SANDIA REPORT

SAND2015-8620
Unlimited Release
Printed October 2015

Stochastic Richardson Extrapolation Based Numerical Error Estimation for Kinetic Plasma Simulations

Gregg A. Radtke, Keith L. Cartwright, and Lawrence C. Musson

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.

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

SAND2015-8620
Unlimited Release
Printed October 2015

Stochastic Richardson Extrapolation Based Numerical Error Estimation for Kinetic Plasma Simulations

Gregg A. Radtke
Electromagnetic Theory
Sandia National Laboratories
P.O. Box 5800
Albuquerque, NM 87185-1152
garadtk@sandia.gov

Keith L. Cartwright
Electromagnetic Theory
Sandia National Laboratories
P.O. Box 5800
Albuquerque, NM 87185-1449
klcartw@sandia.gov

Lawrence C. Musson
Electrical Models and Simulation
Sandia National Laboratories
P.O. Box 5800
Albuquerque, NM 87185-1177
lcmusso@sandia.gov

Abstract

We present a numerical error estimation technique specifically formulated to deal with stochastic code output with multiple discretization parameters. This method is based on multiple fits to an error model with arbitrary convergence rates and cross-coupling terms, performed using nonlinear optimization. The fitting approach varies by the type of residual norm which influences the importance of outliers, and weights which influences the relative importance of data points in the coarse and refined regions of discretization parameter space. To account for the influence of stochastic noise, these fits are performed on multiple bootstrap values based on the underlying response data set. Using an automated discretization domain selection scheme, the fits are performed on a series of reduced sets of discretization levels in order to find an optimal fully-converged result estimate in the minimum variance sense; this automated approach enables straightforward analysis of multiple quantities of interest and/or time and spatially-dependent response data. The overall numerical error analysis method is useful for verification and validation problems for stochastic simulation methods and forms a key component in the overall uncertainty quantification process. The method was demonstrated for steady and unsteady electron diode problems simulated using a particle-in-cell kinetic plasma code, demonstrating excellent results.

Acknowledgment

The authors would like to thank Bill Rider (1446), Rich Hills (1544), and Chris Moore (1352) for helpful discussions and critical feedback during the development of the method presented in this report.

Contents

Nomenclature **12**

1 Introduction **15**

2 StREEQ Numerical Error Estimation Method **19**

Response Data 19

Bootstrap Sampling 20

Variance Model Fitting 21

Discretization Error Model 22

Fitting Models 23

Error Estimation 26

Discretization Parameter Domain Reduction 27

Credibility Assessment 30

3 Application to Verification Problems **31**

Engineered Data Set Example 31

Stationary Diode Verification Problem 41

Code Verification 42

Solution Verification 49

Periodic Diode Verification Problem 50

4 Summary and Future Directions **53**

References **54**

Appendix

A Vlassov-Poisson-Child-Langmuir Diode Solution Derivation 59

List of Figures

2.1 An example sequence of reduced domains for a $D = 2$ dimensional discretization parameter space. The core region Ω^* boundary is shown in gray, while subsequent reduced domains $(\Omega_0, \Omega_1, \Omega_2)$ are shown in blue, red, and green respectively. Also shown, are points of maximum $|\Xi|$ for each StREEQ analysis. 29

3.1 Plot of of generated data points for $N_1 = 5$ (symbols) compared to the first four terms of Equation 3.1 (lines). 32

3.2 Fit coefficient distributions for full data set (blue) and reduced domains (green, red) for normally distributed noise with $N_1 = 5$. Solid lines are raw β values, while the dashed lines for the $\tilde{\beta}$ coefficients include the lack-of-fit error corrections (see Equations 2.14–2.16). The horizontal error bars are 95% confidence intervals and vertical lines indicate exact values. The dots on the left side of the β_0 plot are the response data used to perform the fits, where the blue and green dots were eliminated after the Ω_0 and Ω_1 fits respectively. 34

3.3 Residuals for normally distributed noise with $N_1 = 5$ by regression model: **(a)** full data set Ω_0 and **(b)** and reduced domain Ω_1. Abscissa values represent different discretization levels, while repeated points are obtained from fits for multiple bootstrap samples. 35

3.4 F-test probability value distributions for full data set (blue) and reduced domains (green, red) for normally distributed noise with $N_1 = 5$. 36

3.5 Estimated converged result distribution for normally distributed noise for various numbers of replications N_1. 37

3.6 Alternate noise distributions. 38

3.7 Estimated converged result distribution for various noise distribution choices with $N_1 = 100$. 39

3.8 Estimated converged result distribution with heteroscedastic stochastic noise: full data set (blue), and reduced domains (green, red). The horizontal error bars are 95% confidence intervals and the vertical lines indicate exact values. The dots on the left side are the response data used to perform the fits, where the blue, dark green, light green, and brown dots were eliminated after the Ω_0, Ω_1, Ω_2, and Ω_3 fits respectively. 40

3.9 VPCL diode verification problem points in discretization space: full data set (blue) and minimum variance domain (red). 42

3.10 VPCL diode code verification problem StREEQ prediction for diode current (red) compared to the data set used for the fits for the minimum-variance discretization domain. Vertical line is the exact value. 43

3.11 Convergence rates for VPCL diode code verification problem: **(a)** mesh, **(b)** time step, and **(c)** MPW. Vertical black lines are the theoretical values. 44

3.12 Residuals for VPCL diode code verification problem. 45

3.13 Time step convergence for VPCL diode code verification problem. Error bars are 95% confidence intervals for the mean values (based on all 700 replications) and the horizontal black line is the exact value. 46

3.14 Time-converged VPCL diode code verification problem StREEQ prediction for diode current (red) compared to the data set used for the fits. 47

3.15 Residuals for the time-converged VPCL diode code verification problem. 48

3.16 VPCL diode solution verification problem StREEQ prediction for diode current (red) compared to the data set used for the fits. 49

3.17 Simulated discretization levels for the periodic electron diode verification problem. 51

3.18 Periodic electron diode anode current: comparison between exact result (red) and StREEQ predicted 95% confidence intervals (blue). 52

List of Tables

2.1 Fitting model L_p norm and weighting strategy $a(s)$ for the various fitting models $s = 1, 2, \ldots, 9$ 25

Nomenclature

a discretization parameter dependence exponent for residual weight

A surface area [m^2]

B number of bootstrap samples

D number of discretization parameters

e elementary charge [C] $(= 1.6022 \times 10^{-19})$

$\mathbf{E}$ electric field vector [V/m]

f velocity distribution function

G objective function

$\mathbf{J}$ current flux density vector [A/m^2]

L spatial domain size [m]

m particle mass [kg] $(= 9.1095 \times 10^{-31}$ for electrons)

n electron number density [m^{-3}]

N_j number of replications at discretization level j

N total number of replications $(= \sum_j N_j)$

M total number of discretization levels

q particle charge [C]

t time [s]

T temperature [V]

$\mathbf{v}$ velocity vector [m/s]

V volume [m^3]

w residual weights

W diagonal matrix of squared residual weights

x spatial coordinate [m]

X discretization parameter

$\bar{X}$ scaled discretization parameter

Y response data

Z matrix of auxiliary variables defined by Equation (2.8)

Subscripts

0 reference condition

lof lack-of-fit

t thermal

Greek

α discretization parameter dependence exponents for variance model

β discretization error model coefficients

$\tilde{\beta}$ lack-of-fit error corrected discretization error model coefficients

γ convergence rates

δ dimensionless discretization

Δ discretization

ε residual

ε_0 permittivity of free space [F/m] ($= 8.8542 \times 10^{-12}$)

λ_D electron Debye length [m] ($= \sqrt{\varepsilon_0 T/(en)}$)

μ mean values

ν degrees of freedom in the F distribution

Ξ variable defined by Equation (2.18)

ρ charge density [C/m^3]

ϕ electric potential [V]

Ψ dimensionless MPW

ω_p electron plasma frequency [s^{-1}] ($= \sqrt{e^2 n/(\varepsilon_0 m)}$)

Ω domain in discretization parameter space

Acronyms

CDF cumulative distribution function

CFD computational fluid dynamics

GCI grid convergence index

LAD least absolute deviation

LS least squares

PCMM predictive capability maturity model

PDF probability distribution function

PIC particle-in-cell

QoI quantity of interest

StREEQ stochastic Richardson extrapolation based error quantification

VPCL Vlasov-Poisson-Child-Langmuir

V&V verification and validation

MPW macro particle weight

Chapter 1

Introduction

This report focuses on numerical error estimation techniques for stochastic simulation methods having multiple discretization parameters with arbitrary, but theoretically predicted, convergence rates. Numerical error estimation lies within the purview of verification and validation (V&V), which are the primary means to estimate the accuracy of computational simulations [24]. Verification and validation are key components of the Predictive Capability Maturity Model (PCMM) [23], which is used to assess the level of maturity in modeling and simulation for high consequence applications.

Verification is the process of assessing correctness with which mathematical models are reproduced in a computational simulation. There are two distinct forms of verification: code verification and solution verification. Code verification generally involves the simulation of problems for which an analytic or otherwise high precision solution of known accuracy is available. From simulation responses at multiple discretization levels, the estimated fully-converged solution along with uncertainty bounds is compared directly to the known solution; here, the most common method for estimating the fully-converged solution is Richardson extrapolation [26]. Ideally, the full suite of code verification problems used tests the relevant physics of the target application. [32]. Solution verification refers to performing numerical error analysis of code responses for a problem without a known solution. The overall credibility of solution verification (as well as validation) is fundamentally built upon a solid foundation of code verification problems.

In contrast to verification, validation [24, 32] is the process of assessing the predictive capability of a code by directly comparing simulated response data to experimental data. Uncertainty quantification is a related activity to the validation process, which involves independently estimating the uncertainty in both experimental and computational results due to such factors as experimental error and input parameter uncertainty. Numerical error is clearly an essential component of the total uncertainty in the reported simulation results, although in many cases of practical interest, the numerical error can be shown to be small relative to the other uncertainties, and therefore may be ignored in the overall uncertainty estimate.

While error estimation is principally concerned with finding the uncertainty in a simulation response (either at a given discretization level, or with respect to the fully-converged estimate), estimating the convergence behavior is also useful, especially in the context of code verification. Key parameters used to quantify convergence behavior are the convergence rates for various discretization parameters; these are defined as the exponents in the discretization parameter dependence for the leading numerical error terms which approach zero in the limit as the discretization parameters approach zero. (Convergence rates are also called order-of-convergence [32] and order-of-accuracy [24] by other authors.) Changes to the computational algorithms and code bugs can cause reduction in the observed convergence rates, and therefore convergence rate estimates are a useful tool in the code development process. However, care must be exercised in interpreting deviations from theoretical convergence rates as convergence rates may be problem dependent.

Previous methods for numerical error analysis were developed primarily for deterministic simulation methods with the grid size being the sole discretization parameter. The most well known of these is the Grid Convergence Index (GCI) method, which provides an estimation of the amount of discretization error in the finest grid solution relative to the extrapolated estimate of the converged numerical solution [31, 32, 24]. The GCI method has been used by the computational fluid dynamics (CFD) community since the 1990s, has recently become a standard for publishing in the Journal of Fluids Engineering [8], and was incorporated in the ASME standard in 2009 [1]. In addition, the method is gaining acceptance in the field of solid mechanics [33, 2, 4]. The GCI method, based on Richardson extrapolation [26], estimates solution error using a series of simulations with systematically finer meshes. The Least Squares version of GCI [12, 25, 14, 13, 15] improves upon the error estimation by obtaining an improved estimation of the uncertainty at the expense of needing more simulations.

In the robust verification approach developed by Rider et al. [28, 29], a multi-fitting approach was used to estimate the numerical error for deterministic code results in up to two discretization parameters. The key idea in this work was to estimate the numerical uncertainty by using a suite of fitting models. These fitting models included various residual norms and alternate residual weightings, regularization methods, and predictions based on theoretical and extrema bounds on the convergence rates. He also developed a detailed workflow procedure which starts from the most converged region of the discretization parameter space and expands to include progressively less converged regions, and demonstrated this approach on CFD and neutronics simulation response data. In his recent work [30], he has extended his method to treat ill-behaved simulation data by applying constraints within the nonlinear optimization used to fit the regression models.

The method developed in this report was significantly inspired by robust verification [28, 29], and follows this multi-fitting approach. However, as it is targeted to stochastic code output, it is based on a different and smaller suite of fitting models. It also incorporates bootstrap sampling [16] based on a set of multiple code evaluations at the same discretization level to capture the effects of stochastic noise in the numerical uncertainty estimates. In addition, we have implemented a discretization domain reduction scheme which, starting at the full set of simulation data, progressively reduces the domain of discretization parameter space to exclude data outside the asymptotic region. Evidence for the credibility of a particular such analysis is obtained by performing an F-test and/or direct examination of fit residuals. For convenience, we will refer to this method as Stochastic Richardson Extrapolation based Error Quantification, or StREEQ.

The flagship computational application considered here is the particle-in-cell (PIC) kinetic plasma simulation method, although the methods developed here are directly applicable to any stochastic computational code. Response data for various quantities of interest (QoIs) for these and other Monte Carlo based computational techniques have a stochastic noise variance that is inversely proportional to the number of random samples; in PIC-plasma codes, the number of samples scales directly with the number of computational particles used in the evaluation of the QoI as well as the number of statistically independent, replicated computational responses which are ensemble averaged. For these codes, there are generally three relevant discretization parameters: grid size, time step, and macroparticle weight (MPW).

The StREEQ methodology is demonstrated on both stationary and time-periodic electron diode problems simulated using the Aleph PIC plasma code [3]. These diode problems are code verification examples, since they have a known solution. StREEQ error analysis was also performed on a smaller subset of the response data for the stationary diode problem in order to demonstrate the procedure for solution verification. In all cases, the method produced converged solution estimates which bracketed the known solution and convergence rate estimates consistent with the theoretical values.

The basic outline of this report is as follows: a full description of the StREEQ error estimation method is presented in the first chapter. Several test cases, including steady and time-periodic diode simulations, are analyzed in Chapter 3. A summary of the report, with directions for future research, is presented in Chapter 4; and finally the derivation of the exact solution for the steady diode verification problem is derived in Appendix A.

This page intentionally left blank.

Chapter 2

StREEQ Numerical Error Estimation Method

In this chapter, the overall StREEQ method is introduced. The first sections, which describe the form of response data assumed by the method and how it is processed for use in the error analysis, consist of: generation of response data for use in the StREEQ fits, the bootstrap sampling technique, and the variance model fitting procedure. Subsequent chapters address the core of the method, which involve the choice of an error model and the suite of fitting models, estimation of the converged result and convergence rates with uncertainty, the discretization parameter domain reduction algorithm used to determine the optimal set of data to fit, and credibility assessment.

Response Data

A StREEQ simulation requires response data from multiple simulations at a number of discretization levels. Multiple replications are required at each discretization level in order to account for the stochastic noise, while multiple discretization levels are required to account for the discretization error. Here, we denote the response data for a specific QoI as Y_{jk} which is assumed to be a function of discretization parameters X_{qj}, where $j = 1,2,\ldots,M$, $k = 1,2,\ldots,N_j$, and $q = 1,2,\ldots,D$. Here, j is one of M discretization levels, while q indexes different discretization parameters, where D is the total dimensionality of discretization parameter space. The index k represents a specific replication at give discretization level; for each j, there are N_j computational replications; in total representing $N = \sum_j N_j$ independent simulations. These replications can be obtained in general by repeating the simulation with a different random number seed. In the steady state case, one can generate replications by a series of time-averages that are taken over sufficiently long time intervals to produce uncorrelated, independent values. A sufficient number of discretization levels is required to make the regressions overdetermined ($M > 3,6,10$ are required for $D = 1,2,3$ respectively), as well as ensuring that at least three different values for each discretization parameter are represented to ensure reliable estimation of the convergence rates.

