# An Analysis-by-Synthesis Approach to Tracking of Textiles

N. Hasler[1], B. Rosenhahn[1], M. Asbach[2], J.-R. Ohm[2], H.-P. Seidel[1]

Max-Planck-Institut für Informatik
Department 4: Computer Graphics
66123 Saarbrücken, Germany

Chair and Institute of Communications Engineering
RWTH Aachen
52056 Aachen, Germany

## Abstract

*Despite strong interest in cloth simulation on the one hand and tracking of deformable objects on the other, little effort has been put into tracking cloth motion by modelling the fabric. Here, an analysis-by-synthesis approach to tracking textiles is proposed which, by fitting a simulated textile to a set of contours, is able to reconstruct the 3D-cloth configuration. Fitting is accomplished by optimising the parameters of the mass-spring model that is used to simulate the textile on the one hand and the positions of a limited number of constrained points on the other. To improve tracking accuracy and to overcome the inherently chaotic behaviour of the real fabric several techniques for tracking features on the cloth's surface and the best way for them to influence the simulation are evaluated.*

## 1. Introduction

Cloth simulation has secured an important spot in today's film industry. The employed systems, however, mostly focus on producing believable animations, rather than realistic simulations. Simulations used for cloth tracking, on the other hand, should produce a motion as realistic as possible. The possibility of tracking fabrics by synthesising an observed scene with a mass-spring based cloth simulation is demonstrated in the following. It was found that, as a result of the inherently instable behaviour of cloth, improvements can be gained by introducing non-physical forces to bias the simulation towards following the observed behaviour.

The general approach to synthesising a scene can be summarised in following steps: Firstly, silhouettes are extracted and features on the cloth's surface are tracked and secondly, cloth properties and the positions of constrained points on the fabric are optimised to fit the simulation to observed silhouettes and tracked features. A configuration is evaluated by simulating it, using the difference between captured and simulated silhouettes as error function.

The remainder of this work is structured as follows. Section 2 briefly summarises the state of the art in cloth tracking and the development of the cloth simulation literature. Section 3 introduces the cloth model and summarises the contributions presented in this work. Section 4 contrasts several tracking procedures. The employed optimisation strategy is detailed in section 5, results are presented in section 6, and a summary is given in section 7.

## 2. Previous Work

Even though great interest has been shown in both cloth simulation and cloth tracking only little effort has been put into combining the techniques. In a previous publication we presented an algorithm that fits a cloth simulation to an observed rectangular cloth by optimising the parameters of the simulated cloth model [12].

Jojic and Huang [13] also employed an analysis-by-synthesis approach to estimate the hidden points a cloth is resting on. Their procedure depends on 3D-range data of the real cloth and the two-phased nature of their algorithm restricts its application to static situations.

The approach of Bhat et al. [2] is quite similar to our own, however, their goal is fundamentally different. They attempted to extract the static and dynamic fabric parameters from video images assuming knowledge of the exact positions of constrained points of the observed square fabric. Utilising a simulated annealing optimiser they were able to show that their setup allows them to reconstruct the parameters for a number of different fabric types. Their research, however, lacks verification with a mechanical cloth parameter estimation method such as the Kawabata Evaluation System [15]. The method has also been criticised because it is apparently unable to accurately estimate the bend resistance parameter [23].

Charfi et al. proposed a procedure for estimating the damping parameters of a textile by analysing the trajectories of a cloth that is dropped from a framework [7]. The movement of markers stuck on the surface of the fabric is recorded by a motion capture system. The viscous parameters can then be obtained by adjusting a simulation of the fabric to fit the observed behaviour. Since they found the global optimisation to be instable they iteratively estimate these parameters independently for every frame.

Other cloth tracking and reconstruction approaches do not make use of a cloth simulation to guide the reconstruction. The algorithm by Pritchard and Heidrich [23] can be divided into three stages. First, stereo correspondence is used to reconstruct most of the textured cloth. Then, holes are interpolated and Lowe's SIFT descriptor [17] is employed to map points on the world-space cloth to points on the two dimensional reference cloth. Identified points are finally connected by a seed-and-grow algorithm, rejecting spurious points in the process.

Several other algorithms to recover the three dimensional layout of a cloth were published. Scholz and Magnor [25] presented one approach that used optical flow to calculate the three dimensional scene flow. Holes in the model are not interpolated as in Pritchard's approach. Instead a deformable cloth model is matched to the surface, minimising the deformation energy of the patch. Drift is countered by constraining the edge of the simulation to the silhouette of the real cloth. Unfortunately, their algorithm was only demonstrated on synthetic data. Their work was continued with a publication on tracking cloth marked with a pseudo random coloured dots pattern [26]. The proposed algorithm detects coloured ellipses using colour and brightness information and identifies the exact position on the cloth by examining their local neighbourhood. The identified locations are connected in a way similar to Pritchard's approach. Three dimensional coordinates are reconstructed by using a multi-camera setup. As a last step holes are filled by means of a thin-plate spline interpolation technique.

