A study of the effect of higher order spatial discretizations in SAMR (Structured Adaptive Mesh Refinement) simulations

Sophia Lefantzi [*], Christopher A. Kennedy, Jaideep Ray and Habib N. Najm
Sandia National Laboratories, Livermore, CA, 94550

Abstract

Generally, adaptively refined meshes aim to concentrate grid points in regions of interest while leaving the bulk of the domain sparsely meshed, thus reducing the computational time. Structured adaptively refined meshes in particular overlay grids of different refinement until the required accuracy is achieved. If the depth of the grid hierarchy is an issue (e.g. computational overheads) and has to be kept below a certain maximum, higher accuracy can be achieved either by increasing the mesh refinement factor or by using higher order spatial discretizations and interpolants. In this study, we present our findings for various orders of spatial discretizations as applied to two SAMR simulations: one for the FitzHugh-Nagumo equation and the other for a H_2-Air reaction-diffusion system. We investigate in detail the pairing of the orders of the spatial discretizations and the interpolants and their effect on the overall order of accuracy. We also demonstrate that higher-order SAMR approaches can be computationally more economical compared to second-order approaches.

Keywords : Structured Adaptive Mesh Refinement (SAMR), reacting flows high performance computing, higher order spatial interpolations.

1 Introduction

Essential to any computational method is the manner in which the continuous domain of interest is discretized into a mesh with many individual grid points. The mesh may be static or dynamic. A static mesh is established at the beginning of the computation and does not change afterwards, whereas a dynamic mesh changes during the computation based on the features of the result as the computation progresses. In general if a computation has features that are much smaller than the overall scale of the problem and change in time, there are two options in order to capture these features accurately: the first is to increase the number of static grid points in the domain with a grid density that will accommodate adequate resolution of the finest structures that will appear in the computation, thus creating a fine uniform mesh; the second is to track the small scale features and adopt a dynamic-adaptive griding scheme (Adaptive Mesh Refinement). The advantages of an adaptive meshing scheme compared to a static grid approach are: increased computational savings; increased storage savings; reduced amount of manual labor required for mesh regenerations, since the most efficient discretization is not known in advance and changes with the evolution of the computation. However, adaptive approaches have their drawbacks.

Structured Adaptive Mesh Refinement (SAMR) [1, 2] is a multiscale algorithm for Cartesian meshes employed to achieve higher grid resolution locally, in regions of the domain where it is required. Numerical simulations of reacting flows, multidimensional simulations in particular, can be very expensive in terms of computing power. Structured adaptively refined meshes overlay grids of different refinement until the

[*] Corresponding author; slefant@ca.sandia.gov

Paper **03F-67** ; Proceedings of the Fall Meeting of the Western States Section of the The Combustion Institute, October 20–21, 2003, Los Angeles, CA.

required accuracy is achieved, thus reducing the total number of mesh points compared to uniform mesh approaches. The SAMR method can be summarized as follows: a coarse Cartesian mesh is overlaid over a rectangular domain, and based on a suitably defined error the grid points which require further refinement are identified. These grid points are flagged and collated into rectangular children patches on which another more dense Cartesian mesh is imposed. The refinement factor between the parent and the child mesh is usually kept constant for a given problem. In our paper, we will assume the refinement ratio to be two. The more accurate solution from the finest meshes is periodically interpolated onto the coarser ones.

Although SAMR techniques are fairly simple and straightforward, in practice there are issues associated with their implementation. For example, in explicit schemes the global timestep is restricted by the stability criterion on the finest mesh resulting in an inefficient very small global timestep. To address this issue one can use time-refined explicit schemes [1] with hierarchies of patches, where each patch has a different stable timestep. With this scheme, in order to integrate for a time interval equal to the largest stable timestep (that of the coarsest patches) a recursive integration of all the patches is needed, and most of the time is spent on the finest patches. This sub-cycling increases the complexity of time-integration scheme substantially. Furthermore, time-refined integration incurs significant recursion overheads if the grid hierarchy is deep (large number of refinement levels). Interpolations from coarse-to-fine patches (and vice versa) also add a significant cost to SAMR approaches, especially in 3D. The criteria for adaptation (refining and coarsening regions in space) pose a challenge in most adaptive simulations. While a rigorous approach based on Richardson extrapolation is outlined in [2], it is expensive and is frequently replaced by heuristic work-arounds, thus creating a degree of uncertainty in the accuracy of the solutions. Another important problem arises from parallel utilization of SAMR, that of load balancing. There is a requirement of keeping parents and children together (on the same processor) and that results in poor domain-partitioning which deteriorates with deep hierarchies. There are two possible approaches in order to avoid deep hierarchies: the first is to increase the refinement factor (*e.g.* instead of having a refinement factor of two have a factor of four), thus reducing the required levels of refinement; the second is to implement a higher order method to discretize the spatial derivatives thus obtaining higher accuracy with coarser grids, *i.e.* with less levels of refinement. The first approach reduces the hierarchy depth but does not necessarily reduce the number of grid points, therefore we chose to explore the second approach.

