CASL-U-2013-0192-000

Power uprates
and plant life extension

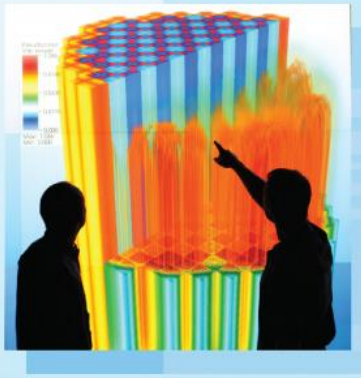

Engineering design
and analysis

Science-enabling
high performance
computing

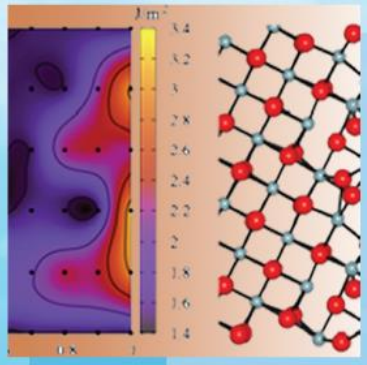

Fundamental science

Plant operational data

Consortium for Advanced Simulation of LWRs

L3:THM.CFD.P7.08
A Study of the Role of Adjoint-Equipped CFD in VUQ Analysis of Channel Boiling Simulations

Isaac A. Asher,
Krzysztof J. Fidkowski
UMich

Completed: September 30, 2013

U.S. DEPARTMENT OF ENERGY | Nuclear Energy

A Study of the Role of Adjoint-Equipped CFD in VUQ Analysis of Channel Boiling Simulations

Isaac A. Asher[*],
Krzysztof J. Fidkowski[†],

September 30, 2013

CASL Milestone L3-THM-CFD-P7-08

[*]Graduate Student, Department of Aerospace Engineering, University of Michigan
[†]Assistant Professor, Department of Aerospace Engineering, University of Michigan

CASL-U-2013-0192-000

Contents

1 Executive Summary **3**

2 Introduction **4**

3 Approach **4**
3.1 Model Equations 4
3.2 Primal Discretization 6
3.3 Adjoint Discretization 7
3.4 Error Estimation 7
3.5 Adaptation 8

4 Results **8**
4.1 Test Problem 8
4.2 Primal Solutions 9
4.3 Adjoint Solutions 9
4.4 Error Estimates 11
4.5 Adaptively Refined Meshes 12

5 Conclusions and Ongoing Work **12**

CASL-U-2013-0192-000

1 Executive Summary

This report documents the completion of milestone L3-THM-CFD-P7-08. The objective of this milestone was to study the use of adjoint methods in numerical simulations of boiling flow in order to efficiently compute output sensitivities and to estimate numerical errors. The scope of the physics considered was restricted to a one-dimensional drift-flux model, which was discretized using a high-order discontinuous Galerkin finite element method. Output adjoint solutions were computed using a discrete adjoint method, which built on an implicit primal formulation. Numerical errors were then computed using an adjoint-weighted residual, which when localized to elements also drove a mesh refinement strategy. The adjoint solver was verified by checking the convergence rate of output sensitivities computed with the adjoint. In driving mesh refinement, for the one-dimensional cases considered the benefits of using an adaptive strategy were marginal compared to uniform refinement. Similar behavior has been observed for other equations, such as the Euler and Navier-Stokes equations of gas dynamics. Based on experience with the latter equations, we expect improved performance of the adaptive algorithm in problems of higher spatial dimensions.

CASL-U-2013-0192-000

2 Introduction

Large numerical simulations in fluid dynamics generate massive amounts of data which can be daunting to post-process and verify. Indeed the expense of these calculations usually leaves few resources for verifying whether the solutions are sufficiently accurate. Inaccuracies due to the numerical discretization, which include effects of mesh size, polynomial order, and time step, can degrade results. Moreover this can be a "silent" degradation in that the user may not see anything obviously wrong with a computed solution (e.g. contour plots may be mostly smooth), even though the solutions contain unacceptably large errors for engineering purposes. The problem is most pronounced for equations with convective transport, which include many thermal hydraulics simulations.

