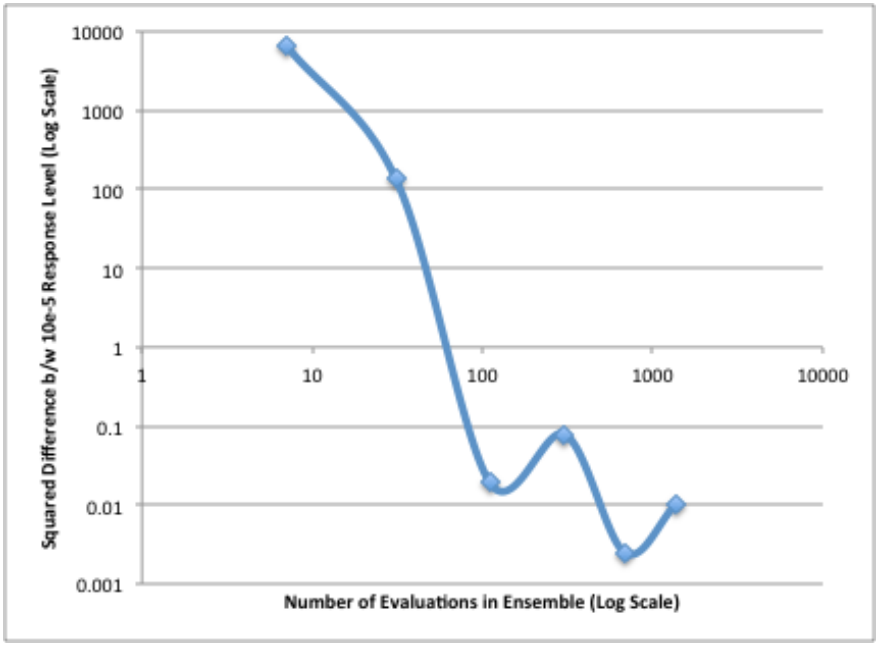

Figure 5.19. Convergence of the calculated margins from PCE, based on sparse grids, as the number of DAKOTA evaluations increases.

Stochastic Collocation: Smolyak Sparse Grids

The final set of ensembles we present are those driven by stochastic collocation using Smolyak sparse grids. The observed number of evaluations to generate the basis functions in these ensembles are provided in Table 5.16. These numbers are the same as those seen in Table 5.11 showing the evaluations using polynomial chaos expansion driven by Smolyak sparse grids.

Expansion Level	Observed Evaluations
1	7
2	31
3	111
4	303
5	687
6	1375
7	2335

Table 5.16. Observed samples for SC ensembles using Smolyak sparse grids.

Probability Level	SC-SGRID1	SC-SGRID2	SC-SGRID3	SC-SGRID4	SC-SGRID5	SC-SGRID6	SC-SGRID7
0.000001	215.1997	282.0272	294.5557	294.7355	294.3713	294.4860	294.4860
0.00001	215.4527	283.6524	295.6558	295.7960	295.5427	295.6197	295.6197
0.0001	216.8676	287.4172	298.1725	298.3573	298.1440	298.2376	298.2376
0.001	224.7357	299.7396	307.2922	307.2480	307.2669	307.2667	307.2667
0.01	260.7595	337.1838	337.4512	337.4404	337.4778	337.4576	337.4576
0.05	351.2758	405.2556	403.2983	403.3370	403.2968	403.2940	403.2940
0.1	432.8379	462.1519	461.9793	462.0225	462.0265	462.0376	462.0376
0.2	562.5460	557.9495	559.0601	559.1064	559.0883	559.1044	559.1044
0.3	670.3572	644.7091	645.6418	645.6468	645.6724	645.6817	645.6817
0.4	767.1856	727.5014	727.5031	727.5075	727.5117	727.5266	727.5266
0.5	856.2118	809.0456	808.9890	809.0090	809.0131	809.0023	809.0023
0.6	939.3272	889.7782	889.9113	889.9093	889.9239	889.9229	889.9229
0.7	1019.2579	984.8661	985.6989	985.6995	985.6995	985.7096	985.7096
0.8	1130.3803	1123.1674	1121.2601	1121.2550	1121.2817	1121.2884	1121.2884
0.9	1297.6923	1347.0931	1342.4535	1342.4377	1342.4662	1342.4729	1342.4729
0.95	1424.8485	1535.7777	1533.9106	1533.9889	1534.0031	1533.9975	1533.9975
0.99	1604.9770	1833.6019	1846.6140	1846.6518	1846.6336	1846.6646	1846.6646
Evaluations	7	31	111	303	687	1375	2335

Table 5.17. Stochastic collocation, sparse grid expansion levels 1–7: Response level values.

