

TITLE: PARALLELIZED INCOMPLETE POISSON PRECONDITIONER IN CLOTH SIMULATION

ABSTRACT:

Efficient cloth simulation is an important problem for interactive applications that involve virtual humans, such as computer games. A common aspect of many methods that have been developed to simulate cloth is a linear system of equations, which is commonly solved using conjugate gradient or multi-grid approaches. In this paper, we introduce to the computer gaming community a recently proposed preconditioner, the incomplete Poisson preconditioner, for conjugate gradient solvers. We show that the parallelized incomplete Poisson preconditioner (PIPP) performs as well as the current state-of-the-art preconditioners, while being much more amenable to standard thread-level parallelism. We demonstrate our results on an 8-core Apple* Mac* Pro and a 32-core code name Emerald Ridge system.

PARALLELIZED INCOMPLETE POISSON PRECONDITIONER IN CLOTH SIMULATION:

Simulating flexible materials, such as cloth, is an important task for applications involving virtual humans such as computer games and visual effects. High quality offline simulations are achieved by using implicit methods for simulating cloth [20–22]. Real-time applications, on the other hand, use explicit or semi-explicit methods for cloth simulation in order to meet time constraints [5, 17]. Despite the decades of research on simulating flexible materials, the efficient simulation of cloth remains an important challenge for computer animation. A large amount of research exists that addresses algorithmic optimizations for speeding up implicit integration methods for simulating cloth. The use of preconditioners [4, 9, 13] has been shown to greatly reduce the number of iterations of the conjugate gradient method in an effort to achieve convergence.

In this paper, we explore the use of a novel preconditioning scheme – the incomplete Poisson preconditioner – that has not been used before in clothing simulation. Using a variety of standard benchmarks, we first demonstrate that this preconditioner is just as good as, if not better, than currently used methods. A major advantage of this method is that it is extremely easy to parallelize and can take advantage of the processing power available in current and next generation multi-core hardware. Current state-of-the-art

preconditioners [13] are not as suitable for parallelization and do not scale well with increase in computational resources. This paper makes the following contributions:

1. To our knowledge, we propose for the first time the use of the incomplete Poisson preconditioner for clothing simulation.
2. We compare the incomplete Poisson preconditioner to the most commonly used preconditioning methods in terms of efficiency, quality and ease of parallelization.
3. We demonstrate that a parallel implementation of the incomplete Poisson preconditioner (PIPP) achieves significant performance improvement on multi-core computers.
4. We demonstrate the scalability of the PIPP on a state-of-the-art, 32-core compute server and show that it is ready for the next generation of hardware resources.

The rest of this document is organized as follows. We start by reviewing related work. Subsequently we present an overview of the method we use for simulating cloth. We describe the Jacobi preconditioner, the symmetric successive over-relaxation, and the incomplete Cholesky preconditioner which are commonly used to accelerate convergence. In addition, we propose the use of the incomplete Poisson preconditioner for cloth simulation. We compare the four preconditioning methods on four standard benchmarks and also demonstrate the effectiveness of parallelizing the incomplete Poisson preconditioning scheme. Finally, we conclude with a discussion of future work.

Related Work

Early work by [20–22] has applied techniques from mechanical engineering and finite element communities to cloth simulation. Since then, there has been an extensive amount of work by different research groups [5, 7, 10, 23] that have addressed several aspects of simulating cloth. An extensive overview of cloth simulation techniques can be found in two survey papers [8, 16]. Preconditioners play a very important part in implicit cloth simulation as they can greatly speed up convergence of numerical methods. The work in [3] used a simple diagonal preconditioner for the modified preconditioned conjugate gradient method (MPCG). The work in [9] demonstrated 20% speedup by using a 3×3 block diagonal preconditioner. This work was extended in [4] by proposing an approximation of the filter matrix A of the MPCG. The work in [13] demonstrates the effectiveness of the incomplete Cholesky and successive symmetric over-relaxation (SSOR) preconditioning schemes by reducing the number of iterations by 20%.

Relation to Prior Work. In this paper, we first examine the fitness of three commonly used preconditioning schemes [3, 13] in comparison to the proposed incomplete Poisson preconditioner. Our simulation method is similar to the implicit simulation method described in [2,3]. We perform collision detection using distance fields [11]. Collision resolution is performed using the techniques described in [6] and [18].

Cloth Simulation Overview

There are many aspects to cloth simulation. A cloth simulator is required to solve a linear system of equations which is used to step the simulator forward by one time step. This system of equations is derived taking into account the specifics of the internal forces and their derivatives. Different soft and hard constraints are imposed on the simulation which must be met. Collision detection and resolution is another area of research that has many contributions. We refer the reader to excellent works in cloth simulation research [3, 12, 14, 15] for more information. In this paper, we focus on the methods of preconditioning that are used to accelerate the preconditioned conjugate gradient solver. In this section, we present an overview of the preconditioned conjugate gradient solver and describe the different methods of preconditioning for cloth simulation.

