UC Berkeley

UC Berkeley Previously Published Works

Title

View-Dependent Adaptive Cloth Simulation with Buckling Compensation.

Permalink

<https://escholarship.org/uc/item/79h5740c>

Journal

IEEE transactions on visualization and computer graphics, 21(10)

ISSN

1077-2626

Authors

Koh, Woojong
Narain, Rahul
O'Brien, James F

Publication Date

2015-10-01

DOI

10.1109/tvcg.2015.2446482

Supplemental Material

<https://escholarship.org/uc/item/79h5740c#supplemental>

Peer reviewed

View-Dependent Adaptive Cloth Simulation with Buckling Compensation

Woojong Koh, Rahul Narain, and James F. O'Brien

Abstract—This paper describes a method for view-dependent cloth simulation using dynamically adaptive mesh refinement and coarsening. Given a prescribed camera motion, the method adjusts the criteria controlling refinement to account for visibility and apparent size in the camera's view. Objectionable dynamic artifacts are avoided by anticipative refinement and smoothed coarsening, while locking in extremely coarsened regions is inhibited by modifying the material model to compensate for unresolved sub-element buckling. This approach preserves the appearance of detailed cloth throughout the animation while avoiding the wasted effort of simulating details that would not be discernible to the viewer. The computational savings realized by this method increase as scene complexity grows. The approach produces a $2\times$ speed-up for a single character and more than $4\times$ for a small group as compared to view-independent adaptive simulations, and respectively $5\times$ and $9\times$ speed-ups as compared to non-adaptive simulations.

Index Terms—Physically based modeling, animation, cloth simulation, adaptive remeshing

1 INTRODUCTION

CLOTH simulation for visual effects has reached a mature state where the use of virtual characters wearing simulated clothing is now widespread. However, cloth simulation remains computationally expensive, particularly when films require high-quality, realistic results computed at high resolution. For characters that are far from the camera, or otherwise less visible in a shot, fine details will not be visible to the viewer and work spent on computing those details is wasted. In most film production settings, both the camera and character motions are known before the simulations are run, and one could use this information to substitute cheaper low-resolution simulations on distant or out-of-frame characters. Unfortunately, manually swapping some characters to low-resolution simulations is cumbersome, particularly when a single character's clothing requires high resolution for some parts of a shot but not others, yet the cloth motion must nevertheless appear coherent throughout. For closeup shots where only part of a character is in frame, savings could also be realized by reducing the computation devoted to out-of-frame cloth, so long as such savings do not result in dynamic artifacts that affect visible parts of the cloth.

In this paper we describe a method for *view-dependent simulation* using dynamically adaptive mesh refinement and coarsening. Instead of using a fixed-resolution simulation mesh, the mesh undergoes local adaptation based on the geometric and dynamic details of the simulated cloth. The degree to which this detail is resolved is adjusted locally based on the view of the simulated cloth, as shown in Fig. 1. Areas that appear large in the camera will be refined to

Fig. 1. The clothing on this character uses view-dependent simulation. Visible areas (blue) are simulated at high resolution, as measured in screen-space. Areas not visible to the camera (red) are simulated at reduced resolution. The first two images show views from the camera's perspective. The rightmost image shows an outside perspective with the camera's view frustum drawn in black wireframe.

show finely detailed dynamics. Areas that are out of frame, facing away from the camera, or at a distance will be correspondingly coarser.

The goal of this work is to preserve the appearance of detailed simulation throughout the animation while avoiding the wasted effort of simulating details that will not be apparent to the viewer. Further, there should be no visible dynamic artifacts created due to varying refinement as the camera and objects move about. Finally, cloth that leaves and reenters visibility should appear to have coherent and consistent dynamics.

Our work builds on the publicly available ARCSim framework (<http://graphics.berkeley.edu/resources/ARCSim>) which can be used to animate sheets of deformable materials such as cloth, paper, plastic, and metal. ARCSim adaptively refines and coarsens anisotropic triangle meshes to efficiently resolve the geometric and dynamic details of the simulated objects. Our method modifies the metrics used by ARCSim so that local mesh visibility is accounted for during

• W. Koh and J.F. O'Brien are with the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, CA 94720-1776. E-mail: {wjkoh,job}@berkeley.edu.

• R. Narain is with the Department of Computer Science and Engineering, University of Minnesota, Minneapolis, MN. E-mail: narain@umn.edu.

Manuscript received 4 Nov. 2014; revised 14 June 2015; accepted 4 May 2015.
Date of publication 0 . 0000; date of current version 0 . 0000.

Recommended for acceptance by E. Sifakis and V. Koltun.

Digital Object Identifier no. 10.1109/TVCG.2015.2446482

refinement. Given the existing framework for adaptivity, our view-dependent refinement is easy to implement and has negligible overhead.