This report will address cases of two- and three-dimensional discretization parameter space for PIC simulations; in the $D = 3$ case, X_{qj} are dimensionless versions of grid size δx, time step δt, and macroparticle weight (MPW, defined as the number of physical particles represented by a single computational particle) Ψ, while in the $D = 2$ case, time step convergence is omitted. The definitions of the dimensionless discretization parameters are defined by specific case in Chapter 3 because they vary slightly by the verification problem being simulated. Generally δx and Ψ related to their raw values Δx and MPW by a simple scaling with (constant) reference values, while δt is understood as a cell-based Courant condition with respect to physical time scale Δt. The overall theory is applicable to an arbitrary set of discretization parameters, although we assume that they are defined such that the discretization error decreases with decreasing X_q in the remainder of the report.

For multiple discretization parameters, it is good practice to keep the rate of refinement consistent between different discretization parameters. This was discussed by Richards [27] for the specific case of spatial and temporal convergence, but is extended in the formula below for the present application. Here γ_x, γ_t, and γ_Ψ are the theoretical grid, time, and MPW convergence rates respectively, while the index j' refers the discretization level where each parameter is the next smallest value after j.

$$\left(\frac{\delta x_{j'}}{\delta x_j}\right)^{\gamma_x} = \left(\frac{\delta t_{j'}}{\delta t_j}\right)^{\gamma_t} = \left(\frac{\Psi_{j'}}{\Psi_j}\right)^{\gamma_\Psi} \tag{2.1}$$

For time dependent QoIs, an independent StREEQ analysis will be performed for each QoI at each time step. When feasible, the set of simulations should be constructed in order to report at times evenly divisible by the time step at each discretization level. Otherwise, interpolation can be performed to a specific time step at the cost of introducing interpolation error into the response data. For spatially dependent QoIs, the situation is similar in principle but more complex, and may require interpolation between nodal values. For this reason, integrated QoIs are generally easier to deal with and better behaved than spatially-distributed or time-dependent QoIs, and also have a reduced level of stochastic noise.

Bootstrap Sampling

Multiple bootstrap samples of the mean μ_j^b and standard deviation (square root of the variance) of the mean σ_j^b, where b indexes separate bootstrap samples, are taken from the underlying response data and used directly to perform the multiple fits in a StREEQ analysis. These mean and deviation samples are obtained from independent bootstrap response samples Y_{jk}^b according to the formulas below. For each bootstrap response sample Y_{jk}^b, the response values are randomly sampled with replacement from the available data ($k = 1, 2, \ldots, N_j$) independently for each discretization level j. A basic property of bootstrapping [16] is that μ_j^b and σ_j^b are approximate samples of the true mean and standard deviation in the mean.

$$\mu_j^b = \frac{1}{N_j} \sum_k Y_{jk}^b \tag{2.2}$$

$$\sigma_j^b = \sqrt{\frac{1}{N_j(N_j-1)} \sum_k \left(Y_{jk}^b - \mu_j^b\right)^2} \tag{2.3}$$

It is generally advantageous to produce a set of data where σ_j^b are approximately constant with respect to discretization level j, although the degree of computational expense necessary to accomplish this is not always feasible. While this condition was maintained for most of the examples analyzed in this report (see Chapter 3), we included a single example with non-constant σ_j^b which achieved excellent results; based on this example and other experience, using response data with non-constant σ_j^b does not seem to be problematic.

In cases where a small number N_j of replications exist at one (or more) discretization levels, the stochastic noise distributions are poorly-resolved in a pointwise sense. While this fundamentally limits the number of unique bootstrap samples per discretization level, this situation is not as grave as it may seem. Assuming that Y_{jk} are unique for each level j at fixed k, there are $\binom{2N_j-1}{N_j}$ unique bootstrap samples. And, since response data is required for multiple discretization levels, the actual number of unique values is $\prod_j \binom{2N_j-1}{N_j}$; or $\binom{2N_1-1}{N_1}^M$ for the special case $N_j = N_1$ for all j. Assuming that the stochastic noise distribution differs only in location and scale between discretization levels, it is this total number of unique values that is the important parameter. In particular, assuming $D = 2$ for which a minimum of $M = 7$ discretization levels are required in order to perform the fits, the absolute worst case scenario for which StREEQ may be performed is for $N_1 = 2$ replications; yet this results in 2187 unique bootstrap samples. For the more reasonable case with $N_1 = 5$, this results in 126 unique samples per level and a total of 5.0×10^{14} unique samples!

Variance Model Fitting

Although the consequences of of small N_j do not generally cause StREEQ analyses to fail (with one notable exception, discussed below), there are two consequences that can be improved by an alternate representation of the variance over the entire data set. First the bootstrap estimator (Equation 2.3) will be unduly noisy, resulting in noisier numerical uncertainty estimation. Another effect of small N_j is that the bootstrap variance estimator $\left(\sigma_j^b\right)^2$ has a $N_j^{1-N_j}$ probability of obtaining a bootstrap sample with a single repeated value leading to a zero variance estimation, thus causing the the fitting process to fail due to the variance dependent weighting (see Equation 2.12). One way to alleviate the second effect is to replace zero variance realizations with the overall sample variance estimation. However, a superior approach is variance model fitting as it adequately addresses both problems when the variance scaling is theoretically known.

The variance model fitting approach builds on the fact that the variance has theoretically known dependencies on discretization parameters in most cases. For PIC kinetic plasma simulations, the variance in volume-integrated quantities will scale according to $\sigma_j^2 \propto \mathrm{MPW}\,\Delta x^3/N_j$, while the variance of surface-integrated quantities scale as $\sigma_j^2 \propto \mathrm{MPW}\,\Delta x^2/N_j$; here we are assuming that the volume and surface areas are defined independently of the mesh resolution level. Generalizing this concept, we assume a variance model given below, in which each dimension of discretization parameter space has an arbitrary exponential dependence specified by α_q coefficients.

$$\hat{\sigma}_j^2 = \frac{\hat{\sigma}_0^2 \prod_q X_{qj}^{\alpha_q}}{N_j} \tag{2.4}$$

Assuming that the stochastic noise in the data comes from the same underlying distribution with only differences in the variance which are captured by the variance model, the single fit coefficient $\hat{\sigma_0^2}$ can be readily obtained in the least squares sense, as shown.

$$\hat{\sigma}_0^2 = \frac{1}{N-1} \sum_j \sum_{k=1}^{N_j} \frac{\left(Y_{jk} - \mu_j\right)^2}{\sum_q X_{qj}^{\alpha_q}} \tag{2.5}$$

The variance model hypothesis can be tested using the nonparametric Brown-Forsythe test for equality of variances [6] for some significance level (we used a significance level of 5% for the validation cases analyzed in Chapter 3); where the use of nonparametric test avoids any assumption on the form of the stochastic noise distribution. This variance model based approach provides an alternative variance estimator for reducing the noise in the residual weights by replacing the bootstrap version σ_j^b with $\hat{\sigma}_j$ from Equation (2.4). For the engineered data set verification problem analyzed in Chapter 3, we show excellent results for estimating the numerical error by using as few as $N_1 = 2$ replications per discretization level when using the variance model fitting approach.

Discretization Error Model

Oberkampf and Roy [24] distinguish between four sources of numerical error: round-off error, statistical sampling error, iterative error, and discretization error. In choosing the error model used in this report, we assumed negligible round-off error and iterative error. The stochastic noise inherent in Monte Carlo based algorithms is an example of statistical sampling error, which is included in our present approach using bootstrap sampling. We are assuming that all remaining sources of numerical error consist of discretization errors, which have a parametric dependence that is explicitly accounted for in the discretization error model. We also assume that the method is convergent, i.e. that the true solution of the equations being simulated (the fully-converged solution) is recovered (for a sufficient amount of ensemble averaging) as the discretization parameters approach zero.

A commonly used model for discretization error is $Y = \beta_0 + \beta_1 X^{\gamma} + \varepsilon$ where X is the (single) discretization parameter (typically indicated by h for grid convergence studies) and γ is the convergence rate.[1] Our discretization error model is given by

---

[1] Although most authors use p for convergence rate, we used a different nomenclature in order to avoid conflict

$$\mu_j = \beta_0 + \sum_q \beta_q X_{qj}^{\gamma_q} + \sum_q \sum_{r>q} \beta_{qr} X_{qj}^{\gamma_q} X_{rj}^{\gamma_r} + \varepsilon_j. \tag{2.6}$$

This an extension of the common approach which assumes an arbitrary number D of discretization parameters and also includes second-order (in $X_q^{\gamma_q}$) cross-coupling between combinations of two discretization parameters. The bootstrap index (b) superscripts from the previous discussion are omitted for clarity, and indices $q, r = 1, 2, \ldots, D$. The β and γ are fit (optimized) using a suite of fitting models. Here, β_0 is an estimate of the converged response, and the γ_q are estimates of the convergence rates. An earlier version of this model retained the β_{qq} terms in the error expansion; this approach proved problematic because the $\beta_{qq} X_q^{2\gamma_q}$ term competes with $\beta_q X_q^{\gamma_q}$ but tending toward a convergence rate of half the true value, and therefore subsequent work has retained only the cross-coupling terms to second-order in the expansion. The total number of model coefficients (β) is $N_\beta = 1 + D(D+1)/2$, and the total number of fit coefficients is $N_\beta + D$.

The matrix form of the above expression is given below. Here, I is a $M \times 1$ column vector of ones, $\mathbf{Z}(\boldsymbol{\gamma})$ is a $M \times N_\beta$ matrix with rows corresponding to unique discretization levels and columns corresponding to fit coefficient (β) terms. Note that each $\mathbf{X}_q^{\gamma_q}$ and $\mathbf{X}_q^{\gamma_q}\mathbf{X}_r^{\gamma_r}$ term is also a $M \times 1$ column vector with entries for all discretization levels, and $\mathbf{I}$ is a $M \times 1$ vector of ones.

$$\boldsymbol{\mu} = \mathbf{Z}(\boldsymbol{\gamma})\boldsymbol{\beta} + \boldsymbol{\varepsilon} \tag{2.7}$$

$$\mathbf{Z}(\boldsymbol{\gamma}) = \begin{bmatrix} \mathbf{I} & \mathbf{X}_1^{\gamma_1} & \ldots & \mathbf{X}_D^{\gamma_D} & \mathbf{X}_1^{\gamma_1}\mathbf{X}_2^{\gamma_2} & \ldots & \mathbf{X}_{D-1}^{\gamma_{D-1}}\mathbf{X}_D^{\gamma_D} \end{bmatrix} \tag{2.8}$$

$$\boldsymbol{\beta}^\top = \begin{bmatrix} \beta_0 & \beta_1 & \ldots & \beta_D & \beta_{12} & \ldots & \beta_{D-1,D} \end{bmatrix} \tag{2.9}$$

$$\boldsymbol{\gamma} = \begin{bmatrix} \gamma_1 & \ldots & \gamma_D \end{bmatrix} \tag{2.10}$$

While we have found this discretization error model to be adequate for the verification problems considered in this report, the overall method is applicable to any parameterized model. We note that Rider [30] used a different form of the cross-coupling term in his work. Verifying the appropriate choice of discretization error model is part of the code verification scheme, and if poor fits are discovered at the solution verification or validation stage (for discretization points within the asymptotic region), this is evidence of an incomplete hierarchy of code verification problems.

Fitting Models

Following the approach used by Rider et al. [29, 28, 30], we use a suite of fitting models to propagate the uncertainties inherent in performing the fits to the discretization error model for multiple realizations of the bootstrap samples. In this report, we used a total of nine fitting models,

---

with L_p norms and p values used later in this chapter.

which varied in both weighting strategy and error norm. This set of fitting models is significantly different than that used by Rider et al., in particular because it chooses a different set of weights and error norms, as well as avoiding the use of regularization models and fits based on fixed convergence rates. We believe this selection of fitting models to be more appropriate for fitting stochastic response data, which is the desired application. Testing the present approach in the limit of vanishing stochastic noise (i.e. for deterministic response data), which would justify its use for general application, is left to a future work.

The overall form of the fitting models is defined by Equation 2.11, where $\mathbf{w}^s$ is an $M \times 1$ column vector of weights (defined in Equation 2.12 and Table 2.1), $||\cdot||_p$ is an L_p error norm, and s indexes the fitting model. The choice of error norm is the mechanism used to control the influence of outliers. At one end of the spectrum, $p = \infty$ denotes the minimax or L_∞ regression with is optimal for uniformly distributed residuals; this approach is maximally sensitive to outliers. In contrast, the least absolute deviation (LAD) or L_1 approach is a robust form of regression [19] as it is minimally sensitive to outliers. The LAD approach is optimal for Laplace (double exponential) distributed residuals, and does not necessarily yield a unique solution for the fit coefficients. The standard least squares or L_2 regression is taken as the middle approach and is optimal for normally distributed residuals. We use the $p = 1, 2, \infty$ norms to represent the full range of behavior. The principle justification for this approach is that least squares regression represents a median expected behavior and the remaining two cases bound the range of behavior. While one can invent any number of ways to distribute p within this range, the effect of the choice of distribution is unknown and using more than three norms to span this space would involve ad-hoc assumptions without more rigorous theoretical justification.

$$G^s(\boldsymbol{\beta}, \boldsymbol{\gamma}) = ||\,(\mathbf{w}^s)^\top [\boldsymbol{\mu} - \mathbf{Z}(\boldsymbol{\gamma})\boldsymbol{\beta}]\,||_{p(s)} \tag{2.11}$$

Weighting strategies are commonly used to control the degree of heteroscedasticity (non-homogeneous variance of the response variable with respect to the discretization parameters). The weighted least squares approach uses variance based weighting, which uses the sample variance to emphasize regions with lower variance in the objective function (Equation 2.11). We use weighting for the additional purpose of emphasizing different regions of discretization parameter space. For stochastic code output, using weights which are a function of the refinement level generally involves a tradeoff between the accuracy of the error model and the strength of the convergence signal (variation of the response as a function of the discretization parameters). Specifically, in the less refined regions of discretization parameter space, the convergence signal is larger and more easily distinguished from the stochastic noise, but may be outside the asymptotic region where the error model is expected to hold. In refined regions, although the convergence signal response is likely to be within the asymptotic region, it can be dominated by the stochastic noise. In this work, we used three different weighting strategies, for which an expression for the individual elements of $\boldsymbol{w}^s$ are shown in Equation 2.12. All cases include the standard variance-based weighting; here σ_j is ideally represented using the variance model fit ($\hat{\sigma}_j$ from Equation 2.4), although the bootstrap estimate (σ_j^b from Equation 2.3) can also be used when the variance model fit is poor or the fit parameters α_q cannot be determined. Additional weighting to emphasize different regions of discretization parameter space are specified by $a(s) \in 0, \pm 1$: standard variance weighting is indi-

Table 2.1. Fitting model L_p norm and weighting strategy $a(s)$ for the various fitting models $s = 1, 2, \ldots, 9$.

s	$p(s)$	$a(s)$
1	1	-1
2	1	0
3	1	1
4	2	-1
5	2	0
6	2	1
7	∞	-1
8	∞	0
9	∞	1

cated by $a = 0$, while $a = \pm 1$ emphasizes less refined and more refined regions of parameter space respectively. Three weighting strategies and three L_p error norms results in a total of nine fitting models with the parameters specified in Table 2.1.

$$w_j^s = \frac{\left(\prod_q X_{qj}\right)^{a(s)}}{\sigma_j} \tag{2.12}$$

The fits were performed using nonlinear optimization. Due to the highly nonlinear nature of the error model (2.7), we adopted a nested approach for efficiency. Specifically, the fitting is performed by minimizing $G^s(\boldsymbol{\beta}'(\boldsymbol{\gamma}), \boldsymbol{\gamma})$, where $\boldsymbol{\beta}'(\boldsymbol{\gamma}) = \operatorname{argmin}_{\boldsymbol{\beta}} G^s(\boldsymbol{\beta}, \boldsymbol{\gamma})$. For robustness in the presence of noisy and/or ill-behaved response data, the convergence rates were constrained according to $\gamma_q^0/4 \leq \gamma_q \leq 2\gamma_q^0$, where γ_q^0 are the theoretical convergence rates. Since the convergence rate can be problem dependent, and a value of half the theoretical rate is commonly encountered for certain classes of degenerate problems (e.g., the presence of discontinuities), the lower bound was set to include this possibility. Setting a lower bound closer to zero can lead to erroneous results as convergence rates may cause the $\beta_q X_q^{\gamma_q}$ to duplicate the effect of β_0, thus corrupting the estimate of the converged response estimate.

