

A Unified Algorithm for Adaptive Spacetime Meshing with Nonlocal Cone Constraints

Shripad Thite*

1 Introduction

Motivation Wave propagation is modeled by hyperbolic partial differential equations (PDEs) in both space and time variables, e.g., the wave equation $u_{tt} - \omega^2 u_{xx} = 0$ in 1D space \times time. A *solution* to the PDE is a function $u(x, t)$ that satisfies the equation and the given initial and boundary conditions. The *wave velocity* ω , the velocity at which changes in physical parameters at a point (x, t) propagate to other points in the domain, may be a function of x and t as well as of u and its derivatives. The spacetime discontinuous Galerkin (SDG) finite element method is a numerical method to approximate the exact solution to the PDE. The SDG method approximates the solution within each spacetime element of a *mesh* of the spacetime domain as a linear combination of simple basis functions. The SDG method allows basis functions to be discontinuous across adjacent elements, which means that the mesh can even be nonconforming. A mesh constructed by standard techniques, such as a Delaunay triangulation, cannot be solved efficiently in general. Points in spacetime, and spacetime elements, are partially ordered by *causal dependence*—a point P influences another point Q if and only if changing the solution at P could possibly change the solution at Q . Spacetime elements must be solved in an order that respects this partial order. The efficiency of the solution technique depends on the number of elements that must be solved together because they depend on each other and are therefore *coupled*. In a Delaunay mesh, there is no guarantee on the size of a coupled system; we, on the other hand, construct an efficient mesh where this size is bounded.

Previous work Üngör and Sheffer [6] gave the first advancing front algorithm, TentPitcher, for meshing directly in 2D \times Time given an initial *acute* triangulation of the space domain M . TentPitcher advances the solution over a piecewise linear triangulated *front*, a terrain over M ; the initial front corresponds to $t = 0$ everywhere in space. The front at any step is *causal* which means that points on the front depend only on points in the past. At each step, the algorithm greedily advances in time a local neighborhood of a causal front τ to get a new causal front τ' and a set

of new spacetime tetrahedra; causality of τ and τ' implies that the solution over the spacetime volume between τ and τ' can be computed immediately. Different parts of the front advance at different rates, depending on the local geometry, unlike with uniform time-stepping schemes. The result is a tetrahedral unstructured mesh Ω of $M \times [0, T]$ for any target time T . Erickson *et al.* [2] extended this algorithm to arbitrary spatial domains and higher dimensions, by imposing additional gradient constraints called *progress constraints* on each front.

Abedi *et al.* [1] gave an algorithm to adapt the size and—because of the causality constraint—also the duration of future spacetime elements to a *posteriori* estimates of numerical error. If the elements created at any step are too coarse, they are rejected and the front is refined by repeated bisection; the resulting smaller triangles lead to smaller spacetime elements. Coarsening or derefinement is performed when allowed by the error indicator.

Previous algorithms assumed a fixed upper bound on the wavespeed everywhere in spacetime or that the wavespeed function was Lipschitz. When the PDE is nonlinear, the wavespeed is not constant and also depends on the solution; the wavespeed at a given point in space can change discontinuously with time. Anisotropy of the medium means that waves propagate asymmetrically, with different speeds in different directions. This author [4] gave an algorithm that works even when the wavespeed increases discontinuously and in the presence of anisotropy. However, this algorithm did not adapt the spatial diameter of spacetime elements to numerical error estimates.

New results In this paper, we prove bounds on the worst-case *temporal aspect ratio* of spacetime tetrahedra constructed by TentPitcher; this ratio, together with the spatial aspect ratio, is likely to be correlated with the *quality* of the numerical solution. We also prove bounds on the size of the final mesh relative to a size optimal mesh.

Additionally, we give a unified algorithm that adapts the size of spacetime elements to error estimates while simultaneously adapting their duration to changing wavespeeds. This new algorithm meshes a given spacetime volume with many fewer tetrahedra in general than either of the previous two algorithms [1, 4].

