Applied Geometry

DOE DE-FG02-04ER25657
Final Report

Mathieu Desbrun
California Institute of Technology

The primary purpose of this 3-year DOE-funded research effort, now completed, was to develop *consistent, theoretical foundations of computations on discrete geometry*, to realize the promise of predictive and scalable management of large geometric datasets as handled routinely in applied sciences. Geometry (be it simple 3D shapes or higher dimensional manifolds) is indeed a central and challenging issue from the modeling and computational perspective in several sciences such as mechanics, biology, molecular dynamics, geophysics, as well as engineering. From digital maps of our world, virtual car crash simulation, predictive animation of carbon nano-tubes, to trajectory design of space missions, knowing how to process and animate digital geometry is *key* in many cross-disciplinary research areas.

Scientific Interest and Context

Geometry has been extensively studied for centuries, almost exclusively from a *differential* point of view. However, with the advent of the digital age, the interest directed to smooth surfaces has now partially shifted due to the growing importance of *discrete geometry*. Whether it be in scientific computation or reverse engineering, in remote sensing or medicine, data sets with exquisite geometric details are created daily. The usefulness of such geometric datasets rests on our ability to process them efficiently: from 3D surfaces in graphics to higher dimensional manifolds in mechanics, computational science must deal with sampled geometric data on a daily basis. Alas, current data processing technology does not provide the efficient and geometrically faithful representations demanded by applications. In fact, only a strong grasp of classical differential geometry as well as a sound understanding of the computational needs in research and industry can bring forth the novel theoretical and practical foundations necessary to establish an accurate, structure-preserving discrete differential calculus.

Recently, there has been a noticeable conjunction of approaches in many different fields (mimetic discretizations of continuum mechanics, discrete variational mechanics, etc.) that have all independently come to the conclusion that one must not discretize a continuous model arbitrarily---one must instead directly develop fully-discrete techniques based on discrete variational principles and discrete invariants in order to guarantee reliable computations. Independently of these applied sciences, the graphics community started in the digital age, and has been focusing on discrete techniques (triangle mesh processing, subdivision surfaces, etc.) since its early beginning, using invariants and defining variational approaches for discrete, piecewise-linear 3D geometry. A common research effort *unifying* these different findings and approaches is timely and will be at the core of our work.

The proposed research effort is to study this promising field that we call **Applied Geometry**. It has its roots in a classical and deep connection between mechanics and geometry in the continuous world that goes back hundreds of years, to the work of Newton, Euler and Lagrange and Poincaré: exploiting this link has benefited both fields ever since. The development of similar connections in the discrete world should therefore have a huge impact in the development of algorithms on both ends of the mechanics-geometry spectrum. For the three year duration of this grant, we plan to focus our attention on these two inter-related topics, that cover most of the spectrum between graphics and physical sciences:

- **Discrete Exterior Calculus**: we believe that a ``correct" way to proceed with computations on digital geometry (simplicial complexes in particular) is to develop, *ab initio*, a calculus on discrete manifolds which parallels the calculus on smooth manifolds of arbitrary finite dimension. We are developing such a discrete exterior calculus, defining discrete differential forms along with vector fields and operators (exterior derivative, Hodge star, wedge product, flat and sharp, contraction, and Lie derivative), starting from our work on geometry processing. A formal derivation of all these basic elements in a non-smooth setting guarantees proper preservation of global invariants (through Stoke's theorem), while local discretization leads to superior numerical quality and good convergence properties.

- **Discrete Variational Geometry:** many physical behaviors are dictated by variational principles: Hamilton's principle and Hodge decomposition in mechanics, or Clairaut's theorem in geometry are well-known examples. We are studying how these important principles can be rewritten in the discrete setting. This should lead to discrete constitutive equations, for which time integration will intrinsically preserve the fundamental discrete invariants. It should, additionally, exhibit intricate relationship with our discrete exterior calculus.

Our proposed synergetic mixture of discrete geometry and mechanics, when both are viewed from a unifying variational geometric point of view, also allows us to find common grounds for both communities. As a consequence, we have been able to discuss potential projects on subject ranging from black hole collision simulation, to numerical integration of Schrödinger equation: we have started such a mutually beneficial flow of ideas and results between physical sciences, mathematics, engineering and information technology with teams at Caltech (PI, Jerrold E. Marsden, Peter Schröder, Anthony Leonard), Columbia U. (Eitan Grinspun), NYU (Denis Zorin), Rice U. (Joe Warren), MIT (Fredo Durand), Duke U. (Herbert Edelsbrunner), and abroad (Pierre Alliez and David Cohen-Steiner, INRIA in France, and people from Microsoft Research in China).

Applications Studied and Milestones Achieved

We have been studying discrete differential calculus from a variational, geometric standpoint as sketched in our initial proposal, i.e., through various aspects varying from graphics to mechanics. Each project is one element in the entire “pipeline” needed for discrete differential calculus, going from how to start from a nice mesh to do computations on, to how to reliably handle differential forms on discrete meshes. Our research so far has focused on the following projects:

- **Variational Quadrangle Meshing:** Partitioning a surface into quadrilateral regions is a common requirement in computational science, computer aided geometric design and reverse engineering. Such quad tilings are amenable to a variety of subsequent applications due to their tensor-product nature, such as B-spline fitting, simulation with *finite elements* or *finite differences*, texture atlasing, and addition of highly detailed modulation maps. Quad meshes are particularly useful in modeling as they aptly capture the symmetries of natural or man-made geometry, allowing artists to design simple surfaces using a quite intuitive placement of quad elements. Automatically converting a triangulated surface (issued from a 3D scanner for instance) into a quad mesh is, however, challenging. Stringent topological conditions make quadrangulating a domain or a surface a rather constrained and global problem. Application-dependent meshing requirements (edge orthogonality, alignment of the elements with the geometry, sizing, and mesh regularity) add further hurdles.
  In order to circumvent these issues, we recently introduced a framework for quadrangle meshing of discrete manifolds. Based on discrete differential forms (already detailed above), our method hinges on extending the discrete Laplacian operator (used extensively in modeling and animation) to allow for line singularities and singularities with fractional indices. When assembled into a singularity graph, these line singularities are shown to considerably increase the design flexibility of quad meshing. In particular, control over edge alignments and mesh sizing are unique features of our

novel approach. Another appeal of our method is its robustness and scalability from a numerical viewpoint: we simply *solve a sparse linear system* to generate a pair of piecewise-smooth scalar fields whose isocontours form **a pure quadrangle tiling, with no T-junctions**. This research work, presented at the ACM Symposium on Geometry Processing 2006, further indicates the importance of discrete forms in computational endeavors as shown in Figure 1.