Recently, an extension to the approach was presented by White et al. [29] who proposed a stereo-setup to reconstruct a random pattern of coloured triangles printed on a cloth. Their principal contributions are an extension to the seed-and-grow algorithm introduced by Pritchard and a strain minimisation technique that allows them to reconstruct points that are visible in one camera only.

The other important area of research employed in this work, particle system based cloth simulation, was pioneered by Terzopoulos et al. [27]. In their work a number of techniques that are common now such as semi-implicit integration, hierarchical bounding boxes, and adaptive time-step control were proposed. Until Baraff and Witkin reintroduced semi-implicit integration [1], decreasing the computational cost of cloth simulation significantly, explicit integration techniques were common.

In the last few years two major strands of development can be made out in the cloth simulation community. One, aiming for real-time simulation, focusses on computation speed alone, sacrificing expressiveness and accuracy if necessary. Desbrun et al. simplified the equation system that needs to be solved every step by precomputing parts of it [10]. Kang and Choi used a coarse mass-spring discretisation and added wrinkles in a post-processing step by interpolating with a cubic spline [14]. Oh et al. introduced a new semi-implicit integration technique that, besides sidestepping the unnatural damping of Baraff and Witkin's integrator, is reportedly able to run in real-time [21].

The other strand attempts to simulate cloth as realistically as possible. The use of nonlinear cloth properties has been introduced by Eberhardt et al. [11]. Simplified nonlinearities have since been integrated into a number of systems such as [9, 4]. Impressive results have been presented by Volino and Magnenat-Thalmann [28]. The fabric properties employed in their system are not only nonlinear but exhibit hysteretic behaviour.

# 3. Overview

Unlike most other cloth tracking approaches an analysis-by-synthesis technique is proposed here. That is a scene containing cloth motion is reconstructed by optimising the parameters of a cloth simulation so that it matches the observed data. This approach has several advantages. Firstly, parts of the real cloth that are temporarily hidden are still modelled by the simulation. Secondly, no interpolation of holes has to be performed. Thirdly, since a full cloth simulation is implemented, strain release as introduced by White et al. [29] becomes superfluous. A full cloth simulation has the additional advantage that the dynamic behaviour of the textile is integrated automatically into the reconstruction.

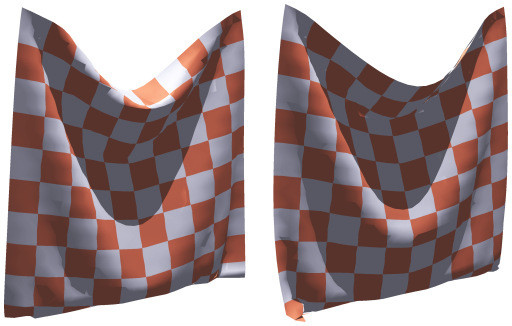

Figure 1: On the right a single mass-point deviates 2cm from the original configuration, severely impacting the final outcome of the simulation.

Unfortunately, the instability inherent in real as well as simulated cloth leads to the eventual divergence of real world and simulation even if initial condition and simulation are correct [9]. The problem is demonstrated in figure 1. The only difference between the two simulations of $1m \times 1m$ fabrics is that a single of 900 mass-points was offset by $2cm$ in the initial configuration of the right simulation run. This simple setup demonstrates that even minor deviations from a given original cloth configuration can

severely impact the final outcome. Besides inaccurate initial condition other influences such as coarse quantisation and numerical accuracy limit the achievable accuracy of the simulation.

This problem can be addressed in several ways but only one was deemed feasible and is consequently explained here. Virtual springs are attached to points on the simulated surface, attracting them to tracked world space coordinates. Although introducing non-physical forces into the simulation should be discouraged, because the ability of the system to reconstruct hidden geometry is lost, we still follow the approach here because this property is not needed within the scope of this project.

Linear springs ($\mathbf{f} \sim d$) as well as springs proportional to the inverse of the squared distance $d$ ($\mathbf{f} \sim 1/d^2$) are explored. The advantages of reciprocal forces are their inherent outlier rejection and the fact that the strongest forces hold points in place instead of accelerating them which could cause the simulation to become instable.

