High-order spatial discretizations and extended stability methods for reacting flows on structured adaptively refined meshes

J. Ray[1], C. A. Kennedy, S. Lefantzi and H. N. Najm
Sandia National Laboratories, Livermore, CA 94551

January 26, 2003

Abstract

Structured adaptively refined meshes (SAMR) ensure efficient resolution of large computational domains by clustering mesh points only where required. We explore the possibility of using higher-order spatial discretizations in such a setting so as to achieve acceptable accuracies with relatively few levels of refinement and greater overall efficiency. To overcome the stringent temporal stability limitations posed by such discretizations when used with a time-integrator, we present an operator-split time-integration method using a Runge-Kutta-Chebyshev scheme on SAMR meshes with time-refinement. The complete formulation is tested on a reaction-diffusion problem, using a 9-species, 19 reactions H2-Air chemical mechanism.

Introduction

Structured Adaptive Mesh Refinement (SAMR) [1, 2] is a powerful technique for concentrating grid resolution in regions of steep spatial gradients. It is a particularly appealing approach to geometrically simple spatial domains. Briefly, the method consists of laying a relatively coarse Cartesian mesh over a rectangular domain. Based on some suitable metric, regions requiring further refinement are identified, the grid points flagged and collated into *rectangular* children patches on which a denser Cartesian mesh is imposed. The refinement factor between parent and child mesh is usually kept constant for a given problem. The process is done recursively, so that one ultimately obtains a hierarchy of patches with different grid densities, with the finest patches overlaying a small part of the domain. The more accurate solution from the finest meshes is periodically interpolated onto the coarser ones. Time integration of PDEs on such meshes can be done with both implicit [3] and explicit [2] schemes. Explicit schemes may be modified (time-refinement [2]) or else their global timestep is restricted by the stability constraint on the finest mesh. Typically, this modification involves the *recursive* integration of the patches, starting with the coarsest level (at a timestep determined by the stability constraint of the patch itself), followed recursively by the integration of its children patches, which are processed more often but at their (smaller) stability-constrained timesteps. For example, in a viscous CFL dominated problem on a 4-level SAMR mesh, the finest mesh takes 2^8 steps for every one on the coarsest. Generally, this constitutes what is known as *subcycling* but may also be approached more rigorously using partitioned multi-rate [4] methods. A- and L-stable [5] implicit schemes may ignore stability restrictions and select step-sizes based exclusively on accuracy and iteration concerns. Such methods usually result in linear systems which are solved by iterative techniques adapted to handle patches which do not span the entire domain.

The simplicity and elegance of SAMR hide a number of issues which arise in practical situations. If a constant refinement factor is used (e.g. 2) , the number of grid points rises geometrically as one refines. When used with time-refined explicit integration (where the finest grids undergo the largest number of timesteps), SAMR-based simulations spend almost all their time in the finest levels. Further, in parallel simulations, the requirement of keeping parents and children together (on the same processor) result in poor domain-partitioning by run-of-the-mill partitioners. Thus in practice time-refined simulations usually have shallow hierarchies (e. g. 4 deep) with a fairly fine "coarse" mesh. The timestep size, determined by the stability constraints of the fine "coarse" mesh is needlessly small.

Specific Objectives

To date, second-order spatial discretizations have been the standard for SAMR simulations on vertex- and cell-centered grids. Accuracy has been achieved mainly by increasing mesh density, which has exacerbated the problem of a rather fine starting mesh (henceforth called Level0 mesh). An alternative approach, which we explore in this paper, is to use higher-order discretizations of both derivative and interpolant operators on the patches. To illustrate the potential efficiency gains of higher-order spatial discretizations, one

[1] Corresponding author: jairay@ca.sandia.gov

Proceedings of the Third Joint Meeting of the U.S. Sections of The Combustion Institute