2 Problem Description

Higher order approaches utilize a given grid resolution far more efficiently than lower order ones, but they do have their drawbacks. Higher-order stencils involve more computations to evaluate a derivative. While this is relatively inconsequential with derivative operators, interpolant operators are evaluated by using a linear combination of as many as p^d (where p is the order of accuracy and d is the dimensionality) values and their respective coefficients, thus forming a potential source of costs, especially in 3D. Further, derivative schemes above fourth-order may at best be closed to a lower order at the domain boundaries to remain time-stable. Finally, higher-order derivative schemes need to be coupled with higher-order prolongation and restriction operators to preserve the accuracy at all levels of the grid hierarchy.

In this paper we address some of the issues pertaining to the practical use of higher-order methods in simulations. While higher-order methods can be used in a cell-center or nodal manner, we will adopt the latter. We begin with a summary of higher-order discretizations, interpolants and filters. Details can be obtained from [3]. Then we obtain higher order (4th) convergence on a SAMR mesh and address the issue of the correct pairing of order of discretization and interpolations. We investigate whether using higher order approaches in SAMR settings offer any economy in simulations, *i.e.* whether the higher computational costs (on a per-grid-point basis) of a higher-order simulation is offset by the savings in the number of grid points. These problems are studied within the framework of the FitzHugh-Nagumo equations where an analytical solution exists. Our method is then used to study a stiff, non-linear reaction-diffusion problem involving detailed H_2-Air chemistry.

2.1 Higher Order Spatial Discretizations

A high-order spatial discretization for vertex-centered AMR with a refinement factor of two requires derivative, coarse-to-fine interpolant, and filter operators. We investigate 4^{th}- and 6^{th}-order discretizations.

All derivatives used in this work will use first-derivative operators, repeatedly if necessary. A fourth-order derivative will be of the form

$$\begin{aligned} f'_i &= \frac{c_L f_{i-3}}{(\Delta x)} + \frac{b_L f_{i-2}}{(\Delta x)} + \frac{a_L f_{i-1}}{(\Delta x)} + \frac{\Upsilon f_i}{(\Delta x)} \\ &+ \frac{a_R f_{i+1}}{(\Delta x)} + \frac{b_R f_{i+2}}{(\Delta x)} + \frac{c_R f_{i+3}}{(\Delta x)}. \end{aligned} \tag{1}$$

Several examples of varying accuracies are listed in Table 1

S	c_L	b_L	a_L	Υ	a_R	b_R	c_R	L.O.T.E.
2E	0	0	-1/2	0	1/2	0	0	-(1/6)ξ^3
3U	0	0	-1/3	-1/2	1	-1/6	0	-(1/12)ξ^4
4U	0	0	-1/4	-5/6	3/2	-1/2	1/12	+(1/20)ξ^5
4E	0	1/12	-2/3	0	2/3	-1/12	0	-(1/30)ξ^5

Table 1: Stencil coefficients for centered explicit (E) and upwinded (U) first-derivative operators of different orders of accuracy on uniform grids. S is the name of the stencil. L.O.T.E is an abbreviation of Leading Order Truncation Error and ξ is a scaled wavenumber. $O(p)$ stencils have $O(\xi^{p+1})$ LOTE.

At the domain boundaries, the derivatives will be closed, to lower order using skewed stencils. Table 2 lists several possible stencils.

c_L	b_L	a_L	Υ	a_R	b_R	c_R	L.O.T.E.
0	0	0	-1	1	0	0	-(1/2)ξ^2
0	0	0	-11/6	3	3/2	-1/3	+(1/4)ξ^4
0	0	-1/3	-1/2	1	-1/6	0	-(1/12)ξ^4
0	0	-1	1	0	0	0	+(1/2)ξ^2
-1/3	3/2	-3	11/6	0	0	0	-(1/4)ξ^4
0	1/6	-1	1/2	1/3	0	0	+(1/12)ξ^4

Table 2: Stencil coefficients of skewed lower-order, first-derivative operators near boundary points. These are to be used with second- and fourth-order derivative discretizations.

2.2 Interpolants

SAMR simulations interpolate data from coarse meshes to a halo of points around their finer children (*prolongation*) and then *restrict* accurate solutions from the finer children onto themselves. In 2D, interpolation blocks use a four-squared block of coarse points and interpolate to the center of a cell. At domain boundaries, these interpolation blocks are skewed relative to the interpolation point. For fourth-order interpolants, there are 3 unique stencil coefficients; $(i \pm 1/2, j \pm 1/2) = 81/256$, $(i \pm 1/2, j \pm 3/2) = (i \pm 3/2, j \pm 1/2) = -9/256$, and $(i \pm 3/2, j \pm 3/2) = 1/256$. At higher order, the number of unique stencils increases.