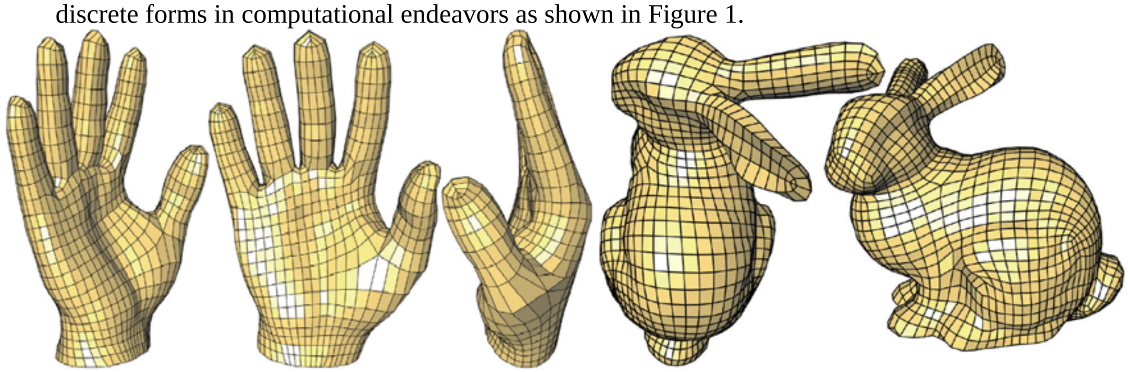

Figure 1: Results of our Quadrangle Meshing technique based on discrete differential forms. By controlling the discontinuities of two gradient fields through a simple linear system, one can remesh a hand (left) or a bunny (right) with only quads, and no T-junctions. Notice the regularity (very few vertices have more (or less) than four neighbors) of the mesh, and the geometrically-pertinent placement of the discontinuities.

- **Variational Tetrahedral Meshing:** Three-dimensional simplicial mesh generation aims at tiling a bounded 3D domain with tetrahedra so that any two of them are either disjoint or sharing a lower dimensional face. Such a discretization of space is **required for most physically-based simulation techniques**: realistic simulation of deformable objects in computer graphics, as well as more general numerical solvers for differential equations in computational science, often needs a discrete domain to apply finite-element or finite-volume methods. Most applications have specific requirements on the size and shape of simplices in the mesh. *Isotropic meshing* is desirable in the common case where nearly-regular tetrahedra (almost equal edge length) are preferred. However, creating high quality tetrahedral meshes is a difficult task for a variety of reasons. First, the mere complexity and size of the resulting meshes requires disciplined and robust data structures and algorithms. There are also basic mathematical difficulties which make tetrahedral meshing significantly harder than its 2D counterpart: the most isotropic 3D element, the regular tetrahedron, does *not* tile 3D space (let alone specific domains), while the equilateral triangle does tile the plane; unlike the 2D case, even well-spaced vertices can create degenerate 3D elements such as slivers. Dealing with boundaries is also fundamentally more difficult in 3D: while there are 2D triangulations conforming to any set of non-intersecting constraints, this is no longer true in 3D. All these facts conspire to make both the development of algorithms and suitable error analyses for the optimal 3D meshing problem very challenging. Given that one can often observe in applications that the worst element in the domain dictates accuracy and/or efficiency, it is clear that great care is required to design the underlying meshes and ensure that they meet the needed quality standards.

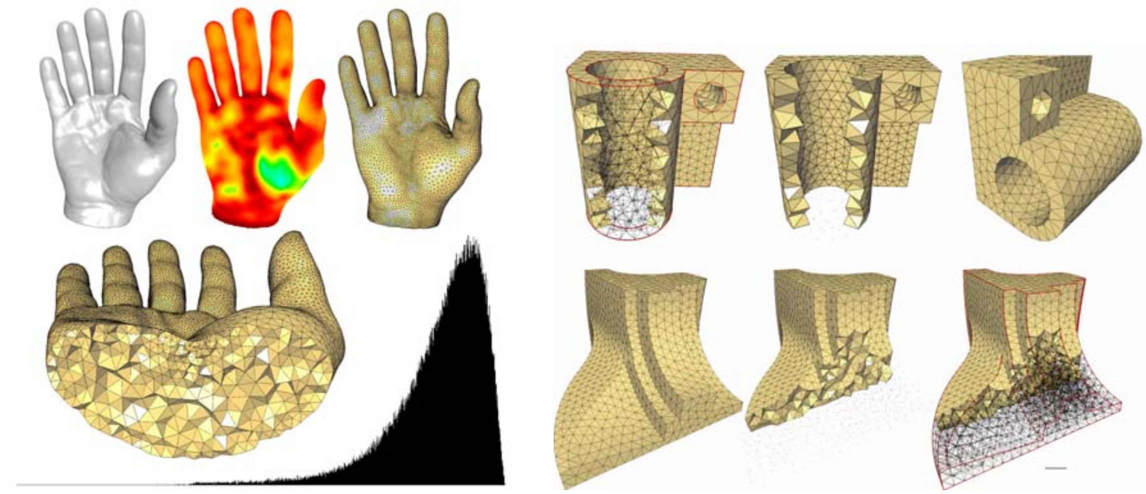

Figure 2: Tetrahedral meshing: (Left) Our newly developed approach allows the tetrahedralization of any 3D region like this scanned hand, with a size of nearly-isotropic tets depending on the local feature size of the shape (i.e., large tets inside, small tets near curved boundaries). (Right) The same technique can handle mechanical parts by conforming to the features. Notice that arbitrary genus is easy to treat with the same algorithm.

In this project, we developed a novel approach to isotropic tetrahedral meshing of complex 3D domains. To achieve robustness, efficiency, and flexibility, our technique consistently minimizes a simple quadratic energy through global updates of both vertex positions *and* connectivity. Mesh design can be controlled easily through a gradation smoothness parameter and a priori selection of the desired number of vertices. The theoretical background of our approach is a simple, underlying *variational principle*, for which we also found an intuitive geometric interpretation tying the shape of the elements to their inertia moments. Our new technique results in high quality meshes as demonstrated in Figure 2. The paper describing this method was published in the ACM Trans. of Graphics journal (special issue of SIGGRAPH '05), and we recently extended this work in a paper that will be presented at SIGGRAPH 2009.

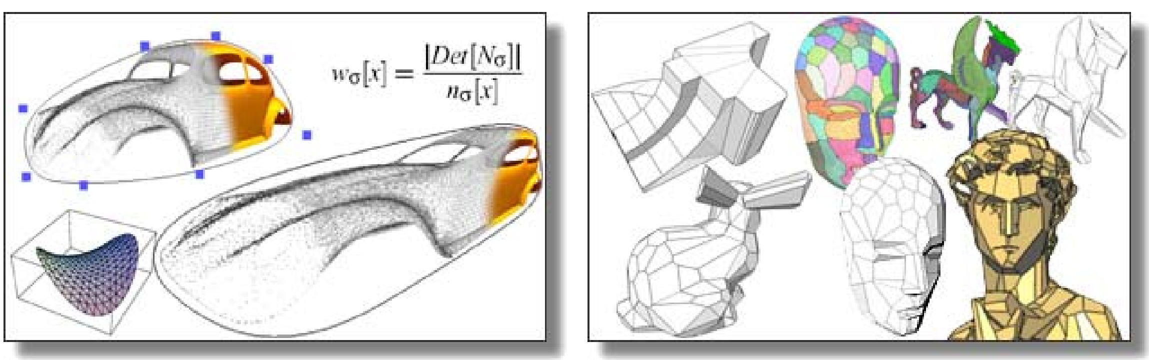

Figure 3: (Left) Barycentric coordinates are a very versatile tool in computations. We have been exploring their use for free-form deformation in a recent submission. (Right) Variational Geometry Approximation allows us to find a near-optimal approximation of curved manifolds in a principled manner. Such technique can be also used to simplify higher-dimensional datasets, such as motion (in space-time) or vector fields.

