

View-Dependent Adaptive Cloth Simulation

Woojong Koh



Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2016-5

<http://www.eecs.berkeley.edu/Pubs/TechRpts/2016/EECS-2016-5.html>

January 9, 2016

Copyright © 2016, by the author(s).
All rights reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

Acknowledgement

I would like to thank Prof. James F. O'Brien and Dr. Rahul Narain for their valuable advice and help on the thesis. I also thank Prof. Jonathan R. Shewchuk for reviewing the manuscript and being a reader for my thesis. I was supported by a Samsung Scholarship and this work is also supported by gifts from Amazon, Intel, and Pixar Animation Studios. The images in this thesis were rendered with the Mitsuba renderer by Wenzel Jakob.

View-Dependent Adaptive Cloth Simulation

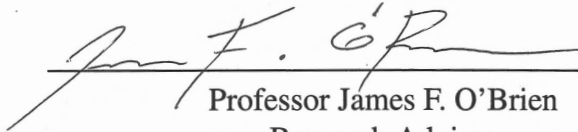
by Woojong Koh

Research Project

Submitted to the Department of Electrical Engineering and Computer Sciences,
University of California at Berkeley, in partial satisfaction of the requirements for
the degree of **Master of Science, Plan II.**

Approval for the Report and Comprehensive Examination:

Committee:

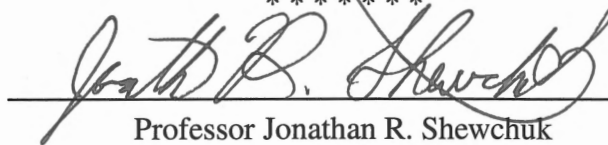


Professor James F. O'Brien

Research Advisor

May 14, 2014

(Date)



Professor Jonathan R. Shewchuk

Second Reader

9 May 2014

(Date)

View-Dependent Adaptive Cloth Simulation

Woojong Koh*

University of California, Berkeley

May 9, 2014

Abstract

This thesis 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. 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 increases as the scene complexity grows, producing a $2\times$ speed-up for a single character and more than $4\times$ for a small group.

*wjkoh@berkeley.edu



Figure 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.

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, most fine details will not be visible to the viewer and work spent computing those details is wasted. In most film production settings, both the camera and character motion 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, but the cloth motion must 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 don't result in dynamic artifacts that affect visible parts of the cloth.

In this thesis 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 detail of the simulated cloth. The degree to which this detail is resolved is adjusted locally based on the view of the simulated cloth. Areas that appear large in the camera will be refined to show finely detailed dynamics. Areas that are out of frame, facing away from the camera, or at a distance will be correspondingly less refined.

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 framework¹ which can be used to animate sheets of deformable materials such as cloth, paper, plastic, and metal. ARCSim adaptively refines anisotropic triangle meshes to efficiently resolve the geometric and dynamic detail of the simulated objects. Our method modifies the metrics used by ARCSim so that local mesh visibility is accounted for during 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, we realize modest savings of roughly a $2.4\times$ speed-up in comparison with an adaptive simulation of ARCSim due to

¹Adaptive Refining and Coarsening Simulator (ARCSim) is available at <http://graphics.berkeley.edu/resources/ARCSim>.

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.

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 [LGF04], adaptively sampled particles [APKG07], tall-cell grids [CM11], or tetrahedral meshes [KFCO06, CFL*07, ATW13]. Adaptive refinement and simplification techniques have also been proposed for articulated bodies [RGL05] and finite element models [GKS02]. Most relevant to our work is the approach of Narain et al. [NSO12], which uses adaptive anisotropic remeshing to resolve detailed features in thin sheets such as cloth. This approach has been extended to efficiently model plastic deformation and sharp creases [NPO13] as well as complex fracture patterns [PNdJO14]. However, all the techniques described above are view-independent and rely only on 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 [CH97] 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, whereas ours does not. Subsequent work has sought

to generate approximated models automatically, for example for particle systems [OFL01], plants [BK04], and hair [WLL*03].

Alternatively, one can modify the simulation resolution based on the viewpoint without changing the underlying model. Barran [Bar06] 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 [KIC06, BK11] and adaptive particles [SG11].