In this milestone we address the error quantification problem through the use of output-based methods, which are used to quantify and minimize errors in scalar outputs of interest. The robustness and efficiency of these methods arise from the fact that they specifically target for refinement those and only those areas of the computational domain that are important for predicting the output. They are well-suited for convection-dominated flows because through the use of adjoint solutions they properly account for error propagation effects, which can be troublesome for many heuristic adaptive indicators.

In this work we demonstrate the application of output-based methods to drift-flux models of thermal hydraulic flows. Although the study is restricted to one spatial dimension, many of the lessons learned in the adjoint discretization, error localization, adjoint consistency, and adaptive refinement will extend to multiple dimensions.

3 Approach

3.1 Model Equations

Drift flux models have been applied to study multiphase systems in one dimension, often in the context of pipe flow. These models are relatively simple compared to more complex 2- or 3-D simulations, and are therefore used to exemplify the benefits of error estimation and

CASL-U-2013-0192-000

adaptation. Following [1], we have a 4-equation model:

$$\frac{\partial \rho_m}{\partial t} + \frac{\partial}{\partial z}(\rho_m v_m) = 0 \tag{1}$$

$$\frac{\partial \alpha\rho_d}{\partial t} + \frac{\partial}{\partial z}(\alpha\rho_d v_m) = \Gamma_d - \frac{\partial}{\partial z}\left(\frac{\alpha\rho_d\rho_c}{\rho_m}V_{dj}\right) \tag{2}$$

$$\begin{aligned}\frac{\partial(\rho_m v_m)}{\partial t} + \frac{\partial}{\partial z}(\rho_m v_m^2) = &-\frac{\partial}{\partial z}p_m + \frac{\partial}{\partial z}\left(\mu_m\frac{\partial v_m}{\partial z}\right) - \rho_m g_z - \frac{f_m}{2D}\rho v_m^2 \\ &- \frac{\partial}{\partial z}\left(\frac{\alpha\rho_d\rho_c}{(1-\alpha)\rho_m}V_{dj}^2\right) - \frac{\partial}{\partial z}\mathrm{COV}_m\end{aligned} \tag{3}$$

$$\begin{aligned}\frac{\partial(\rho_m h_m)}{\partial t} + \frac{\partial}{\partial z}(\rho_m h_m v_m) = &-\frac{\partial}{\partial z}\left(\frac{k_m}{c_{pm}}\frac{\partial h_m}{\partial z}\right) + \frac{q_w}{A} \\ &- \frac{\partial}{\partial z}\left(\frac{\alpha\rho_d\rho_c}{\rho_m}\Delta h_{dc}V_{dj}\right) + \left(v_m + \frac{\alpha(\rho_c-\rho_d)}{\rho_m}V_{dj}\right)\frac{\partial p_m}{\partial z}\end{aligned} \tag{4}$$

The subscript $_m$ denotes a mixture-averaged quantity, while $_d$ and $_c$ denote the diffused (gas) phase and continuous (liquid) phase, respectively. Thermodynamic properties (k_c, k_d, c_{pd}, c_{pc}, μ_c, μ_d, Δh_{dc}) are fixed to a reference state of saturated water at 5 MPa. The reference state is denoted by $_{ref}$. Thermodynamic properties of the mixture are calculated as in [1].

A compressible formulation is used to calculate the void fraction, phase densities, and pressure from the conserved states. First, we construct $\alpha_{ref} = \alpha\rho_d/\rho_{d,ref}$, an approximate void fraction to determine whether the flow is in a regime of "large" or "small" void fraction. For small void fractions ($\alpha_{ref} < 0.5$), the liquid is assumed to be compressible and ρ_d is assumed to be the reference density. Pressure is calculated from a linear fit to NIST data for the liquid, $p(\rho_c)$. Conversely, for large void fractions, the gas is assumed to be compressible, $\rho_c = \rho_{c,ref}$, and $p = p(\rho_d)$.

Two correlations were implemented for the drift velocity, one from Ishii and Hibiki (for undistorted bubbles) and the Chexal-Lellouche model [2]. A simple correlation for the covariance term, which is valid across many flow regimes, was found in Ishii and Hibiki. Wall friction was included (the f_m term in 3) using a mixture-Reynolds number and an explicit form of the Darcy friction factor from [3]. The wall heat term q_w is the heat added per unit length of the pipe.