Preconditioned Conjugate Gradient Solver

An overview of the preconditioned conjugate gradient solver is shown in Algorithm 1. A detailed description of the algorithm can be found here [19]. The preconditioned conjugate gradient method takes as input the following: (a) a symmetric positive semi-definite matrix A , (b) a symmetric positive definite preconditioning matrix P of the same dimension as A , and (c) a vector b . The algorithm iteratively solves the linear system of equations, $Ax = b$ and the iterations stop when $|b - Ax| < \epsilon |b|$, where ϵ is a user-defined tolerance value. The preconditioning matrix P , which must be easily invertible, speeds convergence to the extent that P^{-1} approximates A .

Preconditioning Methods

We examine the performance of three commonly used preconditioning methods: (1) diagonal, (2) symmetric successive over-relaxation (SSOR), and (3) incomplete Cholesky against the unconditioned conjugate gradient method. We also examine a new preconditioning scheme, the incomplete Poisson preconditioner, proposed by Ament et al. [1] for the Poisson problem. Their motivation was to find an easily parallelizable preconditioner for simulations on multi-GPU systems. To the best of our knowledge, this is the first time this preconditioning scheme has been applied to cloth simulation. The mathematical formulation of these preconditioners is as follows:

Diagonal (Jacobi) Preconditioner:

$$P = \text{diag}\{A\}^{-1} \text{ or } P_{i,i} = \frac{1}{A_{i,i}}$$

This simple preconditioning scheme approximates the inverse of a diagonal matrix. Although lacking in quality, it can be computed quickly and provides increase in performance in many cases. The computation of P^{-1} is relatively simple and this preconditioner can be subsequently applied using SpMV.

Procedure **Preconditioned Conjugate Gradient Solver**(A, x, b,P, ϵ)

Input: A: Left hand side of linear system of equations $Ax = b$.

Input: x: Input constraint.

Input: b: Right hand side of linear system of equations $Ax = b$.

Input: P: Preconditioner

Input: ϵ : Maximum tolerance

Output: x: Result.

// Initialization

$\mathbf{r} = \mathbf{b} - \mathbf{Ax}$; // residual

$\mathbf{d} = \mathbf{P}^{-1} \cdot \mathbf{r}$;

$\mathbf{d}_{\text{new}} = \mathbf{r} \cdot \mathbf{d}$;

while $i < \text{MAX} \wedge \mathbf{d}_{\text{new}} > \epsilon^2$ **do**

$\mathbf{q} = \mathbf{A} \cdot \mathbf{d}$;

$c = \mathbf{d} \cdot \mathbf{q}$; // curvature

if $c < 0$ **then**

return FAIL;

else if $c == 0$ **then**

break;

end

$\alpha = \frac{\mathbf{d}_{\text{new}}}{c}$;

$\mathbf{x} = \mathbf{x} + \alpha \cdot \mathbf{d}$;

$\mathbf{r} = \mathbf{r} - \alpha \cdot \mathbf{q}$;

$\mathbf{s} = \mathbf{P}^{-1} \cdot \mathbf{r}$;

$\mathbf{d}_{\text{old}} = \mathbf{d}_{\text{new}}$;

$\mathbf{d}_{\text{new}} = \mathbf{r} \cdot \mathbf{s}$;

if $\mathbf{d}_{\text{new}} < 0$ **then**

break;

end

$\beta = \frac{\mathbf{d}_{\text{old}}}{\mathbf{d}_{\text{new}}}$;

$\mathbf{d} = \mathbf{s} + \beta \cdot \mathbf{d}$;

$i = i + 1$;

end

if $\mathbf{d}_{\text{new}} < 0 \vee i == \text{MAX}$ **then**

return FAIL;

else

return SUCCESS;

Algorithm 1: *Preconditioned Conjugate Gradient Solver*

Incomplete Cholesky Preconditioner:

$$P = (LL^T)^{-1}$$

where L is the Cholesky factorization defined as follows:

$$L_{i,i} = \sqrt{A_{i,i} - \sum_{k=1}^{i-1} L_{i,k}^2}$$
$$L_{i,j} = \frac{1}{L_{i,i}} (A_{i,j} - \sum_{k=1}^{i-1} L_{i,k} L_{j,k}), i > j$$

with the additional constraint to keep the original sparsity pattern of A . The incomplete Cholesky is derived from the Cholesky decomposition method. A symmetric positive-definite matrix can be decomposed into the product of a lower triangular matrix and its conjugate transpose. These triangular matrices can quickly be inverted in order to solve linear systems. In that sense, the incomplete Cholesky preconditioner approximates the full inverse of A without incurring the cost of actually inverting it. It should be noted that P^{-1} is calculated using expensive forward and backward substitutions, which are inherently serial processes because of the triangular structures of L and L^T .

Incomplete Poisson Preconditioner:

$$P = HH^T$$

Where

$$H = I - L * \text{diag}\{A\} - 1$$

and L is the strictly lower triangular matrix of A . This novel preconditioner has a simple structure and is kind of an approximate inverse. As a result, no substitutions are required and this preconditioner can be applied efficiently with SpMV and thread-level parallelism.

