Computers & Graphics 45 (2014) 75–85

Contents lists available at ScienceDirect

Computers & Graphics

journal homepage: www.elsevier.com/locate/cag

Special Section on Graphics Interaction

Velocity-based adaptivity of deformable models

Maxime Tournier [a],*, Matthieu Nesme [b], François Faure [b], Benjamin Gilles [a]

[a] INRIA, LIRMM-CNRS, Université de Montpellier 2, France
[b] INRIA, LJK-CNRS, Université de Grenoble, France

ARTICLE INFO

Article history:
Received 27 June 2014
Received in revised form
27 August 2014
Accepted 27 August 2014
Available online 8 September 2014

Keywords:
Computer animation
Physically based animation
Deformable solids
Adaptive kinematics

ABSTRACT

A new adaptive model for viscoelastic solids is presented. Unlike previous approaches, it allows seamless transitions, and simplifications in deformed states. The deformation field is generated by a set of physically animated frames. Starting from a fine set of frames and mechanical energy integration points, the model can be coarsened by attaching frames to others, and merging integration points. Since frames can be attached in arbitrary relative positions, simplifications can occur seamlessly in deformed states, without returning to the original shape, which can be recovered later after refinement. We propose a new class of velocity-based simplification criterion based on relative velocities. Integration points can be merged to reduce the computation time even more, and we show how to maintain continuous elastic forces through the levels of detail. Such meshless adaptivity allows significant improvements of computation time during simulations. It also provides a natural approach to coarse-to-fine deformable mesh registration.

The stunning quality of high-resolution physically based animations of deformable solids requires complex deformable models with large numbers of independent Degrees Of Freedom (DOF) which result in large equation systems for solving dynamics, and high computation times. On the other hand, the thrilling user experience provided by interactive simulations can only be achieved using fast computation times which preclude the use of high-resolution models. Reconciling these two contradictory goals requires adaptive models to efficiently manage the number of DOFs, by refining the model where necessary and by coarsening it where possible. Mesh-based deformations can be seamlessly refined by subdividing elements and interpolating new nodes within these. However, seamless coarsening can be performed only when the fine nodes are back to their original position with respect to their higher-level elements, which only happens in the locally undeformed configurations (*i.e.*with null strain). Otherwise, a popping artifact (*i.e.*an instantaneous change of shape) occurs, which not only violates the laws of physics, but it is also visually disturbing for the user. Simplifying objects in deformed configurations, as presented in Fig. 1(c), has thus not been possible with previous adaptive approaches, unless the elements are small or far enough from the user. This may explain why extreme coarsening has rarely been proposed, and adaptive FEM models typically range from moderate to high complexity.

We introduce a new approach of adaptivity to mechanically simplify objects in arbitrarily deformed configurations, while exactly maintaining their current shape and controlling the velocity discontinuity, which we call seamless adaptivity. It extends a frame-based meshless method and naturally exploits the ability to attach frames to others in arbitrary relative positions, as illustrated in Fig. 2. In this example, a straight beam is initially animated using a single moving frame, while another control frame is attached to it. We then detach the child frame to allow the beam to bend as needed. If the beam deformation reaches a steady state, the *velocity* field can again be obtained from the moving frame alone, and the *shape* can be frozen in the deformed state by applying an offset to the child frame reference position relative to the moving frame. Setting the offset to the current relative position removes mechanical DOFs without altering the current shape of the object. This deformation is reversible. If the external loading applied to the object changes, we can mechanically refine the model again (*i.e.*activate the passive frame) to allow the object to recover its initial shape or to undergo new deformations. The ability to dynamically adapt the deformation field even in non-rest configuration is the specific feature of our approach, which dramatically enhances the opportunities for coarsening mechanical models compared with previous methods.

Our specific contributions are (1) a deformation method based on a generalized frame hierarchy for dynamically tuning the complexity of deformable solids with seamless transitions; (2) a novel simplification and refinement criterion based on velocity, which allows us to simplify the deformation model in deformed configurations, and (3) a method to dynamically adapt the integration points and enforce the continuity of forces across changes of resolution.

---

* Corresponding author.

http://dx.doi.org/10.1016/j.cag.2014.08.004
0097-8493/© 2014 Elsevier Ltd. All rights reserved.