For least squares fits ($p = 2$), the inner optimization loop can be evaluated directly by solving a matrix equation using a standard approach [11] for linear regression. In this case, the coefficients $\boldsymbol{\beta}'$ can be obtained by solving $[\mathbf{Z}^\top \mathbf{W} \mathbf{Z}]\boldsymbol{\beta}' = \mathbf{Z}^\top \mathbf{W} \boldsymbol{\mu}$, where $\mathbf{W}$ is a diagonal $M \times M$ matrix with (nonzero) elements given by $W_{jj} = (w_j^m)^2$; this procedure results in significant computational cost savings for the least-squares inner regression. Biconjugate gradient stablization [35] was used to solve the above matrix equation, as poorly-conditioned matrices are encountered for practical problems. This procedure is also useful for initializing the nonlinear optimization scheme for general L_p norms by using an L_2 initial guess. For L_1 and L_∞ inner optimizations and all outer optimizations, Nelder-Mead optimization [22] was used. Due to the potential for finding local minima in the outer optimization step, a multi-start global optimization scheme is employed which compares multiple fits with different initial guesses for the convergence rates and keeps the result which achieves the smallest objective function. In this procedure, the theoretical convergence rates is used as one of the initial convergence rate guesses, while the remaining are obtained as Latin hypercube samples from a uniform distribution bounded by the constraints ($\gamma_q^0/4 \leq \gamma_q \leq 2\gamma_q^0$). In this work, we found that using ten sets of initial guesses for the convergence rates was sufficient.

Error Estimation

In a typical StREEQ analysis, the fitting procedure is repeated for each of nine fitting models for each bootstrap sample. In this work, we used $B = 100$ bootstrap samples, which corresponds to 900 total samples and provided reasonably smooth statistics. The key output quantities are the pooled fit parameters—$\beta_0^{bs}, \beta_q^{bs}, \beta_{qr}^{bs}, \gamma_q^{bs}$ with $b = 1, 2, \ldots, B$ and $s = 1, 2, \ldots, 9$—from which the overall confidence intervals can be established. Since the β parameters appear with linear dependence in the error model (2.6), it is straightforward estimate the lack-of-fit uncertainty from the residuals and the Jacobian for the error model for a single regression (b, s); the standard approach is to included the uncertainty (variance) shown by example below for β_q^{bs}, where the variance in the remaining β parameters is similarly constructed.

$$\sigma_{\text{lof}}^2(\beta_q^{bs}) = \frac{1}{1 - N_\beta - D} \sum_j \left(\frac{\partial \varepsilon_j}{\partial \beta_q}\right)^{-2} \left(\varepsilon_j^{bs}\right)^2 \tag{2.13}$$

For this work, we adopted an alternate approach consisting of lack-of-fit corrections to the pooled fit parameters shown below. Here, the $\tilde{\beta}$ are lack-of-fit corrected discretization error model coefficients (β). In this approach, the set of $9B$ parameters is expanded to a set of $9BM$ (9 fit models, B bootstrap samples, and M discretization levels) parameters which substantially smooths the pooled sample distribution. When the fit quality is high (small, well-behaved residual distributions), this correction is negligible, especially for β_0. However, when the fits are poor, this correction helps retain the conservative nature of the estimates.

$$\tilde{\beta}_{0,j}^{bs} = \beta_0^{bs} + \frac{M-1}{M - N_\beta - D}\, \varepsilon_j^{bs} \tag{2.14}$$

$$\tilde{\beta}_{q,j}^{bs} = \beta_q^{bs} + \frac{M-1}{M - N_\beta - D} \left(\frac{\partial \varepsilon}{\partial \beta_q}\right)^{-1} \varepsilon_j^{bs} \tag{2.15}$$

$$\tilde{\beta}_{qr,j}^{bs} = \beta_{qr}^{bs} + \frac{M-1}{M - N_\beta - D} \left(\frac{\partial \varepsilon}{\partial \beta_{qr}}\right)^{-1} \varepsilon_j^{bs} \tag{2.16}$$

Due to the highly nonlinear dependence of the error model on the convergence rates, no lack-of-fit correction was applied to the pooled γ values since the Jacobian can tend toward zero at specific discretization levels in practical problems. Moreover, the linearization step introduced in the derivation for such expressions may not be a valid approximation when the lack-of-fit error is large. However, lack-of-fit corrections for the convergence rates are not generally needed for verification problems. Good code verification problems are computationally cheap enough to allow sufficient data well into the asymptotic region, and therefore lack-of-fit corrections will be small. In contrast, solution verification problems can be very expensive, which usually limits the size of the response data set to the point where precise estimation of the convergence rates is not possible.

The StREEQ approach generates a large amount of data samples, enabling a wide variety of error estimates. Typically, the most important result is obtained from the pooled statistics of $\tilde{\beta}^{bs}_{0,j}$; from which numerically derived 95% confidence intervals provide a useful estimate of the fully converged QoI with numerical uncertainty. For code verification problems with a sufficiently large response data set, estimates of the convergence rates may also be obtained and are useful for diagnosing code bugs and algorithmic performance. Evidence for the credibility of these results can be obtained by examining the residuals plot or performing an F-test for lack-of-fit for the variance-weighted least-squares regressions; both are discussed in a later section within this chapter.

Discretization Parameter Domain Reduction

Parameterized discretization error models, like the one used in this report (Equation 2.7), generally only hold asymptotically in the limit $X_q \to 0$, which is referred to as the asymptotic region of discretization parameter space. As this is violated, the quality of the error model fits suffer. Generally, accuracy and numerical stability considerations for the type of solution being solved will provide a starting point for the coarsest discretization levels to use. In PIC plasma simulations, nonphysical self-heating (and other) undesirable effects are encountered [18] for $\Delta x \gtrsim \lambda_D$, where $\lambda_D = \sqrt{\varepsilon_0 T_0/(e n_0)}$ is the electron Debye length in terms of the permittivity of free space ε_0, elementary charge e and reference density n_0 and temperature T_0. The most commonly-used dimensionless form of the time step is the cell-based Courant number defined by $v_0 \Delta t/\Delta x$ where v_0 is a characteristic speed. For electrostatic PIC methods, v_0 is typically the thermal speed $v_\mathrm{t} = \sqrt{2eT_0/m}$ or equivalently (neglecting the factor of $\sqrt{2}$) $v_0 = \omega_p \lambda_D$ where $\omega_p = \sqrt{e^2 n_0/(\varepsilon_0 m)}$ is the electron plasma frequency and m is the electron mass. For electromagnetic PIC methods, choosing v_0 as the speed of light is generally appropriate. Choosing Courant numbers of greater than unity can lead to gross inaccuracies, while numerical instabilities can occur for $\omega_\mathrm{p} \Delta t \gtrsim 1$ when explicit-time stepping is used. [5, 18]

In the best case scenario for a given set of simulation responses, the StREEQ method will obtain fits with a reasonable fit quality in the first application. Poor fit quality may indicate an inappropriately chosen discretization error model, the presence of other sources of numerical error (e.g. round-off or iterative error), or that discretization parameter domain Ω (i.e. the total set of discretization levels spanned by the response data set) contains points outside the asymptotic region. In the latter case, the fit quality may be substantially improved by suitably restricting Ω by removing a subset of the coarser discretization levels. This restriction can be performed by trial and error, although this becomes highly cumbersome, especially for multiple discretization parameters and/or time/spatial dependent data. To address this, we developed a discretization parameter domain reduction scheme to sequentially reduce the size of discretization parameter domain by identifying and excluding regions where the fit is poor. By performing StREEQ analyses on each reduced domain in discretization parameter space, an optimal domain can be identified.

The first component of the discretization domain reduction scheme is the identification of a core region, defined as the minimum set of the most refined points in discretization parameter space for which the fits are overdetermined. Therefore, the core region Ω^* of the entire domain Ω will have the following properties: (i) it will consist of $M^* = N_\beta + D + 1$ unique points in discretization space, and (ii) it will have a minimum of three distinct points X_q for each discretization parameter q (in order to enable convergence rate estimates). Determining this core region involves first producing a scaled version of the set of discretization levels $\bar{X}_q$ such that $\bar{X}_q = 1$ when X_q is the third smallest unique value in for each q; thus, retaining all points with $\bar{X}_q \leq 1$ would be guaranteed to include a minimum of three unique values of X_q. Using these scaled points, the core region Ω^* is defined as the set of discretization levels obeying the inequality below for positive real numbers (S, Q) which are determined by the algorithm.

$$\left(\sum_q \bar{X}_q^S \right)^{1/S} \leq Q \tag{2.17}$$

In the above, for the limit $S \to \infty$ expression reduces to $\max_q \bar{X}_q \leq Q$. The algorithm proceeds as follows:

1. Set $S = \infty$ and $Q = 1$. If inequality (2.17) contains exactly M^* points, exit. If the number of points is $< M^*$ proceed to step 2. Otherwise, proceed to step 3.
2. Find $Q > 1$ such that inequality (2.17) contains exactly M^* points and exit.
3. Find bounded S such that inequality (2.17) contains exactly M^* points and exit.

The resulting (S, Q) will generally result in a core region Ω^* with the desired number of points M^*. For specific cases, when exactly M^* cannot be achieved with the above algorithm, the approach is to find optimal parameters (S, Q) resulting in the smallest discretization level greater than M^* and eliminate points satisfying $\left(\sum_q \bar{X}_q^S \right)^{1/S} = Q$ taken at random until precisely M^* points are retained in the core region.

The second component of the discretization domain reduction scheme is a technique for determining the points with the largest residuals using LAD (or L_1) regression. For the L_1 regression weighted to overrepresent the more refined regions of discretization parameter space (model $s = 1$ in Table 2.1), we expect the corresponding fits will have maximal ignorance of outliers in the complement of the core region Ω^*, and therefore will act to emphasize discretization levels with large X that have a poor fit to the error model. Since multiple bootstrap residuals values exist for each position $j \in \Omega$, we can identify the discretization levels with the poorest fit on the basis of median and 95% confidence interval bounds using the Ξ_j variable defined below. Here, ε_j^M and $\varepsilon_j^\pm$ correspond to median and upper/lower 95% confidence bounds respectively.

$$\Xi_j = \frac{\pm\varepsilon_j^M}{\varepsilon_j^M - \varepsilon_j^{\mp}} \text{ for } \varepsilon_j^M \gtrless 0 \tag{2.18}$$

The Ξ_j variable has a simple interpretation: when $|\Xi_j| < 1$ the zero value is contained within 95% confidence bounds, while the sign (±) indicates that the median residual is above/below the zero value. Thus, the largest $|\Xi_j|$ indicates the point $j' = \text{argmax}_j |\Xi_j|$ which is the most poorly centered about zero. When this point occurs in the core region, it indicates that the fits are unlikely to improve by further reducing the domain; this defines a stopping point for the algorithm. Otherwise, when the largest value of $|\Xi_j|$ occurs outside the core region, a reduced domain Ω' is formed by using j' as a pivot to remove points from the domain. Specifically, Ω' is the set of points j for which $X_{qj} < X_{qj'}$ in for all q.

Using this process, a full set of StREEQ fits is performed on a sequence of reduced domains $(\Omega_0, \Omega_1, \ldots)$; a schematic representation of such a sequence is shown in Figure 2.1. The expected behavior is that the derived uncertainty estimate for the various fit parameters decreases as points outside the asymptotic region are removed up to a point where the lack of data leads to increased uncertainty. Thus, the typical work process is to perform a StREEQ analysis for the entire sequence of subdomains and choose the optimal domain as that resulting in the minimum pooled variance in the estimated fully-converged result (Equation 2.14); although an alternate choice for the reduced domain can be used (for example, a smaller domain than the minimum variance domain may be optimal when credibility is of paramount importance). This procedure is easily automated, and results in a substantially simplified workflow for performing numerical error analysis for cases with multiple or time/spacial dependent QoIs.

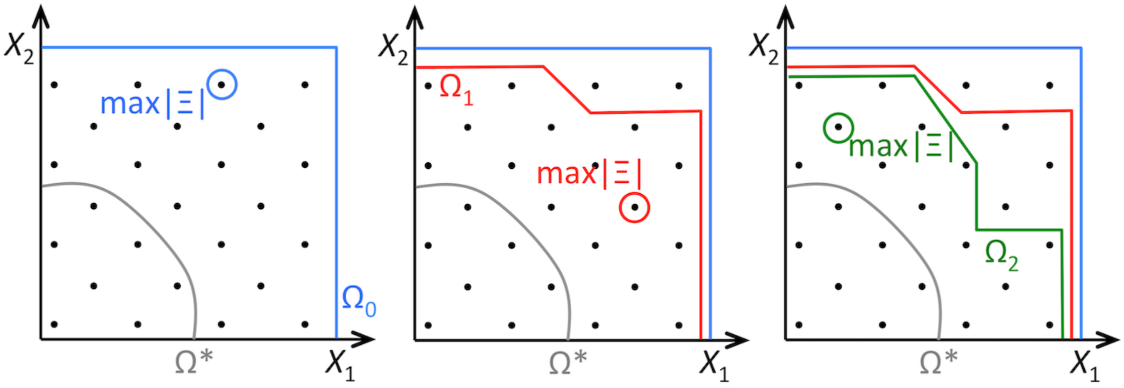

Figure 2.1. An example sequence of reduced domains for a $D = 2$ dimensional discretization parameter space. The core region Ω^* boundary is shown in gray, while subsequent reduced domains $(\Omega_0, \Omega_1, \Omega_2)$ are shown in blue, red, and green respectively. Also shown, are points of maximum $|\Xi|$ for each StREEQ analysis.

Credibility Assessment

The error estimation provided by the StREEQ method is is fundamentally built on the assumption that the discretization error model provides a reasonable description of the convergence behavior. Thus, evidence to support the credibility of the error estimate can be obtained by examining the residuals of the discretization error model fits. Although this examination can be performed by visual inspection in simple cases (i.e., by looking for significant outliers or obvious trends), it is not feasible for multiple QoIs or time/spatial dependent results.

The credibility can also be assessed by the F-test, which is a goodness-of-fit statistical test for regression [21]. Since the F-test is based on the assumption of normally distributed residuals, the test is only valid for the variance-weighted least-squares fitting model ($s = 5$ in Table 2.1). Since this fitting model is the "median approach" for the error norm and weighting strategy, this provides a reasonable metric for highlighting poor model fits. Note that a poor model fit might be due response data occurring outside the asymptotic region (addressed by obtaining more response data in the asymptotic region or by reducing the discretization domain), or a fault with the simulation code or the problem being simulated (e.g. shocks).

The F statistic for a bootstrap sample b is defined as

$$F^b = \frac{N-M}{M-N_\beta-D} \, \frac{\sum_j N_j \left(w_j \varepsilon_j\right)^2}{\sum_j w_j^2 \sum_k \left(Y_{jk}^b - \mu_j^b\right)^2}, \tag{2.19}$$

where the w_j, μ_j, and ε_j are for the variance-weighted least-squares regression model exclusively. This F statistic is evaluated for each bootstrap sample, yielding a distribution of values, each of which are compared to the critical values for the F distribution with degrees of freedom $\nu_1 = N - M$ and $\nu_2 = M - N_\beta - D$ at a specified significance level p_{crit}. Here, the critical values F_{crit} are obtained by inverting the cumulative F distribution $\mathscr{F}$ according to $\mathscr{F}(F_{\text{crit}}; \nu_1, \nu_2) = 1 - p_{\text{crit}}$. Alternatively, the probability values p for each F statistic can be evaluated directly via $p^b = 1 - \mathscr{F}(F^b; \nu_1, \nu_2)$. For p values larger than a chosen significance level (e.g., with a median greater than 5%), evidence for the credibility of the error estimate is obtained.

It is important to note that the F test alone is not proof of the credibility for solution verification and validation problems, but rather further evidence built on top of the code verification activities (see Chapter 1). However, even in the absence of systematic V&V activities to support solution credibility, the StREEQ method tends to be surprisingly robust and conservative. In these cases, using StREEQ error estimation for solution verification still generally results in a better estimate of the converged result than the most refined case.