Symmetric Successive Over-Relaxation:

$$P = (M1 * M2)^{-1}$$

Where

$$M1 = \frac{1}{\omega} * D + L$$

$$M2 = \frac{1}{(2 - \omega)} * (I + \omega * D^{-1} + U)$$

and L,U,D are the strictly lower triangular, the strictly upper triangular and the diagonal matrix of A respectively. Symmetric successive over-relaxation is a variant of the Gauss-Seidel method but with improved convergence speed. As with incomplete Cholesky, a relatively expensive forward and backward substitution step occurs to calculate P-1. It should also be noted that the choice of ω influences convergence. We use the following ω :

$$\omega = \frac{1}{\max([1, \max(L), \max(U)])}$$

Evaluation Results

In this section we compare the proposed incomplete Poisson preconditioner to the most commonly used preconditioners. First, we describe the test cases we use for the comparison. Second, we evaluate the fitness of each of the preconditioning methods. Third, we provide the results of parallelizing the incomplete Poisson preconditioner on next-generation multi-core hardware. A visual comparison of using each of the preconditioners on the benchmarks can be seen in the accompanying video. All preconditioners seem to produce results of similar quality.

Benchmarks

We use four benchmarks for the purpose of exercising the preconditioners on a variety of challenging scenarios that are frequently encountered in simulating cloth. These four benchmarks are described below.

- 1. Free Fall.** This is more of a baseline case, where a piece of cloth falls under gravity and come to rest on a static sphere with no tangling (Figure 3(a)).
- 2. Curtain.** This case extends the previous benchmark by including fixed point constraints (Figure 3(b)).
- 3. Moving Collider.** Further extending the previous case, a cloth patch hung as a curtain interacts with a moving spherical collider (Figure 3(c)). This benchmark is used to test the behavior of the simulator in a dynamic environment.
- 4. Tangling.** Tangling is one of the toughest cases for cloth simulators to handle because of the complexities introduced by the multiple self-collisions (Figure 3(d)). As far as the conjugate gradient solver is concerned, for tangled states the number of nonzeros (thus the stiffness) of the matrix A($Ax = b$) increases significantly. The increased matrix density can significantly affect performance.

Preconditioner Evaluation

We test the performance of the preconditioners by simulating 200 frames for each of the benchmarks described above. The parameters used for the cloth simulator are described in Table 1. The inter-particle forces were shear, bend and stretch. The evaluation results are illustrated in Figure 4. From the results, it is apparent that the incomplete Poisson preconditioner performs on par with incomplete Cholesky for cloth simulation. Table 2 and Figure 1 illustrate the performance results of all preconditioning schemes with increase in number of nodes on cloth patch. Here, we see that PIPP scales well with increase in resolution of cloth patch, but the incomplete Cholesky preconditioner does not. The main advantage of this novel preconditioner is that it can be easily parallelized, whereas incomplete Cholesky is inherently a serial algorithm.

Simulation Parameter	Value
Grid Resolution	50 x 50
Spring Constant	1000
Inter-particle distance	0.005
Damping Factor	2
Time Step	0.01
Error Threshold	10^{-15}
Mass of particle	1

Table 1: *Simulation Parameters.*

#Nodes	Cholesky		Poisson		SSOR		Jacobi		None	
	#iter	Time(s)	#iter	Time(s)	#iter	Time(s)	#iter	Time(s)	#iter	Time(s)
2500	10	0.079	11	0.032	11	0.035	18	0.062	13	0.028
3600	10	0.099	11	0.046	11	0.046	18	0.080	13	0.049
4900	10	0.160	11	0.054	11	0.064	18	0.090	13	0.058
6400	10	0.133	11	0.062	11	0.078	18	0.101	13	0.064
8100	9	0.150	11	0.076	11	0.096	18	0.121	13	0.083
10000	9	0.151	11	0.084	11	0.110	18	0.138	13	0.091
19600	9	0.241	11	0.157	11	0.217	18	0.241	13	0.182
30625	9	0.482	12	0.326	11	0.359	18	0.388	13	0.337
40000	9	1.649	12	0.396	11	0.490	18	0.567	14	0.542

Table 2: *Performance results (number of iterations and simulation time in seconds) for all preconditioning schemes with increase in number of nodes.*

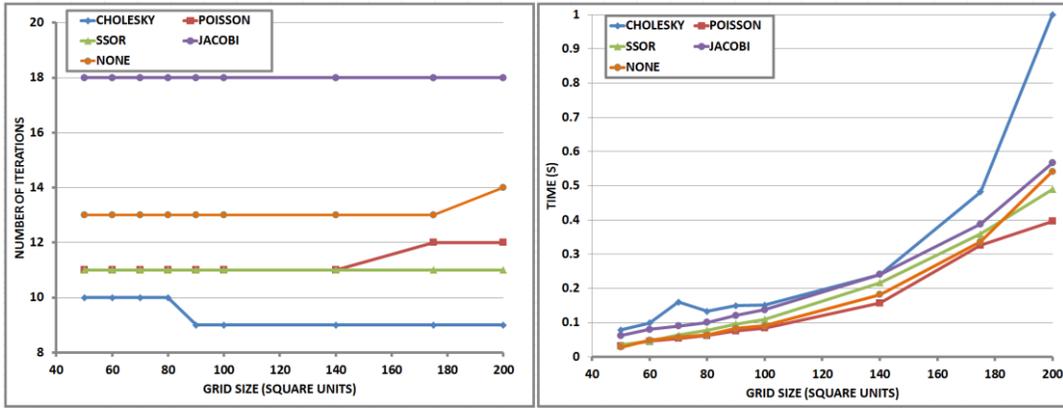


Figure 1: Performance results (number of iterations and simulation time in seconds) for all preconditioning schemes with increase in number of nodes.

