

Reducing reflections from mesh refinement interfaces in numerical relativity

John G. Baker and James R. van Meter

Laboratory for Gravitational Astrophysics, NASA Goddard Space Flight Center, Greenbelt, Maryland 20771, USA

(Received 19 May 2005; published 17 November 2005)

Full interpretation of data from gravitational wave observations will require accurate numerical simulations of source systems, particularly binary black hole mergers. A leading approach to improving accuracy in numerical relativity simulations of black hole systems is through fixed or adaptive mesh refinement techniques. We describe a manifestation of numerical interface truncation error which appears as significant, artificial reflections from refinement boundaries in a broad class of mesh refinement implementations, potentially compromising the effectiveness of mesh refinement techniques for some numerical relativity applications (if left untreated). We elucidate this numerical effect by presenting a model problem which exhibits the phenomenon, but which is simple enough that its numerical error can be understood analytically. Our analysis shows that the effect is caused by variations in finite differencing error generated across low and high resolution regions, and that associated difficulties in demonstrating convergence at modest resolutions are caused by the presence of dramatic speed differences among propagation modes typical of $3 + 1$ relativity. Last, to further verify our understanding of this problem, we present a class of finite differencing stencils of the same order of accuracy as the desired order of convergence, termed mesh-adapted differencing (MAD), which eliminate this pathology in both our model problem and in numerical relativity examples.

DOI: [10.1103/PhysRevD.72.104010](https://doi.org/10.1103/PhysRevD.72.104010)

PACS numbers: 04.25.Dm, 02.70.Bf, 04.30.Db, 04.70.-s

I. INTRODUCTION

Recent years have seen a dramatic rise in opportunities for observing strong-field gravitational dynamics. New observations of dense black-hole-like objects, at stellar, intermediate, and supermassive scales are increasingly frequent. Anticipated gravitational wave observations by ground-based and space-based detectors are expected to capture information about these objects at moments of the strongest gravitational interactions [1,2]. Interpretation of data from any such observations will depend on theoretical modeling of the strong-field interactions of dense black-hole-like objects in the process of generating gravitational radiation. General relativity is the standard model for describing gravitational interactions and wave generation. However, the predictions of general relativity for such cases are not yet fully understood, and will depend on 3D numerical relativity computer simulations [3].

While numerical relativity has progressed markedly in recent years [4], significant improvements in the fidelity of models for events such as binary black hole coalescence will be essential for the full interpretation of upcoming observations. There are many facets to the problem of improving such simulations, including optimally formulating Einstein's equations [5–8], properly handling boundaries, handling black hole singularities, controlling constraint violations, and making judicious gauge choices [9]. There are also basic numerical issues concerning how to, with finite resources, perform such high-fidelity 3D simulations with strong short-wavelength gravitational features near the sources, and weak but critical long-wavelength gravitational (radiation) features emerging in a large domain. Approaches to this latter class of issues include developing higher-order accurate finite differenc-

ing methods [10], spectral methods [11], mesh refinement techniques [12–14] and other forms of numerical patching techniques [15,16].

We focus here on resolving a limitation which has arisen in our work on numerical relativity simulations of binary black hole systems with mesh refinement, but which may have analogues in other numerical patching treatments as well. Mesh refinement techniques divide the computational domain into regions with separate computational grids which can be of higher resolution in some regions than others. Such approaches involve mesh-structure interfaces across which the details of the finite differencing treatment suddenly change. Inevitably, these interfaces contribute to computational error, manifesting such effects as “reflections” off the interfaces. Considerable attention is given to implementing “clean” interfaces, which generate small error, compared to error generated in the bulk regions. At minimum, this requires interface-induced error to *demonstrably* converge at least as rapidly as the bulk error. For practical control of numerical error, simulations at some finite resolution must manifest demonstrable convergence, meaning that, as the resolution is increased from that point, the size of the error will decrease at some expected rate.¹ For some classes of black hole evolutions, with nonvanishing shift advection terms, we have typically observed large error propagating from mesh interfaces which is not demonstrably convergent at achievable resolutions.

¹We note here distinction between the abstract property of *convergence* which implies that beyond some unknown arbitrarily high resolution the error will decrease predictably, and our notion here of *demonstrable convergence* which must be realized in the results at the specified finite resolutions.

Understanding and resolving this problem forms the focus of this paper. For specificity we will focus on second-order calculations, although our analysis will directly generalize to higher-order treatments. While our analysis will show that demonstrable second-order convergence would be achievable through the use of higher-order spatial finite differencing, such higher-order differencing can be problematical in strong field regions. We demonstrate an alternative solution, mesh-adapted differencing (MAD), which provides demonstrable convergence at modest resolutions without formally increasing the convergence order of the spatial differencing.

