

Theory to Practice: Generalized Minimum-Norm Perspective Shadow Maps for Anti-aliased Shadow Rendering in 3D Computer Games

Fan Zhang[†], Hanqiu Sun[†], Chong Zhao[†], Lifeng Wang[‡]

[†] Department of Computer Science and Engineering,
The Chinese University of Hong Kong, Shatin N.T., Hong Kong, China
{fzhang, hanqiu, czhao}@cse.cuhk.edu.hk

[‡] Autodesk Software (China) Co., Ltd. (Shanghai Office),
399 Pu Dian Road, Shanghai Pudong District, Shanghai 200122, China
lifeng.wang@autodesk.com

Abstract. Shadow mapping has been extensively used for real-time shadow rendering in 3D computer games, though it suffers from the inherent aliasing problems due to its image-based nature. The aliasing errors in shadow mapping consist of perspective aliasing error and projection aliasing error. In this paper, we propose a novel shadow-map reparameterization to reduce perspective aliasing for varying viewer and/or light. This reparameterizing technique keeps the perspective aliasing distribution optimal in possible general cases. Our experiments have shown the enhanced shadow quality using our algorithm in dynamic scenes.

Key words: Image-based rendering, shadow algorithms, computer games

1 Introduction

Shadowing effects are very important for the realism of virtual scenes in 3D computer games, which significantly enhance players' perception of the surrounding environment with useful depth clues. For game developers, the reality that they are facing is that the commonly used shading models do not directly support the global illumination effects like shadows.

Shadow mapping[16] is one of the most popular dynamic shadowing techniques. This image-based approach doesn't require the knowledge of the scene's geometry, and requires only one extra rendering pass per light. Such generality and efficiency make shadow mapping extensively adopted in 3D games. There are two rendering passes involved in shadow mapping. In the first pass, the scene is rendered into the *shadow map* texture from the light's point of view. In the second pass, the scene is rendered from the eye's point of view. For each point being rendered, it is transformed into the light's coordinates again for depth comparison. If the corresponding depth value in shadow map is less than that of the transformed point, this point is shadowed.

When using shadow mapping, we need to handle aliasing problem which mainly resulted from insufficient shadow-map resolution (i.e. *under-sampling*). Numerous approaches have been proposed to address this issue in different ways.

Partitioning techniques[7][15][19][2][10][11] help the anti-aliasing in shadow mapping with multiple shadow maps, the performance cost is however increased due to extra rendering passes. Hybrid techniques[8][3][4][1] combine shadow mapping and *shadow volumes*[6] to achieve a better tradeoff of shadow quality and performance. For a thorough review of shadow rendering algorithms refer to [18].

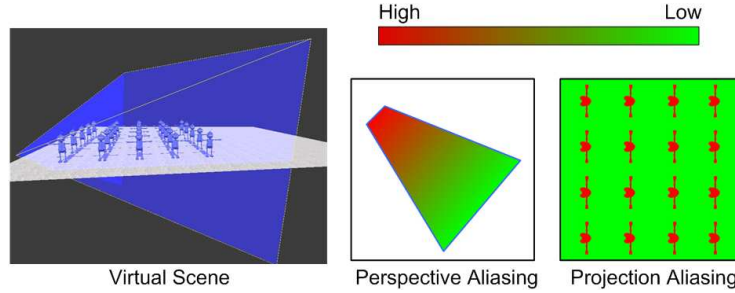


Fig. 1. Perspective aliasing v.s. projection aliasing. The blue frame in the virtual scene represents the view frustum.

The aliasing errors in shadow mapping are usually decomposed into *perspective aliasing errors* and *projection aliasing errors*[14], as visualized in Fig. 1. Projection aliasing depends on local geometry details such that reducing this kind of aliasing requires an expensive scene analysis at each frame. On the other hand, perspective aliasing comes from the foreshortening effect introduced by the perspective projection. Perspective aliasing is the only kind of aliasing that can be alleviated by warping the shadow plane using a perspective transform. This is based on the observation that the distribution of even-spaced texels on the warped shadow map plane better matches the distribution of even-spaced pixels on the screen. Hence, the sampling density at the light in the post-transformed space will better accommodate the requirements for the reconstruction of shadowed images (Fig. 2). This idea was first introduced in *perspective shadow maps* (PSMs)[14]. As pointed out by Wimmer et al[17], the optimal distribution of sampling points on the shadow map plane is a logarithmic function with respect to the depth value. However, this logarithmic reparameterization is not directly supported on current hardware because all input data are interpolated hyperbolically in the stage of rasterization¹. For insightful analysis of perspective aliasing refer to [17][11].

