Natural Sciences and Engineering Research Council of Canada | Conseil de recherches en sciences naturelles et en génie du Canada

FORM 101
Application for a Grant
PART I

Institutional Identifier	
System-ID (for NSERC use only)	210850401
Date	2012/10/26

Family name of applicant	Given name	Initial(s) of all given names	Personal identification no. (PIN)
Bridson	Robert	RE	**Valid** 214193

Department	Institution that will administer the grant
Computer Science	British Columbia

Language of application: [X] English [ ] French

Time (in hours per month) to be devoted to the proposed research / activity: 100

Type of grant applied for: Discovery Grants - Individual

For Strategic Projects, indicate the Target Area and the Research Topic; for Strategic Networks indicate the Target Area.

Title of proposal: Numerical and Geometric Algorithms for Virtual Practical Effects

Provide a maximum of 10 key words that describe this proposal. Use commas to separate them.

computer graphics, visual effects, computer animation, geometric algorithms, exact geometric computation, numerical methods, continuum mechanics, fluid simulation, collision and contact

Research subject code(s) Primary	Secondary	Area of application code(s) Primary	Secondary
2707	2955	801	1205

CERTIFICATION/REQUIREMENTS

If this proposal involves any of the following, check the box(es) and submit the protocol to the university or college's certification committee.

Research involving : Humans [ ] Human pluripotent stem cells [ ] Animals [ ] Biohazards [ ]

Indicate if the proposed research takes place outdoors and if you answered YES to a), b) or c) – Appendix A (Form 101) must be completed• •

[X] NO [ ] YES

TOTAL AMOUNT REQUESTED FROM NSERC

Year 1	Year 2	Year 3	Year 4	Year 5
79,640	79,640	79,640	79,640	79,640

SIGNATURES (Refer to instructions "What do signatures mean?")

It is agreed that the general conditions governing grants as outlined in the NSERC *Program Guide for Professors* apply to any grant made pursuant to this application and are hereby accepted by the applicant and the applicant's employing institution.

Applicant

Head of department

Applicant's department, institution, tel. and fax nos., and e-mail

Computer Science
British Columbia
Tel.: (604) 8221993
FAX: (604) 8225485
rbridson@cs.ubc.ca

Dean of faculty

President of institution (or representative)

The information collected on this form and appendices will be stored in the Personal Information Bank for the appropriate program.

PROTECTED B WHEN COMPLETED

Version française disponible

Canada

Personal identification no. (PIN)	Family name of applicant
Valid 214193	Bridson

SUMMARY OF PROPOSAL FOR PUBLIC RELEASE (Use plain language.)
This plain language summary will be available to the public if your proposal is funded. Although it is not mandatory, you may choose to include your business telephone number and/or your e-mail address to facilitate contact with the public and the media about your research.

Business telephone no. (optional): 604 (822) 1993
E-mail address (optional): rbridson@cs.ubc.ca

In our 2010 Science article, we coined the term 'virtual practical effects' to describe an emerging methodology for technical artists to create digital effects. Supported by physics-based simulation technology, artists can use their knowledge of (and natural intuition for) the real world to virtually build 'physical' mechanisms that create the desired effects. Virtual practical effects has already shown much greater productivity compared to laboriously modelling every phenomenon at a lower level. For example, synchronized crashing ocean waves in 'Avatar' were produced by virtual wave generators modelled after the real thing, running in a physical simulation of the water we wrote. This approach is still in its infancy: my research program tackles the algorithmic challenges still standing in the way of a visual effects revolution.

Part of the research will be dedicated to core low level algorithms, such as efficiently solving the linear and nonlinear systems arising in (multi-)physics simulations, or accurately and robustly tracking a detailed surface as it evolves through time. The other part of the research is concerned with more efficient or more capable numerical models of physical phenomena, like deep ocean waves interacting with boats, solid objects fracturing, or the deformation of thin film bubbles.

The visual effects industry is important and rapidly growing within Canada, and Canada already leads the world in developing 3D animation software (e.g. Maya, Houdini). This program will build on my past research success in these industries, fuelling further growth. In the longer term, it also represents an important step in bringing visual effects capability to anyone with artistic talent and a story to tell - but not the time or inclination to wrestle with low-level technical effects work. Ultimately, the same goals of efficient fidelity to the real world and robustness in the hands of non-numerically-inclined users apply to much broader problems, in particular bringing numerical simulation prototyping to the do-it-yourself/maker revolution.