Response Level

Table 5.17 shows the estimated responses for the specified probability levels that make up the CDFs plotted in Figure 5.20. The estimated responses for the first-level SC grid expansion are the same as those for the first-level PCE grid expansion. The estimated responses using SC for the remaining grid expansion levels very closely resemble those using polynomial chaos expansion. Thus, using Smolyak sparse grids, the same trends we earlier observed for PCE hold here for SC. Namely, estimations for the margin converged by the third-level expansion. Although we would expect this convergence to take longer with a more complicated model where a larger number of uncertain variables might be studied, the early convergence of these CDFs is encouraging because it suggests we may be able to get accurate estimates from the UQ methods using relatively low-level studies that are not exorbitantly expensive.

L^2 Norms

The results from this set of ensembles with respect to the L^2 norms, summarized in Table 5.18 and Figure 5.21, also closely resemble those from previous ensembles. For the first four expansion

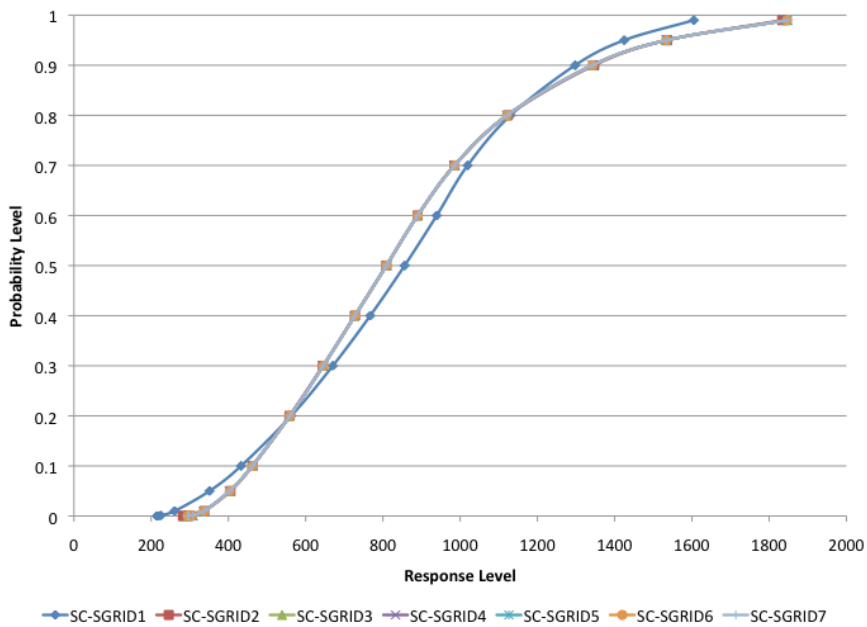

Figure 5.20. Polynomial chaos expansion, sparse grid expansion levels 1–7: Response level values.

levels for the grids, the CDFs generated from stochastic collocation with Smolyak sparse grids do not converge particularly well, as can be seen in Figure 5.22. However, by the fifth-level expansion, the CDFs converge quite well except at the left-most tails, which can be seen in Figure 5.23.

Table 5.18 also shows negative estimates for the L^2 norms for the first- through fourth-level grid expansions. These negative estimates disappear by the fifth-level expansion, which is approximately the same point where they no longer appeared in the previous sets of ensembles.

L^∞ Norms

The L^∞ norms for this set of ensembles, shown in Table 5.19 and Figure 5.24, also resemble those in previous ensembles. Negative estimates are seen in Table 5.19 in the column for the second-level grid expansion (SC-SGRID2), but no where else in the table. It can be seen in Figure 5.25 that the cumulative distribution functions for the first four expansion levels vary noticeably, but converge quite well by the fifth-level expansion (see Figure 5.26).

Probability-level Convergence Plots

The convergence plots for the stochastic collocation ensembles driven by Smolyak sparse grids, presented in Figures 5.27 and 5.28 show a similar trend as seen for the sparse-grid-based PCE