This paper mainly addresses *perspective reparameterization* (PR) techniques that reparameterize the shadow map using a perspective transform as shown in Fig. 2. Perspective reparameterizations generate shadow maps in the post-perspective space of the warping frustum, to cancel out the foreshortening effect introduced by the view frustum. Since the perspective warping transformation enlarges the objects close to the viewer and shrinks the distant objects, perspec-

¹ Strictly speaking, this reparameterization can be discretely approximated using multiple shadow maps[12]. However, the performance drop caused by extra rendering passes is not suited for complex scenes.




Fig. 2. A schematic illustration of perspective reparameterizations, where V and P stand for the eye’s view frustum and perspective warping frustum respectively. Both V and the light are transformed into P ’s post-perspective space, in which the shadow map is generated.

tive aliasing is significantly reduced. Currently existing PRs[14][9][17][13][5][20] exploit two major factors: the near-plane distance of the warping frustum and the warping direction. The essential difference among perspective reparameterizations is the selection of the first factor, because the near plane distance of the warping frustum determines how strong the warping effect is. On the other hand, the selection of the second factor strongly influences the implementation complexity. With an inappropriate warping direction, the frequently changed types of lights after applying the perspective transform may cause mapping singularities. A smart choice of this factor is proposed in *light space perspective shadow maps* (LiSPSMs)[17], which use a warping direction in parallel with the shadow plane to avoid producing mapping singularities. It greatly simplifies the implementation and analysis of perspective reparameterizations.

In this paper, we propose the *generalized minimum-norm perspective reparameterization* (GMNPR) based on the generalized representation of aliasing errors proposed by Zhang et al.[20]. Our algorithm can be regarded as a generalization of LiSPSMs in non-ideal cases where the angle between the light and view directions are not orthogonal. A few advantages offered by our algorithm include:

- In comparison with most of prior reparameterizations in which the optimal aliasing distribution is achieved only in the ideal case, GMNPR keeps the aliasing distribution optimal in possible general cases. Such direction-invariant feature (partially) preserves the optimized shadow quality for dynamic scenes.
- GMNPR inherits the flexibility of LiSPSMs. The warping directions in both two methods are the same. We thus don’t need special treatments such as

the "inverse perspective matrix" [9] for the mapping singularity problem in the post-perspective space.

The remainder of the paper is organized as follows: Section 2 briefly reviews the theory of perspective aliasing in shadow mapping. Section 3 derives our GMNPR in detail. Section 4 presents the experimental results. Finally, the conclusion and further work go to Section 5.

2 Generalized Representation of Aliasing Errors

We first briefly review the aliasing problem in shadow mapping for the overhead light [14][17][11]. In Fig. 3, the light beams through a texel with the size ds (in the z direction) fall on a surface with the length dz in world space. Here, the local parameterization of the shadow map $s(z)$ assumes that the shadow map is accurately focused on the view frustum and no resolution is wasted on invisible parts of the scene. The size of view beams dp on the normalized screen projected from the surface is $dy/z \tan \phi$, where 2ϕ is the field-of-view of the view frustum. α and β denote the angles between the surface normal and vector to the screen and shadow map plane respectively.




Fig. 3. Shadow map aliasing errors.

By restricting the small surface on the z direction, we have $dy \approx dz \cos \alpha / \cos \beta$. The aliasing error \mathbf{E} for the small surface is then defined as

$$\mathbf{E} = \frac{dp}{ds} = \frac{1}{\tan \phi} \frac{dz}{z ds} \times \frac{\cos \alpha}{\cos \beta} \triangleq \mathbf{F} \times \frac{\cos \alpha}{\cos \beta} \quad (1)$$

The above shadow map aliasing consists of *perspective aliasing* dz/zds and *projection aliasing* $\cos \alpha / \cos \beta$. Note that ϕ is a constant once given the view matrix. As we mentioned in the introduction, PRs focus on the reduction of perspective aliasing, i.e. reparameterizing \mathbf{F} .