The pressure term in the enthalpy equation 4 is calculated by treating the pressure gradient as a new state variable. Thus, the term is considered a source term in the equations. The gas source term in 2 is calculated using subcooled and bulk boiling models. The subcooled boiling model is based on the model in the PATHS thermal-hydraulics code. In PATHS, the void fraction is computed directly from various heat transfer correlations for multiphase flow. This void fraction is considered a "target" void fraction that the code should reach. The source term is then made to bring the void fraction toward this target,

$$\Gamma_d = \rho_d v_m \frac{\partial \alpha_{\text{target}}}{\partial z}. \tag{5}$$

As with the pressure, the gradient of α_{target} is treated as a new state variable. For the case

of bulk boiling, the mass transfer is computed using the heat transfer deficit. Interphase heat transfer is computed using standard correlations for Nusselt numbers.

Although time-dependent simulations are possible, the current work focuses on steady state solutions.

3.2 Primal Discretization

The 4-equation model is discretized using the Discontinuous Galerkin method, which allows for high-order solutions and fast convergence. The above equations are re-cast into the following general form

$$\begin{aligned} \frac{\partial U}{\partial t} + \frac{\partial}{\partial z} F(U, Q, \tilde{Q}) &= S(U, Q, \tilde{Q}) \\ \frac{\partial U}{\partial z} - Q &= 0 \\ \frac{\partial \tilde{U}}{\partial z} - \tilde{Q} &= 0 \end{aligned} \tag{6}$$

Here, U are the state variables ρ_m, $\alpha\rho_d$, $\rho_m v_m$, $\rho_m h_m$. The auxiallary state variables $\tilde{U}$, for which derivatives are computed, are pressure and target void (p and α_{target}). The derivatives of the auxiliary states are denoted by $\tilde{Q}$.

The 1D domain is discretized into elements, in which the solution is assumed to be a polynomial. At the interfaces between elements, a numerical flux must be computed since the solution is discontinuous at the interface. Here, the inviscid fluxes (only functions of U) are upwinded using the flux Jacobian of the average state. Since the problem is small (only four state variables), the complete eigen-decomposition of the flux Jacobian can be performed cheaply. Note, the eigen-decomposition was not done analytically due to the many complicated correlations in the inviscid fluxes (covariance, drift velocity, etc.). The Lax-Friedrichs flux, which is more stable but less accurate, was also implemented. In addition, it was also verified that when the covariance and drift velocity are ignored, the wave speeds (eigenvalues of the flux Jacobian) correspond to previously published results on two-phase flow (see [4]), with values of $v_m, v_m, v_m + a, v_m - a$, where the speed of sound a can be computed from the pressure-density relation $p(\rho)$.

The viscous fluxes are determined by first solving for the gradient quantities Q and $\tilde{Q}$ using a local discontinuous Galerkin (LDG) method for a linear problem. The DG scheme then computes averages of Q and $\tilde{Q}$ that are used to evaluate the viscous fluxes at the element boundaries. Finally, all of the source terms are added to an overall residual. Using exact derivatives almost everywhere[1], under-relaxed Newton iterations drive the residual to zero. When Newton iterations fail to converge, BDF pseudo-time stepping is used since it is more stable.

Multiphase models are almost always difficult to converge, and the drift-flux model presented here is no exception. In order to find a solution, a sequence of solves are performed with

[1]The complex-step method is used to find the derivatives of the eigenvectors/values with respect to the state, since this is quite difficult to do analytically

increasingly complex models. For example, the initial solve neglects all source terms except gravity, and sets diffusion coefficients, covariance, and drift velocity to zero. The solution is recomputed as models are slowly incorporated, using the most recent solution as an initial guess. If one is concerned with adaptation or uncertainty quantification, then this is not a significant hindrance, since one can often start with a baseline, converged solution and work from there to either adapt the mesh or tweak parameter values.

3.3 Adjoint Discretization

The continuous adjoint equations for a complex multiphase model such as this one have rarely, if ever, been derived. It has been shown, however, that the above DG discretization is asymptotically adjoint consistent[2][3]. Thus, the discrete adjoint method will suffice to compute an accurate adjoint solution. The adjoint for an output $J(U)$ is then

$$\Psi = \left(\frac{\partial R^T}{\partial U}\right)^{-1} \frac{\partial J}{\partial U}, \tag{7}$$

where $R(U)$ is the DG residual.