Other Language Version of Summary (optional).

Personal identification no. (PIN)	Family name of applicant
Valid 214193	Bridson

Before completing this section, **read the instructions** and consult the *Use of Grant Funds* section of the NSERC Program Guide for Professors concerning the eligibility of expenditures for the direct costs of research and the regulations governing the use of grant funds.

TOTAL PROPOSED EXPENDITURES (Include cash expenditures only)

	Year 1	Year 2	Year 3	Year 4	Year 5
1) Salaries and benefits					
a) Students	60,247	60,247	60,247	60,247	60,247
b) Postdoctoral fellows	0	0	0	0	0
c) Technical/professional assistants	0	0	0	0	0
d)	0	0	0	0	0
2) Equipment or facility					
a) Purchase or rental	6,000	6,000	6,000	6,000	6,000
b) Operation and maintenance costs	1,000	1,000	1,000	1,000	1,000
c) User fees	1,593	1,593	1,593	1,593	1,593
3) Materials and supplies	0	0	0	0	0
4) Travel					
a) Conferences	10,800	10,800	10,800	10,800	10,800
b) Field work	0	0	0	0	0
c) Collaboration/consultation	0	0	0	0	0
5) Dissemination costs					
a) Publication costs	0	0	0	0	0
b)	0	0	0	0	0
6) Other (specify)					
a)	0	0	0	0	0
b)	0	0	0	0	0
TOTAL PROPOSED EXPENDITURES	79,640	79,640	79,640	79,640	79,640
Total cash contribution from industry (if applicable)					
Total cash contribution from university (if applicable)					
Total cash contribution from other sources (if applicable)	0	0	0	0	0
TOTAL AMOUNT REQUESTED FROM NSERC (transfer to page 1)	79,640	79,640	79,640	79,640	79,640

Budget Justification

Salaries and Benefits

My main cost is student salaries: top-ups to scholarships and Teaching Assistantships, and full Research Assistantships for part of the year. Based on the extent of the proposed research, the demand from applicants to work with me, the demand from industry and other labs for my graduates, and my own limits as an effective supervisor, I expect to support from Discovery each year the equivalent of two PhD students, one PhD track student, and one MSc student, working one term as a teaching assistant and two terms as a research assistant (scholarships and internships providing for the rest of my group). At minimum department levels, this works out as:

PhD RA: $7,333.33	× 2 terms	× 2 students =	$29,333.32
PhD TA top-up: $1,662.33	× 1 term	× 2 students =	$3,324.66
PhD track RA: $6,667.67	× 2 terms	× 1 student =	$13,335.34
PhD track TA top-up: $1,210.67	× 1 term	× 1 student =	$1,210.67
MSc RA: $6,166.67	× 2 terms	× 1 student =	$12,333.34
MSc TA top-up: $709.67	× 1 term	× 1 student =	$709.67
		total:	$60,247.00

Equipment

I am making do with quite limited hardware currently, and will urgently need to replace workstations with appropriately powerful computers — even if one of our goals is drastically more efficient algorithms for certain cases, developing them and benchmarking against more traditional methods will require both powerful processors and large amounts of memory. I am budgeting one new computer per year, estimated at $6000 (which currently buys a dual processor HP z600 with 48GB RAM and NVIDIA Quadro 2000 graphics card with one LCD monitor). In addition, I budget $1000 per year to maintain older hardware, replacing hard disks, power supplies, etc. or buying additional networked, backed-up, high-performance disk space.

I include $1593 in user fees as our department charges for the direct costs of technical support to research grants. This includes installation and support of equipment; technical support for researchers; print, file and network servers; printing and copying; and similar direct costs.

Travel

In graphics, the best venues for dissemination are conferences such as SIGGRAPH, and attending SIGGRAPH in particular is vital careerwise — with great value from presenting work at the ACM/Eurographics Symposium on Computer Animation as well. I budget an average of $1,800 for air travel, hotel, and registration costs for six people per year, working out to $10,800 total, taking into account limited additional travel support available through my institution for grad students.

Relationship to Other Research Support

Interactive Cinematic-Quality Fluid Animation

This is a one-time hardware donation I received from NVIDIA in 2012, consisting of one Quadro 2000 graphics card, which will be used by my PhD student Todd Keeler to investigate GPU acceleration of the Fast Multipole Method (FMM) as applied to fluid simulation. The $600 figure provided is an estimate of the cost of this hardware: no actual money was supplied.