*Department of Computer Science, University of Illinois at Urbana-Champaign; thite@cs.uiuc.edu

2 Temporal aspect ratio and mesh size

The initial front corresponds to time $t = 0$ everywhere in the space domain $M \subset \mathbb{E}^2$. Imagine time increasing upwards. At each step, TentPitcher lifts up a local minimum vertex P of a causal front τ to the vertex P' on a new causal front τ' . For every front triangle PQR incident on P , this creates a new spacetime tetrahedron $P'PQR$ in the volume between τ and τ' , with a causal *inflow* facet PQR on τ and a causal *outflow* facet $P'QR$ on τ' . The volume between τ and τ' is called a *tent* and PP' is the *tentpole*. Causality implies that the solution everywhere in the tent can be computed immediately. The number of new spacetime tetrahedra is equal to the degree of P ; these tetrahedra are coupled because they share non-causal facets. Only elements in a single tent are coupled; tents pitched at different local minima of τ are independent and are solved in parallel.

Suppose the wavespeed everywhere in spacetime is constant, equal to ω ; let σ denote the *slope*, i.e., $\sigma = 1/\omega$. Causality means that the slope of PQR and of $P'QR$ must be less than σ . To guarantee nondegeneracy of tetrahedra, the front at each step must also satisfy so-called *progress constraints*. The causality and progress constraints limit the amount of progress made by the front at each step of the algorithm, i.e., the height of each tentpole; these constraints are functions of the slope σ and the shape of the local triangulation. The progress constraint that each front τ must satisfy depends on the causal slope encountered by the new front τ' in the *next* step which, in this case, is the same slope σ .

The *duration* of a spacetime element is the length of the shortest time interval that contains it. The *height* of a spacetime element is the length of the longest time interval contained in it. Our algorithms maximize the height of each spacetime element subject to causality and progress constraints. The progress guarantee of Erickson *et al.* [2] can be rephrased as follows: the height of each spacetime tetrahedra in the tent pitched at P is at least $\varepsilon\sigma w_p$ where $\varepsilon \in (0, \frac{1}{2}]$ is a fixed parameter and w_p denotes the distance of p from the boundary of the kernel of $\text{link}(p)$ in the spatial projection. Thus, the height of the tentpole at P is bounded from below by a positive function of ε , the slope σ , and the shape of the triangles in $\text{star}(p)$.

Temporal aspect ratio The *temporal aspect ratio* of a spacetime element is the ratio of the height of the element to its duration; this ratio is always in the range $(0, 1]$ with a larger value corresponding to a “better” element. The duration of the tetrahedron $P'PQR$ can be at most $2\sigma w_p$ because both facets PQR and $P'QR$ are causal. Together with the lower bound on the height of the tetrahedron, this implies the following theorem.

Theorem 1 *The temporal aspect ratio of any tetrahedron in Ω is at least $\varepsilon/2$.*

On size optimality TentPitcher constructs groups of coupled tetrahedra inside each tent such that the boundary of the tent is causal. This guarantees convergence of the DG solution [3]. Each tetrahedron constructed by TentPitcher has both a causal inflow facet and a causal outflow facet; additionally, the spatial projection of each tetrahedron $P'PQR$ is the triangle pqr in the original space mesh. Given a triangulation M of the space domain and a target time T , we say that a tetrahedral spacetime mesh of $M \times [0, T]$ is *valid* if (i) each tetrahedron has both a causal inflow facet and a causal outflow facet; and (ii) for every point x in the spatial projection Δ of each tetrahedron, the diameter of Δ does not exceed the diameter of the triangle of M containing x .

Fix an arbitrary point x in space. The *size* of a spacetime mesh of $M \times [0, T]$ is the maximum over $x \in M$ of the number of spacetime elements that intersect the temporal segment $x \times [0, T]$.

Theorem 2 *The size of the mesh constructed by TentPitcher is $\tilde{O}(1/\varepsilon^2)$ times the minimum size of any valid mesh of the spacetime volume $M \times [0, T]$.*

Proof. Let D and ρ denote the diameter and inradius respectively of the triangle pqr of M containing x . By causality, any temporal segment of duration $2\sigma D$ must intersect at least two distinct tetrahedra in a valid mesh; therefore, the number of spacetime tetrahedra in a valid mesh that intersect $x \times [0, T]$ is at least $\lceil T/(2\sigma D) \rceil$.

