

DETC2013-13239

**DRAFT: CHRONO: A PARALLEL PHYSICS LIBRARY FOR RIGID-BODY,
FLEXIBLE-BODY, AND FLUID DYNAMICS**

**Toby Heyn
Hammad Mazhar
Arman Pazouki
Daniel Melanz
Andrew Seidl
Justin Madsen**

Andrew Bartholomew

Dan Negrut

Simulation Based Engineering Lab
Department of Mechanical Engineering
University of Wisconsin
Madison, WI, 53706
Email: heyne@wisc.edu

David Lamb

US Army TARDEC
Warren, MI 48397
Email: david.lamb@us.army.mil

Alessandro Tasora

Department of Industrial Engineering
University of Parma
V.G.Usberti 181/A, 43100, Parma, Italy
Email: tasora@ied.unipr.it

ABSTRACT

The last decade witnessed a manifest shift in the microprocessor industry towards chip designs that promote parallel computing. Until recently the privilege of a select group of large research centers, Teraflop computing is becoming a commodity owing to inexpensive GPU cards and multi to many-core x86 processors. This paradigm shift towards large scale parallel computing has been leveraged in Chrono, a freely available C++ multi-physics simulation package. Chrono is made up of a collection of loosely coupled components that facilitate different aspects of multi-physics modeling, simulation, and visualization. This contribution provides an overview of Chrono::Engine, Chrono::Flex, Chrono::Fluid, and Chrono::Render, which are modules that can capitalize on the processing power of hundreds of parallel processors. Problems that can be tackled in Chrono include but are not limited to granular material dynamics, tangled large flexible structures with self contact, particulate flows, and tracked vehicle mobility. The paper presents an overview of each of these modules and illustrates through several examples the potential of this multi-physics library.

INTRODUCTION

Over the last decade there has been a manifest trend in the hardware industry to increase flop/s rates by increasing the number of cores available on a processor. To a very large extent, the tide that has risen sequential computing for several decades is subsiding. The frequency at which cores are operated today has at best plateaued; in many cases, it went down in an attempt to tame power dissipation and overheating. Instruction level parallelism advances that ensured respectable gains through pipelining and out of order execution have largely fulfilled their potential. The bright spot in this evolving hardware landscape has been the growing impetus behind parallel computing hardware. If anything has held steady over the last four decades, it has been the pace at which transistors are packed per unit area in computer chips. This trend allows today chip designs that draw on 22 nm feature length. Intel's road map calls for 14 nm technology in 2014, 10 nm in 2016, 7 nm in 2018, and 5 nm in 2020. In other words, the number of transistors per unit area will continue to double every two years for the current decade. This will translate into immediate access to commodity chips that host multiple

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 24 JAN 2013		2. REPORT TYPE Journal Article		3. DATES COVERED 18-10-2012 to 15-01-2013	
4. TITLE AND SUBTITLE CHRONO: A PARALLEL PHYSICS LIBRARY FOR RIGID-BODY, FLEXIBLE-BODY, AND FLUID DYNAMICS				5a. CONTRACT NUMBER W911NF-12-1-0395	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) David Lamb; Hammad Mazhar; Arman PAzouki; Andrew Seidl; Justin <adsem				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Department of Mechanical Engineering, University of Wisconsin, 1308 W. Dayton St., Madison, WI, 53715-1149				8. PERFORMING ORGANIZATION REPORT NUMBER ; #23631	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) U.S. Army TARDEC, 6501 East Eleven Mile Rd, Warren, MI, 48397-5000				10. SPONSOR/MONITOR'S ACRONYM(S) TARDEC	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S) #23631	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES Proceedings of the ASME 2013 International Design Engineering Technical Conferences & Computers and Information in Engineering Conference IDETC/CIE 2013 August 4-7, 2013, Portland, Oregon, USA					
14. ABSTRACT The last decade witnessed a manifest shift in the microprocessor industry towards chip designs that promote parallel computing. Until recently the privilege of a select group of large research centers, Teraflop computing is becoming a commodity owing to inexpensive GPU cards and multi to many-core x86 processors. This paradigm shift towards large scale parallel computing has been leveraged in Chrono, a freely available C++ multi-physics simulation package. Chrono is made up of a collection of loosely coupled components that facilitate different aspects of multi-physics modeling, simulation, and visualization. This contribution provides an overview of Chrono::Engine, Chrono::Flex, Chrono::Fluid, and Chrono::Render, which are modules that can capitalize on the processing power of hundreds of parallel processors. Problems that can be tackled in Chrono include but are not limited to granular material dynamics, tangled large flexible structures with self contact, particulate flows, and tracked vehicle mobility. The paper presents an overview of each of these modules and illustrates through several examples the potential of this multi-physics library.					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Public Release	18. NUMBER OF PAGES 15	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

compute cores. Given the stagnation in processor operating frequency, an ever growing gap between CPU speed and memory speed, and the waning of instruction level parallelism gains, it becomes apparent that the only way we can continue to enjoy reduced simulation times or ability to rely on refined models is to fall back on parallel computing. There are two major directions in which parallel computing has evolved. The x86 architecture has defined a solution that evolved as a steady and predictable process in which the number of cores on a chip increased over time: AMD produces today 16 core chips, while Intel has 12 core processors. Leveraging these chips requires a low entry point that calls for programming against relatively mature libraries such as OpenMP, MPI, pthreads, cilk, TBB, etc. At memory bandwidths of 75 GB/s and flop/s rates of 0.3 TFlop/s, this has traditionally represented the conservative choice for entering the parallel computing arena. With the release of CUDA 1.0 in 2006, Nvidia offered a second alternative to leveraging parallel computing by programming the ubiquitous video cards available on millions of desktops worldwide. This path to parallel computing is less conventional as it requires one to get familiar with the hardware layout and memory hierarchy associated with GPUs. Today, an Nvidia GPU has close to seven billion transistors. Priced at about \$6000, an Nvidia Kepler K20x delivers a memory bandwidth of 250 GB/s and 1.3 TFlop/s by virtue of using more than 2800 Scalar Processors. It is used side by side with a regular CPU processor, which means that heterogeneous computing, on the CPU and GPU, can lead to substantial speed gains. In this framework, the GPU plays the role of an accelerator by boosting the floating point performance of the CPU. A similar setup is offered now by Intel; i.e., CPU plus accelerator, owing to its recent release of the Knights Corner architecture. A Knights Corner chip has about 60 cores, can deliver up to 320 GB/s and 1 TFlop/s, and uses the x86 instruction set architecture, which translates into an easier adoption path provided one is familiar with OpenMP or MPI.