In our work, we also take inspiration from geometric level of detail techniques, such as those for real-time rendering of terrain [DWS*97, Hop98] or complex scenes like architectural walkthroughs [FS93]. These techniques inform our understanding of the important view-dependent criteria for representing geometrical detail. Hoppe [Hop97] used surface orientation and screen-space geometric error as refinement criteria. Xia et al. [XESV97] further propose the use of local illumination gradients, projected lengths of edges in screen space, and silhouette boundaries.

3 Method

Our view-dependent adaptive remeshing scheme builds on the adaptive anisotropic remeshing framework described by Narain et al. [NSO12]. 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} ,

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

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 anisotropic 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 size of each face 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 detail.

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 scale the sizing field to uniformly coarsen the mesh. 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 field $v \leq 1$, which we call the *view factor*, and modify the sizing field as

$$\mathbf{M}_{\text{vd}} = v^2 \mathbf{M}. \quad (2)$$

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

In general, we would like to use full resolution ($v = 1$) for faces that are visible in the current view, and coarser resolution for back-facing and out-of-frame faces based on user-specified parameters $v_{\text{back}}, v_{\text{out}} < 1$. However, simply defining v in this piecewise-constant fashion 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, let d be its distance from the view frustum in world space. We define the spatially smoothed view factor \tilde{v} by

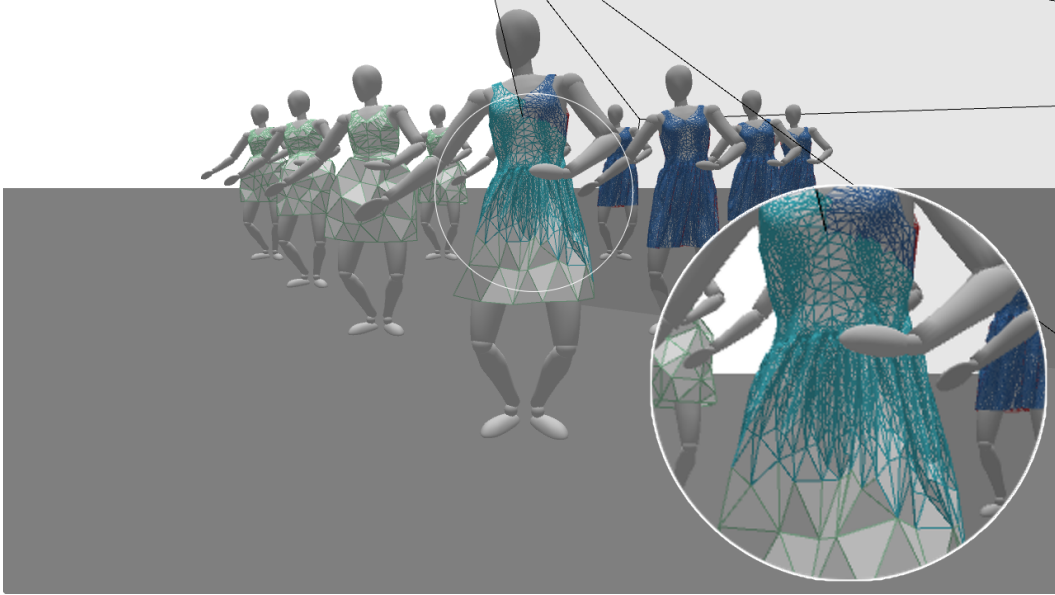


Figure 2: An example of spatial smoothing. Cyan edges indicate out-of-view faces under spatial smoothing, and green edges indicate out-of-view faces far from the view frustum.

linearly interpolating to v_{out} over a user-specified margin length m :

$$\tilde{v} = \begin{cases} v_{\text{fb}} & \text{if } d = 0, \\ v_{\text{fb}} - \frac{d}{m}(v_{\text{fb}} - v_{\text{out}}) & \text{if } 0 < d < m, \\ v_{\text{out}} & \text{if } d \geq m, \end{cases} \quad (3)$$

where v_{fb} is 1 or v_{back} depending on the direction of the face normal. Thus, we have $\tilde{v} = 1$ or v_{back} for visible faces and $\tilde{v} = v_{\text{out}}$ for faces far from the view frustum, with a continuous transition region in between, as Figure 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 normal of back-faces to create a smooth transition, we find that normals can vary too rapidly across silhouettes to offer any useful smoothing this 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{v} as a function of time, holding the face fixed and considering the prescribed motion of the viewpoint. We smooth \tilde{v} 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

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

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

Unlike smoothing by, say, moving averages, our approach is conservative in that $v(t) \geq \tilde{v}(t)$; in particular, visible regions always have $v = 1$. Further, in the presence of discontinuities such as jump cuts, v increases continuously from v_{out} to 1 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 Figure 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.

