

A Unified Treatment of Elastostatic Contact Simulation for Real Time Haptics

Doug L. James* and Dinesh K. Pai*†

University of British Columbia

Abstract

We describe real-time, physically-based simulation algorithms for haptic interaction with elastic objects. Simulation of contact with elastic objects has been a challenge, due to the complexity of physically accurate simulation and the difficulty of constructing useful approximations suitable for real time interaction. We show that this challenge can be effectively solved for many applications. In particular global deformation of linear elastostatic objects can be efficiently solved with low run-time computational costs, using pre-computed Green's functions and fast low-rank updates based on *Capacitance Matrix Algorithms*. The capacitance matrices constitute exact force response models, allowing contact forces to be computed much faster than global deformation behavior. *Vertex pressure masks* are introduced to support the convenient abstraction of localized scale-specific point-like contact with an elastic and/or rigid surface approximated by a polyhedral mesh. Finally, we present several examples using the CyberGlove™ and PHANToM™ haptic interfaces.

1 Introduction

Discrete linear elastostatic models (LEMs) are important physically-based elastic primitives for computer haptics because they admit a very high-degree of precomputation, or “numerical compression” [1], in a way that affords cheap force response models suitable for force feedback rendering of stiff elastic objects during continuous contact. The degree of useful precomputation is quite limited for many types of nonlinear and/or dynamical elastic models, but LEMs are an exception, mainly due to the precomputability of time-independent *Green's functions* (GFs) and the applicability of linear superposition principles. Intuitively, GFs form a basis for describing all possible deformations of a LEM. Thus, while LEMs form a relatively simple class of elastic models in which geometric and material linearities are an ultimate limitation, the fact that the model is linear is also a crucial enabling factor. We conjecture that LEMs will remain one of the best runtime approximations of stiff elastic models for simulations requiring stable high-fidelity force feedback.

A central idea for LEMs in computer haptics is the formulation of the boundary value problem (BVP) solution in terms of suitable precomputed GFs using *Capacitance Matrix Algorithms* (CMAs). Derived from the Sherman-Morrison-Woodbury formula for low-rank updating of matrix inverses (and factorizations), CMAs have a long history in linear algebra [30, 16], where they have been commonly used for static reanalysis [22], to efficiently solve LEM contact mechanics problems [12, 25] and more recently for interactive LEM simulations [6, 21].

For computer haptics, a fundamental reason for choosing to compute the LEM elasticity solution using a CMA formulation, is that

the *capacitance matrix*¹ is the main quantity of interest: it is the *compliance matrix* which relates the force feedback response to the imposed contact displacements. Also, the precomputation of GFs effectively decouples the global deformation and force response calculations, so that the capacitance matrix can be extracted from the GFs at no extra cost; this is the fundamental mechanism by which a haptic interface can efficiently interact with a LEM of very large complexity. The user can feel no difference between the force response of the complete system and the capacitance matrix, because none exists. Lastly, CMAs are direct matrix solvers whose deterministic operation count is appealing for real time applications.

The second part of this paper addresses the special case of point-like interaction. It has long been recognized that point contact is a convenient abstraction for haptic interactions, and the PHANToM™ haptic interface is a testament to that fact. While it is possible to consider the contact area to be truly a point for rigid models, infinite contact pressures are problematic for elastic models and tractions need to be distributed over finite surface areas. We propose to do this efficiently by introducing nodal traction distribution masks which address at least two core issues. First, having a point contact with force distributed over a finite area is somewhat contradictory, and the traction distribution is effectively an underdetermined quantity without any inherent spatial scale. This is resolved by treating the contact as a single displacement constraint whose traction distribution enters as a user (or manipulator) specified parameter. The distribution of force on the surface of the model can then be consistently specified in a fashion which is independent of the scale of the mesh. Second, given the model is discrete, special care must be taken to ensure a sufficiently regular force response on the surface, since irregularities are very noticeable during sliding contact motions. By suitably interpolating nodal traction distributions, displacement constraints can be imposed which are consistent with regular contact forces for numerous discretizations.

The remainder of the paper is organized as follows. After a discussion of related work (§2), the notation and definitions for a wide class of linear elastostatic models used herein are given (§3). Fast CMAs for general BVP solution using precomputed GFs of a particular reference BVP type are described in detail in §4. Particular attention is given to the role of the capacitance matrix for the construction of globally consistent stiffness matrices for use in local haptic buffer models (§5). The special case of point-like contacts are considered in detail, and we introduce (runtime computable) vertex masks for haptic presentation of surface stiffness (§6). Some results are given for our implementations (§7) followed by conclusions and a discussion of future work (§8).

2 Related Work

While a significant amount of work has been done on interactive simulation of physically-based elastic models, e.g., in computer

¹The term “capacitance” is due to historical convention [16].

*Institute of Applied Mathematics

†Dept. of Computer Science, 2366 Main Mall, Vancouver, B.C., V6T 1Z4, Canada, {djames|pai}@cs.ubc.ca

graphics [14], only a relatively small subset of work has addressed the simulation requirements of force feedback computer haptics. Much of the computer haptics literature on deformable models has been concerned with surgical simulation, while our focus is more general. It is not our intent to survey all related haptic work here, but simply to mention some relevant work combining kinesthetic force feedback with elastic models.

2.1 Elastostatic Models

There are several instances in the literature of real time simulation of linear elastostatic models based on precomputed GFs methods. These models were used because of their low runtime costs, and desirable force feedback properties. To date only polygonal models based on FEM and BEM discretizations have been considered, although other variants are possible within the linear systems description presented in the following sections.

Of particular relevance is the work done by researchers at INRIA who have made extensive use of a real time elastostatic FEM model for liver related surgical simulations [6, 5, 9]. During a precomputation phase they have used condensation [35] as well as iterative methods [8] to compute displacement responses due to unit forces applied to vertices on the “free” boundary. At run time, they solve a small system of equations to determine the correct superposition of responses to satisfy the applied surface constraints, which may be identified as a case of the capacitance matrix approach. We note that the point-like contact approach used in [8] could benefit from pressure mask concepts (§6). Recently, they have used anisotropic material properties for the liver [29]. Other groups have also used the precomputed elastostatic FEM approach of [6] for surgical simulation, e.g., the KISMET surgical simulator [23] incorporates precomputed models to provide high-fidelity haptic force feedback.

A limitation of the GF precomputation strategy is that incremental runtime modifications of the model require extra runtime computations. While it may be too costly for interactive applications, this can also be efficiently performed using low-rank updating techniques such as for static reanalysis in the engineering community [22]. For surgical simulation, a practical approach has been to use a hybrid domain decomposition approach in which a more easily modified dynamic model is used in a smaller region to be cut [9, 17].

Finally, the authors presented a interactive animation technique in [21] which combined precomputed GFs of boundary element models with matrix-updating techniques for fast boundary value problem (BVP) solution. Although computer haptics was an intended application, no force feedback implementation was mentioned. This paper generalizes that approach with a broad GF-based linear systems framework that subsumes the discretization issues of both [21] and the FEM approaches of [6, 8].

2.2 Other Elastic Models

Various approaches have been taken to simulate dynamic elastic models, by addressing commonly encountered difficulties such as the computational complexity of time-stepping 3D models, and numerical time integration issues, e.g., stiffness and stability. In order to meet the intense demands of force feedback rendering rates, most have opted for a multirate simulation approach. It is worth noting explicitly that methods for interactively simulating soft dynamic objects are in many ways complementary to the CMA methods presented here for simulating relatively stiff LEM. A few notable examples are now mentioned.

Local buffer models were presented by Balaniuk in [2] for rendering forces computed by e.g., deformable object, simulators which can not deliver forces at fast rendering rates. An application of the technique was presented for a virtual echographic exam

training simulator in [10]. While we do not use the same approach here, the local buffer model concept is related to our capacitance matrix method for force computation.

Astley and Hayward [1] introduced an approximation for linear viscoelastic FEM models that also exploits linearity, in this case by precomputing multilevel Norton equivalents for the system’s stiffness matrix. By doing so, haptic interaction is possible by employing an explicit multirate integration scheme wherein a model associated with the contact region is integrated at a higher rate than the remaining coarser model.

Çavuşoğlu and Tendick [7] also use a multirate approach. Balanced model reduction of a linearization of their nonlinear dynamical lumped element model is used to suggest a spatially localized dynamic model approximation for force feedback rendering. While promising, the example considered is a very special case of a supported model, and it is unclear how the local model would be derived in more generic geometric cases, as well as in the presence of nonlocal influences such as for multiple changing contacts, e.g., with surgical tools.