Chapter 3

Application to Verification Problems

In this section, we apply the theory developed in Chapter 2 to several verification problems, including: (i) an engineered data set with a known bias for less refined points in convergence space, (ii) a steady electron diode problem, and (iii) a time-periodic electron diode problem. These problems were chosen both to test various aspects of the StREEQ method, as well as to demonstrate key steps in the code and solution verification process which are required to establish credibility for computer simulations relevant to a particular application [24].

The theory in Chapter 2 was implemented in MATLAB[1] and used to analyze all cases presented in this report; we used $B = 100$ independent bootstrap samples for each fitting model with weights based on the variance model fits (rather than the bootstrap sample deviation in the mean), and achieved excellent results.

Engineered Data Set Example

In order to test the discretization domain reduction algorithm, we created artificial data sets using the following model.

$$Y = 1 - 0.1X_1^2 - 0.05X_2 - 0.1X_1^2X_2 + \varepsilon + 0.1X_1^7X_2^{7/2}\sin\left(2\pi\left[\log\left(X_1\sqrt{X_2}\right) + 0.25\right]\right) \tag{3.1}$$

In the above expression, the first four terms are clearly in the form of the discretization error model (Equation 2.6), while the last term is a bias term which disappears for small X. The stochastic noise term ε is modeled with zero mean and fixed variance of $4 \times 10^{-6}/N_1$ where N_1 is the number of replications simulated for each point in convergence space. The points in convergence parameter space were chosen as $X_1 = 1, \frac{2}{3}, \frac{1}{2}, \frac{2}{5}, \frac{1}{3}$ and $X_2 = 1, \frac{1}{2}, \frac{1}{3}, \frac{1}{4}, \frac{1}{5}$. For this set of parameters, the bias term is significant for only the coarsest point ($X_1 = X_2 = 1$).

We first consider the case with normally distributed error with $N_1 = 5$ replications at each point in convergence space, which is plotted in Figure 3.1; the large bias at $X_1 = X_2 = 1$ is very apparent. Since the data was chosen with constant pointwise variance in the mean, we employed the

---

[1] MATLAB 2014b, The MathWorks Inc., Natick, MA

variance model fit with $\alpha_1 = \alpha_2 = 0$, which was confirmed using the Brown-Foresythe test. After performing initial StREEQ fits to the entire data set, the discretization domain reduction algorithm identified the point $X_1 = X_2 = 1$ as having the largest $|\Xi|$ indicating that it is the worst outlier in the coarse region as expected. Subsequent fits further reduced the size of the domain, although the lowest uncertainty in the converged result estimate is obtained by removing only the single worst point in this case. The full fit coefficient distributions are shown in Figure 3.2 for the full domain (Ω_0) and subsequent reductions (Ω_1 and Ω_2) of the data set. Clearly, the StREEQ technique based on Ω_1 does an excellent job estimating the true fit coefficients; and surprisingly, the fits for Ω_0 nearly capture all the true fit coefficients, albeit within much larger uncertainty bounds. Note that in Figures 3.2(a) – 3.2(d), both the raw (solid) and lack-of-fit corrected (dashed) coefficient distributions are shown. In the problems considered in this report, the lack-of-fit correction was found to be a relatively minor effect for β_0, and for all subsequent plots shown in this report, only the lack-of-fit corrected results will be plotted. For the full data set fits, clustering around the fit constraints in γ_1 and γ_2 is observed, which is indicates very large uncertainty in the convergence rates; although, since the majority of the fit values occur between the fit constraints, this does not necessarily indicate deviation from theoretical convergence rates.

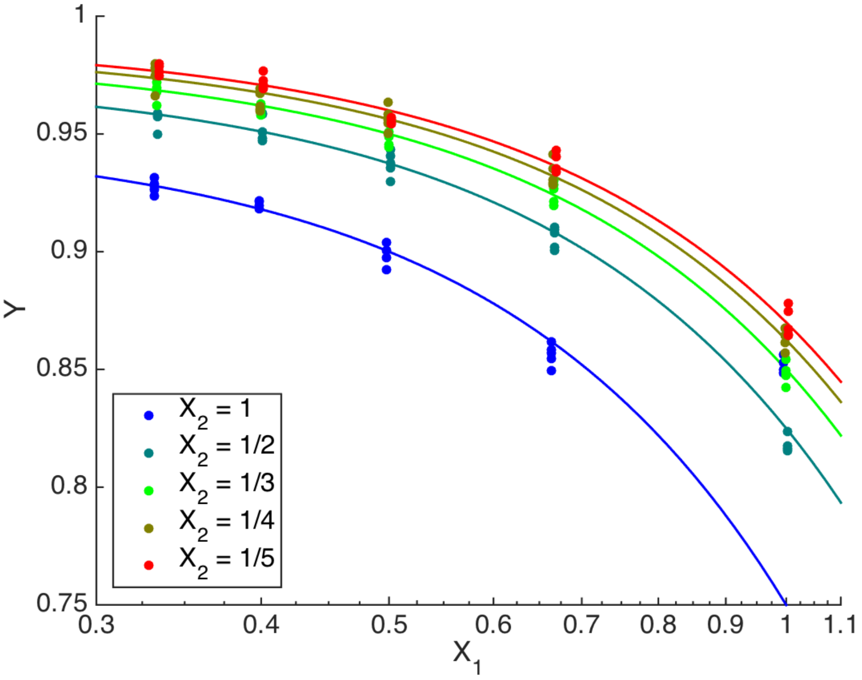

Figure 3.1. Plot of of generated data points for $N_1 = 5$ (symbols) compared to the first four terms of Equation 3.1 (lines).

Evidence supporting the credibility of the StREEQ estimates is obtained by analyzing the fit residuals. The raw fit residuals for the full data set Ω_0 and reduced domain Ω_1 are shown in Figure 3.3 for each fitting model. Clearly, significant outliers appear for the full data set, while the residuals are reasonably centered about zero for the reduced domain Ω_1. The F-test for the variance-weighted least-squares also indicate a poor overall credibility for the full data set, but reasonable results are obtained for reduced domains which remove the point at $X_1 = X_2 = 1$. The median p values are negligible, 5.2%, and 11% for discretization domains Ω_0, Ω_1, and Ω_2 respectively; the latter two values exceeding a significance level of 5%. This is also shown in the full distribution of p values in Figure 3.4; here a clear separation between Ω_0 and subsequent distributions is observed.

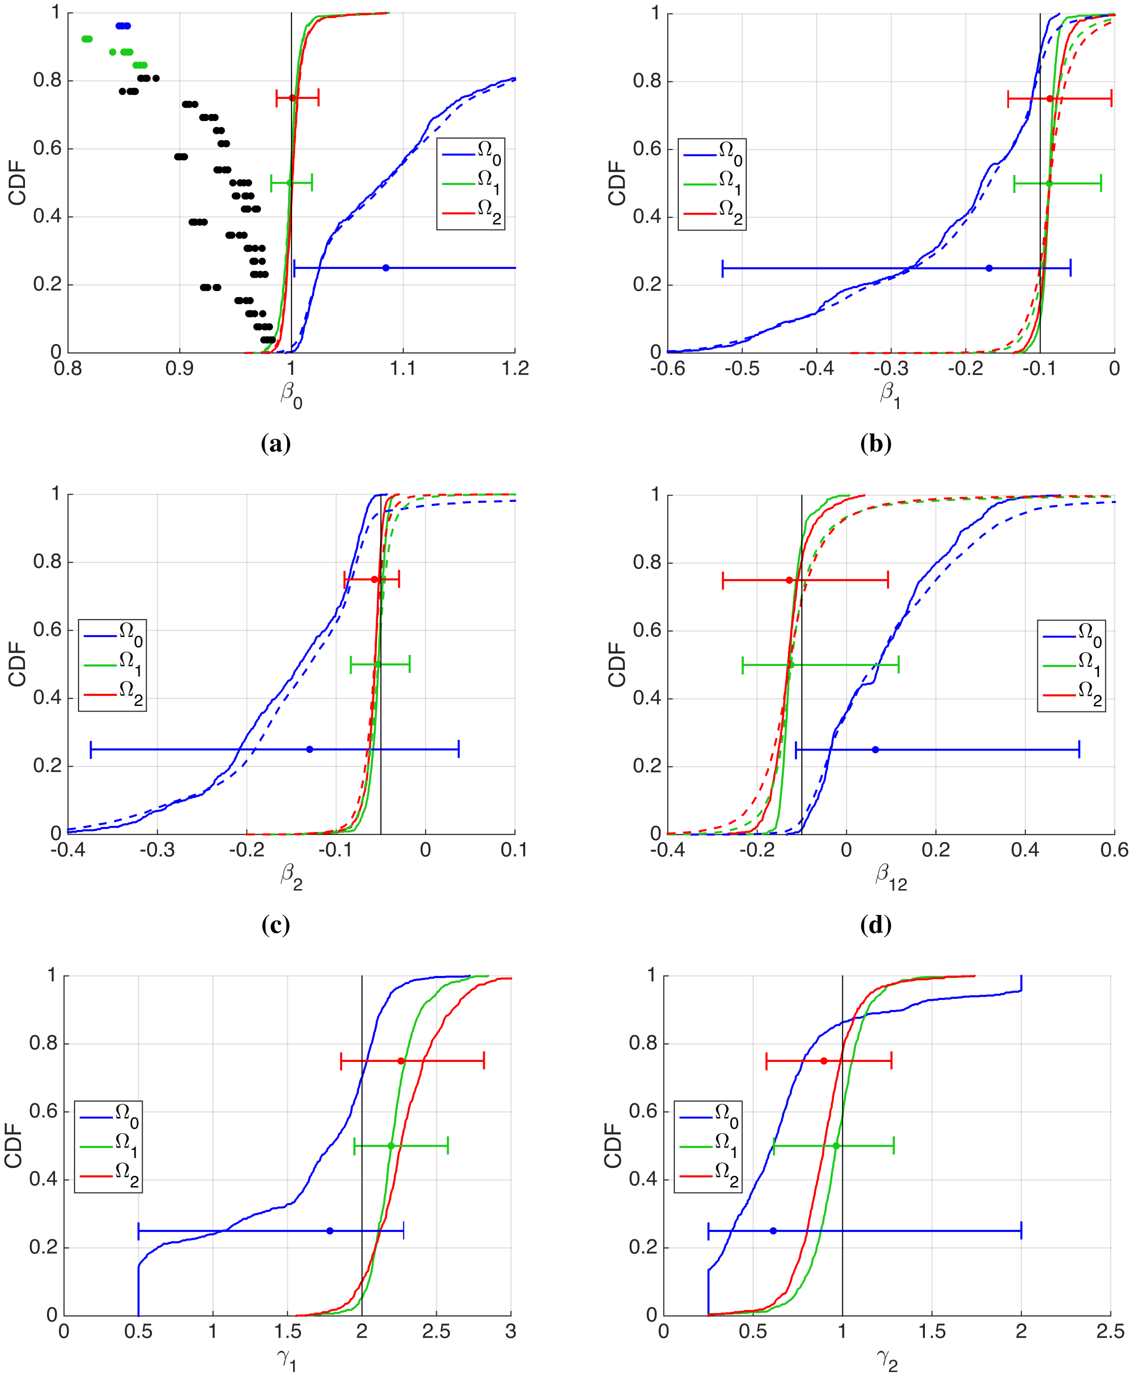

(a) (b)

(c) (d)

(e) (f)

Figure 3.2. Fit coefficient distributions for full data set (blue) and reduced domains (green, red) for normally distributed noise with $N_1 = 5$. Solid lines are raw β values, while the dashed lines for the $\tilde{\beta}$ coefficients include the lack-of-fit error corrections (see Equations 2.14–2.16). The horizontal error bars are 95% confidence intervals and vertical lines indicate exact values. The dots on the left side of the β_0 plot are the response data used to perform the fits, where the blue and green dots were eliminated after the Ω_0 and Ω_1 fits respectively.

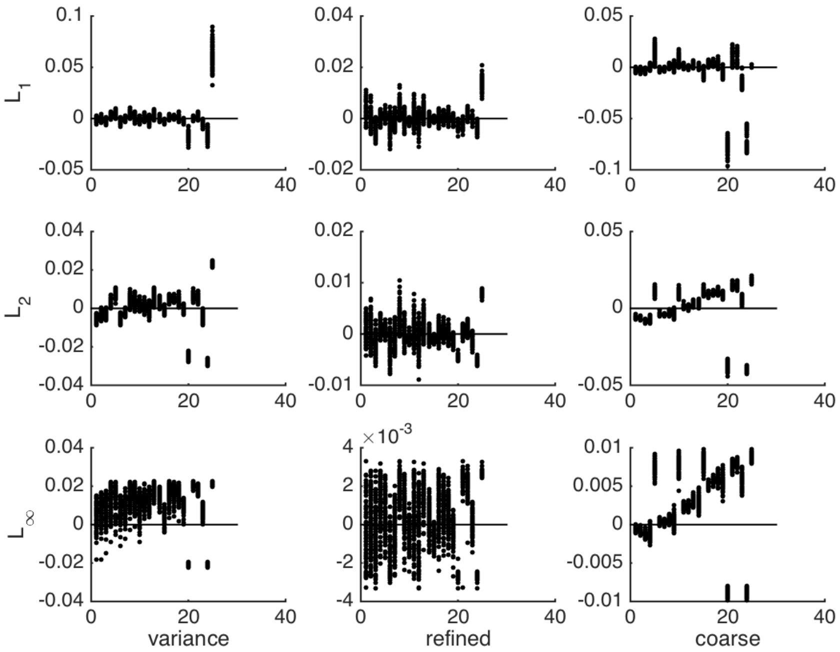

(a) Full data set

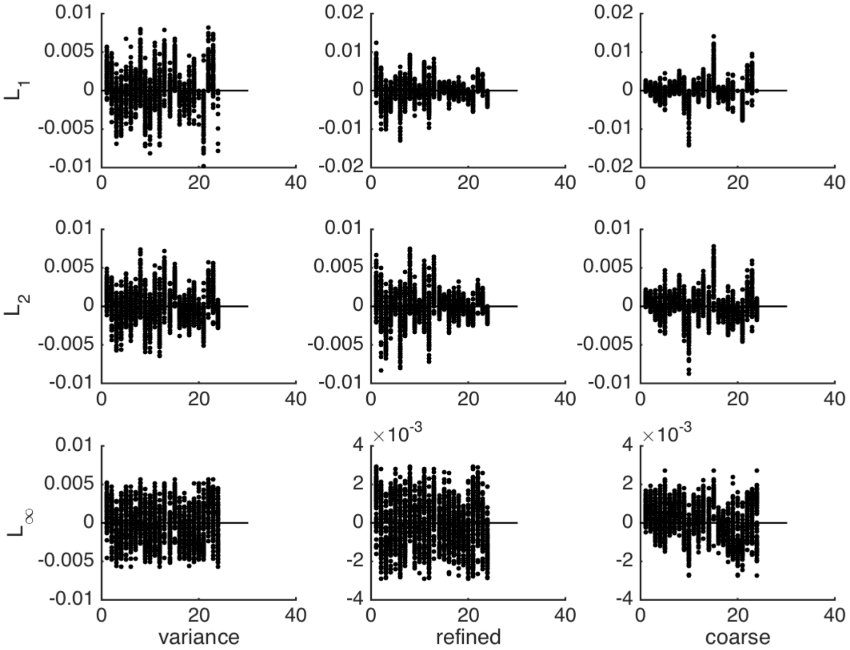

(b) Reduced domain Ω_1

Figure 3.3. Residuals for normally distributed noise with $N_1 = 5$ by regression model: **(a)** full data set Ω_0 and **(b)** and reduced domain Ω_1. Abscissa values represent different discretization levels, while repeated points are obtained from fits for multiple bootstrap samples.