While an effectively accelerated FMM implementation will be useful in advancing some of the projects in the program, it is not in itself part of the research proposed in this application. Therefore there is no overlap.

GPU-Based Fast Fluids for Video Games

I am in the middle of arranging two coupled MITACS Accelerate internships for the fall of 2012, supporting my PhD student Todd Keeler and Masters student Ryan Goldade at SFU for whom I will serve as supervisor during the internship. The goal of the project is to investigate practical issues associated with using dynamic triangle mesh surface tracking in video games, and in particular making our vortex sheet smoke simulations practical in games. It will have wrapped up before the proposed NSERC Discovery program begins.

This is related but orthogonal to my goals concerning advancing mesh-based surface tracking: for the NSERC Discovery program, I'm chiefly concerned with better algorithms and supporting new applications such as fracture, not the performance of an implementation.

Explanatory Notes

Some explanation may be in order for *Support held in the past 4 years*. Our industrial partners for the MITACS Seed grant with P.I. Dinesh Pai ultimately could not provide the necessary funding, thus there was no money available for my research. The NSERC Strategic Project with P.I. Sidney Fels went ahead, but the overriding goal of the project — to build the ArtiSynth biomechanical modeling software and use it to investigate various medical conditions and interventions — was not compatible with the best interests of the HQP (Christopher Batty) I had speculatively involved, and thus while I helped guide part of the project I did not take any funding personally.

Also, while I have an extremely strong industrial presence, this is clearly not reflected in my research funding. Partly this is due to visual effects studios generally embracing direct collaboration (consulting, buying software) rather than funding external research, and partly my philosophy of strictly separating foundational discovery-oriented research at the university (openly published, patent-free, with software released open-source whenever possible) from systems/implementation efforts in industry.

Proposal

Recent Progress (published)

My Form 100 describes several recent publications directly relevant to the projects outlined below; I will begin by highlighting a few in particular.

We made a breakthrough in continuous collision detection (CCD) [j3], exactly and efficiently determining if mesh elements moving along constant velocity paths collide. Our method involves several strategies to reduce existence of roots of the underlying bound-constrained multilinear system to analysis of more tractable linear systems, a significant departure from approaches taken in prior non-robust tests [22,26,7].

Our new implicit and coupled (and significantly more stable) discretization of surface tension forces between fluids on meshes, based directly on minimizing surface area [c2], promises much greater robustness over prior methods based on time-splitting and/or curvature estimates.

We also created the first fully Lagrangian 3D vortex sheet method with topology change [c1], in the context of buoyant smoke simulation, giving extremely sharp results for inviscid dynamics compared to any other simulation method. Moreover, this runs in linear time w.r.t. the surface area of interest, independent of the volume of fluid, a major advance over prior smoke animation approaches, and with our robust topology change algorithms finally cracked the previous limitation (unbounded sheet growth due to vortex roll-up) of Lagrangian vortex sheets.

Recent Progress (unpublished)

The bottleneck for many simulations, particularly when robustness and stability at large time steps is crucial, is solving large linear systems. One of the best matches for modern architectures is domain decomposition [31], which with the right coarse grid solver can be optimally scalable — but automatically constructing an effective coarse grid approximation to a coupled multiphysics problem with irregular geometry is an open problem. While the usual approach is to use straightforward Galerkin projection with a cleverly built coarse grid basis (which appears very difficult in general), we have drawn inspiration from the Discontinuous Galerkin methodology [3] and used a straightforward subdomain-by-subdomain basis (e.g. low-degree polynomials, with discontinuities across boundaries) with a clever modification of the Galerkin projection to properly treat the interactions between subdomains. Assuming very little user input beyond just the sparse matrix (no need to know the underlying continuous problem, hence we call it Discrete Discontinuous Galerkin), we have an automatic "algebraic" method which scales optimally on a wide variety of elliptic partial differential equations (PDEs) including elasticity and certain forms of the biharmonic equation, though we still are investigating multiphysics problems and more general shell dynamics.