It becomes apparent that in the immediate future, any increase in simulation speed or model complexity in Computational Science will be fueled by parallel computing. This paper outlines an ongoing effort in the area of computational multi-body dynamics that is motivated by this belief. It starts with a description of a core simulation engine that aims at simulation of many-body dynamics problems with friction and contact. Chrono::Engine handles both rigid and flexible bodies and draws on MPI and/or GPU computing. It then discusses Chrono::Fluid, a GPU parallel simulation tool that aims at fluid-solid interaction problems, which is singled out as an application area that has been largely ignored until recently due to an excessive computational burden incurred by the simulation of systems of practical relevance. Finally, the paper outlines a rendering pipeline that is used for postprocessing of big data. Chrono::Render is capable of using 320 cores and is built around Pixar's RenderMan. All these components combine to produce Chrono, a multi-physics simulation environment that is designed to take advantage of

commodity parallel computing made available by many-core and GPU architectures.

Chrono::Engine

The Chrono::Engine software is a general-purpose simulator for three dimensional multi-body problems [1]. Specifically, the code is designed to support the simulation of very large systems such as those encountered in granular dynamics, where the number of interacting elements can be in the millions. Target applications include tracked vehicles operating on granular terrain [2] or the Mars Rover operating on discrete granular soil. In these applications, it is desirable to model the granular terrain as a collection of many thousands or millions of discrete bodies interacting through contact, impact, and friction. Note that such systems also include mechanisms composed of rigid bodies and mechanical joints. These challenges require an efficient and robust simulation tool, which has been developed in the Chronosimulation package. Chrono::Engine was initially developed leveraging the Differential Variational Inequality (DVI) formulation as an efficient method to deal with problems that encompass many frictional contacts - a typical bottleneck for other types of formulations [3, 4]. This approach enforces non-penetration between rigid bodies through constraints, leading to a cone-constrained quadratic optimization problem which must be solved at each time step [5]. Chrono::Engine has since been extended to support the Discrete Element Method (DEM) formulation for handling the frictional contacts present in granular dynamics problems [6, 7]. This formulation computes contact forces by penalizing small interpenetrations of colliding rigid bodies. Various contact force models can be used depending on the application [8, 9].

The remainder of this section describes the features of Chrono::Engine, starting with the structure of the code. Next, several sub-sections describe the use of GPU computing in the collision detection task, the use of MPI for distributed solution of large systems, and validation work which has been done to assess the accuracy of the simulation tool.

Code Structure of Chrono::Engine

The core of Chrono::Engine is built around the concept of middleware, namely a layer of classes and functions that can be used by third-party developers to create complex mechanical simulation software with little effort [10]. Because of this, graphical user interfaces and end-user tools are not the main focus of the CHRONO::Engine core project; it is assumed that programs with graphical interfaces are built on top of such middleware, or should be considered as additional, or optional, modules.

Given the complexity of the project, approaching half a million lines of code, the software is organized in classes and namespaces as recommended by the Object Oriented Program-

relieves the programmer from the burden of taking care of object's lifetime, given that the relationships between objects can be quite complex as illustrated in Fig.3. A large portion of the C++ classes are available also as Python modules; this enables the use of most simulation features in a scripted environment. Since novice users are more comfortable with Python than with C++, the Python interface proved to be optimal for teaching purposes. The Python interface was produced using the SWIG utility, a process that automatically generates the code for the Python wrapper.

The software architecture has been designed to accommodate an expandable system for handling assets (meshes, textures, CAD models), with multiple paths from pre-processing to post-processing. To this end, we also provide a C# add-in for a parametric 3D CAD package (SolidWorks) that can be used to export models into Chrono::Engine without programming efforts (see Fig.4).




FIGURE 4. Network of asset workflows.

Collision Detection in Chrono::Engine

This section describes the collision detection algorithm designed and implemented for the Chrono::Engine package. Recall that problems of interest are focused on granular dynamics, such as sand flowing inside an hourglass, a rover running over sandy terrain, an excavator/frontloader digging/loading granular mate-




FIGURE 5. Example of AAB generation for 3D cylinder.

rial, etc. In this context, the collision detection task is performed on a rather small collection of rigid and/or deformable bodies of complex geometry (hourglass wall, wheel, track shoe, excavator blade, dipper), and a very large number of bodies (millions to billions) that make up the granular material. On this scale, the collision detection task, particularly when dealing with the granular material, fits perfectly the Single Instruction Multiple Data (SIMD) computation paradigm. Specifically, the same sequence of instructions needs to be applied to every individual body and/or contact in the granular material. Therefore, a collision detection algorithm capable of leveraging the SIMD compute power of commodity Graphics Processing Units (GPUs) was developed and implemented to remove collision detection as the bottleneck in large granular dynamics simulations.

The parallel collision detection algorithm is separated into two phases, broadphase, and narrowphase. The broadphase algorithm quickly determines a list of potential contact pairs while the narrowphase algorithm determines actual contact information. A brief outline of the parallel collision detection algorithm is presented below, for more details see [11, 12, 13].

Broad-Phase Algorithm The Broad-Phase algorithm is used to compute whether two bodies might be in contact at a given time. The purpose of the broad-phase algorithm is not to find actual contact information, but rather to determine if a contact could potentially occur based on the AABs of the bodies involved.

An Axis Aligned Bounding Box (AAB) is a special case of a bounding box that is always aligned to the global reference frame, simplifying collision detection as the bounding box cannot rotate. Because of this, the volume enclosed by the bounding box will always be equal to or greater than the volume of the shape it encloses. AAB generation is simple and can be easily parallellized on a per object basis. See Fig. 5 for an example of AAB computation for a cylinder in 3D space.

Spatial Subdivision Algorithm A high-level overview of the GPU-based collision detection is as follows. The collision detection process starts by identifying the intersections between AABs and bins (see Fig. 6 for a visual representation of a bin).




FIGURE 6. Example of 3D space divided into bins.