In cases where only a single character is being modeled, our approach realizes modest savings of roughly a $2.4\times$ speed-up in comparison with ARCSim's default adaptive simulation. These savings are due to coarsening out-of-view and back-facing parts of the character's clothing. For small crowd scenes, the savings are larger, up to $4.5\times$, as background and out-of-view characters are automatically coarsened. Compared to a non-adaptive simulation, the speed-up is more than $9\times$. For massive crowd scenes with thousands of agents, we expect that even greater savings could be realized with our approach.

2 RELATED WORK

Adaptive discretizations have been found to give significant performance and scalability benefits for a number of computationally intensive simulation tasks. In fluid simulation, detailed liquid surfaces can be animated efficiently by refining the spatial resolution near the surface using octrees [1], adaptively sampled particles [2], tall-cell grids [3], or tetrahedral meshes [4], [5], [6]. Adaptive refinement and simplification techniques have also been proposed for mass-spring systems [7], articulated bodies [8] and finite element models [9]. Most relevant to our work are techniques for adaptive cloth simulation [10], [11], [12], [13], [14], which use remeshing to resolve detailed wrinkles and folds. The approach of Narain et al. [14] has also been extended to efficiently model plastic deformation and sharp creases [15] as well as complex fracture patterns [16]. However, all those techniques are view-independent and element sizing is controlled only by geometrical and dynamical properties of the simulated system.

For animation applications, a number of techniques have also been proposed that take into account the current viewpoint and attempt to expend less computational effort in regions that are visually less important. One approach, known as simulation level of detail, is to switch between dynamical models of varying degrees of simplification, depending on the viewpoint. Carlson and Hodgins [17] introduced such a method for real-time animation of large groups of legged creatures. Their approach required the user to manually design the simplified models for each level of detail. Subsequent work has sought to generate approximated models automatically, for example for particle systems [18], plants [19], and hair [20]. Chenney and Forsyth [21] took a different approach, culling out-of-view objects by not solving their equations of motion, and using knowledge of the statistics of the motion to predict their plausible states when they reenter the view.

Alternatively, one can modify the simulation resolution based on the viewpoint without changing the underlying model. Barran [22] performed fluid simulation on a non-Cartesian grid based on cylindrical coordinates centered at the viewer, thus directly taking the distance from the viewer into account in the discretization. Viewpoint information has also been used to vary simulation resolution in traditional adaptive discretizations for fluids, such as octrees [23], [24] and adaptive particles [25].

In our work, we take inspiration from geometric level of detail techniques, such as those for real-time rendering of terrain [26], [27] or complex scenes like architectural walkthroughs [28] and polygonal CAD models [29]. These techniques inform our understanding of the important view-dependent criteria for representing geometrical details. Hoppe [30] used surface orientation and screen-space geometric error as refinement criteria. Xia et al. [31] further propose the use of local illumination gradients, projected lengths of edges in screen space, and silhouette boundaries.

3 METHODS

Our view-dependent adaptive remeshing scheme builds on ARCSim, the adaptive anisotropic remeshing framework described by Narain et al. [14]. We introduce a new view-dependent refinement strategy that complements their use of dynamical and geometrical refinement criteria. This approach realizes significant speed improvements by extending the domain of adaptive remeshing to include *perceptual* properties as well as physical ones.

The method of Narain et al. defines a *sizing field* that specifies the desired spatially varying resolution of the simulation mesh, taking into account various geometric and dynamic criteria such as local curvature, velocity gradient, compressive strain, and obstacle proximity. It is represented as a 2×2 symmetric tensor field $\mathbf{M}$ which is first computed on faces and then transferred to vertices via area-weighted averaging. Once the sizing field is defined, an edge between vertices i and j is considered valid if its size with respect to $\mathbf{M}$,

$$s(i, j)^2 = \mathbf{u}_{ij}^T \left(\frac{\mathbf{M}_i + \mathbf{M}_j}{2} \right) \mathbf{u}_{ij}, \quad (1)$$

does not exceed 1, where $\mathbf{u}_{ij} = \mathbf{u}_i - \mathbf{u}_j$ is the vector between the two vertices in material space. If $s(i, j)^2 > 1$, the edge is deemed invalid and must be split. The remeshing algorithm proceeds by splitting all invalid edges, collapsing as many edges as possible without introducing new invalid edges, and flipping edges to maintain an anisotropically Delaunay triangulation. This procedure produces a mesh that is as coarse as possible while containing no invalid edges and remaining Delaunay in the non-Euclidean space of the metric.

We modify their algorithm so that, rather than using purely physical and geometrical criteria to determine the mesh resolution, we vary the desired degree of refinement over space and time based on visual importance relative to a specified camera motion. We implement this variation by modifying the sizing field $\mathbf{M}$ so that the resolution of the mesh is no more than what is needed to resolve visually important features. This modification reduces computational effort in regions that are less visible from the camera, bringing about a more efficient simulation without losing significant visual details.