ε	2E	4E	6E	8E	10E
$10^{-2.0}$	25.65	8.344	5.712	4.696	4.156
$10^{-2.5}$	45.53	11.22	7.012	5.516	4.753
$10^{-3.0}$	81.07	15.03	8.572	6.444	5.412
$10^{-3.5}$	144.1	20.20	10.44	7.507	6.136
$10^{-4.0}$	256.5	26.85	12.69	8.727	6.943
$10^{-4.5}$	455.3	35.70	15.44	10.12	7.834
$10^{-5.0}$	816.0	47.60	18.70	11.72	8.837

Table 1: Number of grid points per wavelength required for a chosen error tolerance ε.

may consider first-derivative operators. The qualitative conclusions for derivative operators may then be applied to interpolant operators provided Runge phenomena are not experienced. Jameson[6] has attempted to quantify the relative efficiencies of different orders-of-accuracy. Table 1 lists the number of grid points needed by explicit, centered derivative operators of various order to resolve a particular wavenumber mode to a given error tolerance. Even a modest increase in order (from 2^{nd} to 4^{th}) drops the resolution requirements in 1D by a factor of 3 for a 10^{-2} error tolerance; for tighter tolerances a factor of 20 is achieved; and at higher orders, one achieves almost 2 orders of magnitude. Given that these savings will be squared and cubed as one proceeds to 2D and 3D, respectively, the advantages of a higher-order scheme become evident.

The use of higher-order stencils is not without its 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. However, the sparser grid (small data size) of a higher-order formulation will benefit tremendously from the higher cache-hit rate and it is unclear whether the "time-to-solution" of a higher-order scheme will be vastly different. In an explicit time integration procedure, the larger eigenvalues of the higher-order matrix derivative operator may reduce the maximum step-size by a third or so. Further, derivative schemes above fourth-order may 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 study we will first empirically attempt to achieve higher-order spatial convergence on a multiple-level SAMR mesh to identify correct discretization-interpolation sets. This will be followed by an adaptation of a Runge-Kutta-Chebyshev scheme [7] for time-refined integration on SAMR meshes. The paper ends with a demonstration in a reaction-diffusion problem with a stiff chemical term. We will employ 4^{th}-order discretizations in our exposition but 6^{th}- and 8^{th}-order results will be presented too.

Higher-Order 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}-, 6^{th}- and 8^{th}-order discretizations.

Derivative Operators

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

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

Several examples of varying accuracies are listed in Table 2 At the domain boundaries, the derivatives can be calculated either by creating a halo of grid points around the domain and using the same stencils or by closing, possibly to lower order, using skewed stencils. Table 3 lists several possible stencils. The approach for higher orders is identical, but it omitted for brevity. We will use the second approach because of its better stability characteristics.

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 2: Stencil coefficients for centered (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.

Interpolation Operators

SAMR simulations periodically *prolong* data from coarse meshes to a halo of points around their finer children so that centered stencils may be used throughout. This process is skipped on a patch boundary if it abuts a domain boundary. More accurate data are also *restricted* from a finer patch to its coarser parent. In two-dimensions, interpolation blocks take data off of a four-squared block of points (for fourth

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 3: Stencil coefficients of skewed lower-order, first-derivative operators near boundary points. These can be used with fourth-order derivative discretizations.

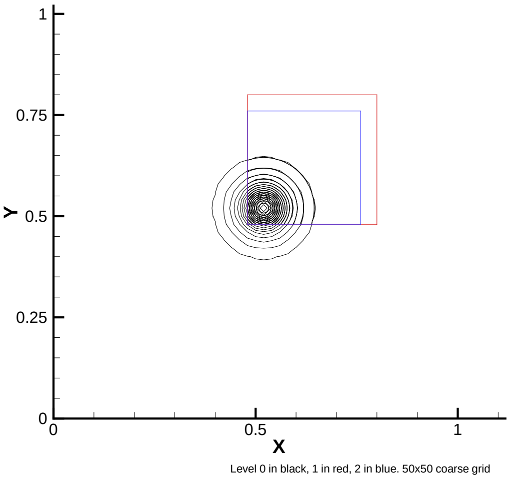

Figure 1: The initial condition and the subdomain undergoing refinement. The red and blue grid are Level1 and Level2 meshes respectively. They were purposely chosen in a skewed manner w.r.t. the field so that a well defined gradient would be present at the subdomain boundary, posing a credible test for the prolongation operators