Consider a minimal sequence of tent pitching steps, called a *superstep*, in which each vertex of Δpqr is lifted at least once. When p is pitched, the amount of progress made by x is proportional to $\text{dist}(x, \overleftarrow{qr})$. Since $\text{dist}(x, \overleftarrow{qr}) + \text{dist}(x, \overleftarrow{rp}) + \text{dist}(x, \overleftarrow{pq}) \geq \rho$, the amount of progress made by x during a superstep is at least $\varepsilon\sigma\gamma\rho$, where $\gamma \in (0, 1]$ denotes the minimum of $w_p/\text{dist}(p, \overleftarrow{qr})$, $w_q/\text{dist}(q, \overleftarrow{rp})$, and $w_r/\text{dist}(r, \overleftarrow{pq})$. Hence, after at most $\lceil T/(\varepsilon\sigma\gamma\rho) \rceil$ supersteps, the point x achieves or exceeds the target time T .

By causality, any two vertices of Δpqr can advance in fewer than $4\sigma D$ consecutive steps without also advancing the third vertex. Therefore, the number of steps in each superstep is at most $\lfloor (4\sigma D)/(\varepsilon\sigma w) \rfloor$ where $w = \min\{w_p, w_q, w_r\}$. It follows that the number of tetrahedra in the spacetime mesh constructed by TentPitcher intersected by $x \times [0, T]$ is at most $\lceil T/(\varepsilon\sigma\gamma\rho) \rceil \cdot \lfloor (4\sigma D)/(\varepsilon\sigma w) \rfloor$.

The ratio of the upper bound to the lower bound on the size is $O\left(\frac{1}{\varepsilon^2} \frac{1}{\gamma} \frac{D^2}{\rho w}\right)$. \square

3 Nonlocal cone constraints

The *cone of influence* of a point P is the set of points that depend on P . This cone has its apex at P and its slope in any spatial direction is the reciprocal of the wavespeed at P in that direction; fast waves correspond to cones with smaller slope. A front τ is *causal* if and only if τ lies below the *cone of influence* of every point P on τ ; each such cone is a *causal cone constraint*. When the medium is *anisotropic*, the cones are asymmetric, e.g., with elliptical cross-sections. When the PDE is nonlinear or when the medium is anisotropic, a distant but fast wave, i.e., a *nonlocal cone*, can overtake a slower wave and hence limit the duration of new elements. Therefore, maximizing the progress of P , and thus the duration of new tetrahedra, requires querying the lower hull of the cones of influence. After the solution is computed on the new front, we obtain a new set of cone constraints. Maintaining the entire arrangement of cones of influence is expensive and unnecessary for our purpose; it suffices to obtain a cone that bounds (tightly) the actual cone of influence at P , i.e., to estimate a lower bound $\tilde{\sigma}(P)$ on the actual slope $\sigma(P)$. We assume the *absence of focusing*, which means that the cone of influence of any point P is contained in the cone of influence of every other point Q such that P depends on Q . Thus, the slope $\sigma(P)$ at P is bounded so that $0 < \sigma_{\min} \leq \sigma(P) \leq \sigma_{\max} < \infty$.

When the wavespeed at a point in space can increase discontinuously, cone constraints must be enforced on the front at every step that depend on the front in the *next* step. We give an algorithm that looks ahead k steps of the algorithm to estimate the slope on future fronts. The lookahead parameter k can be fixed or chosen adaptively. When $k = 0$, we assume that the minimum slope σ_{\min} occurs on the front in the next step, so our estimate of future wavespeed is $\tilde{\sigma} = \sigma_{\min}$. When $k > 0$, we can use the current estimate to compute the next front and the actual slope on this new front to refine our previous estimate. We repeat this process either until subsequent iterations fail to improve the estimate $\tilde{\sigma}$ of future slope or until sufficient progress has already been ensured by the current estimate. (See [4] for the case $k = 1$.) To efficiently query the arrangement of cones, we use a discrete *bounding cone hierarchy* induced by a balanced partition of the triangular front, similar to a bounding volume hierarchy used in collision detection.

We will prove a minimum progress guarantee of $T_{\min} > 0$, a function of the local shape of the triangulation and the global minimum slope σ_{\min} , similar to that proved by Abedi *et al.* [1].

