ITR/ACS: Simulation of Flows with Dynamic Interfaces on Multi-Teraflop Computers

Guy E. Blelloch, Omar Ghattas, Gary L. Miller, Noel J. Walkington
Carnegie Mellon University

James F. Antaki, Bartley P. Griffith, Marina V. Kameneva
Robert L. Kormos, William R. Wagner, ZhongJun Wu
McGowan Center for Artificial Organ Development, University of Pittsburgh Medical Center

George M. Turkiyyah
University of Washington

1. Overview

We propose to develop advanced parallel geometric and numerical algorithms and software for simulating complex flows with dynamic interfaces. The development of scalable, parallel high-accuracy algorithms for simulating such flows poses enormous challenges, particularly on systems with thousands of processors. We will use the resulting tools to simulate blood flow in artificial heart devices. This application provides an excellent testbed for the methods we develop: simulation-based artificial organ design is extremely computationally challenging and of critical societal importance.

Flows with dynamic interfaces arise in many fluid-solid and fluid-fluid interaction problems, and are among the most difficult computational problems in continuum mechanics. Examples abound in the aerospace, automotive, biomedical, chemical, marine, materials, and wind engineering sciences. These include large-amplitude vibrations of such flexible aerodynamic components as high aspect ratio wings and blades; flows of mixtures and slurries; wind-induced deformation of towers, antennas, and lightweight bridges; hydrodynamic flows around offshore structures; interaction of biofluids with elastic vessels; and materials phase transition problems. We are particularly interested in modeling the flow of blood, which is a mixture of interacting solid cells and fluid plasma. Current blood flow models are macroscopic, treating the mixture as a homogeneous continuum. Microstructural models resolve individual cell deformations and interactions with the surrounding fluid plasma. Because of the computational difficulties of resolving tens of thousands of deforming cellular interfaces, no one to date has simulated realistic blood flows at the microstructural level. Yet such simulations are necessary in order to gain a better understanding of blood damage—which is central to improved artificial organ design—and for the development of more rational macroscopic blood models.

Parallel flow solvers on fixed domains are reasonably well understood. In contrast, simulating flows with dynamic interfaces is much more difficult. The central challenges are to develop numerical algorithms that stably and accurately couple the moving fluid and solid domains and resolve the deforming interfaces, and geometric algorithms for evolving and managing the resulting dynamic particle/mesh systems. The associated dynamic data structures are particularly troublesome on highly parallel computers, which are made necessary by the complexity of many applications. Most current methods approach the difficulties of dynamic interfaces by computing the flow on a fixed, regular grid. The effect of the dynamic interfaces is then incorporated either through some

type of constraint or force representing the interface, or through an auxiliary field variable that signifies the presence of fluid or solid material at a spatial point. Parallelizing these methods is relatively straightforward, since the flow is computed on a fixed grid. However, the resulting fixed resolution is a serious disadvantage if one wants to vary resolution sharply within the grid. This is the case for example when local interfacial dynamics are critical, as in blood flow or phase change problems.

Our approach is radically different. We will treat the fluid and solid domains as collections of particles, with associated meshes, that evolve over time, and devise numerical algorithms that couple the fluid and solid together seamlessly. We will attack the difficulty of generating and managing a constantly evolving mesh/particle system by creating fundamentally new highly parallel and scalable algorithms for the convex hull, Delaunay triangulation, meshing, partitioning, and N-body components. Our preliminary 2D work demonstrates that the resulting geometric computations can be made very cheap compared to numerical computations. Despite the conventional wisdom on parallel dynamic mesh methods, we believe that—with careful attention to fundamental algorithmic issues—flow simulations on constantly evolving domains can be made to scale to the thousands of processors that characterize multi-teraflop systems.

While microstructural blood flow modeling will serve as our first application, the computational algorithms and software we create will be more widely applicable to a variety of fluid–solid interaction problems. More generally, the core parallel computational geometry kernels—convex hull, Delaunay triangulation, coarsening/refinement, partitioning, N-body—provide generic support for the geometric computations underlying many dynamic irregular problems. We will create and publically distribute a portable library of efficient implementations of these algorithms. Much as the PETSc library has greatly simplified the task of programming parallel PDE solvers by providing many of the necessary *numerical* kernels, we envision a library of parallel *geometric* kernels being of great benefit across a wide range of scientific computing problems that involve dynamic meshes.

We have assembled a multidisciplinary team that combines Carnegie Mellon's leadership in computer and computational science with the University of Pittsburgh Medical Center's world-class program in artificial organs. This project will support 11 graduate students and a group of undergraduates. These students will be part of a new program at CMU in Computational Science and Engineering that we are in the process of establishing. The proposed project will also be part of that program, and we believe will serve as an archetype of how applications, computational, computer, and mathematical scientists can work together to tackle societal problems that cannot be addressed solely from the vantage of any one discipline. Moreover, we intend to communicate our work to the broader public (as we have done in the past), in the process demonstrating how high end computing can contribute to improving the health of our society.