Our objective is that faces and edges in visible regions should be large enough to affect the visual appearance of cloth in the final image. The farther a face is from a camera, the smaller it appears under perspective projection. The projected area of a triangle also becomes smaller as the face normal becomes perpendicular to the viewing direction. Therefore, in regions that are distant, or are not normal to the view direction, a high-resolution mesh will be wasteful as any detailed features resolved by the fine elements will appear extremely small in the final image.

In this section, we describe how to limit the refinement of mesh elements based on their projected sizes in screen space. We do so by constraining the allowed edge lengths to a specified minimum visual angle in screen space. This approach takes both orientation and depth into account in unified and consistent way. Previous work [NSO12] constrains the allowed edge lengths in material space to a range $[\ell_{\min}, \ell_{\max}]$ by clamping the eigenvalues of the sizing tensor to lie between ℓ_{\max}^{-2} and ℓ_{\min}^{-2} . In our work, however, we constrain the sizing tensor with respect to screen space rather than material space. We transform the sizing tensor \mathbf{M}_{vd} 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 vertices to their screen-space coordinates. On visible faces, this function is locally invertible and its Jacobian \mathbf{S} is a full-rank 2×2 matrix that can be evaluated locally for each face. As the sizing tensor \mathbf{M}_{vd} 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

$$\mathcal{S}(\mathbf{M}_{\text{vd}}) = \mathbf{S}^{-\text{T}} \mathbf{M}_{\text{vd}} \mathbf{S}^{-1}. \quad (5)$$

We seek the minimum change to the screen-space sizing tensor $\tilde{\mathbf{M}}_{\text{vd}} = \mathcal{S}(\mathbf{M}_{\text{vd}})$ such that edges of screen-space length ℓ_{\min} will not be refined further. This modification can be

achieved by clamping the eigenvalues of $\tilde{\mathbf{M}}_{\text{vd}}$ to not exceed ℓ_{\min}^{-2} :

$$\tilde{\mathbf{M}}_{\text{vd}} = \mathbf{Q} \begin{bmatrix} \tilde{\lambda}_1 & 0 \\ 0 & \tilde{\lambda}_2 \end{bmatrix} \mathbf{Q}^\top, \quad (6)$$

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

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

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

3.3 Transferring sizing field from faces to vertices

The sizing field defined by the procedure above is represented as a 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 [NSO12] has used a simple area-weighted averaging procedure. However, we have found that that approach tends to lose detail in regions with rapid variation in the sizing field, such as at silhouettes where the view factor changes from 1 to v_{out} . 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 on view factors, we first resample per-face view factors before the simple area-weighted averaging procedure. If the values of v 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 Figure 4. As we don't want to change the view factors of the visible faces from 1, 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

	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×

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 thesis. 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.

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 [NSO12] 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. [VMT06], which we found to be robust to large and complex interpenetrations.

4 Results and Discussion

We have implemented the methods described in this thesis as an extension to the ARCSim code. 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 $v_{\text{back}} = 0.2$, $m = 0.4$ m, $T = 5$ frames, and $v_{\text{out}} = 0.01$. For the example shown in Figure 1, we used larger $v_{\text{out}} = 0.1$ to avoid intersection problems with the layered garments.

Figure 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 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. Figure 5 plots the time per frame and total cumulative computation time for this example for both the basic adaptive simulation and the view-dependent version.

Figure 6 shows a character wearing a simple dress, showing a similar degree of view-dependent refinement and coarsening to Figure 1. A group of ten such characters is shown in Figure 7. 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 viewpoint. In practical applications, multiple characters in a crowd would have different garments and motions and would therefore have to be simulated individually even without view dependence.

Timings for these three examples are reported in Table 1. The single characters 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, it tends to obscure

others 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 observations would hold more strongly for larger groups. Accordingly we hypothesize that massive scenes involving thousands of characters would likely realize speed-up 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 gentle swaying that looks fairly natural. If the subject is moving, then the artificial 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 [NPO13] or fracture [PNdJO14].