Our implementation considers two visibility criteria. In regions that are not visible from the current camera position, that is, those that are out of frame or facing away from the camera, we uniformly coarsen the mesh by

Fig. 2. An example of spatial smoothing. Cyan edges indicate out-of-view faces under spatial smoothing, and green edges indicate out-of-view faces far from the view frustum.

scaling the sizing field with constants. In regions that are visible, we control the sizes of elements in terms of their projected lengths in screen space so that distant or foreshortened elements are coarser. This approach is roughly equivalent to adaptivity based on screen-space metrics. However, to avoid artifacts that would occur due to fast-moving cameras or cuts between views, our algorithm applies conservative spatial and temporal smoothing to the sizing field.

3.1 Coarsening of non-visible regions

For non-visible regions, we seek to uniformly coarsen the mesh relative to the original view-independent sizing field. To do so, we define a scalar $\nu_i \leq 1$ for every face i , which we call the *view factor*, and modify the per-face sizing field $\mathbf{M}$ before transferring to vertices as

$$\tilde{\mathbf{M}}_i = \nu_i^2 \mathbf{M}_i. \quad (2)$$

As the sizing criterion (1) is quadratic, this scaling increases the target edge lengths by a factor of ν_i^{-1} .

In general, we would like to use high resolution (ν_{front}) for faces that are visible in the current view, and low resolution for back-facing and out-of-frame faces based on user-specified parameters ν_{back} and ν_{out} both of which are less than ν_{front} . However, simply defining ν in a piecewise-constant fashion in both space and time causes severe artifacts because of the discontinuous change in the sizing field. First, the discontinuity in sizes at the boundary between in-frame and out-of-frame faces leads to noticeable motion artifacts such as popping due to the influence of spurious forces from the out-of-frame region. Second, rapid camera movements and jump cuts can cause previously coarse regions with inaccurate geometry to suddenly become visible. To eliminate these discontinuities and obtain artifact-free results, we define the view factor in a way that is continuous over both space and time.

3.1.1 Spatial smoothing

Instead of using a spatially discontinuous field, we enforce a continuous falloff of the view factor between in-frame and out-of-frame faces. For a given mesh face i , let d_i be its distance from the view frustum in world space. We define the

spatially smoothed view factor $\tilde{\nu}_i$ by linearly interpolating to ν_{out} over a user-specified margin length m :

$$\tilde{\nu}_i = \begin{cases} \nu_{\text{fb}} & \text{if } d_i = 0, \\ \nu_{\text{fb}} - \frac{d_i}{m}(\nu_{\text{fb}} - \nu_{\text{out}}) & \text{if } 0 < d_i < m, \\ \nu_{\text{out}} & \text{if } d_i \geq m, \end{cases} \quad (3)$$

where ν_{fb} is ν_{front} or ν_{back} depending on the direction of the face normal. Thus, we have $\tilde{\nu} = \nu_{\text{front}}$ or ν_{back} for in-frame faces and $\tilde{\nu} = \nu_{\text{out}}$ for out-of-frame faces far from the view frustum, with a continuous transition region in between, as Fig. 2 shows.

There is still a discontinuity on the boundary between front-facing faces and backward-facing faces. While it is tempting to use the direction of the face normals to create a smooth transition, we find that normals can vary too rapidly across silhouettes to offer any useful smoothing that way. Instead, we address this discontinuity with a different approach, described later in Section 3.3.

3.1.2 Temporal smoothing and anticipation

We use temporal smoothing to avoid visibly discontinuous changes in mesh size due to camera motion, which may cause noticeable popping artifacts. We include anticipation that ensures the cloth gets refined to sufficient resolution *before* it appears in the frame, preventing undesirable transients.

For any given face, we can treat the view factor before temporal smoothing $\tilde{\nu}$ as a function of time, holding the face fixed and considering the prescribed motion of the camera. We smooth $\tilde{\nu}$ over a time interval $[t, t + T]$ based on the current time t as follows. Define a temporal window function $w(\tau)$ which satisfies $w(0) = 1$ and falls off to zero at $\tau = T$. The temporally smoothed view factor is

$$\nu(t) = \max_{\tau \in [0, T]} w(\tau) \tilde{\nu}(t + \tau). \quad (4)$$

This is analogous to dilation by a non-flat structuring element in mathematical morphology. In our implementation, we use $w(\tau) = 1 - \tau/T$.

Unlike smoothing by, say, moving averages, our approach is conservative in that $\nu(t) \geq \tilde{\nu}(t)$; in particular, visible regions always have $\nu = \nu_{\text{front}}$. Further, in the presence of discontinuities such as jump cuts, ν increases continuously from ν_{out} to ν_{front} or ν_{back} over a time period T in advance of the jump. This anticipatory refinement allows enough time for the system to settle into a feasible high-resolution state before becoming visible, as shown in Fig. 3.

3.2 Screen-space resolution of visible regions

In non-visible regions, it is sufficient to uniformly coarsen the mesh resolution as above. However, for visible regions, we wish to preserve the geometrical and dynamical detail resolved by the original sizing field as much as possible, only coarsening when such detail would not be visually important, as illustrated in Fig. 4.