3.4 Error Estimation

The dual-weighted-residual method is used for output error estimation. Local residuals of the solution are weighted with the adjoint to determine their effects on the output J. The error estimates for each test function are summed to compute an estimate of the output error (the absolute values can be summed for a conservative error estimate). The residuals and adjoint are computed with respect to a fine discretization; in this way, the discretization error of the solution is approximated.

$$J_{\text{true}} - J(U) \approx \sum_{\text{elements}} \Psi_{\text{fine}}^T R_{\text{fine}}(U)$$

In particular, the error associated with p (order) or h (element size) refinement can be computed. The "fine" adjoint is computed by either solving exactly, approximately (linearizing about a smoothed U), or simply by smoothing the original Ψ with a block-implicit smoother. Currently, only h-refinement and simple smoothing of Ψ are used.

A cheaper error estimate can be computed by using the original adjoint solution rather than a fine version; the original Ψ is just injected into the fine discretization (Ψ_{inj}). This is often called the "computable correction." Then, the "remaining error" can be estimated in

[2]The discretization is only asymptotically adjoint consistent because of the use of the auxiallary state variables $\tilde{U}$; otherwise it would be fully adjoint consistent.

[3]Boundary conditions for the DG method, while not discussed, require great care to keep the method adjoint consistent.

this way:

$$
\begin{aligned}
J_{\text{true}} - J(U) &= \sum_{\text{elements}} \Psi_{\text{fine}}^T R_{\text{fine}}(U) = \sum_{\text{elements}} \left[\Psi_{\text{inj}} + (\Psi_{\text{fine}} - \Psi_{\text{inj}})\right]^T R_{\text{fine}}(U) \\
\epsilon_{\text{comp}} &= \sum_{\text{elements}} \Psi_{\text{inj}}^T R_{\text{fine}}(U) \\
\epsilon_{\text{rem}} &= \sum_{\text{elements}} (\Psi_{\text{fine}} - \Psi_{\text{inj}})^T R_{\text{fine}}(U)
\end{aligned} \tag{8}
$$

Thus, in cases when strict Galerkin orthogonality does not hold, one can compute a "corrected" output, $J_{\text{corr}} = J(U) + \epsilon_{\text{comp}}$, with the associated error ϵ_{rem}.

3.5 Adaptation

In this work, the metric used for adaptation is $|\epsilon_{\text{comp}}|$. A simple growth rate in the number of elements is prescribed, and the elements with the most error are refined by splitting those elements in half. Thus, at each iteration the number of elements $n_{t,i}$ grows like $n_{t,i+1} = (1+\gamma)n_{t,i}$, where γ is the fraction of elements refined. Currently, elements are refined at most once per adaptive iteration, and coarsening is not used. These choices make the adaptive process more stable at the cost of extra computation in some cases.

In order to alleviate some of the extra computation due to refining elements only once per iteration, a modification of the above strategy was used. Suppose there were n_{most} elements that contributed over 90% of the error in the total error measure. If $n_{\text{most}} < \gamma n_{t,i}$, i.e. most of the error is concentrated in a small number of elements, only n_{most} elements are refined. This is analogous to allowing multiple levels of refinement within one iteration, except that here the error estimates are recomputed after the mesh is refined at each level. For the results presented in this work, $\gamma = 0.3$.

4 Results

4.1 Test Problem

The code is set up to simulate water flowing upwards through a 2cm diameter, 5m long heated pipe as shown in Figure 1. Physical properties are taken from the NIST database (webbook.nist.gov) for isothermal water at 5MPa. Heat is applied to a portion of the tube according to

$$
q_w(z) = 10 \left[\frac{\text{erf}((z-1)) - \text{erf}((z-3))}{2}\right] 2\pi r_{\text{tube}} \tag{9}
$$

The heat profile is a smoothed version of a discontinuous heat source for $1 < z < 3$. The inflow enthalpy is set just below saturation, ensuring that boiling occurs and that all of the bubbles do not condense. The output of interest J is chosen to be the twice the average void

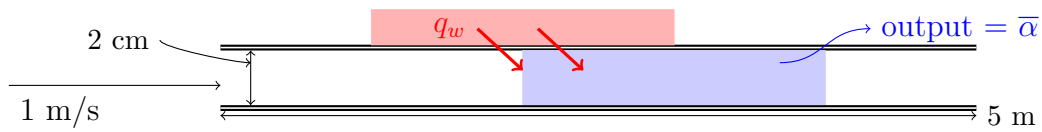

Figure 1: Problem setup for sample solutions. Note, for the sake of a compact diagram, the pipe is depicted as horizontal even though the flow is upwards (the z direction is to the right in the picture).