Definition 1 (k -progressive) Let $\triangle ABC$ be a given triangle with A as its lowest vertex. We inductively define $\triangle ABC$ as k -progressive if

- (1) $\triangle ABC$ is causal;
- (2) Let $\triangle A'BC$ denote the triangle after lifting A by T_{\min} to A' . Then, $\triangle ABC$ must satisfy progress constraint $\sigma(A'BC)$ and $\triangle A'BC$ must be $\max\{0, k - 1\}$ -progressive.

Base case $k = 0$: $\triangle ABC$ is 0-progressive iff it satisfies progress constraint σ_{\min} .

4 Adaptive refinement and coarsening

Abedi *et al.* [1] gave an adaptive algorithm by strengthening the progress constraints due to Erickson *et al.* [2]. (This author later [5] corrected an oversight in their proofs, also slightly improving Tent-Pitcher.) The adaptive algorithm refines a triangle on the front using newest-vertex bisection; repeated bisection of a single triangle gives rise to at most 8 predictable homothetic shapes of front triangles. The front is coarsened by undoing previous refinement. The adaptive algorithm enforces a progress constraint on $\triangle PQR$ at every step that anticipates all the different shapes that can be obtained from $\triangle pqr$ by refinement. However, the algorithm of Abedi *et al.* [1] does not anticipate changes in the slope and assumes the minimum slope at every step.

The crucial observation we make here is that a cone of influence that limits the progress of $\triangle PQR$ may not limit the progress of a smaller triangle $\triangle ABC$, a descendant of $\triangle PQR$ by refinement. (The converse is also true.) Specifically, we observe that $\tilde{\sigma}(ABC) \geq \tilde{\sigma}(PQR)$. We obtain a unified algorithm by relaxing the progress constraints of Abedi *et al.* to allow a *child* triangle after by newest-vertex bisection to satisfy a potentially weaker progress constraint than its larger *parent* triangle. A potential drawback is that coarsening two *sibling* triangles coplanar on a front may require both triangles to satisfy a progress constraint *stricter* than their individual progress constraints. This can make coarsening requests harder to satisfy in practice; however, refinement can always be performed when necessary.

A key property we use in deriving the new progress constraint is that the boundary of a cone is a *ruled surface*; if any triangle $\triangle ABC$ intersects a given cone this intersection is a line segment in the plane of $\triangle ABC$. Therefore, if the bisector edge AD intersects a cone of influence then so do at least two of the edges of $\triangle ABC$ (Figure 1). Therefore, if a fast wavespeed in the future intersects AD but not an edge of $\triangle ABC$, then it can do so only if $\triangle ABC$ has been bisected at least once.

Fix ε, φ such that $0 < \varepsilon < \varphi < (1 + \varepsilon)/2 < 1$. For any triangle $\triangle abc$ with newest-vertex a , we define the *diminished width* of $\triangle abc$, denoted by $\tilde{w}(abc)$, as the minimum of $(1 - \varepsilon) \text{dist}(a, \overleftrightarrow{bc})$, $(1 - \varphi) \text{dist}(b, \overleftrightarrow{ac})$, and $(1 - \varphi) \text{dist}(c, \overleftrightarrow{ab})$.




Figure 1: $\triangle abc$ after newest-vertex bisection.

Definition 2 (Progress constraint σ) A triangle $\triangle ABC$ (Figure 1) satisfies progress constraint σ if and only if the applicable constraint from the following list is satisfied:

$$\begin{aligned} \text{If } a \text{ is the lowest vertex: } & |t(b) - t(c)| \leq 4\tilde{w}(fda)\sigma \\ \text{If } b \text{ is the lowest vertex: } & |t(a) - t(c)| \leq 2\tilde{w}(dab)\sigma \\ \text{If } c \text{ is the lowest vertex: } & |t(a) - t(b)| \leq 2\tilde{w}(dca)\sigma \end{aligned}$$

Definition 3 ((k, l) -progressive) We inductively define a triangle $\triangle PQR$ as (k, l) -progressive if it is k -progressive (Definition 1) and any child $\triangle ABC$ obtained by newest-vertex bisection is $(k, \max\{0, l - 1\})$ -progressive.

Base case $l = 0$: $\triangle PQR$ is $(k, 0)$ -progressive if an arbitrary descendant $\triangle ABC$ obtained by newest-vertex bisections satisfies progress constraint $\tilde{\sigma}(PQR)$.