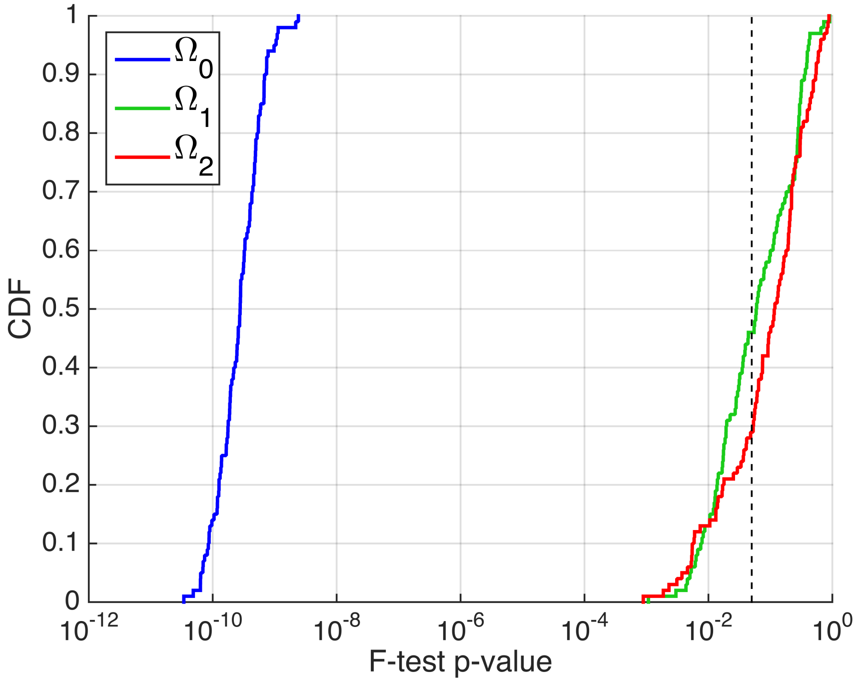

Figure 3.4. F-test probability value distributions for full data set (blue) and reduced domains (green, red) for normally distributed noise with $N_1 = 5$.

We also used the engineered data model for exploring the effect of the number of replications and the form of the stochastic noise distribution. To investigate the first effect, we performed a full StREEQ analysis on data sets constructed by varying the number of replications per point in convergence space N_1 (for normally-distributed noise with identical variances in the mean). In all cases, the automated discretization domain reduction algorithm identified the point $X_1 = X_2 = 1$ as the worst outlier in the first iteration, and the best results (minimum uncertainty in β_0) were found for domain Ω_1. The lack-of-fit corrected distribution β_0 is shown for many values of N_1 in Figure 3.5, demonstrating only random differences. The median p values for the F-test varied widely with no discernible trend from 0.6% to 11% with an average of 5.5%. This result is significant because it provides evidence that the StREEQ method (with variance fitting) gives reasonable numerical error estimates even for a very small number of replications.

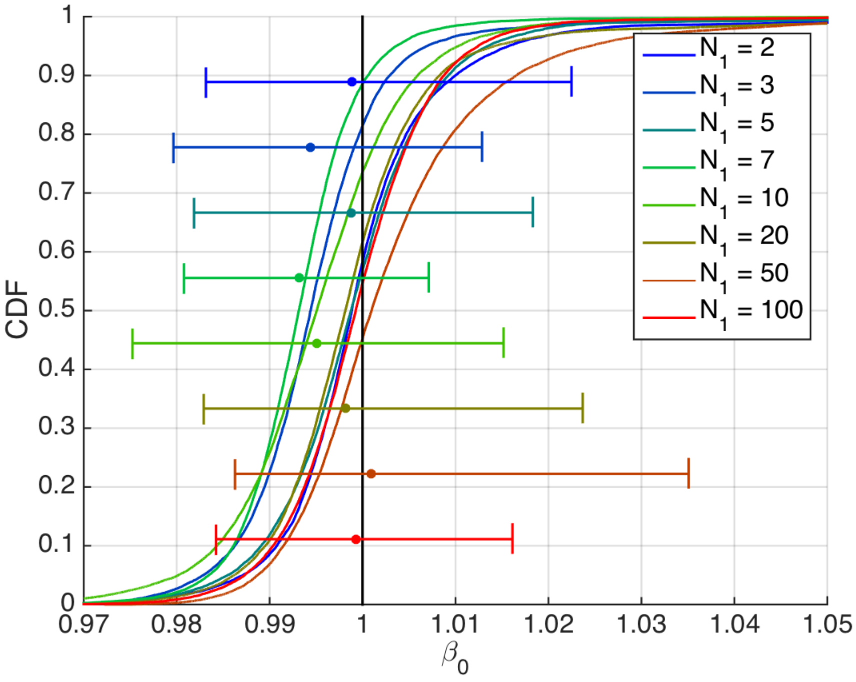

Figure 3.5. Estimated converged result distribution for normally distributed noise for various numbers of replications N_1.

To investigate the effect of the noise distribution, we used the engineered data model for five distributions, each with the same variance, mean, and $N_1 = 100$. In addition to normal, we created data with Laplace, uniform, log-normal, and beta distributed noise. The additional parameters for the log-normal and beta distributions were chosen to result in highly skewed distributions, as shown in Figure 3.6. The resulting StREEQ analysis predicts the uncertainties in the converged result shown in Figure 3.7. While all distributions capture the true solution well, the effect of skew in the log-normal distribution clearly results in a greatly increased uncertainty for $\beta_0 > 1$. The F-test results are not reported here, as the choice of clearly non-normal error distributions violates the inherent assumptions. Nevertheless, at least for the present test case, the validity of the StREEQ method does not seem to be dependent on normality of the noise distribution.

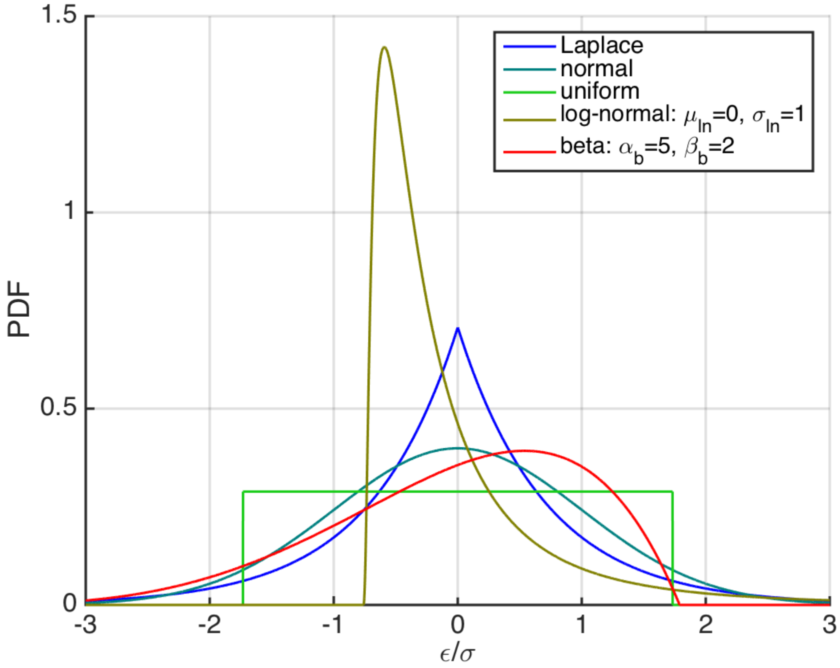

Figure 3.6. Alternate noise distributions.

Finally, we tested the effect of heteroscedastic response data by generating data of the form of Equation 3.1 with normally distributed stochastic noise with variance $2 \times 10^{-5} X_2/N_1$ with $N_1 = 5$. Note that this new specification maintains the original variance at $X_2 = \frac{1}{5}$, but results in higher variance for other values of X_2. The variance model fit was used with $\alpha_1 = 0$, $\alpha_2 = 1$, which was again confirmed using the Brown-Foresythe test. As before, StREEQ determines the initial reduced discretization domain Ω_1 by omitting the point $X_1 = X_2 = 1$, and subsequent reduced domains containing 15, 11, and 9 out of the original 25 points. The converged result estimates are shown in Figure 3.8; note that the true solution is captured within estimated 95% confidence intervals for every choice of discretization parameter domain. The minimum variance reduced domain is Ω_1 for which the median F-test probability value is 9%. The other fit coefficients are also captured very well, but aren't shown here.

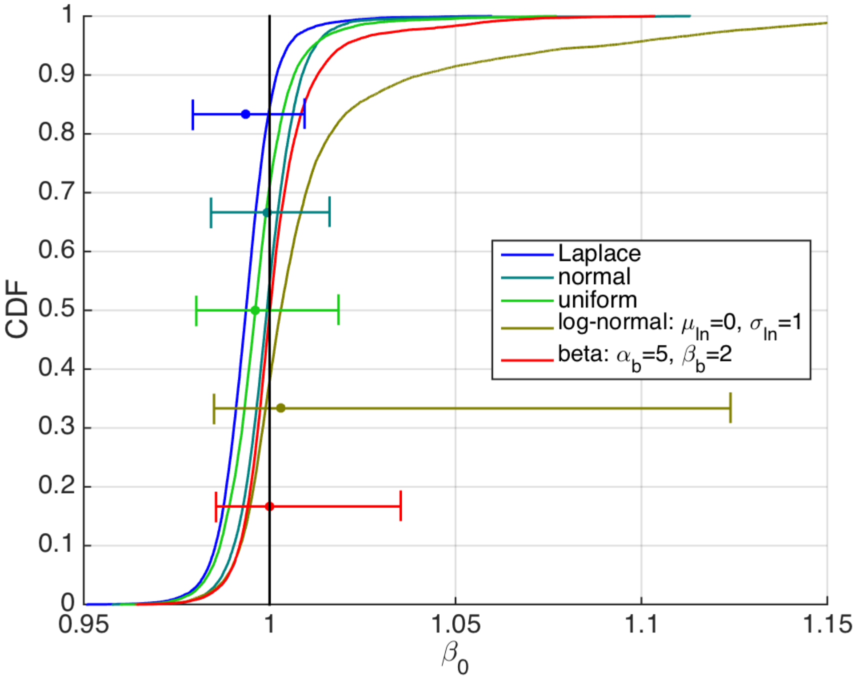

Figure 3.7. Estimated converged result distribution for various noise distribution choices with $N_1 = 100$.

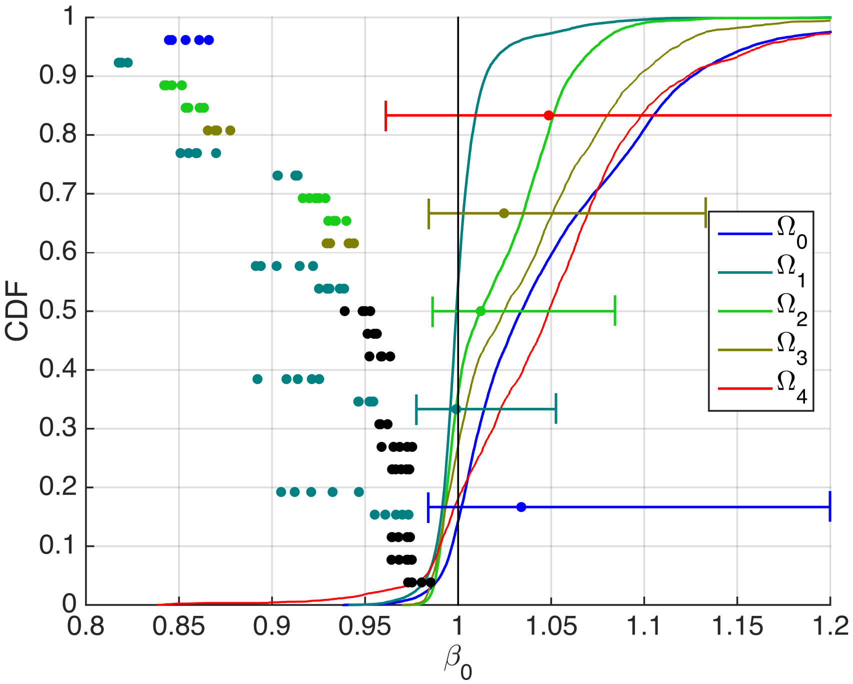

Figure 3.8. Estimated converged result distribution with heteroscedastic stochastic noise: full data set (blue), and reduced domains (green, red). The horizontal error bars are 95% confidence intervals and the vertical lines indicate exact values. The dots on the left side are the response data used to perform the fits, where the blue, dark green, light green, and brown dots were eliminated after the Ω_0, Ω_1, Ω_2, and Ω_3 fits respectively.

Stationary Diode Verification Problem

The Vlasov-Poisson-Child-Langmuir (VPCL) diode consists of two electrodes maintained at zero potential (voltage). At one electrode, Maxwellian (normally-distributed in velocity) electrons are injected into the physical domain. Only the electrons injected with sufficiently high energy to overcome the induced potential field will make it to the other electrode to be collected, while the rest are turned around and collected at the injection electrode. A full solution to the steady-state electron current is derived in Appendix A. For the simulations, we used the Aleph [3] PIC plasma simulation code, which is believed to be second-order convergent in space and time, and first-order convergent in macroparticle weight.

Dimensionless forms of the discretization parameters were chosen based on standard practice for collisionless kinetic plasma simulations, and are defined below. The grid size Δx is scaled by the reference electron Debye length λ_{D}, where $\delta x = 1$ corresponds to a numerical stability limit. Likewise, time step Δt is scaled in terms of its number of cell traversals per time step. Since Aleph is an electrostatic PIC code, we choose $\lambda_{\mathrm{D}}\omega_{\mathrm{p}}$ as the reference electron velocity. Note that the definitions for Debye length and plasma frequency were defined in the previous chapter. Restricting our analysis to one-dimensional problems, the dimensionless form of the MPW, Ψ, is scaled by the number of physical particles in a Debye length,[2] where A is the cross-sectional area normal to the current flow direction (electrode surface area), which is scaled out of the problem in one-dimensional simulations.

$$\delta x = \Delta x / \lambda_{\mathrm{D}} \ (= X_1) \tag{3.2}$$

$$\delta t = \lambda_{\mathrm{D}}\omega_{\mathrm{p}}\Delta t / \Delta x \ \ (= X_2) \tag{3.3}$$

$$\Psi = \mathrm{MPW}/(n_0 A \lambda_{\mathrm{D}}) \ \ (= X_3) \tag{3.4}$$

For the VPCL diode problem, we used an injection number density of $n_0 = 10^{16}\ \mathrm{m}^{-3}$, injection temperature $T_0 = 10$ V, and a domain size $L = 20\lambda_{\mathrm{D}}$. The corresponding total electron current is $-J = 77.0596\ \mathrm{A/m}^2$ (see Appendix A). This was simulated in Aleph using a one-dimensional problem domain. The dimensionless parameters were chosen in the range $\frac{1}{8} \leq \delta x, \delta t \leq 1$, where subsequent increments differ by a factor of approximately $\frac{1}{\sqrt{2}}$, and $4096 \leq \Psi^{-1} \leq 64$, where subsequent increments differ by a factor of $\frac{1}{2}$ (note that this method of selecting discretization levels satisfies Equation 2.1). The steady state current values were recorded after 1000 electron traversals (i.e. $t \geq 1000 \times L/\lambda_{\mathrm{D}}\omega_{\mathrm{p}}$) and averaged over time bins of size $(2000/\omega_{\mathrm{p}}) \times \Psi$; the variable time bins were used to generate samples with constant variance. In total, $N_1 = 700$ independent replications were generated by sampling the particles injected into and exiting out of the simulation domain. The various StREEQ error estimations performed below used variance model fits with $\alpha_x = \alpha_t = \alpha_\Psi = 0$, which was confirmed using the Brown-Foresythe test.

---

[2] For two- and three-dimensional simulations, the weight may be indexed to specify a number of computational particles per an area or volume indexed to the Debye length; e.g. within a Debye circle or Debye sphere.

Code Verification

The VPCL diode problem is a good code verification problem because it exercises a useful physics model (Maxwellian distributed electron injection above the space charge limit), is computationally inexpensive, and most importantly, has an analytical solution. Applying the StREEQ error estimation technique to the full set of available data resulted in multiple reductions of the original data set with the best results (minimum variance) achieved for reduced domain Ω_{21}. This reduced data set is shown in comparison to the full data set Ω_0 in Figure 3.9. This drastic reduction in the data set (from 343 to 103 points) occurred principally because the small pointwise variance (due to the large number of replications) made the fits more sensitive to discrepancies with the error model for coarser points in convergence space. The principle discrepancy noticed was oscillatory convergence in time step for coarse points in convergence space.[3] The resulting prediction for the diode current, shown in Figure 3.10, has 95% confidence intervals $77.019 \leq -J \leq 77.076$ A/m^2 and a median value of 77.054 A/m^2; this is an excellent prediction of the true result (77.0596 A/m^2).