fraction for a portion of the pipe in and upstream of the heat source:

$$J = \int_2^4 \alpha(z) dz \tag{10}$$

4.2 Primal Solutions

The solution to the test problem is shown in Figure 2. The plots are consistent with some bubbles being produced in the tube mostly around $z = 2$. The density decreases as the bubbles are produced, and the velocity increases due to bouyancy. The temperature decreases initially because of some inaccuracies in the subcooled boiling model in which too much void is produced at low temperatures. The excess void acts as a heat sink on the liquid. Note, the saturation temperature is 537.09 K.

The convergence criterion is based on the infinity norm of the relative residual $|R|_{\max}/|U|_{\max}$. The convergence tolerance was set to 10^{-7} because it was found that the relative residuals oscillate around 10^{-7} after many iterations. Since this corresponds to a decrease of around 7 to 9 orders of magnitude, the solutions are well-converged even though the residuals do not reach machine zero (around 10^{-14} here). In any case, the output convergence is still at or above the expected rates. Figure 3 shows that the $p = 0$ solution (corresponding to a first order finite volume method) converges even faster than first order (black line on the plot). The higher order $p = 1, 2, 3$ solutions converge at the expected rate of $(\Delta z)^{p+1}$.

4.3 Adjoint Solutions

The adjoint solution for the output J is shown in Figure 4. The adjoint has four components because it is the sensitivity of the output to the residuals in the (four) governing equations. The output is not directly dependent on the second two states, $\rho_m v_m$ and $\rho_m h_m$. Therefore, the output is not sensitive to residuals in the momentum or energy equations. This is shown in the bottom plots of Figure 4. The upper plots show that the output is sensitive to the accuracy of the solution for $0 < z < 4$. This range includes both the range over which the output has support ($2 < z < 4$) and also the upstream portion $0 < z < 2$. However, solution accuracy downstream of the output support does not affect the output, as one would expect.[4]

[4]The speed of sound in air-water mixtures is generally quite small compared to the speed in pure substances. Indeed, the speed of sound in this example is always less than the mean velocity, so there are no backwards

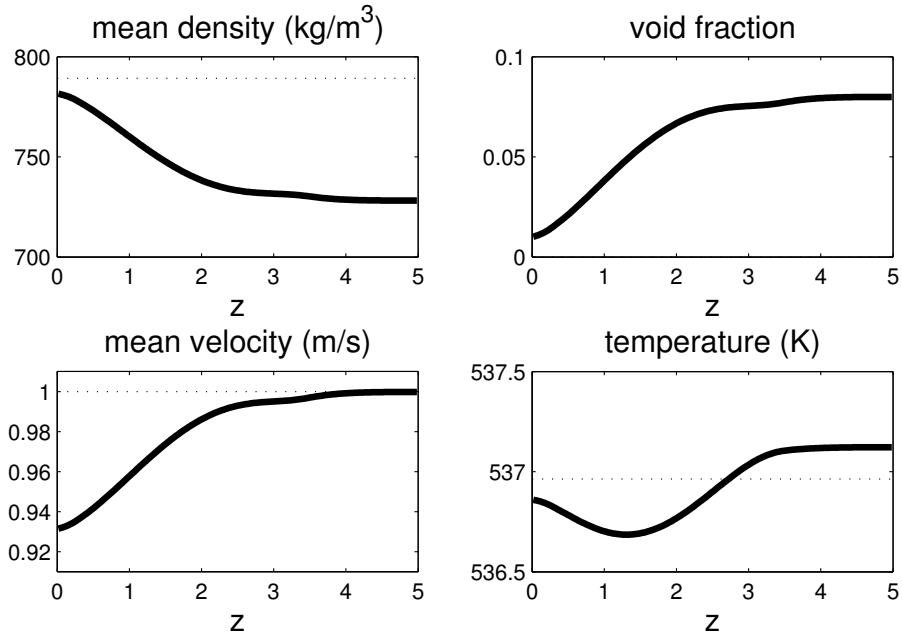

Figure 2: Primal solution for the test problem. Dotted lines indicate liquid density (top left), nominal velocity (bottom left), and inlet temperature (bottom right). Mesh has 16 p=2 elements.