3 The Algorithm

Because the perspective warping effect decreases as the light direction goes to parallel with the view direction, λ should be *direction-adaptive* to satisfy the application-specific constraint as θ changes. With the representation expressed by Eq. (2), users are capable of designing the direction-adaptive $\lambda(\theta)$ to produce the expected warping strength in the general case. Although λ_{LiSPSM} also takes into account θ , as we explain later, it actually equalizes the aliasing distribution only in the ideal case. Based on the observations above, we proposed the *generalized minimum-norm perspective reparameterization* (GMNPR)² to satisfy the following two requirements,

- keep the optimal aliasing distribution in possible general cases.
- converge to standard shadow maps as the light direction goes to parallel with the view direction.

There’s nearly no difference for implementing and applying GMNPR in comparison with LiSPSMs, except for the λ selection in the warping frustum, to control the strength of the warping effect. In the following, we thus focus on deriving appropriate λ values to produce the robust aliasing distribution for dynamic scenes.

In mathematical sense, we can define variant norms. Like LiSPSMs, GMNPR is designed to minimize the \mathbf{L}^∞ -norm of errors. Once given λ and θ , Eq. (2) shows that $\mathbf{F}(z)$ has a single positive local extremum (a minimum) at the location $z = n(1 - \frac{\tan\phi}{\tan\theta})(f - n + (f + n)\frac{\tan\phi}{\tan\theta})^{-1}$. Therefore, the maximum $\mathbf{L}^\infty(\mathbf{F})$ is achieved at the boundaries of the depth range. Let $\mathbf{L}_{\max}^\infty(\mathbf{F})$ be the maximum $\mathbf{L}^\infty(\mathbf{F})$ within $[n, f]$. Minimizing $\mathbf{L}_{\max}^\infty(\mathbf{F})$ requires that the errors at near and far planes are equal, or

$$\mathbf{F}(\lambda, \theta, n) = \mathbf{F}(\lambda, \theta, f). \quad (3)$$

The above maxima can be directly calculated from Eq. (2) as follows,

- **Case 1:** if $\phi \leq \theta \leq \pi/2$,

$$\begin{aligned} \mathbf{F}(\lambda, \theta, n) &= \frac{(f - n + (f + n)\frac{\tan\phi}{\tan\theta})}{n \tan\phi} \frac{(\lambda + \mu)^2}{\lambda(\lambda + 1)} \\ \mathbf{F}(\lambda, \theta, f) &= \frac{(f - n + (f + n)\frac{\tan\phi}{\tan\theta})}{f \tan\phi} \frac{(\lambda + \nu)^2}{\lambda(\lambda + 1)} \\ \mu &= \frac{n}{(f - n)\frac{\tan\theta}{\tan\phi} + f + n} \quad \nu = 1 - \frac{f}{(f - n)\frac{\tan\theta}{\tan\phi} + f + n} \end{aligned}$$

- **Case 2:** if $0 \leq \theta \leq \phi$,

$$\begin{aligned} \mathbf{F}(\lambda, \theta, n) &= \frac{2f}{n \tan\theta} \frac{(\lambda + \mu)^2}{\lambda(\lambda + 1)} \quad \mathbf{F}(\lambda, \theta, f) = \frac{2}{\tan\theta} \frac{(\lambda + \nu)^2}{\lambda(\lambda + 1)} \\ \mu &= \frac{1}{2} - \frac{f - n}{2f} \frac{\tan\theta}{\tan\phi} \quad \nu = \frac{1}{2} \end{aligned}$$

² To keep consistent with the naming convention in [20], we use the term "perspective reparameterization" instead of "perspective shadow maps" in this paper. Furthermore, the two terms are usually exchangeable in this research line.

Leading the above results to Eq. (3) gives the optimal λ selection,

$$\lambda = \frac{\nu - \mu}{\sqrt{f/n} - 1} - \mu = \begin{cases} \lambda_1(\theta) = \frac{(n + \sqrt{fn}) - n \frac{\tan \phi}{\tan \theta}}{f - n + (f+n) \frac{\tan \phi}{\tan \theta}} & \theta \in [\phi, \frac{\pi}{2}] \\ \lambda_2(\theta) = \frac{f + \sqrt{fn}}{2f} \frac{\tan \theta}{\tan \phi} - \frac{1}{2} & \theta \in [0, \phi] \end{cases} \quad (4)$$




Fig. 5. Adjustment of the λ_{GMNPR} selection in practice.