order interpolations) and interpolate a value onto the geometric center a cell. Within the interior of the domain, the interpolated point is the geometric center of the block but near physical boundaries, the block will skew relative to the interpolation point. Whereas second-order coarse-to-fine interpolation requires no boundary closure, fourth-order does. At second-order, the stencil coefficient for the points $(i \pm 1/2, j \pm 1/2) = 1/4$. At fourth order, there are 3 unique stencil coefficients; $(i \pm 1/2, j \pm 1/2) = 8/256$, $(i \pm 1/2, j \pm 3/2) = (i \pm 3/2, j \pm 1/2) = -9/256$, and $(i \pm 3/2, j \pm 3/2) = 1/256$. At higher order, the number of unique stencils increases. The two required boundary blocks for fourth-order and details of higher-order interpolants are too long to present here.

Filter Operators

Filtering [8] 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 [9]. Filters are an important element in the overall numerical method.

Results and Discussion

Validation of Discretizations

We explore the feasibility of using higher-order discretization in a SAMR setting by employing them to solve the heat equation. We solve Eqn. 2 on a unit square with mesh-nodal data. ϕ_0, the initial condition is a Gaussian with a characteristic size $\delta = 0.05$. A $lim_{r \to \infty} \nabla \phi = 0$ boundary condition is imposed at the edges of the unit square. A second-order, explicit Runge-Kutta method is used for time-advancement.

$$\phi_t = \nu \nabla^2 \phi, \phi_0 = exp(-r^2/\delta^2). \tag{2}$$

The analytical solution in 2D

$$\phi(r,t) = \frac{\delta^2}{\delta^2 + 4\nu t} exp(-\frac{r^2}{\delta^2 + 4\nu t}),$$

is used to evaluate the error in the numerical solution. A 50×50 Level0 mesh was used and was refined twice by a factor of 2 in an arbitrarily defined subdomain. Fig. 1 shows the initial condition and the two levels of refinement. ν was set to 1.

We integrate this equation in 20 (Level0 mesh) timesteps to $t = 10^{-5}$ using $4^{th}-, 6^{th}-$ and $8^{th}-$ order derivative stencils and interpolants. The error is evaluated only in the refined section of the domain and is identical on all levels (since the solution is restricted from the finest to the coarsest levels). In Fig. 2 we plot the error on the finest level for tests with zero, one and two levels of refinement and compare them with the ideal case of $4^{th}-, 6^{th}-$ and $8^{th}-$ order convergence. We see expected convergence characteristics, except for the highest order methods at their lowest errors. In each case, discretizations and interpolations of identical orders are used in pairs.

Extended Stability Explicit Methods

As mentioned above, shallow grid hierarchies benefit from explicit time-integrators capable of large timesteps. Given

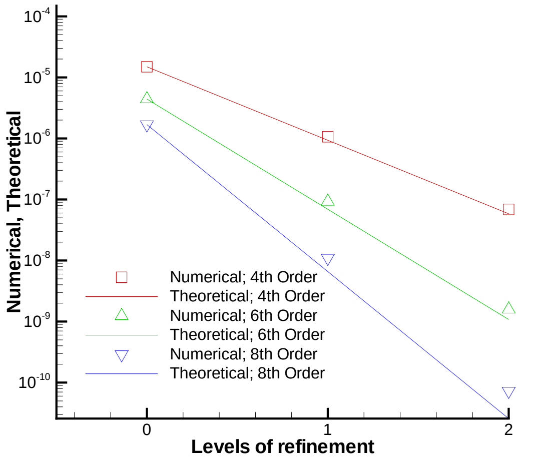

50x50 coarse mesh, t = 1e-5, 20 timesteps.

Figure 2: Convergence results using $4^{th}-$, $6^{th}-$ and $8^{th}-$ order discretizations and interpolations. The solid lines are ideal results, the symbols are experimental results. Red, green and blue denote $4^{th}-$, $6^{th}-$ and $8^{th}-$ order methods. We see that at very low errors (10^{-10}), the convergence for the $8^{th}-$ order discretization degrades to about 7.2.