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 ten times faster than 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.

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

I would like to thank Prof. James F. O'Brien and Dr. Rahul Narain for their valuable advice and help on the thesis. I also thank Prof. Jonathan R. Shewchuk for reviewing the manuscript and being a reader for my thesis. I was supported by a Samsung Scholarship and this work is also supported by gifts from Amazon, Intel, and Pixar Animation Studios.

The images in this thesis were rendered with the Mitsuba renderer by Wenzel Jakob.

References

- [APKG07] ADAMS B., PAULY M., KEISER R., GUIBAS L. J.: Adaptively sampled particle fluids. *ACM Trans. Graph.* 26, 3 (July 2007).
- [ATW13] ANDO R., THÜREY N., WOJTAN C.: Highly adaptive liquid simulations on tetrahedral meshes. *ACM Trans. Graph.* 32, 4 (July 2013), 103:1–103:10.
- [Bar06] BARRAN B. A.: *View dependent fluid dynamics*. Master's thesis, Texas A&M University, 2006.

- [BK04] BEAUDOIN J., KEYSER J.: Simulation levels of detail for plant motion. In *Proceedings of the 2004 ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (2004), SCA '04, pp. 297–304.
- [BK11] BUNLUTANGTUM R., KANONGCHAIYOS P.: Enhanced view-dependent adaptive grid refinement for animating fluids. In *Proceedings of the 10th International Conference on Virtual Reality Continuum and Its Applications in Industry* (2011), VRCAI '11, pp. 415–418.
- [CFL*07] CHENTANEZ N., FELDMAN B. E., LABELLE F., O'BRIEN J. F., SHEWCHUK J. R.: Liquid simulation on lattice-based tetrahedral meshes. In *ACM SIGGRAPH/Eurographics Symposium on Computer Animation 2007* (Aug. 2007), pp. 219–228.
- [CH97] CARLSON D. A., HODGINS J. K.: Simulation levels of detail for real-time animation. In *Proceedings of the Conference on Graphics Interface '97* (1997), pp. 1–8.
- [CM11] CHENTANEZ N., MÜLLER M.: Real-time Eulerian water simulation using a restricted tall cell grid. In *ACM SIGGRAPH 2011 Papers* (2011), SIGGRAPH 2011, pp. 82:1–82:10.
- [DWS*97] DUCHAINEAU M., WOLINSKY M., SIGETI D. E., MILLER M. C., ALDRICH C., MINEEV-WEINSTEIN M. B.: Roaming terrain: Real-time optimally adapting meshes. In *Proceedings of the 8th Conference on Visualization '97* (1997), VIS '97, pp. 81–88.
- [FS93] FUNKHOUSER T. A., SÉQUIN C. H.: Adaptive display algorithm for interactive frame rates during visualization of complex virtual environments.

In *Proceedings of the 20th Annual Conference on Computer Graphics and Interactive Techniques* (1993), SIGGRAPH '93, pp. 247–254.

- [GKS02] GRINSPUN E., KRYSL P., SCHRÖDER P.: Charms: A simple framework for adaptive simulation. In *Proceedings of the 29th Annual Conference on Computer Graphics and Interactive Techniques* (2002), SIGGRAPH 2002, pp. 281–290.
- [Hop97] HOPPE H.: View-dependent refinement of progressive meshes. In *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques* (1997), SIGGRAPH '97, pp. 189–198.
- [Hop98] HOPPE H.: Smooth view-dependent level-of-detail control and its application to terrain rendering. In *Proceedings of the Conference on Visualization '98* (1998), VIS '98, pp. 35–42.
- [KFCO06] KLINGNER B. M., FELDMAN B. E., CHENTANEZ N., O'BRIEN J. F.: Fluid animation with dynamic meshes. In *Proceedings of ACM SIGGRAPH 2006* (Aug. 2006), pp. 820–825.
- [KIC06] KIM J., IHM I., CHA D.: View-dependent adaptive animation of liquids. *ETRI Journal* 28, 6 (Dec. 2006), 697–708.
- [LGF04] LOSASSO F., GIBOU F., FEDKIW R.: Simulating water and smoke with an octree data structure. *ACM Trans. Graph.* 23, 3 (Aug. 2004), 457–462.
- [NPO13] NARAIN R., PFAFF T., O'BRIEN J. F.: Folding and crumpling adaptive sheets. *ACM Trans. Graph.* 32, 4 (July 2013), 51:1–51:8.
- [NSO12] NARAIN R., SAMII A., O'BRIEN J. F.: Adaptive anisotropic remeshing for cloth simulation. *ACM Transactions on Graphics* 31, 6 (Nov. 2012), 147:1–10. Proceedings of ACM SIGGRAPH Asia 2012, Singapore.