The modified sizing field $\tilde{\mathbf{M}}$ gives the mesh resolution needed to accurately capture the dynamics of all cloth features. From the camera's perspective, however, wrinkles and folds that are very distant or foreshortened will appear

Fig. 3. A sequence of frames with a jump cut, showing a dress simulated with temporal anticipation (above), and without (below). Blue edges indicate faces visible in the current view. Temporal anticipation allows the cloth to reach the correct high-resolution state before it becomes visible.

extremely small in the final image. Such features are not likely to affect the visual appearance of the cloth and need not be represented accurately; they can simply be coarsened away. On the other hand, features that appear large in the image should remain accurately represented at the resolution determined by the original sizing field.

In the previous work of Narain et al. [14], the allowed edge lengths were constrained to a range $[\ell_{\min}, \ell_{\max}]$ in material space by clamping the eigenvalues of the sizing tensor to fall between $\ell_{\max}^{-2}$ and $\ell_{\min}^{-2}$. As a result, features that are extremely small in absolute geometrical size are coarsened, while larger ones remain unaffected. In our work, we take the same approach, but apply it to the projected lengths of edges in screen space: edges that *appear* too small, relative to the viewer, are coarsened. To do so, we transform the sizing tensor $\tilde{\mathbf{M}}$ to screen space, apply the bound on the shortest allowable edge length, and transform it back to material space.

For a given configuration of the sheet, consider the function from material-space coordinates of points to their screen-space coordinates. On a visible face, this function is locally invertible and its Jacobian $\mathbf{S}$ is a full-rank 2×2 matrix that can be computed from the screen-space coordinates of the three vertices. As the sizing tensor $\tilde{\mathbf{M}}$ acts as a quadratic form acting on vectors in material space, the corresponding tensor that acts on screen-space vectors can be obtained via the transformation

$$\mathcal{S}(\tilde{\mathbf{M}}) = \mathbf{S}^{-\top} \tilde{\mathbf{M}} \mathbf{S}^{-1}. \quad (5)$$

We require the minimum change to the screen-space sizing tensor $\mathcal{S}(\tilde{\mathbf{M}})$ such that edges of screen-space length $\ell_{\min}$ will not be refined further. This change can be achieved by modifying the eigendecomposition of $\mathcal{S}(\tilde{\mathbf{M}})$ so that its

Fig. 4. As the distance between the camera and each character increases from right to left, the simulation mesh becomes progressively coarser (below). However, the minimum edge length measured in screen-space remains roughly constant. As a result the figures shown in the rendered image (above) have a uniform level of visual detail.

eigenvalues do not exceed $\ell_{\min}^{-2}$:

$$\mathcal{S}(\tilde{\mathbf{M}}) = \mathbf{Q} \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix} \mathbf{Q}^\top, \quad (6)$$

$$\hat{\lambda}_i = \min(\lambda_i, \ell_{\min}^{-2}), \quad (7)$$

$$\hat{\mathbf{M}} = \mathbf{Q} \begin{bmatrix} \hat{\lambda}_1 & 0 \\ 0 & \hat{\lambda}_2 \end{bmatrix} \mathbf{Q}^\top. \quad (8)$$

We transform this modified screen-space tensor back into material space to obtain the final sizing field we use for remeshing visible regions, $\mathbf{M}_{\text{vd}} = \mathcal{S}^{-1}(\hat{\mathbf{M}}) = \mathbf{S}^\top \hat{\mathbf{M}} \mathbf{S}$.

3.3 Transferring sizing field from faces to vertices

The sizing field defined by the procedures above is represented as a metric tensor on each face. This tensor field must be resampled onto mesh vertices so that it can be used in the sizing criterion (1). Previous work [14] has used a simple area-weighted averaging procedure. However, we have found that the approach tends to lose detail in regions with rapid variation in the sizing field, such as at silhouettes where the view factor changes from ν_{front} to ν_{back} . The issue is exacerbated because coarser regions, which have larger faces, are given higher weight, leading to excessive coarsening at boundary vertices.

In order to handle this discontinuity in view factors, we first resample per-face view factors before the simple area-weighted averaging procedure. If the values of ν differ by more than a specified threshold across any edge, we do simple averaging between two adjacent view factors, and assign the averaged value to one of the faces as shown in Fig. 5. As we do not want to change the view

Fig. 5. Locally interpolating discontinuous view factors at silhouettes to ensure smooth silhouettes.

factors of the visible faces from ν_{front} , we always assign the averaged view factor to a non-visible face.

This approach ensures that silhouette boundaries are refined to the same extent as other visible regions of the cloth, improving the appearance of silhouettes. This change affects the simulation mesh only in a limited area near silhouette boundaries, so it does not hinder overall performance.

3.4 Interpenetration handling

Remeshing invariably produces changes in the geometry of the cloth mesh, and can introduce interpenetrations of the cloth with itself or with obstacles. We found that the simple approach for interpenetration handling used in previous remeshing work [14] does not always converge to an interpenetration-free configuration in the presence of the aggressive coarsening we perform in non-visible regions. Instead we use the approach of intersection contour minimization proposed by Volino et al. [32], which we found to be robust to large and complex interpenetrations.