- **Barycentric Coordinates:** Barycentric coordinates are one of the most basic mathematical tools in many computational sciences: they are a convenient and **coordinate-free** interpolation method. Although the formulas for simplices (triangles, tetrahedra and so on) are widely known and routinely used, there has been no satisfactory extension of these to *arbitrary convex polytopes* despite a plethora of potential applications. Since the beginning of our award, we have worked on a simple, computationally convenient formula of a canonical form of barycentric coordinates valid in arbitrary dimension (in collaboration with Joe Warren from Rice University). The resulting functions are rational, smooth and provably of the lowest possible degree. We have also extended the formulas for convex polytopes to smooth, convex domains, which led to new Green functions with interesting applications in PDE problems. The largely-geometric component of this work has already led to various applications (such as free-from deformation). This recent work has recently been published in the *Advances in Computational Mathematics* journal, and an extension was presented at the Symposium on Geometry Processing. We plan (in the near future—although this opens a vast array of research questions, so it may not be extremely soon) to use these coordinates as basis functions on Voronoi regions, for computational purposes.

- **Discrete Differential Forms for Computational Purposes**: given the overwhelming geometric nature of the most fundamental calculus of these last few centuries, it seems relevant to approach computations from a *geometric standpoint*. This will allow us to clearly express and separate the topological (metric-independent) and geometrical (metric-dependent) components of equations, rendering them amenable to direct and proper discretization. In particular, the resulting discrete treatment will be formally identical to–and will partake of the same properties as—the continuum model: their implementation on a discrete domain will *respect* the intrinsic structure, even at a numerical level. Also, it is interesting to notice that our discrete versions of forms or manifolds are often much simpler to define that their continuous counterparts; basically we can define a *k-form* as a certain type of mapping (integration) from oriented k-dimensional submanifold to a real number. Likewise, even if the continuous idea of orientation is the equivalence class of atlas determined by the Jacobian (intuitively, a point can have two orientations as positive or negative; a curve have 2 directions; a surface can be clockwise or counterclockwise; a volume can have a direction as a right-handed helix or a left-handed one), the discrete counterpart is more straightforward as no notion of atlas is needed. Using values on simplices as integrals of forms offers a principled calculus that generalizes the traditional method of computations (finite element, finite differences, finite volume). An introductory chapter for a book on Discrete Differential Geometry will appear (by the end of the year) on this integral-based idea to foster the use of discrete differential forms.

  We have also developed *novel high-order basis functions* for differential forms. Vertex- and face-based subdivision schemes are now routinely used in geometric modeling and computational science, and their primal/dual relationships are well studied. We have managed to interpret these schemes as defining bases for *discrete differential 0- resp. 2-forms*, and complete the picture by introducing *edge-based* subdivision schemes to construct the missing bases for discrete differential 1-forms. Such subdivision schemes map scalar coefficients on edges from the coarse to the refined mesh and are *intrinsic* to the surface. Our construction is based on treating vertex-, edge-, and face-based subdivision schemes as a *joint triple* and enforcing that subdivision commutes with the topological exterior derivative. We demonstrated our construction for the case of arbitrary topology triangle meshes. Using Loop's scheme for 0-forms and generalized half-box splines for 2- forms results in a *unique* generalized spline scheme for 1-forms, easily incorporated into standard subdivision surface codes. Once a metric is supplied, the scalar 1-form coefficients define a smooth tangent vector field on the underlying subdivision surface. Design of tangent vector fields is made particularly easy with this machinery as Figure 4 shows. This work was presented at ACM SIGGRAPH 2006 venue.

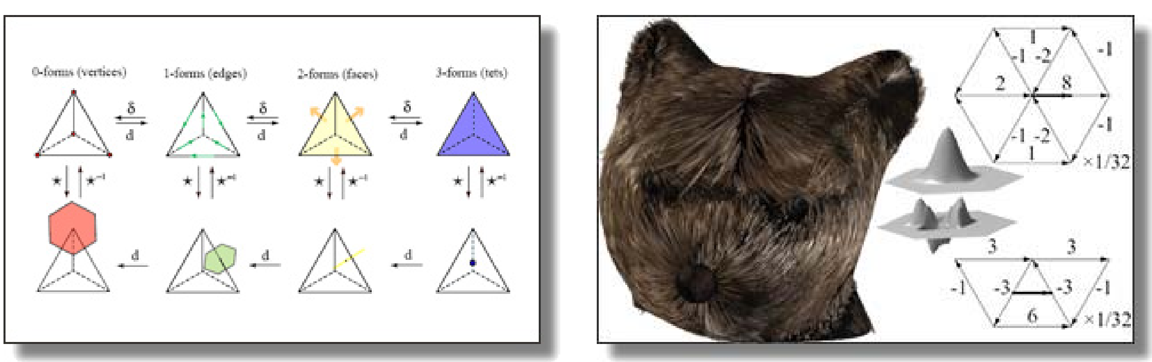

Figure 4: (Left). Discrete forms, equipped with a discrete exterior derivative (dual to the boundary operator) and a discrete Hodge star, provide a direct discretization of the DeRham complex. (Right). We have also defined smooth basis forms to interpolate these discrete forms in the ambient continuous space.

- **Design and manipulation of Tangent Vector Fields:** Smoothly varying tangent vector fields appear in many applications that either deal with flow on curved surfaces, or simply need to control the appearance of surfaces. Examples in graphics include anisotropic, texture synthesis, non-photorealistic rendering, line integral convolution, and spot noise among many others. A huge majority of approaches use coordinates of these vectors in the embedding, leading to both continuity issues on the manifold and the need for constraining these vectors to remain in the tangent plane. In contrast, we have been able to treat these tangent vectors as proxies of one-forms (i.e., one value per edge regardless of the dimension of the embedding space). This allowed us to formulate vector field design as a linear problem by using an intrinsic, coordinate-free approach based on discrete differential forms and the associated Discrete Exterior Calculus (DEC). By representing the field as scalars over mesh edges (i.e., discrete 1-forms), we obtain an intrinsic, coordinate-free formulation in which field smoothness is enforced through discrete Laplace operators. Unlike previous methods, such a formulation leads to a linear system whose sparsity permits efficient pre-factorization. Constraints are incorporated through weighted least squares and can be updated rapidly enough to enable interactive design, as demonstrated in the context of anisotropic texture synthesis (see Figure 5). This work was published at ACM SIGGRAPH 2007.

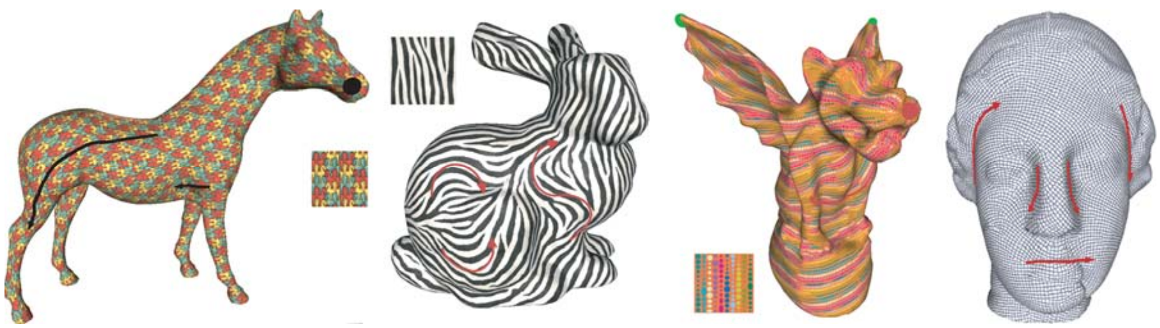

Figure 5: Design of Tangent Vector Fields. Gallery of surface texture synthesis results based on vector fields specified with a variety of constraints, demonstrating that even just a few constraints can quickly build overall fields with pleasing flows. The resulting vector fields are defined through *one value per edge*, and thus do not require the use of coordinates.