Problems such as realistically simulating an incompressible liquid with fully implicit surface tension, or a fluid coupled with a large-deformation elastic solid, can feature a system of equations with a comparatively small nonlinear component. Nonlinear Conjugate Gradient is arguably the optimal solver in this case, as it reduces to the desired standard Conjugate Gradient (CG) method if the nonlinear part is negligible, avoiding the overhead of Newton-Krylov or other nonlinear solvers. However, CG is only applicable to positive definite problems; in visual effects we frequently encounter situations with incompatible constraints such as an animated solid cavity forcing an "incompressible" fluid contained within to compress, resulting in an inconsistent singular system for which CG blows up. The usual solution of projecting out a known null-space doesn't extend to inconsistent nonlinear constraints, and as we move to multiphysics problems we also hit indefinite problems where CG doesn't work. We have thus extended preconditioned Conjugate Residual [20] to the nonlinear case, including a novel test for stagnation due to inconsistency (and then switch to our new nonlinear extension of the preconditioned Conjugate Gradient Least-Squares method to handle conflicting constraints in a natural least-squares manner).

Objectives

The theme of the program is better enabling "virtual practical effects" (see summary) through better algorithms for physical simulation. The short-to-medium term objectives is progress on multiple projects to remove limitations on simulation. What stands in the way now? Some phenomena of visual interest haven't yet been well modeled for physics-based animation, e.g. thin film bubbles (going beyond simple Voronoi cell or spherical particle representations, and with an eye towards finally creating a well-founded bulk models of foam, an essential part of almost every large water effect). For others, like 3D solids interacting with heavy deep-water ocean waves, we can do it already but the workflow is extremely manually intense, blending expensive 3D fluid simulation with tuned heightfield models and *ad hoc* procedural particle systems — a new approach is needed for progress. In some cases the models are there but the solvers are too slow for an effective design cycle, demanding radically different algorithmic approaches to get to interactivity. Multiphysics simulations, particularly fluids coupled with soft bodies, still haven't seen much adoption in the industry in part due to lack of fast (non-)linear solvers for such systems. Many algorithms, particularly for geometric problems such as collision and contact or surface tracking of fluids, involve non-physical parameters (tolerances, iteration counts, etc.) which may require inconvenient or unintuitive adjustment by users to avoid failure cases — and occasionally fail under all parameter settings.

As my research program is bound up with applications, the anticipated impact discussed below is a primary component of my long-term objectives. There is also an admittedly nebulous motivation intrinsic to my research, curiosity centred on the algorithmic possibilities and limitations of discrete approximations to continuous physics. I focus particularly on linear systems derived from dynamics, and on dynamic surfaces arising in continuum mechanics (both material boundaries and conceptual surfaces such as vortex sheets). For example, for the former I hope to gain insight into the relationship between the quality of a vertex separator of a sparse matrix, and how effectively its (dense) Schur complement can be approximated by a sparse matrix, via analogy to boundary element methods — helping to bridge the divide between direct and iterative solvers. In terms of dynamic surfaces, uncovering the implied regularization of the continuum behind discrete topology changes is another case of a problem I continue to work on while engaging in more concrete projects.

Literature Review

The idea of animating various continuous natural phenomena by direct numerical simulation of the dynamics goes back at least to the fluid dynamics of Yaeger et al. for the film *2010* in 1986 [36] and Terzopoulos et al.'s elastic solids in 1987 [29]. Among many other milestones in the subsequent decades I'd highlight the first forays into continuous collision detection for deformable objects by Moore and Wilhelms in 1988 [22], the first animation use of the full 3D Navier-Stokes equations by Foster and Metaxas in 1996 [13], and a full continuum mechanics approach to fracture by O'Brien and Hodgins in 1999 [24]. These and many, many more contributions laid the groundwork for virtual practical effects. However, the workflow methodology of virtual practical effects — for animation — has not been well documented in the archival literature. Wiebe and Houston's 2004 description of the Tar Monster [32], in essence making a creature out of an simulated fountain of tar, is certainly a precedent, but anecdotally I have heard Industrial Light & Magic applied the concept as early as 2000, building up a stormy ocean for *The Perfect Storm* by simulating wind driving the waves from rest.

Complementary to virtual practical effects is the advent of physics-based lighting and rendering in film. While realistic simulation of the physics of light transport goes back even earlier in computer graphics (the rendering equation was established already in 1986 by Immel et al. [16] and Kajiya [18]), exploiting it fully in lighting design for film is far more recent — see for example McAuley et al.'s SIGGRAPH 2010 course [21]. While the process of lighting used to require placing myriads of nonphysical virtual lights in a scene (e.g. ambient lights that don't correspond to a physical source but provide a desired 'fill') and tweaking purely algorithmic parameters in shaders, the emergence of practical physics-based global illumination renderers allows the lighting artist to instead work intuitively, placing "real" lights and "real" board reflectors off camera just as they would on a real set.