2.3 Filters

Filtering [4] is a way to cleanly remove high-wavenumber information from the grid. There are two compelling reasons to filter. The first is that any finite difference numerical method has limited accuracy and higher-wavenumber information that is unresolvable by the numerical method needs to be removed before it interferes with the resolved wave-numbers. Secondly, in order to avoid failure of the interpolant operators, no wavenumber that represents less than approximately six grid points per wavenumber may be present [5]. Filters are an important element in the overall numerical method.

Details on higher-order SAMR discretizations, interpolants and filters are in [3].

2.4 Strategy

We solve two problems with higher order spatial interpolations. The first is the FitzHugh-Nagumo equation which has an analytical traveling wave solution. We use $2^{\rm nd}$, $4^{\rm th}$ and 6^{th} order discretizations on SAMR meshes to obtain the theoretical convergence rate. A variety of interpolants is used to determine the appropriate (discretization, interpolant) pair, and a computational cost analysis is performed. The second problem we address is a stiff reaction-diffusion problem. We first perform a convergence analysis with respect to the grid in order to identify the mesh resolution required to resolve the problem. Then we use SAMR with both $2^{\rm nd}$ and $4^{\rm th}$ order approaches to solve the same problem with enough levels of refinement to obtain the resolution indicated by the grid convergence study. We then compare the two solutions.

3 Test cases

3.1 FitzHugh-Nagumo equation

We solve the following 1D problem:

$$U_t = DU_{xx} + AU(1-U)(U-\alpha) \tag{2}$$

subject to the boundary condition $U_x = 0$ at $x = \pm\infty$ This equation admits an analytical traveling wave solution

$$U(\xi, t) = \frac{1}{2}(1 + \tanh\frac{\xi}{2\epsilon})$$

where $\xi = x + st$,

$$\epsilon = \sqrt{\frac{2D}{A}} \text{ and } s = \sqrt{\frac{AD}{2}}(1-2\alpha)$$

Details of the derivation are in [6]. We chose $D = 1.0, A = 2 \times 10^4$ and $\alpha = 0.3$,giving us $\epsilon = 10^{-2}$ and $s = 40$.

The problem is solved, numerically for $0 \leq x \leq 1$ domain discretized with 100 points. The front is initialized at $x = 0.5$. The region $0.26 \leq x \leq 0.61$ is covered by a Level1 patch and $0.3 \leq x \leq 0.55$ by a Level2 patch. The problem is integrated, in a time-refined manner, using a second-order explicit Runge-Kutta scheme (Heun's method) with a Δt on the coarse mesh of 6.25×10^{-7}, chosen so that the $O((\Delta t)^2)$ temporal errors are far smaller than the spatial errors. The problem is integrated up to $t = 1.25 \times 10^{-3}$, corresponding to a 5ϵ traversal of the wave. Fig. 1 shows the profile on the Level0 mesh at the start and end

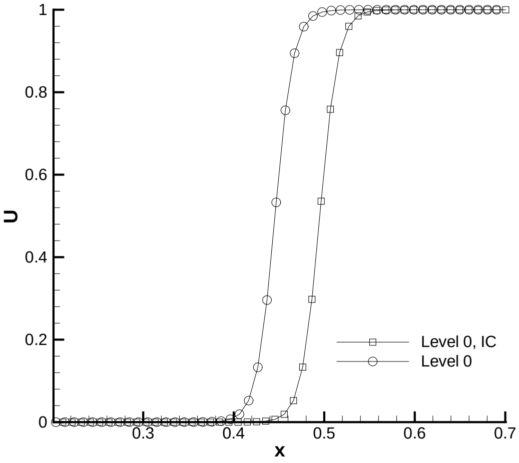

Figure 1: Profile of U at $t = 0, 1.25 \times 10^{-3}$ (graphs with □ and ∘, respectively) on the coarse mesh. The data at the Level0 mesh points are marked with symbols; the levels refine by a factor of 2. This particular run had 3 levels and the duration of integration was chosen to that the front would move a distance equal to 5ϵ, as shown.

of the run. We see that the front is defined by about 10 grid points on the coarse mesh and better as higher level patches are added. The refinement factor is 2.

In Fig. 2 we plot the RMS error with respect to the analytical solution on the individual levels on the grid hierarchy. The errors on levels 0 to 3 are indicated by □, ∘, △ and ⋄ while the ideal convergence is given by the solid lines. We see that second and fourth-order convergence is obtained. No filtering is done; with 10 coarse grid points in the front, Runge phenomenon was not observed. In Fig. 2 we see that, as expected, the fourth-order simulation is far more accurate than the second order one and the difference in accuracy increases dramatically with resolution. Further, a given level of error (a horizontal line in Fig. 2) can be achieved with a coarser mesh (fewer levels of refinement) when the fourth-order (as opposed to a second-order) approach is employed. Thus if numerical accuracy, to the exclusion of everything else, is the objective, a higher-order approach is the obvious choice.