Parallelization Results

To evaluate parallelization options for the incomplete Poisson preconditioner, we implemented a parallel version using Pthreads. The columns of matrix A are equally distributed among the available threads, which means that each thread operates on $\text{No.Columns} / \text{No.threads}$ columns of A. The input of the parallelization function is A and the output $P = H^*H^T$, where $H = I - L*\text{diag}\{A\}^{-1}$. We tested for different numbers of threads and the times we report include construction of the threads as well as thread synchronization. Our tests were performed on an 8-core Apple* Mac* Pro running Mac* OS X* 10.6 with 12GBs of RAM (Figure 2(a)) and a 32-core code name Emerald Ridge server with Intel® Xeon® processors X7560 and 32GB of RAM running OpenSUSE* Linux* 11.3 (Figure 2(b)). Both systems have Intel® Hyper-Threading Technology. To compute execution time, we used system-specific high-resolution timers: mach absolute time() on Mac OS X* and clock_gettime() on Linux*. We further refer the reader to [1] for GPU parallel implementations of the incomplete Poisson preconditioner. Figure 2 shows that PIPP scales very well with the number of available cores. For both systems the performance of the PIPP increases significantly and reaches saturation after the number of threads equals the number of available hyper-cores. In the case of the 8-core system the performance of a single thread is about 2 seconds, while the performance of 16 threads is about 0.5 seconds. Similarly, in the case of the 32-core system a single thread takes more than a second, while 64 threads run at about 0.06 seconds.

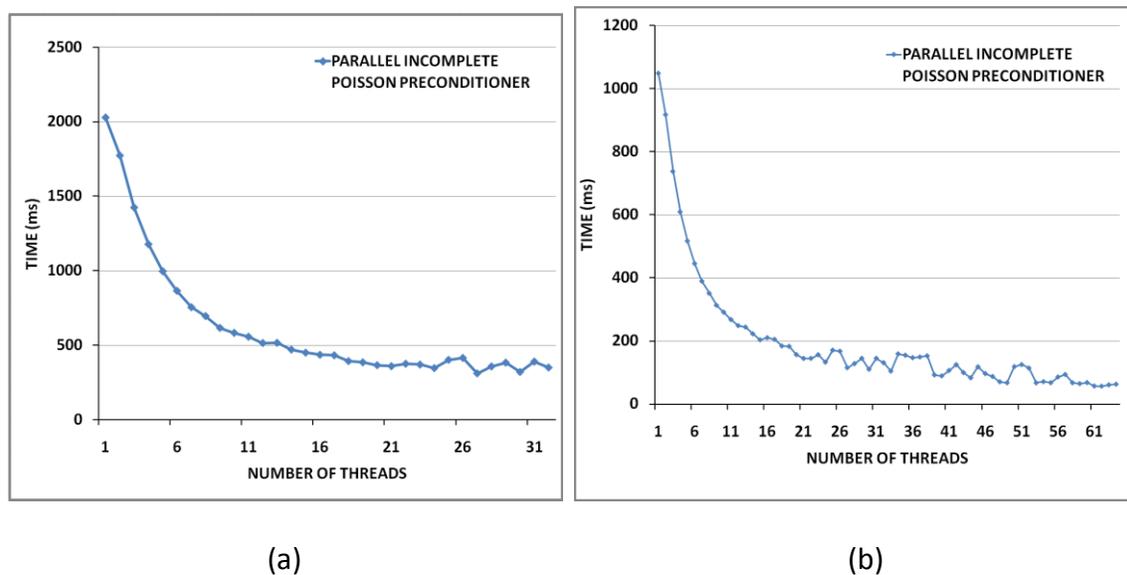


Figure 2: Parallelization results on multi-core hardware. (a) 8-core machine. (b) 32-core machine.