Debunne et al. [11] presented a space-time approach for simulating a hierarchical multirate dynamic linear-strain model. Zhuang and Canny [34] use a dynamic lumped finite element model exhibiting nonlinear (Green’s) strain. It is capable of being time-stepped at graphics frame rates for sufficiently soft objects using an explicit integration scheme. Interactive simulation of dynamic elastic models exclusively for superquadric shapes was considered by Ramanathan and Metaxas [31]. Volumetric and voxel-based modeling approaches for surgical simulation have also been considered, e.g., by Gibson et al. [13].

3 Linear Elastostatic Boundary Model Preliminaries

Linear elastostatic objects are essentially three-dimensional linear springs, and as such they are useful modeling primitives for physically-based simulations. The unfamiliar reader might consult a suitable background reference before continuing [3, 18, 35, 4, 21].

In this section, background material for a generic discrete GF description for a variety of precomputed linear elastostatic models is provided. Conceptually, GFs form a basis for describing all possible deformations of a LEM subject to a certain class of constraints. This is useful because it (1) provides a common language to describe all discrete LEMs, (2) subsumes extraneous discretization details by relating only physical quantities, and (3) clarifies the generality of the force feedback algorithms described later.

Another benefit of using GFs is that they provide an efficient means for exclusively simulating only boundary data (displacements and forces) if desired. While it is possible to simulate various internal volumetric quantities (see §3.5), simulating only boundary data involves less computation. This is sufficient since we are primarily concerned with interactive simulations that impose surface constraints and provide feedback via surface deformation and contact forces.

3.1 Geometry and Material Properties

Given that the fast solution method is based on linear systems principles, essentially any linear elastostatic model with physical geometric and material properties is admissible. We shall consider models in three dimensions, although many arguments also apply to lower dimensions. Suitable models would of course include bounded volumetric objects with various internal material properties, as well as special subclasses such as thin plates and shells. Since only a boundary or interface description is utilized for specifying user interactions, other exotic geometries may also be easily

considered such as semi-infinite domains, exterior elastic domains, or simply any set of parametrized surface patches with a linear response. Similarly, numerous representations of the surface and associated displacement shape functions are also possible, e.g., polyhedral, NURBS or subdivision surfaces [32].

3.2 Nodal Displacements and Traction

Let the undeformed boundary be denoted by Γ . The change in shape of the surface is described by the surface *displacement field* $\mathbf{u}(\mathbf{x})$, $\mathbf{x} \in \Gamma$, and the surface force distribution is called the *traction² field* $\mathbf{p}(\mathbf{x})$, $\mathbf{x} \in \Gamma$. We will assume that each surface field is parametrized by n nodal variables (see Figure 1), so that the discrete displacement and traction vectors are

$$\mathbf{u} = [\mathbf{u}_1, \dots, \mathbf{u}_n]^T \quad (1)$$

$$\mathbf{p} = [\mathbf{p}_1, \dots, \mathbf{p}_n]^T, \quad (2)$$

respectively, where each nodal value is a vector in \mathbb{R}^3 . This description admits a very large class of surface displacement and traction distributions.




Figure 1: *Illustration of discrete nodal displacements* \mathbf{u} defined at vertices on the undeformed boundary Γ (solid blue line), that result in a deformation of the surface (to dashed red line). Although harder to illustrate, a similar definition exists for the traction vector, \mathbf{p} .

In order to relate traction distributions to forces, define a scalar function space, \mathcal{L} , on the model’s boundary:

$$\mathcal{L} = \text{span} \{ \phi_j(\mathbf{x}), \quad j = 1 \dots n, \quad \mathbf{x} \in \Gamma \}, \quad (3)$$

where $\phi_j(\mathbf{x})$ is a scalar basis function associated with the j^{th} node. The continuous traction field is then a 3-vector function with components in \mathcal{L} ,

$$\mathbf{p}(\mathbf{x}) = \sum_{j=1}^n \phi_j(\mathbf{x}) \mathbf{p}_j, \quad (4)$$

The force on any surface area is equal to the integral of $\mathbf{p}(\mathbf{x})$ on that area. It then follows that the nodal force associated with any nodal traction is given by

$$\mathbf{f}_j = a_j \mathbf{p}_j \quad \text{where} \quad a_j = \int_{\Gamma} \phi_j(\mathbf{x}) d\Gamma_{\mathbf{x}} \quad (5)$$

defines the area associated with the j^{th} node.

Our implementation uses linear boundary element models, for which the nodes are vertices of a closed triangle mesh. The mesh is modeled as a Loop subdivision surface [24] to conveniently obtain multiresolution models for rendering as well as uniformly parameterized surfaces ideal for BEM discretization and deformation depiction. The displacement and traction fields have convenient vertex-based descriptions

$$\mathbf{u}_j = \mathbf{u}(\mathbf{x}_j), \quad \mathbf{p}_j = \mathbf{p}(\mathbf{x}_j),$$

²Surface traction describes force per unit area.

where $\mathbf{x}_j \in \Gamma$ is the j^{th} vertex. The traction field is a piecewise linear function, and $\phi_j(\mathbf{x})$ represents a “hat function” located at the j^{th} vertex with $\phi_j(\mathbf{x}_j) = 1$. Given our implementation, we shall often refer to node and vertex interchangeably.

3.3 Discrete Boundary Value Problem (BVP)

At each step of the simulation, a discrete BVP must be solved which relates specified and unspecified nodal values, e.g., to determine deformation and feedback forces. Without loss of generality, it shall be assumed that either position or traction constraints are specified at each boundary node, although this can be extended to allow mixed conditions, e.g., normal displacement and tangential tractions. Let nodes with prescribed displacement or traction constraints be specified by the mutually exclusive index sets Λ_u and Λ_p , respectively, so that $\Lambda_u \cap \Lambda_p = \emptyset$ and $\Lambda_u \cup \Lambda_p = \{1, 2, \dots, n\}$. In order to guarantee an equilibrium constraint configuration we will require that there is at least one displacement constraint, i.e., $\Lambda_u \neq \emptyset$. We shall refer to the (Λ_u, Λ_p) pair as the *BVP type*.

Typical boundary conditions for a force feedback loop consist of specifying some (compactly supported) displacement constraints in the area of contact, with “free” boundary conditions (zero traction) and other (often zero displacement) support constraints outside the contact zone. The solution to (7) yields the rendered contact forces and surface deformation.

Denote the unspecified and complementary specified nodal variables by

$$\mathbf{v}_j = \begin{cases} \mathbf{p}_j : j \in \Lambda_u \\ \mathbf{u}_j : j \in \Lambda_p \end{cases} \quad \text{and} \quad \bar{\mathbf{v}}_j = \begin{cases} \bar{\mathbf{u}}_j : j \in \Lambda_u \\ \bar{\mathbf{p}}_j : j \in \Lambda_p \end{cases}, \quad (6)$$

respectively. By linearity of the discrete elastic model, there formally exists a linear relationship between all nodal boundary variables

$$0 = \mathbf{A}\mathbf{v} + \bar{\mathbf{A}}\bar{\mathbf{v}} = \mathbf{A}\mathbf{v} - \mathbf{z} \quad (7)$$

where the BVP system matrix \mathbf{A} and its complementary matrix $\bar{\mathbf{A}}$ are, in general, dense block n -by- n matrices [18]. Body force terms associated with other phenomena, e.g., gravity, have been omitted for simplicity, but can be included since they only add an extra contribution to the \mathbf{z} term.

A key relationship between BVP system matrices $(\mathbf{A}, \bar{\mathbf{A}})$ of different BVP types (Λ_u, Λ_p) is that they are related by exchanges of corresponding block columns, e.g., $(\mathbf{A}_j, \bar{\mathbf{A}}_j)$, and therefore small changes to the BVP type result in low-rank changes to the BVP system matrices (see §4.2.1).

While the boundary-only system matrices in (7) could be constructed explicitly, e.g., via condensation for FEM models [35] or using a boundary integral formulation (see next section), it need not be in practice. The discrete integral equation in Equation 7 is primarily a common starting point for later definition of GFs and derivation of the CMA, while GFs may be generated with any convenient numerical method, or even robotically scanned from real objects [28].