Unlike suggested in our previous work [12] tracked points are not just employed to establish attraction forces but become part of the evaluation function and consequently directly influence the optimisation. The other contribution presented in this work is an attempt to simplify the experimental setup by using fewer cameras. The loss of achievable accuracy is quantitatively examined. Additionally, primarily as a measure to speed up the optimisation, the method suggested by Charfi et al. [8] is explored. That is cloth properties are estimated independently for every frame. Figure 2 contrasts the two approaches. Our previous optimisation strategy [12] shown on the left has two loops, optimising cloth properties in an outer loop, because they are invariant during a sequence and positions of constrained points in an inner loop because they change every frame. The new strategy uses only one loop and estimates all parameters for every frame.

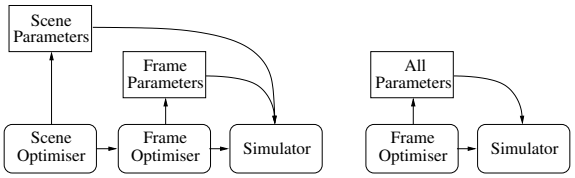


Figure 2: Two optimisation paradigms are evaluated.

### 3.1. Cloth Modelling

At the heart of the proposed tracking system a cloth simulation is responsible for integrating fabric dynamics into the tracking procedure, performing stress reduction and ensuring that all proposed configurations are legal. It is consequently an important module of the tracking system. A damped mass-spring system as used in most cloth simulations forms the basis of the system. The force $\mathbf{f_{spring}}$ on two connected mass-points is given by

$$\mathbf{f_{spring}} = \pm \left( k_s \frac{|\mathbf{\Delta x}| - l_0}{|\mathbf{\Delta x}|} + k_d \frac{\mathbf{\Delta x} \cdot \mathbf{\Delta v}}{|\mathbf{\Delta x}|^2} \right) \mathbf{\Delta x}$$

where $k_s$ and $k_d$ are stiffness and damping constants, $\mathbf{\Delta x}$ and $\mathbf{\Delta v}$ denote the differences between the two involved particles in position and speed respectively, while $l_0$ signifies the restlength of the connecting spring. Unlike proposed for some simple systems bending is not modelled by springs connecting every other particle but by quad-based bend resistances as proposed by Bridson et al. [6]. Additionally, a sophisticated air-resistance model, originally introduced by Bhat et al. [2], is employed by adding a non-linear normal and a tangential force to the corners of moving triangles,

$$\mathbf{f_{air}} = -A \left( k_n \frac{|\mathbf{v_n}|^2}{1 + k_f |\mathbf{v_n}|^2} \cdot \frac{\mathbf{v_n}}{|\mathbf{v_n}|} + k_t \mathbf{v_t} \right),$$

where $k_n$, $k_t$, and $k_f$ are material dependent coefficients and $A$ denotes the area of the triangle. Normal and tangential speeds $\mathbf{v_n}$ and $\mathbf{v_t}$ are computed by decomposing the average speed of a triangle normal and tangential to its surface. An additional measure which was originally introduced by Provot [24] to speed up the simulation can also be interpreted to provide a piecewise linear approximation of the non-linear behaviour of cloth. The method consists of applying additional impulses to particles connected by springs that have moved too far apart or too close together.

The collision detection is closely related to the system proposed by Bridson et al. [5]. The only major difference between our systems is that higher order discrete oriented polytopes (k-DOPs) are used instead of axis aligned bounding boxes. At heart k-DOPs are bounding volumes that use a fixed number of pairwise parallel planes to bound an arbitrary volume. The simplest form is the 6-DOP which is equivalent to an axis aligned bounding box. Higher order k-DOPs are ordinarily constructed by adding additional planes to the 6-DOP. The most common k-DOPs are displayed in figure 3. As discovered by Mezger et al. [20] 14-DOPs exhibit the highest performance when they are used for bounding cloth.

## 4. Tracking

Attraction forces can be defined to attract points of the simulated cloth either directly to features identified and tracked in the pattern of the cloth in one or more views of the scene or to an explicit sparse three-dimensional reconstruction of the cloth. Attraction to 3D-points is simpler and thus considered first. Features are detected and descriptors are cal-

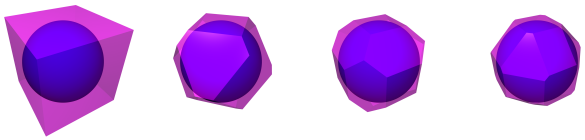

Figure 3: The most common discrete oriented polytopes are constructed by adding additional bounding planes to the axis aligned bounding box. Here 6, 14, 18, and 26-DOPs are displayed bounding a sphere.

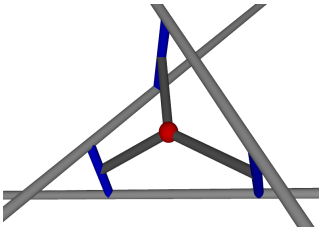

Figure 4: 3D-markers are generated by intersecting rays shot from camera origins through the corresponding image planes into the scene.

culated for them using the Scale Invariant Feature Transform (SIFT) [18].