Probability Level	**SC-SGRID1**	**SC-SGRID2**	**SC-SGRID3**	**SC-SGRID4**	**SC-SGRID5**	**SC-SGRID6**	**SC-SGRID7**
0.000001	0.0046	-0.1286	-0.0058	-0.0355	0.0039	0.0050	0.0051
0.00001	0.0047	-0.1239	-0.0054	-0.0295	0.0045	0.0051	0.0051
0.0001	0.0048	-0.1171	-0.0041	-0.0145	0.0059	0.0051	0.0054
0.001	0.0055	-0.0947	-0.0011	0.0058	0.0067	0.0054	0.0058
0.01	0.0076	-0.0438	0.0080	0.0081	0.0095	0.0076	0.0074
0.05	0.0144	0.0107	0.0161	0.0126	0.0135	0.0111	0.0112
0.1	0.0229	0.0222	0.0238	0.0163	0.0175	0.0144	0.0145
0.2	0.0398	0.0400	0.0351	0.0245	0.0256	0.0249	0.0249
0.3	0.0567	0.0517	0.0437	0.0361	0.0365	0.0376	0.0377
0.4	0.0736	0.0581	0.0535	0.0507	0.0502	0.0508	0.0507
0.5	0.0905	0.0672	0.0665	0.0672	0.0655	0.0653	0.0651
0.6	0.1075	0.0830	0.0839	0.0842	0.0820	0.0823	0.0822
0.7	0.1257	0.1061	0.1044	0.1011	0.1012	0.1017	0.1016
0.8	0.1483	0.1347	0.1252	0.1236	0.1252	0.1251	0.1250
0.9	0.1788	0.1623	0.1523	0.1566	0.1562	0.1563	0.1564
0.95	0.2056	0.1839	0.1793	0.1800	0.1802	0.1799	0.1799
0.99	0.2529	0.2185	0.2260	0.2256	0.2264	0.2261	0.2261
Evaluations	7	31	111	303	687	1375	2335

Table 5.18. Stochastic collocation, sparse grid expansion levels 1–7: L^2 norm for response level values.

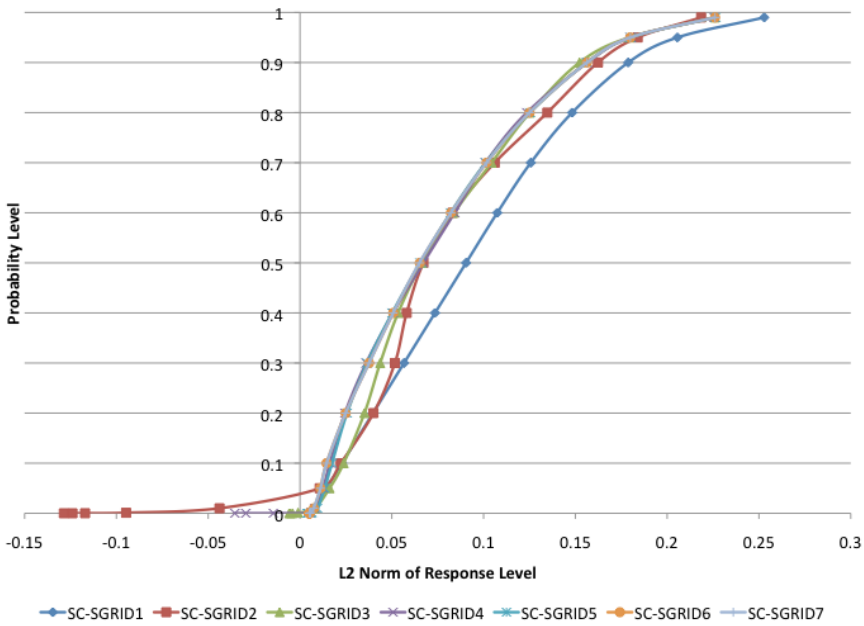

Figure 5.21. Stochastic collocation, sparse grid expansion levels 1–7: L^2 norm for response level values.

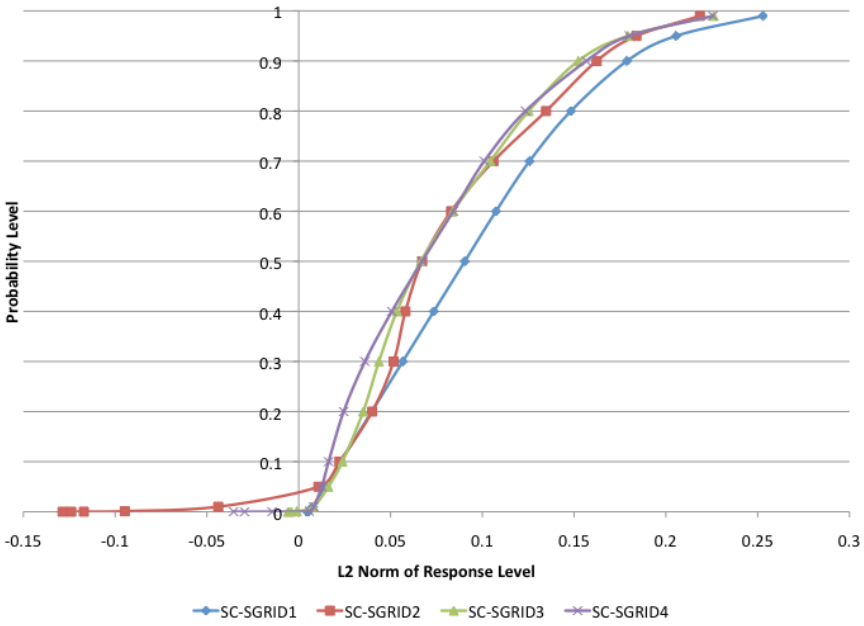

Figure 5.22. Stochastic collocation, sparse grid expansion levels 1–4: L^2 norm for response level values.