Now let's consider the critical angle $\theta = \phi$. Fig. 5 plots λ against θ . We can see that λ decreases slowly as θ goes from $\pi/2$ to ϕ , thus the perspective warping effect becomes stronger. At the singularity $\theta = \phi$, λ takes a nose-dive and rapidly plunges towards 0, which reaches well for $\lambda > 0$. But $\lambda = 0$ is obviously disastrous, as it practically concentrates the whole shadow map into one texel. Due to the fact that the perspective projection is a hyperbolic mapping, the depth region $[\lambda \leq z \leq 2\lambda]$ occupies about half of the depth range $[0 \leq s \leq \sim 0.5]$ in the post-perspective space. In practice, it's never a good idea to let λ be too small in any perspective reparameterization, as this introduces strong warping such that the shadow-map texels are usually biased to few pixels in the near plane. Hence, the rapid decrease to zero for λ makes Eq. (4) hard to be used in practice when $\theta < \phi$. To avoid this problem, we have to replace λ_2 by a positive function λ_2^* which should satisfy two criteria,

- smooth transition of shadow qualities at $\theta = \phi$.

Note the singularity at $\theta = \phi$ may result in the discontinuous transition of shadow qualities in practice. To avoid this problem, the continuous λ transition at $\theta = \phi$ should be satisfied.

$$\lambda_2^*(\phi) = \lambda_1(\phi) \quad (5)$$

- convergence to SSM as θ goes to 0.

When $\theta < \phi$, from Fig. 4, the warping frustum bounds two lines on the far plane of the view frustum. The available depth range distorted by the warping frustum is significantly narrowed. In such cases, perspective reparameterizations should weaken the warping strength in order to avoid the unbalanced distribution of shadow details on the screen. In particular, when $\theta = 0$, all perspective reparameterizations converge to standard shadow maps.

$$\lim_{\theta \rightarrow 0} \lambda_2^*(\theta) = \infty \quad (6)$$

Many candidates can fulfill Eqs. (5) and (6). In our experiments, λ_2^* is chosen to be

$$\lambda_2^*(\theta) = \lambda_1(\phi) / \tan\left(\frac{\theta}{4\phi}\pi\right) = \frac{\sqrt{fn}}{2f} \frac{1}{\tan\left(\frac{\theta}{4\phi}\pi\right)}. \quad (7)$$

Substituting Eq. (7) into Eq. (4) gives our final λ_{GMNPR} selection as follows,

$$\lambda_{\text{GMNPR}} = \begin{cases} \lambda_{\text{GMNPR}}^{(1)}(\theta) = \lambda_1(\theta) = \frac{n + \sqrt{fn} - n \frac{\tan\phi}{\tan\theta}}{f - n + (f+n) \frac{\tan\phi}{\tan\theta}} & \theta \in [\phi, \frac{\pi}{2}] \\ \lambda_{\text{GMNPR}}^{(2)}(\theta) = \lambda_2^*(\theta) = \frac{\sqrt{fn}}{2f} \frac{1}{\tan\left(\frac{\theta}{4\phi}\pi\right)} & \theta \in [0, \phi] \end{cases} \quad (8)$$

Eq. (8) gives us a direction-adaptive λ selection for *minimum-norm* aliasing distribution, which essentially extends LiSPSMs to general cases. Note that all λ values discussed in this paper are for the warping frustum P with a normalized depth range, namely $P_f - P_n = 1$, where P_f and P_n represent P 's near-plane and far-plane values respectively. In LiSPSMs, the near-plane distance of the warping frustum P_n in general cases is approximated as

$$P_n = \frac{n + \sqrt{fn}}{\sin\theta}$$

By normalizing the depth range of the warping frustum, we have

$$\lambda_{\text{LiSPSM}}(\theta) = \frac{P_n}{P_f - P_n} = \begin{cases} \frac{\frac{n + \sqrt{fn}}{\sin\theta}}{\left(\frac{n \frac{\tan\phi}{\tan\theta} + f - n + f \frac{\tan\phi}{\tan\theta}\right) \sin\theta} & \theta \in [\phi, \frac{\pi}{2}] \\ \frac{\frac{n + \sqrt{fn}}{\sin\theta}}{2f \tan\phi \cos\theta} & \theta \in [0, \phi] \end{cases} \quad (9)$$

Refer to [20] for the derivation of $P_f - P_n$. Eq. (4) tells us that the above approximation is not optimal in non-ideal cases (at least for $\theta > \phi$). The λ selection in LiSPSM makes \mathbf{F} optimal only in the ideal case where $\theta = \pi/2$.