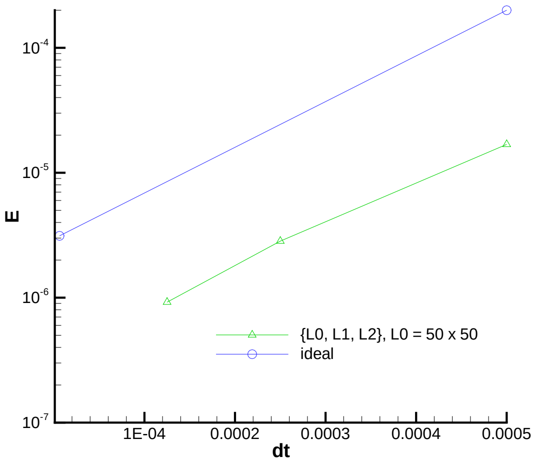

For {L0, L1, L2} run, errors are for the L2 mesh. Dt corresponds to L0's dt

Figure 3: Convergence of the temporal error is the solution of the heat equation. A second-order spatial discretization was used on a mesh with 2 levels of refinement. Errors show 2^{nd} order convergence in time.

an IVP $\phi_t = F(\phi)$, one can design Runge-Kutta schemes with extended stability regions along the negative real axis, i. e. they are suitable for problems where the eigenvalues of $\partial F/\partial\phi$ are near the negative real axis. We use one such second-order Runge-Kutta-Chebyshev scheme [7] with a stability region which grows as $\approx 0.6s^2$ along the negative real axis, s being the number of stages in the scheme. Local error analysis [7] reveals that the error term is $O(\tau^3)$, but also a function of s. Convergence results are in [10]. Briefly, each timestep can be written as :

$$\begin{aligned} Y_0 &= \phi^n \qquad (3)\\ Y_1 &= Y_0 + \tilde{\mu_1}\tau F_0 \\ Y_j &= (1-\mu_j-\nu_j)Y_0 + \mu_j Y_{j-1}\nu_j Y_{j-1} \\ &\quad +\tilde{\mu_j}\tau F_{j-1} + \tilde{\gamma_j}\tau F_0, j = 2\ldots s \end{aligned}$$

$$\phi^{n+1} = Y_s$$

where the time-advancement of ϕ^n to ϕ^{n+1} is over a timestep τ. $\tilde{\mu_j}$, $\tilde{\gamma_j}$, μ_j and ν_j are analytical functions of s. The stage timestep size varies as $\Delta t_j/\tau \approx (j^2-1)/(s^2-1)$ where j is the stage number.

The extension of a traditional explicit Runge-Kutta method (second order) for time-refined stepping on SAMR grids was outlined in [2]. The extension of RKC is conceptually similar. Each RKC stage is first-order, requiring only linear interpolation (from a parent mesh's data) at the child patch edges (where it becomes impossible to evaluate a higher-order spatial stencil due to lack of grid points on one side). Implementation details involve significant booking and temporal interpolations from coarse to fine meshes based on the stage number of the integration process on the fine mesh.

In Fig. 3 we use RKC (modified for time-refined time-stepping on SAMR meshes) for the heat equation (Eqn. 2). A second order discretization was chosen. Spatial refinement was not limited to the subdomain shown above, rather it included the entire Gaussian. Two levels of refinement were allowed. The timestep size dt was larger than the one used in the spatial convergence test but much smaller than the stable limit of an 8-stage RKC. We see that the error convergence is 2^{nd}- order in time.

Reaction-Diffusion Problem

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. Mixture-averaged transport will be used. We use SAMR so that we can resolve detailed flame structure for laboratory-sized unsteady flames. Higher-order discretizations and RKC (both validated above) will be used in this setup. The code is being developed as CCA components [11] which allows a modular architecture.