2. Motivating Application: Microstructural Blood Flow Modeling in Artificial Heart Devices

We will apply the geometric and numerical algorithms for simulation of flows with dynamic interfaces to the problem of modeling blood flow through artificial heart devices. This application serves as an excellent testbed for two reasons. First, its extreme computational complexity results in a need for terascale computing, and will tax our algorithms, software, and hardware to their

limits. Second, it permits us to contribute to a problem of great societal importance. Below we begin by recalling some essential statistics highlighting an urgent need for prosthetic organs, and some preliminary work already completed by our bioengineering team members, both of which motivate the need for more accurate blood flow modeling. We finish with a brief discussion of blood flow models and indicate how the proposed research is essential for their development.

2.1. The need for simulation-based artificial organ design

The shortage of donor organs has reached alarming levels. A recent article in the *New York Times* proclaims that transplant programs in the New York region "virtually ground to a halt" in early January due to a precipitous drop in donors [40]. The United Network for Organ Sharing estimates that the number of Americans awaiting transplantation has grown to 60,000, of whom 8% can expect to die while waiting [40]. Despite recent advances in such progressive therapies as tissue engineering, trans-species transplantation, and gene therapy, an artificial organ represents the only hope for the majority of these patients in the foreseeable future. As the population continues to age, not only will the demand for artificial organs increase, but the devices themselves will be expected to last longer and be more efficient.

Despite decades of research and the critical demand for artificial organs, their development has progressed at an unacceptably slow pace. It is insightful to compare and contrast the design of prosthetic organs with the design of aircraft. Both require sophisticated fluid dynamic computations, are expensive to build and test, and place a premium on safety and reliability. However, the aircraft industry has sophisticated mathematical and numerical models that facilitate the design of safe, economical aircraft, while the analogous models for artificial organs are still in an embryonic stage. Significant research will be required to develop the necessary models. The resulting software tools will allow extensive computational testing and optimization of proposed designs *prior to* the initiation of expensive animal and clinical trials. The result will be a reduction in the incidence of suboptimal devices entering the clinical trials phase, where failure is catastrophic and redesign is not economically feasible. Computer modeling will not only reduce the design cycle time, but provide greater insight into the behavior of such systems, leading to superior designs.

The development by our bioengineering team members of a next-generation rotary artificial heart assist device (Figure 1a) demonstrates the tremendous potential of high-end simulation in artificial organ design. Figure 1b depicts the results of design improvements that were enabled by computer simulation [7, 8]. A prototype of the improved design was implanted recently in a cow and exhibited an order-of-magnitude reduction in blood damage compared to a previous device. We could never have achieved this performance improvement without computer modeling. As important as this milestone was, much significant work remains: the simulations were based on a homogeneous, Newtonian model of blood, which is a very simplified approximation of its true character. Any design to be used in humans must be based on more realistic blood models in order to predict actual flows and physical damage.

2.2. Blood flow models

Blood is fundamentally a complex, non-Newtonian fluid-solid mixture (which is why blood simulations lag those in aerodynamics): it consists of deformable cellular bodies, primarily red blood