II. INTERFACE PERFORMANCE

The pertinent features of our numerical scheme are as follows. We are solving the 3 + 1 BSSN formulation of Einstein's equations [5,12,17,18]. Our gauge condition is numerically determined, typically with some variation of the 1 + log slicing and hyperbolic Gamma-driver shift evolution equations [9]. We integrate in time with the iterative Crank-Nicholson method [19]. All spatial derivatives are computed by second-order accurate, centered differencing, save for advection derivatives, for which we use second-order accurate upwinded differencing.

We are particularly interested in simulating gravitational radiation generated in black hole collisions. To resolve the black hole sources adequately, while pushing the computational grid boundary sufficiently far away, we use fixed mesh refinement, as implemented by a software package for this purpose called PARAMESH [20]. With this implementation, the resolutions of two adjacent refinement regions always differ by a factor of 2; i.e. the grid spacing of the coarser region is double what it is in the finer region. "Ghostzones" or "guard cells," typically two layers, are required to provide buffering between refinement levels. These guard cells are filled in by interpolation.

Whether the simulation performs adequately in the presence of refinement interfaces is a question of particular concern to us. Inevitably, refinement boundaries are sources of numerical errors, although these reflections are often satisfactorily convergent and negligible. An exception has plagued us in the case of a nonnegligible shift, β^i , at refinement boundaries. In this case we observe a reflection pulse that propagates at a velocity of $-\beta^i$.

An example can be seen in the case of single Schwarzschild black hole (in isotropic Schwarzschild coordinates) centered in a nested-box arrangement of refinement regions. The Hamiltonian constraint provides a measure of the error in the runs and is plotted in Fig. 1. These "reflected" error waves, or "bumps," notably propagating toward the black hole (at $x = 0$) from each interface, are strongest when the refinement interface is nearest to the black hole (where β^i is largest). We have also noticed that the reflections seem to originate coincidentally with the passing of an initial gauge pulse through the




FIG. 1 (color online). Convergence plot for the Hamiltonian constraint error H at time $t = 4M$. There is a single puncture black hole centered at the origin, and refinement boundaries at $|x_i| = 1M, 2M, \text{ and } 4M$. The finest grid spacing h_f for each simulation is indicated in the figure, and the highest resolution is multiplied by a factor of 4 while the lowest resolution is divided by a factor of 4. For second-order errors the curves should superpose in the limit $h \rightarrow 0$.

interface. Such gauge pulses are typical in black hole simulations with "1 + log" type lapse conditions, and propagate in this case at $\sqrt{2}$ times the speed of light asymptotically. In Fig. 1 we test for demonstrable convergence of the bumps by comparing the error in the Hamiltonian constraint at three resolutions, a low resolution run of $h_f = M/16$, a moderate resolution run of $h_f = M/32$, and a run with uniformly doubled resolution, $h_f = M/64$, where h_f refers to the resolution of the finest grid. The curves have been rescaled so that they should superpose if the errors are second-order convergent and the grid spacing is sufficiently small. Disturbingly, the figure fails to demonstrate that these errors converge at reasonable grid resolutions, $M/32$ and $M/64$. Whether these reflection errors would demonstrably converge at sufficiently high resolution is difficult to determine in our black hole applications of limited achievable grid size.

In any case, the practical effect is that it is a challenge to effectively control these errors which may dominate the error around refinement boundaries in the strong field region. Interestingly, however, we have not observed any adverse effect, such as poor convergence, imprinted by these bumps on such phenomena as gravitational radiation measured far from a binary black hole system [13]. However, other physical quantities of interest, such as the event horizon, seem likely to be more sensitive to these near-field errors.