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

$$\frac{\partial\Phi}{\partial t} = K\nabla\cdot(F\nabla\Phi) + R, \qquad (4)$$

$$\Phi = \left\{\begin{array}{c} T \\ Y_1 \\ \vdots \\ Y_{N-1}\end{array}\right\}, K = \left\{\begin{array}{c}\frac{1}{\rho c_p}\\ \frac{1}{\rho}\\ \vdots \\ \frac{1}{\rho}\end{array}\right\}, F = \left\{ \begin{array}{c}\lambda \\ \rho D_1 \\ \vdots \\ \rho D_M\end{array}\right\},$$

$$R = \left\{\begin{array}{l} \frac{1}{\rho c_p}[-\sum_{i=1}^N h_i\omega_i] \\ \frac{1}{\rho}\omega_1 \\ \vdots \\ \frac{1}{\rho}\omega_M \end{array}\right\}, \sum_{i=1}^N Y_i = 1$$

where R is the reactive production of heat and species while $K\nabla \cdot (F\nabla\Phi)$ is the diffusive transport (by Fick's Law) of the heat and the species. $M = N - 1$, where N is the number of chemical species. λ is the thermal conductivity and D_i are the diffusion coefficients. This is a very coarse model of a flame. It includes chemistry and the diffusion of heat and species; it neglects momentum and the transport of heat by diffusion of species. Further, pressure is assumed to be constant in time and space (*i.e.* burning in an open domain). Radiation is neglected too. The species are assumed to diffuse independently into the mixture at a mesh point *i.e.* the diffusion coefficient D_i of the i^{th} species is mixture averaged. We employ operator-splitting [12] to advance diffusion using RKC while using a BDF5 scheme [13] to integrate chemistry. The general procedure is similar to that in [14], except for the projection solution, which is omitted here.

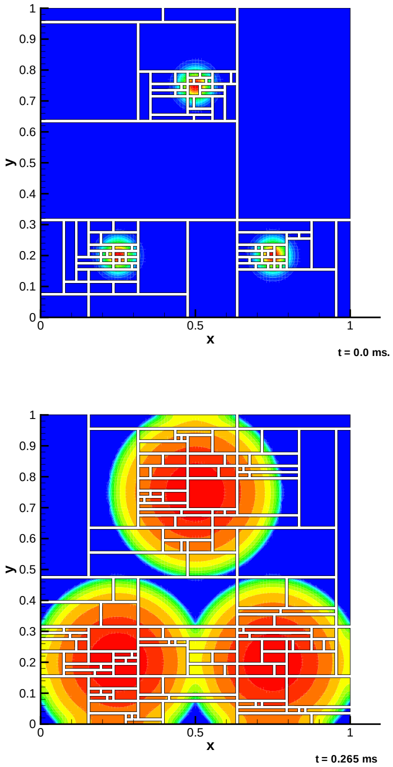

Figure 4: Temperature field at t = 0 and 0.265 ms. The black lines show domain decomposition across 24 processors.

We simulate the ignition of a stoichiometric $\mathrm{H_2} - \mathrm{Air}$ mixture with three hot spots. A 9-species 19 reversible reaction mechanism [15] is used for the chemistry. Second-order derivative stencils with bicubic interpolations were used. The domain was $1\mathrm{cm} \times 1\mathrm{cm}$ with a 100×100 Level0 mesh. 3 levels of refinement were allowed, with a refinement factor of 2. RKC timestep size was limited to 500 nanoseconds to reduce splitting errors. This is still more than a order of magnitude than what could be achieved by a traditional RK2 method, adapted for time-refinement.

In Fig. 4 we see the got spots ignite the mixture. The ignition fronts are seen to move through the mixture and annihilate sections of each other as they interact. In Fig. 5 we see $\mathrm{H_2O_2}$ profiles plotted on the finest mesh as well as the distribution of patches over different levels of refinement. In Fig. 5 we see the resolution capability of the SAMR resolving the fine $\mathrm{H_2O_2}$ profile by about 10 points.