The choice of a numerical scheme has to be tempered with its cost. We measured the floating point operation count on Intel Pentium III processors with 256 kB L2 cache and running at 1 GHz. The interpolation and discretizations were written in Fortran77 and compiled with Portland Group F77 compiler `pgf77`. The overarching control code was in C++ compiled with the GNU suite of compilers `gcc-2.96`. The problems were run on the IA-32 cluster at NCSA (`platinum.ncsa.uiuc.edu`) and the floating point operations were measured using PAPI [7].

Fig. 3 shows the total number of floating point operations count, normalized by the count from a 1-level second-order run, as a function of the error. The data was obtained from the runs plotted in Fig. 2. It is clear, that for this problem, the fourth order simulation is *cheaper* than the second order approach, and it gets progressively *more* economical as the error tolerance become stringent.

We note that if, for some reason, one chooses to operate in a marginally resolved manner (approaching the Trefethen and Weideman limit [5]), we see from the left extreme of Fig. 2 that the (large) errors achieved by the two approaches may very well be comparable. Further, judging from the right limit of Fig. 3, the higher-order approach would not offer much economy either.

We now address the issue of the correct choice of interpolant with a given discretization. Since U on

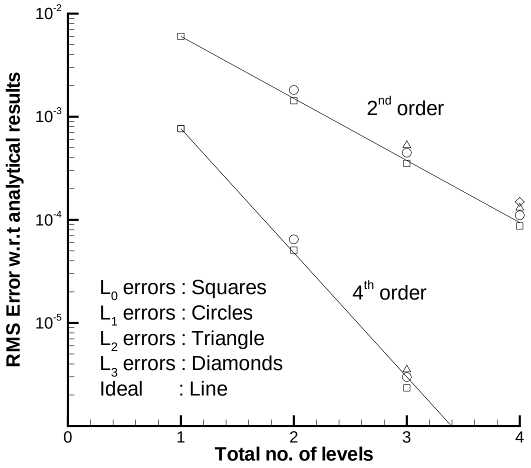

Figure 2: The RMS error on the individual levels, as the simulation is run on a 1-, 2-, 3- and 4-level grid hierarchy. The errors on levels 0 to 3 are indicated by □, ○, △ and ◇. The ideal behavior is given by the solid line. Results for both second and fourth order discretizations are shown. A 1-level grid hierarchy has an effective resolution of 100 in the $[0, 1]$ domain; 2-, 3- and 4-level grid hierarchies are correspondingly equivalent to a 200-, 400- and 800- grid point uniform mesh.

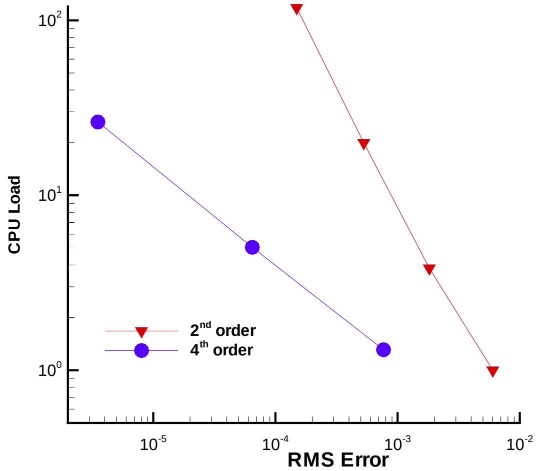

Figure 3: The computational load versus RMS error for the second and fourth order approaches. Results have been normalized by the computational load of a second-order, 1-level grid hierarchy run (1,208,728,001 floating point operations). It is clear, that for this problem, the fourth order simulation is *cheaper* than the second order approach, and it gets progressively *more* economical to use the higher-order approach for a given level of error as the error tolerances become stringent.

each grid is eventually the result of interpolation based on local polynomials, one cannot differentiate this interpolated data indefinitely. Interpolated data of order p_I which is differentiated k times will not preserve its original accuracy and consequently care has to be exercised to ensure the the final differentiated term does not fall below the desired accuracy.