- **Spherical Parameterizations:** There is by now a rich literature on the construction of energy-based parameterizations for surface. While much of this work has focused on the planar case, i.e., the mapping of a topological disk region of a given mesh to the plane, spherical parameterizations have been singled out as a special case occurring frequently enough in practice to warrant their own methods. These approaches often extend a specific method known in the planar setting to the sphere. Unfortunately, the intrinsic non-linearity of these extensions coupled with the lack of boundary vertices to anchor the parameterization seriously hinder practical implementation: tailored solvers are often used in conjunction with several vertex constraints in order to obtain non-degenerate solutions. Alas, these fixes do not lead to optimal parameterizations as constraints often introduce severe distortion.
  Consequently, our approach began with the observation that it would be desirable to construct a general procedure to take existing methods (we remained deliberately agnostic as to the particular weights being used) from the planar parameterization case and adapt them to the sphere. In this work, we thus introduced a straightforward technique for easing the computation of spherical parameterizations by a *simple modification of traditional planar parameterization methods*: our spherical energies differ from the usual planar quadratic energies only in the multiplication by a simple factor based on the inverse distance of each triangle from the sphere center. The simplicity of the method (Figure 6) resulted in a paper in the *Journal of Graphics Tools* published in 2006.

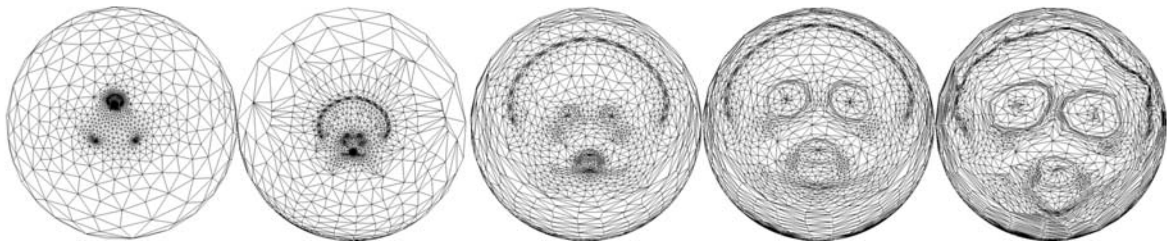

Figure 6: Results of spherical parameterization of a triceratops model using various schemes (leftmost: simplest, Tutte-like embedding) that preserve either angles (2nd from left) or areas (last), or even linear combinations of them (middle ones).

- **Circulation-Preserving Integration Scheme for Fluids:** Physical accuracy, numerical stability, and low computational cost are foremost goals in applied mathematics and computational science. An important ingredient in achieving these goals is the conservation of fundamental motion invariants. For example, rigid or deformable body simulation have benefited greatly from conservation of linear and angular momenta. In the case of fluids, however, none of the current techniques focuses on conserving invariants, and consequently, they often introduce a visually and physically disturbing numerical diffusion of *vorticity*. Just as important is the resolution of complex simulation domains: regular (even if adaptive) grid techniques can be computationally delicate.

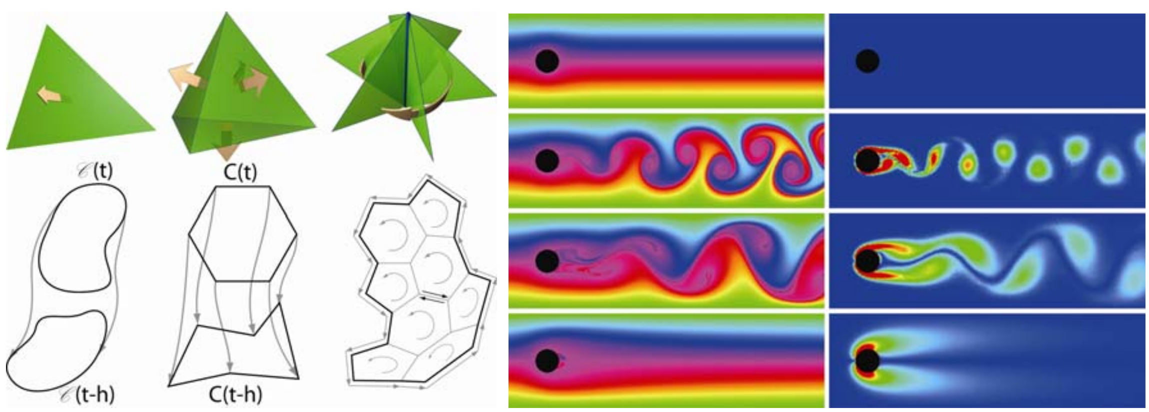

Figure 7: Geometry-Based Simulation of Fluids: using a flux-based discretization (naturally Eulerian and coordinate-free), we can enforce a discrete Kelvin theorem to integrate Euler's equations. Once viscosity is added, this simple, geometry-based integration scheme exhibits the well-known von Karmann vortex streets in the typical flow past an obstacle (rightmost: vorticity plot of the simulation, making evident the creation of vorticity on boundaries).

It is known that it is not possible to exactly preserve momenta *and* total energy simultaneously in the discrete setting as shown in Zhong and Marsden in 1988. One can, however, design integration scheme to keep underlying geometric structures intact as one goes from the continuous to the discrete formulation. More precisely, appropriate geometric discretization of the physics allows one to construct discrete analogs of momenta and energy. Equipped with these discrete structure-preserving quantities, integration schemes can then be designed to enforce their invariant nature. We followed this path by using *vorticity* as our primary simulation variables and designing a time integration scheme which will conserve circulation through vorticity advection—i.e., through a **discrete Kelvin theorem**. As a by-product our velocity fields are divergence free *without* any need to continually re-project to keep this property. For comparison, and to the best of our knowledge, none of the integration schemes proposed in CFD have been designed to satisfy the conservation properties of the underlying equations (aside from the limited case of linearized NS equations).

Consequently, our method offers several new and desirable properties: (1) arbitrary simplicial meshes (triangles in 2D, tetrahedra in 3D) can be used to define the fluid domain; (2) the computations are efficient due to discrete operators with small support; (3) the method is stable for arbitrarily large time steps; (4) it preserves a *discrete circulation*, reducing significantly the usual numerical diffusion of vorticity; and (5) its implementation is straightforward. This work (Figure 7) was published in the ACM *Transactions on Graphics* journal in 2007.

- **Spectral Conformal Parameterization**: We have also introduced a spectral approach to automatically and efficiently obtain discrete free-boundary conformal parameterizations of triangle mesh patches, without the common artifacts due to positional constraints on vertices and without undue bias introduced by sampling irregularity. Due to its central importance in geometry processing, the subject of mesh parameterization has been researched for a number of years. "Parameterizing" a triangle mesh traditionally means computing a correspondence between a discrete, triangulated surface patch (possibly with holes) and a homeomorphic planar mesh through a piecewise linear map. Finding this piecewise linear mapping amounts to assigning each mesh node a pair of coordinates *(u,v)* referring to its position in the planar region (see Figure 8, left). Such (ideally one-to-one) mappings provide a flat parametric space, allowing complex mesh processing operations such as surface fitting and remeshing to be performed directly on a flat domain rather than on the curved, original surface patch. Planar coordinates are also particularly useful to dramatically enhance the visual richness of a 3D surface through texture mapping, both for overly simplified character meshes in game engines and for incredibly detailed surfaces in computer-generated feature films. Consequently, fast methods generating less distortion than current tools are still in high demand.