3.4 Example: Boundary Element Models

A simple closed-form definition of $(\mathbf{A}, \bar{\mathbf{A}})$ is possible for models discretized with the boundary element method (BEM) [4, 21]; BEM discretizations are possible for models with homogeneous and isotropic material properties. The surface-based nodal quantities are related by the dense linear block matrix system

$$0 = \mathbf{H}\mathbf{u} - \mathbf{G}\mathbf{p} = \sum_{j=1}^n \mathbf{h}_{ij} \mathbf{u}_j - \sum_{j=1}^n \mathbf{g}_{ij} \mathbf{p}_j \quad (8)$$

where G and H are n -by- n block matrices, with each matrix element, g_{ij} or h_{ij} , a 3-by-3 influence matrix with known expressions [4]. In this case, the j^{th} block columns of A and \bar{A} may be identified as column exchanged variants of G and H :

$$A_{:j} = \begin{cases} -G_{:j} & : j \in \Lambda_u \\ H_{:j} & : j \in \Lambda_p \end{cases} \quad (9)$$

$$\bar{A}_{:j} = \begin{cases} H_{:j} & : j \in \Lambda_u \\ -G_{:j} & : j \in \Lambda_p \end{cases} \quad (10)$$

While we use BEM models for our implementation, we reiterate that the CMA is independent of the method used to generate the GFs (explained next).

3.5 Fast BVP Solution with Green's Functions

GFs of a single BVP type provide an economical means for solving (7) for that BVP, and when combined with the CMA (§4) will also be useful for solving other BVP types. From (7), the general solution of a BVP type (Λ_u, Λ_p) may be expressed in terms of discrete GFs³ as

$$\mathbf{v} = \Xi \bar{\mathbf{v}} = \sum_{j=1}^n \xi_j \bar{\mathbf{v}}_j = \sum_{j \in \Lambda_u} \xi_j \bar{\mathbf{u}}_j + \sum_{j \in \Lambda_p} \xi_j \bar{\mathbf{p}}_j, \quad (11)$$

where the discrete GFs of the BVP system are the block column vectors

$$\xi_j = -(\mathbf{A}^{-1} \bar{\mathbf{A}})_{:j} \quad (12)$$

and

$$\Xi = -\mathbf{A}^{-1} \bar{\mathbf{A}} = [\xi_1 \xi_2 \cdots \xi_n]. \quad (13)$$

Equation (11) may be taken as the definition of the discrete GFs (and even (7)), since it is clear that the j^{th} GF simply describes the linear response of the system to the j^{th} node's specified boundary value, $\bar{\mathbf{v}}_j$. Once the GFs have been computed for one BVP type, that class of BVPs may be solved easily using (11). An attractive feature for interactive applications is that the entire solution can be obtained in $18ns$ flops⁴ if only s boundary values (BV) are nonzero (or have changed since the last time step). Temporal coherence may also be exploited by considering the effect of individual changes in components of $\bar{\mathbf{v}}$ on the solution \mathbf{v} .

Further leveraging linear superposition, each GF system response may be augmented with other additional information in order to simulate other precomputable quantities. Volumetric stress, strain and displacement data may also be simulated at preselected locations. Applications could use this to monitor stresses and strains to determine, e.g., if fracture occurs or that a nonlinear correction should be computed.

3.6 Precomputation of Green's Functions

Since the GFs for a single BVP type only depend on geometric and material properties of the deformable object, they may be precomputed for use in a simulation. This provides a dramatic speed-up for simulation by determining the deformation basis (the GFs) ahead of time. While this is not necessary a huge amount of work (see Table 2), the principal benefits for interactive simulations are the availability of the GF elements via cheap look-up table operations, as well as the elimination of redundant runtime computation when

³Note on GF terminology: we are concerned with discrete numerical approximations of continuous GFs, however for convenience these GF vectors may simply be referred to as GFs.

⁴counting both + and *

computing solutions, e.g., using a haptic device to grab a vertex of the model and move it around simply renders a single GF.

Once a set of GFs for a LEM are precomputed, the overall stiffness can be varied at runtime by scaling BVP forces accordingly, however changes in compressibility and internal material distributions do require recomputation. In practice it is only necessary to compute the GF corresponding to nodes which may have changing or nonzero boundary values during the simulation.

4 Fast Global Deformation using Capacitance Matrix Algorithms (CMAs)

This section presents an algorithm for using the precomputed GFs of a relevant *Reference BVP* (RBVP) type to efficiently solve other BVP types. With an improved notation and emphasis on computer haptics, this section unifies and extends the approaches presented in [21] exclusively for BEM models, and for FEM models in, e.g., [6], in a way that is applicable to all LEMs regardless of discretization, or origin of GFs [28]. Haptic applications are considered in §5.

4.1 Reference Boundary Value Problem (RBVP) Choice

A key step in the GF precomputation process is the initial identification of a RBVP type, denoted by $(\Lambda_u^0, \Lambda_p^0)$, that is representative of the BVP types arising during simulations. For interactions with an exposed free boundary, a common choice is to have the uncontacted model attached to a rigid support as shown in Figure 2. The n -by- n block system matrices associated with the RBVP are identified with a subscript as A_0 and \bar{A}_0 , and the corresponding GFs are hereafter always denoted by Ξ .

Note that the user's choice of RBVP type determines which type of nodal constraints (displacement of traction) are commonly specified (in order to define Ξ), but is independent of the actual numerical boundary values $\bar{\mathbf{v}}$ used in practice. For example, there are no requirements that certain boundary values are zero, although this results in fewer summations (see (11)).




Figure 2: *Reference Boundary Value Problem (RBVP) example:* The RBVP associated with a model attached to a flat rigid support is shown with boundary regions having fixed (Λ_u^0) or free (Λ_p^0) nodal constraints indicated. A typical simulation would impose contacts on the free boundary via displacement constraints with the CMA.

4.2 Capacitance Matrix Algorithm (CMA) for BVP Solution

Precomputed GFs speed-up the solution to the RBVP, but they can also dramatically reduce the amount of work required to solve related BVP when used in conjunction with CMAs. This section describes the CMA and presents the derivation of related formulae.

4.2.1 Relevant Formulae

Suppose the constraint-type changes, e.g., displacement↔traction, with respect to the RBVP at s nodes specified by the list of nodal indices $S = \{S_1, S_2, \dots, S_s\}$. As mentioned earlier, it follows from (6) and (7) that the new BVP system matrices (A, \bar{A}) are related to those of the RBVP (A_0, \bar{A}_0) by s block column swaps. This may be written as

$$A = A_0 + (\bar{A}_0 - A_0) EE^T \quad (14)$$

$$\bar{A} = \bar{A}_0 + (A_0 - \bar{A}_0) EE^T \quad (15)$$

where E is an n -by- s block matrix

$$E = [I_{S_1} \ I_{S_2} \ \dots \ I_{S_s}]$$

containing columns of the n -by- n identity block matrix, I , specified by the list of updated nodal indices S . Postmultiplication by E extracts columns specified by S . Throughout, E is used to write sparse matrix operations using dense data, e.g., Ξ , and like the identity matrix, it should be noted that there is no cost involved in multiplication by E or its transpose.

Since the BVP solution is

$$v = A^{-1}z = -A^{-1}\bar{A}\bar{v}, \quad (16)$$

substituting (15) for \bar{A} and the Sherman-Morrison-Woodbury formula [15] for A^{-1} (using the GF definition $\Xi = -A_0^{-1}\bar{A}_0$),

$$A^{-1} = A_0^{-1} + (I + \Xi)E(-E^T\Xi E)^{-1}E^T A_0^{-1}, \quad (17)$$

into (16), leads directly to an expression for the solution in terms of the precomputed GFs⁵. The resulting *capacitance matrix formulae* are

$$v = \underbrace{v^{(0)}}_{n \times 1} + \underbrace{(E + (\Xi E))}_{n \times s} \underbrace{C^{-1}}_{s \times s} \underbrace{E^T v^{(0)}}_{s \times 1} \quad (18)$$

where C is the s -by- s *capacitance matrix*, a negated submatrix of Ξ ,

$$C = -E^T \Xi E, \quad (19)$$

and $v^{(0)}$ is the response of the RBVP system to $z = -\bar{A}\bar{v}$,

$$v^{(0)} = A_0^{-1}z = [\Xi(I - EE^T) - EE^T] \bar{v}. \quad (20)$$

4.2.2 Algorithm for BVP Solution

With Ξ precomputed, formulae (18)-(20) immediately suggest an algorithm given that only simple manipulations of Ξ and inversion of the smaller capacitance submatrix are required. An algorithm for computing *all* components of v is as follows:

- For each new BVP type (with a different C matrix) encountered, construct and temporarily store C^{-1} (or LU factors) for subsequent use.
- Construct $v^{(0)}$.
- Extract $E^T v^{(0)}$ and apply the capacitance matrix inverse to it, $C^{-1}(E^T v^{(0)})$.
- Add the s column vectors $(E + (\Xi E))$ weighted by $C^{-1}(E^T v^{(0)})$ to $v^{(0)}$ for the final solution v .