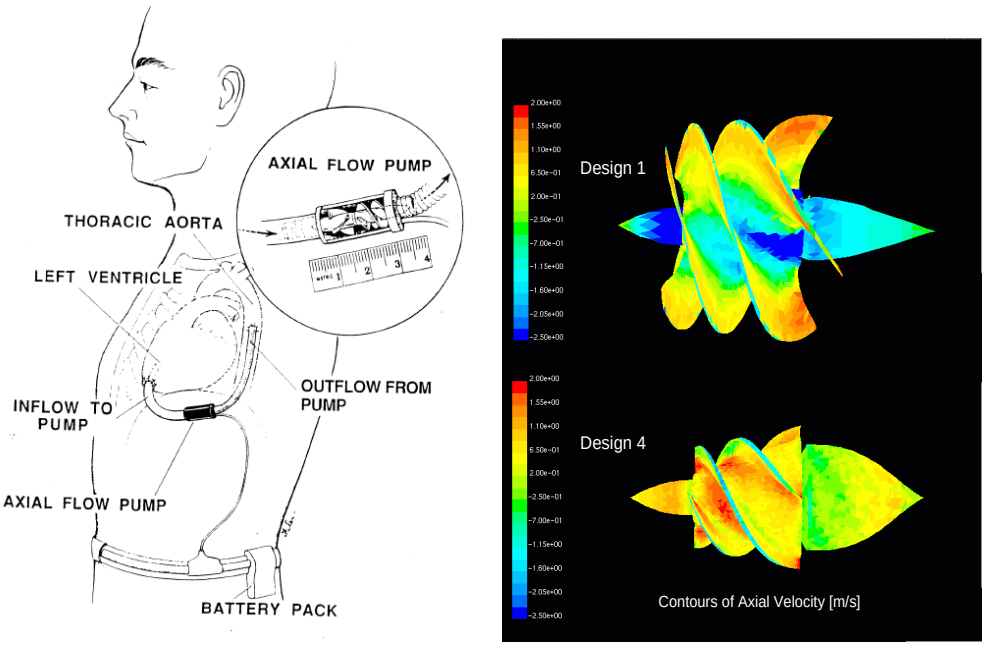

(a) (b)

Figure 1: (a) Schematic of artificial heart device. (b) Initial and computer-optimized design, showing improvement in flow field near surface. Electronic version in color.

cells (RBCs), suspended in an essentially Newtonian plasma. This is illustrated in Figure 2a, which depicts blood flow through a 12μm arteriole. The red blood cells are severely deformed—in a hydrostatic state they assume the shape of axisymmetric biconcave disks. Experimental work, employing cardiovascular devices under physiological operating conditions, clearly demonstrates significant non-Newtonian behavior of blood flows [28, 32, 30].

Blood's microstructure, particularly the RBCs, gives rise to non-Newtonian behavior, including (1) shear-thinning—due to the breakup of red blood cell aggregates at low shear, and the deformation of cells at high shear; (2) viscoelasticity—due to two microscopic mechanisms of storage and release of energy, namely cell deformation and cell aggregation, which contribute an elastic component to the rheological behaviour of the suspension; (3) shear-induced anisotropy—due to the elongation and alignment of cells with the flow at high strain rates, thereby defining a "preferred direction;" and (4) margination—inhomogeneous distribution of the cellular phase in the flow (for example, depletion of red cells near a solid boundary.)

Our microscopic model will represent blood as a suspension of highly deformable, aggregable red cells within a Newtonian, incompressible medium, the plasma. The cells will be modeled as Newtonian gel-filled biconcave discoid sacculi, with finitely-deforming incompressible elastic membranes. The elastic properties will be obtained experimentally by micropipette aspiration [34]. The attraction laws will be derived from a force balance considering the bridging force caused by macromolecular adsorption (like Velcro), repulsive electrostatic force, and elastic surface tractions. This model will require microscopic experimental data, to be performed at UoP, derived from techniques pioneered by Chien et al. [10, 9] and Healy et al. [15].

The resulting microstructural numerical simulations, as illustrated in Figure 2b, will permit a rigorous account of non-Newtonian blood behavior, and provide an extreme example of the class

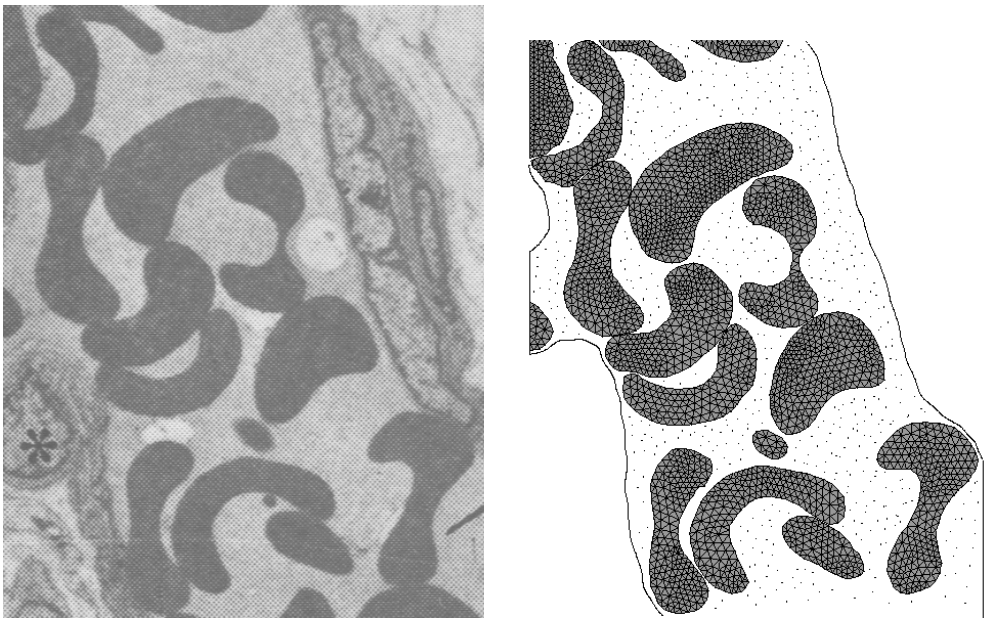

(a) (b)