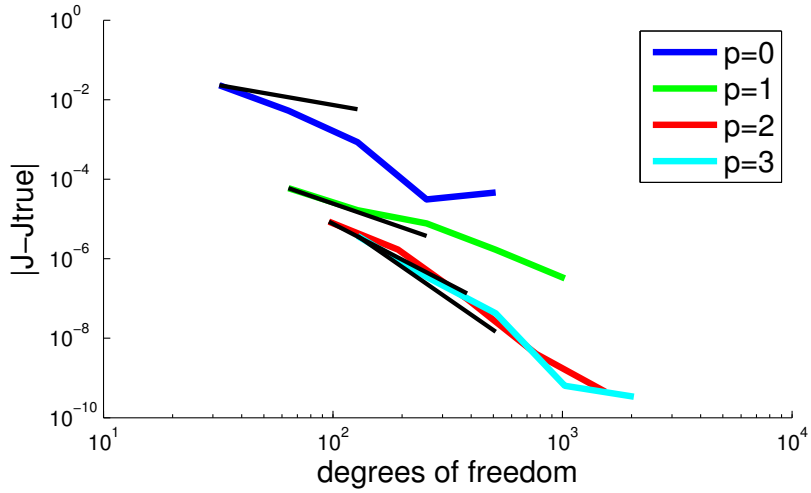

Figure 3: Mesh convergence of J for the test problem. Black lines indicate the expected convergence rate of $p+1$ for the DG method. The relative residual tolerance for each run was 10^{-7}. Meshes ranged from 8 to 256 elements.

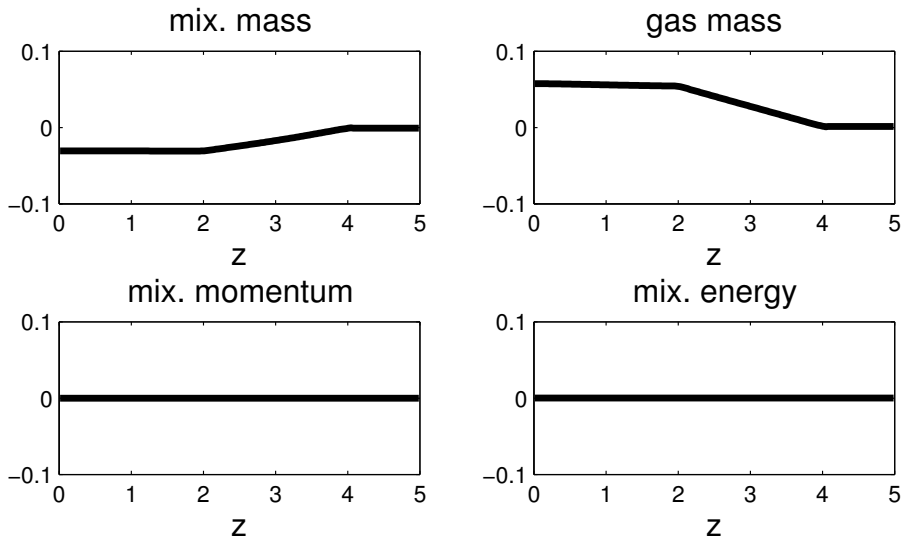

Figure 4: Four components of the dual solution for the test problem. Mesh has 16 quadratic elements (p=2).

The adjoint can be used to compute errors in the output. Thus, the accuracy of the adjoint can be computed by perturbing the solution $\tilde{U} = U + \Delta U$ and re-evaluating the output. Figure 5 shows that the error in the adjoint-based estimate of $J(\tilde{U}) - J(U)$ converges as $O((\Delta U)^2)$. This is expected because the adjoint captures linear sensitivity information, leaving a quadratic error term:

$$\begin{aligned} J(U + \Delta U) - J(U) &= \sum_{\text{elements}} \frac{\partial J}{\partial R}\frac{\partial R}{\partial U}\Delta U + O((\Delta U)^2) \\ &= \sum_{\text{elements}} \Psi^T \left(R(U + \Delta U) - R(U)\right) + O((\Delta U)^2) \end{aligned} \tag{11}$$

4.4 Error Estimates

The adjoint can be used to estimate more than just ΔJ due to arbitrary changes in the solution. By taking a particular "solution change," the adjoint can estimate the numerical error in the simulation. This is described above in Equation 8. The error estimates for the $p = 2$ meshes (number of elements from 8 to 256) are shown in Figure 6. The original output is shown in blue, with the corrected version J_{corr} in green. The "true" output was calculated on a much finer mesh with $p = 3$. The conservative estimate of ϵ_{rem} is the dashed green line[5]. As expected, the corrected output has less error than the original output and the error

propagating waves.