In order to be able to attach attraction forces to the simulated cloth its initial position must be known. For all experiments it is assumed that the cloth is initially resting flat on the ground. Thus, an unambiguous mapping from camera coordinates into world coordinates can be defined and 2D-coordinates in cloth space $(u, v)$ can be assigned to features detected in the initial frame. Combining features from different views enforcing minimal distances and SIFT descriptor differences these markers can be re-identified in subsequent frames either using a tracking algorithm or the first frame as the reference.

The tracking procedure can be improved by restricting the distance a marker may move from one frame to the next in world and camera coordinates which is trivial to assert. It is additionally possible to enforce that markers cannot move apart. This constraint is based on the assumption that cloth cannot stretch more than a few percent. As the initial distances between tracked features are known, it is possible to reject markers that do not fit into the distance grid. Assuming that there are substantially less outliers than there are inliers, a voting approach can be taken to identify the former. That is the difference between the distances in world coordinates $d_{wij}$ between markers $i$ and $j$, $(i \neq j)$ within one frame and the corresponding distances in 2D-cloth coordinates $d_{cij}$ are used to reject markers which violate too many distance constraints.

Additionally, restricting the maximum distance a marker may move between consecutive frames eliminates a few spurious markers. This filter can be applied either in world coordinates to markers or in camera coordinates to features tracked from one frame to the next or to both. Applying it to features in camera coordinates has the advantage that it can be employed before markers are formed. An affected marker may then possibly be preserved which would otherwise have been rejected by the same criterion in world space. However, it may not be possible to catch all markers in camera coordinates so it is desirable to run the filter in world space as well.

Rays identifying features are cast from camera origins through the corresponding image planes. Several such rays from different cameras identifying the same feature ideally intersect at exactly one point in world space. In practice, however, due to inaccuracies capturing the images, calculating camera parameters, and detecting the features, they do not necessarily do so. Thus, marker position $\mathbf{p}$ is derived from involved rays by computing

$$\mathbf{p} = \frac{1}{N} \sum_{i=1}^{N} \sum_{j=1}^{N} \mathbf{c_{ij}} \qquad i \neq j$$

where $N$ denotes the number of involved rays and $\mathbf{c_{ij}}$ is the position of the closest point on ray $i$ to ray $j$,

$$\mathbf{c_{ij}} = \frac{\mathbf{d_i} \cdot (\mathbf{d_j} \times \mathbf{m_j}) - (\mathbf{d_i} \cdot \mathbf{d_j})\,(\mathbf{d_i} \cdot (\mathbf{d_j} \times \mathbf{m_i}))}{(\mathbf{d_i} \times \mathbf{d_j}) \cdot (\mathbf{d_i} \times \mathbf{d_j})} \mathbf{d_i} + \mathbf{d_i} \times \mathbf{m_i}.$$

Here $\mathbf{d}$ and $\mathbf{m}$ denote direction and momentum of rays defined in Plücker Coordinates [3].

### 4.1. 2D-Tracking

Due to the excessive loss of markers which is primarily a result of the difficulty of combining features from distinct cameras a different approach is described here which completely avoids this step. The difficulty of identifying features between cameras is a result of the large angles between the utilised cameras which causes substantial changes to the local feature descriptors when they are transformed from one camera into an other on the one hand and the repetitive nature of the fabric's pattern on the other.

Thus, a different class of approaches is proposed here which entirely omits the explicit formation of 3D-markers. The result is that a considerably larger number of features can be tracked. Although each of them only constrains a point on the cloth's surface to lie on a ray, the increased number of features and thus constraints improves the result.

The procedure works as follows: For the first frame the same method for identifying the cloth-space coordinates of the identified features is used as introduced above for the three dimensional tracker. Then these features are tracked

independently in every view using a variant of the common Lucas-Kanade algorithm [19].

Due to the weaker constraint of this algorithm a possibly less intuitive approach has to be taken when calculating attraction forces. Simulated points are not attracted to fixed world coordinates any more but to arbitrary rays in world space. However, when using Plücker Coordinates to represent the feature rays the calculation of forces perpendicular to these rays becomes achievable. Assume for example that the ray, a feature at world coordinates $\mathbf{p}$ is attracted to, in Plücker Coordinates is represented by a normalised direction $\mathbf{d_0}$ and a moment $\mathbf{m_0}$. The shortest vector $\mathbf{v}$ pointing from the ray to the point can then simply be calculated by

$$\mathbf{v} = (\mathbf{m_0} - \mathbf{p} \times \mathbf{d_0}) \times \mathbf{d_0}.$$

# 5. Optimisation

After the general approach and the procedure for introducing attraction forces have been presented an optimisation strategy is proposed in this section.