The AABB-bin pairs are subsequently sorted by bin id. Next, each bin's starting index is determined so that the bins' AABBs can be traversed sequentially. All AABBs touching a bin are subsequently checked against each other for collisions.

Narrow-Phase Algorithm Once potential contacts have been determined from the broad-phase collision detection stage, the Narrow-Phase algorithm needs to process each possible contact and determine if it actually occurs. To this end an algorithm capable of determining contacts between convex geometries was implemented on the GPU. This algorithm, called "XenoCollide" [14], is based upon Minkowski Portal Refinement (MPR) [15].

Using MPI for distributed Chrono

Chrono has been further extended to allow the use of CPU parallelism for certain problems. To efficiently simulate large systems, a domain decomposition approach has been developed to allow the use of many-core compute clusters. In this approach, we divide the simulation domain into a number of sub-domains in a lattice structure. Each sub-domain manages the simulation of all bodies contained therein. Note that bodies may span the boundary between adjacent sub-domains. In this case, the body is considered shared and its dynamics may be influenced by the participating sub-domains. The implementation leverages the MPI standard [16] to implement the necessary communication and synchronization between sub-domains.

This approach enables the simulation of large systems in two ways. First, it relies on the power of parallel computing since one compute core can be assigned to each MPI process (and therefore to each sub-domain). These processes can execute in parallel, constrained only by the required communication and synchronization. Second, it allows access to the larger memory pool available on distributed memory architectures. Whereas a single node or GPU card may have about 6 GB of memory, a distributed memory cluster may have on the order of 1TB of memory, enabling the simulation of vastly larger problems.

Note that the domain decomposition approach currently uses

the discrete element method to resolve friction and contact forces between elements in the system. The approach also supports constraints between bodies in the simulation by considering an assembly of constrained rigid bodies as a unit which must always be kept together. Therefore, if any body in a chain of constrained bodies is contained in a given sub-domain, all bodies in the chain are considered by that sub-domain and used to correctly solve the constraint equations.

Sub-division and Set-up A pre-processing step is used to discretize the simulation domain into a specified number of sub-domains, set up the communication conduits between processes, and initialize the sub-domains as appropriate. The sub-division is based on a cubic lattice with support for arbitrary sized divisions. The sub-domain boundaries are aligned with the global cartesian coordinate system, and their locations are user-specified. Separate MPI processes are mapped to each sub-domain. Note that at this time, the sub-division is static and does not change during the simulation. Therefore, the user should be careful to set up the discretization to maintain the best possible load balancing.


In terms of communication, each sub-domain in the grid can communicate with all other sub-domains. These communication pathways are set up and initialized during the pre-processing step and persist throughout the simulation.

Note that this implementation relies heavily on inheritance and the class-based structure of Chrono. For example, `ChSystem` is extended to `ChSystemMPI` by including the code to perform communication and synchronize the sub-domains.

Simulation and Communication Each sub-domain is now represented by a `ChSystemMPI` object and an associated MPI process. For example, assume a simulation is discretized into a set S of m sub-domains. In this case, let $S = \{A, B, C, D\}$ and $m = 4$, and map an MPI rank to each sub-domain so that A is mapped to MPI rank 0, $B \rightarrow 1$, $C \rightarrow 2$, and $D \rightarrow 3$. Each sub-domain maintains at all times $m + 1$ lists of objects. The first list contains all objects which are even partially contained in the associated sub-domain. These are the objects which must be considered when computing contact forces, for example. The next m lists contain bodies which are shared with other sub-domains.

In our example, sub-domain B maintains the lists B_O , B_A , B_B , B_C , and B_D . B_O is the list of all objects that intersect (touch) sub-domain B , while B_A is the list of objects which are in sub-domain A and B . Note that sub-domain A has a list A_B which should contain the same objects as B_A . Further, list B_B is not used but is created for the sake of generality. All lists are maintained in order sorted by object ID number (see Fig.7).

The sub-domains are now ready for time-stepping. Each sub-domain X performs collision detection among all objects in



A	B	C	D
A _O : 1,2,3	B _O : 3,4,7	C _O : 2,3,5,6	D _O : 3,6,7
A _A :	B _A : 3	C _A : 2,3	D _A : 3
A _B : 3	B _B :	C _B : 3	D _B : 3,7
A _C : 2,3	B _C : 3	C _C :	D _C : 3,6
A _D : 3	B _D : 3,7	C _D : 3,6	D _D :

FIGURE 7. (TOP) Sample 2D simulation domain with four sub-domains and seven objects. (BOTTOM) Corresponding object lists for each sub-domain.

list X_O and computes the associated collision forces based on the DEM force model. Then, sub-domain X computes the net force on each object in list X_O , taking into account the contact forces, gravitational forces, and applied forces.

Next, mid-step communication occurs. Sub-domain X should send to each sub-domain Y the net force on each body in list X_Y . Similarly, X should receive from each Y the net force on each body in list X_Y . Finally, X should compute the total force on each body in list X_O . Note that X may receive force contributions for a given body from any or all of the other sub-domains in the system.

At this point, each sub-domain X has the true net force on each body in its list X_O . Each sub-domain can advance the state of its bodies in time by one time step by computing the new accelerations, velocities, and positions of all objects in the sub-domain given their mass/inertia properties and the set of applied forces. We perform an end-of-step communication to synchronize object states among sub-domains. All sub-domains which share a given body should compute its new state identically, but due to the potential for round-off error we synchronize the state from the master sub-domain (where the center-of-mass is located) to all others. The final stage is to process the $m + 1$

lists in each sub-domain, as objects may enter or leave a given sub-domain or be shared between a different set of sub-domains, necessitating updates of the contents of the lists.

Example Simulation In this example we simulate a Mars Rover type wheeled vehicle operating on granular terrain. The vehicle is composed of a chassis and six wheels connected via revolute joints. The wheels are driven with a constant angular velocity of π rad/sec. The granular terrain is composed of 2,016,000 spherical particles. The simulation is divided into 64 sub-domains and uses a time step of 10^{-5} sec. A snapshot from the simulation can be seen in Fig.8. In the figure, note that the wheels of the rover are checkered blue and white. This signifies that the master copy of the rover assembly is in the blue sub-domain and the rover spans into adjacent sub-domains.




FIGURE 8. Snapshot of Mars Rover simulation with 2,016,000 terrain particles using 64 sub-domains. Bodies are colored by sub-domain, with shared bodies (those which span sub-domain boundaries) colored white.