Other possible generalized PRs also can be derived from Eq. (2), such as the *generalized linear perspective reparameterization* (GLPR)[20] which requires $\mathbf{F}(z) \sim z$ for given λ and θ values, where \sim stands for the linear dependency. This reparameterization essentially extends PSMs to non-ideal cases. The main problem of GLPR is that the warping strength might be over-strong in some cases. We compare our GMNPR with GLPR in Section 4.

Analysis of Aliasing Distribution

Fig. 6 plots the aliasing distributions in LiSPSM (top row) and GMNPR (bottom row) for the configuration $n = 1$, $f = 100$ and $\phi = \pi/6$. Due to the critical angle $\theta = \phi$ in Eq. (2), we compare them in the following two aspects.

- non-degenerate cases $\theta > \phi$.

The aliasing distribution in GMNPR, $\mathbf{F}_{\text{GMNPR}}$, is direction-invariant (in terms of minimum-norm) as the θ value changes. The aliasing errors at boundaries of the depth range remain equal for all possible θ values, while $\mathbf{F}_{\text{LiSPSM}}$ fluctuates with the varying of θ . As θ decreases, the aliasing errors

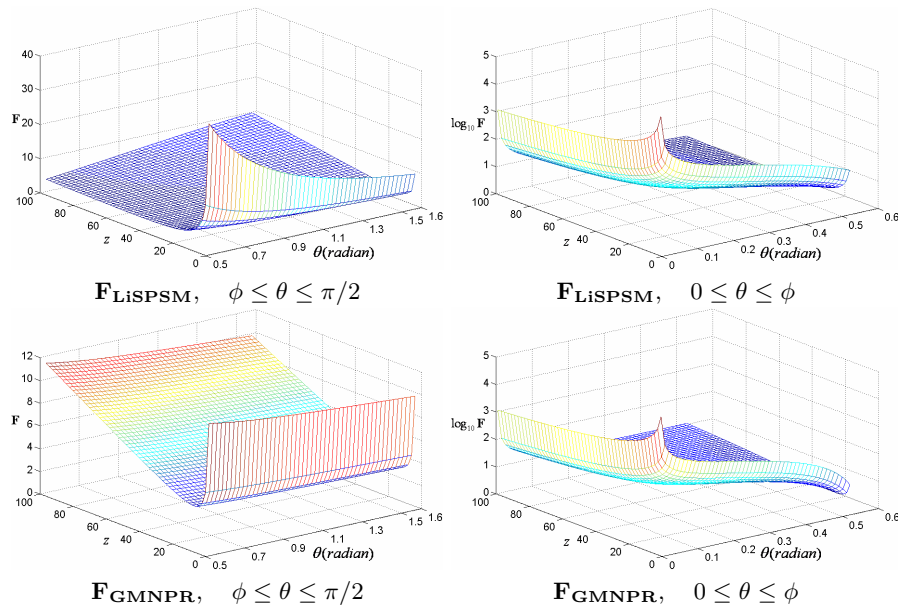


Fig. 6. The aliasing distributions.

at near regions in LiSPSM increase, but the shape of $\mathbf{F}_{\text{GMNPR}}$ keeps unchanged for varying θ values. Such direction-invariant feature is very useful for shadow rendering in dynamic scenes, because the consistent transition of shadow quality should be preserved as possible during the user’s navigation/interaction.

- degenerate cases $\theta \leq \phi$.

The values of both $\mathbf{F}_{\text{LiSPSM}}$ and $\mathbf{F}_{\text{GMNPR}}$ become extremely large for small θ values. For better illustration, a \log_{10} scale is used in Fig. 6 now. As θ goes from ϕ to 0, all PRs can not make \mathbf{F} optimal and converge to SSMs. The effectiveness of the perspective warping dramatically decreases such that no PR has any special advantage over others. This problem is well known as the *dueling frusta* case in which the anti-aliasing (using single shadow map with given resolution) still remains unsolved in computer graphics.

4 Results

We have developed our GMNPR using Microsoft DirectX SDK 9.0. We run the shadow rendering tests using 800*600 image resolution, and the performance is measured by an Intel Pentium 4 2.8GHz CPU with 1G RAM, a NVidia GeForce 6800Ultra GPU with 256M video memory. In addition, we have compared our algorithm with SSM, LiSPSM and GLPR[20]. Two virtual scenes are composed and tested in our experiments, as shown in Fig. 7. In our tests, a dynamic directional lighting source and field of view 60° are set up. $n = 1m$, $f = 800m$, $\phi = \pi/6$ are configured for all of tests. All algorithms achieve real-time performance in our experiments.