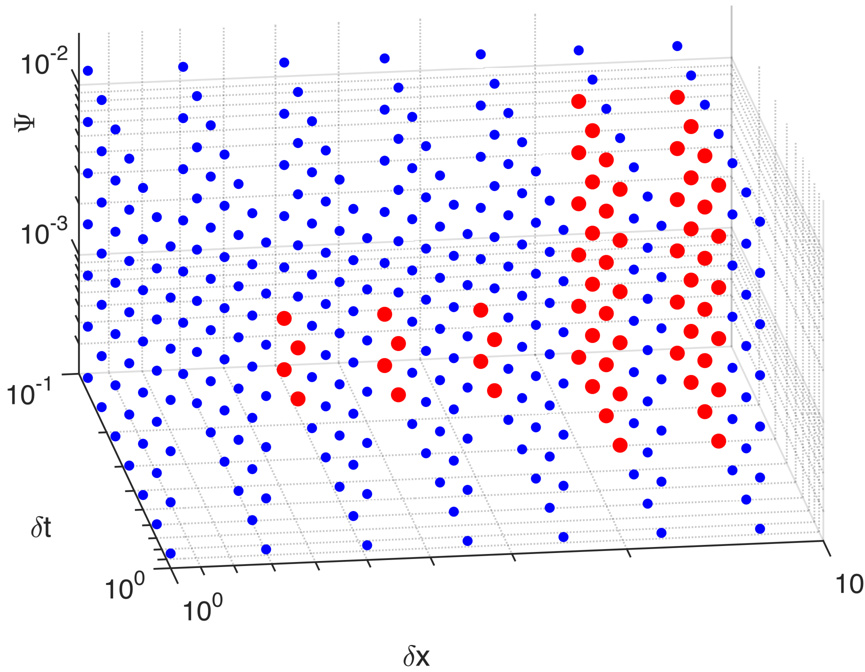

Figure 3.9. VPCL diode verification problem points in discretization space: full data set (blue) and minimum variance domain (red).

---

[3] The reason for the observed oscillatory convergence in δt is unknown. However, the electric potential field and particle motion is tightly coupled in space charge limited current flow, and the oscillatory convergence behavior may be due to cross-coupling between field and particle numerical error due to the time-splitting algorithm in PIC.

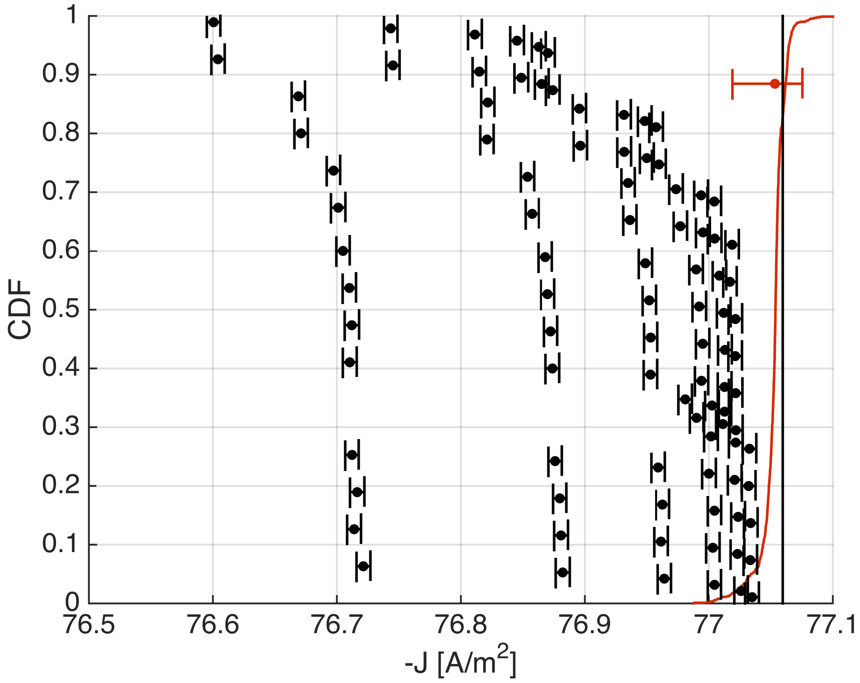

Figure 3.10. VPCL diode code verification problem StREEQ prediction for diode current (red) compared to the data set used for the fits for the minimum-variance discretization domain. Vertical line is the exact value.

The convergence rates are shown in Figure 3.11. While the mesh and MPW results do confirm the theoretically predicted second- and first-order convergence rates, minor clustering of the fit values around the constraints for time step convergence is observed. This clustering is a result of a very small time convergence signal which causes many fits to register unreasonable results for time convergence rates that are excluded by the fit constraints. On the basis of this StREEQ analysis, we can only claim that these results are consistent with second-order time convergence. For MPW convergence, a plateau occurs at approximately $0.87 < \gamma_\Psi < 0.95$ which is due to the pooled distribution of γ_Ψ being a composite of nine different fit models; near this plateau, the primary numerical uncertainty contribution for $\gamma_\Psi < 1$ comes from the refined- and coarse-weighted L_∞ norm fitting models alone.

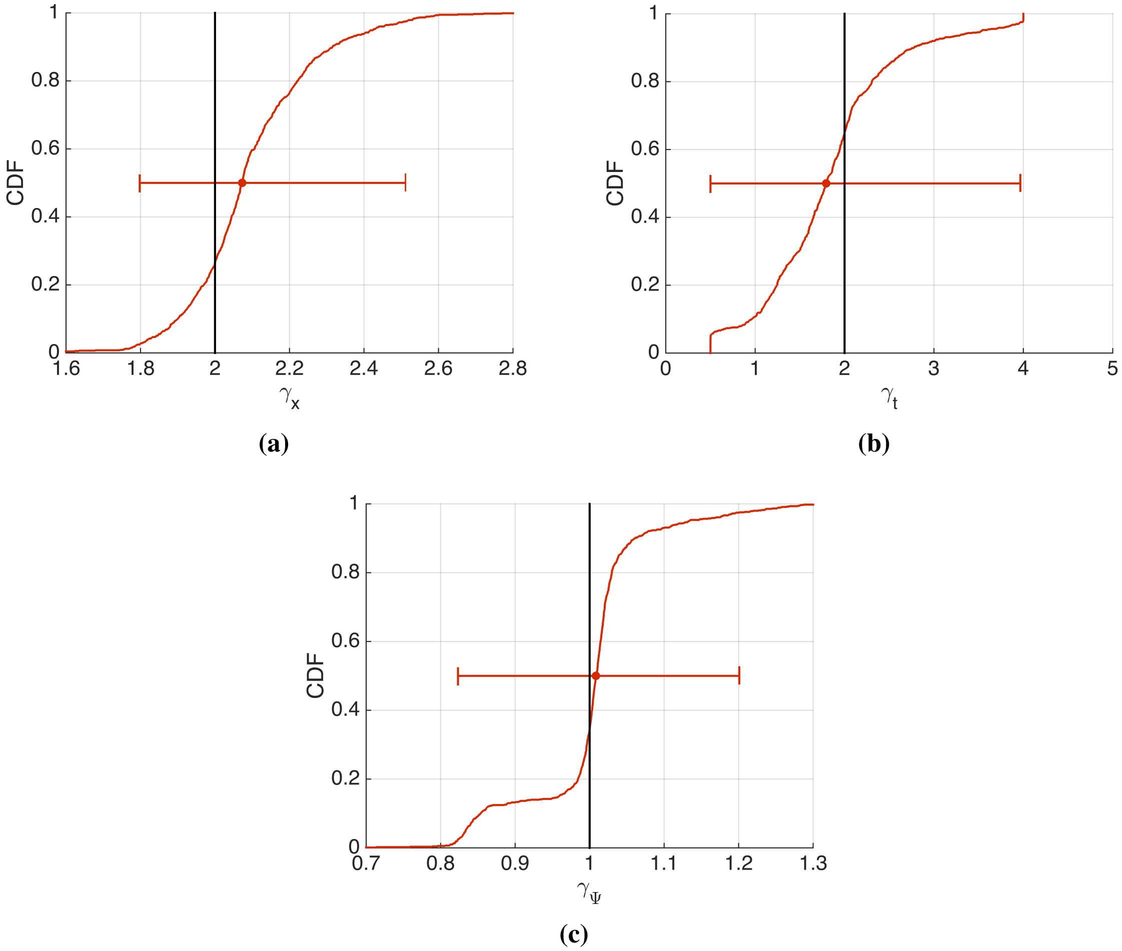

(a) **(b)**

(c)

Figure 3.11. Convergence rates for VPCL diode code verification problem: **(a)** mesh, **(b)** time step, and **(c)** MPW. Vertical black lines are the theoretical values.

The residuals, shown in Figure 3.12, are reasonably clustered around zero for the reduced domain. Note that the obvious heteroscedasticity in the refined- and coarsely-weighted fits is expected as it introduces a deliberate bias in the data set. The median *p* value for the F-test in this case is 0.07%, which is very poor. This is likely due to the very large number of samples which provides enough statistical evidence to find even the smallest issue with the discretization error model (e.g., the oscillatory δt convergence behavior).

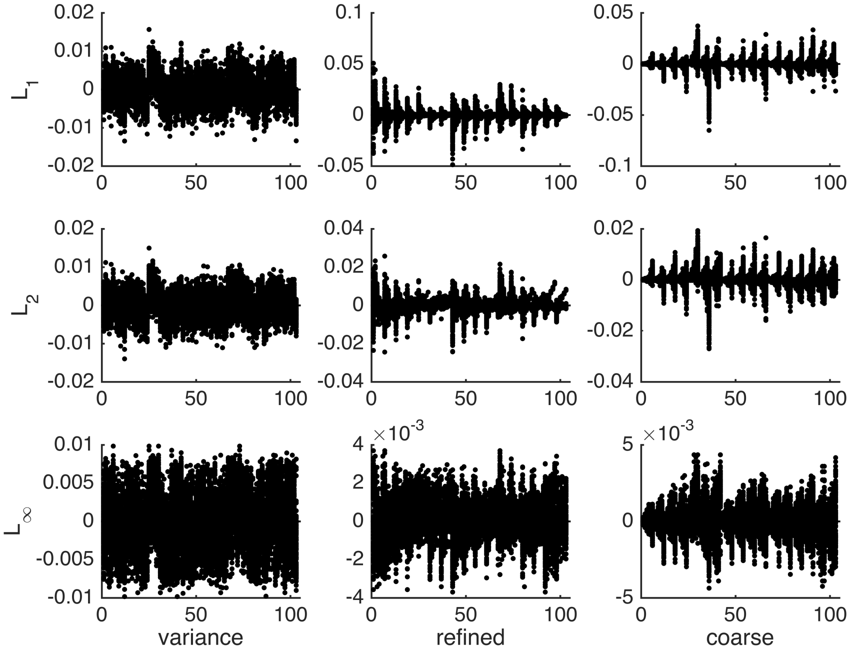

Figure 3.12. Residuals for VPCL diode code verification problem.

Since the largest discrepancies with the discretization error model were observed for time convergence, the obvious approach is to reduce the dimension of convergence space by keeping only the data with the most refined time step. The validity of this approach requires that $\delta t = \frac{1}{8}$ be fully converged in time; this is demonstrated in Figure 3.13 for three different levels of convergence in $(\delta x, \Psi)$. This plot also demonstrates the oscillatory convergence signal in δt, which is sustained until fully converged in time step. Using this approach, the minimum variance reduced domain (containing 27 out of the original 49 points) yields 95% confidence intervals of $77.040 \leq -J \leq 77.085$ A/m^2 and a median value of 77.056 A/m^2, where the full distribution is compared to the underlying data in Figure 3.14. Most importantly, the median value for the F-test p value is now 17%, indicating less discrepancy with the discretization error model and therefore substantial evidence for the credibility of the result. This is also consistent with the improved behavior of the residuals distributions shown in Figure 3.15 The spatial and MPW convergence rates again confirm the theoretical second- and first-order convergence, however, these are not reproduced here as they are similar to Figures 3.11(a) and 3.11(c).

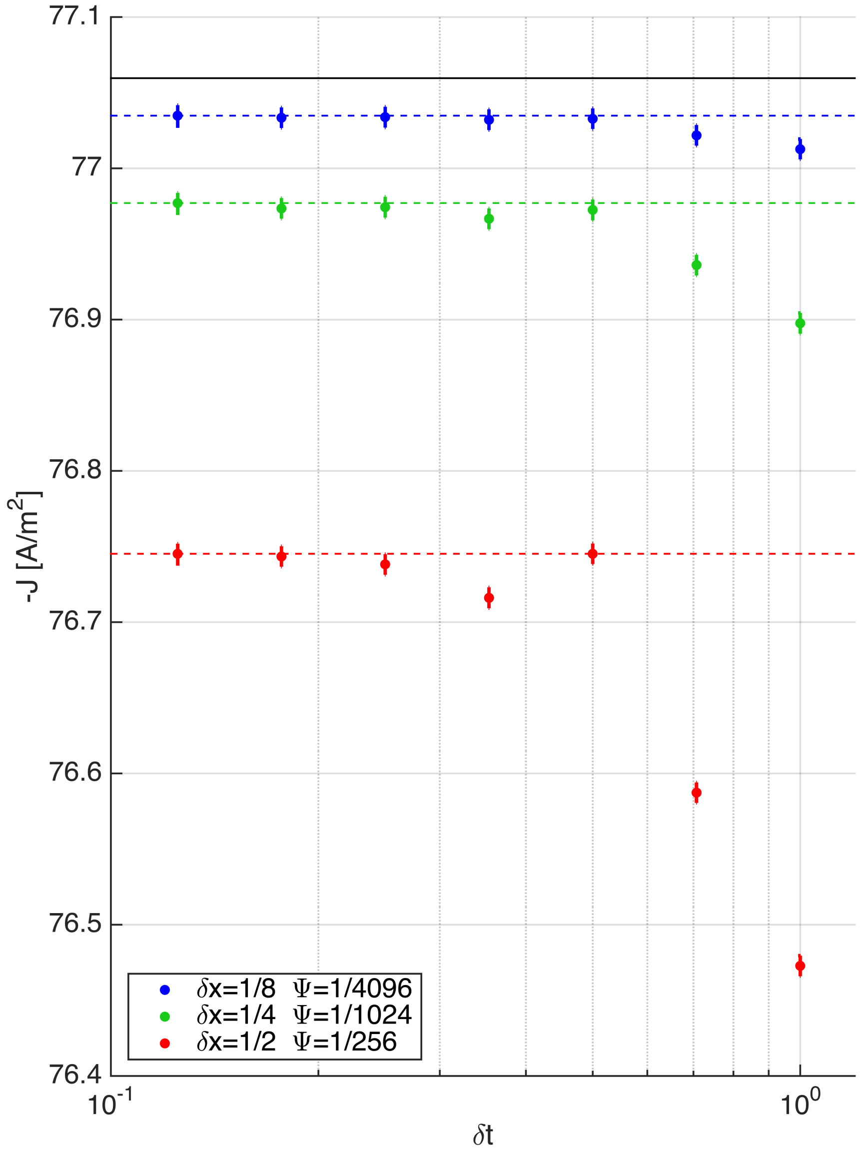

Figure 3.13. Time step convergence for VPCL diode code verification problem. Error bars are 95% confidence intervals for the mean values (based on all 700 replications) and the horizontal black line is the exact value.