Validation and Demonstration of Technology

This section describes a validation effort in which experimental results were compared to simulation results obtained from Chrono::Engine. To this end, a test rig was designed and fabricated to measure the rate at which granular material flowed out of a slit due to gravity. Chrono::Engine was used to set up a corresponding simulation to match the experimental results. For more detail, see [17].

Experimental Model The experimental set-up consisted of a fixed base, a movable wall (angled at 45°), a translational stage, a linear actuator, and a scale (see schematic in Fig. 9). The linear actuator was capable of quickly opening a precise gap, out of which the granular material would flow due to gravity. The scale recorded the mass of collected granular material as a function of time. The granular material consisted of approximately 40,000 uniform glass disruptor beads with diameter of 500 microns. Experiments were performed for gap sizes of 1.5 mm, 2 mm, 2.5 mm, and 3 mm. At least 5 experiments were performed for each gap size.




FIGURE 9. Schematic of validation experiment. A linear actuator and translational stage moved the left angled side a fixed amount, opening a precise gap from which the particles flowed. The mass flow rate was measured by the scale. Schematic not to scale.

Simulation Model Chrono::Engine was used to build a model representing the experimental set up described above. In the model, the trough was represented by four rectangular boxes of finite dimensions. The motion of the box representing the angled side was captured from the data sheet of the translational stage. The granular material was modeled as perfect, identical spheres with the same mass and coefficient of friction.

The load cell measured the outflow through the gap. In the simulation, the scale was modeled by counting the number of spheres below a certain height. The number of spheres multiplied by the mass and gravity yielded the weight which was compared with experimental results. A plane was used to contain the spheres after they had been counted.

In order to save computational time, the simulation was split into two parts: one representing the process of filling the trough

and the other the opening and measuring process. In this way, the trough was filled with randomly positioned spheres which were allowed to settle. Once the kinetic energy of the system was below 0.001 Joules and had reached a relatively constant value, the x-, y-, and z-position of each sphere was saved to a file.

The same initial conditions from the settling simulation were used to perform all of the necessary simulations. At the beginning of each simulation the position data set of the spheres was loaded into the model and the spheres were created at the same positions they appeared in the filling process. The motion was applied to the translating side to achieve the desired gap size, and the material began to flow.

The simulations setup consisted of 39,000 rigid body spheres with a radius of 2.5×10^{-4} m and a mass of 1.631×10^{-7} kg. The following parameters were set for this simulation. A time step of 10^{-4} [s] with 500 CCP iterations, and a tolerance of 10^{-7} for the maximum velocity correction. Simulations were generally run for 8 seconds. SI units were used for all parameters.

Procedure Used to Select the Friction Coefficient

The friction coefficient of a certain material is not a constant value. It can depend on various environmental influences such as humidity, surface quality, temperature etc. The friction coefficient of glass was an unknown in the validation process and needed to be determined before further observations could be done. To achieve this, one experiment at a gap size of 1.5 mm was performed and multiple simulations with the same setup and different friction coefficients were performed. The simulation results were compared to the experimental test results to determine which friction coefficient resulted in the best match. It was determined that $\mu = 0.15$ most closely matched the experimental results. This value was used for all subsequent simulations.

Results The weight of the collected granular material is plotted versus time in Fig. 10 through Fig. 13. The mass flow rate is proportional to the slope. For example, Fig. 10 shows the results with a gap size of 3mm and 5 experimental runs (dashed lines represent experimental data, while the solid red line represents simulation data). The average slope of the experimental runs for this test was $1.41\text{E-}2$ [N/s] and the slope of the simulation was $1.40\text{E-}2$ [N/s].

The experimental flow rate is based off of the average of several experimental runs and the uncertainty of the rates is determined by a Student's T distribution with 95% confidence. Based on an error of less than 2%, the simulated results match the experimental results well. Despite this, there are several factors that could be improved upon. Namely, the design of the rig did not exactly match the simulated trough. The rig was machined from aluminum, whereas the trough in the simulation had the same coefficient of friction as the granular material. Likewise, the sim-




FIGURE 10. Weight vs time for a gap size of 3 mm.




FIGURE 11. Weight vs time for a gap size of 2.5 mm.




FIGURE 12. Weight vs time for a gap size of 2 mm.




FIGURE 13. Weight vs time for a gap size of 1.5 mm. This was the test case that was used for calibration.

Chrono::Flex

The Chrono::Flex software is a general-purpose simulator for three dimensional flexible multi-body problems and provides a suite of flexible body support. The features included in this module are multiple element types, the ability to connect these elements with a variety of bilateral constraints, multiple solvers, and contact with friction. Additionally, Chrono::Flex leverages the GPU to accelerate solution of large problems.

Element Types

Chrono::Flex includes two element types implemented using the Absolute Nodal Coordinate Formulation (ANCF) [18, 19]. The gradient-deficient beam element and the gradient-deficient plate element are described below.

Gradient-Deficient Beam Elements This implementation uses gradient deficient ANCF beam elements to model slender beams, examples of which are shown in Figure 14. These are two node elements with one position vector and only one gradient vector used as nodal coordinates. Each node thus has 6 coordinates: three components of the global position vector of the node and three components of the position vector gradient at the node. This formulation displays no shear locking problems for thin and stiff beams and is computationally more efficient compared to the original ANCF due to the reduced number of nodal coordinates [20]. The gradient deficient ANCF beam element does not describe a rotation of the beam about its own axis so the torsional effects cannot be modeled.

Gradient-Deficient Plate Elements Much like beams, numerical difficulties are encountered in the fully parameterized plate element when the system has very thin and stiff components [21]. The high frequencies that are induced along the thin direction of the element require an extremely small time step, re-

ulation was conducted in essentially vacuum, ignoring any aerodynamic forces, while the experiment was performed at ambient conditions. Lastly, the effects of humidity were not taken into account by the simulation.




FIGURE 14. Two models with friction and contact using Chrono::Flex beam elements: a ball sitting on grass-like beams and a ball hitting a net.

sulting in longer simulation times. In the case where the aspect ratio (length divided by thickness) of the element is high, plane stress assumptions can be made that allow a reduced- order element to be accurate. Specifically, Kirchhoff's plate theory, which does not account for shear deformation, is used and results in an element with 36 degrees of freedom, or nodal coordinates, shown in Figure 15.




