

# 2004 NSF Division of Materials Research ITR Workshop Discontinuous Galerkin Methods in Materials Modeling

R. Abedi, S.-H. Chung, J. Erickson, Y. Fan, M. Garland, D. Guoy, R. Haber, M. Hawker, M. Hills, K. Jegdic, R. Jerrard, L. Kale, J. Palaniappan, B. Petracovici, J. Sullivan, S. Thite, Y. Zhou

Center for Process Simulation and Design • University of Illinois at Urbana-Champaign • Departments of Computer Science, Mathematics and Theoretical & Applied Mechanics

## Hyperbolic PDEs and Conservation Laws in Materials Modeling

Hyperbolic systems play an important role in continuum models of materials systems. Examples include:

- shock wave propagation in solids
- evolution equations for state variables in inelastic constitutive models (porosity, dislocation density, etc.)
- Hamilton-Jacobi level set models for interface kinetics

Hyperbolic systems are among the most difficult problems for numerical simulation.




Figure 1: Martensite–Austenite phase boundary (Credit: Thomas Shield, University of Minnesota)

Their solutions exhibit shocks (or discontinuities) that are difficult to capture on numerical grids. Available numerical methods are imperfect, and the search

for better methods is an active area of research.

We are developing a new analysis technique, spacetime discontinuous Galerkin (SDG) finite element methods, to address this class of problems. SDG methods offer a number of desirable features:

- exact balance on every element
- no global oscillations due to shocks
- $O(N)$  complexity on causal grids
- supports nonconforming, hp-adaptive spacetime meshes
- rich parallel structure, modest communication requirements
- track moving boundaries and interfaces

This poster reports progress in formulating and implementing new SDG methods for elastodynamics and describes on-going work to apply them in multiscale modeling of materials microstructures.




Figure 2: Shock wave propagation in a star-shaped rocket grain

## SDG Finite Element Model

The distinguishing features of the SDG finite element method are:

- inter-element discontinuous basis
- direct spacetime model (in lieu of time-marching in semi-discrete methods)

These lead to several advantages when the method is implemented on causal grids.




Figure 3: Discontinuous Galerkin finite element basis functions




Figure 7: Unstructured spacetime mesh showing adaptive refinement for crack-tip scattering problem. The trajectories of the main shock front, Rayleigh waves, as well as scattered dilatational and shear waves are evident.

## Tent Pitcher: Constructing and Solving Adaptive Spacetime Grids

Our solution relies on a novel spacetime meshing algorithm called “Tent Pitcher”. Given a space mesh  $M$  and a target time  $T$ , Tent Pitcher constructs an unstructured mesh on the spacetime domain  $M \times [0, T]$  using a local advancing front method. The advancing front is a *terrain* in spacetime, initially the space mesh  $M$  at time  $t = 0$ . Tent Pitcher repeatedly chooses a vertex of the front, advances that vertex forward in time to create a “tent”, solves the PDE within that tent, and finally updates the front.