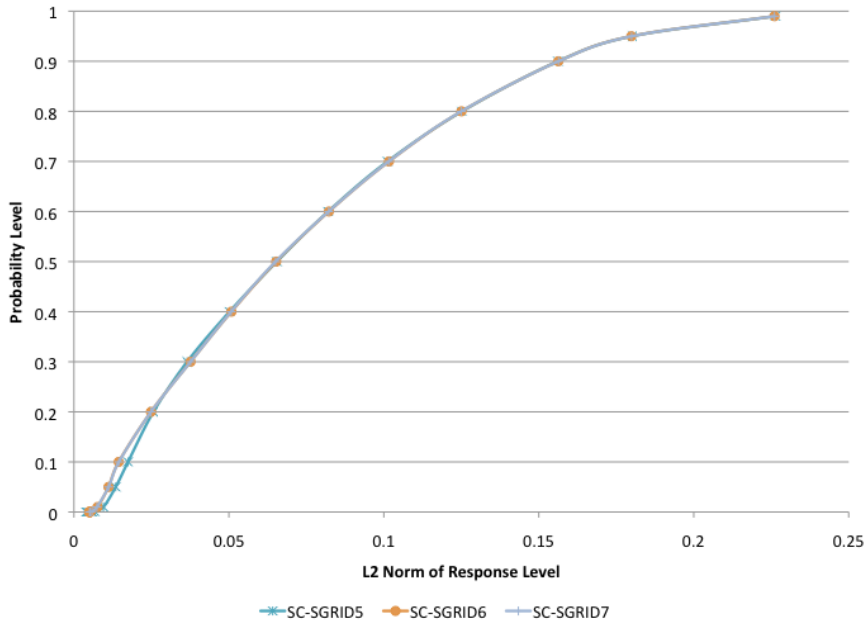

Figure 5.23. Stochastic collocation, sparse grid expansion levels 5–7: L^2 norm for response level values.

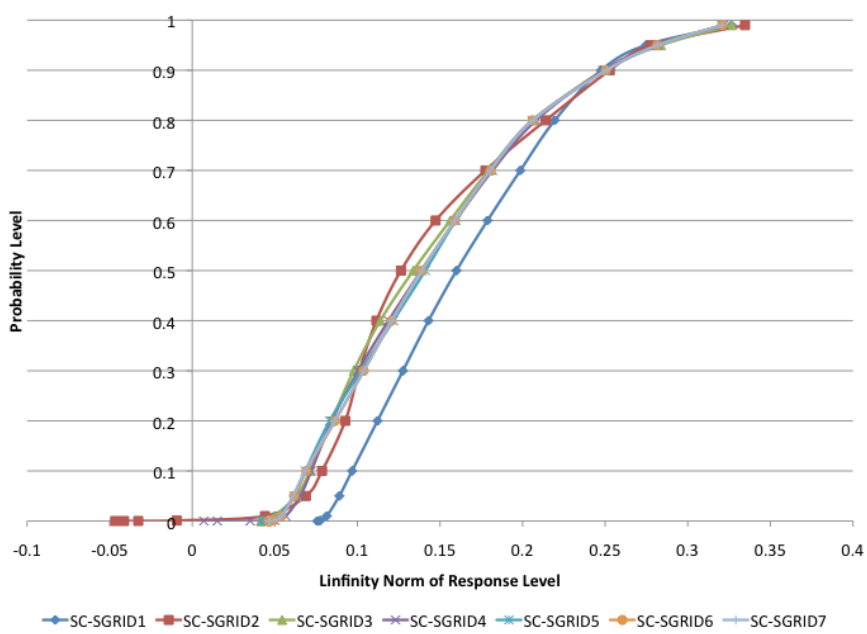

Figure 5.24. Stochastic collocation, sparse grid expansion levels 1–7: L^∞ norm for response level values.