⁵Similarly from [21] with $\delta A_S = (\bar{A}_0 - A_0)E$.

4.2.3 Complexity Issues

Given s nonzero boundary values, each new capacitance matrix LU factorization involves at most $\frac{2}{3}s^3$ flops, after which each subsequent solve involves approximately $18ns$ flops ($s \ll n$). This is particularly attractive when $s \ll n$ is small, such as often occurs in practice with localized surface contacts.

An important feature of the CMA for interactive methods is that it is a direct matrix solver with a deterministic operation count. It is therefore possible to predict the runtime cost associated with each matrix solve and associated force feedback subcomputations (see §5), thus making CMAs predictable for real-time computations.

4.3 Selective Deformation Computation

A major benefit of the CMA direct BVP solver is that it is possible to just evaluate selected components of the solution vector at runtime, with the total computing cost proportional to the number of components desired. This is a key enabling feature for force feedback where, e.g., contact forces are desired at different rates than the geometric deformations. Selective evaluation would also be useful for optimizing (self) collision detection queries, avoiding simulation of occluded or undesired portions of the model, as well as rendering adaptive level of detail representations.

In general, any subset of solution components may be determined at a smaller cost than computing v entirely. Let the solution be desired at nodes specified by the set of indices D , with the desired components of v extracted by E_D^T . Using (18), the selected solution components may be evaluated as

$$E_D^T v = E_D^T v^{(0)} + E_D^T (E + (\Xi E)) C^{-1} E^T v^{(0)}$$

using only $\mathcal{O}(s^2 + s|D|)$ operations. The case where $S = D$ is especially important for force feedback and is discussed exclusively in the following section.

5 Capacitance Matrices as Local Buffer Models

For force feedback enabled simulations in which user interactions are modeled as displacement constraints applied to an otherwise free boundary, the capacitance matrix has a very important role: it constitutes an exact contact force response model by describing the compliance of the contact zone. Borrowing terminology from [2], we say that the capacitance matrix can be used as a local buffer model. While the capacitance matrix is used in §4.2.2 to determine the linear combination of GFs required to solve a particular BVP and reconstruct the global deformation, it also has the desirable property that it effectively decouples the global deformation calculation from that of the local force response. The most relevant benefit for haptics is that the local contact force response may be computed at a much faster rate than the global deformation.

5.1 Capacitance Matrix Local Buffer Model

From (18), the S components of the solution v are

$$\begin{aligned} E^T v &= E^T [v^{(0)} + (E + (\Xi E)) C^{-1} E^T v^{(0)}] \\ &= E^T v^{(0)} + \underbrace{(E^T E)}_I C^{-1} E^T v^{(0)} + \underbrace{(E^T \Xi E)}_{-C} C^{-1} E^T v^{(0)} \\ &\downarrow \\ &= E^T v^{(0)} + C^{-1} E^T v^{(0)} - E^T v^{(0)} \\ &= C^{-1} (E^T v^{(0)}). \end{aligned} \quad (21)$$

Consider the situation, which naturally arises in haptic interactions, in which the only nonzero constraints are updated displacement constraints, i.e.,

$$\bar{\mathbf{v}} = \mathbf{E}\mathbf{E}^T\bar{\mathbf{v}} \Rightarrow \mathbf{v}^{(0)} = -\bar{\mathbf{v}} \quad (\text{using (20)}). \quad (22)$$

In this case, the capacitance matrix completely characterizes the local contact response, since (using (22) in (21))

$$\mathbf{E}^T\mathbf{v} = -\mathbf{C}^{-1}\mathbf{E}^T\bar{\mathbf{v}}. \quad (23)$$

This in turn parametrizes the global response since these components (not in \mathbf{S}) are

$$\begin{aligned} (\mathbf{I} - \mathbf{E}\mathbf{E}^T)\mathbf{v} &= (\mathbf{I} - \mathbf{E}\mathbf{E}^T) \left[\mathbf{v}^{(0)} + (\mathbf{E} + (\Xi\mathbf{E}))\mathbf{C}^{-1}\mathbf{E}^T\mathbf{v}^{(0)} \right] \\ &= (\mathbf{I} - \mathbf{E}\mathbf{E}^T)(\Xi\mathbf{E})(\mathbf{E}^T\mathbf{v}) \end{aligned} \quad (24)$$

where we have used (23) and the identity $(\mathbf{I} - \mathbf{E}\mathbf{E}^T)\mathbf{E} = 0$. Such properties allow the capacitance matrix and Ξ to be used to derive efficient local models for surface contact.

For example, given the specified contact zone displacements

$$\mathbf{u}_S = \mathbf{E}^T\bar{\mathbf{v}}, \quad (25)$$

the resulting tractions are

$$\mathbf{p}_S = \mathbf{E}^T\mathbf{v} = -\mathbf{C}^{-1}(\mathbf{E}^T\bar{\mathbf{v}}) = -\mathbf{C}^{-1}\mathbf{u}_S, \quad (26)$$

and the rendered contact force is

$$\mathbf{f} = \mathbf{a}_S^T\mathbf{p}_S = (-\mathbf{a}_S^T\mathbf{C}^{-1})\mathbf{u}_S = \mathbf{K}_S\mathbf{u}_S, \quad (27)$$

where \mathbf{K}_S is the effective stiffness of the contact zone used for force feedback rendering,

$$\mathbf{a}_S = (a_{S_1}, a_{S_2}, \dots, a_{S_s})^T \otimes \mathbf{I}_3 \quad (28)$$

represents nodal areas (5), and \mathbf{I}_3 is the scalar 3-by-3 identity matrix. A similar expression may be obtained for torque feedback. The visual deformation corresponding to solution components outside the contact zone is then given by (24) using $\mathbf{p}_S = \mathbf{E}^T\mathbf{v}$.

5.2 Example: Single Displacement Constraint

A simple case, which will be discussed in much greater detail in §6, is that of imposing a displacement constraint on single a node k which otherwise had a traction constraint in the RBVP⁶. The new BVP therefore has only a single constraint switch with respect to the RBVP, and so $s = 1$ and $\mathbf{S} = \{k\}$. The capacitance matrix here is just $\mathbf{C} = -\Xi_{kk}$ so that the k^{th} nodal values are related by

$$\mathbf{p}_k = -\mathbf{C}^{-1}\bar{\mathbf{u}}_k = (\Xi_{kk})^{-1}\bar{\mathbf{u}}_k \quad \text{or} \quad \bar{\mathbf{u}}_k = \Xi_{kk}\mathbf{p}_k.$$

The capacitance matrix can generate the force response, $\mathbf{f} = a_k\mathbf{p}_k$, required for haptics in $\mathcal{O}(1)$ operations, and for graphical feedback the corresponding global solution is $\mathbf{v} = \xi_k\mathbf{p}_k$.

5.3 Force Feedback for Multiple Displacement Constraints

When multiple force feedback devices are interacting with the model by imposing displacement constraints, the force and stiffness felt by each device are tightly coupled in equilibrium. For example, the stiffness felt by the thumb in Figure 3 will depend on how the other fingers are supporting the object. For multiple contacts like this, the capacitance matrix again provides an efficient force response model for haptics. Without presenting the equations in detail, we shall just mention that the force responses for each of the contact patches can be derived from the capacitance matrix in a manner similar to equations (25)-(28).

⁶This case occurs, for instance, when the tip of a haptic device comes into contact with the free surface of an object.




Figure 3: *Grasping simulation*: Using a CyberTouch data input device from Virtual Technologies Inc. (Top), a virtual hand (Bottom) was used to deform an elastostatic BEM model with approximately 900 surface degrees of freedom (dof) at graphical frame rates (> 30 FPS) on a personal computer. The capacitance matrix algorithm was used to impose displacement constraints on an otherwise free boundary, often updating over 100 dof per frame. While force feedback was not present, the capacitance matrices computed could also have been used to render contact forces at a rate higher than that of the graphical simulation.