A front is *progressive* if every triangle on the front is (k, l) -progressive for some $k, l \geq 0$. Our unified algorithm greedily maximizes the progress such that each front is (k, l) -progressive for some choice of k and l . The algorithm can be as complicated as desired. Definition 3 stresses the fact that our algorithm can optimize the choice of k and l , likely doing better than the theoretical guarantee; however, a simple choice of $k = l = 1$ may suffice in practice.

Lemma 3 (1) If a front τ is progressive, then the front after bisecting a triangle of τ is also progressive. (2) If a front τ is progressive, then for any local minimum vertex P the front τ' , obtained from τ by advancing P in time by T_{min} , is progressive.

Proof. (1) By Definition 3, if a triangle PQR of the front τ is (k, l) -progressive, then either of the two smaller triangles after bisecting $\triangle PQR$ is (k, l') -progressive for $l' = \max\{l - 1, 0\}$.

(2) This part was essentially proven by Abedi *et al.* (see [5]) when each triangle PQR on the front τ satisfies progress constraint $\sigma(P'QR)$. Our algorithm ensures $\triangle PQR$ satisfies progress constraint $\tilde{\sigma}(PQR)$, where $\tilde{\sigma}(PQR) \leq \sigma(P'QR)$. Because the progress constraint is monotonic in the slope σ , $\triangle PQR$ automatically satisfies progress constraint $\sigma(P'QR)$. The algebraic details are straightforward [5] and are omitted here for lack of space, to appear in a forthcoming paper. \square

By induction on the number of tent pitching and refinement steps, we have the following theorem.

Theorem 4 Given a triangular mesh $M \in \mathbb{E}^2$ and a target time value T , our algorithm builds a finite tetrahedral mesh of the spacetime domain $M \times [0, T]$, provided each triangle is refined only a finite number of times.

5 Conclusion and open problems

We gave an advancing front spacetime meshing algorithm that unifies previous algorithms [1, 5, 4] which tackled nonlinearity and adaptivity separately. The unified algorithm constructs a $2D \times \text{Time}$ mesh with provable guarantees, for the efficient solution of nonlinear and anisotropic problems. We will report experimental results in the near future.

We would like to extend adaptivity to 3D and higher dimensions. The space domain often changes with time; for instance, in the simulation of rocket fuel combustion, the shape of the residual fuel changes with time. We propose to include other front modification operations, such as edge flips, in addition to pitching inclined tentpoles, into a new meshing algorithm. The new algorithm will track features in spacetime, such as phase and domain boundaries and shock fronts, by aligning mesh facets along such features. The SDG method accurately captures the discontinuous solution when mesh facets align exactly with such singular surfaces.

Acknowledgments The author thanks the other members of the CPSD spacetime group, especially Jeff Erickson and Robert Haber.

References

- [1] R. Abedi, S.-H. Chung, J. Erickson, Y. Fan, M. Garland, D. Guoy, R. Haber, J. M. Sullivan, S. Thite, Y. Zhou. Spacetime meshing with adaptive refinement and coarsening. In *Proc. Symp. Comput. Geom. (SoCG)*, pp. 300–309, 2004.
- [2] J. Erickson, D. Guoy, J. M. Sullivan, A. Üngör. Building space-time meshes over arbitrary spatial domains. In *Proc. 11th Int'l. Meshing Roundtable (IMR)*, pp. 391–402, 2002.
- [3] K. Jegdic, R. Jerrard, R. Haber. Convergence of a causal spacetime discontinuous Galerkin method for Temple 2×2 systems. Preprint.
- [4] S. Thite. Efficient spacetime meshing with nonlocal cone constraints. In *Proc. 13th Int'l. Meshing Roundtable (IMR)*, pp. 47–58, 2004.
- [5] S. Thite. Spacetime meshing for discontinuous Galerkin methods. Ph.D. thesis proposal. Online at <http://www.uiuc.edu/~thite/thesis>, 2004.
- [6] A. Üngör, A. Sheffer. Tent-Pitcher: A meshing algorithm for space-time discontinuous Galerkin methods. In *Proc. 9th Int'l. Meshing Roundtable (IMR)*, pp. 111–122, 2000.