Figure 2: (a) Electron micrograph of blood flow in a 12μm arteriole, showing interaction of cells with surrounding plasma and walls. (b) Numerical discretization of the fluid-solid mixture, using triangular finite elements for the viscoelastic membrane and interior viscous fluid of the cell and vortex particles (shown subsampled) for the plasma.

of dynamic interface flow simulations we target in this proposal. Simulations like these will be critical for modeling blood flow in regions of hemodynamic devices in which clearances are on the order of tens of cell diameters, as in bearings and rotor tips. Significant blood damage can occur in these "hot spots." Here, macroscopic models fail, and microstructural computations are both necessary and feasible (on state-of-the-art supercomputers) to quantify important effects. An example of such regions are the blade tips of the artifical heart device of Figure 1b, in which the blade-stator clearance is of the order of tens of microns. To resolve the cell–plasma–blade tip interactions, a computational domain covering $200 \times 200 \times 25$ RBCs is necessary. Adequate spatial discretization results in $\mathcal{O}(10^{10})$ grid points, necessitating multi-teraflop computers.

On the other hand, blood flow modeling at the centimeter-scales of the entire device in Figure 1b requires macroscopic models, since explicit resolution of micron-scale microstructure for the whole device is both undesirable and unattainable even on teraflop computers. Moreover, solutions of the macroscopic models determine initial and boundary conditions for microscopic modeling of local hot spots having length scales on the order of cell dimensions. Conversely, macroscopic models (such as mixture theories) depend on information (such as cell-plasma interface forces, partition of elastic and kinetic energies, etc.) that is extremely difficult to obtain experimentally. In this situation, microstructural blood flow models can provide essential dynamic data for development of rational macroscopic models, both for validation as well as suggesting the forms they should take.

Antaki and his collaborator Rajagopal have introduced recently a rigorous macroscopic single-phase continuum model for blood, predicated on a generalized Oldroyd-B fluid that accounts for shear-thinning and shear history [43]. The microstructural sumulations we propose will assist

greatly in validating how well this theory homogenizes microstructural behavior. The group is also developing a macroscopic multiphase blood model, based on Rajagopal's general continuum theory for mixtures of solid particulates in a viscous fluid [19, 20, 33, 18, 17]. One of the difficulties with this approach is that the fluxes coupling the conservation laws for each component must be specified through constitutive relations. The microscopic models we propose will be indispensable for estimation of these quantities.

In addition to blood flows models, the UoP PIs have been developing blood failure and thrombosis models. Antaki, Burgreen, Kameneva, and Rajagopal have presented a particle-specific failure theory that accounts for the exposure history of an individual cell [21, 42, 24]. In addition, significant inroads have been made into macroscopic modeling of agonist-mediated platelet activation and transport through the use of coupled convective-diffusion laws [36, 37, 39, 38]. In both of models, the availability of dynamic microstructural simulation data will prove invaluable in assessing these models.

Significant effort will be made to validate the blood flow models. We will verify that they reproduce important features such as shear thinning, viscoelasticity, cell migration, etc. and will calibrate them with published data and experimental results to be conducted at UoP. This experimental program leverages other funding sources, including the NIH and private foundations. Experimental facilities in our state-of-the-art laboratory at the McGowan Center include laser-based flow visualization [27, 25, 26], flow cytometry and platelet aggregometry [41], blood rheometry [23, 22, 21], and emerging biocompatibility assays for thrombogenesis, coagulation, inflammation, and cellular activation.

In conclusion, the proposed microstructural simulations are necessary for understanding blood flow dynamics and damage in small-clearance regions of hemodynamic devices where macroscopic theory fails; for providing closure data and validation for mixture and other macroscopic flow models; and for developing more rational blood damage models. In all these cases, the number of RBCs contained in sufficiently large representative regions will result in up to the order of billions of grid points, placing the simulations among the largest CFD simulations ever attempted. The need to resolve up to a million deforming cell-plasma interfaces within these flows makes our target simulations more complex and challenging than any flow problem we are aware of, and calls for completely novel algorithms that can scale up to parallel systems with thousands of processors.

3. Parallel Algorithms for Simulation of Flows with Dynamic Interfaces

Simulating fluids flows in domains with deforming, dynamic interfaces is significantly more difficult than doing so on static domains. Yet numerous scientific and industrial flow problems involve the interaction of a fluid with a moving or deforming solid. Many of these problems are sufficiently complex to require highly parallel supercomputers, which is particularly challenging due to the need to scalably manage and evolve the dynamic and irregular data structures. New fundamental research in computer science and computational science will be required to create dynamic-interface flow solvers that scale to the multi-teraflop parallel systems that are beginning to come on line at ASCI and NSF supercomputing centers.