FIGURE 15. Two models with friction and contact using Chrono::Flex plate elements: a cloth hanging on a sphere and a closed contour shaped like a tire.

Kinematic Constraints

Several types of mechanical joints are modeled in Chrono::Flex. A spherical joint [22] between two nodes of any two bodies will require the position vector of each node to be identical. A revolute joint will have two additional constraints to the spherical joint constraints. In this case, the gradient vectors of the two nodes will remain in a plane perpendicular to the axis of revolute joint. There are also additional constraints due to the element connectivity in each beam. The element connectivity can be modeled as a fixed joint between the nodes. Here the common node between two elements is treated as two different nodes attached to each other through the fixed joint. This fixed joint requires all the nodal coordinates of the two nodes be iden-

tical. The generalized coordinates of the system change in time under the effect of applied forces such that these constraint equations are satisfied at all times. The time evolution of the system is governed by the Lagrange multiplier form of the constrained equations of motion.

Solvers

The equations shown in Figure 16 form a system of index-3 Differential Algebraic Equations (DAEs). Although several low order numerical integration schemes have been effectively used to solve index-3 DAEs, Chrono::Flex utilizes the Newmark integration scheme [23]. Originally used in the structural dynamics community for the numerical integration of a linear set of second order ODEs, it was adapted for the discretization of DAEs. This implicit solver was proved to have convergence of order 1 or 2, depending on the choice of parameters γ and β .

$$\begin{array}{l}
 \text{Generalized Accelerations} \rightarrow \mathbf{M}\ddot{\mathbf{q}} + \underbrace{\Phi_{\mathbf{q}}^T(\mathbf{q},t)\lambda}_{\text{Reaction Force}} + \underbrace{\mathbf{Q}_{\text{int}}(\mathbf{q})}_{\text{Elastic Force}} = \underbrace{\mathbf{Q}_{\text{ext}}(\dot{\mathbf{q}},\mathbf{q},t)}_{\text{Applied Force}} \\
 \text{Generalized Mass Matrix} \rightarrow \mathbf{M} \\
 \text{Force Balance Equations} \rightarrow \mathbf{M}\ddot{\mathbf{q}} + \Phi_{\mathbf{q}}^T(\mathbf{q},t)\lambda + \mathbf{Q}_{\text{int}}(\mathbf{q}) = \mathbf{Q}_{\text{ext}}(\dot{\mathbf{q}},\mathbf{q},t) \\
 \text{Generalized Positions} \rightarrow \mathbf{q} \\
 \text{Holonomic Kinematic Constraints} \rightarrow \Phi(\mathbf{q},t) = 0
 \end{array}$$

FIGURE 16. The equations of motion for Chrono::Flex.

At each time step, the numerical solution commences by solving the nonlinear set of equations shown in Figure 17. The numerical solution of the nonlinear algebraic system falls back on a Newton-type iterative algorithm that requires the computation of its sensitivity matrix. Advancing the numerical solution in time draws on three loops: the outer-most loop marches forward in time, while at each time step the second loop solves the algebraic discretization problem in Figure 17. Each iteration in this second loop launches a third loop whose role is that of producing a vector of corrections for the acceleration and Lagrange multipliers. The corrections are computed using the BiCGStab iterative solver [24], which also provides for a matrix-free solution. A serial solver was implemented using the Armadillo Matrix Algebra Library [25] and a GPU parallel solver was implemented using CUSP [26], a linear algebra library built on top of CUDA. Chrono::Flex was validated in [27] as well as in [28] against the commercial code ADAMS [29], and the nonlinear finite element analysis code ABAQUS [30].

Chrono::Fluid

The Chrono::Fluid component aims at leveraging GPU computing to efficiently simulate fluid dynamics and fluid-solid interaction problems. Fluid-Solid Interaction (FSI) covers a wide range of applications, from blood and polymer flow to




FIGURE 17. The discretized equations of motion for Chrono::Flex (fully implicit).

tanker trucks and ships. Simulation of the FSI problem requires two components: Fluid and Solid simulations. Simulation of the Solid phase, either rigid or flexible, in an HPC fashion, is described in previous sections. To leverage the existing solid phase simulation, the fluid flow simulation should satisfy some conditions, introduced by the aforementioned target problems. First, the fluid flow may experience large domain deformation due to the motion of the solid phase. Second, the two phases should be coupled via an accurate algorithm. Third, target problems may experience free surface as well as internal flow. Finally, the whole simulation should be capable of an HPC implementation to maintain the scalability of the code.

Fluid flow can be simulated in either an Eulerian or a Lagrangian framework. Provided that the interfacial forces are captured thoroughly, the Lagrangian framework is capable of tracking the domain deformation introduced by the motion of the solid phase at almost no extra cost. Smoothed Particle Hydrodynamics (SPH) [31, 32, 33], its modifications [34, 35], and variations [36] have been widely used for the simulation of the fluid domain in a Lagrangian framework. The main evolution equations of the fluid flow using SPH are expressed as

$$\frac{d\rho_a}{dt} = \rho_a \sum_b \frac{m_b}{\rho_b} (v_a - v_b) \cdot \nabla_a W_{ab} \quad (1)$$

$$\frac{dv_a}{dt} = - \sum_b m_b \left(\left(\frac{p_b}{\rho_a^2} + \frac{p_a}{\rho_b^2} \right) \nabla_a W_{ab} - \frac{(\mu_a + \mu_b) r_{ab} \cdot \nabla_a W_{ab}}{\bar{\rho}_{ab}^2 (r_{ab}^2 + \epsilon \bar{h}_{ab}^2)} v_{ab} \right) + f_a \quad (2)$$

which are solved in conjunction with

$$dx/dt = v \quad (3)$$

to update the fluid properties. In Eqs. (1) to (3), ρ , v , and p are local fluid density, velocity, and pressure, respectively, m

is the representative fluid mass assigned to the SPH marker, W is a kernel function which smooths out the local fluid properties within a resolution length $l = \kappa h$, and r_{ab} is the distance between two fluid markers denoted by a and b . Fluid flow evolution equations, defined by Eqs. (1) to (3), are solved explicitly, where pressure is related to density via an appropriate state equation to maintain the compressibility below 1%. To increase the accuracy and stability of the simulation, an XSPH modification [34] and Shephard filtering [37] were applied.