As the fabric properties do not change during a sequence but the positions of constrained points of the cloth are valid for a single frame only, the problem can intuitively be divided into two loops (see figure 2). In the outer loop parameters that are scene invariant, are optimised. The inner loop optimises parameters that change every frame. For both optimisers a simulated annealing algorithm [22] was initially used. However, as only a sequential implementation was available and Kolda's asynchronous parallel pattern search [16] was found to produce similar results in significant less time, it was employed for all experiments described below.

The error of the inner optimisation loop $E_{sf}$ is calculated by summing the differences between the silhouettes of the simulated cloth and the measured silhouettes of all employed cameras. The outer optimisation uses the error function $E_s$ that sums the error of the inner error function $E_{sf}$ for all frames. Here $\mathbf{x}(f)$ is the state vector of a particular local configuration containing the values of all global parameters and the local parameters of frame $f$ whereas $\mathbf{x}$ comprises all global and local parameters.

$$E_{sf}(\mathbf{x}, f) = \sum_{c=1}^{N_c} |\mathbf{S_s}(f, \mathbf{T}(c), \mathbf{x}) - \mathbf{S_m}(f, \mathbf{T}(c), \mathbf{x})|$$

$$E_s(\mathbf{x}) = \sum_{f=1}^{N_f} E_{sf}(\mathbf{x_f})$$

Here $\mathbf{S_s}$ and $\mathbf{S_m}$ are the binary simulated and measured silhouettes respectively as a function of the frame number $f$, the transformation matrix $\mathbf{T}(c)$ of camera $c$, and $N_f$ and $N_c$ are the number of frames and cameras respectively.

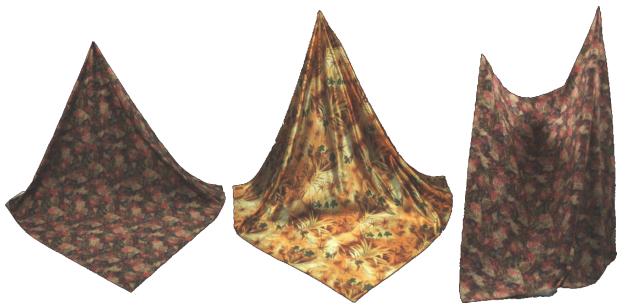

Figure 5: Characteristic frames from the three experiments

The initial condition of the optimisation and the number and positions of constrained particles are assumed to be known a priory. Three experiments were conducted using two different fabrics. Experiments 1 and 2 consisted of lifting the cloth at one of its corners whereas in experiment 3 it was lifted at two of its corners. One of two different square fabrics each with a side-length of $1.4m$ were captured by 7 calibrated, synchronised cameras at a framerate of $30fps$. Characteristic frames are shown in figure 5.

# 6. Results

In figure 6 the remaining errors after performing optimisations with different attraction force stiffnesses normalised by the error of optimising without attraction forces are shown. The shown results were obtained using the 3D and 2D-tracking approaches described in section 4.

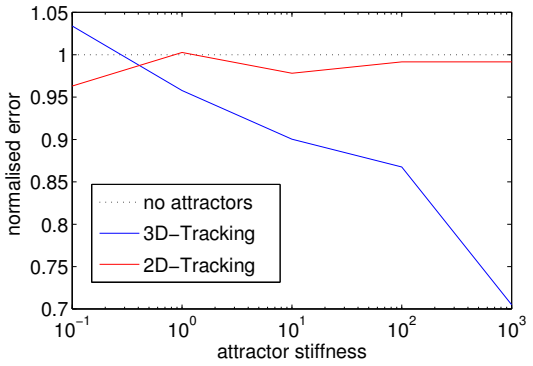


Figure 6: Comparison of different tracking mechanisms employing linear attraction forces.

Strong linear attraction forces using 3D-tracking produce significantly better fits. However, these results do not look natural as shown in figure 7(a). An optimisation that employs reciprocal attraction forces on the other hand is visually a lot more pleasing (see figure 7(c)) even though the

absolute error of the optimisation is larger. As a comparison two different views of the frame are displayed in figure 8. An additional drawback of attraction forces in general and the linear variant in particular is that they negatively influences the ability of the system to estimate underlying geometry. However, this problem is inconsequential in the current project as there is no underlying geometry that needs to be estimated.