  In our method, high-quality parameterizations are computed through a constrained minimization of a discrete weighted conformal energy by finding the largest eigenvalue/eigenvector of a generalized eigenvalue problem involving sparse, symmetric matrices. We demonstrated that this novel and robust approach improves on previous linear techniques both quantitatively and qualitatively, as indicated in Figure 8, middle. This work was published and presented at the *Symposium on Geometry Processing* '08.

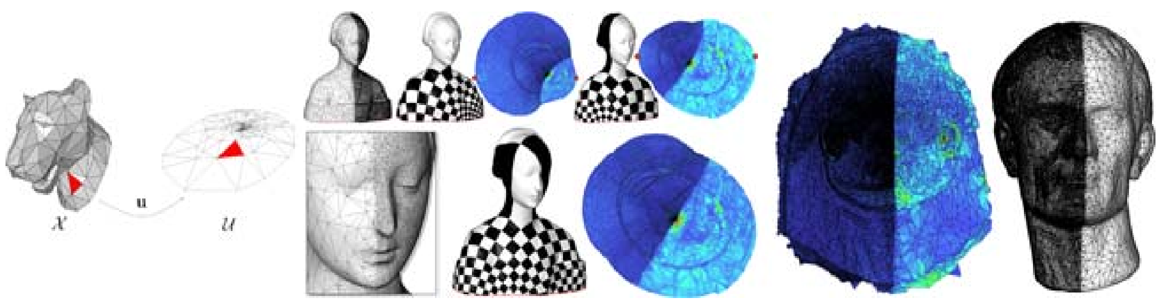

Figure 8: When paramaterizing a triangle mesh (left) with varying sampling rates (middle), previous linear methods (middle, top) fail to capture the symmetry of the mesh in the parameterization. In contrast, our spectral approach (middle/bottom and right) automatically computes a low-distortion conformal map.

- **Discrete Time Integrators**: We have also introduced a general-purpose numerical time integration scheme for Lagrangian dynamical systems—an important computational tool at the core of most physics-based animation techniques. Several features make this particular time integrator highly desirable for computational science and simulation. It numerically preserves important invariants, such as linear and angular momenta; the symplectic nature of the integrator also guarantees a correct energy behavior, even when dissipation and external forces are added; holonomic constraints can also be enforced quite simply; finally, our simple methodology allows for the design of high-order accurate schemes if needed. Two key properties set the method apart from earlier integrators. First, the nonlinear equations that must be solved during an update step are replaced by a minimization of a novel functional, speeding up time stepping. Second, the formulation introduces additional variables that provide key flexibility in the implementation of the method. These properties are achieved using a discrete form of a general variational principle called the Pontryagin-Hamilton principle, expressing time integration in a geometric manner. Non-linear elasticity can be directly simulated as indicated in Figure 9. This work was presented at the *Symposium on Computer Animation* '06.

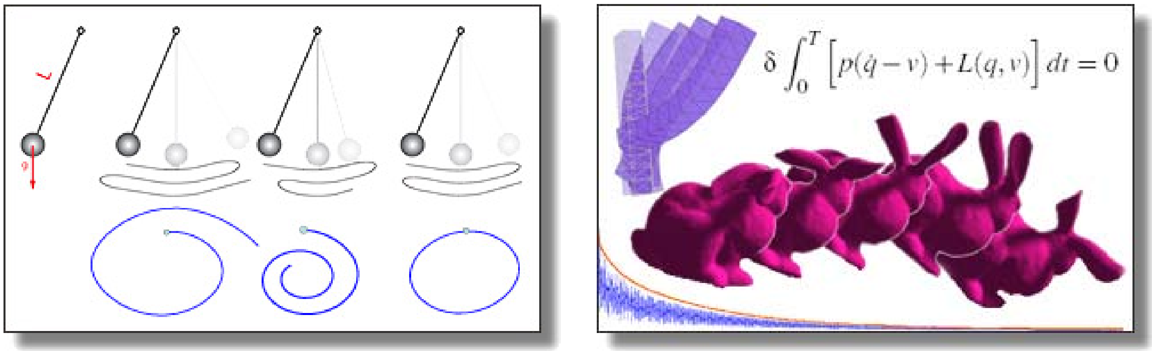

Figure 9: (Left). For a physical system as simple as pendulum, different integrators exhibit very diverse resulting behaviors; symplectic integrators, however, captures the periodic nature at no exra computational cost. (Right). Our geometric, Pontryagin-based symplectic integrators allow fast and predictive simulation of non-linear elasticity, with dampling or not.

- **Reconstruction of Unoriented Point-Sets:** Surface reconstruction from point clouds is motivated by a number of CAGD, point-based graphics, and reverse engineering applications where scattered point samples of a surface need to be turned into a proper, watertight surface mesh. Particularly challenging are point sets generated by laser scanners and hand-held digitizers, as they are often noisy (due to the inherent uncertainty of measurements), unorganized (due to the merging of several scans), and possibly containing large holes (due to occlusions during the acquisition process). In such a context, surface reconstruction can only be approximating—instead of interpolating—as data points are more of an indication of proximity to the surface than actual sample points. While a number of algorithms can now efficiently reconstruct oriented points (i.e., point sets where a normal is provided at each sample), fewer methods are able to approximate raw (unoriented) point sets, with controllable smoothness.

  We have introduced an algorithm for reconstructing watertight surfaces from unoriented point sets. Using the Voronoi diagram of the input point set, we first deduce a tensor field whose principal axes and eccentricities locally represent respectively the most likely direction of the normal to the surface, and the confidence in this direction estimation. An implicit function is then computed by solving a generalized eigenvalue problem such that its gradient is most aligned with the principal axes of the tensor field, providing a best-fitting isosurface reconstruction. Our approach possesses a number of distinguishing features. In particular, the implicit function optimization provides resilience to noise, adjustable fitting to the data, and controllable smoothness of the reconstructed surface. Finally, the use of simplicial meshes (possibly restricted to a thin crust around the input data) and (an)isotropic Laplace operators renders the numerical treatment simple and robust. This work was presented at the *Symposium on Geometry Processing* in 2007 (and won the *best paper* award).

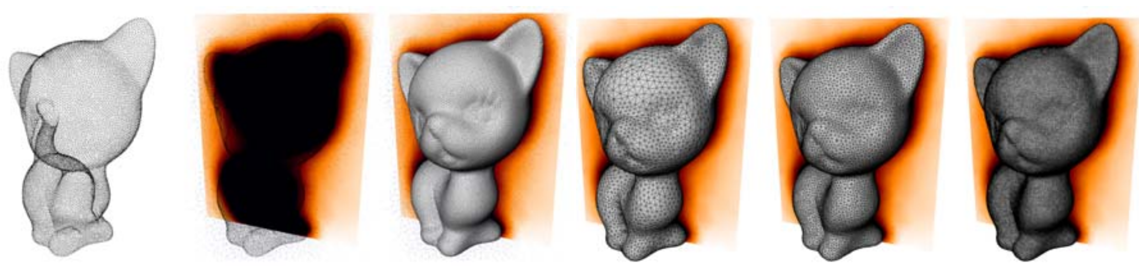

Figure 10: From Point Sets to Surface. Our variational method allows high-fidelity reconstruction of unprocessed point sets (3D scanned kitten statuette, 20K points). The final mesh is extracted via isocontouring of a scalar function computed through optimization involving anisotropic Dirichlet energy and biLaplacian energy. From left to right: input point set, optimized implicit function, and three output meshes at increasing resolutions.