III. LINEARIZED BSSN

To understand the source of the β -speed error exhibited in the last section, we begin by considering a linearized

BSSN system with 1 + log slicing and hyperbolic Gamma-driver shift. The system of equations is

$$\dot{a} = -2K \quad (1)$$

$$\dot{B}^i = \dot{\Gamma}^i \quad (2)$$

$$\dot{\beta}_1^i = \frac{3}{4}B \quad (3)$$

$$\dot{\phi} = -\frac{1}{6}(K - \partial_i \beta_1^i) + \beta_0^k \partial_k \phi \quad (4)$$

$$\dot{K} = -\partial_i \partial_i a + \beta_0^k \partial_k K \quad (5)$$

$$\dot{h}_{ij} = -2\tilde{A}_{ij} + \partial_i \beta_1^j + \partial_j \beta_1^i - \frac{2}{3}\delta_{ij} \partial_k \beta_1^k + \beta_0^k \partial_k h_{ij} \quad (6)$$

$$\begin{aligned} \dot{\tilde{A}}_{ij} = & [-\partial_i \partial_j a - \frac{1}{2} \partial_k \partial_k h_{ij} + \frac{1}{2} \partial_i \tilde{\Gamma}^j + \frac{1}{2} \partial_j \tilde{\Gamma}^i - 2\partial_i \partial_j \phi]^{TF} \\ & + \beta_0^k \partial_k \tilde{A}_{ij} \end{aligned} \quad (7)$$

$$\dot{\tilde{\Gamma}}^i = -\frac{4}{3} \partial_i K + \partial_k \partial_k \beta_1^i + \frac{1}{3} \partial_i \partial_j \beta_1^j + \beta_0^k \partial_k \tilde{\Gamma}^i \quad (8)$$

where $a \equiv \alpha - 1$, $\beta_1^i \equiv \beta^i - \beta_0^i$ (with β_0^i assumed spatially uniform for simplicity), and $h_{ij} \equiv \tilde{\gamma}_{ij} - \delta_{ij}$.

For a problem that varies only in one dimension, with no initial transverse components, if we assume plane-wave solutions (generalizable by Fourier analysis), then the above system of equations can be written in the form

$$\partial_t |u\rangle = ikM|u\rangle \quad (9)$$

where

$$|u\rangle = \begin{pmatrix} \hat{a} \\ \hat{B} \\ \hat{\beta}_1 \\ \hat{\phi} \\ \hat{K} \\ \hat{h} \\ \hat{A} \\ \hat{\Gamma} \end{pmatrix} e^{i(kx - \omega t)} \quad (10)$$

with \hat{a} , \hat{B} , $\hat{\beta}_1$, $\hat{\phi}$, \hat{K} , \hat{h} , \hat{A} , and $\hat{\Gamma}$ the amplitudes of a , B^x , β_1^x , ϕ , K , h_{xx} , \tilde{A}_{xx} , and $\tilde{\Gamma}^x$ respectively, and

$$M = \begin{pmatrix} 0 & 0 & 0 & 0 & \frac{2}{ik} & 0 & 0 & 0 \\ 0 & 0 & \frac{4ik}{3} & 0 & \frac{4}{3} & 0 & 0 & -\beta_0 \\ 0 & \frac{3}{4ik} & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & -\frac{1}{6} & -\beta_0 & -\frac{1}{6ik} & 0 & 0 & 0 \\ -ik & 0 & 0 & 0 & -\beta_0 & 0 & 0 & 0 \\ 0 & 0 & -\frac{4}{3} & 0 & 0 & -\beta_0 & -\frac{2}{ik} & 0 \\ -\frac{2ik}{3} & 0 & 0 & -\frac{4ik}{3} & 0 & -\frac{ik}{2} & -\beta_0 & -\frac{2}{3} \\ 0 & 0 & \frac{4ik}{3} & 0 & \frac{4}{3} & 0 & 0 & -\beta_0 \end{pmatrix}. \quad (11)$$

The eigenvalues are 0, $-\beta_0$, $-\beta_0 + 1$, $-\beta_0 - 1$, $-\frac{1}{2}\beta_0 + \frac{1}{2}\sqrt{\beta_0^2 + 8}$, $-\frac{1}{2}\beta_0 - \frac{1}{2}\sqrt{\beta_0^2 + 8}$, $-\frac{1}{2}\beta_0 + \frac{1}{2}\times$

$\sqrt{\beta_0^2 + 4}$, and $-\frac{1}{2}\beta_0 - \frac{1}{2}\sqrt{\beta_0^2 + 4}$. If $|\lambda\rangle$ is the eigenvector associated with eigenvalue λ , and $\langle\lambda|$ is defined such that $\sum_\lambda |\lambda\rangle\langle\lambda| = I$, then

$$M = \sum_\lambda \lambda |\lambda\rangle\langle\lambda|. \quad (12)$$

By substituting Eq. (12) into Eq. (9) the system of evolution equations can be decomposed into a series of advection terms, each associated with a characteristic velocity equal to one of the eigenvalues. Note that, assuming $\beta_0 \ll 1$, the β -speed mode is much slower than any of the other nonzero-speed modes. As we are particularly concerned with this mode, it is instructive to write the evolution equation thus:

$$\begin{aligned} \partial_t \begin{pmatrix} a \\ B \\ \beta_1 \\ \phi \\ K \\ h \\ A \\ \Gamma \end{pmatrix} = & -\beta_0 \partial_x \begin{pmatrix} 0 \\ 0 \\ 0 \\ -\frac{3}{8} \\ 0 \\ 1 \\ 0 \\ 0 \end{pmatrix} \left(-\frac{8}{3} \phi - \frac{1}{3ik} \Gamma \right) \\ & + \sum_{\lambda \neq -\beta_0} \lambda \partial_x |\lambda\rangle\langle\lambda| u. \end{aligned} \quad (13)$$

This equation indicates that disturbances originating in ϕ and $\tilde{\Gamma}^i$ can propagate in ϕ and h_{ij} at β speed and, further, that this is the only means of β -speed propagation allowed by this system. Thus ϕ is uniquely significant in both generating and propagating β -speed modes. Note that Eq. (13) resembles the $\dot{\phi}$ evolution equation in the BSSN system.

IV. A MODEL PROBLEM

Motivated by the discussion of the last section, a simple model for the generation and propagation of the β -speed modes is a one-dimensional advection problem with an additional driving term,

$$\dot{\phi}(x, t) = \beta \partial_x \phi(x, t) + f(x, t). \quad (14)$$

In our numerical simulations, it appears that the reflected bumps may be triggered by a rapidly propagating gauge pulse which propagates outward early in the simulations. For our model problem, which we will term ‘‘Bumpy’’ because of its most salient feature, we will drive the advection equation with a pulse propagating at speed ν , significantly faster than the advection speed β . We thus let $f(x, t) = f(x - \nu t)$ where ν is larger than β and both are assumed positive. This model equation is simple enough to be solved directly,

$$\begin{aligned}\varphi(x, t) &= \int_0^t f(x + (t - t')\beta, t') dt' \\ &= \frac{-1}{v + \beta} (F(x - vt) - F(x + \beta t)),\end{aligned}\quad (15)$$

where $\partial_x F(x) = f(x)$. For definiteness we can take the driving term to be a Gaussian pulse, $f(x) = \exp(-x^2)$, implying $F(x) = (\sqrt{\pi}/2)\text{erf}(x)$. For a localized pulse, such as this, the $F(x + \beta t)$ term in the exact solution will be negligible in the region of interest, near $x = x_0$.

To test if this model problem is sufficient, we have numerically evolved Eq. (14) using a 1D finite difference code with a resolution $dx = h$ for $x < x_0$ and $dx = 2h$ for $x > x_0$, realizing a refinement boundary at $x = x_0 = 50$. We evolved over the domain $0 < x < 100$ with periodic boundary conditions, using three-point upwind finite difference stencil and a mesh refinement scheme similar to that in our numerical simulations with PARAMESH. In Fig. 2 we show the errors in our evolution variable φ in the vicinity of the refinement interface for runs with several resolutions, $h = 5/32, 5/64, 5/128$ and $5/256$. In each case there is clearly a reflection error bump propagating to the left, away from the refinement interface, which we find propagates at speed β . As before the errors have been rescaled so that that second-order converging error should superpose in the limit of small grid spacing. Comparing the three lower resolution runs, we do not see good superposition, and the peaks appear to demonstrably converge at roughly first order. The comparative appearance of these errors is similar to that seen for black hole simulations in Fig. 1. However, running at higher resolution, which is readily allowed by this 1D model, we find that the reflected error is indeed second-order convergent, as is suggested by




FIG. 2 (color online). Convergence plot for the error in φ . There is a refinement boundary at $x = 50$. The coarse grid spacing h of each simulation is indicated in the figure, and each curve has been multiplied by a respective factor as appropriate to demonstrate second-order convergence. The convergence is demonstrable only for resolutions of $h \lesssim 5/128$.

comparing the two higher resolution runs in Fig. 2. This convergence is manifest only at relatively high resolution (high relative to the wavelength—and higher than can be easily achieved with respect to a gravitational wavelength in a binary black hole run). Further experimentation has indicated that these results are not strongly affected by variations in the time-integration method or by changing among mesh refinement interfacing schemes which are consistent with the overall second-order finite differencing accuracy.