- [OFL01] O'BRIEN D., FISHER S., LIN M.: Automatic simplification of particle system dynamics. In *The Fourteenth Conference on Computer Animation. Proceedings* (2001), pp. 210–257.
- [PNdJO14] PFAFF T., NARAIN R., DE JOYA J. M., O'BRIEN J. F.: Adaptive tearing and cracking of thin sheets. *ACM Trans. Graph.* (2014). To appear.
- [RGL05] REDON S., GALOPPO N., LIN M. C.: Adaptive dynamics of articulated bodies. *ACM Transactions on Graphics (SIGGRAPH 2005)* 24, 3 (2005).
- [SG11] SOLENTHALER B., GROSS M.: Two-scale particle simulation. *ACM Trans. Graph.* 30, 4 (July 2011), 81:1–81:8.
- [VMT06] VOLINO P., MAGNENAT-THALMANN N.: Resolving surface collisions through intersection contour minimization. *ACM Trans. Graph.* 25, 3 (July 2006), 1154–1159.
- [WLL*03] WARD K., LIN M. C., LEE J., FISHER S., MACRI D.: Modeling hair using level-of-detail representations. In *Proceedings of the 16th International Conference on Computer Animation and Social Agents (CASA 2003)* (2003), CASA '03, pp. 41–47.
- [XESV97] XIA J. C., EL-SANA J., VARSHNEY A.: Adaptive real-time level-of-detail-based rendering for polygonal models. *IEEE Transactions on Visualization and Computer Graphics* 3, 2 (Apr. 1997), 171–183.