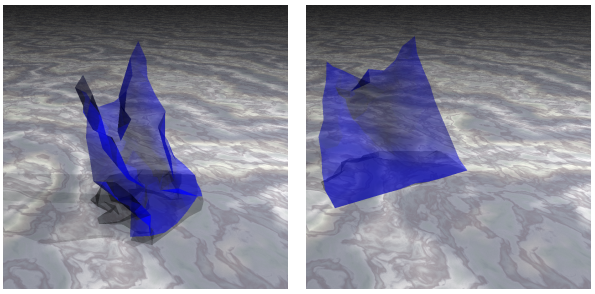

(a) The cloth crumbles when strong linear forces are used.

(b) Using only one camera causes the simulation to drift.

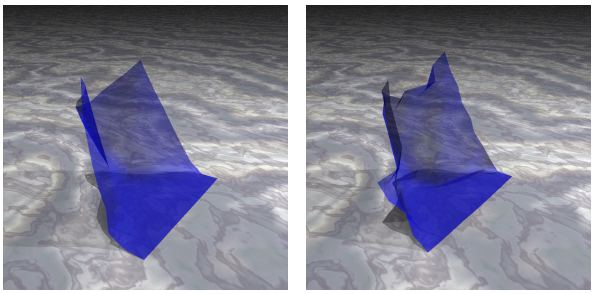

(c) Reciprocal forces produce reasonable results.

(d) Results created with three cameras are almost equivalent.

Figure 7: In this experiment a square piece of cloth is lifted from the ground at two of its corners. Reconstructions of frame 410 are given using different variants of the presented algorithm.

## 6.1. Fabric Parameters Estimated Locally

Estimating global parameters independently for every frame can be expected to decreases the overall error of a sequence because the space of possible configurations is larger than that of the original optimiser and in some frames better configurations can probably be found if global parameters may be modified locally. The proposed setup is in this way similar to the experiments conducted by Charfi et al. [8] who

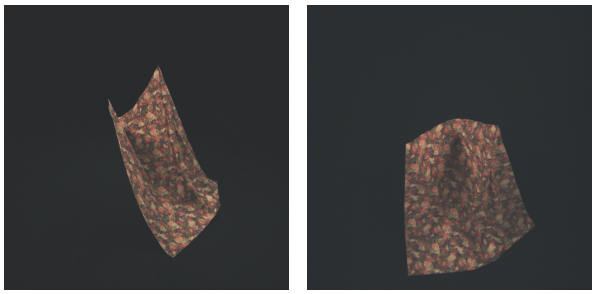

(a) side view

(b) top view

Figure 8: Two original views of the same frame as shown in figure 7.

tried to estimate dynamic parameters of cloth. Another expected result is that the optimisation is significantly faster because, although the number of parameters considered in a single frame is larger, every frame has to be optimised only once instead of, on average, approximately 50 times.

Changing fabric properties during the course of a sequence is obviously physically not justifiable. The primary objective of this setup is, however, to determine experimentally whether the optimiser chooses to change the global parameters considerably during the course of a sequence and whether the expected improvements are attainable at all. Additionally, although the inner optimisation loop has to handle more variables the optimisation is expected to be significantly faster because the outer optimisation loop can be dropped entirely.

Employing no attraction forces it was, against expectation, found that the attainable errors are almost twice as high as those obtained with the two-loop version. This effect can be attributed to the increased complexity of the search space which probably overwhelmed the optimiser. It could, however, be established that the optimisation was significantly faster (factor 10) than the original setup.

## 6.2. Augmented Evaluation Function

Apart from attracting simulated points to them, tracked features can more passively be used by augmenting the objective function that guides the optimisation process. The error function can be written as

$$E_{\alpha,\beta}(\mathbf{x}) = \alpha E_s(\mathbf{x}) + \beta E_a(\mathbf{x}) \qquad \alpha, \beta \geq 0$$

$\alpha$ and $\beta$ are weighting factors of the silhouette error function $E_s(\mathbf{x})$ and the attractor error function $E_a(\mathbf{x})$. Again $\mathbf{x}$ is a vector of global and local parameters. $E_a(\mathbf{x})$ measures the error of a state by calculating the average length of all attractors available in one frame.

The optimal combination of $\alpha$ and $\beta$ is searched for by optimising the test sequences using different combinations of $\alpha$ and $\beta$. Figure 9 shows the normalised error $\hat{E}_{\alpha,\beta}$ which is defined as

$$\hat{E}_{\alpha,\beta} = \frac{E_{\alpha,\beta}}{\alpha \cdot E_{1,0} + \beta \cdot E_{0,1}}$$

for different values of $\alpha$. This function is obviously one when either $\alpha$ or $\beta$ are one while the other is zero. For other combinations, however, it shows whether the achievable result are better (smaller) than expected by a linear combination of $E_{1,0}$ and $E_{0,1}$ or worse.