Conclusions

We have presented results of higher-order discretized solutions of the heat equation. For moderate order (fourth and sixth), the results are as expected on a 3-level SAMR mesh. No numerical instability was noticed due to the lower-order closure of higher-order derivative schemes at boundaries. The sixth order scheme (closed to fourth order at domain boundaries) could drop an order in the presence of gradients near boundaries. However, the test was stopped before this could happen.

The extension of RKC to SAMR did not reveal any hidden numerical issues. Its similarity with second-order Runge-Kutta allowed the treatment in [2] to be applied successfully in a straightforward manner.

The reaction-diffusion problem was solved successfully. It is safe to say that resolution of structures was achieved. A stable global timestep of $\approx$ 1000 ns (achievable by a 8-stage RKC scheme) was never tested because of the splitting errors at such a timestep - an artificial limit of 500 ns (at which splitting errors were empirically found to be acceptable) was used.

Future work will concentrate on the mesh-density / order trade-offs as well as a more rigorous convergence analysis in the presence of thin fronts. Incorporation of convective terms into the equations (which introduce a small but purely imaginary component to the eigenvalues of the Jacobian) will be used to test the stability of the RKC.

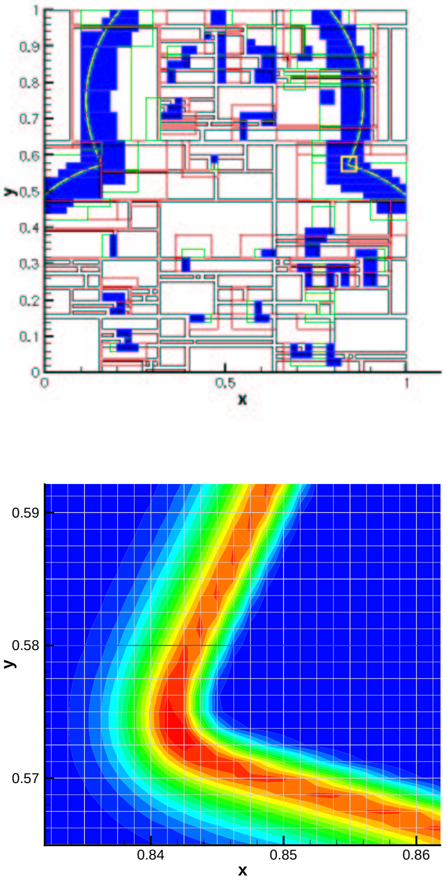

Figure 5: Above : H_2O_2 profiles plotted on the finest mesh level. Patches on Levels 0, 1 and 2 are shown in black, red and green respectively. Below: A zoom into the yellow box in the picture above. We see the H_2O_2 profile being resolved by about 10 grid points.

Acknowledgments

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

References

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

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

[3] Ann S. Almgren, John B. Bell, Phillip Collela, Louis H. Howell, and Michael L. Welcome. A conservative adaptive projection method for the variable density incompressible Navier-Stokes equations. *J. Comp. Phys.*, 142:1–46, 1998.

[4] M. Gunther, A. Kværno and P. Rentrop. Multirate partitioned Runge-Kutta methods. *BIT*, 41:504–514, 2001.

[5] E. Hairer and G. Wanner. *Solving Ordinary Differential Equations I; Stiff and Differential Algebraic Equations*. Springer-Verlag, Berlin, 1996.

[6] L. Jameson. AMR vs. high-order schemes, wavelets as a guide. Technical Report UCRL-ID-141537, Lawrence Livermore National Laboratory, Livermore, California, 2000.

[7] L. F. Shampine, B. P. Sommeijer and J. G. Verwer. RKC: an explicit solver for parabolic PDEs. *J. Comp. Appl. Math.*, 88:315–326, 1998.

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

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

[10] J. G. Verwer, W. H. Hunsdorder and B. P. Sommeijer. Convergence properties of the Runge-Kutta-Chebyshev method. *Num. Math.*, 57:157–178, 1990.

[11] 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.

[12] G. Strang. On the construction and comparison of difference schemes. *SIAM J. Numer. Anal.*, 5(3):506–517, 1968.

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

[14] 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.

[15] 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.