The usual approach to solving flow problems with interfaces that undergo large motion is to relax the idea of conforming to the moving interfaces, and instead employ a regular, fixed grid throughout

the domain. The effect of the interface is then incorporated through other means. Notable examples are the *immersed boundary method* [31], in which interface forces appear as body forces in the fluid; *fictitious domain methods* [14], in which interface conditions are satisfied weakly with Lagrange multipliers; *level set methods* [35], in which the transition in a smooth level set function delineates the interface; and *volume-of-fluid* methods [16], which use a volume-fraction field variable to distinguish different phases. The attraction of all of these methods is that meshes are static and need not evolve, and that fast solvers are possible on the regular grids. The major drawback, on the other hand, is the resulting fixed spatial resolution, particularly disadvantageous if one needs to vary resolution sharply throughout the domain.

A more attractive option we are pursuing employs purely Lagrangian methods in the fluid-solid mixture. Since these methods use a particle-based representation of field variables, and the particles freely convect with the flow, the methods are ideal for complex flow problems that require resolution of moving boundaries and interfaces. Furthermore, since the adoption of a Lagrangian description eliminates the convective term in the momentum equation, we need to solve just Stokes-like problems at each time step, which simplifies the numerics considerably. Finally, Lagrangian methods are naturally adaptive—mesh points convect and diffuse with the flow and naturally concentrate in regions of high vorticity.

However, a number of computational challenges must be overcome to realize the potential of Lagrangian methods. Their great advantage—that the field approximation diffuses and convects with the flow—has also been the major barrier to their widespread use, especially on parallel computers. When the approximation is mesh-based, the mesh distorts with the flow, leading to the need for frequent retriangulation even if the particles remain well-spaced, and an additional need to coarsen/refine the particle set if they do not. Dynamic unstructured mesh algorithms—those needed to coarsen, refine, retriangulate, and repartition the particle system as it evolves over time—are regarded as particularly difficult on highly parallel computers. Indeed, the 1997 Petaflops Algorithms workshop assessed the prospects of a number of high performance computing algorithms scaling to petaflops computers [29] (here we assume the assessmemt is valid for multi-teraflops systems as well). Algorithms were classified according to Class 1 (scalable with appropriate effort), Class 2 (scalable provided significant research challenges are overcome), and Class 3 (possessing major impediments to scalability). The development of dynamic unstructured grid methods (as we propose)—including mesh generation, mesh adaptation and load balancing—were designated Class 2, and thus pose a significant, yet we believe achievable, research challenge. In contrast, static-grid PDE solvers (an example of which is flow solvers on fixed grids) were classified as Class 1.

Lagrangian methods essentially trade off reduced numerical solution difficulties for increased geometric complexity. Given the challenges of parallel dynamic mesh algorithms, the scarcity of mesh-based Lagrangian methods is not surprising. Nevertheless, these methods have many desirable properties, and the recent development by computer science members of our team of a fast scalable parallel Delaunay triangulation algorithm [5, 6] encourages us to pursue a parallel mesh-based Lagrangian flow solver. The Blelloch–Miller et al. algorithm is a new divide-and-conquer parallel algorithm for 2D Delaunay triangulation. The algorithm is theoretically optimal in work, requires polylogarithmic depth, and is very efficient in practice. The main idea is to use a divide-and-conquer technique in which all the work of partitioning is done on the divide step, as compared to the standard divide-and-conquer technique which does the work on the merge

step. The advantage is that our divide step, which is based on a convex hull, is much easier to parallelize than the standard merge step, which involves sequentially stitching the components together. The algorithm is also very efficient: for one million points in the highly non-uniform Kuzmin distribution, for example, it runs in 4 seconds on 64 processors of a Cray T3D, as compared to 110 seconds for the best sequential implementation of Delaunay triangulation we could find. The ideas in principle extend to 3D.

We have incorporated the Blelloch–Miller Delaunay algorithm into an MPI-based 2D parallel dynamic-mesh Lagrangian solver for Navier Stokes flows [4]. Figure 3 shows some snapshots from a simulation (on 16 processors of a Cray T3E) of the benchmark problem of vortex shedding around a cylinder. The basic computational entities are the particles, which move freely with the flow. The particle set is triangulated at each time step, resulting in a completely new mesh and completely new partition of particles onto processors. (Contrast this with most parallel adaptive PDE solvers, which make incremental changes to the mesh, and at much larger intervals.) Particles are removed and added (i.e. the mesh is coarsened/refined) to maintain a well-spaced criterion (leading to quality meshes) and to satisfy error-based local maximum edge lengths. Since the particle description is Lagrangian, newly inserted particles acquire values of velocity and pressure by simple interpolation on the current mesh—there is no need to project the mesh back to that of the previous time step. Figure 3 shows fluid particles and their corresponding meshes at early (while the flow is still steady) and late (after the initiation of vortex shedding) time instants. The number of particles increases by a factor of 3.3 over the course of the simulation, and as can be seen, particles tend to concentrate in regions of high vorticity. Our code is built on top of the MPI-based PETSc library for numerical solution of PDEs [1, 2, 3]. All components of the flow solver are parallel: particle transport, finite element approximation, linear Krylov solver and preconditioner (provided by PETSc), time stepper, and computational geometry algorithms (convex hull, Delaunay, partitioning, coarsening, and refinement). Timings for a problem with a maximum of 127,000 unknowns on a Cray T3E reveal that the computational geometry components consume less than 5% of the time, and the code maintains 93% parallel efficiency (on the maximum number of processors the resolution warrants). Although PETSc was designed for static grids, these results suggest that dynamic grids can be integrated into its framework efficiently (in fact, we intend to make our software library callable from PETSc, thus making it available to thousands of current and future PETSc users).