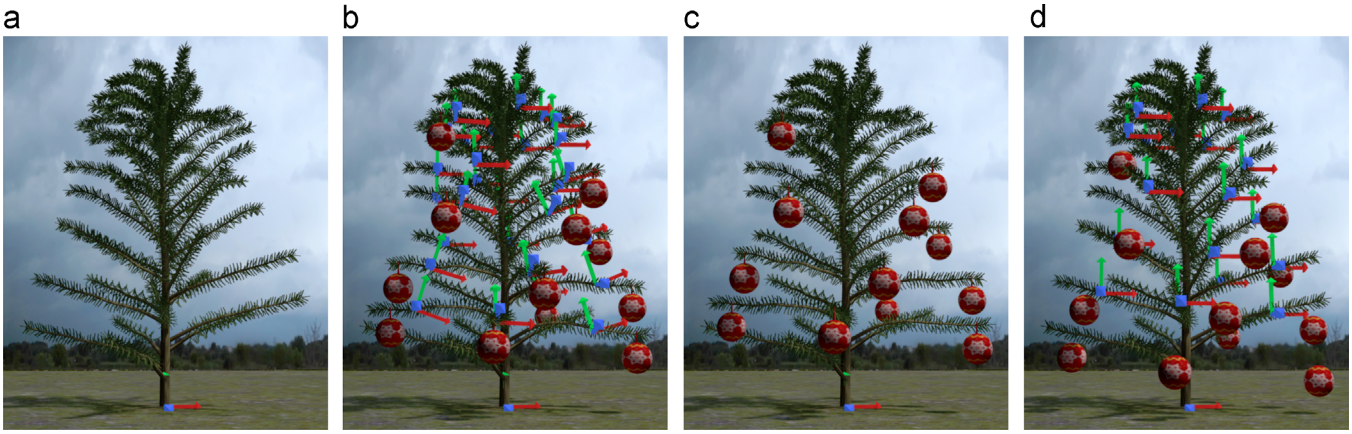

Fig. 1. Deformable Christmas tree with our proposed adaptive deformation field. (a) One frame is sufficient in steady state. (b) When ornaments are attached, additional frames are activated to allow deformations. (c): The velocity field can be simplified again when the equilibrium is reached. Note that our method can simplify locally deformed regions. (d) Once the branches are released, the velocity field is refined again to allow the branches to recover their initial shape.

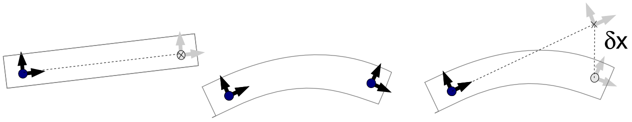

Fig. 2. Seamless coarsening in a deformed state. Left: Reference shape, one active frame in black, and a passive frame in grey attached using a relative transformation (dotted line). Middle: Activating the frame let it to move freely and deform the object. Right: Deactivated frame in a deformed configuration using an offset $\delta \mathbf{X}$.

The present article extends an earlier conference version [1] by adding new results on deformable mesh registration (6.4), more derivations concerning metrics (4.1), and providing more details based on reviewer comments. The remainder is organized as follows. We summarize the original frame-based simulation method and introduce notations in Section 2. An overview of our adaptive framework is presented in Section 3. We formalize and discuss different criteria for nodal adaptivity in Section 4. The adaptivity of the integration points is then introduced in Section 5. Results obtained with our method are presented and discussed in Section 6, including an application to deformable mesh registration, and we conclude in Section 7 with future work.

1. Related work

The simulation of viscoelastic solids is a well-studied problem in computer graphics, starting with the early work of Terzopoulos et al. [2]. A survey can be found in [3]. Frame-based models have been proposed [4–7], and the impressive efficiency of precomputed reduced models has raised a growing interest [8–13], but run-time adaptivity remains a challenge. The remainder of this review focuses on this issue.

Hutchinson et al. [14] and Ganovelli et al. [15] first combined several resolutions of 2D and 3D solids dynamically deformed by mass-springs. Cotin et al. [16] combined two mechanical models to simulate various parts of the same object. Most adaptive methods are based on meshes at multiple resolutions. Mixing different mesh sizes can result in T-nodes that are mechanically complex to manage in the Finite Element Method (FEM). Wu et al. [17] chose a decomposition scheme that does not generate such nodes. Debunne et al. [18] performed the local explicit integration of non-nested meshes. Grinspun et al. [19] showed that hierarchical shape functions are a generic way to deal with T-nodes. Sifakis et al. [20] constrained T-nodes within other independent nodes. Martin et al. [21] solved multi-resolution junctions with polyhedral elements. Several authors proposed to generate on the fly a valid mesh with dense and fine zones. Real-time remeshing is feasible for 1D elements such as rods and wires [22–24] or 2D surfaces like cloth [25]. For 3D models, it is an elegant way to deal with cuttings, viscous effects and very thin features [26–28]. A mesh-less, octree-based adaptive extension of shape matching has been proposed [29]. Besides all these methods based on multiple resolutions, Kim and James [30] take a more algebraic approach, where the displacement field is decomposed on a small, dynamically updated, basis of orthogonal vectors, while a small set of carefully chosen integration points are used to compute the forces. In contrast to these works, our method relies on velocity field analysis and a meshless discretization.

Numerous error estimators for refinement have been proposed in conventional FEM analysis. For static analysis, they are generally based on a precomputed stress field. This is not feasible in real-time applications, where the current configuration and corresponding stress must be used. Wu et al. [17] proposed four criteria based on the curvature of the stress, strain or displacement fields. Debunne et al. [18] considered the Laplacian of the displacement. Lenoir et al. [22] refined parts in contact for wire simulation. These approaches refine the objects where they are the most deformed, and they are not able to save computation time in equilibrium states different from the rest state. The problems relative to the criterion thresholds are rarely discussed, even though potential popping artifacts can be problematic: the smaller the thresholds, the smaller the popping artifacts, but also the more difficult to simplify and thus the less efficient.

While our adaptive scheme is primarily targeted at physically based animation, it can also be interesting to improve the robustness of deformable mesh registration schemes. Finding correspondences between a source (template) mesh and a target mesh or point cloud is a fundamental task in shape acquisition [31] and analysis [32]. As reviewed in [33], local or global correspondence search is generally regularized using a deformation method, that constrains the displacement of the template to a set of feasible transformations. Iterative closest point (ICP) algorithm [34] is the most common procedure to align a source mesh to a target mesh. At each iteration, source point correspondences are locally found by an optimized closest point search [35]. In the original ICP algorithm, the best global linear transformation is found by minimizing distances between source points and their corresponding points. Instead, elastic ICP [36,37] can be easily performed by treating distance gradients as external forces (*i.e.* springs) applied to a given deformable model. Deformable