Probability Level	**SC-SGRID1**	**SC-SGRID2**	**SC-SGRID3**	**SC-SGRID4**	**SC-SGRID5**	**SC-SGRID6**	**SC-SGRID7**
0.000001	0.0757	-0.0467	0.0416	0.0070	0.0423	0.0463	0.0464
0.00001	0.0758	-0.0412	0.0421	0.0152	0.0426	0.0466	0.0468
0.0001	0.0762	-0.0327	0.0431	0.0351	0.0434	0.0473	0.0472
0.001	0.0775	-0.0094	0.0454	0.0500	0.0457	0.0490	0.0496
0.01	0.0815	0.0440	0.0497	0.0566	0.0517	0.0539	0.0541
0.05	0.0891	0.0688	0.0630	0.0653	0.0618	0.0617	0.0611
0.1	0.0968	0.0786	0.0714	0.0721	0.0685	0.0691	0.0684
0.2	0.1122	0.0926	0.0844	0.0844	0.0832	0.0865	0.0862
0.3	0.1276	0.1012	0.0977	0.1005	0.1021	0.1037	0.1040
0.4	0.1431	0.1115	0.1138	0.1188	0.1217	0.1206	0.1208
0.5	0.1600	0.1263	0.1338	0.1386	0.1411	0.1388	0.1389
0.6	0.1788	0.1473	0.1562	0.1594	0.1595	0.1591	0.1590
0.7	0.1987	0.1775	0.1795	0.1818	0.1803	0.1806	0.1804
0.8	0.2195	0.2140	0.2065	0.2081	0.2059	0.2060	0.2062
0.9	0.2475	0.2528	0.2494	0.2488	0.2506	0.2504	0.2508
0.95	0.2747	0.2775	0.2837	0.2807	0.2820	0.2814	0.2820
0.99	0.3265	0.3347	0.3269	0.3218	0.3206	0.3209	0.3208
Evaluations	7	31	111	303	687	1375	2335

Table 5.19. Stochastic collocation, sparse grid expansion levels 1–7: L^∞ norm for response level values.

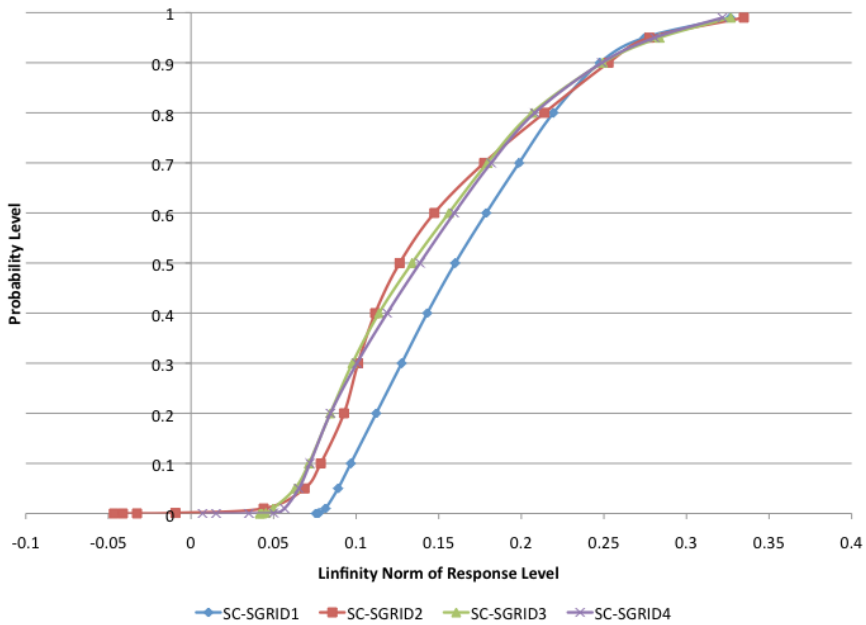

Figure 5.25. Stochastic collocation, sparse grid expansion levels 1–4: L^∞ norm for response level values.