6 Surface Stiffness Models for Point-like Contact

The second part of this paper concerns a special class of boundary conditions describing point-like contact interactions. Such interactions are commonly in the haptics literature for *rigid* surface models [26, 19]. Unlike their rigid counterparts, special care must be taken with elastic models to define *finite contact areas* for point-like interactions since point-like contacts defined only as single-vertex (§5.2) or nearest neighbour [8] constraints lead to mesh-related artifacts, and ambiguous interactions as the mesh is refined (see Figure 4). However, the benefit of point-like contacts comes from the convenience of the point-like parameterization of the contact and not because the contact is highly concentrated or “pin-like”. We present an approach using *vertex pressure masks* which maintains the single contact description yet distribute forces on a specified

scale. This allows point contact stiffnesses to be consistently defined as the mesh scale is refined, and provides an efficient method for force feedback rendering of forces with regular spatial variation on irregular meshes.




Figure 4: *Point contact must not be taken literally for elastic models:* This figure illustrates the development of a displacement singularity associated with a concentrated surface force as the continuum limit is approached. In the left image, an upward unit force applied to a vertex of a discrete elastic model results in a finite vertex displacement. As the model's mesh is refined (middle and right image), the same concentrated force load eventually tends to produce a singular displacement at the contact location, and the stiffness of any single vertex approaches zero (see Table 1). Such point-like constraints are mathematically ill-posed for linear models based on a small-strain assumption, and care must be taken to meaningfully define the interaction.

6.1 Vertex Pressure Masks for Distributed Point-like Contacts

In this section, the distribution of force is described using compactly-supported per-vertex pressure masks defined on the free boundary in the neighbourhood of each vertex.

6.1.1 Vertex Pressure Mask Definition

Scalar pressure masks provide a flexible means for modeling vector pressure distributions associated with each node. This allows a force applied at the i^{th} node to generate a traction distribution which is a linear combination of $\{\phi_j(\mathbf{x})\}$ and not just $\phi_i(\mathbf{x})$.

In the continuous setting, a scalar surface density $\rho(\mathbf{x}) : \Gamma \rightarrow \mathbb{R}$ will relate the localized contact force \mathbf{f} to the applied traction \mathbf{p} via⁷

$$\mathbf{p}(\mathbf{x}) = \rho(\mathbf{x})\mathbf{f}$$

which in turn implies the normalization condition

$$\int_{\Gamma} \rho(\mathbf{x}) d\Gamma \mathbf{x} = 1. \quad (29)$$

In the discrete setting, the piecewise linear surface density on Γ is

$$\rho(\mathbf{x}) = \sum_{j=1}^n \phi_j(\mathbf{x}) \rho_j \in \mathcal{L}, \quad (30)$$

and is parameterized by the discrete scalar vertex mask vector,

$$\rho = [\rho_1, \rho_2, \dots, \rho_n]^T.$$

Substituting (30) into (29), the discrete normalization condition satisfied becomes

$$a^T \rho = 1, \quad (31)$$

where a are the vertex areas from (5). Notice that the mask density ρ has units of $\frac{1}{\text{area}}$.

In practice, the vertex pressure mask ρ may be specified in a variety of ways. It could be specified at runtime, e.g., as the byproduct of a physical contact mechanics solution, or be a user specified quantity. We shall consider the case where there is a compactly

⁷Tensor-valued masks for torque feedback can also be computed.

supported scalar function $\rho(\mathbf{x})$ specified at each vertex on the free boundary. The corresponding discrete vertex mask ρ may then be defined using nodal collocation (see Figure 5),

$$\rho_j = \begin{cases} \rho(\mathbf{x}_j), & j \in \Lambda_p^0, \\ 0, & j \in \Lambda_u^0, \end{cases},$$

followed by suitable normalization,

$$\rho := \frac{\rho}{a^T \rho},$$

to ensure the satisfaction of (31).




Figure 5: *Collocated scalar masks:* A direct means for obtaining a relative pressure amplitude distribution about each node, is to employ a user-specified scalar functional of the desired spatial scale. The scalar pressure mask is then given by nodal collocation (left), after which the vector traction distribution associated with a nodal point load is then computed as the product of the applied force vector and the (compactly supported) scalar mask (right).

In the following, denote the density mask for the i^{th} vertex by the n -vector ρ^i , with nonzero values being indicated by the set of masked nodal indices \mathcal{M}_i . Since the intention is to distribute force on the free boundary, masks will only be defined for $i \in \Lambda_p^0$. Additionally, these masks will only involve nodes on the free boundary, $\mathcal{M}_i \subset \Lambda_p^0$, as well as be nonempty, $|\mathcal{M}_i| > 0$.

6.1.2 Example: Spherical Mask Functionals

Spherically symmetric radially decreasing mask functionals with a scale parameter were suitable candidates for constructing vertex masks via collocation on smooth surfaces. One functional we used (see Figure 6 and 7) had linear radial dependence,

$$\rho^i(\mathbf{x}; r) = \begin{cases} 1 - \frac{|\mathbf{x} - \mathbf{x}_i|}{r}, & |\mathbf{x} - \mathbf{x}_i| < r, \\ 0, & \text{otherwise.} \end{cases},$$

where r specifies the radial scale⁸. The effect of changing r is shown in Figure 6.




Figure 6: *Illustration of changing mask scale:* An exaggerated pulling deformation illustrates different spatial scales in two underlying traction distributions. In each case, pressure masks were generated using the linear spherical mask functional (see §6.1.2) for different values of the radius parameter, r .

⁸ r may be thought of as the size of the haptic probe's tip.




Figure 7: *Effect of pressure masks on surface stiffness:* Even models with reasonable mesh quality, such as this simple BEM kidney model, can exhibit perceptible surface stiffness irregularities when single-vertex stiffnesses are used. A plot (a) of the vertex area, a , clearly indicates regions of large (dark red) and small (light blue) triangles. In (b) the norm of the single-vertex surface stiffness, $\|\mathbf{K}(\mathbf{x})\|$, reveals a noticeable degree of mesh-related stiffness artifacts. On the other hand, the stiffness plotted in (c) was generated using a pressure mask (collocated linear sphere functional (see §6.1.2) of radius twice the mesh’s mean edge length) and better approximates the regular force response expected of such a model. Masks essentially provide anti-aliasing for stiffnesses defined with discrete traction distributions, and help avoid “soft spots.”

6.2 Vertex Stiffnesses using Pressure Masks

Having consistently characterized point-like force loads using vertex pressure masks, it is now possible to calculate the stiffness of each vertex. In the following sections, these vertex stiffnesses will then be used to compute the stiffness at any point on model’s surface for haptic rendering of point-like contact.

6.2.1 Elastic Vertex Stiffness, \mathbf{K}^E

For any single node on the free boundary, $i \in \Lambda_p^0$, a finite force stiffness, $\mathbf{K}_i \in \mathbb{R}^{3 \times 3}$, may be associated with its displacement, i.e.,

$$\mathbf{f} = \mathbf{K}_i \mathbf{u}_i, \quad i \in \Lambda_p^0.$$

As a sign convention, it will be noted that for any single vertex displacement

$$\mathbf{u}_i \cdot \mathbf{f} = \mathbf{u}_i \cdot (\mathbf{K}_i \mathbf{u}_i) \geq 0, \quad i \in \Lambda_p^0$$

so that positive work is done deforming the object.

Given a force \mathbf{f} applied at vertex $i \in \Lambda_p^0$, the corresponding distributed traction constraints are

$$\mathbf{p}_j = \rho_j^i \mathbf{f}.$$

Since the displacement of the i^{th} vertex is

$$\mathbf{u}_i = \sum_{j \in \mathcal{M}_i} \rho_j^i \Xi_{ij} \mathbf{f},$$

therefore the effective elastic stiffness of the masked vertex is

$$\mathbf{K}_i = \mathbf{K}_i^E = \left(\sum_{j \in \mathcal{M}_i} \rho_j^i \Xi_{ij} \right)^{-1}, \quad i \in \Lambda_p^0. \quad (32)$$

Some examples are provided in Table 1 and Figure 7.

Therefore, in the simple case of a single masked vertex displacement constraint \mathbf{u}_i , the local force response model exactly determines the resulting force, $\mathbf{f} = \mathbf{K}_i \mathbf{u}_i$, distributed in the masked region. The corresponding globally consistent solution is

$$\mathbf{v} = \zeta_i \mathbf{f} = \left(\sum_{j \in \mathcal{M}_i} \rho_j^i \xi_j \right) \mathbf{f}$$

where ζ_i is the convolution of the GFs with the mask ρ , and characterizes the distributed force load. The limiting case of a single vertex constraint corresponds to $\mathcal{M}_i = \{i\}$ with $\rho_j^i = \delta_{ij}/a_i$ so that the convolution simplifies to $\zeta_i = \xi_i/a_i$.