- **Generalized Gradient Flows:** Matching (or registration) of deformable surfaces is a fundamental problem in medical image analysis and computational anatomy. One particularly challenging instance of the problem arises in the field of human brain mapping, where deformable registration of two cortical surfaces is required for intersubject comparisons and intrasubject analysis of neuroanatomical surface data. Related studies include progression of disorders such as Alzheimer's disease, brain growth patterns, genetic influences and the effects of drug abuse on the structure and function of the brain. The challenge in registering two cortices lies in the wide inter-subject variability and the convoluted geometry of the cortical surface, representing a real "stress test" for any general deformable registration technique. Various landmark-based and landmark-free methods have been developed. Parameterization-based techniques first find a mapping between the cortical surface and a plane or a sphere, then align in the parameter domain cortical features such as mean curvature or sulcal landmarks. The often large change in metric due to the mapping needs to be accounted for while performing the alignment process in the parameter domain, adding to the computational costs. Another class of techniques operates directly in the ambient space by finding a 3D warping field that aligns the cortical features. Most of these methods are volume-based, aiming to align image features such as intensities or invariant geometric moments, rather than surfaces. As a result, their matching of the cortices often exhibits inaccuracies.

Despite being routinely required in medical applications, deformable surface registration is notoriously difficult due to large intersubject variability and complex geometry of most medical datasets. We introduced a general and flexible deformable matching framework based on generalized surface flows that efficiently tackles these issues through tailored deformation priors and multiresolution computations. The value of our approach over existing methods was demonstrated for automatic and user-guided cortical registration in a *MICCAI* paper (Figure 11) in 2007, while the geometric foundations were detailed in a *Symposium of Geometry Processing* paper the same year.

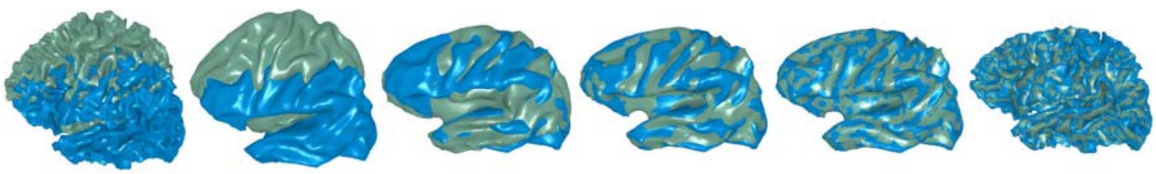

Figure 11: Brian Mapping: Automatically matching a template (grey) to the subject cortex (blue). Partially flattened representations of both surfaces are iteratively aligned using a Hausdorff flow with a smoothing prior. The obtained alignment yields a correspondence between the original surfaces. The final color mix is due the fact that the surfaces lie on each other.

- **Discrete Geometrical Optimal Control Framework:** We have also ventured into the design of optimal motion control algorithms for robotic systems with symmetries. We considered the problem of computing the controls f(t) necessary to move a finite-dimensional mechanical system with configuration space Q from its initial state (position and velocity) to a goal state (position and velocity), while minimizing a cost function of the form:

$$J(q, \dot{q}, f) = \int_0^T C(q(t), \dot{q}(t), f(t))\mathrm{dt}$$

  Minimum control effort problems as well as minimum-time problems can be implemented through the design of the function C. Additional nonlinear equality or inequality constraints on the configuration (such as obstacle avoidance in the environment) and velocity variables can be imposed as well. Systems of interest captured by this formulation include autonomous vehicles such as unmanned helicopters, micro-air vehicles or underwater gliders.

  We have studied the optimal motion control of mechanical systems through a discrete geometric approach. At the core of our formulation is a discrete Lagrange-d'Alembert- Pontryagin variational principle, from which are derived discrete equations of motion that serve as constraints in our optimization framework. We apply this discrete mechanical approach to holonomic systems with symmetries and, as a result, geometric structure and motion invariants are preserved. We illustrate our method by computing optimal trajectories for a simple model of an air vehicle flying through a digital terrain elevation map (see Figure 12), and point out some of the numerical benefits that ensue. These results were presented at the *Robotics: Science and Systems* conference in 2007.

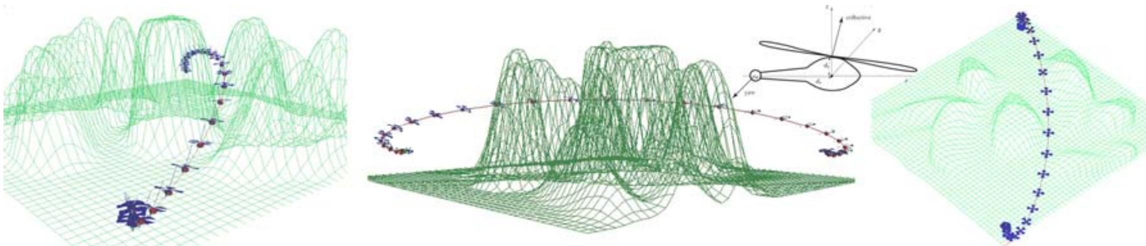

Fig **re**

shows a helicopter path through an outdoor canyon; top and side views of the helicopter trajectory.

- **Eulerian Processing of Surfaces and Foliations:** Evolving surfaces, be it for modeling or animation purposes, is a routine task in Computer Graphics. Over the last decade, the method of choice to process geometry has consisted of a Lagrangian setup where the surface is explicitly stored as a piecewise-linear mesh, and vertices are moved so as to achieve the desired deformation. Great success with this approach has been reported for editing, smoothing, and parameterization, often using variational formulations. Nevertheless, Lagrangian methods come with their share of difficulties, including mesh element degeneracies, self-intersections, and topology changes, all of which require delicate treatment. While some of these issues can be addressed with point sets, the problem of continuous (fine) resampling remains, and the concern of a proper, topologically-sound surface reconstruction arises. Consequently, Eulerian methods emerged as a great alternative to meshes in several applications. One particularly successful Eulerian approach is the Level Set Method (LSM). The LSM methodology has proven very useful in vision, image processing, as well as graphics since the traditional hurdles that mesh processing faces are nicely circumvented due to a parameterization-free treatment. However, other problems arise in this particular Eulerian approach. From a numerical point of view, the LSM relies on finite difference methods applied to a distance function. This specific setup has significant consequences, first and foremost being that volume loss cannot be prevented without using additional (most often Lagrangian) computational devices. A less obvious consequence is that the variational nature of useful flows such as mean curvature motion, which was numerically exploited and proven crucial for mesh processing, is no longer respected.