We solve the same problem, but only the region $0.4 \leq x \leq 0.49$ is covered by a fine mesh. As mentioned before (Sec. 1) interpolants are used to initialize a halo of cells at the edges of the fine mesh so that centered discretizations may be used everywhere. These interpolation errors manifest themselves in the Right Hand Side term of the equation. It thus becomes necessary that these interpolation-generated be comparable to the discretization error so that their effect is observable in the evaluation of the RHS. In order to do so, we terminate the fine mesh at $x = 0.49$ so that these interpolations are performed in a region of high gradient (and consequently large errors). The correct choice of interpolant will ensure that the order of convergence of the errors governed by interpolant is equal to or larger than that of the discretization. The error introduced by this interpolation changes as the wave moves into the fine mesh and the interpolation is done in a region of small gradient. Thus, during our tests, we will integrate for only a small period in time, so that the gradient at the edge of the fine mesh remains large. Further, the interaction of interpolation and discretization errors occurs only in the U_{xx} term, which in certain regions may be overwhelmed by the "reaction term", $AU(1-U)(U-\alpha)$. Thus, in order to measure convergence, we will measure the error in U_{xx} vis-a-vis the exact solution. Experiments will be done on a hierarchy with 2 levels. Errors will be calculated (for convergence testing) by varying the coarse mesh resolution.

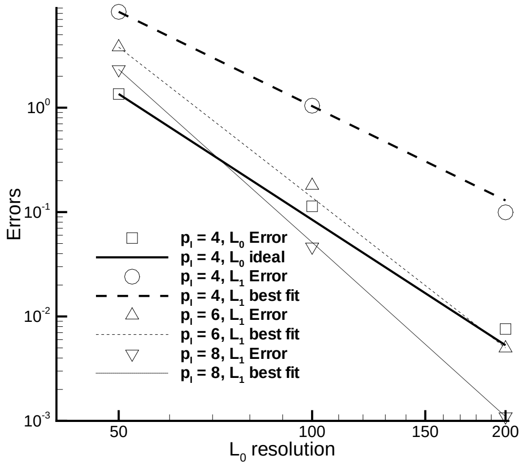

Figure 4: Above: We plot the error on L_0 and L_1, as a function of the coarse-mesh resolution. The ideal and actual convergence of the errors in U_{xx} are plotted in the heavy black line and □ respectively. 4^{th} order discretizations are used. Actual errors on L_1 mesh from a 4^{th} order interpolant run are plotted as ○ and the heavy dashed line is the best linear fit. The convergence (for ○) is less than 4^{th} order. △ and ▽ are the results for 6^{th} and 8^{th} order interpolants. The linear fits show a convergence greater than the coarse mesh, *i.e.* 4^{th} order.

In Fig. 4 and 5, we plot the convergence of the U_{xx} term as a function of the resolution of the coarse mesh when different discretization and interpolant pairs are used. For fourth order interpolants (Fig. 4) 6^{th} and 8^{th} order interpolants are seen to preserve the order while for 6^{th} order discretizations, only 8^{th} order interpolants were found to be adequate. We did not use odd-ordered interpolants to restrict ourselves to dissipative errors. Thus, for problems where the largest spatial derivative is of order 2, the sufficient

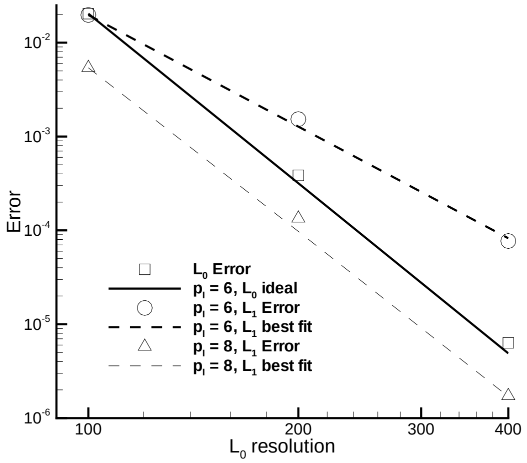

Figure 5: Above: We plot the error on L_0 and L_1, as a function of the coarse-mesh resolution. The ideal and actual convergence of the errors in U_{xx} are plotted in the heavy black line and □ respectively. 6^{th} order discretizations are used. Actual errors on L_1 mesh from a 6^{th} order interpolant run are plotted as ○ and the heavy dashed line is the best linear fit. The convergence is less than 4^{th} order. △ are the results for 8^{th} order interpolants. The linear fit shows a convergence greater than the coarse mesh, *i.e.* 6^{th} order.

condition for ${p_D}^{th}$ order convergence appears to be $p_I \geq p_D + 1$.

3.2 Hydrogen-Air reaction/diffusion

We now investigate whether the benefits of a higher-order approach (Sec. 3.1) are applicable to more complex, coupled nonlinear problems.

Our ultimate aim is to design a code for the simulation of flames with detailed chemistry. We intend to solve the low Mach number approximation of the Navier-Stokes equation in open domains. We use SAMR so that we can resolve detailed flame structure for laboratory-sized unsteady flames. The code is being developed as CCA components [8] which allows a modular architecture.