Mesh Level	Vertices	$\ \mathbf{K}\ _F$, Single	$\ \mathbf{K}\ _F$, Masked
1	34	7.3	13.3
2	130	2.8	11.8
3	514	1.1	11.2

Table 1: *Vertex stiffness dependence on mesh resolution:* This table shows vertex stiffness (Frobenius) norms (in arbitrary units) at the top center vertex of the BEM model in Figure 10(a), as geometrically modeled using Loop subdivision meshes for three different levels of resolution. The stiffness corresponding to a single vertex constraint exhibits a large dependence on mesh resolution, and has a magnitude which rapidly decreases to zero as the mesh is refined. On the other hand, the stiffness generated using a vertex pressure mask (collocated linear sphere functional (see §6.1.2) with radius equal to the coarsest (level 1) mesh’s mean edge length) has substantially less mesh dependence, and quickly approaches a nonzero value.

6.2.2 Rigid Vertex Stiffness, \mathbf{K}^R

For rigid surfaces a finite force response may be defined using an isotropic stiffness matrix,

$$\mathbf{K}^R = k^{\text{Rigid}} \mathbf{I}_3 \in \mathbb{R}^{3 \times 3}, \quad k^{\text{Rigid}} > 0.$$

This is useful for defining responses at position constrained vertices of a deformable model,

$$\mathbf{K}_i = \mathbf{K}^R, \quad i \in \Lambda_u^0, \quad (33)$$

for at least two reasons. First, while it may seem physically ambiguous to consider contacting a constrained node of a deformable object, it does allow us to define a response for these vertices without introducing other simulation dependencies, e.g., how the haptic interaction with the elastic object support is modeled. Second, we shall see in §6.3 that defining stiffness responses at these nodes is important for determining contact responses on neighbouring triangles which are not rigid.

6.3 Surface Stiffness from Vertex Stiffnesses

Given the vertex stiffnesses, $\{K_i\}_{i=1}^n$, the stiffness of any location on the surface is defined using nodal interpolation

$$K(\mathbf{x}) = \sum_{i=1}^n \phi_i(\mathbf{x})K_i, \quad \mathbf{x} \in \Gamma, \quad (34)$$

so that $(K(\mathbf{x}))_{kl} \in \mathcal{L}$. Note that there are no more than three nonzero terms in the sum of (34), corresponding to the vertices of the face in contact. In this way, the surface stiffness may be continuously defined using only $|\Lambda_p^0|$ free boundary vertex stiffnesses and a single rigid stiffness parameter, k^{Rigid} , regardless of the extent of the masks. The global deformation is then visually rendered using the corresponding distributed traction constraints.

For a point-like displacement constraint \bar{u} applied at $\mathbf{x} \in \Gamma$ on a triangle having vertex indices $\{i_1, i_2, i_3\}$, the corresponding global solution is

$$\mathbf{v} = \sum_{i \in \{i_1, i_2, i_3\} \cap \Lambda_p^0} \zeta_i \phi_i(\mathbf{x})\mathbf{f}. \quad (35)$$

This may be interpreted as the combined effect of barycentrically distributed forces, $\phi_i(\mathbf{x})\mathbf{f}$, applied at each of the triangle's three masked vertex nodes, which is consistent with (38).

6.4 Rendering with Finite Stiffness Haptic Devices

Similar to haptic rendering of rigid objects, elastic objects with stiffnesses greater than some maximum renderable magnitude (due to hardware limitations) have forces displayed as softer materials during continuous contact. This can be achieved using a *haptic vertex stiffness*, K_i^H , which is proportional to the elastic vertex stiffness, K_i^E . While the stiffnesses could all be uniformly scaled on the free boundary, this can result in very soft regions if the model has a wide range of surface stiffness. Another approach is to set

$$K_i^H = \eta_i K_i^E \quad \text{where} \quad \eta_i = \min \left(1, \frac{\|K^R\|}{\|K_i^E\|} \right),$$

so that the elastic haptic model is never more stiff than a rigid haptic model. The surface's haptic stiffness $K^H(\mathbf{x})$ is then determined using (34), so that $\|K^H(\mathbf{x})\| \leq \|K^R\|, \forall \mathbf{x} \in \Gamma$.

In accordance with force reflecting contact, the deformed elastic state corresponds to the haptic force applied at the contact location \mathbf{x}^C . This produces geometric contact configurations similar to that shown in Figure 8, where the haptic displacement \mathbf{u}^H can differ from the elastic displacement \mathbf{u}^E . The geometric deformation is determined from the applied force \mathbf{f} and equation (35). Note that when the haptic and elastic stiffnesses are equal, such as for soft materials, then so are the elastic and haptic displacements. In all cases, the generalized “god object” [36] or “surface contact point” [33] is defined as the parametric image of \mathbf{x}^C on the deformed surface.

7 Experimental Results

GFs were precomputed using the boundary element method (BEM) with piecewise linear boundary elements. Table 2 provides timings for the BEM precomputation stages as well as the submillisecond cost of simulating point-like deformations using GFs. Further timings of CMA suboperations are shown in Table 3, and reflect interactive performance for modest numbers of constraint type changes, s . All timings were performed using unoptimized Java code on a single processor Pentium III, 450MHz, 256MB computer with




Figure 8: *Geometry of point-like contact*: The surface of the static/undeformed geometry (curved dashed line) and that of the deformed elastic model (curved solid line) are shown along with: applied force (\mathbf{f}), static contact location (\mathbf{x}^C), deformed elastic model contact location (\mathbf{x}^E), haptic probe-tip location (\mathbf{x}^H), haptic contact displacement ($\mathbf{u}^H = \mathbf{x}^H - \mathbf{x}^C$), elastic contact displacement ($\mathbf{u}^E = \mathbf{x}^E - \mathbf{x}^C$), static contact normal (\mathbf{n}^C) and elastic contact normal (\mathbf{n}^E). Once the contact is initiated by the collision detector, the sliding frictional contact can be tracked in surface coordinates at force feedback rates.

Sun's JDK 1.3 client JVM for Windows. These times can be significantly reduced by using hardware-optimized matrix libraries, and current computing hardware.


An application of the CMA for multiple distributed contacts with unilateral contact constraints was the grasping task illustrated in Figure 3 using the LEM from Figure 10(a). A short video clip is also available online [20].

Our current force feedback implementation is based on the point-like contact approach discussed in the previous section. Forces are rendered by a 3 dof PHANTOM™ haptic interface (model 1.0 Premium), on a dual Pentium II computer running Windows NT. The haptic simulation was implemented in C++, partly using the GHOST® toolkit, and interfaced to our ARTDEFO elastostatic object simulation written in Java™ and rendered with Java 3D™. The frictional contact problem is computed by the haptic servo loop at 1 kHz, which then prescribes boundary conditions for the slower graphical simulation running at 25–80 Hz. For a point-like contact, it was only necessary to perform collision detection on the undeformed model, so this was done using the GHOST® API. A photograph of the authors demonstrating the simulation is shown in Figure 9, and a number of screen shots for various models are shown in Figure 10. A short video clip is also available online [20].


We observed that the vertex masks were successful in producing noticeable improvements in the smoothness of the sliding contact force, especially when passing over regions with irregular triangulations (see Figure 7). We have not conducted a formal human study of the effectiveness of our simulation approach. However, the haptic simulation has been demonstrated to hundreds of users at two conferences: the 10th Annual PRECARN-IRIS (Institute for Robotics and Intelligent Systems) Conference (Montreal, Quebec, Canada, May 2000) and in the ACM SIGGRAPH 2000 Exhibition (New Orleans, Louisiana, USA, July 2000). Users reported that the simulation felt realistic. In general, the precomputed LEM approach was found to be both stable and robust.

8 Summary and Discussion

We have presented a detailed approach for real time solution of boundary value problems for discrete linear elastostatic models (LEM), regardless of discretization, using precomputed GFs in con-



(a) A simple nodular shape with a fixed base region.



(b) A kidney-shaped model with position constrained vertices on part of the occluded side.



(c) A plastic spatula with a position constrained handle.