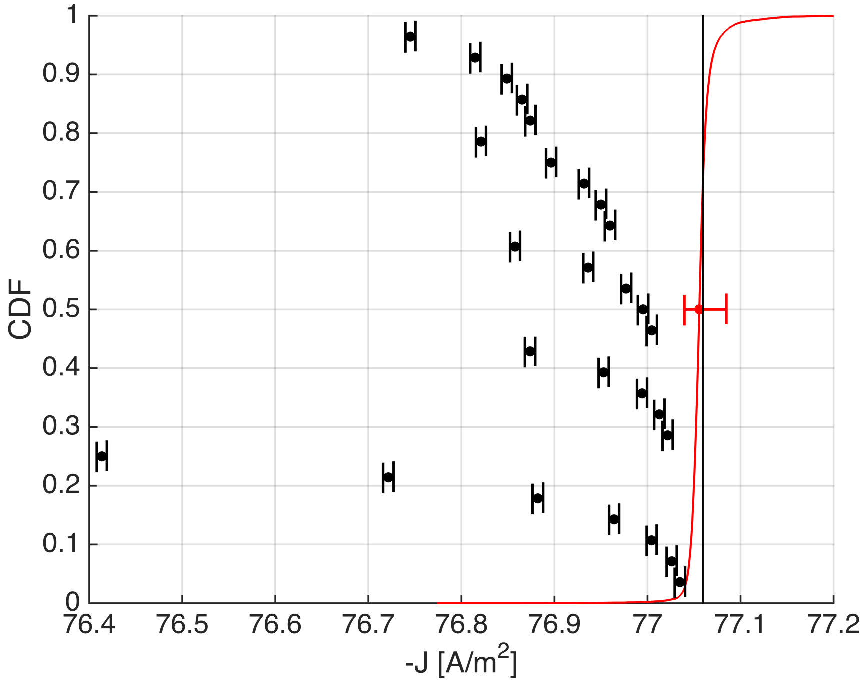

Figure 3.14. Time-converged VPCL diode code verification problem StREEQ prediction for diode current (red) compared to the data set used for the fits.

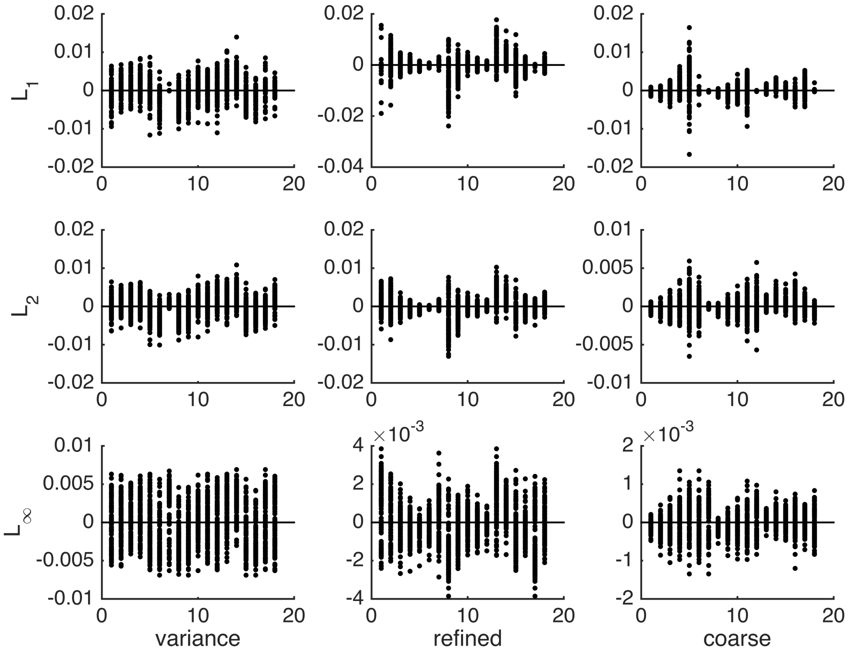

Figure 3.15. Residuals for the time-converged VPCL diode code verification problem.

Solution Verification

In contrast to code verification problems, solution verification problems are typically much more computationally intensive with an unknown solution. However, the lack of a known solution makes such problems a poor test of the StREEQ method. Instead of a real solution verification problem, here we use a drastically reduced subset of the data for the VPCL diode problem as a test solution verification problem in order to adequately test the method. In this case, only seven replications (out of the original 700) for points $\delta x, \delta t \geq \frac{1}{2}$ and $\Psi^{-1} \leq 256$ were retained. The resulting prediction for the converged diode for this case is shown in Figure 3.16, which used only 19 of the original 27 data points. Due to the drastically reduced size of the data, there is significantly more uncertainty in the prediction, although it still nicely brackets the true solution. Likewise, due to the increased pointwise variance of the underlying data, the lack-of-fit F-test obtained a median p value of 8.3% which supports the credibility of this result.

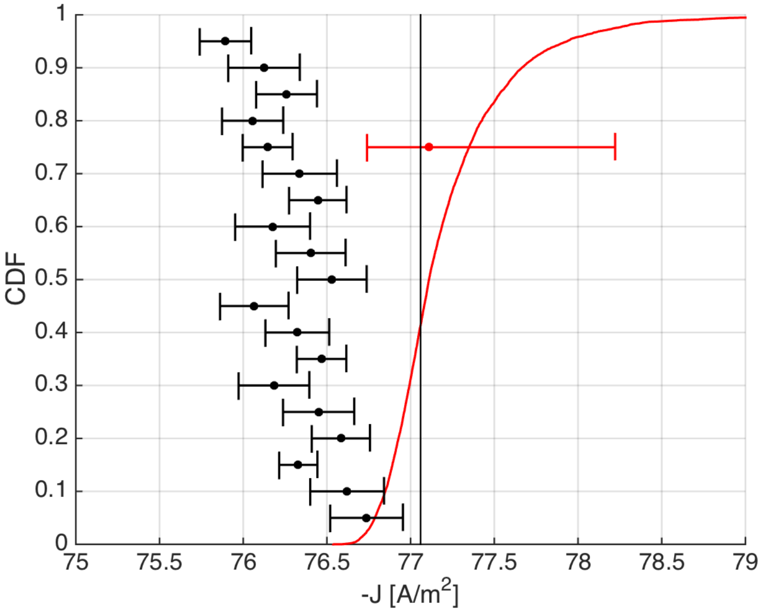

Figure 3.16. VPCL diode solution verification problem StREEQ prediction for diode current (red) compared to the data set used for the fits.

Periodic Diode Verification Problem

The final verification problem considered in this report is a time-periodic electron diode with an implicit analytic formulation derived recently by Caflisch et al. [7] This consists of a simple diode in which cold electrons are emitted from a grounded cathode with a sinusoidally varying density, but the anode is driven with a specific (non-sinusoidal) periodic variation in electric potential. The time-dependent current density at the anode J_1, which is QOI in this case, was found to exceed the Child-Langmuir limit [20] on average. Using the dimensionless parameters reported in Caflisch et al. [7] and scaling the electron injection density and velocity by 10^{12} m^{-3} and 10^6 m/s respectively (in order to simulate physically realistic values); this results in a $L = 31.5$ mm diode driven with a period $P = 57.1$ ns. We fixed the time step to 1.12 ns (512 time steps per period) which is sufficiently small to ensure a fully-converged time step for all meshes. QOI data was recorded after 100 simulated periods in order to omit the early transient effects.

Due to the cold injection condition, defining a reasonable Debye length was problematic, and therefore alternate definitions of the convergence parameters were defined. Dimensionless grid size was defined as $\delta x = \Delta x / L$, where δx^{-1} ranged between 8 and 64. The dimensionless MPW Ψ was defined as the inverse average number of particles in the entire simulation and Ψ^{-1} varied from 512 to 32768. The discretization levels followed a staggered grid pattern as shown in Figure 3.17. Using a staggered pattern can allow a more efficient covering of discretization parameter space with fewer simulated values, and was used for that purpose in the present example. The number of replications for each point j in convergence space was set to $N_j = \Psi^{-1}/4096$ in order to preserve a pointwise constant variance in the mean.

This verification example was chosen principally to demonstrate the utility of the discretization domain reduction approach for performing numerous StREEQ error estimations at multiple time periods. Here, we found an independent error estimate at 17 different time periods spaced 32 time steps apart. The variance model fit was used with $\alpha_x = 0$ and $\alpha_\Psi = 1$ to account for the variability in the number $\sim \Psi^{-1}$ of computational particles in each simulation, and the fit of the variance model was confirmed using the Brown-Forsythe test. This example was chosen primarily due to the time-dependence in the QOI. Performing StREEQ error estimation for this problem is greatly simplified by the discretization domain reduction algorithm, which independently determines a minimum-variance domain for each time step. In this case, the minimum-variance domain ranged from nine to 25 discretization levels for various time steps. The results, obtained at intervals of 32 time steps, are shown in Figure 3.18, which show excellent agreement with the result of Caflisch et al. In this case, the stochastic noise level was the dominant contribution to error estimate for most points in time, but grid and MPW convergence also influenced the results. The median p values for the F-test were greater than 50% for all times shown in the figure.

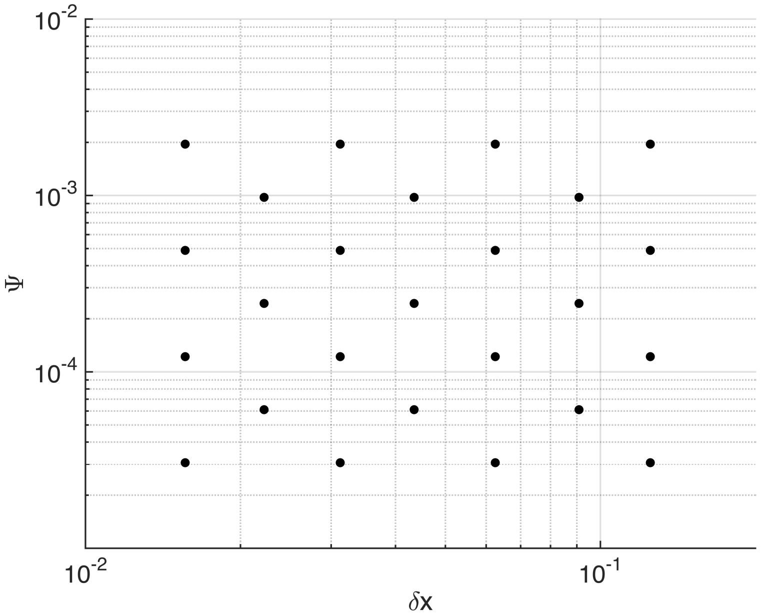

Figure 3.17. Simulated discretization levels for the periodic electron diode verification problem.

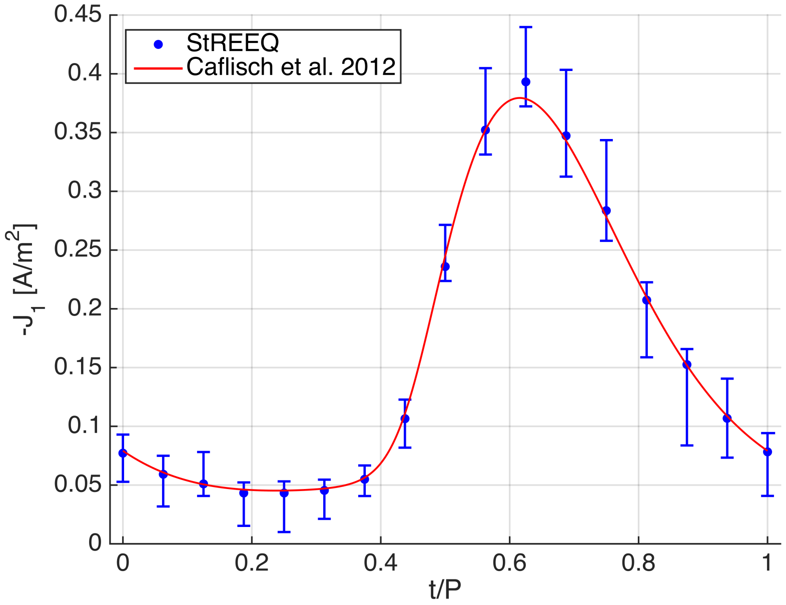

Figure 3.18. Periodic electron diode anode current: comparison between exact result (red) and StREEQ predicted 95% confidence intervals (blue).

Chapter 4

Summary and Future Directions

The stochastic Richardson extrapolation based error quantification (StREEQ) method is a powerful numerical error analysis technique for stochastic code response data. This method is based on using set of nine fitting models to represent uncertainty in the fitting process (influence of outliers, relative importance of subregions of discretization parameter space), while multiple bootstraps represent the influence of stochastic noise inherent in stochastic particle simulations. These features result in a method which appears to produce robust and conservative estimated confidence intervals for the fully-converged result, even in cases when credibility (as estimated by a lack-of-fit F-test) is highly suspect. By using engineered data sets and steady and time-periodic electron diode simulation data using a PIC kinetic plasma code, we demonstrated its usefulness for code verification and solution verification problems, which are key components in the overall verification and validation (V&V) process.

Future work is expected in two key areas: (i) performing StREEQ analysis for verification problems using existing and new Sandia analysis codes, and also (ii) development of a method to efficiently combine StREEQ numerical error estimation with input parameter uncertainty. In the first area, the existing Sandia plasma codes of interest include the Aleph [3] electrostatic PIC and the EMPHASIS [34] electromagnetic PIC codes; in addition, EMPIRE, an electromagnetic PIC plasma code which is currently in the early phases of development. For these codes, a subset of the regression test suite can be modified to perform StREEQ analysis code verification tests. This introduces several important advantages, including: allowing for a direct comparison between exact results by incorporating the numerical error estimation, and by using the convergence behavior as a diagnostic for changes in algorithmic behavior.

Additional future work will entail the development a combined uncertainty estimation approach which effectively combines numerical error estimations obtained from StREEQ analyses with input parameter uncertainties obtained by another method (such as Latin Hypercube sampling). This may be complicated in certain applications, as the convergence behavior may be radically different in different regions of input parameter space, as was observed in a recent validation exercise [9].

References

[1] The American Society of Mechanical Engineers. Standard for verification and validation in computational fluid dynamics and heat transfer, ASME V&V 20-2009, 2009.

[2] M. Barsotti. Comparison of FEM and SPH for modeling a crushable foam aircraft arrestor bed. In *11th International LS-DYNA Conference*, 2010.

[3] M. T. Bettencourt, J. J. Boerner, P. S. Crozier, R. W. Hooper, M. M. Hopkins, T. P. Hughes, H. E. Meyer, C. H. Moore, S. G. Moore, and L. C. Musson. Aleph manual. Technical report SAND2014-xxxx, Sandia National Laboratories, Albuquerque, New Mexico 87185, October 2014.

[4] B. Bewick, I. Flood, and Z. Chen. A neural-network model-based engineering tool for blast wall protection of structures. *International Journal of Protective Structures*, 2(2):159–176, 2011.

[5] C. K. Birdsall and A. B. Langdon. *Plasma Physics via Computer Simulation*. Taylor and Francis, 2005.

[6] M. B. Brown and A. B. Forsythe. Robust tests for equality of variances. *Journal of the American Statistical Association*, 69:364–367, 1974.

[7] R. E. Caflisch and M. S. Rosin. Beyond the child-langmuir limit. *Physical Review E*, 85(056408), 2012.

[8] I. B. Celik, U. Ghia, and P. J. Roache. Procedure for estimation and reporting of uncertainty due to discretization in CFD applications. *Journal of Fluids Engineering-Transactions of the ASME*, 130(7), 2008.

[9] P. Crozier, J. J. Boerner, G. V. Weirs, C. H. Moore, and E. V. Barnat. Ion beam simulation validation. Technical presentation SAND2015-7064PE, Sandia National Laboratories, Albuquerque, New Mexico 87185, August 2015 (OUO).

[10] P. Degond. The Child-Langmuir law in the kinetic theory of charged particles. Part 1, electron flows in vacuum. In *Advances in Kinetic Theory*. World Scientific, Singapore, 1994.

[11] N. R. Draper and H Smith. *Applied Regression Analysis*. Wiley, third edition, 1998.

[12] L. Eca and H. Hoekstra. An evaluation of verification procedures for CFD applications. In *24th Symposium on Naval Hydrodynamics*, 2002.

[13] L. Eca, M. Hoekstra, A. Hay, and D. Pelletier. Verification of RANS solvers with manufactured solutions. *Engineering with Computers*, 23(4):253–270, 2007.

[14] L. Eca, M. Hoekstra, and P. J. Roach. Verification of calculations: an overview of the Lisbon workshop. *AIAA paper*, (4728), 2005.

[15] L. Eca and M. M. Hoekstra. Evaluation of numerical error estimation based on grid refinement studies with the method of the manufactured solutions. *Computers and Fluids*, 38(8):1580–1591, 2009.