Fig. 7. Testing scenes: Robots (11K) and Ceiling (23K).

Robots: to verify the theoretical results shown in Fig. 6, we compare different algorithms in this complex scene with a large number of triangles. Fig. 8 shows two views at $\theta = 55^\circ$ and 35° . As we can see, SSM produces the worst shadow quality for near objects. Since GMNPR preserves the optimal aliasing distribution in most general cases, the shadow quality in GMNPR is better than that in LiSPSM. The associated shadow maps are listed as well. As we can see, the shadow maps in GMNPR preserves that near regions consume the majority of shadow map resolution as the viewer or light moves.

Ceiling: we use this scene to compare GMNPR with GLPR. Fig. 9 shows two views and corresponding light’s views, where the blue frame represents the view frustum. In View-1, both GLPR and GMNPR produce satisfactory shadow quality at near regions. From the associated light’s views, we can see the warping strength in GLPR is stronger than that in GMNPR. However, the over-strong warping strength in GLPR sometimes may cause the extremely unbalanced distribution of shadow details. When the viewer moves closer to the object, λ_{GLPR} becomes too small such that the warped object only occupies a very small portion of shadow map in practice. The shadow quality thus sharply decreases as shown in View-2. This problem is hard to be solved even using some optimization techniques like ”unit cube clipping” [9].

5 Conclusion

In this paper, we have proposed the *generalized minimum-norm perspective reparameterization* (GMNPR) to enhance the shadow rendering in 3D computer games. Our algorithm preserves the optimal aliasing distribution in possible general cases. Such direction-invariant feature is very important for the consistent transition of shadow quality during user’s interaction with dynamic scenes. Our experimental results have shown the improved shadow quality using our algorithm.

As the computational model used in this paper (Fig. 4) assumes that the warping direction is parallel to the shadow plane and lights are directional, we will further our work to derive other direction-adaptive reparameterizations using arbitrary warping direction, and extend our analysis to other types of lights. It would be interesting to take into account the errors in the x direction [11] for better understanding of perspective reparameterizations.

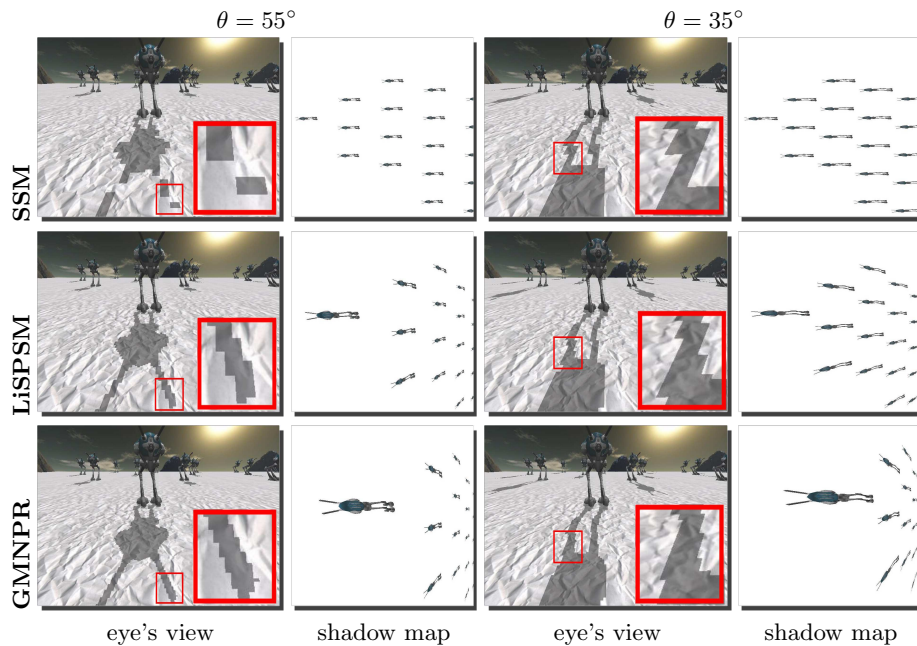


Fig. 8. Comparison of the shadow quality and associated shadow maps. The red rectangles zoom in the local shadow details.