  To avoid this typical numerical flaws present in current Eulerian methods, we have introduced a new, purely Eulerian framework for geometry processing of surfaces and foliations. Contrary to current Eulerian methods used in graphics and vision, we use conservative methods and a variational interpretation, offering a unified framework for routine surface operations such as smoothing, offsetting, and animation. Computations are performed on a fixed volumetric grid without recourse to Lagrangian techniques such as triangle meshes, particles, or path tracing. At the core of our approach is the use of the Coarea Formula to express area integrals over isosurfaces as volume integrals. This enables the simultaneous processing of multiple isosurfaces, while a single interface can be treated as the special case of a dense foliation. We show that our method is a powerful alternative to conventional geometric representations in delicate cases such as the handling of high-genus surfaces, weighted offsetting, foliation smoothing of medical datasets, and incompressible fluid animation. This work was published in the ACM Transactions on Graphics (SIGGRAPH 2007), see Figure 13.

Figure 13: Our fully-Eulerian discrete variational approach to 3D geometry can be used for (left) outward and inward surface offset (here spatially varying by height), (center) simultaneous smoothing of foliations (all isosurfaces of volumetric medical data), and (right) conservative mass advection for incompressible fluid simulation.

- **Geometric Character of Stress in Mechanics:** We also proposed a reformulation of classical continuum mechanics in terms of bundle-valued exterior forms. Our motivation was to provide a geometric description of force in continuum mechanics leading to an elegant geometric theory that, at the same time, *may enable the development of discrete, space-time integration algorithms for elasticity that respect the underlying geometric structure at the discrete level*. In classical mechanics the traditional approach is to define all the kinematic and kinetic quantities using vector and tensor fields. For example, velocity and traction are both viewed as vector fields and power is defined as their inner product, which is induced from an appropriately defined Riemannian metric. On the other hand, it has long been appreciated in geometric mechanics that *force* should not be viewed as a vector, but rather a one-form. This fits naturally with one of the main properties of a force, namely that when paired with a displacement (a vector), one gets *work*. No metric is needed for this operation of course when force is thought of as a one form. One also sees the same thing when one looks at the tensorial nature of the Euler–Lagrange equations: the equations themselves are natually one-form equations, not vector equations. Despite this, the notion of force as a one-form has not properly been put into the foundations of continuum mechanics. In the geometric approach to continuum mechanics we proposed, traction is defined as an exterior one-form. Consequently, one also has a *metric-independent* notion of *power* as the natural pairing between the velocity vector field and the traction one-form. Although the importance of the geometric character of these fields is already known in mechanics, the classical derivation of the balance laws presented in most works does *not* reflect this geometric understanding.
  In this work, we showed that the stress field in the classical theory of continuum mechanics may be taken to be a covector-valued differential two-form. The balance laws and other fundamental laws of continuum mechanics can be nicely rewritten in terms of this geometric stress. A geometrically-attractive and covariant derivation of the balance laws from the principle of energy balance in terms of this stress was also presented. This work was published in Z. angew. Math. Phys in 2007.

- **Discrete Geometry Processing**: During our geometry-based research on discrete differential calculus, we have also made contributions to graphics and discrete geometry processing where 3D Euclidean space is assumed. In particular, we introduced mesh quilting, a geometric texture synthesis algorithm in which a 3D texture sample given in the form of a triangle mesh is seamlessly applied inside a thin shell around an arbitrary surface through local stitching and deformation. We show that such geometric textures allow interactive and versatile editing and animation, producing compelling visual effects that are difficult to achieve with traditional texturing methods. Unlike pixel-based image quilting, mesh quilting is based on stitching together 3D geometry elements. Our quilting algorithm finds corresponding geometry elements in adjacent texture patches, aligns elements through local deformation, and merges elements to connect texture patches seamlessly. For mesh quilting on curved surfaces, a critical issue is to reduce distortion of geometry elements inside the 3D space of the thin shell. To address this problem we introduce a low-distortion parameterization of the shell space so that geometry elements can be synthesized even on very curved objects without the visual distortion present in previous approaches. We demonstrated how mesh quilting generates convincing decorations for a wide range of geometric textures as illustrated in Figure 14.

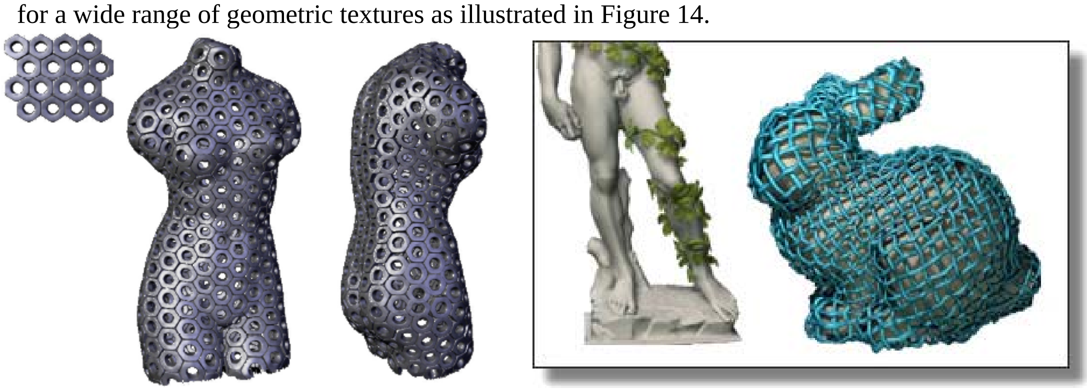

Figure 14: Our research has pioneered the quilting of complex, mesh-based texture-like details over meshes to construct exquisite details. (Left): the Venus model is densely covered with nut elements. (Right): other examples of 3D quilting.

We also presented *mesh puppetry*, a variational framework for detail-preserving mesh manipulation through a set of high-level, intuitive, and interactive design tools. Our approach builds upon traditional rigging by optimizing skeleton position and vertex weights in an integrated manner. New poses and animations are created by specifying a few desired constraints on vertex positions, balance of the character, length and rigidity preservation, joint limits, and/or self-collision avoidance. Our algorithm then adjusts the skeleton and solves for the deformed mesh simultaneously through a novel cascading optimization procedure, allowing real-time manipulation of meshes for fast design of pleasing and realistic poses (see Figure 15). We demonstrated the potential of our framework through an interactive deformation platform and various applications such as deformation transfer and motion retargeting. Both pieces of work were published at SIGGRAPH, the first one in 2006, the second one in 2007.

Figure 15: Mesh Puppetry: The Armadillo mesh (top left) can be deformed to take various sport poses in a matter of seconds. Its body automatically leans forward and raises its left leg backward to keep balance when trying to reach the user-specified (red) target (shot put, bottom left). Fixing only the positions of its hands and feet is enough to make the Armadillo look like it is bouncing off a springboard (high diving, top right); or the pose of a sprint athlete on the finish line (100m, bottom right).

- **Nonholonomic Integrators for Vehicle Motions**: A *vehicle* is an actuated mechanical system that moves and interacts with its environment, such as a car, helicopter, or boat. While vehicles constitute a highly visible component of the world around us, the topic of vehicle dynamics has received little attention in the computer animation literature, and only a few off-the-shelf solutions for vehicle animation and control exist. This deficiency is in sharp contrast with other animation and control tasks such as rigid/articulated/deformable body simulation, fluid phenomena, and character animation for which a plethora of techniques are available. Additionally, human familiarity with vehicles' highly idiosyncratic trajectories makes it difficult (or simply tedious) for artists to capture the essence of vehicle motion. Although locomotion and actuation have been thoroughly studied by roboticists, very few numerical integrators have been developed for complex, *non-holonomic* mechanical systems.