As a preliminary model to this problem, we consider a reaction-diffusion problem.

$$\begin{aligned}
\frac{\partial \Phi}{\partial t} &= K\nabla \cdot (F \nabla \Phi) + R, \qquad (3)\\
\Phi &= \left\{ \begin{array}{c} T \\ Y_1 \\ \vdots \\ Y_M \end{array} \right\}, K = \left\{ \begin{array}{c} \frac{1}{\rho c_p} \\ \frac{1}{\rho} \\ \vdots \\ \frac{1}{\rho} \end{array} \right\}, F = \left\{ \begin{array}{c} \lambda \\ \rho D_1 \\ \vdots \\ \rho D_M \end{array} \right\}, \\
R &= \left\{ \begin{array}{l} \frac{1}{\rho c_p}[-\sum_{i=1}^{N} h_i \omega_i] \\ \frac{1}{\rho}\omega_1 \\ \vdots \\ \frac{1}{\rho}\omega_M \end{array} \right\}, \sum_{i=1}^{N} Y_i = 1
\end{aligned}$$

where R is the thermo-chemical source terms while $K\nabla \cdot (F\nabla\Phi)$ is the diffusive transport terms. $M = N - 1$, where N is the number of chemical species. λ is the thermal conductivity and D_i are the diffusion coefficients. Pressure is assumed to be constant in time and space (*i.e.* burning in an open domain). A constant Lewis number model (different Le for different species) is used for the diffusion coefficients. We employ operator-splitting [9] using RKC to advance diffusion, and a BDF5 scheme [10] to integrate chemistry. The general procedure is similar to that in [11], except for the projection solution, which is omitted here, since we solve a set of ODEs rather than a set of DAEs.

We simulate the ignition of a stoichiometric $\mathrm{H}_2 - \mathrm{Air}$ mixture with a randomly varying initial temperature field. The field is sufficiently hot in certain localized regions so that the mixture ignites. A mechanism with 9-species and 19 reversible reactions [12] is used for the chemistry. The domain is $1\mathrm{cm} \times 1\mathrm{cm}$ with a 100×100 coarse level mesh. Three levels of refinement are allowed, with a refinement factor of two (*i.e.* 12.5 microns on the finest patches). For the explicit time integration we employ a second-order Runge-Kutta-Chebyshev scheme [3]. The global timestep size on the coarsest mesh is limited to 500 nanoseconds in order to reduce splitting errors. This is still larger, by an order of magnitude, than the timestep used with a traditional RK2 method adapted for time-refinement. We test two different orders of spatial discretizations: second-order derivative stencils with bicubic interpolations, with a formal accuracy of order two, and fourth order derivative stencils with sixth order prolongation operators, with a formal accuracy of order 4.

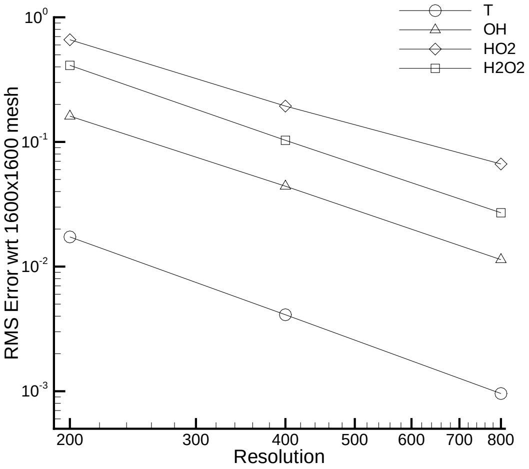

Figure 6: Convergence of difference in the solution w.r.t. grid refinement. Solutions were obtained on $200 \times 200, 400 \times 400$ and 800×800 uniform meshes and compared with solutions obtained by doubling the grid resolution. Temperature (∘), OH (△), $\mathrm{H}_2\mathrm{O}_2$ (□) and HO_2 (◇) mass fractions are compared. We see that the largest discrepancy is in $\mathrm{H}_2\mathrm{O}_2$ and HO_2 solutions. The 800×800 mesh corresponds to a 12.5 microns resolution.

We first identify a resolution sufficiently fine to resolve the solution. We solve the problem on a uniform mesh with $200 \times 200, 400 \times 400$ and 800×800 meshes and compare the difference in these solutions with the solution from a simulation with double the resolution (*i.e.*, 200×200 against a 400×400 run etc.). We plot the RMS difference for some of the variables in Fig. 6. We see that in the 800×800 case, the largest difference is around 6% so we consider this resolution (12.5 microns) acceptable. Henceforth we will refine to this resolution, with accuracy being achieved via higher-order discretizations.