FSI with Smoothed Particle Hydrodynamics: A Quick Overview

A proper choice of fluid-solid coupling should satisfy the no-slip and impenetrability conditions on the surface of the solid obstacles. By attaching Boundary Condition Enforcing (BCE) markers on the surface of the solid objects, the local relative velocity, i.e. at the markers location, of the two phases will be zero (Fig. 18). The position and velocity of the BCE markers are updated according the motion of the solid phase, which results in the propagation of the solid motion to the fluid domain. On the other hand, the interaction forces on the BCE markers are used to calculate the total force and torque exerted by the fluid on the solid object.




FIGURE 18. Coupling of the fluid and solid phases. BCE and fluid markers are represented by black and white circles, respectively.

FSI with Smoothed Particle Hydrodynamics: Proximity Computation

The overall simulation of the FSI framework is performed in parallel, where each thread handles the force calculation of a fluid or BCE marker first, and a rigid body later. Next, the parallel threads perform the kinematics update of

the fluid markers, rigid bodies, and BCE markers, respectively. An essential part of the force calculation stage is the proximity computation, which will be explained briefly herein. Proximity computation used in our work leverages the algorithm provided in CUDA SDK [38], where the computation domain is divided into bins whose sizes are the same as the resolution length of the SPH kernel function. A hash value is assigned to each marker based on its location with respect to the bins. Markers are sorted based on their hash value. The sorted properties are stored in independent arrays to improve the memory access and cache coherency. To compute the forces on a marker, the lists of the possible interacting markers inside its bin and all 26 neighbor bins are called. The hash values of the bins are used to access the relevant segments of the sorted data.

Validation and Demonstration of Technology

The aforementioned FSI simulation engine was used to validate the lateral migration of cylindrical particles in plane Poiseuille flow, spherical particles in pipe flow, and particle distribution in Poiseuille flow of suspension [39,40]. Due to the scalability of Chrono::Fluid in both fluid and solid phases, increasing the number of rigid bodies, which translates into decreasing the number of fluid particles, does not affect the simulation time. Therefore, the simulation of a highly dense suspension is possible. Figure 19 shows the result of the simulation of the flow of suspension including 1,500 particles through a channel. A similar scenario with 13,000 particles in suspension was successfully simulated in Chrono::Fluid.




FIGURE 19. Simulation of rigid bodies inside a fluid flow: Rigid ellipsoids with their BCE markers are shown in the left image while the fluid's velocity contours and rigid ellipsoids at the mid-section of the channel are shown in the right image.

Chrono::Render

Chrono::Render is a software package that enables simple, streamlined, and fast visualization of arbitrary data using Pixar's RenderMan [41]. Specifically, Chrono::Render contains a hybrid of processing binaries and Python scripting modules that seek to abstract away the complexities of rendering with RenderMan. Additionally, Chrono::Render is targeted for providing rendering as an automated post-processing step in a remote simulation pipeline, hence it is controlled via a succinct XML specification for "gluing" together rendering with arbitrary processes. As seen in Figure 20, Chrono::Render combines simulation data, XML describing how to use the data, and optional user-defined Python scripts into a complex, visually-rich scene to be rendered by RenderMan.




FIGURE 20. Chrono::Render architecture.

On the Choice of RenderMan

Using RenderMan for rendering is motivated by the scope of arbitrary data sets and the potentially immense scene complexity that results from big data; REYES, the underlying architecture for RenderMan is ideally suited for this task. REYES works by dividing each surface in the scene into a grid of micropolygons and shades at the grid vertices [42] (see Figure 21).

This results in tractable rendering for complex scenes because: (a) only a small portion of the scene needs to be in memory at any given time; (b) grid-based computation leads to optimal memory access patterns; (c) non-visible objects need not be loaded into memory; (d) fully-rendered objects can be removed from memory; and (e) objects are tessellated according to size on the screen; less complex geometry is dynamically loaded whenever possible.




FIGURE 21. An overview of the REYES Pipeline.

REYES is perfectly suited for parallel processing since it scales linearly with the number of cores. Considering that REYES needs only a handful of relevant scene elements at a time, this data can be parsed into low-memory buckets and distributed amongst cores for parallel rendering; thus REYES' low memory-footprint and efficient concurrent resource usage for the complex scenes makes it a great renderer for a distributed-computing platform.

Accessibility of High-Quality Graphics

Although REYES can manage the issue of scene complexity, leveraging this power is difficult without computer graphics expertise. The guiding principle of Chrono::Render is to make high-quality rendering available to researchers, most of whom don't have the background or bandwidth to spend time learning how to use complex graphics applications or make sense of REYES' intricacies. Consequently, Chrono::Render encapsulates into the XML specification the complicated steps needed to make interesting visual effects, such as multipass rendering. The user must only instance the correct XML components to achieve high-quality renders. The program flow of Chrono::Render is shown in Figure 22.

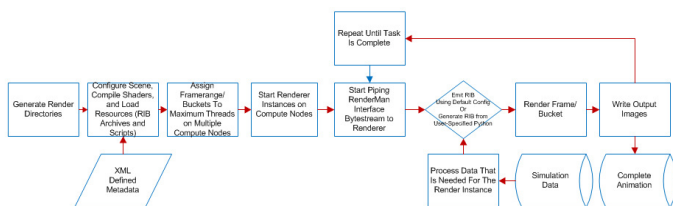


FIGURE 22. Chrono::Render execution workflow.

The XML specification allows for the concise expression of salient features and scene objects. For example, the snippet in Figure 23 illustrates the XML file that translates a single line

from a CSV data file into a RenderMan sphere using two shaders.




FIGURE 23. Simple XML for a sphere with a Surface and Displacement shader.

Although simple, the render is visually rich. This description is often enough to visualize most generic data, but it cannot handle all arbitrary visualizations, so in order to maintain generality we make use of Python scripts and wrappers to enable simplified procedural RenderMan Interface Bytestream generation. Any XML element can be scripted such that at runtime, the script output will be piped into the same rendering context. This makes it possible to perform processing for specialized data as well as modularize the rendering of specific effects. Obviously this adds more complexity for defining the scene, but Chrono::Render provides Python modules with methods and classes intended to ease this programming as much as possible. Additionally, most of the Chrono::Render Python modules wrap C++ functions and classes with the purpose of exploiting speed while still making use of the syntactical/type-free simplicity of Python. Figure 24 gives an example of combining XML with Python scripts to achieve a more complicated render.