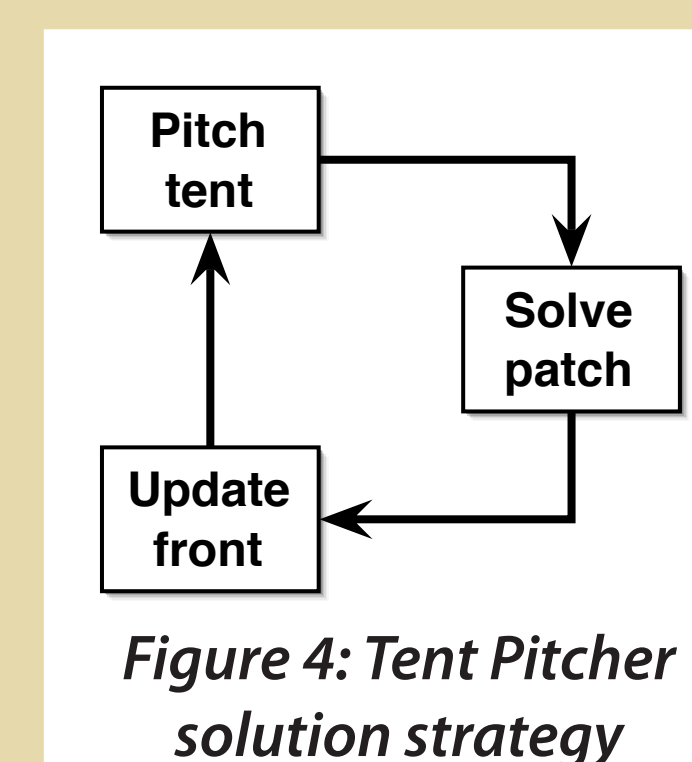


Figure 4: Tent Pitcher solution strategy




Figure 5: Pitching tents in spacetime

The height of each tent is limited in two ways. The *causality* constraint limits the slope of each facet to be less than the inverse of the local wave speed. This constraint ensures that the solution within each tent depends only on the solutions within previous tents. A more technical *progress* constraint ensures that our algorithm makes significant forward progress at every iteration.




Figure 6: Causality constraint: facet separates domains of influence and dependence

We also refine or coarsen the front in response to a posteriori error estimates returned by our spacetime DG solver. If the error within a patch is too large, we refine the front using *newest-vertex refinement*.