V. A FINITE DIFFERENCE ANALYSIS

In this section we attempt to understand the numerical behavior of the Bumpy problem Eq. (14) by constructing here an analytic model for the numerical error in our simulations. Noting that the time discretization, and the details of the interpolation scheme used in applying refinement conditions seem not to be directly linked to the problematic error features we see, we model the error with as few assumptions as possible about these details. We treat the numerical errors continuously in time, and we consider the effects of spatial finite differencing in terms of a continuous field $\varphi(x, t, h)$ representing the numerical solution at (fine-grid) resolution $dx = h$. We then expand φ in orders of h ,

$$\varphi(x, t, h) = \varphi_e(x, t) + h^2 \varphi_2(x, t) + \mathcal{O}(h^3). \quad (16)$$

Where $\varphi_e(x, t)$ is understood to be the exact solution Eq. (15). We consider the effect of our finite differencing scheme by replacing the spatial derivative appearing in Eq. (14) with a suitable finite difference operator \tilde{D}_h . Thus,

$$\dot{\varphi}(x, t, h) = \beta \tilde{D}_h \varphi(x, t, h) + f(x, t). \quad (17)$$

As above, we include a refinement jump in the grid at $x = x_0$, represented here by applying a coarser version of the finite difference stencil in the $x > x_0$ part of the spatial domain. This will be consistent with any refinement interface algorithm which applies the same finite difference stencil in both the coarse and fine regions, and applies an interpolative guard-cell filling algorithm at the interfaces which leads to finite differences at the interface which are consistently second-order accurate. On a uniform grid the second-order error term of a finite first derivative is generally proportional to the third derivative of the field,

$$D_h \varphi(x, t, h) = \partial_x \varphi(x, t, h) + e_2 h^2 \partial_x^3 \varphi(x, t, h) + \mathcal{O}(h^3).$$

For the specific upwind differencing operator used in Sec. IV, the stencil of which is,

$$D_h = \frac{1}{h} \left(-\frac{3}{2} + 2E_h - \frac{1}{2}E_{2h} \right), \quad (18)$$

where E is the spatial translation operator defined such that $E_h f(x) = f(x + h)$, the constant error coefficient turns out to be $e_2 = -1/3$. Including the refinement jump, we have

$$\begin{aligned}\tilde{D}_h\varphi &= D_h\varphi + \Theta(x-x_0)(D_{2h}-D_h)\varphi \\ &= \partial_x\varphi + (1+3\Theta(x-x_0))e_2h^2\partial_x^3\varphi + \mathcal{O}(h^3).\end{aligned}\quad (19)$$

Substituting into Eq. (17), and rearranging, yields

$$\begin{aligned}h^2\dot{\varphi}_2(x,t,h) &= [-\dot{\varphi}_e(x,t) + \beta\partial_x\varphi_e(x,t) + f(x,t)] \\ &\quad + h^2\beta[\partial_x\varphi_2(x,t,h) + e_2(1+3\Theta(x-x_0)) \\ &\quad \times \partial_x^3\varphi_e(x,t)] + \mathcal{O}(h^3).\end{aligned}\quad (20)$$

Then noting that the first term vanishes, and taking the limit $h \rightarrow 0$, we derive

$$\begin{aligned}\dot{\varphi}_2(x,t) &= \beta\partial_x\varphi_2(x,t) \\ &\quad + \beta e_2(1+3\Theta(x-x_0))\partial_x^3\varphi_e(x,t),\end{aligned}\quad (21)$$

where we have used the notation $\varphi_2(x,t) = \varphi_2(x,t,0)$. This is our model for the generation and propagation of finite differencing error in the numerical model problem in Sec. IV.

Since Eq. (21) is of the same form as Eq. (14) we can solve it in like fashion. We leave out the negligible $F(x+\beta t)$ term in φ_e , substituting

$$\begin{aligned}f(x,t) &= \beta e_2(1+3\Theta(x-x_0))\partial_x^3\varphi_e(x,t) \\ &\simeq -\frac{\beta e_2}{v+\beta}(1+3\Theta(x-x_0))\partial_x^3F(x-vt)\end{aligned}$$

into the first line of Eq. (15). Then, performing the integral with careful attention to the presence of the step function, we get the solution

$$\begin{aligned}\varphi_2(x,t) &= (s(x-vt) - s(x+\beta t))(1+3\Theta(x-x_0)) \\ &\quad + 3s\left(\frac{-v}{\beta}\left(x-x_0+\beta\left(t-\frac{x_0}{v}\right)\right)\right) \\ &\quad \times (\Theta(x-x_0+\beta t) - \Theta(x-x_0))\end{aligned}\quad (22)$$

where,

$$s(x) = \frac{\beta e_2}{(v+\beta)^2}\partial_x^2F(x) = -\frac{2\beta e_2x}{(v+\beta)^2}\exp(-x^2).$$