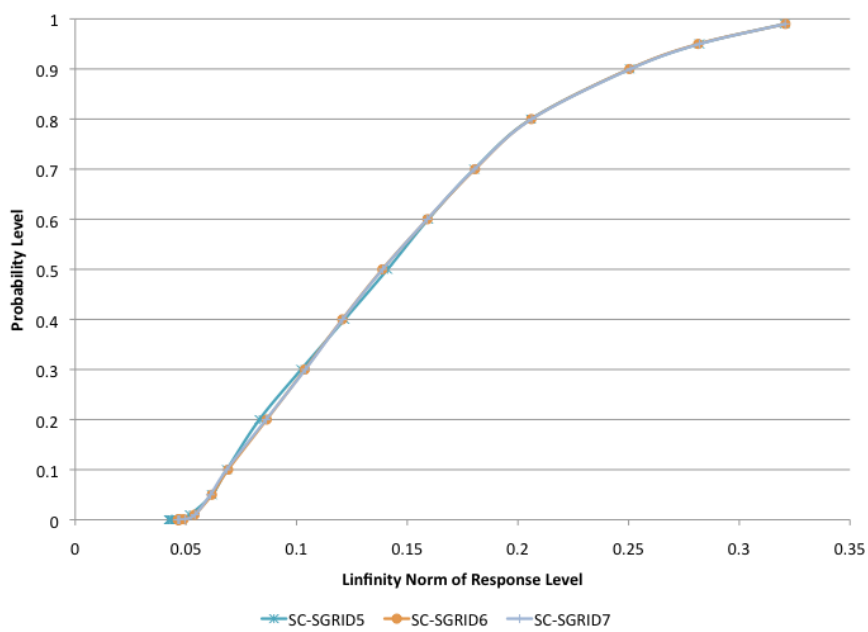

Figure 5.26. Stochastic collocation, sparse grid expansion levels 5–7: L^∞ norm for response level values.

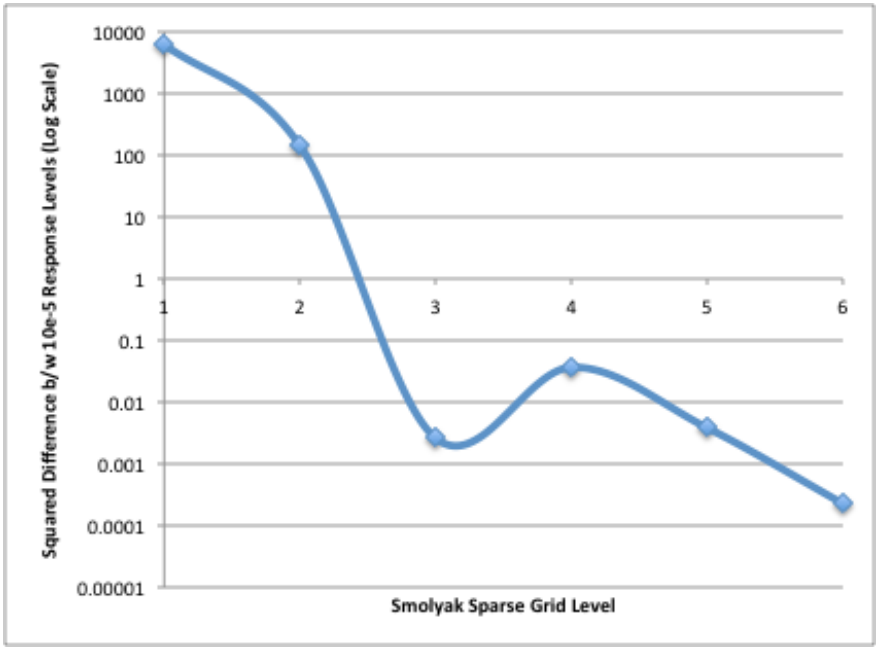

Figure 5.27. Convergence of the calculated margins from SC as the sparse grid level increases.

ensembles: it is not clear that ensembles beyond a third-level expansion add precision to the estimates made concerning the margin. The fact that this was seen for both types of stochastic expansion DAKOTA ensembles, and not for the quadrature-grid-based ensembles, is illuminating and suggests that sparse grid ensembles may achieve measurable convergence around third-level grid expansions. Of course, the ensembles explored in this work considered only three uncertain variables, and it is likely that higher-dimension ensembles with additional uncertain variables may require a higher grid expansions to achieve convergence.