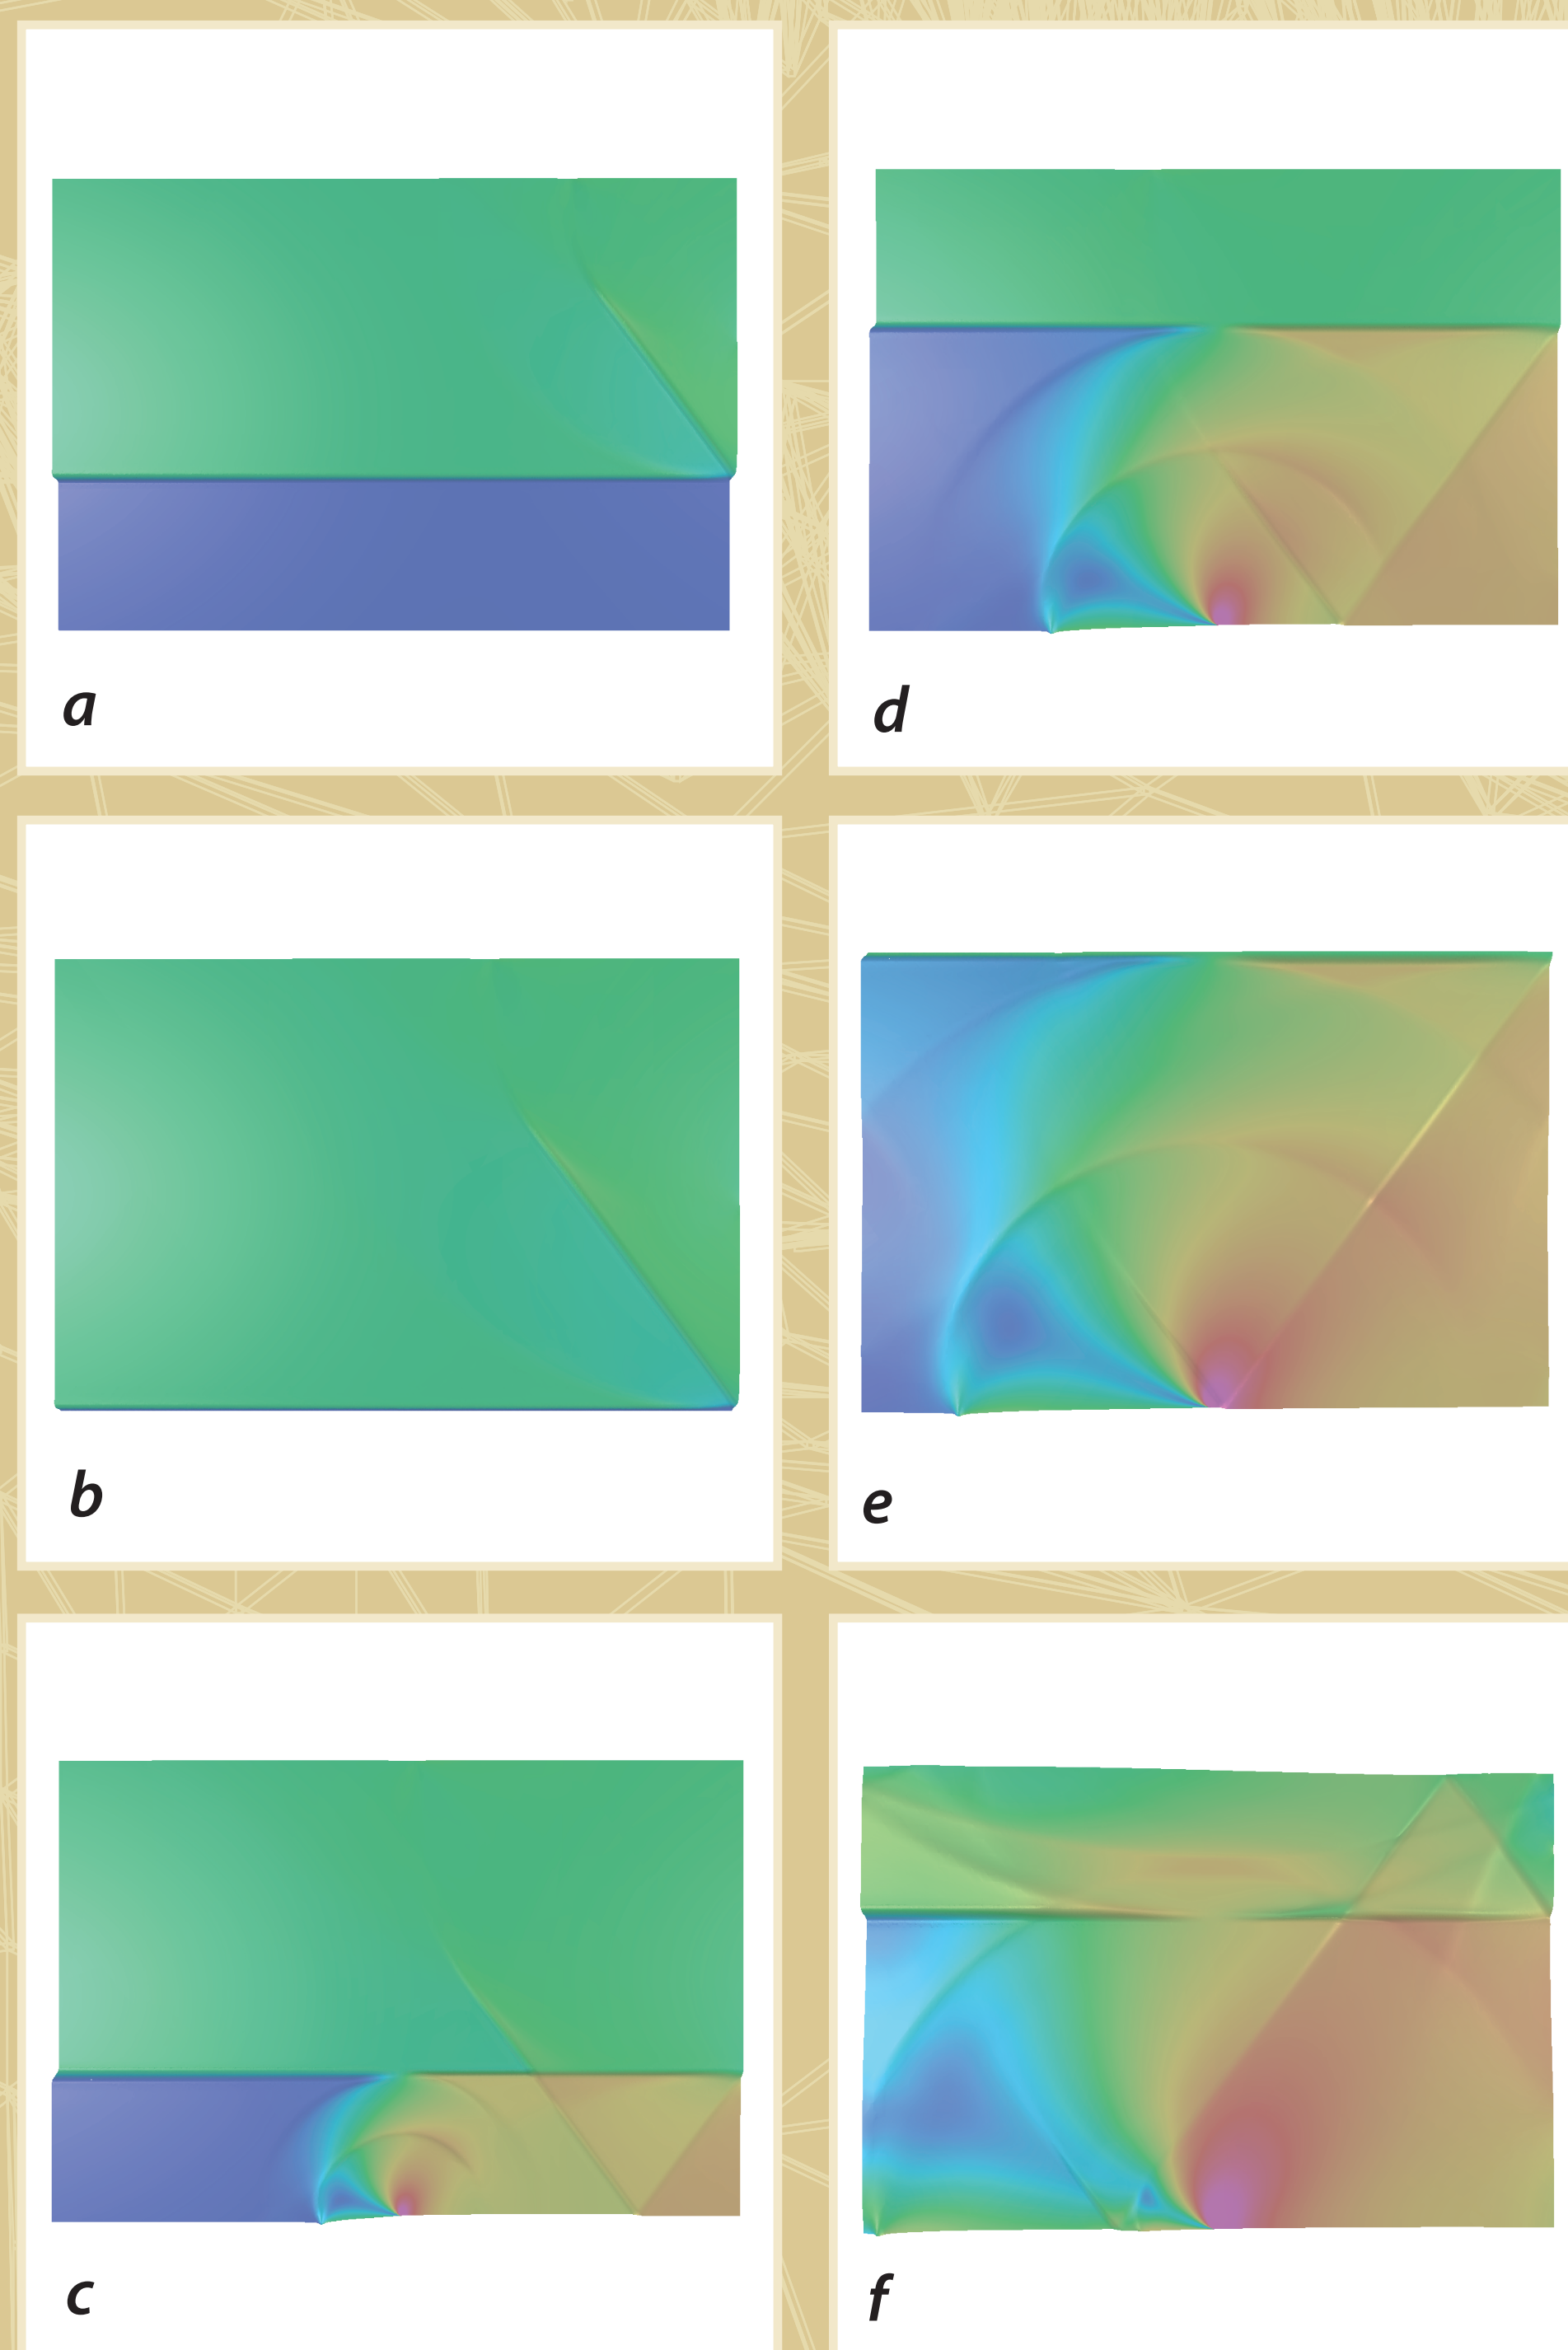
Figure 8: Newest-vertex refinement strategy