As in Sec. IV, the $s(x+\beta t)$ is negligible in the relevant region near the refinement interface. Thus the first term in Eq. (22) is effectively the differencing error associated with the upsweep in φ_e which propagates across the grid at speed v . This part grows 4 times larger in the coarse $x > x_0$ region. The second term is quite interesting, as it propagates in the reverse direction at speed β . It has the same $s(x)$ shape as the forward propagating component, but it is reversed, and contracted by a factor β/v . It is timed to originate at the interface as the first pulse crosses. The two terms combine make the full solution continuous at the interface. Heuristically, one could say that the second term is caused by the discontinuity in the differencing error, generated in order to produce a regular solution, and then, necessarily advecting away as required by the original

model equation, Eq. (14). It is contracted because it propagates at a different speed than the first term, while their time dependences must match at the interface.

Note that although we concretely consider second-order finite differencing here, the finite differencing order is largely irrelevant in this analysis. The calculation can be directly adapted for the leading order in a higher-order finite differencing by changing a few coefficients.

We check that this analysis accurately describes the numerical errors in our Bumpy model evolutions by comparing the predicted error with the numerical error results. As Fig. 3 shows, the prediction of Eq. (22) agrees to high precision with the numerical simulation errors realized in high resolution runs. The figure shows both the rightward propagating wave which travels at velocity v with the driving pulse, and the shorter-wavelength reflected wave. The reason that demonstrable convergence was not readily achieved in our Bumpy simulations is now clear. Because the reflection must propagate at a significantly slower speed than the pulse which generated it, while their frequencies must match, the reflected pulse necessarily has a shortened wavelength. The key ingredients in producing this effect are that we are simulating a system with strongly mismatched propagation speeds, and with the potential for generating mode-mixing reflections off numerical features such as our refinement interfaces. A similar difficulty has recently been noted in the context of reflections from external boundaries [21]. In such situations, with resolution just sufficient to accurately resolve the long-wavelength features in the solution, the shortened wavelength of the numerical reflections makes them unresolvable without a significant increase in the resolution (say by a factor of v/β). We thus have reason to expect that such




FIG. 3 (color online). Numerical error in φ , immediately after a pulse incident on the refinement interface at $x = 50$ has passed through, generating a pulse of transmitted error and a contracted pulse of reflected error. The error is second-order convergent and in agreement with the analytic prediction of the numerical error, as indicated.

an effect is indeed the source of the bumpy reflection in our black hole runs. We will verify this by attempting to eliminate this type of error.

VI. IMPROVING THE DIFFERENCING STENCILS

Errors that are slow to demonstrably converge, of the type exposed in Secs. IV and V, now seem likely to occur under fairly general circumstances. This explains our difficulties in many attempts to eliminate these effects by tweaking the interface conditions in various ways. Our analysis suggests that the only ways to avoid these problems, aside from eliminating the slowly propagating modes by setting $\beta^i = 0$, is to remove the mode-mixing reflections at interfaces. The results of the last section clearly show that the reflected error in the Bumpy system is overwhelmingly due to the discontinuity in the differencing operator. More specifically, the reflection is related to the discontinuity in the second-order truncation error. This observation suggests a solution.

Consider using modified differencing stencils such that the coefficient of the second-order truncation error in the fine-grid region is multiplied by a constant factor q_0 while the coefficient of the second-order truncation error in the coarse grid region is multiplied by a constant factor q_1 . Then the previously constant coefficient e_2 of the last section becomes

$$e_2 \rightarrow [q_0 + (q_1 - q_0)\Theta(x - x_0)]e_2. \quad (23)$$

Repeating the steps of the last section for this new truncation error, the solution for the second-order error in φ becomes

$$\begin{aligned} \varphi_2(x, t) = & (s(x - vt) - s(x + \beta t)) \\ & \times (q_0 + (4q_1 - q_0)\Theta(x - x_0)) \\ & + (4q_1 - q_0)s\left(\frac{-v}{\beta}\left(x - x_0 + \beta\left(t - \frac{x_0}{v}\right)\right)\right) \\ & \times (\Theta(x - x_0 + \beta t) - \Theta(x - x_0)). \end{aligned} \quad (24)$$

Thus for the spatially blueshifted, reflected error, we now obtain,

$$\varphi_{\text{ref}} = (4q_1 - q_0)s\left(\frac{-v}{\beta}\left(x - x_0 + \beta\left(t - \frac{x_0}{v}\right)\right)\right). \quad (25)$$

One possibility would be to eliminate the leading order reflection error simply by choosing $q_0 = q_1 = 0$. This choice corresponds to using higher-order stencils, say third or fourth order accurate, throughout the entire grid. Higher-order differencing methods are clearly valuable in numerical relativity simulations [10]. However, we will focus here on identifying a minimal way to realize effective second-order convergence in a second-order convergent finite differencing scheme. Our result should generalize to arbitrary order.