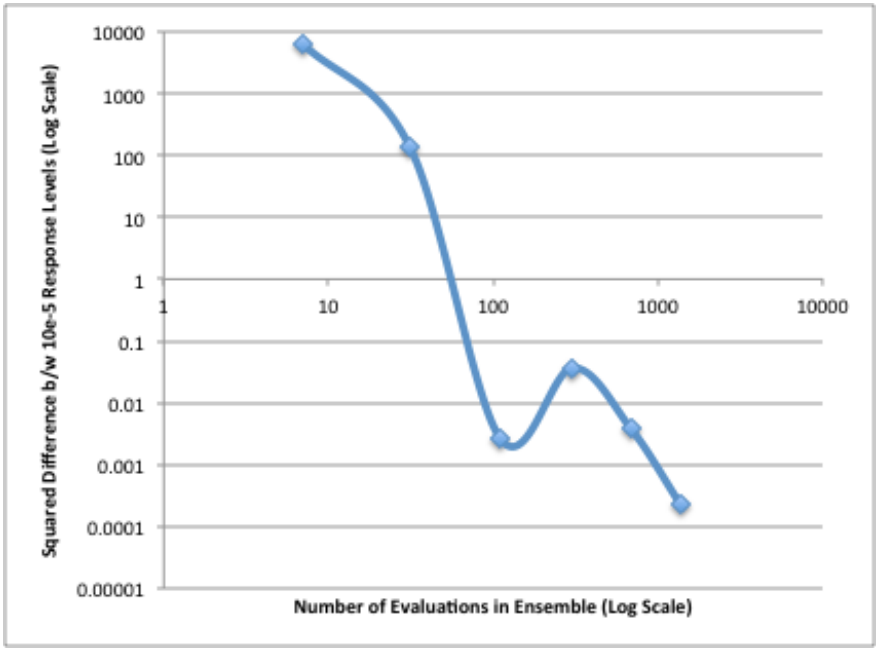

Figure 5.28. Convergence of the calculated margins from SC, based on sparse grids, as the number of DAKOTA evaluations increases.

Expansion Level	DAKOTA Evaluations	Response at 10^{-5} Probability Level	Difference from 7th-level Response	Square of the Difference
1	7	215.45274	80.15192573	6422.569015
2	31	283.6524202	11.9522456	135.5156928
3	111	295.6557654	-0.05109963	0.01899511
4	303	295.7960191	-0.19135333	0.078961292
5	687	295.5427226	0.06194315	0.002491862
6	1375	295.6197297	-0.01506397	0.010260069
7	2335	295.6046658	0.0	0.0

Table 5.20. Convergence of estimated response to a "true" solution at the seventh-level sparse grid expansion using SC.

This page intentionally left blank.

Chapter 6

Concluding Remarks

This report explored the use of the DAKOTA toolkit's polynomial chaos expansion and stochastic collocation methods, using both tensor product quadrature and Smolyak sparse grids, for a model with a mesh containing over 150,000 tetrahedral finite elements. From these four sets of UQ ensembles, a number of observations can be drawn and applied to future uses of these DAKOTA methodologies.

Probability Levels

First, although our response metric was arbitrary in nature, we consistently saw noticeable fluctuations in the estimated response levels for increasingly lower probability levels. For example, the PCE ensembles using tensor product quadrature grids with a fourth-level expansion saw the response level drop from 337.9673 at probability level 0.01 (10^{-2}) to 290.6769 at probability level 10^{-6}. Similar observations can be seen for the other methodologies. Because we observed non-trivial differences in the predicted response levels for probability levels such as 10^{-2} versus 10^{-6}, we believe there is potential benefit to considering these ever-smaller probability levels in other scenarios for high-consequence systems (if it makes sense to predict responses at such small probability levels for the problem in question).

We also observed that there is less change in the estimated responses around the middle of the CDF curves as the grid expansion level of the ensembles increase and more observations are considered, while there are notable changes in the response at both tails of the CDF; we can say this especially for the left tails where we focused our attention in this work. Investigators can use these lessons to help determine how rapidly areas of a CDF that are of interest to them will converge for ensembles with higher-level grid expansions. In many (but presumably not all) cases, investigators will have interest in their responses at the tails of a CDF, and more expensive ensembles will likely be needed for the cumulative distribution functions in those areas to converge.

Convergence of Cumulative Distribution Functions

On the subject of convergence of cumulative distribution functions, we noticed that there was often little change in the CDFs of our thermal race margin response for ensembles with third-

level grid expansions and onward. As such, it may not be necessary to conduct ensembles larger than third- or fourth-level expansions on larger problem sizes. Further analysis on this topic is warranted.

This work also found evidence that the needed grid expansion level for CDF convergence may also depend on the region of a cumulative distribution function that is of interest to investigators. As mentioned earlier, for the problem we investigated, we found that convergence was generally slowest in the tails of the CDFs, especially the left tail. However, we caution that we did not invest the effort to conduct statistical analyses to quantitatively measure the differences in these CDFs. A rigorous quantitative analysis of these functions was beyond the scope of this work. Instead, the judgments we have made with respect to regions of the CDF that converged were qualitative, and therefore limited, in nature.

Finally, we observed that convergence of CDFs was generally slower for the L^2 norms and L^∞ norms than for the response itself. We found that for these norms, the grid expansion level may need to be taken as high as a fifth-level expansion before real convergence is seen in the cumulative distribution functions. This was observed for both polynomial chaos expansion and stochastic collocation methods, using both tensor quadrature grids and Smolyak sparse grids, and contrasts with margin estimations, where third- and fourth-level expansions were generally sufficient.