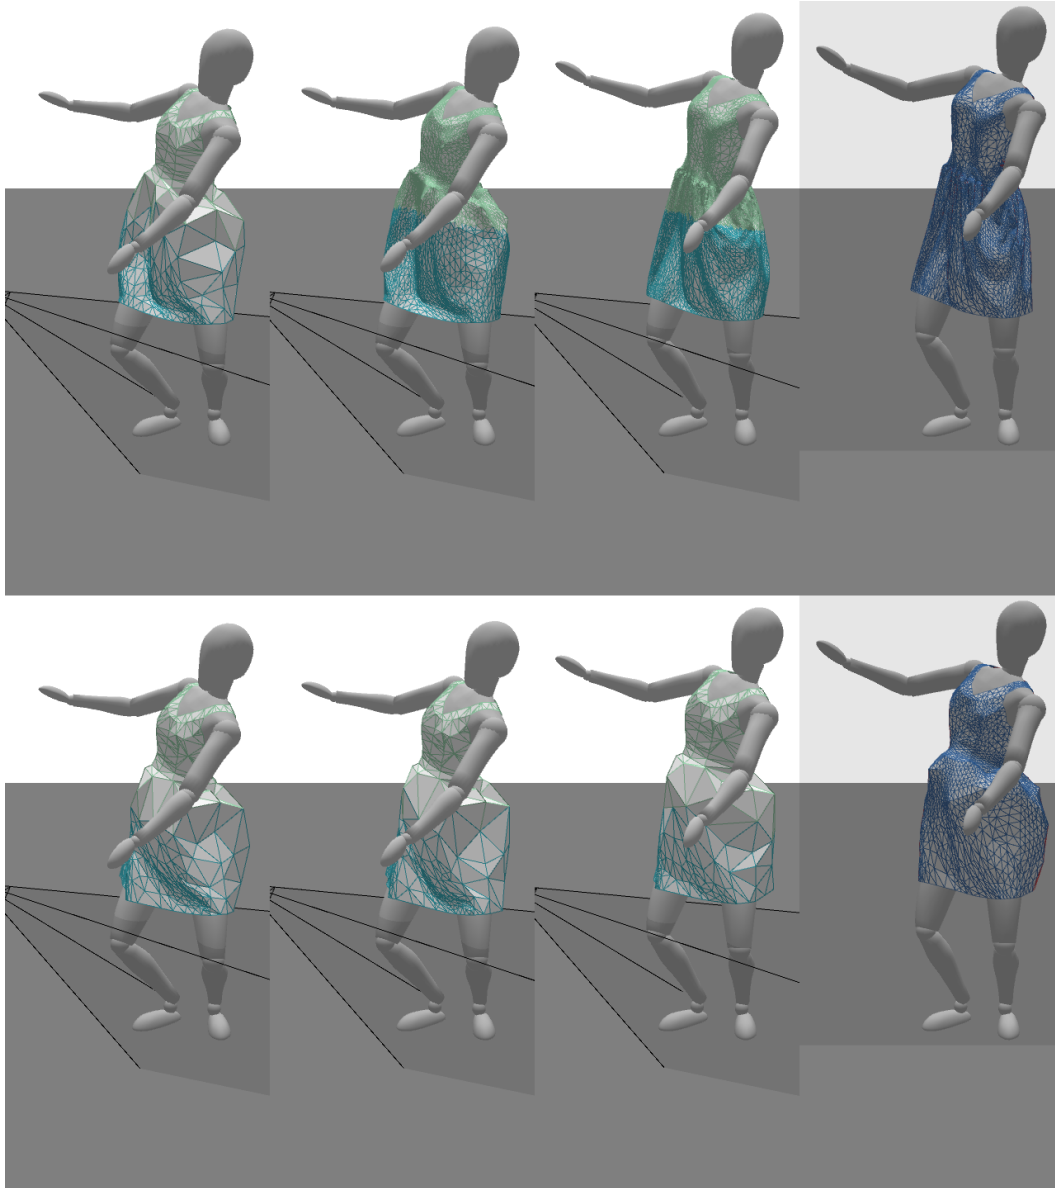


Figure 3: A sequence of frames with a jump cut, showing a dress simulated with temporal anticipation (above), and without (below). Blue edges indicate faces visible in the current view. Temporal anticipation allows the cloth to reach the correct high-resolution state before it becomes visible.

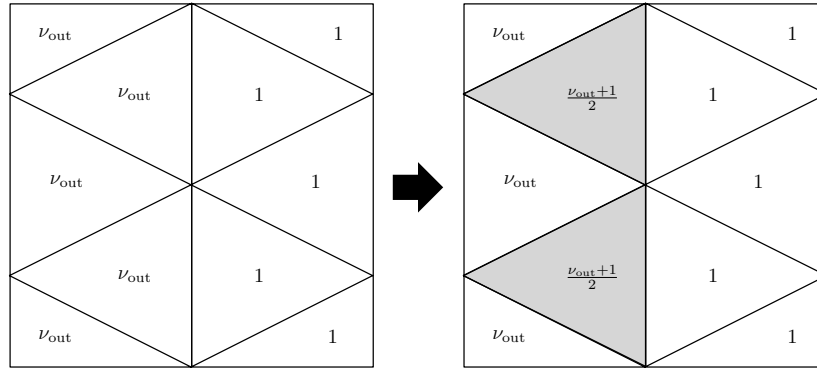


Figure 4: Locally interpolating discontinuous view factors at silhouettes to ensure smooth silhouettes.