(d) A seemingly gel-filled banana bicycle seat with matching metal supports.

Figure 10: *Screenshots from real time haptic simulations:* A wide range of ARTDEFO models are shown subjected to various displacements using the masked point-like contacts of §6. For each model, the middle of the three figures is uncontacted by the user's interaction point (a small green ball).

Elastic Model	# Vertices (n)	# Faces	Precomp (min)	LUD %	Simulate (ms)
Nodule	130v (89 free)	256f	1.1	1.1%	0.05
Kidney	322v (217 free)	640f	7.7	3.1%	0.13
Spatula	620v (559 free)	1248f	45	5.7%	0.34
Banana Seat	546v (245 free)	1088f	25	20.0%	0.15

Table 2: *GF precomputation and simulation times* for the BEM models depicted in Figure 10. All GFs corresponding to moveable free vertices (in Λ_p^0) were computed, and the precomputation time (Precomp) of the largest model is less than an hour. As is typical of BEM computations for models of modest size ($n < 1000$), the $\mathcal{O}(n^2)$ construction of the matrices (H and G in equation 8) is a significant portion of the computation, e.g., relative to the $\mathcal{O}(n^3)$ cost of performing the LU decomposition (LUD %) of the A matrix. The last column indicates that submillisecond graphics-loop computations (Simulate) are required to determine the point contact deformation response of each model’s free boundary.

# Updates, s	LU Factor (ms)	LU Solve (ms)	$(\Xi E)(E^T \bar{v})$ for $n = 100$ (ms)
10	0.54	0.03	0.38
20	2.7	0.15	0.74
40	19	0.58	1.7
100	310	5.7	5.7

Table 3: *Timings of CMA suboperations* such as LU decomposition (LU Factor) and back-substitution (LU Solve) of the capacitance matrix, as well as the weighted summation of s GFs (per 100 nodes) are shown for different sizes of updated nodal constraints, s .




Figure 9: : *Photograph of simulation in use*: Users were able to push, slide and pull on the surface of the model using a point-like manipulandum. Additionally, it was possible to change the surface friction coefficient, as well as the properties of the pressure mask, with noticeable consequences. The PHANTOM™ (here model 1.0 Premium) was used in all force feedback simulations, and is clearly visible in the foreground.

junction with capacitance matrix algorithms (CMAs). The data-driven CMA formulation highlights the special role of the capacitance matrix in computer haptics as a contact compliance useful for generating contact force and stiffness models, and provides a framework for extending the capabilities of these models.

Additionally, the important special case of point-like contact was addressed with special attention given to the consistent definition of contact forces for haptics. While this topic has been discussed before, we have introduced vertex masks to specify the distribution of contact forces in a way which leads to physically consistent force feedback models which avoid the numerical artifacts which lead to nonsmooth rendering of contact forces on discrete models, as well as ill-defined contacts in the continuum limit.

There are several issues to be addressed by future work on the simulation of LEMs for computer haptics.

One of the promises of linear GF models is that it should be possible to precompute and haptically touch large-scale models even if they are too large to be graphically rendered. However, the CMA presented here is very efficient for small models (small n) and limited constraints (small s), but further optimizations and required for precomputing, storing and simulating large-scale LEMs. Extending LEMs to accommodate geometric and material nonlinearities is also an area of study. Results addressing these problems will appear shortly in subsequent publications.

A key challenge for interactively rendering elastic models is the plausible approximation of friction in the presence of multiple distributed elastic-rigid and elastic-elastic contacts. While large contact areas are a potential problem for LEM haptics, i.e., due to large update costs, the accompanying collision detection and friction problems appear to be at least as difficult. Incorporation of LEMs into hybrid interactive dynamical simulations is also a relatively unexplored area.

Finally, the same issues (perceptible force regularity and spatial consistency) which motivated our approach for a single point-like contact, also arise for multiple point-like contacts and in general with multiple distributed contacts. The tight coupling of force stiffnesses between all contact zones, and therefore each (networked) force feedback device, can make this a difficult problem in practice.

A Justification of Interpolated Traction Distributions for Point Contact

This section derives the nodal boundary conditions associated with a localized point contact at an arbitrary mesh location. The practical consequence is that the discrete traction distribution may be conveniently interpolated from suitable nearby nodal distributions or masks.

Given a continuous surface traction distribution, $\mathbf{p}(\mathbf{x})$, a corresponding discrete distribution $\Phi(\mathbf{x})\mathbf{p}$ may be determined by a suitable projection into \mathcal{L} of each Cartesian component of $\mathbf{p}(\mathbf{x})$. For example, consider the projection of a scalar function on Γ defined as the minimizer of the scalar functional $\mathcal{E} : \mathbb{R}^{3n} \mapsto \mathbb{R}$,

$$\mathcal{E}(\mathbf{p}) = \int_{\Gamma} [\|\mathbf{p}(\mathbf{x}) - \Phi(\mathbf{x})\mathbf{p}\|_2^2 + \|B\Phi(\mathbf{x})\mathbf{p}\|_2^2] d\Gamma_{\mathbf{x}},$$

where $B : \mathcal{L} \mapsto \mathbb{R}$ is some linear operator that can be used, e.g., to

penalize nonsmooth functions, and $\Phi(\mathbf{x}) : \mathbb{R}^{3n} \mapsto \mathbb{R}^3$ is a nodal interpolation matrix defined on the surface,

$$\Phi(\mathbf{x}) = [\phi_1(\mathbf{x})\phi_2(\mathbf{x}) \cdots \phi_n(\mathbf{x})] \otimes I_3, \quad \mathbf{x} \in \Gamma,$$

where I_3 is the 3-by-3 identity matrix. The Euler-Lagrange equations for this minimization are

$$\begin{aligned} \sum_{j=1}^n \left(\int_{\Gamma} [\phi_i(\mathbf{x})\phi_j(\mathbf{x}) + (B\phi_i(\mathbf{x})) (B\phi_j(\mathbf{x}))] d\Gamma_{\mathbf{x}} \right) \mathbf{p}_j \\ = \int_{\Gamma} \phi_i(\mathbf{x})\mathbf{p}(\mathbf{x})d\Gamma_{\mathbf{x}} \quad i = 1, 2, \dots, n, \end{aligned}$$

which, in an obvious notation, is written as the linear matrix problem

$$\mathcal{A}\mathbf{p} = \mathbf{f} \quad (36)$$

to be solved for the nodal traction values \mathbf{p} . Note that \mathcal{A} has units of area.

The relevant traction distribution for point-like contact is a scale-independent concentrated point load

$$\mathbf{p}(\mathbf{x}) = \mathbf{p}^{\delta}(\mathbf{x}) = \mathbf{f}^{\delta} \delta(\mathbf{x} - \mathbf{x}^{\delta})$$

which models a force \mathbf{f}^{δ} delivered at $\mathbf{x}^{\delta} \in \Gamma$. The force n -vector in equation (36) has components

$$f_i = \phi_i(\mathbf{x}^{\delta})\mathbf{f}^{\delta}$$

and the corresponding pressure distribution's nodal values are

$$\mathbf{p} = \mathcal{A}^{-1}\mathbf{f}.$$

For compactly supported basis functions, $\phi_i(\mathbf{x})$, \mathbf{f} has only a small number of nonzero components for any given \mathbf{x} . Hence $\phi_i(\mathbf{x}^{\delta})$ are the interpolation weights describing the contribution of the nearby nodal pressure distributions, here specified by the columns of \mathcal{A}^{-1} .

As an example, consider the important case where \mathcal{L} is a continuous piecewise linear function space with $\phi_i(\mathbf{x}_j) = \delta_{ij}$. This was the space used in our implementation. In this case, at most only three components of \mathbf{f} are nonzero, given by the indices $\{i_1, i_2, i_3\}$ which correspond to vertices of the contacted triangle τ^{δ} , i.e., for which $\mathbf{x}^{\delta} \in \tau^{\delta}$. The values $\phi_i(\mathbf{x}^{\delta})$ are the barycentric coordinates of \mathbf{x}^{δ} in τ^{δ} . The pressure distribution's nodal values are then

$$\mathbf{p} = \mathcal{A}^{-1}\mathbf{f} \quad (37)$$

$$\begin{aligned} &= \sum_{k=1}^3 (\mathcal{A}^{-1})_{:i_k} f_{i_k} \\ &= \sum_{k=1}^3 \phi_{i_k}(\mathbf{x}^{\delta}) \left[(\mathcal{A}^{-1})_{:i_k} \mathbf{f}^{\delta} \right] \\ &= \sum_{k=1}^3 \phi_{i_k}(\mathbf{x}^{\delta}) \mathbf{p}^{(i_k)}, \end{aligned} \quad (38)$$

where $\mathbf{p}^{(i_k)}$ is the pressure distribution corresponding to the application of the load directly to vertex i_k , and $(\cdot)_{:i_k}$ refers to block column i_k of the matrix. Therefore the *piecewise linear pressure distribution for a point load applied at a barycentric location on a triangle is equal to the barycentric average of the pressure distributions associated with the point load applied at each of the triangle's vertices*. This may be recognized as an elastic generalization of *force shading* [27] for rigid models.