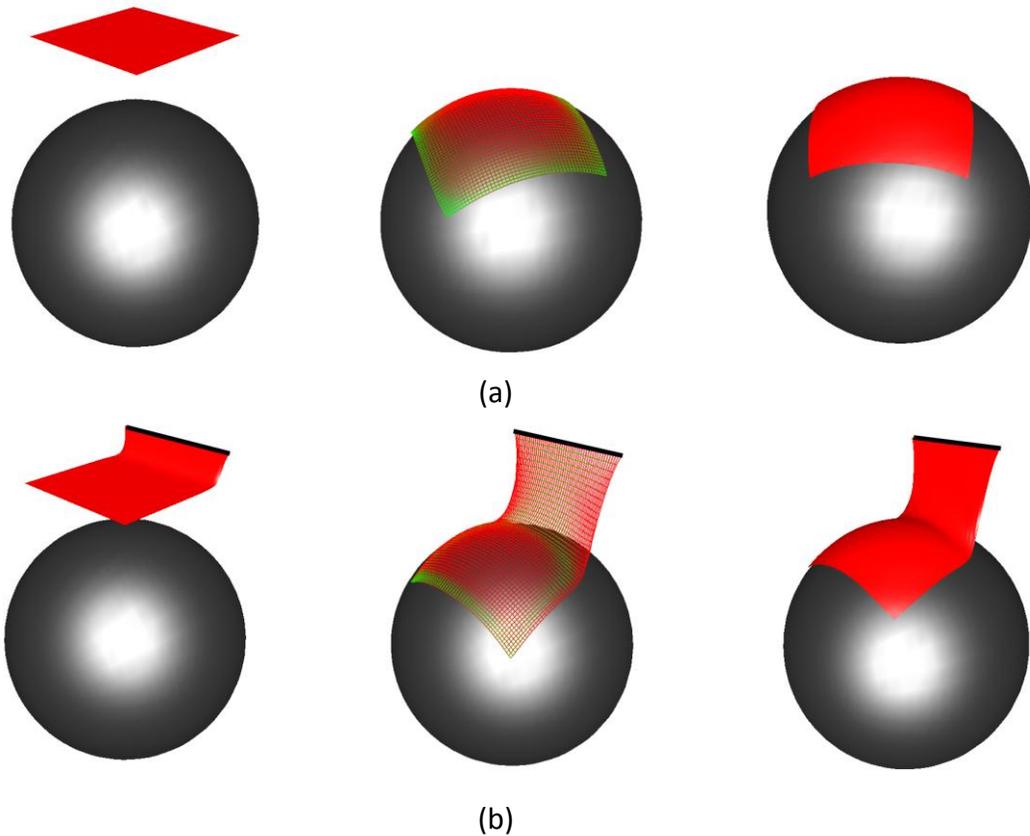
Conclusion and Future Work

We have presented a recently proposed preconditioner, the incomplete Poisson preconditioner, for conjugate gradient solvers. We have analyzed the fitness of the proposed preconditioning scheme on several benchmarks and compared its

performance to commonly used methods. We have showed that PIPP performs as well as the current state-of-the-art preconditioners, while being much more amenable to standard thread-level parallelism. Our experiments on two multi-core systems show that PIPP scales very well with the number of available processing cores.

Acknowledgements

The work in this paper was partially supported by Intel through a Visual Computing grant and the donation of the 32-core code name Emerald Ridge system with Intel® Xeon® processors X7560. In particular we would like to thank Randi Rost and Scott Buck from Intel for their support. We would like to thank Rhythm & Hues Studios and in particular Peter Huang and Tae-Yong Kim for their support through grants and software donations. We would also like to extend our gratitude to Thanasis Voggiannou for providing an open source cloth simulation engine which was used in part to generate the results for this paper.