The scope of the specific projects mentioned here is too broad to permit a detailed review, but I will highlight some pertinent papers.

For the Discrete Discontinuous Galerkin domain decomposition effort, an important inspiration was Brezina et al.'s AMGe, an algebraic multigrid method exploiting the unassembled local stiffness matrices [6]: it indicated how a purely algebraic, black-box linear solver can be given readily available and useful information beyond just the matrix, without getting into geometry and rediscretization. We also build on Bassi et al.'s agglomeration-based coarse grid correction for domain decomposition of Discontinuous Galerkin problems [4], but at a purely discrete algebraic level making it applicable to a much wider class of matrices.

Our nonlinear extension of preconditioned MINRES is closely connected to Luenberger's original Conjugate Residual methods [20], and De Sterck's recent extension of GMRES to nonlinear optimization [11]. Several works have also considered the problem of inconsistent singular systems for Krylov solvers, notably Calvetti et al.'s range-restricted GMRES [8] and Choi et al.'s MINRES-QLP [10], but preconditioners haven't yet been satisfactorily included short of going to a full least-squares solve.

The paradigm of Exact Geometric Computation is well surveyed by Yap [37] and the use of exact predicates, such as efficiently implemented by Shewchuk [27], is now common in graphics, e.g. when calculating the intersection of two meshes [5,9].

Work on efficiently and reliably resolving collisions, going beyond resolving instantaneous contact, has recently progressed significantly. Harmon et al. [14] showed the possibility of resolving collisions between elastica guaranteeing basic invariants such as non-interpenetration and causality while also being sure of the simulation progressing to the final time, through careful use of penalties and asynchronous time steps. Ainsley et al. have recently made a leap forward in performance [1]. It remains to be seen if this can be done with a "sharp" method not based on penalties.

Wojtan et al. have advanced mesh-based surface tracking for graphics applications [33,34], extending the front-tracking algorithms from Glimm's group [12]. Müller took another approach to surface-tracking based on extended Marching Cubes, achieving impressive real-time performance [23].

The Boundary Element Method (BEM) first appeared in animation for wind-cloth interaction (as the panel method) by Ling et al. in 1996 [19] and for interactive deformable solids by James and Pai in 1999 [17], but has been much more widely used in computational science and engineering. Of particular interest to the water projects in this proposal is Hou et al.'s landmark work on implicit integration of surface tension regularizing a vortex sheet separating two phases [15] and Xue et al.'s deep-water waves and solid interaction modeling with potential flow [35]. For fracture, standard reviews such as Aliabadi's [2] provide a survey of the dynamic calculations — our interest is more on algorithms for the discrete geometry side.

Methodology

Our CCD work [j3] has opened up the way to exact treatment of existence of roots to bound-constrained polynomial systems. An obvious next step is to apply the same thinking to higher order collision problems, such as those with spline geometry or (using rational splines and re-parameterized time to avoid trigonometric functions, similar to how conics are captured with rational B-spline surfaces) the helical motions of rigid bodies with constant linear and angular velocity.

While we can now provably detect collisions with (or despite!) floating-point arithmetic [j3], resolving those collisions within a simulation is another question. We are looking for a method that, given candidate trajectories for mesh vertices, reliably and efficiently generates collision-free trajectories which are consistent with contact mechanics, in floating-point. Even in a 2D frictionless scenario with no internal or other external forces, exact integration of a single point mass hitting a single edge can give a curved trajectory, so this is by no means a trivial problem. While there are several straightforward approaches, such as generalizing constrained optimization methods developed for instantaneous contact, I suspect the complication of thoroughly handling rounding error will demand ideas from elsewhere. For example, robust Boolean operations on mesh-bounded volumes have been made tractable by switching to a representation (plane equations for faces, with vertices and edges implied through face adjacency) where the output involves just a new combinatorial structure on the same floating-point values [5], determined by predicates. A transformation/discretization similar in spirit could likewise reduce collision resolution to a discrete combinatorial problem.

There is a lot more potential for our Lagrangian surface tracker [j8], both in incremental improvements (adaptive mesh sizing, parallelism, a richer library of mesh operations) and larger advances. Fixed-grid level set methods are essentially limited to first order accuracy around topology changes, but a Lagrangian method can in principle do better by accurately localizing such events in space and time: we will start our investigation by triggering topology changes to the mesh in response to continuous collision detection rather than proximity at fixed time steps. Achieving second order accuracy or better in the presence of topology change, currently impossible with any surface tracking method, could be a game-changer for many applications in scientific computing beyond this proposal.