3.5 Coarse-scale buckling compensation

Real cloth materials typically have a relatively high resistance to in-plane compression, but a relatively low resistance to out-of-plane bending. As a result, cloth will often drape in configurations that would be non-developable at coarse scale, but where fine-scale wrinkles and folds take the place of coarse-scale compression. These fine wrinkles and folds not only allow commonly observed draping, but also responsible for giving specific materials their distinctive appearances.

However, extremely coarse mesh regions do not have the resolution required to form fine wrinkles or folds, and furthermore those coarse regions tend to suffer from element locking that suppresses bending and wrinkle formation. As a result, when out-of-view regions are heavily coarsened, they often exhibit behavior that is dramatically different from that of more refined regions. In particular, coarse regions have a tendency to balloon out in an undesirable manner that is atypical of real cloth, as shown in Figs. 3 and 10.

Because extreme coarsening mainly occurs for out-of-frame cloth regions, this undesirable behavior is generally not observable by the viewer. Additionally, our temporal smoothing causes cloth regions to be refined before they are visible, allowing transient artifacts to settle out. Nevertheless, the incorrect motion of out-of-frame regions may cause undesirable side effects that affect the behavior of the visible part of the cloth.

Fig. 6. A comparison of no buckling compensation (left), non-adaptive buckling compensation (middle), and adaptive buckling compensation with $k = 10$ (right).

We can avoid this artificial stiffening by switching coarse faces to a material model that permits compression in lieu of sub-element buckling [33], [34]. However, simply removing compression resistance for all out-of-view faces would also create undesirable artifacts, as shown in Fig. 6 (middle). Instead, we allow compression adaptively based on each face's area in material space as follows.

The Green strain tensor of face i is

$$\mathbf{E}_i = \frac{1}{2}(\mathbf{F}_i^T \mathbf{F}_i - \mathbf{I}), \quad (9)$$

where $\mathbf{F}_i$ is the deformation gradient. We reduce the face's compression resistance by simply reducing the amount of compressive strain that the constitutive model sees. Specifically, for each out-of-frustum face, we compute the eigendecomposition of its strain tensor $\mathbf{E}_i$ to obtain principal strains λ_1 and λ_2 . We replace each principal strain with

$$\hat{\lambda}_j = \begin{cases} \lambda_j & \text{if } \lambda_j \geq 0, \\ \alpha_i \lambda_j & \text{if } \lambda_j < 0, \end{cases} \quad (10)$$

with area-based compression factor α_i ,

$$\alpha_i = \left(\frac{A_{\max} - A_i}{A_{\max} - A_{\min}} \right)^k, \quad (11)$$

where A_i is area of face i in material space, $A_{\max}$ and $A_{\min}$ are the areas of the largest and smallest out-of-frustum faces, and $k \geq 1$ is a parameter controlling the amount of compensation. Because this compression factor increases with material-space area, the degree of allowed compression changes adaptively: the bigger the triangle, the more compressible it is. This behavior approximates the fine-scale buckling that would otherwise have occurred in the cloth represented by the coarse element, leading to a more plausible behavior in coarse regions as shown in Figs. 6 and 7.

We note that this approach to approximating the coarse effects of fine-scale buckling is not specific to *view-dependent* adaptivity, and it could be applied to other adaptive algorithms suffering from similar artifacts. The compression factor α is unrelated to the view factor, as buckling is purely a physical phenomenon independent of the viewing direction. Therefore, we have defined the compression factor entirely in terms of the area of faces and a user-defined parameter k .

Fig. 7. No buckling compensation (top) and adaptive buckling compensation with $k = 10$ (bottom) in Multiple Dancers example.

4 RESULTS AND DISCUSSION

We have implemented the methods described in this paper as an extension to the ARCSim simulator. These extensions will be integrated into a future release of the software.

To test the utility of the proposed methods, we ran comparisons with adaptive and non-adaptive simulations on several examples. All simulations were run on a machine with an Intel Xeon E5-2680 v2 processor and 16 GB RAM using 4 execution threads. The constants for view-dependent refinement that were used for these examples are $\nu_{\text{front}} = 1$, $\nu_{\text{back}} = 0.2$, $m = 0.4$ m, $T = 5$ frames, and $\nu_{\text{out}} = 0.01$. However, for the example shown in Fig. 1, we used larger $\nu_{\text{out}} = 0.1$ to avoid intersection problems with the layered garments.

Fig. 1 shows a character wearing a layered outfit consisting of three interacting garments. As the accompanying video shows, the character's clothing exhibits complex dynamic motion with detailed wrinkles. As the camera

Fig. 8. A comparison of view-independent and view-dependent adaptive simulations for the example of Fig. 1 in terms of per-frame and cumulative simulation times.

Fig. 9. A solo dancer is shown from the camera's perspective (left, middle) and from an external perspective with the camera's view frustum drawn in black wireframe (right).