By refining the front, we reduce the size of future spacetime elements. If the error within a patch is below some threshold, we attempt to undo earlier refinements. Adaptivity allows us to track shocks and other subtle features of the evolving solution. Our method creates non-conforming spacetime meshes; fortunately, these are supported by our spacetime DG methods.

Our technique has three key advantages:

- *Adaptive*: The size and duration of each spacetime element depends only on the local spatial geometry and the complexity of the local solution.
- *Fast*: We solve a small system of equations for each tent, instead of one huge system for each time step, so our total solution time is only linear in the number of spacetime elements.
- *Flexible*: Tents with no causal relationship can be pitched and solved in any order, or in parallel (cf. parallel solution method)






**Fig. 9** Sequence of images showing scattering of a shock wave by a crack tip. The shock front approaches the crack from the top edge and scatters from the crack tip at bottom center. The shear and pressure components of the scattered waves are clearly visible. The height field depicts velocity magnitude; color depicts strain energy density on a logarithmic scale.

## Pixel-Exact Visualization of Spacetime Data Sets

Special post-processing and rendering techniques are needed to visualize data sets computed on unstructured space-time grids, such as the ones displayed in Fig. 7. We take the trace of the data on a series of constant-time planes to produce animations and still images, such as the ones shown in Fig. 9.


Traditionally, visualizations are produced using standard polygon rendering algorithms that typically drawing one planar polygon per element, with linearly interpolated color. These approximate rendering methods can obscure the true quality of solutions computed with a high-order polynomial bases. In contrast, we wish to produce *pixel-exact* renderings at interactive rates.



**Fig. 12** Comparison of per-vertex rendering (left) with pixel-exact rendering (right)

To achieve this end, we take advantage of the modern programmable graphics processing units (GPUs); these heavily pipelined devices contain up to 64 parallel pixel units on a single card. We use this power to evaluate solution polynomials and to calculate surface normals and lighting, all on a per-pixel basis. This leads to high-accuracy visualizations (see Fig. 12), and allows us to distinguish solution artifacts from rendering artifacts.

Even with very tight error bounds, our SDG solutions can generate slightly different values in neighboring elements due to floating point imprecision. This can lead to pixel drop-outs that overstate the discontinuity. We overdraw all polygon edges to eliminate the pixel drop-outs while retaining legitimate solution discontinuities (Fig. 13).