For this vortex shedding problem, our Lagrangian dynamic mesh method is overkill; a fixed grid (Eulerian) method suffices since the domain is not deforming. However, we chose this example because it is a severe test of Lagrangian mesh-based methods: the mesh becomes hopelessly twisted and distorted in the wake region unless retriangulation is done at each time step. These proof-of-concept results demonstrate that with proper attention to algorithm design, the cost of creating a completely new mesh at *every time step* (including refining, coarsening, triangulating, and partitioning) can be made small relative to the numerical components—contrary to commonly-held beliefs. Thus, the entire dynamic mesh flow solver (including geometric and numerical components) scales well over a reasonable range of granularity.

Fundamental research in both numerical and geometric algorithms is required to scale the solver up to 3D, to multiple solid bodies interacting with a fluid, and to thousands of processors. The numerical algorithms research task will build on our work on variationally-based finite element methods for coupling finitely-deforming elastic solids and viscous incompressible fluids, for which

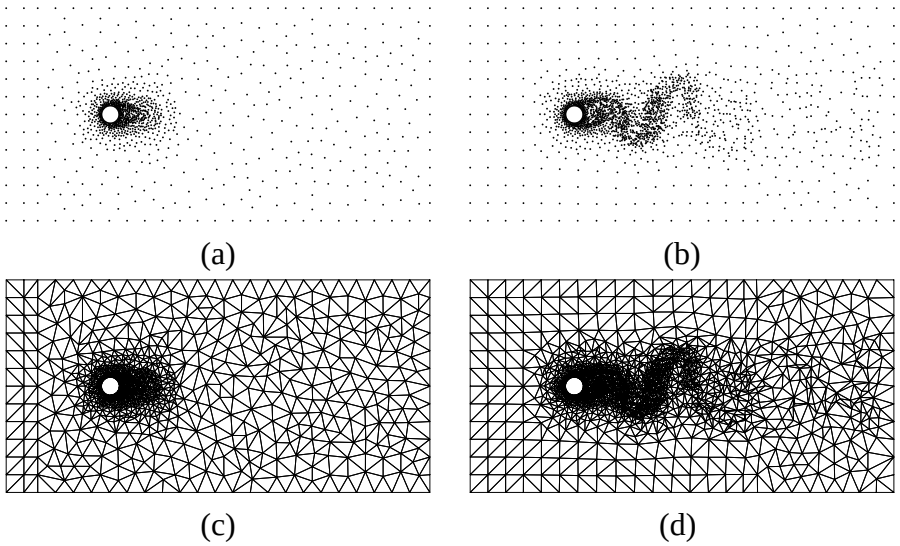

Figure 3: Parallel dynamic mesh Lagrangian flow solver. Particle distribution (a) and (b) and corresponding mesh (c) and (d) at early and late time instants. Computations on Cray T3E at Pittsburgh Supercomputing Center.

we have developed efficient domain-decomposition methods for solution of the resulting coupled discrete systems in 2D [11, 12, 13]. Extending the solvers to 3D and to multiple solid bodies will require significant research. For the geometric algorithms component, the major tasks are to create and implement the 3D parallel computational geometry infrastructure supporting dynamic unstructured meshes: convex hull, Delaunay triangulation, Delaunay coarsening/refinement, and graph partitioning. Again, fundamental algorithms research is required.

References

[1] S. BALAY, W. D. GROPP, L. C. MCINNES, AND B. F. SMITH, *Efficient management of parallelism in object oriented numerical software libraries*, in Modern Software Tools in Scientific Computing, E. Arge, A. M. Bruaset, and H. P. Langtangen, eds., Birkhauser Press, 1997, pp. 163–202.

[2] ——, *PETSc 2.0 users manual*, Tech. Rep. ANL-95/11 - Revision 2.0.24, Argonne National Laboratory, 1999.