In Fig. 7 we see some regions in the mixture ignite. The ignition fronts are seen to move through the mixture and annihilate sections of each other as they interact. The distribution of patches over different levels of refinement at two different time instants is also seen in Fig. 7. As certain regions (initially with large temperature gradients) fail to ignite and lose their gradients via diffusion, the fine patches are removed,

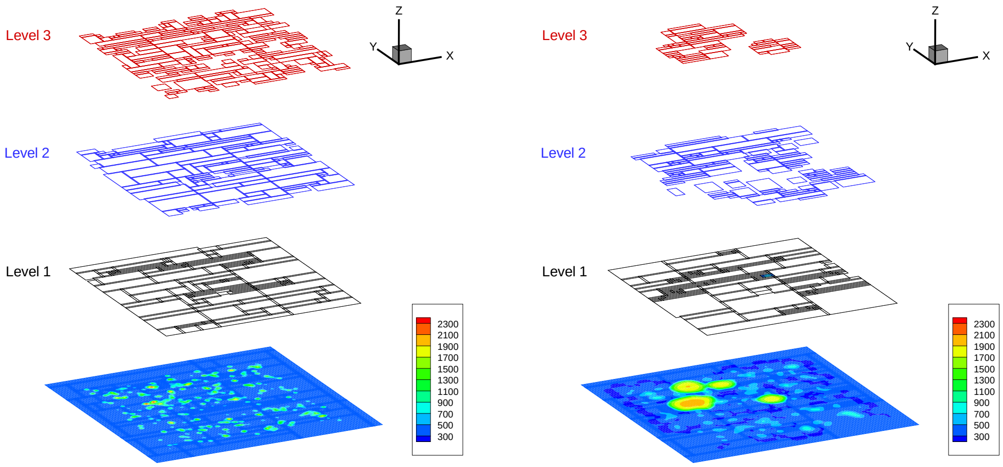

Figure 7: Temperature field at time t = 0 (above) and 90 microseconds (below). Patches at different levels in the grid hierarchy are plotted separately. We see the vastly different hierarchy configuration at the two instants in time. Regions which have ignited and have steep reaction fronts are resolved with Level 3 grids. Regions which failed to ignite have had the fine meshes removed by $t = 90$ microseconds.

leaving the region with coarser meshes. In Fig. 8 we see $\mathrm{H_2O_2}$ profiles plotted on the finest mesh. We see the resolution capability of the SAMR resolving the fine $\mathrm{H_2O_2}$ profile by about 10 points.

In Fig. 9 we plot the $\mathrm{HO_2}$ profiles extracted from the second and fourth order runs. We see that they are similar but are slightly offset along the y-axis, on the unburned side ($y \leq 0.535$ cm). The 20 grid points in the $\mathrm{HO_2}$ profile on the finest mesh indicates that the solution is adequately resolved. The fourth-order run required that an 8^{th} order filter be applied at the beginning of every timestep, a consequence of using the sixth order prolongation operator. Filtering was not required in a separate run (not shown here) where the sixth-order interpolant was replaced by a fourth order one. Filters were not required with 6th order discretizations in Sec. 3.1, probably because the front was quite well resolved even on the coarse mesh. In this particular case, where the fronts *develop* in time, the sixth order interpolants were observed to give unphysical answers when filtering is not applied, during the formative stages of the fronts (when steep, probably unresolved gradients, would exist especially in the coarser levels).

4 Conclusions

It is clear, from Sec. 3.1, that 4^{th} order approaches are achievable on SAMR meshes. Further, as shown in Fig. 3, a higher-order approach could even be *more* economical computationally vis-a-vis second order approaches. We have established, too, that a derivative stencil of order p_D needs to be paired with interpolation of order p_I, where $p_I > p_D$. In our specific case, where the order of the largest derivative was 2, $p_I = p_D + 2$ ensured the correct convergence for 4^{th} and 6^{th} order derivative stencils.

We employed this approach on a non-linear problem in Sec. 3.2. First, a convergence (w.r.t. grid refinement) analysis was performed to identify a grid resolution sufficient to resolve this particular solution. This resolution was then achieved, (far more economically) via SAMR using both second and fourth order strategies, drawing copiously from the experience gained with the FitzHugh-Nagumo problem. The simulation remained stable and a comparison of the 4^{th} order results with 2^{nd} order ones show a good match.

We are currently in the process of augmenting the reaction-diffusion problem with a convection term and the system of equations being solved with equations for the evolution of momentum. This, coupled with the projection scheme (Sec. 3.2), will create a fourth-order SAMR infrastructure for the DNS simulation of

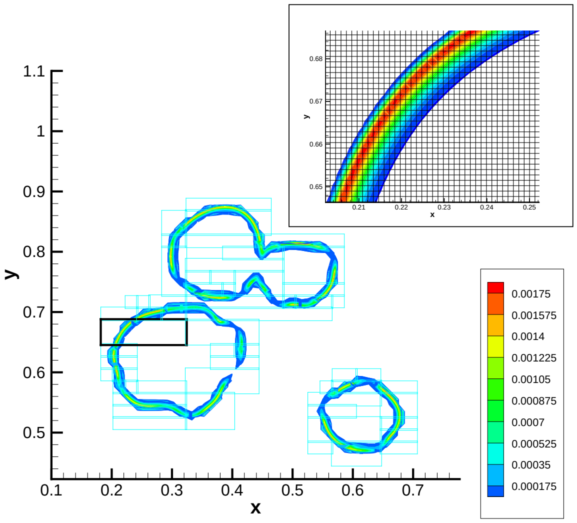

Figure 8: H_2O_2 profiles plotted only on the finest mesh level. Patches on Levels 3 are shown. Inset: A zoom of the patch in black is shown overlayed by the mesh. We see the H_2O_2 profile being resolved by about 10 grid points.

flames. Results from these simulations will be reported in future communications.