Fig. 10. A group of multiple dancers is rendered from the camera's perspective (top), while an external view showing the camera's view frustum is shown below. Characters completely outside the view frustum are simulated at very low resolution.

moves continuously or transits between cuts, the simulation mesh is updated to maintain a constant level of visible detail while out-of-view regions are aggressively coarsened. Fig. 8 plots the time per frame and total cumulative computation time for this example for both the basic adaptive simulation and our view-dependent adaptive version.

Fig. 9 shows a character wearing a simple dress, showing a degree of view-dependent refinement and coarsening similar to Fig. 1. A group of ten such characters is shown in Fig. 10. Note that while these characters are all performing the same actions in the same type of dress, they are simulated individually with different resolutions depending on their location relative to the camera. In practical applications, multiple characters in a crowd would likely have different garments and distinct movements, and would

TABLE 1

Statistics and timing numbers for the examples. Non-adaptive simulations use a fixed high-resolution mesh. Adaptive simulations use the unmodified scheme implemented in ARCSim. View-dependent simulations use the method described in this paper. The adaptive simulations are used as a baseline for comparison. The non-adaptive mesh resolution was selected to match the visual quality of the adaptive simulations.

	Num. Faces			Num. Vertices			Time / Frame (seconds)			Speed-up
	min	max	mean	min	max	mean	min	max	mean	
Karate										
Non-adaptive	123,790	123,790	123,790	63,889	63,889	63,889	136.47	288.01	161.44	0.49×
Adaptive	25,785	71,142	41,973	14,135	37,199	22,358	31.41	184.39	79.04	1×
View-dependent	5,223	39,139	21,501	3,142	20,770	11,670	5.52	83.29	33.02	2.39×
Solo Dancer										
Non-adaptive	43,710	43,710	43,710	22,736	22,736	22,736	27.31	58.09	29.27	0.54×
Adaptive	12,030	21,763	18,041	6,593	11,535	9,659	7.78	22.00	15.80	1×
View-dependent	730	15,951	9,638	560	8,599	5,314	0.83	13.81	7.55	2.09×
Multiple Dancers										
Non-adaptive	437,100	437,100	437,100	227,360	227,360	227,360	273.13	580.85	292.73	0.47×
Adaptive	119,515	184,340	161,028	65,505	98,630	86,897	79.97	178.24	136.99	1×
View-dependent	11,216	102,945	36,339	7,619	56,285	21,228	12.02	82.54	30.76	4.45×

therefore also have to be simulated individually even without view dependence.

Timings for these three examples are reported in Table 1. The single character examples realize a speed-up between $2.1\times$ and $2.4\times$. This speed-up becomes more substantial for the group of ten characters where it is roughly $4.5\times$. The greater speed-up occurs because when a single character fills the screen, requiring full adaptive resolution, other characters tend to go out of frame which then switch to low resolution and generate substantial savings. In order for all characters to be visible, the camera must be fairly far back which yields savings across all characters. These same factors would hold more strongly for larger groups. Accordingly we hypothesize that massive scenes involving thousands of characters would likely realize speed-ups substantially larger than those we have observed for small groups.

Remeshing does introduce small amounts of motion that might be visible if the character is otherwise still. As shown in the video, however, this motion does not manifest as visible popping because the dramatic remeshing generally happens off screen. The motion that appears on screen is only gentle swaying that looks fairly natural. If the subject is moving, then this gentle motion is completely masked by the normal motion of the cloth. Even with a still subject, the artificial motion can be hard to detect due to the movement of the camera.

5 CONCLUSIONS

The methods we have described provide a simple way of achieving computational savings for cloth simulation. For scenes involving multiple characters or large crowds, these savings can be substantial. We have demonstrated our approach in simulations of clothing, but believe that it could equally well be applied to other objects that can be simulated in an adaptive framework, including materials that can crumple [15] or fracture [16].

The main limitation of our method is that it requires an adaptive framework. However, once that framework is in place, view-dependent adaptivity is relatively simple to implement. For the group of dancers, our view-dependent

adaptive simulation is nearly $10\times$ faster than an equivalent non-adaptive simulation. We believe that such large performance gains outweigh any costs associated with changing mesh topology. We also believe that our approach would scale very well to massive scenes with thousands of actors, where it would produce even larger savings.

As mentioned in the introduction, our method targets high-quality offline simulation, and there would be additional challenges in applying it to interactive animation. In particular, our temporal smoothing requires a prescribed camera motion, so it may not be possible to apply it in interactive settings where camera motions are not predetermined.

Our technique for buckling compensation rests on the core idea of modifying the constitutive model to account for unresolved sub-element detail. Such an approach could be useful in a much broader context, for example to improve simulation fidelity in resource-constrained applications, or to improve the correspondence between coarse preview visualizations and the final high-resolution computation. We hope to explore such avenues in future work.

In general, simulations used for physically based animation in graphics have been designed so that they capture visible phenomena for realistic appearance. These simulations typically avoid the type of error analysis that one finds in most engineering contexts because it is difficult to quantify a measure of perceptual error that would be relevant to graphics simulations. The work we've presented here explicitly drives adaption according to a heuristic measure of visible error. Future work in this area could more explicitly take perceptual error into account so that reduced resolution could be used where masking effects due to motion, complexity, or other phenomena cause humans to be less sensitive to apparent detail.