Fig. 9. Comparison of the shadow quality and corresponding light's views. Note that the extremely stretched shadow details in GLPR due to the over-strong warping effect. The optimization technique "unit cube clipping" [9] is used for both algorithms.

Acknowledgements

This work was supported by RGC research grant (NO. 415806). The Hugo robot model used in Fig. 1 was created by Laurence Boissieux at INRIA. The free 3D models used in our experiments are downloaded from www.planetquake.com and www.3dcafe.com.

References

1. T. Aila and T. Akenine-Möller. A hierarchical shadow volume algorithm. In *Proc. of Graphics Hardware 2004*, pages 15–23. Eurographics Association, 2004.
2. J. Arvo. Tiled shadow maps. In *Proc. of Computer Graphics International 2004*, pages 240–247. IEEE Computer Society, 2004.
3. J. Arvo and T. Aila. Optimized shadow mapping using the stencil buffer. *Journal of Graphics Tools*, 8(3):23–32, 2004.
4. E. Chan and F. Durand. An efficient hybrid shadow rendering algorithm. In *Proc. of the Eurographics Symposium on Rendering 2004*, pages 185–195. Eurographics Association, 2004.
5. H. Y. Chong and S. J. Gortler. A lixel for every pixel. In *Proc. of the Eurographics Symposium on Rendering 2004*, pages 167–172. Eurographics Association, 2004.
6. Franklin C. Crow. Shadow algorithms for computer graphics. In *Proc. of SIGGRAPH '77*, pages 242–248. ACM, 1977.
7. R. Fernando, S. Fernandez, K. Bala, and D. P. Greenberg. Adaptive shadow maps. In *Proc. of SIGGRAPH '01*, pages 387–390. ACM, 2001.
8. N. K. Govindaraju, B. Lloyd, S.-E. Yoon, A. Sud, and D. Manocha. Interactive shadow generation in complex environments. *ACM Trans. Graph.*, 22(3):501–510, 2003.
9. S. Kozlov. Perspective shadow maps: care and feeding. *GPU Gems*. Addison–Wesley, 2004.
10. A. E. Lefohn, S. Sengupta, J. Kniss, R. Strzodka, and J. D. Owens. Glift: Generic, efficient, random-access gpu data structures. *ACM Trans. Graph.*, 25(1):60–99, 2006.
11. B. Lloyd, D. Tuft, S. Yoon and D. Manocha. Warping and partitioning for low error shadow maps. In *Proc. of the Eurographics Symposium on Rendering 2006*, pages 215–226. Eurographics Association, 2006.
12. B. Lloyd, N.K. Govindaraju, D. Tuft, S. Molnar and D. Manocha. Practical logarithmic shadow maps. *SIGGRAPH '06 Sketches*, pages 103. ACM, 2006.
13. T. Martin and T.-S. Tan. Anti-aliasing and continuity with trapezoidal shadow maps. In *Proc. of the Eurographics Symposium on Rendering 2004*, pages 153–160. Eurographics Association, 2004.
14. M. Stamminger and G. Drettakis. Perspective shadow maps. In *Proc. of SIGGRAPH '02*, pages 557–562. ACM, 2002.
15. K. Tadamura, X. Qin, G. Jiao, and E. Nakamae. Rendering optimal solar shadows with plural sunlight depth buffers. *The Visual Computer*, 17(2):76–90, 2001.
16. L. Williams. Casting curved shadows on curved surfaces. In *Proc. of SIGGRAPH '78*, pages 270–274. ACM, 1978.
17. M. Wimmer, D. Scherzer, and W. Purgathofer. Light space perspective shadow maps. In *Proc. of the Eurographics Symposium on Rendering 2004*, pages 557–562. Eurographics Association, 2004.
18. A. Woo, P. Poulin, and A. Fournier. A survey of shadow algorithms. *IEEE Comput. Graph. Appl.*, 10(6):13–32, 1990.
19. F. Zhang, H. Sun, L. Xu, and L. K. Lun. Parallel-split shadow maps for large-scale virtual environments. In *Proc. of the 2006 ACM international conference on Virtual reality continuum and its applications*, pages 311–318. ACM, 2006.
20. F. Zhang, L. Xu, C. Tao and H. Sun. Generalized linear perspective shadow map reparameterization. In *Proc. of the 2006 ACM international conference on Virtual reality continuum and its applications*, pages 339–342. ACM, 2006.