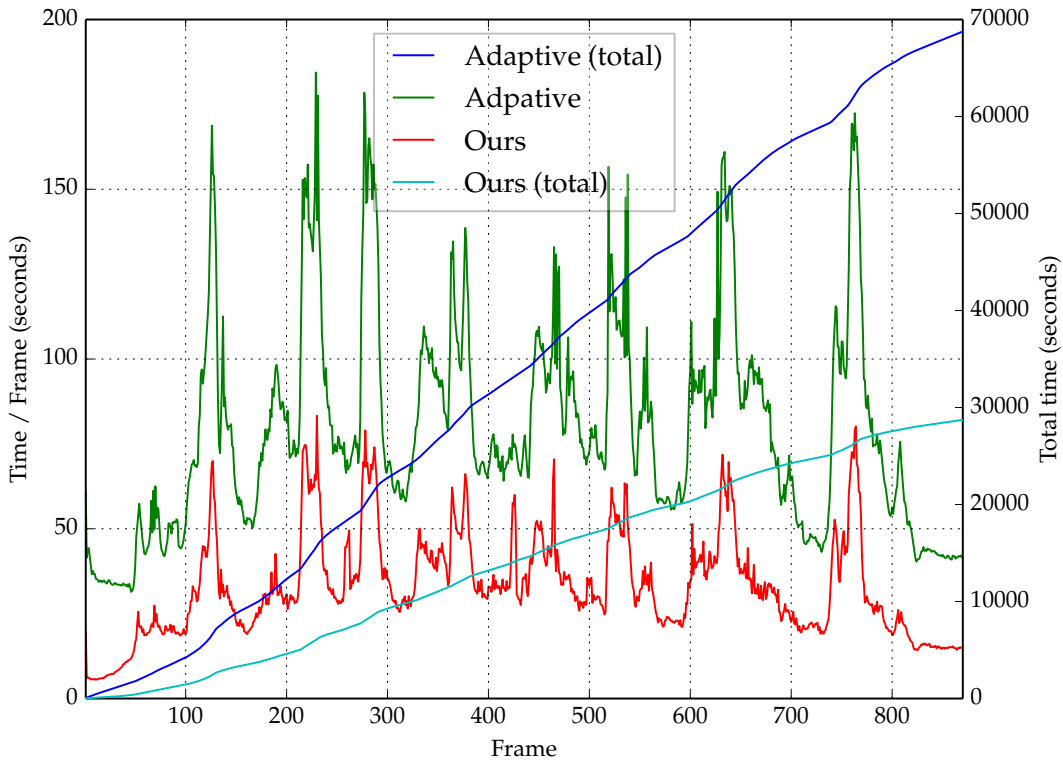


Figure 5: A comparison of view-independent and view-dependent adaptive simulations for the example of Figure 1 in terms of per-frame and cumulative simulation times.

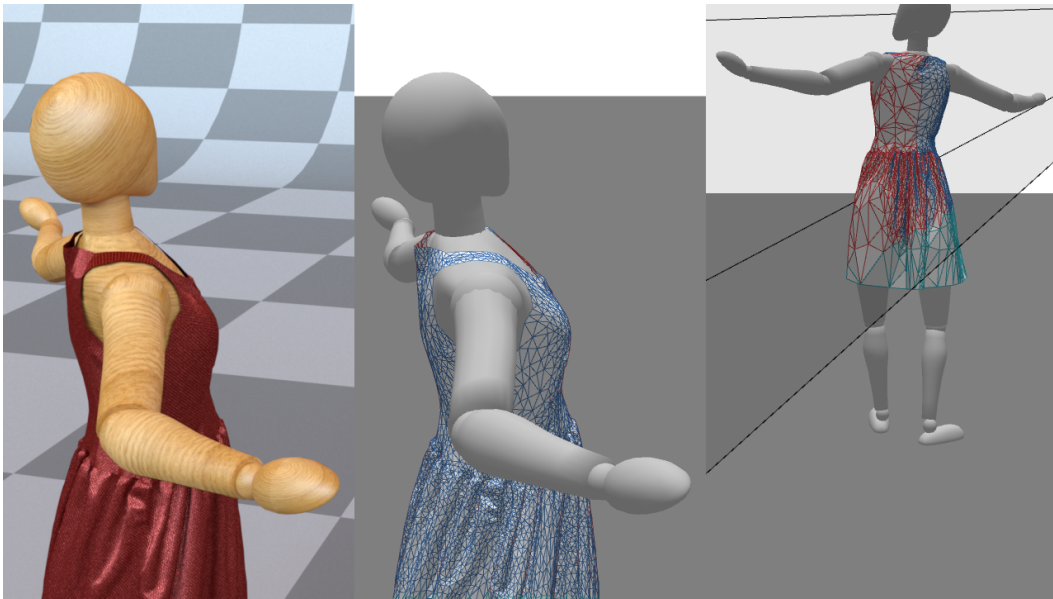


Figure 6: A solo dancer is shown from the camera's perspective (left, middle) and from an external perspective with the camera's view frustum drawn in black wireframe (right).



Figure 7: 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.