[5]The conservative estimate is $\sum |R^{\text{fine}}(U)(\Psi^{\text{fine}} - \Psi^{\text{inj}})|$, where the absolute values prevent error cancellation.

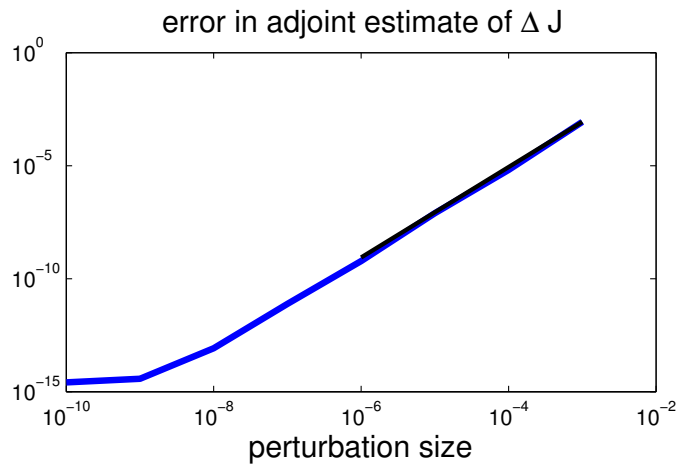

Figure 5: Accuracy of adjoint-based estimate of the output error. Black line is second order. Points are averages of 50 random perturbations.

estimate is larger than the actual error in the solution. In addition, all quantities converge as the mesh is refined.

4.5 Adaptively Refined Meshes

The adaptation scheme presented above was used on the test problem. The output convergence is shown in Figure 7, alongside the uniform-refinement convergence results. The adaptive process does work and in this case is generally no better or worse than uniform refinement. It is expected that the adaptive process will show more significant benefit when used on higher-dimensional problems with more localized outputs of interest. In addition, the convergence of the corrected output J_{corr} is shown in Figure 8. The corrected output has lower error than the uncorrected version but converges at approximately the same rate.

5 Conclusions and Ongoing Work

In this milestone we have applied adjoint-based error estimation and mesh adaptation to a one-dimensional drift-flux model discretized using the discontinuous Galerkin finite element method. While the implementation of the discrete adjoint was relatively straightforward after computing the required derivatives, some using the complex step method, we had to be careful of adjoint consistency considerations when making choices in the primal discretization. Sensitivity calculations showed that the adjoint was computed correctly. In addition, we verified the adjoint-weighted residual error estimates by comparing to the true error in an adaptive simulation. Finally mesh adaptation driven by the error indicator did lead to a decrease of the error, albeit not at an improved rate compared to uniform refinement. Possible reasons for this lack of improvement include the use of the computable correction to drive the

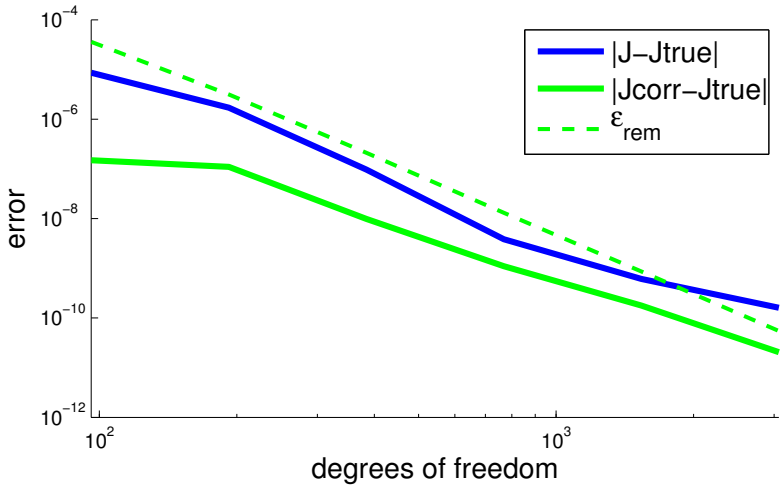

Figure 6: Convergence of output with mesh refinement for $p = 2$. Note the horizontal axis is degrees of freedom.