[3] ——, *PETSc home page*. `http://www.mcs.anl.gov/petsc`, 1999.

[4] G. BLELLOCH, O. GHATTAS, I. MALČEVIĆ, G. MILLER, AND N. WALKINGTON, *A parallel dynamic mesh-based Lagrangian method for incompressible Navier Stokes flows*, in Fifth US National Congress on Computational Mechanics, Boulder, CO, Aug. 1999.

[5] G. BLELLOCH, J. HARDWICK, G. MILLER, AND D. TALMOR, *Design and implementation of a practical Delaunay algorithm*, Algorithmica, 24 (1999), pp. 243–269.

[6] G. BLELLOCH, G. L. MILLER, AND D. TALMOR, *Developing a practical projection-based parallel Delaunay algorithm*, in Proceedings ACM Symposium on Computational Geometry, 1996, pp. 186–195.

[7] G. BURGREEN, J. ANTAKI, AND B. GRIFFITH, *A design improvement strategy for axial blood pumps using computational fluid dynamics*, ASAIO Journal, 42 (1996), pp. M354–M360.

[8] G. W. BURGREEN AND J. F. ANTAKI, *CFD-based design optimization of a three-dimensional rotary blood pump*, in Proceedings of the Sixth AIAA/NASA/ISSMO Symposium on Multidisciplinary Analysis and Optimization, AIAA, September 1996.

[9] S. CHIEN AND K. M. JAN, *Red cell aggregation by macromolecules: roles of surface adsorption and electrostatic repulsion*, Journal of Supramolecular Structure, 1 (1973), pp. 385–409.

[10] S. CHIEN, K. M. JAN, AND S. USAMI, *Roles of electrostatic repulsion and mechanical shearing in red cell aggregation*, Bibliotheca Anatomica, 13 (1975), pp. 93–94.

[11] O. GHATTAS AND X. LI, *A variational finite element method for nonlinear fluid-solid interaction and its sensitivity analysis*, in Proceedings of the Fifth AIAA/NASA/USAF/ISSMO Symposium on Multidisciplinary Analysis and Optimization, AIAA, September 1994.

[12] ——, *A variational finite element method for stationary nonlinear fluid-solid interaction*, Journal of Computational Physics, 121 (1995), pp. 347–356.

[13] ——, *Domain decomposition methods for sensitivity analysis of a nonlinear aeroelasticity problem*, International Journal of Computational Fluid Dynamics, 11 (1998), pp. 113–130.

[14] R. GLOWINSKI, T. PAN, AND J. PERIAUX, *Fictitious domain methods for incompressible viscous flow around moving rigid bodies*, in The Mathematics of Finite Elements and Applications, John Wiley & Sons, 1996, pp. 155–174.

[15] J. C. HEALY, *Physical method of measuring energy in erythrocyte aggregation*, Biorheology, 11 (1974), pp. 185–190.

[16] C. W. HIRT AND B. D. NICHOLS, *Volume of fluid (VOF) method for the dynamics of free boundaries*, Journal of Computational Physics, 39 (1981), pp. 201–225.

[17] J. HUMPHREY, *Thermal mechanics of mixtures*, Mathematics and Mechancis of Solids, (1998). In press.

[18] J. HUMPHREY AND K. RAJAOPAL, *Finite thermoelasticity of constrained elastomers subject to biaxial loading*, Journal of Elasticity, (1998). In press.

[19] G. JOHNSON, *Application of Continuum Theories to Flows of Fluid–Solid mixtures*, PhD thesis, University of Pittsburgh, Feb. 1991.

[20] G. JOHNSON, M. MASSOUDI, AND K. R. RAJAGOPAL, *A review of interaction mechanisms in fluid–solid flows*, Tech. Rep. DOE/PETC/TR-90/9, U.S. Department of Energy, Pittsburgh Energy Technology Center, 1990.

[21] M. KAMENEVA, J. ANTAKI, H. BOROVETZ, B. GRIFFITH, K. BUTLER, K. YELESWARAPU, M. WATACH, AND R. KORMOS, *On mechanisms of red blood cell (RBC) trauma in assisted circulation*. To appear in *ASAIO Journal* 1995.

[22] M. KAMENEVA, J. ANTAKI, M. WATACH, H. BOROVETZ, AND R. KORMOS, *Heparin effect on red blood cell aggregation*, Biorheology, 31 (1994), pp. 297–304.

[23] M. KAMENEVA AND E. A. J.F. ANTAKI, *Sheep model in the study of hemorheology with assisted circulation: Effect of axial flow blood pump*, ASAIO Journal, (1994).