Other Capabilities

Beyond interpreting parameters and data into RenderMan calls, Chrono::Render provides tools for bootstrapping rendering projects. Chrono::Render can: (a) construct directory structures for localizing and managing scene resources; (b) automate distribution of rendering across a multi-node network; (c) convert common graphics file formats into RenderMan file formats such as Wavefront Objs and Mtls to RenderMan RIBs and Shaders; (d) generate XML for automatically adding parameters to the scene description for describing advanced visual effects such as subsurface scattering, ambient occlusion, reflections, etc.; (e) mesh point-clouds, particularly useful for particle-based fluid

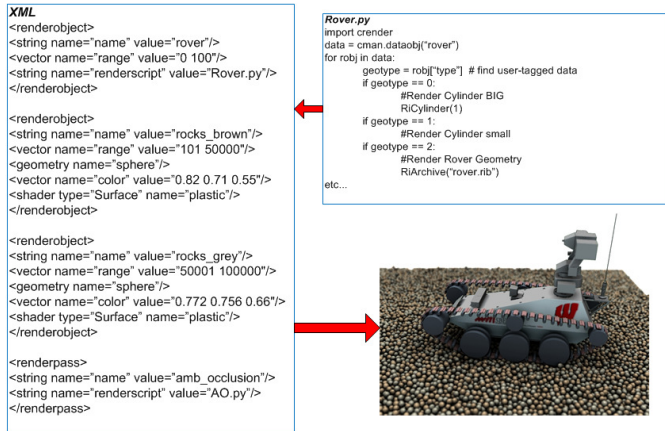


FIGURE 24. General purpose rendering with Chrono::Render. The Rover body contains multiple shape descriptions of which are generated from a Python script. Data is tagged with a name which can be later be accessed using some of Chrono::Render's Python functionality.

simulations; and (f) dump the generated RenderMan calls to disk for reuse.

Chrono:Render is currently available for free download as a pre-built binary for Linux. Members of the Wisconsin Applied Computing Center can use this capability remotely as a service by leveraging 320 AMD cores on which Chrono::Render is currently deployed.

Conclusions and Future Work

The Chronosimulation package is composed of a collection of components designed to perform multi-physics simulations leveraging emerging high-performance computing hardware. Chrono::Engine provides support for rigid body dynamics, focusing on large granular dynamics problems, Chrono::Flex enables simulation of flexible beam and plate elements interacting through contact and bilateral constraints, while Chrono::Fluid allows the simulation of fluid flows and fluid-solid interaction problems. Finally, Chrono::Render provides high-quality visualization of arbitrary simulation data from the other Chrono-components. These components have been designed to leverage high-performance computing hardware whenever possible. Chrono::Engine supports CPU parallelism through a domain-decomposition approach, while Chrono::Engine, Chrono::Flex, and Chrono::Fluid all support GPU parallelism to further improve simulation performance.

While these components provide useful simulation capabilities on their own, ongoing work seeks to further integrate the various Chronocomponents.

ChronoAvailability

Major releases of the Chrono::Engine software are available from the Chrono::Engine website at <http://chronoengine.info>. Chronoin its entirety can be downloaded from <http://sbel.wisc.edu/chrono>. The latter site also displays the nightly build status for various platforms and unit testing results.

ACKNOWLEDGMENT

Financial support for the Wisconsin authors was provided in part by the National Science Foundation Award 0840442 and Army Research Office W911NF-12-1-0395. Financial support for A.Tasora was provided in part by the Italian Ministry of Education under the PRIN grant 2007Z7K4ZB. We thank NVIDIA and AMD for sponsoring our research programs in the area of high-performance computing.

REFERENCES

- [1] Tasora, A., and Anitescu, M., 2011. "A matrix-free cone complementarity approach for solving large-scale, nonsmooth, rigid body dynamics". *Computer Methods in Applied Mechanics and Engineering*, **200**(5-8), pp. 439–453.
- [2] Heyn, T., 2009. "Simulation of Tracked Vehicles on Granular Terrain Leveraging GPU Computing". M.S. thesis, Department of Mechanical Engineering, University of Wisconsin–Madison, http://sbel.wisc.edu/documents/TobyHeynThesis_final.pdf.
- [3] Anitescu, M., and Tasora, A., 2010. "An iterative approach for cone complementarity problems for nonsmooth dynamics". *Computational Optimization and Applications*, **47**(2), pp. 207–235.
- [4] Tasora, A., and Anitescu, M., 2010. "A convex complementarity approach for simulating large granular flows". *Journal of Computational and Nonlinear Dynamics*, **5**(3), pp. 1–10.
- [5] Negrut, D., Tasora, A., Mazhar, H., Heyn, T., and Hahn, P., 2012. "Leveraging parallel computing in multibody dynamics". *Multibody System Dynamics*, **27**, pp. 95–117. 10.1007/s11044-011-9262-y.
- [6] Cundall, P., 1971. "A computer model for simulating progressive large-scale movements in block rock mechanics". In *Proceedings of the International Symposium on Rock Mechanics*. Nancy, France.
- [7] Cundall, P., and Strack, O., 1979. "A discrete element model for granular assemblies". *Geotechnique*, **29**, pp. 47–65.
- [8] Mindlin, R., and Deresiewicz, H., 1953. "Elastic spheres in contact under varying oblique forces". *Journal of Applied Mechanics*, **20**, pp. 327–344.