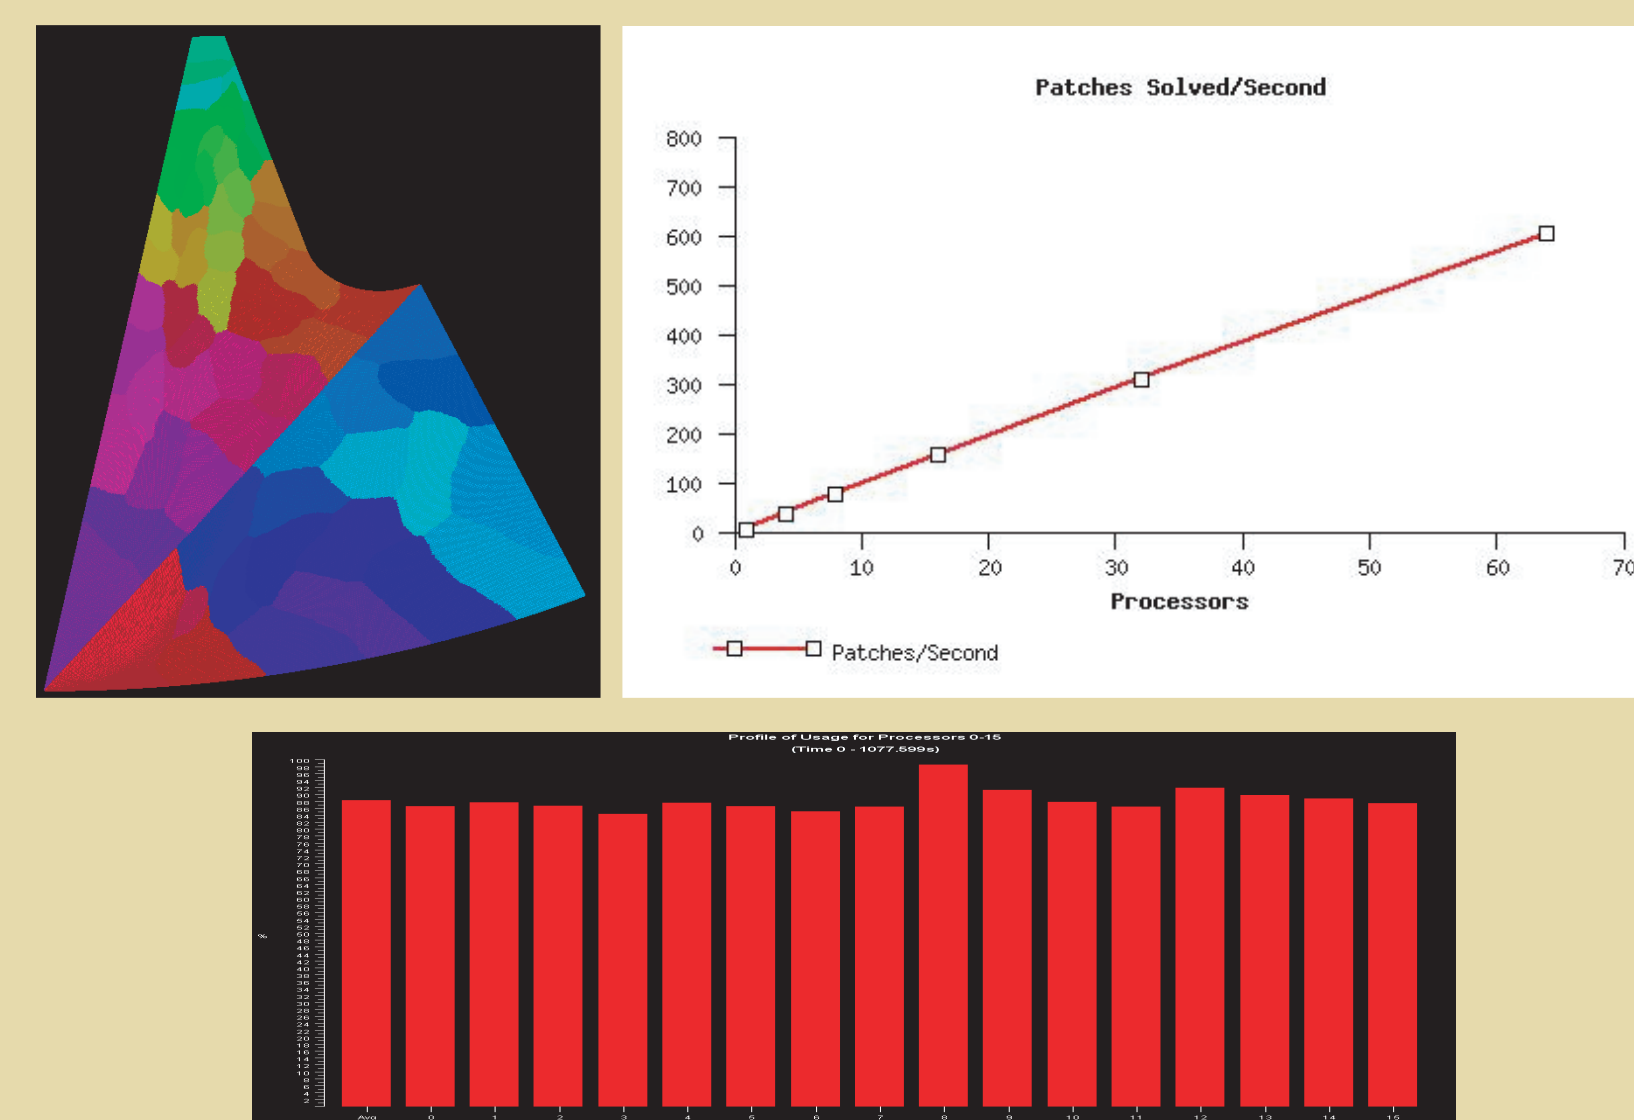
**Fig. 13** Comparison of rendering without (left) and with (right) anti-aliasing along edges