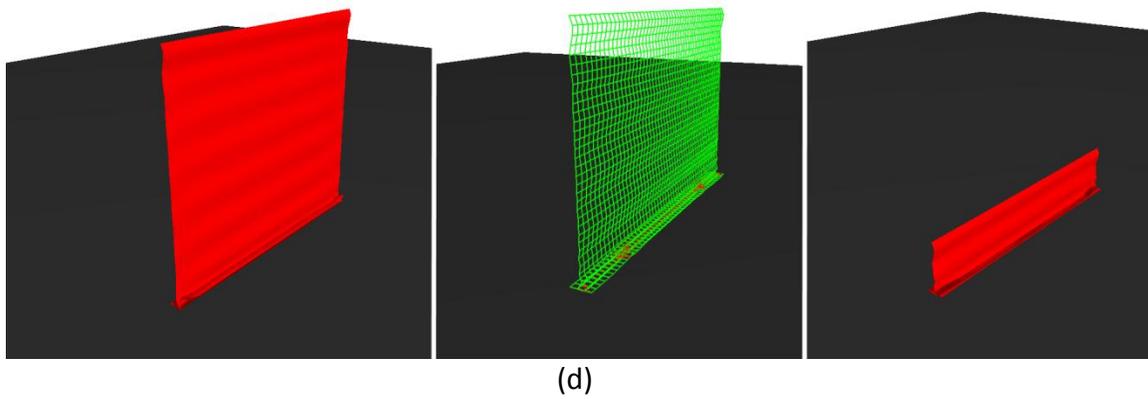
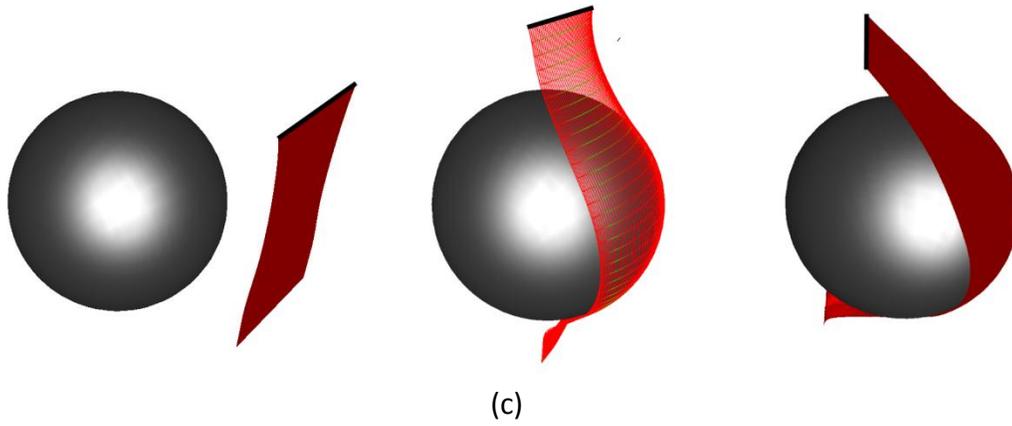
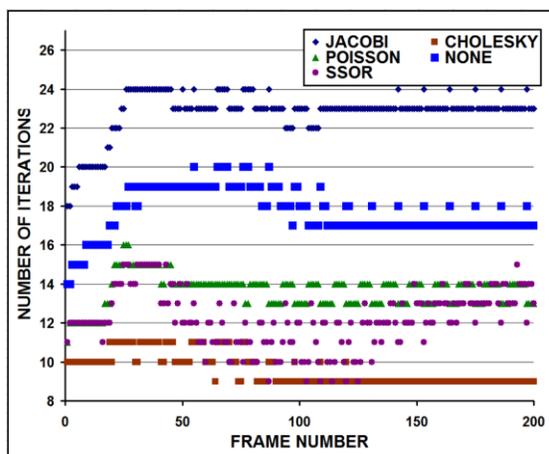
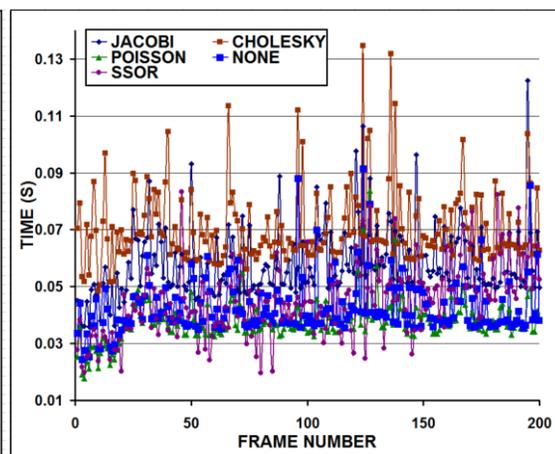


Figure 3: Benchmark Scenes. (a) Cloth falling on a sphere. (b) Cloth hanging as a curtain colliding with sphere. (c) Cloth patch colliding with moving spherical object. (d) Tangling cloth. In the middle frames the red color indicates edges under stress.



(a)



(b)