Tensor Product Quadrature Grids versus Smolyak Sparse Grids

We observed that the growth of the sparse grids was not necessarily slower than for tensor product quadrature grids in our ensembles. The reason for this was that we only had three uncertain variables available to study due to the limited nature of our model. We expect sparse grids will be much more efficient than quadrature grids for the larger ensembles with many more uncertain variables. However, it is notable that sparse grids may not necessarily be the most cost-effective choice for all cases.

We also found little difference between the convergence of the CDFs generated by either the PCE or SC methods, for either quadrature grids or sparse grids. In fact, the estimated response levels between methods were generally either the same, or very close to the same for the problem we studied. The goal would be to utilize the methodology whose expected growth rates are anticipated to be the most cost effective for the problem under study.

Improvements to DAKOTA

One possible improvement to DAKOTA that became apparent to us throughout the course of this work was the ability to generate probability density functions. Although cumulative distribution functions are often more useful in terms of things like probability of exceeding threshold values, in work we have in mind for the future it would be useful to have the ability to generate probability density functions at given locations with specified parameters, as well as projecting

data onto the basis functions. The DAKOTA developers have acted on our request for this capability, and probability distribution function output will be available for both PCE and SC methods in the next release of DAKOTA.

Another potential improvement to DAKOTA would be the ability hold more samples in memory when predicting response levels for specified probability levels. We have also been recently advised that there will be modest improvements to this capability in the next release of DAKOTA.

Global Reliability Methods

Finally, our next anticipated step in future work will be to explore the use of DAKOTA's global reliability methods for uncertainty quantification. One common trend throughout the results presented in this work is that both stochastic expansion methods tended to converge better and more quickly in the middle of the CDF than at the tails (particularly the left tails where we focused our attention). We first focused on stochastic expansion methods because of their potential synergy with multi-physics applications due to their ability to express final solutions as random processes. Now that we better understand the benefits and limitations of these methods, we will explore global reliability methods, which are expected to be more efficient at predicting responses at the tails of the probability distributions.

This page intentionally left blank.

References

[1] B.A. Adams, K.R. Dalbey, M.S. Eldred, D.M. Gay, L.P. Swiler, W.J. Bohnhoff, J.P. Eddy, K. Haskell, and P.D. Hough. DAKOTA, a multilevel parallel object-oriented framework for design and optimization, parameter estimation, uncertainty quantification, and sensitivity analysis. Technical Report SAND2010–2183, Sandia National Laboratories, Albuquerque, NM, December 2009.

[2] S.W. Bova, K.D. Copps, and C.K. Newman. CALORE: A computational heat transfer program, Volume 1: Theory manual. Technical Report SAND2005–6083P, Sandia National Laboratories, Albuquerque, NM, October 2005.

[3] G.A. Gray, R.G. Hills, and P.D. Hough. Some guidance on preparing validation plans for the DART full system models. Technical Report SAND2009–1331, Sandia National Laboratories, Livermore, CA, August 2009.

[4] J.E. Marsden and M.J. Hoffman. *Basic Complex Analysis*. W.H. Freeman and Company, New York, N.Y., U.S.A., 3rd edition, 1999.

[5] W.L. Oberkampf, T.G. Trucano, and M. Pilch. Verification, validation, and predictive capability in computational engineering and physics. Technical Report SAND2003–3769, Sandia National Laboratories, Albuquerque, NM, 2003.

[6] Sandia National Laboratories. CUBIT. http://cubit.sandia.gov/; Last accessed: August 10, 2011.

[7] Sandia National Laboratories. Sierra Mechanics. http://sierra.sandia.gov/; Last accessed: August 10, 2011.

[8] Sandia National Laboratories. The DAKOTA Project. http://dakota.sandia.gov/; Last accessed: August 10, 2011.

[9] L.A. Schoof and V.R. Yarberry. EXODUS II: A finite element data model. Technical Report SAND92-2137, Sandia National Laboratories, Albuquerque, NM, September 1992.

DISTRIBUTION:

1	MS	1318	Brian Adams, 1411
1	MS	1318	Michael Eldred, 1411
1	MS	9013	James Lauffer, 8231
1	MS	9042	Timothy Kostka, 8249
1	MS	9042	Arthur Ortega, 8365
3	MS	9409	Jeremy Templeton, 8365
1	MS	9151	Genetha Gray, 8954
1	MS	9151	Patty Hough, 8954
1	MS	9159	Jerry McNeish, 8954
3	MS	9155	Joseph Ruthruff, 8954
1	MS	0899	RIM — Reports Management, 9532

Sandia National Laboratories