## Parallel Implementation: Finite Element Framework and Charm++

The SDG method is easy to parallelize when it is implemented on patchwise causal meshes, such as those generated by Tent Pitcher, because the solution on each new patch depends only on its immediate neighbors along its inflow boundary. The amount of data per patch is small, making it inexpensive to communicate patch data between processors.

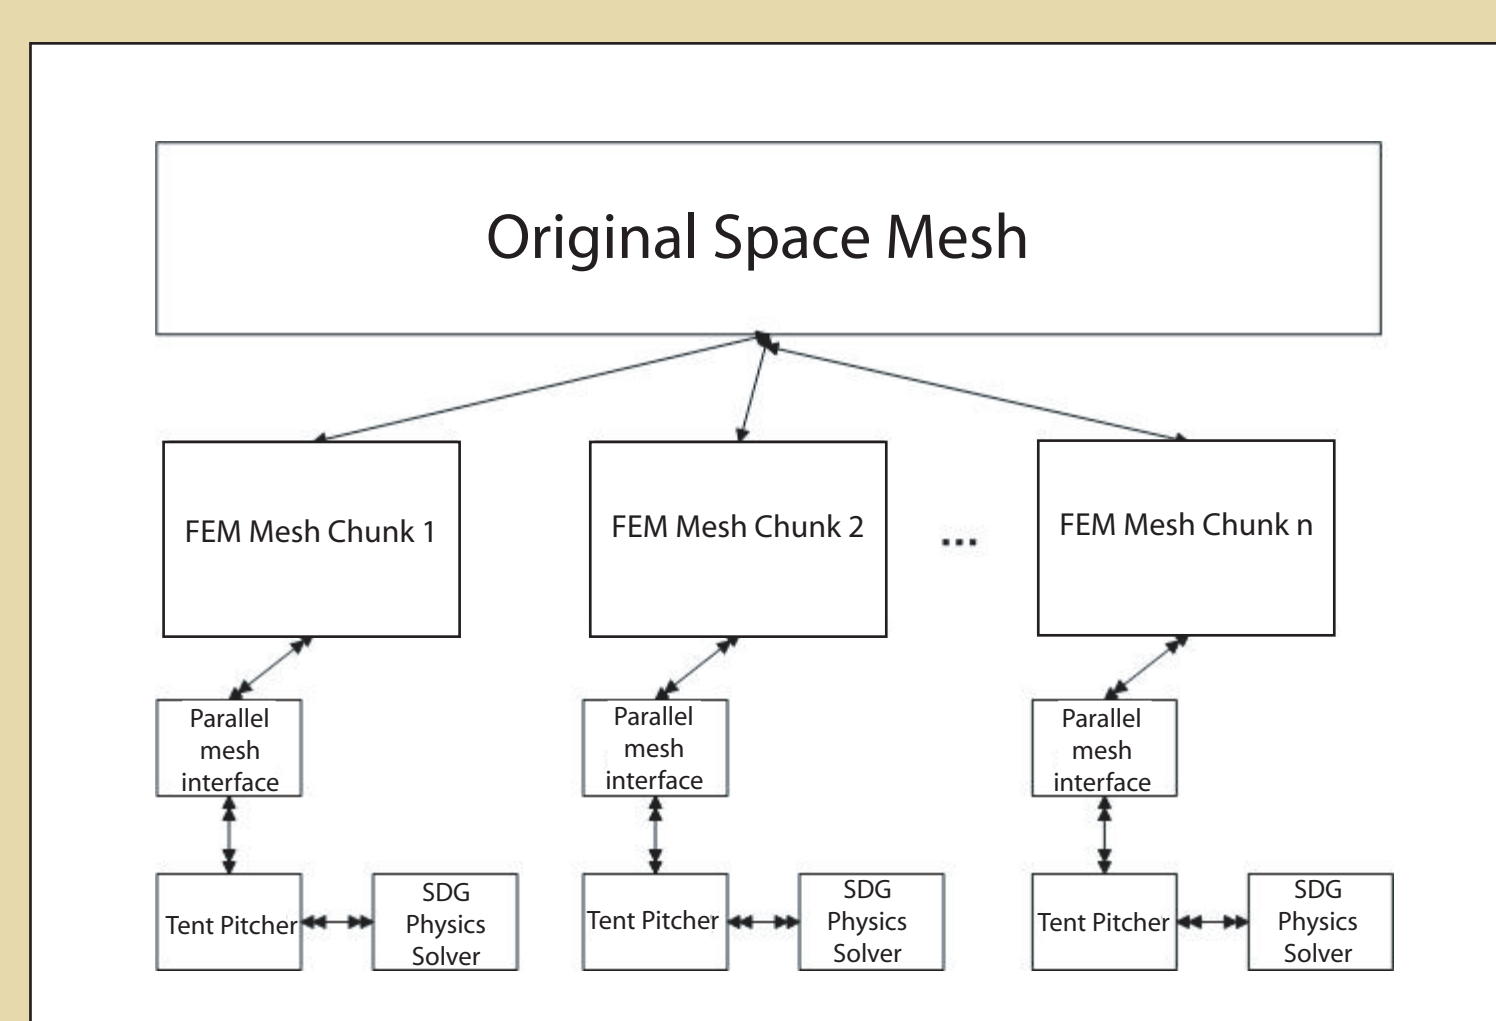
Our parallel implementation is based on the *Finite Element Framework*, a software development environment that handles parallel details with minimal effort by the applications programmer. The framework builds on top of AMPI or native MPI and allows access to *Charm++* (<http://charm.cs.uiuc.edu>).

The framework partitions the space mesh into chunks that run in parallel, each with copies of Tent Pitcher and the SDG physics code, so each chunk can advance its front in spacetime independently.



**Fig. 11** Domain decomposition of space mesh (left), linear scaling of solution rate with number of processors (right), and processor utilization for a 16-processor run (bottom).

The framework assigns each chunk to a virtual processor (VP), with multiple VPs assigned to each real processor. Charm++ provides dynamic load balancing by migrating chunks between real processors as needed. We obtain good processor utilization (95% or better) and nearly linear scaling in the patch solution rate.



**Fig. 10** Parallel software architecture


## Continuing Work

We will continue to develop the spacetime discontinuous Galerkin technology, while increasing our emphasis on its application to a variety of materials systems, including

- Austenite-martensite transitions in shocked shape memory alloys
- Dewetting of inclusions in composites
- Solutions of the time-dependent Schroedinger equation in TD-DFT.
- New atomistic-continuum coupling strategies

Further development of the SDG technology will include:

- Simultaneous parallel/adaptive solutions
- Interface tracking (inclined tent poles)
- Extend to 3 space dimensions x time



**Fig. 14** Scattering around circular inclusions