Having exploited mesh tracking for traditional velocity-pressure liquid simulation [j7,c2,c4] and for vortex sheet smoke dynamics [c1], we will turn to several other fluid applications. The most successful deep ocean wave model in animation is based on free-surface potential flow, but drastically simplified and restricted to periodic height fields [30]. Combining our surface tracking with boundary elements and the Fast Multipole Method we can evolve much more complex ocean waves and include interaction with solids — in time linear with the surface mesh elements, not the volume of the ocean. We will also try to extend our vortex sheet model to thin-flame combustion for fire where the surface evolves with the reaction speed as well as the fluid velocity, and to underwater bubbles, adding surface tension forces. A detailed look at small-scale foam is warranted, by simulating the thin-film dynamics of the walls of each bubble directly with a well-tessellated mesh.

Fully physics-based destruction [24,25] is steadily gaining acceptance in visual effects, but less computationally expensive procedural methods, based on heuristics such as cracking into Voronoi cells, are still more common. For low-deformation fracture, we are adopting a quasi-static boundary element approach to elasticity — using just a surface mesh and thus avoiding the cost of a tetrahedral mesh covering the whole volume — with a new crack propagation algorithm to evolve the surface mesh (deriving from the stress state on the surface vertices whether a crack should proceed, and if so how far and in what direction [2]).

Impact

The most immediate anticipated impact is in visual effects, enabling technical artists to work more productively through more capable and efficient physics simulation. With regards to feasibility, I already have an excellent track record of publishing and/or open-sourcing research then getting it in heavy use in industry either directly via my own implementations or through my colleagues at film studios; I also now have the unique role of architecting the next generation visual effects platform for Autodesk (the dominant software vendor for 3D computer graphics, as well as the engineering/CAD field and more).

Projects such as our linear-time smoke have already garnered interest from the games industry, both for cut-scenes and in-game interaction; the goal of robust simulation running fast enough for effective design is obviously in tune with the needs for game effects.

In the longer term, one of the more interesting promises of this work is democratizing visual effects, so that anyone with creative talent and a story to tell will be able to bring it to cinematic life, without needing a larger studio's help. The simulated virtual world will have high enough fidelity that it automatically produces the intuitively expected effects without need for technical expertise in manually 'cheating' it.

Naturally the algorithms developed in this program are expected to have broader impact in computational science and engineering — particularly linear and nonlinear solvers, exact geometric algorithms for collision and contact calculations, and more advanced surface tracking. Moreover, the demands for virtual practical effects are not so different from what is needed to bring numerical prototyping to everyone. The advent of 3D printing and related flexible manufacturing technology, for example, is promising major strides forward in the do-it-yourself/maker revolution; numerical tools accessible to any user for analyzing and improving their designs before fabrication will take it that much further. The transfer of physical simulation for graphics to numerical prototyping by non-engineers has already begun: my algorithms for coupling liquids with rigid bodies [j11] in our *Naiad* fluid solver [a5] was used in pre-production on *The Hobbit* to prototype real props interacting with real water to good effect. This proposal will help open that door to the third industrial revolution.

Contribution to the Training of Highly Qualified Personnel

Projects

The sampling of projects included in the Proposal are all meant to be primarily lead by students; several have already been matched up with current students (Discrete Discontinuous Galerkin ~ Essex Edwards; thin film bubbles ~ Yufeng Zhu; boundary-based ocean waves and thin flames ~ Todd Keeler; boundary-based fracture ~ Crawford Doran; boundary-based underwater bubbles ~ Xinxin Zhang). Apart from matching projects to the talents and interests of the students, I devise projects with maximum training impact: ones that push the student to deeply learn new material, that can open up a new area and launch a career, and even if practical use in industry may take a lot of further development work are nonetheless tackling a real world problem of significance. I believe this is reflected in my track record.

Training Environment and Opportunities

While each project has a single lead, I encourage a collaborative atmosphere in my group where everyone is aware of everyone's progress and challenges, and can help out at all levels. Recently I've taken this a step further, with joint projects pairing a junior and senior student on one project (e.g. in [j3] and [c1] Brochu was the senior lead, but in the former Edwards took charge of theoretical aspects and in the latter Keeler took on the integral equations and Fast Multipole Method implementation). Apart from greater productivity for all involved, this is improving cohesion and continuity in my group as well as getting new students up to speed that much faster.