ACKNOWLEDGMENTS

The authors thank the other members of the Visual Computing Lab at UC Berkeley for their valuable advice and help. This work was supported by NSF grants CNS-1444840 and IIS-1353155, an AWS in Education Grant award, and gifts from Intel, Pixar, and Qualcomm. Woojong Koh was

supported by a Samsung Scholarship. This project was built using the ARCSim framework developed at UC Berkeley, and the images in this paper and video were rendered with the Mitsuba renderer by Wenzel Jakob.

REFERENCES

- [1] F. Losasso, F. Gibou, and R. Fedkiw, "Simulating water and smoke with an octree data structure," *ACM Trans. Graph.*, vol. 23, no. 3, pp. 457–462, Aug. 2004.
- [2] B. Adams, M. Pauly, R. Keiser, and L. J. Guibas, "Adaptively sampled particle fluids," *ACM Trans. Graph.*, vol. 26, no. 3, pp. 48:1–7, Jul. 2007.
- [3] N. Chentanez and M. Müller, "Real-time Eulerian water simulation using a restricted tall cell grid," in *ACM SIGGRAPH 2011 Papers*, ser. SIGGRAPH 2011, 2011, pp. 82:1–82:10.
- [4] B. M. Klingner, B. E. Feldman, N. Chentanez, and J. F. O'Brien, "Fluid animation with dynamic meshes," in *Proc. of ACM SIGGRAPH 2006*, Aug. 2006, pp. 820–825.
- [5] N. Chentanez, B. E. Feldman, F. Labelle, J. F. O'Brien, and J. R. Shewchuk, "Liquid simulation on lattice-based tetrahedral meshes," in *ACM SIGGRAPH/Eurographics Symp. on Computer Animation 2007*, Aug. 2007, pp. 219–228.
- [6] R. Ando, N. Thürey, and C. Wojtan, "Highly adaptive liquid simulations on tetrahedral meshes," *ACM Trans. Graph.*, vol. 32, no. 4, pp. 103:1–103:10, Jul. 2013.
- [7] D. Hutchinson, M. Preston, and T. Hewitt, "Adaptive refinement for mass/spring simulations," in *7th Eurographics Workshop on Animation and Simulation*. Springer-Verlag, 1996, pp. 31–45.
- [8] S. Redon, N. Galoppo, and M. C. Lin, "Adaptive dynamics of articulated bodies," *ACM Trans. on Graphics (SIGGRAPH 2005)*, vol. 24, no. 3, 2005.
- [9] E. Grinspun, P. Krysl, and P. Schröder, "CHARMS: A simple framework for adaptive simulation," in *Proc. of the 29th Annual Conf. on Computer Graphics and Interactive Techniques*, ser. SIGGRAPH 2002, 2002, pp. 281–290.
- [10] J. A. Thingvold and E. Cohen, "Physical modeling with B-spline surfaces for interactive design and animation," *SIGGRAPH Comput. Graph.*, vol. 24, no. 2, pp. 129–137, Feb. 1990.
- [11] L. Li and V. Volkov, "Cloth animation with adaptively refined meshes," in *Proc. 28th Australasian Computer Science Conference*, vol. 38, 2005.
- [12] J. Villard and H. Borouchaki, "Adaptive meshing for cloth animation," *Eng. with Computers*, vol. 20, no. 4, pp. 333–341, 2005.
- [13] T. Brochu, E. Edwards, and R. Bridson, "Efficient geometrically exact continuous collision detection," *ACM Trans. Graph.*, vol. 31, no. 4, pp. 96:1–96:7, Jul. 2012.
- [14] R. Narain, A. Samii, and J. F. O'Brien, "Adaptive anisotropic remeshing for cloth simulation," *ACM Trans. on Graphics*, vol. 31, no. 6, pp. 147:1–10, Nov. 2012.
- [15] R. Narain, T. Pfaff, and J. O'Brien, "Folding and crumpling adaptive sheets," *ACM Trans. Graph.*, vol. 32, no. 4, pp. 51:1–8, Jul. 2013.
- [16] T. Pfaff, R. Narain, J. M. de Joya, and J. F. O'Brien, "Adaptive tearing and cracking of thin sheets," *ACM Trans. on Graphics*, vol. 33, no. 4, pp. 1–9, Jul. 2014.
- [17] D. A. Carlson and J. K. Hodgins, "Simulation levels of detail for real-time animation," in *Proc. of the Conf. on Graphics Interface '97*, 1997, pp. 1–8.
- [18] D. O'Brien, S. Fisher, and M. Lin, "Automatic simplification of particle system dynamics," in *The Fourteenth Conf. on Computer Animation*, 2001, pp. 210–257.
- [19] J. Beaudoin and J. Keyser, "Simulation levels of detail for plant motion," in *Proc. of the 2004 ACM SIGGRAPH/Eurographics Symp. on Computer Animation*, ser. SCA '04, 2004, pp. 297–304.
- [20] K. Ward, M. C. Lin, J. Lee, S. Fisher, and D. Macri, "Modeling hair using level-of-detail representations," in *Proc. of the 16th International Conf. on Computer Animation and Social Agents (CASA 2003)*, ser. CASA '03, 2003, pp. 41–47.
- [21] S. Cheney and D. Forsyth, "View-dependent culling of dynamic systems in virtual environments," in *Proc. of the 1997 Symp. on Interactive 3D Graphics*, ser. I3D '97, 1997, pp. 55–58.
- [22] B. A. Barran, "View dependent fluid dynamics," Master's thesis, Texas A&M University, 2006.
- [23] J. Kim, I. Ihm, and D. Cha, "View-dependent adaptive animation of liquids," *ETRI Journal*, vol. 28, no. 6, pp. 697–708, Dec. 2006.
- [24] R. Bunlutangtum and P. Kanongchaiyos, "Enhanced view-dependent adaptive grid refinement for animating fluids," in *Proc. of the 10th International Conf. on Virtual Reality Continuum and Its Applications in Industry*, ser. VRCAI '11, 2011, pp. 415–418.
- [25] B. Solenthaler and M. Gross, "Two-scale particle simulation," *ACM Trans. Graph.*, vol. 30, no. 4, pp. 81:1–81:8, Jul. 2011.
- [26] M. Duchaineau, M. Wolinsky, D. E. Sigeti, M. C. Miller, C. Aldrich, and M. B. Mineev-Weinstein, "ROAMing terrain: Real-time optimally adapting meshes," in *Proc. of the 8th Conf. on Visualization '97*, ser. VIS '97, 1997, pp. 81–88.
- [27] H. Hoppe, "Smooth view-dependent level-of-detail control and its application to terrain rendering," in *Proc. of the Conf. on Visualization '98*, ser. VIS '98, 1998, pp. 35–42.
- [28] T. A. Funkhouser and C. H. Séquin, "Adaptive display algorithm for interactive frame rates during visualization of complex virtual environments," in *Proc. of the 20th Annual Conf. on Computer Graphics and Interactive Techniques*, ser. SIGGRAPH '93, 1993, pp. 247–254.
- [29] D. Luebke and C. Erikson, "View-dependent simplification of arbitrary polygonal environments," in *Proc. of the 24th Annual Conf. on Computer Graphics and Interactive Techniques*, ser. SIGGRAPH '97, 1997, pp. 199–208.
- [30] H. Hoppe, "View-dependent refinement of progressive meshes," in *Proc. of the 24th Annual Conf. on Computer Graphics and Interactive Techniques*, ser. SIGGRAPH '97, 1997, pp. 189–198.
- [31] J. C. Xia, J. El-Sana, and A. Varshney, "Adaptive real-time level-of-detail-based rendering for polygonal models," *IEEE Trans. on Visualization and Computer Graphics*, vol. 3, no. 2, pp. 171–183, Apr. 1997.
- [32] P. Volino and N. Magnenat-Thalmann, "Resolving surface collisions through intersection contour minimization," *ACM Trans. Graph.*, vol. 25, no. 3, pp. 1154–1159, Jul. 2006.
- [33] K.-J. Choi and H.-S. Ko, "Stable but responsive cloth," *ACM Trans. Graph.*, vol. 21, no. 3, pp. 604–611, Jul. 2002.
- [34] H. Wang, J. O'Brien, and R. Ramamoorthi, "Multi-resolution isotropic strain limiting," *ACM Trans. Graph.*, vol. 29, no. 6, pp. 156:1–156:10, Dec. 2010.