  In the course of our research effort, we introduced general integrators for vehicles which handle both holonomic and non-holonomic constraints. Our work extends the recently developed "geometric Lie group integrators" to provide a principled approach to the design of structure-preserving integrators for vehicles. Compared to previous methods, our approach to nonholonomic systems with symmetries is very general as it can handle *arbitrary group structure*, constraints, and shape dynamics, and is not restricted to a configuration space that is either solely a group or has a Chaplygin-type symmetry. As a result, our formulation contains an additional discrete momentum equation analogous to the continuous case that explicitly accounts for and respects the interaction between symmetries and constraints in the vehicle dynamics. Our resulting numerical schemes provide several practical benefits directly relevant to computer graphics applications. First, a user can easily apply our framework to any vehicle by supplying its Lagrangian and constraints. Second, there is no need to use local coordinates that require expensive chart-switching, or special handling of singularities and numerical drift as required in previous methods. Additionally, fairly large time steps can be used without affecting numerical stability, making the method practical for the frame rates often used in animation. Finally, motion is computed in the minimum state-space dimension, thereby avoiding the computational burden that the conventional use of Lagrange multipliers induces. Consequently, our formulation allows the design of motions for systems with intricate dynamics through a simple algorithmic procedure, while benefiting from the desirable properties of discrete mechanics and Lie group methods such as robust and predictive numerics. Our results have been accepted in *Transactions of Graphics,* and will be published soon.

Figure 16: A snakeboard (left) is animated using our nonholonomic integrator that realistically accounts for hip and foot motion. Our work presents a general framework for designing variational holonomic integrators and structure-respecting nonholonomic integrators for all sorts of vehicles, including cars, helicopters (middle), and boats (right). These Lie group-based integrators are particularly robust for large time steps, and compete in efficiency with RK methods for small time steps.

- **Asynchronous Variational Integrators for Computational Electromagnetism**: Using our expertise on discrete forms, we also introduced a general family of variational, multisymplectic numerical methods for solving Maxwell's equations, using discrete differential forms in spacetime. In doing so, we demonstrate several new results, which apply both to some well-established numerical methods and to new methods introduced here. First, we showed that Yee's finite-difference time-domain (FDTD) scheme, along with a number of related methods, are multisymplectic and derive from a discrete Lagrangian variational principle. Second, we generalized the Yee scheme to unstructured meshes, not just in space, but in 4-dimensional spacetime. This relaxes the need to take uniform time steps, or even to have a preferred time coordinate at all. Finally, as an example of the type of methods that can be developed within this general framework, we introduced a new asynchronous variational integrator (AVI) for solving Maxwell's equations. As Figure 17 shows, basic simulations on very challenging grids show excellent energy and conservation behavior of our asynchronous integrator.

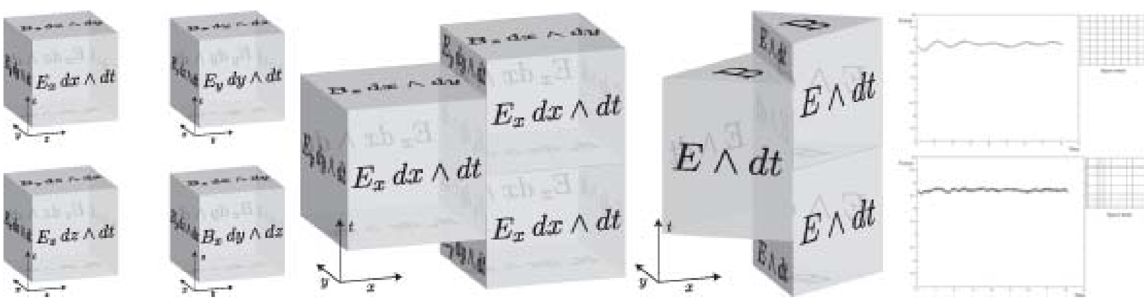

Figure 17: The Yee scheme uses staggered grids to store field values (left). By extending these discrete forms on non-conforming space-time meshes, we were able to extend the Yee scheme to become an asynchronous time integrator (middle), while preserving the good numerical properties (right).

Research Dissemination and Exposure

Publications: Credits to this DOE Early Career funding were included in all the papers (10 SIGGRAPH papers, 1 EuroVis, 2 ACM TOG, 5 SGP, 1 SCA, 1 JGT, 1 RSS, 1 MICCAI, 1 ZAMP, and 1 GI) and ongoing submissions. These papers were original pieces of work, *not* republished results. A book chapter was also published to help disseminate our initial results to the scientific community.

Community Duties : Since the beginning of this grant, we have participated in program committees such as SIGGRAPH 2008/2006/2005, Symposium on Geometry Processing 2004-2009, Shape Modeling International 2004, Symposium on Computer Animation 2004-2009, etc), editorial duties (Associate Editor for ACM Transactions on Graphics), and conference chairmanship (ACM/EG Symposium on Computer Animation ‘04, ACM/EG Symposium on Geometry Processing ‘05).

Advising Duties: Since the beginning of this grant, we have graduated three PhD student (Dr Yiying Tong, November 2004; Ilya Eckstein, August 2007; Roger Donaldson, June2008). We are currently advising three PhD student at Caltech, and co-adivsing two other Caltech graduate students from other departments.

Other Funding: We got awarded (in September 2004) a single-PI NSF grant on Discrete Differential Geometry, that gave us even more leverage on the results from this DOE grant. More recently, we received NSF money from the CPA program to study "eigengeometry", and a NSF grant to study geometric integrators.

Personnel: Funding has been partially used (until 2007) to support a post-doctoral researcher, Dr Yiying Tong, who worked with us on fluid simulation and discrete differential forms; most of the rest of the funding (besides computer equipments and travels) was spent on supporting our graduate students in agreement with our initial budget.

Miscellaneous: We have also been involved with some companies (such as Pixar Animation Studios, Utopia Compression, and recently, Ageia (as part of their Technical Advisory Board) and GM), in an effort to transfer our various technologies to the industry. Although we have spent a limited amount of time on this task, we have built solid connections that will undoubtedly be beneficial in the future.

Summary and Future Endeavors

The results obtained in the course of this grant have, in many ways, exceeded the progress announced in the original proposal—mostly due to our move to Caltech, where cross-disciplinary research can be achieved much more easily. However, we required an extension as our move also prevented us from finding enough graduate students to use up the funds. Even now that this grant is over, we will continue demonstrating the power of Discrete Differential Calculus on a wide number of test cases to prove the relevance of a discrete, geometrically-motivated calculus for computations by building upon the results of this grant.

This DOE award has also resulted in innovative **teaching contributions**: we included our latest theoretical results on discrete manifolds in conventional differential geometry classes, to heighten the students' intuition. In particular, the PI is teaching an Discrete Differential Geometry class where students from Computer Science, Applied Math, and Control & Dynamical Systems can learn more about this topic. Finally, a course at the ACM SIGGRAPH conference (in 2005, 2006, and 2008) on our work was taught to reach our community more efficiently.

As a final comment, we wish to mention that this grant has allowed us to develop enough expertise to continue focusing on simulation. In particular, we have been developing numerical techniques (using applied geometry) to deal with fluids, magnetohydrodynamics, plasma, and charged fluids. We will be seeking DOE funds to pursue and complete these research directions.

For further information, please contact: mathieu@caltech.edu