Note that the j^{th} column of \mathcal{A}^{-1} is a vertex mask that describes the nodal distribution of the load applied to the j^{th} vertex. Modifying the penalty operator B results in masks with varying degrees of smoothness and spatial localization.

References

- [1] Oliver Astley and Vincent Hayward. Multirate haptic simulation achieved by coupling finite element meshes through norton equivalents. In *Proceedings of the IEEE International Conference on Robotics and Automation*, 1998.
- [2] Remis Balaniuk. Building a haptic interface based on a buffer model. In *Proceedings of the IEEE International Conference on Robotics and Automation*, 2000.
- [3] James R. Barber. *Elasticity*. Kluwer Press, Dordrecht, first edition, 1992.
- [4] C. A. Brebbia, J. C. F. Telles, and L. C. Wrobel. *Boundary Element Techniques: Theory and Applications in Engineering*. Springer-Verlag, New York, second edition, 1984.
- [5] Morten Bro-Nielsen. Surgery simulation using fast finite elements. *Lecture Notes in Computer Science*, 1131:529–, 1996.
- [6] Morten Bro-Nielsen and Stephane Cotin. Real-time volumetric deformable models for surgery simulation using finite elements and condensation. *Computer Graphics Forum*, 15(3):57–66, August 1996.
- [7] Murak Cenk Çavuşoğlu and Frank Tendick. Multirate simulation for high fidelity haptic interaction with deformable objects in virtual environments. In *Proceedings of the IEEE International Conference on Robotics and Automation*, San Francisco, USA, 2000.
- [8] S. Cotin, H. Delingette, and N. Ayache. Realtime elastic deformations of soft tissues for surgery simulation. *IEEE Transactions On Visualization and Computer Graphics*, 5(1):62–73, 1999.
- [9] Stephane Cotin, Herve Delingette, and Nicholas Ayache. Efficient linear elastic models of soft tissues for real-time surgery simulation. Technical Report RR-3510, Inria, Institut National de Recherche en Informatique et en Automatique, 1998.
- [10] Diego d'Aulignac, Remis Balaniuk, and Christian Laugier. A haptic interface for a virtual exam of a human thigh. In *Proceedings of the IEEE International Conference on Robotics and Automation*, San Francisco, USA, 2000.
- [11] Gilles Debunne, Mathieu Desbrun, Alan Barr, and Marie-Paule Cani. Interactive multiresolution animation of deformable models. In N. Magnenat-Thalmann and D. Thalmann, editors, *Computer Animation and Simulation '99*, SpringerComputerScience, pages 133–144. Springer-Verlag Wien New York, 1999. Proceedings of the Eurographics Workshop in Milano, Italy, September 7–8, 1999.
- [12] Y. Ezawa and N. Okamoto. High-speed boundary element contact stress analysis using a super computer. In *Proc. of the 4th International Conference on Boundary Element Technology*, pages 405–416, 1989.
- [13] Sarah Gibson, Christina Fyock, Eric Grimson, Takeo Kanade, Rob Kikinis, Hugh Lauer, Neil McKenzie, Andrew Mor, Shin Nakajima, Hide Ohkami, Randy Osborne, Joseph Samosky, and Akira Sawada. Volumetric object modeling for surgical simulation. *Medical Image Analysis*, 2(2):121–132, 1998.
- [14] Sarah F. Gibson and Brian Mirtich. A survey of deformable models in computer graphics. Technical Report TR-97-19, Mitsubishi Electric Research Laboratories, Cambridge, MA, November 1997.

- [15] Gene H. Golub and Charles F. Van Loan. *Matrix Computations*. Johns Hopkins University Press, Baltimore and London, third edition, 1996.
- [16] William W. Hager. Updating the inverse of a matrix. *SIAM Review*, 31(2):221–239, June 1989.
- [17] K. V. Hansen and O. V. Larsen. Using region-of-interest based finite element modeling for brain-surgery simulation. *Lecture Notes in Computer Science*, 1496:305–, 1998.
- [18] Friedel Hartmann. *The mathematical foundation of structural mechanics*. Springer-Verlag, New York, 1985.
- [19] C.-H. Ho, C. Basdogan, and M. A. Srinivasan. Efficient point-based rendering techniques for haptic display of virtual objects. *Presence*, 8(5):477–491, 1999.
- [20] Doug L. James. <http://www.cs.ubc.ca/~djames/deformable>.
- [21] Doug L. James and Dinesh K. Pai. ARTDEFO: Accurate Real Time Deformable Objects. *Computer Graphics*, 33(Annual Conference Series):65–72, 1999.
- [22] A. M. Abu Kassim and B. H. V. Topping. Static reanalysis: a review. *Journal of Structural Engineering*, 113:1029–1045, 1987.
- [23] U. Kühnapfel, H.K. Çakmak, and H. Maaß. 3d modeling for endoscopic surgery. In *Proceedings of IEEE Symposium on Simulation*, pages 22–32, Delft University, Delft, NL, October 1999.
- [24] Charles Loop. Smooth subdivision surfaces based on triangles. Master’s thesis, University of Utah, Department of Mathematics, 1987.
- [25] K. W. Man, M. H. Aliabadi, and D. P. Rooke. Analysis of Contact Friction using the Boundary Element Method. In M. H. Aliabadi and C. A. Brebbia, editors, *Computational Methods in Contact Mechanics*, chapter 1, pages 1–60. Computational Mechanics Publications and Elsevier Applied Science, 1993.
- [26] T. H. Massie and J. K. Salisbury. The phantom haptic interface: A device for probing virtual objects. In *ASME Winter Annual Meeting, Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, Chicago, IL, Nov. 1994.
- [27] Hugh B. Morgenbesser and Mandayam A. Srinivasan. Force shading for haptic shape perception. In *Proceedings of the ASME Dynamics Systems and Control Division*, volume 58, 1996.
- [28] Dinesh K. Pai, Kees van den Doel, Doug L. James, Jochen Lang, John E. Lloyd, Joshua L. Richmond, and Som H. Yau. Scanning Physical Interaction Behavior of 3D Objects. In *Computer Graphics Proceedings (SIGGRAPH 2001)*. ACM Siggraph, 2001.
- [29] G. Picinbono, J. C. Lombardo, H. Delingette, and N. Ayache. Anisotropic elasticity and force extrapolation to improve realism of surgery simulation. In *Proceedings of IEEE International Conference on Robotics and Automation*, San Francisco, USA, 2000.
- [30] William H. Press, Brian P. Flannery, Saul A. Teukolsky, and William T. Vetterling. *Numerical Recipes: The Art of Scientific Computing*, chapter Sherman-Morrison and Woodbury, pages 66–70. Cambridge University Press, Cambridge, 1987.
- [31] R. Ramanathan and D. Metaxas. Dynamic deformable models for enhanced haptic rendering in virtual environments. In *IEEE Virtual Reality Conference*, 2000.
- [32] P. Schröder, D. Zorin, T. DeRose, D. R. Forsey, L. Kobbelt, M. Lounsbery, and J. Peters. Subdivision for modeling and animation. *SIGGRAPH 99 Course Notes*, August 1999.
- [33] Sensable Technologies, Inc. GHOST SDK, <http://www.sensable.com>.
- [34] Yan Zhuang and John Canny. Haptic interaction with global deformations. In *Proceedings of the IEEE International Conference on Robotics and Automation*, 2000.
- [35] O. C. Zienkiewicz. *The Finite Element Method*. McGraw-Hill Book Company (UK) Limited, Maidenhead, Berkshire, England, 1977.
- [36] C. B. Zilles and J. K. Salisbury. A constraint-based god-object method for haptic display. In *ASME Haptic Interfaces for Virtual Environment and Teleoperator Systems*, volume 1, pages 149–150, Chicago, IL (US), 1994.