- [9] Kruggel-Emden, H., Simsek, E., Rickelt, S., Wirtz, S., and Scherer, V., 2007. "Review and extension of normal force models for the discrete element method". *Powder Technology*, **171**, pp. 157–173.
- [10] Tasora, A., Righettini, P., and Silvestri, M., 2007. "Architecture of the chrono::engine physics simulation middleware". In Proceedings of ECCOMAS 2007 Multibody Conference.
- [11] Mazhar, H., Heyn, T., and Negrut, D., 2011. "A scalable parallel method for large collision detection problems". *Multibody System Dynamics*, **26**, pp. 37–55. 10.1007/s11044-011-9246-y.
- [12] Pazouki, A., Mazhar, H., and Negrut, D., 2012. "Parallel collision detection of ellipsoids with applications in large scale multibody dynamics". *Mathematics and Computers in Simulation*, **82**(5), pp. 879 – 894.
- [13] Pazouki, A., Mazhar, H., and Negrut, D., 2010. "Parallel ellipsoid collision detection with application in contact dynamics-DETC2010-29073". In Proceedings to the 30th Computers and Information in Engineering Conference, S. Fukuda and J. G. Michopoulos, eds., ASME International Design Engineering Technical Conferences (IDETC) and Computers and Information in Engineering Conference (CIE).
- [14] Snethen, G., 2007. Xenocollide website, Sept. <http://www.xenocollide.com>.
- [15] Snethen, G., 2008. "Xenocollide: Complex collision made simple". In *Game Programming Gems 7*, S. Jacobs, ed. Charles River Media, pp. 165–178.
- [16] Gropp, W., Lusk, E., and Skjellum, A., 1999. *Using MPI: Portable Parallel Programming with the Message-Passing Interface, Second Edition*. MIT Press.
- [17] Melanz, D., Tupy, M., Smith, B., Turner, K., and Negrut, D., 2010. "On the validation of a differential variational inequality approach for the dynamics of granular material-DETC2010-28804". In Proceedings to the 30th Computers and Information in Engineering Conference, S. Fukuda and J. G. Michopoulos, eds., ASME International Design Engineering Technical Conferences (IDETC) and Computers and Information in Engineering Conference (CIE).
- [18] Berzeri, M., Campanelli, M., and Shabana, A. A., 2001. "Definition of the elastic forces in the finite-element absolute nodal coordinate formulation and the floating frame of reference formulation". *Multibody System Dynamics*, **5**, pp. 21–54.
- [19] von Dombrowski, S., 2002. "Analysis of large flexible body deformation in multibody systems using absolute coordinates". *Multibody System Dynamics*, **8**, pp. 409–432. 10.1023/A:1021158911536.
- [20] Gerstmayr, J., and Shabana, A., 2006. "Analysis of thin beams and cables using the absolute nodal co-ordinate formulation". *Nonlinear Dynamics*, **45**(1), pp. 109–130.
- [21] Dufva, K., and Shabana, A., 2005. "Analysis of thin plate structures using the absolute nodal coordinate formulation". *Proceedings of the Institution of Mechanical Engineers, Part K: Journal of Multi-body Dynamics*, **219**(4), pp. 345–355.
- [22] Shabana, A. A., 2005. *Dynamics of Multibody Systems*, third ed. Cambridge University Press.
- [23] Hussein, B., Negrut, D., and Shabana, A., 2008. "Implicit and explicit integration in the solution of the absolute nodal coordinate differential/algebraic equations". *Nonlinear Dynamics*, **54**(4), pp. 283–296.
- [24] Yang, L., and Brent, R., 2002. "The improved bicgstab method for large and sparse unsymmetric linear systems on parallel distributed memory architectures". In Algorithms and Architectures for Parallel Processing, 2002. Proceedings. Fifth International Conference on, IEEE, pp. 324–328.
- [25] Sanderson, C., 2010. Armadillo: An open source c++ linear algebra library for fast prototyping and computationally intensive experiments. Tech. rep., Technical report, NICTA.
- [26] Bell, N., and Garland, M., 2012. Cusp: Generic parallel algorithms for sparse matrix and graph computations. Version 0.3.0.
- [27] Khude, N., Melanz, D., Stanciulescu, I., and Negrut, D., 2011. "A parallel gpu implementation of the absolute nodal coordinate formulation with a frictional/contact model for the simulation of large flexible body systems". ASME Conference on Multibody Systems and Nonlinear Dynamics.
- [28] Melanz, D., 2012. "On the validation and applications of a parallel flexible multi-body dynamics implementation". M.S. thesis, University of Wisconsin–Madison.
- [29] MSC.Software, 2012. "Adams: Automatic dynamic analysis of mechanical systems". *Ann Arbor, Michigan*.
- [30] ABAQUS, 2004. "User manual - version 6.5, hibbitt, karls-son and sorensen". *Inc., Pawtucket, RI*.
- [31] Lucy, L., 1977. "A numerical approach to the testing of the fission hypothesis". *The Astronomical Journal*, **82**, pp. 1013–1024.
- [32] Gingold, R., and Monaghan, J., 1977. "Smoothed particle hydrodynamics-theory and application to non-spherical stars". *Monthly Notices of the Royal Astronomical Society*, **181**(1), pp. 375–389.
- [33] Monaghan, J., 2005. "Smoothed particle hydrodynamics". *Reports on Progress in Physics*, **68**(1), pp. 1703–1759.
- [34] Monaghan, J., 1989. "On the problem of penetration in particle methods". *Journal of Computational Physics*, **82**(1), pp. 1–15.
- [35] Dilts, G., 1999. "Moving-least-squares-particle hydrodynamics i. consistency and stability". *International Journal for Numerical Methods in Engineering*, **44**(8), pp. 1115–1155.
- [36] Koshizuka, S., Nobe, A., and Oka, Y., 1998. "Numeri-

- cal analysis of breaking waves using the moving particle semi-implicit method”. *International Journal for Numerical Methods in Fluids*, **26**(7), pp. 751–769.
- [37] Dalrymple, R., and Rogers, B., 2006. “Numerical modeling of water waves with the sph method”. *Coastal engineering*, **53**(2), pp. 141–147.
- [38] NVIDIA Corporation, 2012. NVIDIA CUDA Developer Zone. Available online at <https://developer.nvidia.com/cuda-downloads>.
- [39] Pazouki, A., and Negrut, D., 2012. “Direct simulation of lateral migration of bouyant particles in channel flow using gpu computing”. In *Computers and Information in Engineering*, CIE32, ASME.
- [40] Pazouki, A., and Negrut, D., 2012. “Numerical investigation of particle distribution in Poiseuille flow of suspension”. *Langmuir*. submitted.
- [41] Pixar, 1988, 1989, 2000, 2005. The renderman interface. Technical specification, Pixar.
- [42] Cook, R. L., Carpenter, L., and Catmull, E., 1987. “The Reyes Image Rendering Architecture”. *SIGGRAPH 1987 Proceedings*, pp. 95–102.