[24] M. V. KAMENEVA, H. S. BOROVETZ, J. F. ANTAKI, P. LITWAK, W. J. FEDERSPIEL, R. L. KORMOS, AND B. P. GRIFFITH, *Effect of perfluorochemical emulsion on hemorheology and shear induced blood trauma*, in Oxygen Transport to Tissue XVIII, Nemotoa and LaManna, eds., Plenum Press, New York, 1997, pp. 383–390.

[25] J. KERRIGAN, J. ANTAKI, T. MAHER, AND F. SHAFFER, *Quantitative measurement of steady flow patterns in a miniature axial flow blood pump using fluorescent image tracking velocimetry*, in Proceedings of ASME Fluids Engineering Division Summer Meeting, June 19-23 1994, pp. 155–160.

[26] J. KERRIGAN, F. SHAFFER, T. MAHER, T. DENNIS, H. BOROVETZ, AND J. ANTAKI, *Fluorescent image tracking velocimetry (fitv) of the nimbus axipump*, in Proceedings of the American Society for Artificial Internal Organs 39th Annual Meeting, 1993, pp. M639–M643.

[27] J. KERRIGAN, K. YAMAZAKI, R. MEYER, T. MORI, Y. OTAKE, E. OUTA, M. UMEZU, H. BOROVETZ, R. KORMOS, B. GRIFFITH, H. KOYANAGI, AND J. ANTAKI, *Particle tracking flow visualization within an intra-ventricular axial flow assist device*, in Proceedings of 3rd International Society of Rotary Blood Pumps, 1995.

[28] K. MANN, S. DEUTSCH, J. TARBELL, D. GESELOWITZ, G. ROSENBERG, AND W. PIERCE, *An experimental study of Newtonian and non-Newtonian flow dynamics in a ventricular assist device*, J. Biomechanical Eng, 109 (1987), pp. 139–147.

[29] *The 1997 Petaflops Algorithms Workshop summary report*. `http://www.hpcc.gov/cicrd/pca-wg/pal97.html`.

[30] G. PANTALOS AND M. SHARP, *Influence of viscosity and pressure on prosthetic valve regurgitation*, Intl. J. Artificial Organs, 16 (1993), pp. 151–154.

[31] C. PESKIN, *Numerical analysis of blood flow in the heart*, Journal of Computational Physics, 25 (1977), pp. 220–252.

[32] M. POHL, M. WENDT, AND D. LERCHE, *The influence of rheological properties of test fluids on the flow pattern inside the artificial ventricle (type Rostock) and in the aortic outflow tract*, in Biofluid Mechanics: Blood Flow in Large Vessels, D. Liepsch, ed., 1989.

[33] K. RAGAGOPAL, *An introduction to mixture theory*, in Mathematical theory in fluid mechanics, Pitman Research Notes in Mathematics Series, G. Galdi, J. Malek, and J. Necas, eds., vol. 354, Longman, 1996, pp. 86–113.

[34] R. RAND AND A. BURTON, *Mechanical properties of the red cell membrane i and ii*, Biophysical Journal, 4 (1964), pp. 115–135, 303–316.

[35] J. A. SETHIAN, *Theory, algorithms, and applications of level set methods for propagating interfaces*, in Acta Numerica 1996, Cambridge, 1995.

[36] E. SORENSEN, G. BURGREEN, AND J. ANTAKI, *Modelling of platelet-mediated thrombogenesis*, Annals of Biomedical Engineering, (1997), p. 305.

[37] E. SORENSEN, G. BURGREEN, AND J. ANTAKI, *Simulation of platelet-mediated thrombogenesis*, in Proceedings of Biomedical Engineering Society Annual Fall Meeting, October 10-12 1998.

[38] E. SORENSEN, G. BURGREEN, W. R. WAGNER, AND J. F. ANTAKI, *Computational analysis of platelet deposition and thrombus grown on biomaterials*, in Proceedings of American Society of Mechanical Engineering Summer Bioengineering Conference, Big Sky, MT, June 16-20 1999. To appear.

[39] E. N. SORENSEN, G. W. BURGREEN, W. R. WAGNER, AND J. F. ANTAKI, *Computational simulation of platelet deposition and activation*, Annals of Biomedical Engineering, (1999). Accepted for publication.

[40] J. STEINHAUER, *With sudden drop in donors, organ shortage worsens*. New York Times, January 23, 1999.

[41] W. WAGNER, *Biochemical and biophyscial mechanisms of mural thrombosis on natural surfaces*, PhD thesis, University of Texas at Austin, 1991.

[42] K. YELESWARAPU, J. ANTAKI, AND K. RAJAGOPAL, *A mathematical model for shear-induced hemolysis*, Artificial Organs, (1995). To appear.

[43] K. YELESWARAPU, M. KAMENEVA, K. RAJAGOPAL, AND J. ANTAKI, *A generalized Oldroyd-B fluid model for blood i:steady flows.*, Annals of Biomedical Engineering, (1998). To appear.