Woojong Koh is a third-year Ph.D. student in Computer Science at the University of California, Berkeley. He received his B.S. in Computer Science and Engineering from the Seoul National University, Seoul, South Korea.

Rahul Narain is an Assistant Professor in the Department of Computer Science & Engineering at the University of Minnesota. He received his Ph.D. in Computer Science from the University of North Carolina at Chapel Hill and was a postdoctoral scholar at the University of California, Berkeley. His research interests lie in physics-based animation and numerical simulation, particularly focusing on adaptive techniques for modeling complex multiscale phenomena.

James F. O'Brien is a Professor of Computer Science at the University of California, Berkeley. His primary area of interest is computer animation, with an emphasis on generating realistic motion using physically based simulation and motion capture techniques. He has authored numerous papers on these topics. In addition to his research pursuits, Prof. O'Brien has worked with film and game companies on integrating advanced simulation physics into games and special effects. His methods for destruction modeling

have been used in over 70 feature films and AAA game titles. In 2015 the Academy of Motion Picture Arts and Sciences recognized his work in destruction modeling with an Award for Technical Achievement. He received his doctorate from the Georgia Institute of Technology in 2000, the same year he joined the faculty at U.C. Berkeley. Professor O'Brien is a Sloan Fellow and ACM Distinguished Scientist, has been selected as one of Technology Review's TR-100, and has been awarded research grants from the Okawa and Hellman Foundations. He is currently serving as ACM SIGGRAPH Director at Large.