FIG. 4 (color online). Comparison of the error in φ reflected from the refinement boundary at $x = 50$ in the case of a second-order conventional stencil and a second-order MAD stencil.

Note that with the choice $q_0 = 1$ and $q_1 = \frac{1}{4}$, the reflected error vanishes. More generally, of course, any choice of q_0 and q_1 that makes the truncation error continuous across the refinement boundary will remove the second-order reflection. In particular, the choice

$$q_n = \left(\frac{h_0}{h_n}\right)^2 \quad (26)$$

will work in the n th refinement region of a grid with an arbitrary number of refinement levels, where h_0 is assumed to be the grid spacing in the finest region. We call such a mesh-adapted differencing scheme MAD.

A second-order MAD stencil can be obtained simply by linearly combining a second-order accurate stencil with a higher-order accurate stencil, as follows:

$$D_{h_n} = q_n D_{h_n}^{[2]} + (1 - q_n) D_{h_n}^{[j]} \quad (27)$$

where the superscripted numbers in square brackets represent the order of accuracy of the differencing operator, and $j > 2$. (Of course, this stratagem can be readily generalized for higher-order MAD operators as well.) In the particular case of a second-order upwinded stencil combined with a third-order lopsided stencil, the resulting stencil has the form

$$\begin{aligned} D_{h_n} = & \frac{1}{h_n} \left[\left(-\frac{1}{3} + \frac{1}{3} q_n \right) E_{-h_n} + \left(-\frac{1}{2} - q_n \right) \right. \\ & \left. + (1 + q_n) E_{h_n} + \left(-\frac{1}{6} - \frac{1}{3} q_n \right) E_{2h_n} \right]. \end{aligned} \quad (28)$$

We have implemented this MAD stencil in the numerical simulations of the Bumpy system. Figure 4 shows the result for a run at a moderate resolution. With the MAD stencil the dominant reflected error is nearly eliminated.




FIG. 5 (color online). Comparison of the Hamiltonian constraint H in the case of second-order accurate, conventional differencing stencils and second-order accurate, MAD stencils, as well as a demonstration of convergence in the latter case. There is a single puncture black hole centered at the origin, and refinement boundaries at $|x_i| = 1M, 2M,$ and $4M$.

VII. RESULTS FOR BLACK HOLE EVOLUTIONS

We can demonstrate that our analysis does indeed explain the problematic reflections in our black hole simulations, by verifying that the same remedies applied to fix our Bumpy model problem, i.e. higher-order or MAD stenciling, also fix the problem in our black hole simulations.

As in [10], we have found simulations employing third or higher-order accurate differencing of advection terms are unstable in the vicinity of black hole punctures, if the stencil has any points on the “downwind” side. If the third or higher-order accurate differencing stencil does not have any points on the downwind side, it requires three or more layers of guard cells to accommodate points on the upwind side, which is expensive memory wise. Thus, although higher-order spatial differencing should reduce reflection from refinement boundaries, it brings with it a new set of problems to solve.

Use of a second-order accurate MAD stencil, as in Eq. (28), for advection, avoids the above difficulties. As we generally locate the punctures within the finest grid regions, and the MAD stencil automatically reverts to conventional second-order upwinding in this region, the advection derivative does not take any points on the downwind side in the vicinity of the puncture and we find that stability is maintained.

In our Einstein solver, we have implemented a second-order accurate MAD stencil for advection. For all non-advection derivatives, the MAD stencil is constructed from

a linear combination of second-order centered and fourth-order centered stencils, as in Eq. (27). As a result, reflection from refinement boundaries has been dramatically reduced in the case of a single puncture. Figure 5 compares the Hamiltonian constraint error from a run with conventional stencils with the Hamiltonian constraint error in runs with MAD stencils, demonstrating clear improvement. The plot shows that the reflection error for our $h_f = M/32$ case is much reduced with the MAD stencil. By comparing with the $h_f = M/64$, MAD-stencil run, we also see that the remaining bump now converges (demonstrably) away at second order or better.