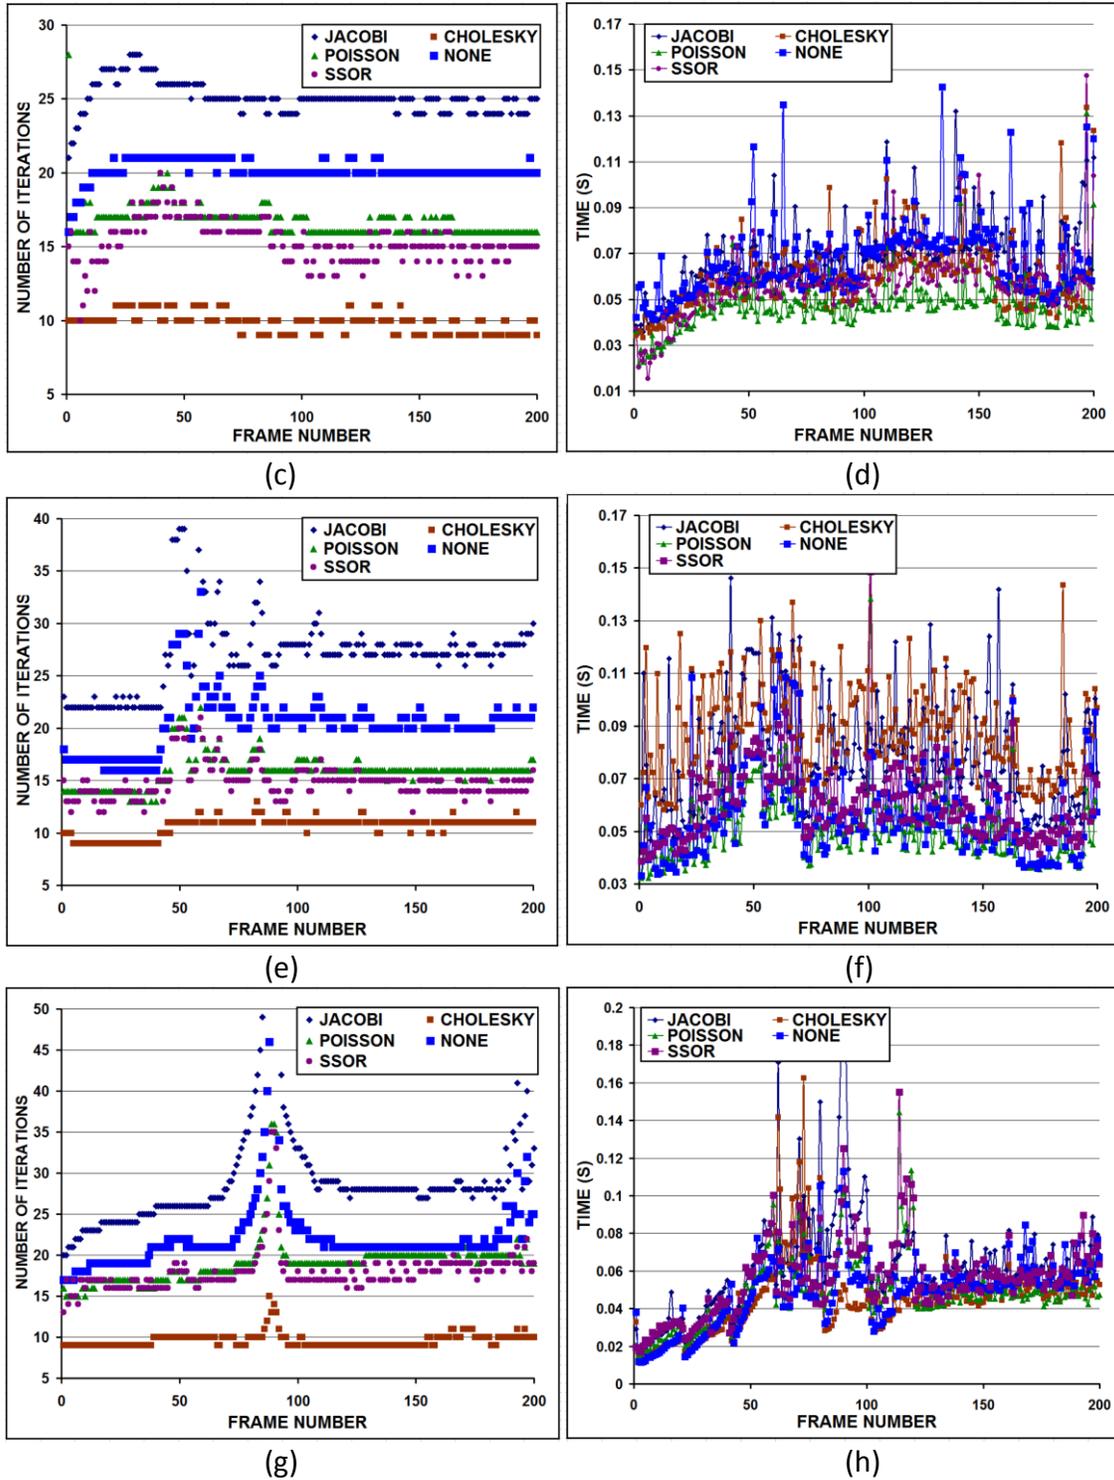


Figure 4: Performance results of different preconditioning schemes on all benchmarks. (a),(b): Number of iterations and simulation time for Free fall benchmark. (c),(d): Curtain benchmark. (e),(f) Tangling benchmark. (g),(h) Moving collider benchmark.

AUTHORS:



Costas Sideris is a PhD student in the UCLA Computer Science Department, advised by Professor Petros Faloutsos. Costas received a Bachelor's degree in Computer & Electrical Engineering from the National Technical University of Athens, Greece in 2010. Costas is a member of the Modeling, Animation, and Graphics Lab at UCLA. His research focuses on physical simulations and animation, as well as vision subjects including 3D reconstruction and perceptual interfaces. For more information, please visit: www.cs.ucla.edu/~costas



Mubbasir Kapadia is the assistant director and a postdoctoral researcher at the SIG Center for Computer Graphics at the University of Pennsylvania. His research interests include applying automated planning and machine learning to simulate autonomous virtual humans. Kapadia has a PhD in computer science from the University of California, Los Angeles. For more information, please visit:

<http://www.seas.upenn.edu/~mubbasir/>

REFERENCES:

- [1]. Ament, M., Knittel, G., Weiskopf, D., Strasser, W.: A parallel preconditioned conjugate gradient solver for the Poisson problem on a multi-GPU platform. In: Proceedings of the 2010 18th Euromicro Conference on Parallel, Distributed and Network-based Processing. pp. 583–592. PDP '10, IEEE Computer Society (2010)
- [2]. Ascher, U., Boxerman, E.: On the modified conjugate gradient method in cloth simulation. *The Visual Computer* 19, 526–531 (2003)
- [3]. Baraff, D., Witkin, A.: Large steps in cloth simulation. In: Proceedings of ACM SIGGRAPH. pp. 43–54 (1998)
- [4]. Boxerman, E.: Speeding up cloth simulation. Ph.D. thesis, The University of British Columbia, BC, Canada (2003)
- [5]. Breen, D.E., House, D.H., Wozny, M.J., Breen, D.E.: Predicting the drape of woven cloth using interacting particles (1994)

- [6]. Bridson, R., Fedkiw, R., Anderson, J.: Robust treatment of collisions, contact and friction for cloth animation. In: ACM SIGGRAPH 2005 Courses. SIGGRAPH '05, ACM, New York, NY, USA (2005)
- [7]. Carignan, M., Yang, Y., Thalmann, N.M., Thalmann, D.: Dressing animated synthetic actors with complex deformable clothes. In: Computer Graphics (Proc. SIGGRAPH. pp. 99–104 (1992)
- [8]. Choi, K., Ko, H.: Research problems in clothing simulation. Computer-Aided Design 37(6), 585–592 (2005)
- [9]. Choi, K.J., Ko, H.S.: Stable but responsive cloth. In: Proceedings of ACM SIGGRAPH. pp. 604–611 (2002)
- [10]. Eberhardt, B., Weber, A., Strasser, W.: A fast, flexible, particle-system model for cloth draping. IEEE Comput. Graph. Appl. 16, 52–59 (September 1996)
- [11]. Fuhrmann, A., Sobottka, G., Grob, C.: Distance fields for rapid collision detection in physically based modeling. In: GRAPHICON (2003)
- [12]. Goldenthal, R., Harmon, D., Fattal, R., Bercovier, M., Grinspun, E.: Efficient Simulation of Inextensible Cloth. SIGGRAPH (ACM Transactions on Graphics) 26(3) (2007)
- [13]. Hauth, M., Etmuss, O., Strasser, W.: Analysis of numerical methods for the simulation of deformable models. The Visual Computer 19, 581–600 (2003)
- [14]. Muller, M.: Hierarchical position based dynamics. In: Proceedings of Virtual Reality Interactions and Physical Simulations (VRIPhys2008). pp. 13–14 (2008)
- [15]. Muller, M., Heidelberger, B., Hennix, M., Ratcliff, J.: Position based dynamics. J.Vis. Comun. Image Represent. 18, 109–118 (April 2007)
- [16]. Nealen, A., Muller, M., Keiser, R., Boxerman, E., Carlson, M.: Physically based deformable models in computer graphics. Computer Graphics Forum 25, 809–836 (2006)
- [17]. Okabe, H., Imaoka, H., Tomiha, T., Niwaya, H.: Three dimensional apparel cad system. In: Proceedings of the 19th annual conference on Computer graphics and interactive techniques. pp. 105–110. SIGGRAPH '92, ACM, New York, NY, USA (1992)
- [18]. Selle, A., Su, J., Irving, G., Fedkiw, R.: Robust high-resolution cloth using parallelism, history-based collisions, and accurate friction. IEEE Transactions on Visualization and Computer Graphics 15, 339–350 (March 2009)

[19]. Shewchuk, J.R.: An introduction to the conjugate gradient method without the agonizing pain. Tech. rep. (1994)

[20]. Terzopoulos, D., Fleischer, K.: Deformable models. *The Visual Computer* 4(6), 306–331 (1988)

[21]. Terzopoulos, D., Fleischer, K.: Modeling inelastic deformation: Viscoelasticity, plasticity, fracture. *Computer Graphics (Proc. SIGGRAPH'88)* 22(4), 269–278 (1988)

[22]. Terzopoulos, D., Platt, J., Barr, A., Fleischer, K.: Elastically deformable models.

ADVISOR: PETROS FALOUTSOS



Petros Faloutsos is an assistant professor at the Department of Computer Science and Engineering at York University. Before joining York, he was a faculty member at the Computer Science Department at the University of California at Los Angeles, where in 2002 he founded the first computer graphics lab at UCLA, called M.A.Gix. He served as the lab's director until 2011. Faloutsos received his PhD degree (2002) and his MSc degree in Computer Science from the University of Toronto, Canada and his BEng degree in Electrical Engineering from the National Technical University of Athens, Greece. Faloutsos' research interests focus on digital media, computer graphics, virtual humans, hardware accelerators for graphics, health informatics and surgical robotics. Professor Faloutsos is also interested in computer networks, and he has co-authored a highly cited paper on the topology of the Internet that received an ACM SIGCOMM Test of Time Award in 2010. Faloutsos is a member of the Editorial Board of the *Journal of The Visual Computer*, and has served as a Program Co-Chair for the ACM SIGGRAPH/Eurographics Symposium on Computer Animation 2005 and for the Motion In Games Conference 2011. He is a member of the ACM and the Technical Chamber of Greece.

FURTHER RESOURCES:

UCLA MAGiX lab: <http://magix.ucla.edu/>