[16] B. Efron and R. J. Tibshirani. *An Introduction to the Bootstrap*. CRC Press, 1998.

[17] C. Greengard and P. Raviart. A boundary-value problem for the stationary Vlassov-Poisson equations: the plane diode. *Communications on Pure and Applied Mathematics*, 43(4):473–507, 1990.

[18] R. W. Hockney and J. W. Eastwood. *Computer Simulation Using Particles*. Taylor and Francis, 1998.

[19] P. J. Huber and E. M Rochetti. *Robust Statistics*. Wiley, second edition, 2009.

[20] M. A. Lieberman and A. J. Lichtenberg. *Principles of Plasma Discharges and Materials Processing*. Wiley, second edition, 2005.

[21] J. W. Neill. Testing for lack of fit in nonlinear regression. *The Annals of Statistics*, 16(2):733–740, 1988.

[22] J. A. Nelder and R. Mead. A simplex method for function minimization. *The Computer Journal*, 7:308–313, 1965.

[23] W. L. Oberkampf, M. Pilch, and T. G. Trucano. Predictive capability maturity model for computational modeling and simulation. Technical report SAND2007-5948, Sandia National Laboratories, Albuquerque, New Mexico 87185, October 2007.

[24] W. L. Oberkampf and C. J. Roy. *Verification and validation in scientific computing*. Cambridge University Press, 2010.

[25] D. Pelletier and P. J. Roache. Verification and validation of computational heat transfer. In *Handbook of Numerical Heat Transfer*, pages 417–442. Wiley, second edition, 2006.

[26] L. F. Ricardson. The approximate arithmetical solution by finite differences of physical problems involving differential equations, with an application to the stresses in a masonry dam. *Philosophical Transactions of the Royal Society of London*, 210:307–357, 1911.

[27] S. A. Richards. Completed Richardson extrapolation in space and time. *Communications in Numerical Methods in Engineering*, 13:573–582, 1997.

[28] W. J. Rider. Preliminary solution verification of Denovo: Focus on spatial-angular convergence. Technical report SAND2013-1421P, Sandia National Laboratories, Albuquerque, New Mexico 87185, February 2013.

[29] W. J. Rider and J. R. Kamm. Advanced solution verification of CFD solutions for LES of relevance to GTRF estimates. Technical report SAND2012-7199P, Sandia National Laboratories, Albuquerque, New Mexico 87185, August 2012.

[30] W. J. Rider, W. R. Witkowski, J. R. Kamm, and T. M. Wildey. Robust verification analysis. *under review for publication in Journal of Computational Physics*, 2015.

[31] P. J. Roache. Perspective: a method for uniform reporting of grid refinement studies. *Transactions-American Society of Mechanical Engineers Journal of Fluids Engineering*, 116:405–413, 1994.

[32] P. J. Roache. *Fundamentals of Verification and Validation*. Hermosa, Albuquerque, NM, 2009.

[33] L. Schwer. Recent and continuing activities in verification and validation by standards and other groups. In *11th Pan-American Congress of Applied Mathematics*, 2010.

[34] C. D. Turner, M. F. Pasik, D. B. Seidel, T. D. Pointon, and K. L. Cartwright. EMPHASIS$^{\text{TM}}$/Nevada UTDEM user guide version 2.1.2. Technical report SAND2015-xxxx, Sandia National Laboratories, Albuquerque, New Mexico 87185, April 2015.

[35] H. A. Van der Vorst. Bi-CGSTAB: A fast and smoothly converging variant of BI-CG for the solution of nonsymmetric linear systems. *SIAM J. Sci. and Stat. Comput.*, 13(2):631–644, 1992.

Appendix A

Vlassov-Poisson-Child-Langmuir Diode Solution Derivation

In this section we derive the solution for a one-dimensional, time-independent Child-Langmuir diode. This solution has been shown with more generality, arbitrary emission distribution, and finite voltage across the gap in Greengard and Raviart [17] and Degond [10]; however, the simpler solution shown here has a Maxwellian emission distribution and zero voltage across the gap. For the basic background on kinetic plasmas, the interested reader is directed to Reference [20]. While an electron diode is our primary consideration, our derivation is made in more general terms to consider a single particle species with arbitrary charge.

For a one-dimensional collisionless non-magnetized plasma consisting of a single charge species, the governing equation for the velocity distribution is the one-dimensional Vlasov equation

$$\frac{\partial f}{\partial t} + v_x \frac{\partial f}{\partial x} + qE_x \frac{\partial f}{\partial v_x} = 0, \tag{A.1}$$

where $f(x, v_x)$ is the distribution function which is a function of spatial coordinate x and the x-coordinate of velocity $\mathbf{v}$. Here, t is time, q is particle charge, $\mathbf{E} = -\nabla\phi$ is the electric field, and ϕ is the electric potential. The first term in Equation A.1 is identically zero for a time-independent problem, which is achieved in steady-state conditions.

The governing equation for the electrical potential inside the diode is the Poisson equation given by

$$\nabla^2\phi = -\rho/\varepsilon_0, \tag{A.2}$$

where $\rho = qn$ is the charge density for a single particle species with number density n, and ε_0 is the permittivity of free space. Two boundaries are defined at $x = 0$ and $x = L$, where L is the length of the diode. The inflow boundary condition is a Maxwellian particle distribution at $x = 0$, while both boundaries are grounded, i.e. $\phi(0) = \phi(L) = 0$.

With $\partial f/\partial t = 0$, the solution to the above equations is an example of a conservative Hamiltonian. Thus, the velocity distribution f can be constructed from single particle trajectory solutions of the Vlasov equation that conserve energy $mv_x^2/2 + q\phi(x)$, and extend the Maxwellian injection condition to all space. The resulting distribution is

$$f(x, v_x) = \frac{n_0}{v_\mathrm{t}\sqrt{\pi}} \exp\left[-\left(\frac{v_x^2}{v_\mathrm{t}^2} + \frac{2q\phi(x)}{mv_\mathrm{t}^2}\right)\right], \tag{A.3}$$

where n_0 is the injected particle density, $v_\mathrm{t} = \sqrt{2eT_0/m}$ is the thermal velocity, e is the elementary charge, T_0 is the injection temperature, and m is the particle mass. Note that for $v_x > 0$ at $x = 0$, this expression reduces to the Maxwellian injection velocity distribution as expected.

From Equation A.2, we know that the potential field solution must have negative/positive curvature on the entire domain for positive/negative particle charge respectively and a single potential maximum/minimum within the domain denoted by ϕ_m and located at x_m. Thus, particles injected with kinetic energies $mv_x^2/2 < q\phi_\mathrm{m}$ are reflected back to the boundary, while those injected with $mv_x^2/2 > q\phi_\mathrm{m}$ pass the potential extremum and are accelerated to the $x = L$ boundary. This gives the appropriate lower velocity bound on distribution (A.3) as $v_x > \mp\sqrt{2q(\phi_\mathrm{m} - \phi)/m}$ for $x \lessgtr x_\mathrm{m}$.

The current is evaluated by integrating the first moment of the distribution function directly using the lower bound discussed above

$$J_x = q\int \mathrm{d}v_x\, v_x f(x, v) = \frac{qn_0v_\mathrm{t}}{2\sqrt{\pi}} \exp\left(-\frac{2q\phi_\mathrm{m}}{mv_\mathrm{t}^2}\right) = \frac{qn_0v_\mathrm{t}}{2\sqrt{\pi}} e^{-\Phi_\mathrm{m}} \tag{A.4}$$

where the (positive) dimensionless potential extremum is defined as $\Phi_\mathrm{m} = q\phi_\mathrm{m}/(eT_0)$. The current flux is then determined by the value of Φ_m, the solution of which is developed below.

The charge density is evaluated by direct integration.

$$\rho = q\int \mathrm{d}v_x\, f(x, v) \tag{A.5}$$

To simplify the equations we introduce the following normalizations: $\eta = v_x/v_\mathrm{t}$, $\Phi = q\phi/(eT_0)$, $\xi = x/L_S$, where $L_S = \sqrt{e\varepsilon_0 T_0/(q^2 n_0)}$. For electrons, L_S reduces to the electron Debye length based on the injection conditions. With this normalization, Poisson's equation is now

$$\frac{\mathrm{d}^2\Phi}{\mathrm{d}\xi^2} = \frac{1}{\sqrt{\pi}}\int \mathrm{d}\eta\, e^{-(\eta^2+\Phi)}, \tag{A.6}$$

where the integration limits in dimensionless form are now $\eta > \mp\sqrt{\Phi_\mathrm{m} - \Phi}$ for $\xi \lessgtr \xi_\mathrm{m}$; here ξ_m is the dimensionless potential maximum position.

Performing the integrations yield

$$\frac{\mathrm{d}^2\Phi}{\mathrm{d}\xi^2} = \frac{1}{2}\frac{\mathrm{d}}{\mathrm{d}\xi}\left(\frac{\mathrm{d}\Phi}{\mathrm{d}\xi}\right)^2 = \pm\frac{e^{-\Phi}}{2}\left[1 \pm \mathrm{erf}\left(\sqrt{\Phi_\mathrm{m} - \Phi}\right)\right] \text{ for } \xi \lesseqgtr \xi_\mathrm{m}. \tag{A.7}$$

An additional integration gives an expression for the potential gradient, where $\mathrm{d}\Phi/\mathrm{d}\xi = 0$ at $\xi = \xi_\mathrm{m}$ was used.

$$\left(\frac{\mathrm{d}\Phi}{\mathrm{d}\xi}\right)^2 = \left[1 \pm \mathrm{erf}\left(\sqrt{\Phi_\mathrm{m} - \Phi}\right)\right] e^{-\Phi} \mp \left(2\sqrt{\frac{\Phi_\mathrm{m} - \Phi}{\pi}} \pm 1\right) e^{-\Phi_\mathrm{m}} \text{ for } \xi \lesseqgtr \xi_\mathrm{m}. \tag{A.8}$$

Equations A.8 are separable integrals that can be integrated from the walls to the potential maximum which is equal to the length of the system.

$$\begin{aligned}\Lambda \equiv \frac{L}{L_S} \int_0^{\Phi_\mathrm{m}} \mathrm{d}\Phi \Bigg(&\left\{\left[1 + \mathrm{erf}\left(\sqrt{\Phi_\mathrm{m} - \Phi}\right)\right] e^{-\Phi} - \left(2\sqrt{\frac{\Phi_\mathrm{m} - \Phi}{\pi}} + 1\right) e^{-\Phi_\mathrm{m}}\right\}^{1/2} \\ &\left\{\left[1 - \mathrm{erf}\left(\sqrt{\Phi_\mathrm{m} - \Phi}\right)\right] e^{-\Phi} + \left(2\sqrt{\frac{\Phi_\mathrm{m} - \Phi}{\pi}} - 1\right) e^{-\Phi_\mathrm{m}}\right\}^{1/2} \Bigg)\end{aligned} \tag{A.9}$$

The dimensionless length Λ is the single parameter that controls the system. Equation A.9 does not have a known analytical solution, but $\Lambda \approx 4\sqrt{\Phi_\mathrm{m} e^{\Phi_\mathrm{m}}}$ is a close approximation obtained by neglecting the $\sqrt{\Phi_\mathrm{m} - \Phi}$ terms, which are small; this approximation bounds the solution from above and has less than 5.6% error at all points. Series expansions of $\Lambda - 4\sqrt{\Phi_\mathrm{m} e^{\Phi_\mathrm{m}}}$ in $\Phi_\mathrm{m}^{1/2}$ or $\Phi_\mathrm{m}^{-1/2}$ can also be developed; however, the convergence properties would be poor near $\Phi_\mathrm{m} = 1$ and high precision numerical evaluation of Equation A.9 is computationally trivial. The relationship between Λ and Φ_m is shown in Figure A.1, where the approximation is also shown.

The analytical result for the verification problem studied in this report was computed for an electron diode ($q = -e$) using $\Lambda = 20$, which obtains $\Phi_\mathrm{m} = 2.3979$. By fixing the injection density to $n_0 = 10^{16}$ m^{-3} and the injection temperature to $T_0 = 10$ V, a current flux of $-J_x = 77.0596$ A/m^2 is obtained.

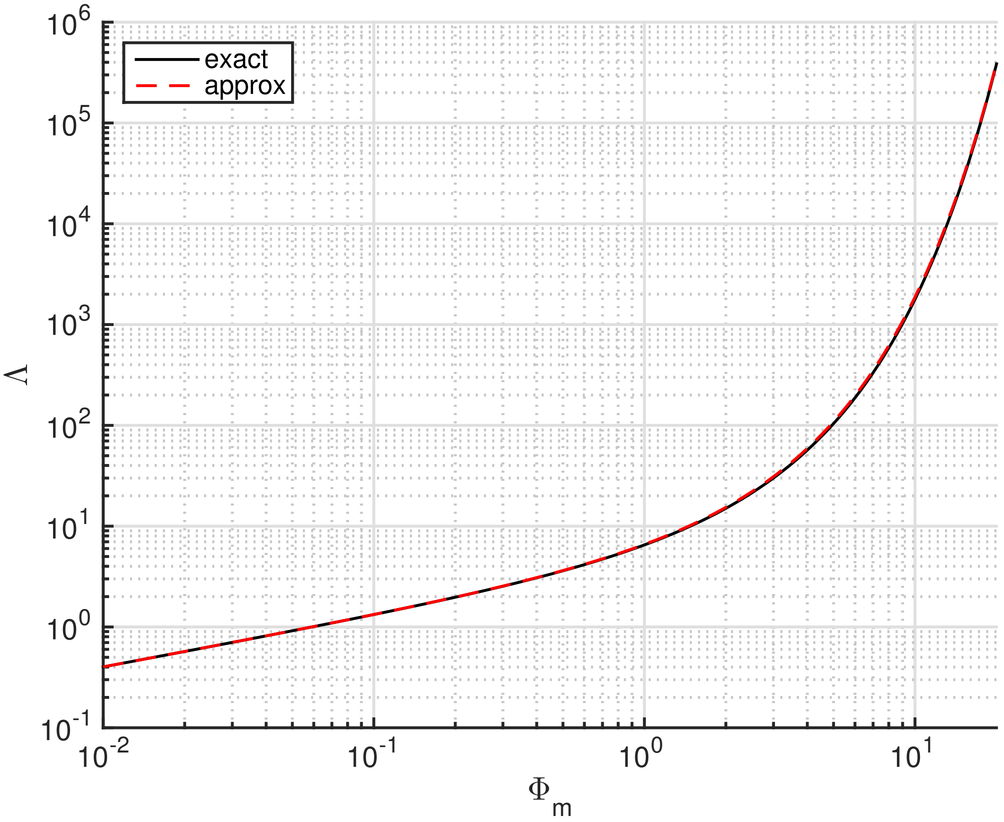

Figure A.1. Plot of dimensionless diode length (black, solid line) and approximation (red, dashed line) as a function of potential extremum. In logarithmic scales, the difference between the exact value and the approximation is difficult to discern.

DISTRIBUTION:

1	MS 0346	Anne Grillet, 1513
1	MS 0386	Richard G. Hills, 2122
1	MS 0828	Paul Crozier, 1541
1	MS 0828	Walter R. Witkowski, 1544
1	MS 0836	Jeremiah J. Boerner, 1513
1	MS 0840	Michail A. Gallis, 1513
1	MS 0840	Daniel J. Rader, 1513
1	MS 0845	Scott A. Hutchinson, 1540
1	MS 1118	Matthew M. Hopkins, 1423
1	MS 1152	Lorena I. Basilio, 1352
1	MS 1152	Keith L. Cartwright, 1352
1	MS 1152	Timothy D. Pointon, 1352
1	MS 1152	Gregg A. Radtke, 1352
1	MS 1152	C. David Turner, 1352
1	MS 1167	Bryan V. Oliver, 1340
1	MS 1168	Matthew T. Bettencourt, 1352
1	MS 1168	Christopher H. Moore, 1352
1	MS 1177	Lawrence C. Musson, 1355
1	MS 1181	Larry X. Schneider, 1350
1	MS 1318	Tim M. Wildey, 1441
1	MS 1320	Eric C. Cyr, 1442
1	MS 1323	William J. Rider, 1446
1	MS 0899	Technical Library, 9536 (electronic copy)

Sandia National Laboratories