Beyond a broad span of research topics, lively group meetings and the collaborative atmosphere of the lab, my students benefit from my flexibility in engagement. For those who have advanced to fully-fledged independent researchers, I'm happy to step back and just provide ideas and guidance; for others that would benefit from it, I relish diving into programming the most difficult parts of the code, taking the lead on the mathematics of some part of the project, debugging code or math, or writing the first draft — always using this as a teaching opportunity.

Students also benefit from my considerable industrial involvement, not just in terms of gaining access to internships at top facilities (exposure which I think is critical for applied research, even in a purely academic career) and job prospects afterwards, but even more so for the insight I can give into what problems matter in the real world, and conversely where the industry doesn't yet realize it needs to head and academia can lead.

Training in the Large

I take the objective of training HQP in a broader sense too, reaching out in multiple ways. One of my frustrations is seeing so few female graduate applicants for my topic of research, despite promising ratios in undergraduate computer graphics classes. In response I have been encouraging female undergraduates in particular to get involved with research early, and have been one of the most active members in my department in terms of outreach to high school students. I firmly believe that with the right approach anyone can appreciate the aesthetics of an algorithm and the thrill of numerically recreating reality, and be inspired to go further. In a similar vein I was invited to give the 2012 I. E. Block Community Lecture, a prestigious SIAM event where the wider intellectual community is invited to gain insight into what applied and industrial mathematics is. My book on fluids [n2] and associated free online course notes have had a wide impact in educating both academics (many graduate-level courses in graphics use them) and practitioners; I'm currently also writing the computer graphics chapter for the Princeton Applied Mathematics Companion which will similarly have a broad audience.

References

[1] S. Ainsley, E. Vouga, E. Grinspun, and R. Tamstorf, 2012, *Speculative parallel asynchronous contact mechanics*, to appear in ACM. Trans. Graph. (Proc. SIGGRAPH Asia), 8 pages.

[2] M. H. Aliabadi and D. P. Rooke, 1991, *The Boundary Element Method*, in Numerical Fracture Mechanics, Springer Netherlands, pp. 90–139.

[3] D. Arnold, F. Brezzi, B. Cockburn, and L. D. Marini, 2002, *Unified analysis of Discontinuous Galerkin methods for elliptic problems*, SIAM J. Numer. Anal., 39, pp. 1749–1779.

[4] F. Bassi, L. Botti, A. Colombo, D.A. Di Pietro, and P. Tesini, 2012, *On the flexibility of agglomeration based physical space Discontinuous Galerkin discretizations*, J. Comp. Phys., 231, pp. 45–65.

[5] G. Bernstein and D. Fussell, 2009, *Fast, exact, linear booleans*, Computer Graphics Forum, 28, pp. 1269–1278.

[6] M. Brezina, A. Cleary, R. Falgout, V. E. Henson, J. Jones, T. Manteuffel, S. McCormick, and J. Ruge, 2000, *Algebraic multigrid based on element interpolation (AMGe)*, SIAM J. Sci. Comput., 22(5), pp. 1570–1592.

[7] R. Bridson, R. Fedkiw, and J. Anderson, 2002, *Robust treatment of collisions, contact and friction for cloth animation*, ACM Trans. Graph. (Proc. SIGGRAPH) 21:3, pp. 594–603.

[8] D. Calvetti, B. Lewis, and L. Reichel, 2000, *GMRES-type methods for inconsistent systems*, Linear Algebra and its Applications, 316, pp. 157–169.

[9] M. Campen and L. Kobbelt, 2010, *Exact and robust (self-)intersections for polygonal meshes*, Computer Graphics Forum, 29, pp. 397–406.

[10] S.-C. T. Choi, C. C. Paige and M. A. Saunders, 2011, *MINRES-QLP: A Krylov subspace method for indefinite or singular symmetric systems*, SIAM J. Sci. Comput., 33(4), pp. 1810–1836.

[11] H. De Sterck, 2012, *Steepest descent preconditioning for nonlinear GMRES optimization*, Numer. Linear Algebra Appl., doi: 10.1002/nla.1837.

[12] J. Du, B. Fix, J. Glimm, X. Jiaa, and X. Lia, 2006, *A simple package for front tracking*, J. Comp. Phys 213:2, pp. 613–628.

[13] N. Foster and D. Metaxas, 1996, *Realistic animation of liquids*, Graph. Models Image Process. 58:5, pp. 471–483.