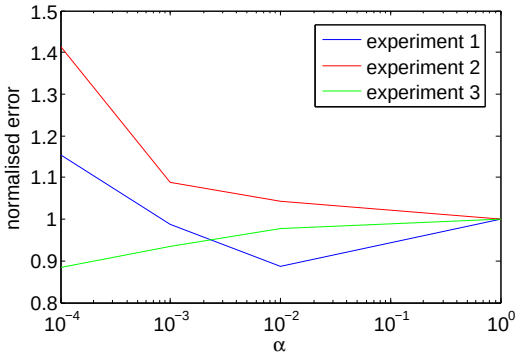


Figure 9: The normalised error $\hat{E}_{\alpha,\beta}$ calculated for different values of $\alpha$ while $\beta$ is held at 1.

The results in figure 9 show that for experiments 1 and 3 improvements are possible by enhancing the evaluation function. For experiment 2 a reflective, repetitive pattern was used, increasing the difficulty of tracking and 3D-reconstruction. The inferior quality of the reconstruction is probably the reason that the greater the influence of these points the worse the result.

### 6.3. Simplification of the Experimental Setup

In the previous experiments seven cameras were used, six positioned along the perimeter of the room and one suspended from the ceiling. This number of cameras already requires a sophisticated hardware infrastructure. So it is interesting to investigate whether using fewer cameras produces acceptable results already.

Figure 10 shows that a reasonable linearity between the number of cameras and the achievable residual errors can be observed, suggesting that a smaller number of cameras is sufficient for reconstructing the observed scene well. Although one camera is, as shown in figure 7(b), not adequate because no depth information is available three cameras already produce convincing results (see figure 7(d)).

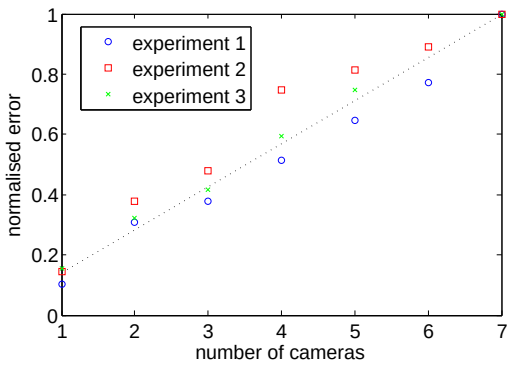


Figure 10: Optimisation error as a function of the number of employed cameras normalised to the error achieved when seven cameras are used.

## 7. Summary

In this work an analysis-by-synthesis approach to tracking of textiles was presented. While it was found that global parameters cannot reliably be estimated for every frame it was possible to show that by entwining attraction forces into the evaluation function some improvements can be gained. Additionally, for the simple experiments that were conducted a good linearity between the number of cameras and the residual optimisation error could be established. This finding obviously does not hold for more complicated scenes where deep folds occlude parts of the textile.

The most critical problem when more complicated scenes and cloth models are tackled is that the computational cost of the current implementation is prohibitive. Running the optimisation in parallel on 7 AMD Opteron processors at 2.2GHz the algorithm takes from 10h to 30h to converge. So at least some of the further effort will have to go into improving the speed of the procedure.

This is how to do an unnumbered section (note asterisk).

## Acknowledgements

We gratefully acknowledge funding by the Max-Planck Center for Visual Computing and Communication.

## References

[1] D. Baraff and A. Witkin. Large steps in cloth simulation. In *Proceedings of ACM SIGGRAPH 98*, pages 43–54. ACM Press, 1998.

[2] K. Bhat, C. Twigg, J. Hodgins, P. Khosla, Z. Popović, and S. Seitz. Estimating cloth simulation parameters from video. In *Proceedings of ACM SIGGRAPH/Eurographics Sympo-*

*sium on Computer Animation (SCA 2003)*, pages 37–51. ACM Press, 2003.

[3] W. Blaschke. *Kinematik und Quaternionen, Mathematische Monographien 4*. Deutscher Verlag der Wissenschaften, 1960.

[4] R. Bridson. *Computational aspects of dynamic surfaces*. PhD thesis, Stanford University, 2003.

[5] R. Bridson, R. Fedkiw, and J. Anderson. Robust treatment of collisions, contact and friction for cloth animation. *ACM Transactions on Graphics (ACM SIGGRAPH 2002)*, 21(3):594–603, July 2002.

[6] R. Bridson, S. Marino, and R. Fedkiw. Simulation of clothing with folds and wrinkles. In *Proceedings of ACM SIGGRAPH/Eurographics Symposium on Computer Animation (SCA 2003)*, pages 28–36. ACM Press, 2003.