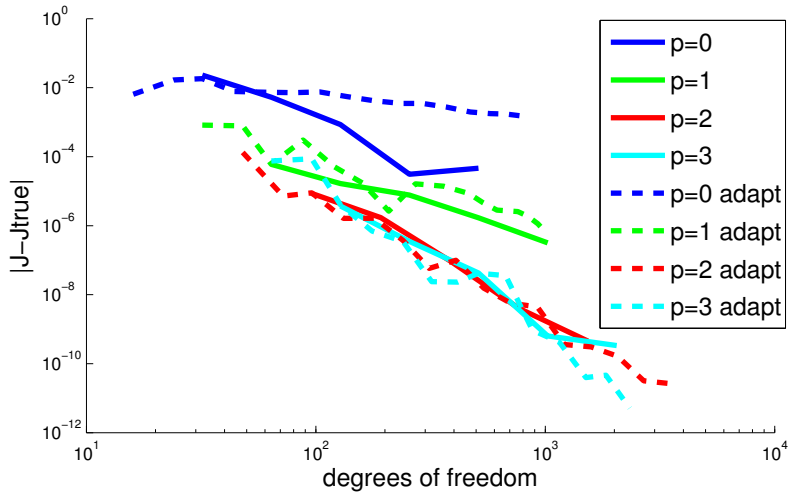

Figure 7: Convergence of J with uniform (solid) and adaptive (dashed) refinement. Note the horizontal axis is total degrees of freedom.

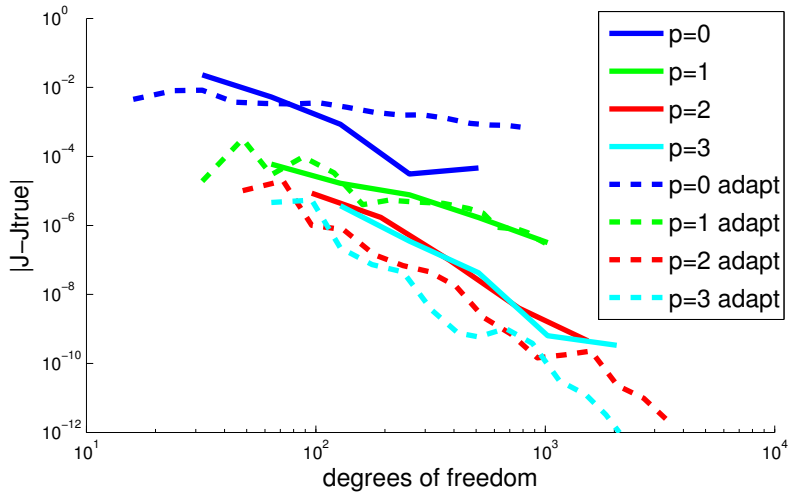

Figure 8: Convergence of corrected output J_{corr} with uniform (solid) and adaptive (dashed) refinement. Note the horizontal axis is total degrees of freedom.

adaptation, and more fundamentally, the restriction to a one-dimensional simulation. In multiple dimensions, regions important to a specific output will likely be more localized so that the adaptive refinement will potentially show significant advantages compared to uniform refinement.

References

[1] M. Ishii and T. Hibiki. *Thermo-Fluid Dynamics of Two-Phase Flow*. Springer Verlag, New York, 2006.

[2] Annalisa Manera, Horst-Michael Prasser, and Tim HJJ VAN DER HAGEN. Suitability of drift-flux models, void-fraction evolution, and 3-d flow pattern visualization during stationary and transient flashing flow in a vertical pipe. *Nuclear technology*, 152(1):38–53, 2005.

[3] Eva Romeo, Carlos Royo, and Antonio Monzón. Improved explicit equations for estimation of the friction factor in rough and smooth pipes. *Chemical engineering journal*, 86(3):369–374, 2002.

[4] Hervé Guillard, Fabien Duval, Jean Claude Latche, and Roxana Panescu. Numerical multiphase modeling of bubbly flows. *ANNALI DELL'UNIVERSITA'DI FERRARA*, 53(2):243–253, 2007.