For simplicity we show only plots from single black hole simulations here, although we have applied MAD in more interesting binary black hole cases as well. We find that the improvements generalize naturally to these cases. In particular we find that errors in the Weyl scalars near refinement boundaries are reduced by at least an order of magnitude.

VIII. CONCLUSIONS

We have studied a problem which occurs in numerical relativity simulations with nonvanishing shift with mesh refinement. These involve slow advection modes across a mesh interface, and produce large numerical reflections that propagate at the speed β of the advection. We have successfully modeled the problem analytically. Our analysis suggests that the effect is a general consequence of discontinuous jumps in the finite differencing stencil in a problem with propagation modes of widely differing speeds. Our proposed solution, to adjust the finite differencing stencils to make the leading order differencing error continuous across mesh interfaces, is shown to be effective in black hole simulations with mesh refinement, but may have wider application in numerical relativity. In addition to grids of nonuniform refinement, mesh-adapted differencing may also be appropriate for grids with multiple coordinate patches, where discontinuous differencing error can also be expected.

ACKNOWLEDGMENTS

We are happy to thank J. David Brown for helpful discussion. We gratefully acknowledge CPU time grants from the Commodity Cluster Computing Project (NASA-GSFC) and Project Columbia (NASA Advanced Supercomputing Division, NASA-Ames). This work was supported in part by NASA Space Sciences Grant No. ATP02-0043-0056. J. v. M. was also supported in part by the Research Associateship Programs Office of the National Research Council.

- [1] G. Gonzalez *et al.* (LIGO Scientific Collaboration), *Pramana* **63**, 663 (2004).
- [2] K. Danzmann and A. Rudiger, *Classical Quantum Gravity* **20**, S1 (2003).
- [3] B. Schutz, in *The Astrophysics of Gravitational Wave Sources*, edited by J.M. Centrella (AIP, Melville, NY, 2003).
- [4] M. Alcubierre, gr-qc/0412019.
- [5] T.W. Baumgarte and S.L. Shapiro, *Phys. Rev. D* **59**, 024007 (1999).
- [6] H. Friedrich, *Classical Quantum Gravity* **13**, 1451 (1996).
- [7] L.E. Kidder, M.A. Scheel, and S.A. Teukolsky, *Phys. Rev. D* **64**, 064017 (2001).
- [8] C. Bona, T. Ledvinka, C. Palenzuela, and M. Zacek, *Phys. Rev. D* **67**, 104005 (2003).
- [9] M. Alcubierre, B. Brügmann, P. Diener, M. Koppitz, D. Pollney, E. Seidel, and R. Takahashi, *Phys. Rev. D* **67**, 084023 (2003).
- [10] Y. Zlochower, J.G. Baker, M. Campanelli, and C.O. Lousto, *Phys. Rev. D* **72**, 024021 (2005).
- [11] L.E. Kidder, M.A. Scheel, S.A. Teukolsky, E.D. Carlson, and G.B. Cook, *Phys. Rev. D* **62**, 084032 (2000).
- [12] B. Imbiriba, J. Baker, D.-I. Choi, J. Centrella, D.R. Fiske, J.D. Brown, J. van Meter, and K. Olson, *Phys. Rev. D* **70**, 124025 (2004).
- [13] D.R. Fiske, J.G. Baker, J.R. van Meter, D.-I. Choi, and J.M. Centrella, *Phys. Rev. D* **71**, 104036 (2005).
- [14] E. Schnetter, S.H. Hawley, and I. Hawke, *Classical Quantum Gravity* **21**, 1465 (2004).
- [15] G. Calabrese and D. Neilsen, *Phys. Rev. D* **71**, 124027 (2005).
- [16] J. Thornburg, *Classical Quantum Gravity* **21**, 3665 (2004).
- [17] M. Shibata and T. Nakamura, *Phys. Rev. D* **52**, 5428 (1995).
- [18] T. Nakamura, K. Oohara, and Y. Kojima, *Prog. Theor. Phys. Suppl.* **90**, 1 (1987).
- [19] S. Teukolsky, *Phys. Rev. D* **61**, 087501 (2000).
- [20] P. MacNeice, K.M. Olson, C. Mobarri, R. deFainchtein, and C. Packer, *Comput. Phys. Commun.* **126**, 330 (2000).
- [21] L.E. Kidder, L. Lindblom, M.A. Scheel, L.T. Buchman, and H.P. Pfeiffer, *Phys. Rev. D* **71**, 064020 (2005).