[14] D. Harmon, E. Vouga, B. Smith, R. Tamstorf, and E. Grinspun, 2009, *Asynchronous contact mechanics*, ACM Trans. Graph. (Proc. SIGGRAPH), 28:3, Article 87, 12 pages.

[15] T. Y. Hou, J. S. Lowengrub, and M. J. Shelley, 1994, *Removing the stiffness from interfacial flows with surface tension*, J. Comput. Phys. 114:2, pp. 312–338.

[16] D. S. Immel, M. F. Cohen, and D. P. Greenberg, 1986, *A radiosity method for non-diffuse environments*, Proc. ACM SIGGRAPH, pp. 133–142.

[17] D. L. James and D. K. Pai, 1999, *ArtDefo: accurate real time deformable objects*, Proc. ACM SIGGRAPH, pp. 65–72.

[18] J. T. Kajiya, 1986, *The rendering equation*. Proc. ACM SIGGRAPH, pp. 143–150.

[19] L. Ling, M. Damodaran, and R. K. L. Gay, 1996, *Aerodynamic force models for animating cloth motion in air flow*, The Visual Computer 12:2, pp. 84–104.

[20] D. G. Luenberger, 1970, *The Conjugate Residual method for constrained minimizations problems*, SIAM J. Numer. Anal., 7, pp. 390-398.

[21] S. McAuley, S. Hill, N. Hoffman, Y. Gotanda, B. Smits, B. Burley, and A. Martinez, 2012, *Practical physically-based shading in film and game production*, ACM SIGGRAPH Courses, Article 10, 7 pages.

[22] M. Moore and J. Wilhelms, 1988, *Collision detection and response for computer animation*, Proc. ACM SIGGRAPH, pp. 289–298.

[23] M. Müller, 2009, *Fast and robust tracking of fluid surfaces*, Proc. ACM/Eurographics Symp. Comp. Anim., pp. 237–245.

[24] J. F. O'Brien and J. K. Hodgins, 1999, *Graphical modeling and animation of brittle fracture*, Proc. ACM SIGGRAPH, pp. 137–146.

[25] E. G. Parker and J. F. O'Brien, 2009, *Real-time deformation and fracture in a game environment*, Proc. ACM/Eurographics Symp. Comp. Anim., pp. 156–166.

[26] X. Provot, 1997, *Collision and self-collision handling in cloth model dedicated to design garment*, Proc. Graphics Interface, pp. 177–189.

[27] J. R. Shewchuk, 1996, *Robust adaptive floating-point geometric predicates*, Proc. Symp. Comput. Geom., pp. 141–150.

[28] J. Stam, 2009, *Nucleus: Towards a unified dynamics solver for computer graphics*, IEEE CADCG, pp. 1–11

[29] D. Terzopoulos, J. Platt, A. Barr, and K. Fleischer, 1987, *Elastically deformable models*, Proc. ACM SIGGRAPH, pp. 205–214.

[30] J. Tessendorf, 2002, *Simulating Ocean Water*, ACM SIGGRAPH Courses, Article 9.

[31] A. Toselli and O. Widlund, 2004, *Domain Decomposition Methods - Algorithms and Theory*, vol. 34 of Springer Series in Computational Mathematics, Springer.

[32] M. Wiebe and B. Houston, 2004, *The tar monster: creating a character with fluid simulation*, ACM SIGGRAPH Sketches, p. 64.

[33] C. Wojtan, N. Thürey, M. Gross, and G. Turk, 2009, *Deforming meshes that split and merge*, ACM. Trans. Graph. (Proc. SIGGRAPH), 10 pages.

[34] C. Wojtan, N. Thürey, M. Gross, and G. Turk, 2010, *Physics-inspired topology changes for thin fluid features*, ACM. Trans. Graph. (Proc. SIGGRAPH), 8 pages.

[35] M. Xue, H. Xu, Y. Liu, and D. K. P. Yue, 2001, *Computations of fully nonlinear three-dimensional wave–wave and wave–body interactions*, J. Fluid Mech., 438, pp. 11–39 & 41–66.

[36] L. Yaeger, C. Upson, and R. Myers, 1986, *Combining physical and visual simulation: creation of the planet Jupiter for the film "2010"*, Proc. ACM SIGGRAPH, pp. 85–93.

[37] C. Yap, 2004, *Robust geometric computation*, In CRC Handbook of Computational and Discrete Geometry, J. E. Goodman and J. O'Rourke, Eds., 2 ed. CRC Press LLC, ch. 41.