[7] H. Charfi, A. Gagalowisz, and R. Brun. Determination of fabric viscosity parameters using iterative minimization. In *Computer Analysis of Images and Patterns: 11th International Conference*, pages 789–798, Versailles, France, September 2005.

[8] H. Charfi, A. Gagalowisz, and R. Brun. Measurement of fabric viscosity. In *Mirage*, pages 261–268, Rocquencourt, France, March 2005.

[9] K.-J. Choi and H.-S. Ko. Stable but responsive cloth. *ACM Transactions on Graphics (ACM SIGGRAPH 2002)*, 21(3):604–611, July 2002.

[10] M. Desbrun, P. Schröder, and A. Barr. Interactive animation of structured deformable objects. In *Proceedings of Graphics Interface (GI 1999)*, pages 1–8. Canadian Computer-Human Communications Society, 1999.

[11] B. Eberhardt, A. Weber, and W. Straßer. A fast, flexible, particle-system model for cloth draping. *IEEE Computer Graphics and Applications*, 16(5):52–59, September 1996.

[12] N. Hasler, M. Asbach, B. Rosenhahn, J.-R. Ohm, and H.-P. Seidel. Physically based tracking of cloth. In *Proceedings of the International Workshop on Vision, Modeling and Visualization 2006*, Aachen, Germany, November 2006.

[13] N. Jojic and T. Huang. On analysis of cloth drape range data. In *ACCV '98: Proceedings of the Third Asian Conference on Computer Vision-Volume II*, pages 463–470, London, UK, 1997. Springer-Verlag.

[14] Y.-M. Kang and H.-G. Cho. Bilayered approximate integration for rapid and plausible animation of virtual cloth with realistic wrinkles. In *Proceedings of Computer Animation*, pages 203–214. IEEE Computer Society, 2002.

[15] S. Kawabata. *The standardization and analysis of hand evaluation*. The Textile Machinery Society of Japan, 1980.

[16] T. Kolda. Revisiting asynchronous parallel pattern search. Technical Report SAND2004-8055, Sandia National Laboratories, Livermore, CA 94551, February 2004.

[17] D. Lowe. Object recognition from local scale-invariant features. In *International Conference on Computer Vision*, pages 1150–1157, September 1999. SIFT first introduced.

[18] D. Lowe. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, 60(2):91–110, 2004.

[19] B. Lucas and T. Kanade. An iterative image registration technique with an application to stereo vision (darpa). In *Proceedings of the 1981 DARPA Image Understanding Workshop*, pages 121–130, April 1981.

[20] J. Mezger, S. Kimmerle, and O. Etzmuß. Progress in collision detection and response techniques for cloth animation. In *Proceedings of 10th Pacific Conference on Computer Graphics and Applications (PG 2002)*, pages 444–445. IEEE Computer Society, 2002. taugt nix - zu kurz.

[21] S. Oh, J. Ahn, and K. Wohn. Low damped cloth simulation. *The Visual Computer (to appear)*, 2005.

[22] W. Press, W. Vetterling, S. Teukolsky, and B. Flannery. *Numerical Recipes in C++: the art of scientific computing*. Cambridge University Press, 2nd edition, 2002.

[23] D. Pritchard and W. Heidrich. Cloth motion capture. *Computer Graphics Forum (Eurographics 2003)*, 22(3):263–271, September 2003.

[24] X. Provot. Deformation constraints in a mass-spring model to describe rigid cloth behavior. In *Proceedings of Graphics Interface (GI 1995)*, pages 147–154. Canadian Computer-Human Communications Society, 1995.

[25] V. Scholz and M. Magnor. Cloth motion from optical flow. In B. Girod, M. Magnor, and H.-P. Seidel, editors, *Proc. Vision, Modeling and Visualization 2004*, pages 117–124, Stanford, USA, November 2004. Akademische Verlagsgesellschaft Aka GmbH.

[26] V. Scholz, T. Stich, M. Keckeisen, M. Wacker, and M. Magnor. Garment motion capture using color-coded patterns. *Computer Graphics Forum (Proc. Eurographics EG'05)*, 24(3):439–448, August 2005.

[27] D. Terzopoulos, J. Platt, A. Barr, and K. Fleischer. Elastically deformable models. In *Computer Graphics (Proceedings of ACM SIGGRAPH 87)*, pages 205–214. ACM Press, July 1987.

[28] P. Volino and N. Magnenat-Thalmann. Accurate garment prototyping and simulation. *Computer-Aided Design Applications*, 2(5):645–654, 2005.

[29] R. White, D. Forsyth, and J. Vasanth. Capturing real folds in cloth. Technical Report UCB/EECS-2006-10, EECS Department, University of California, Berkeley, February 2006.