Acknowledgments

This work was supported by the Department of Energy, Office of Basic Energy Sciences, SciDAC Computational Chemistry Program and SciDAC Advanced Scientific Computing (ISIC) Programs.

References

[1] M. Berger and J. Oliger. Adaptive mesh refinement for hyperbolic partial differential equations. *Journal of Computational Physics*, 53:484–512, 1984.

[2] M.J. Berger and P. Colella. Local adaptive mesh refinement for shock hydrodynamics. *Journal of Computational Physics*, 82:64–84, 1989.

[3] J. Ray, C. Kennedy, S. Lefantzi and H. N. Najm. High-order spatial discretizations and extended stability methods for reacting flows on structured adaptively refined meshes. In *Third Joint Meeting of the U.S. Sections of The Combustion Institute*, Chicago, Illinois, March 2003. Distributed as on a CD.

[4] C. A. Kennedy and M. H. Carpenter. Several new numerial methods for compressible shear layer simulations. *Appl. Num. Math.*, 14:397–433, 1994.

[5] L. N. Trefethen and J. A. C. Weideman. Two results on polynomial interpolation in equally spaced points. *J. Approx. Theory*, 65:247–260, 1991.

[6] C. P. Fall, E. S. Marland, J. M. Wagner and J. J. Tyson, editors. *Computational Cell Biology*, volume 20 of *Interdisciplinary Applied Mathematics, Mathematical Biology*. Springer, 2001.

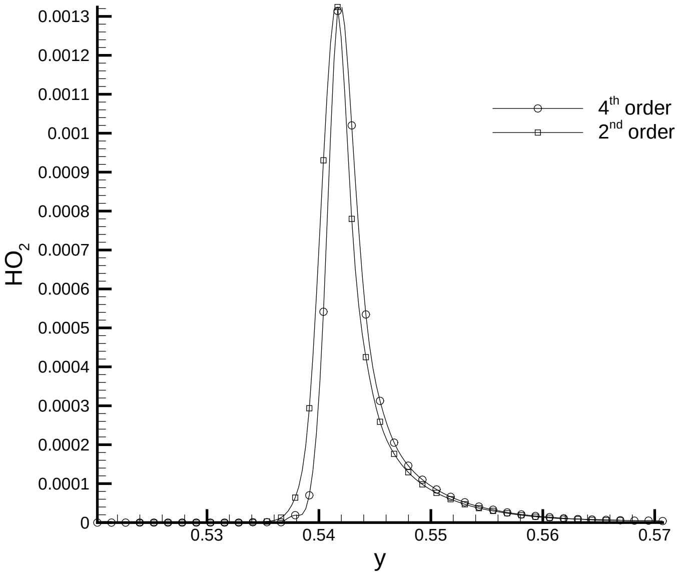

Figure 9: HO_2 profiles on the 12.5 micron mesh level using the second and fourth order approaches, at $x = 0.26$ cm. The results are for $t = 93$ microseconds. We can see that the profiles, though slightly offset along the y-axis, are similar. There are about 20 points in the profile on the finest mesh.

[7] S. Browne, J. Dongarra, N. Garner, G. Ho and P. Mucci. A portable programming interface for performance evaluation on modern processors. *International Journal of High Performance Computing Application*, 14(3):189–204, 2000. http://http://icl.cs.utk.edu/projects/papi/.

[8] S. Lefantzi, J. Ray and H. N. Najm. Using the common component architecture to design high performance scientific simulation codes. In *Proceedings of the International Parallel and Distributed Processing Symposium*, Nice, France, April 2003.

[9] G. Strang. On the Construction and Comparison of Difference Schemes. *SIAM J. Numer. Anal.*, 5(3):506–517, 1968.

[10] S. D. Cohen and A. C. Hindmarsh. Cvode, a stiff/nonstiff ode solver in c. *Computers in Physics*, 10(2):138–143, 1996.

[11] O.M. Knio, H.N. Najm and P.S. Wyckoff. A Semi-Implicit Numerical Scheme for Reacting Flow. II. Stiff, Operator-Split Formulation. *J. Comp. Phys.*, 154:428–467, 1999.

[12] F. L. Dryer, R. A. Yetter and H. Rabitz. A comprehensive reaction mechanism for carbon monoxide/hydrogen/oxygen kinetics. *Combust. Sci. and Tech.*, 79:97–128, 1991.