

# Details on the Demo for “Improved Moment Shadow Maps for Translucent Occluders Soft Shadows and Single Scattering”

Christoph Peters

## Contents

<b>1</b>	<b>Controls</b>	<b>1</b>
<b>2</b>	<b>Settings</b>	<b>2</b>
2.1	HUD and scene . . . . .	2
2.2	Surface shadows . . . . .	3
2.3	Single scattering . . . . .	4
<b>3</b>	<b>Recommended Experiments</b>	<b>5</b>
<b>4</b>	<b>Notes</b>	<b>6</b>
<b>5</b>	<b>Credits</b>	<b>7</b>
	<b>References</b>	<b>7</b>

## 1 Controls

Before starting the demo please edit *ShadowSettings.cfg* to reflect your screen resolution. You also need to install the redistributable for Visual Studio 2015 (bundled as *VC\_redist.x64.exe*). In the demo use:

- *Arrow keys* to move horizontally (right shift moves faster, right control slower),
- *Page up/down* to move vertically,
- *Mouse* to look,
- *Left mouse button* to orient the directional light parallel to the viewing direction,
- Hold a *letter key* and press a *digit key* to change a setting (as shown in the text-based HUD),

- Hold a *letter key* and *left shift* and press a *digit key* to change a setting to an option with index  $> 10$ ,
- Press *F3* to save the current camera and light, press *F4* to restore the saved state.
- Press *F8* to reload shaders related to shadows (files *Shadow.fx*, *Participating-Media.fx* and *Lighting.fx* in the directory *Shaders*). This is useful if you want to experiment with the shaders and iterate rapidly.

## 2 Settings

The demo relies on a text-based head-up display (HUD) to show available and active options. In each line the name of a setting is followed by the key used to edit the setting. Hold this letter key and press a digit key shown next to an option to activate this option. The active option is marked with , others with . Note that changed settings often necessitate compilation of shaders and therefore the demo may freeze for several seconds. Settings that are irrelevant for the currently active technique are hidden.

**Submenu (M+digit):** To keep the HUD clear, settings for different aspects of the demo are part of different submenus. This setting allows you to switch between these submenus or to display all settings at once. Settings which are not displayed can still be changed.

### 2.1 HUD and scene

**HUD (H+digit):** Allows you to specify the amount of information shown in the HUD.

**Environment (E+digit):** The demo ships with a variety of different environments providing different challenges for the presented techniques. For more details about the used assets please refer to Section 5.

**Animation mode (I+digit):** Defines which parts of the environment should be animated. You can animate the flying dragon (seaport environment only), let the sunlight follow a path much like that of the sun, both or nothing.

**Light size (L+digit):** For soft shadows it is assumed that the directional light spans a particular angle within two planes defined by the view space of the shadow map. This option defines the size of this angle. Larger values produce softer shadows.

**Shader output (U+digit):** By default the demo shows a shaded environment. Alternatively, you can visualize the shadow intensity or disable shading. The latter option is useful if you want to look at single scattering results without getting distracted by surface shading.

## 2.2 Surface shadows

**Hard shadow preset (S+digit):** Presets for different techniques implementing filtered hard shadows. When changed, presets override several options at once to demonstrate a particular technique with the recommended parameters. You can still change the individual settings, which makes the preset jump back to “Custom”.

**Soft shadow preset (F+digit):** Presets for soft shadowing techniques. All of them approximate the result of percentage closer soft shadows in some way.

**Shadow technique (T+digit):** Available shadow mapping techniques for hard and soft shadows.

**Translucent occluders (O+digit):** If this option is enabled, surfaces are rendered to the shadow map with alpha blending. This provides an approximation to shadows for translucent occluders. Otherwise the shadow map is created from a depth buffer. Note that this option is only available for some kinds of filterable shadow maps.

**Shadow map resolution (R+digit):** The resolution of the used shadow maps in texels. This option applies to the shadow maps used for surface shadows and single scattering.

**Shadow map quantization (Q+digit):** The number of bits per channel per texel to be used for shadow maps. For some soft shadow mapping techniques both floating point and integer formats are available. Otherwise the adequate format is chosen automatically.

**Shadow map antialiasing (A+digit):** The quality of shadows can be improved by rendering the shadow map with multisample antialiasing. The option defines the number of samples. This setting is only available for filterable shadow maps.

**Kernel size (K+digit):** For filtered hard shadows this setting defines the radius (and standard deviation) of the Gaussian filter applied to the shadows. For soft shadows it defines the size of the search region for the blocker search and the largest possible filter size. Either way greater values lead to softer shadows but for filtered hard shadows the hard-coded upper limit is  $15 \cdot 15$ .

**Soft shadow interpolation (J+digit):** For some soft shadow techniques bilinear interpolation imposes a significant performance hit. For comparison this setting allows you to switch between nearest-neighbor interpolation and bilinear interpolation.

**Depth bias (Z+digit):** The demo uses a primitive depth bias to avoid surface acne. For surface shadows it simply subtracts the constant defined by this setting from the depth of each fragment before entering shadow computations. For soft shadows the depth bias is proportional to the filter size and the the

setting defines the maximal value. Depth values are in the range  $[-1, 1]$ . Large values introduce detached shadows known as Peter-Panning.

**Moment bias (V+digit):** For four moment shadow mapping a bias is needed to compensate for quantization artifacts in the moment shadow map. The strength is defined by this setting. For 16-bit quantization the recommended value is  $6 \cdot 10^{-5}$ , for 32-bit quantization it is  $3 \cdot 10^{-7}$ .

**Exponential variance factors (X+digit):** Exponential variance shadow maps warp depth exponentially before computing two moments. This setting defines the two factors that are applied to the depth in  $[-1, 1]$  before taking the exponential.

**Overdarkening (D+digit):** A simple way to reduce light leaking is to remap the domain of shadow intensity values. Surfaces receiving the amount of light indicated by this setting will still be treated as fully shadowed. For brighter surfaces the shadow intensity is remapped linearly.

**Thread group size (G+digit):** Soft shadow techniques use compute shaders generating prefix sums of rows and columns. The number of threads per group used for dispatch of these shaders is defined by this setting. The results are independent of the choice, so the option providing the shortest frame time should be chosen.

## 2.3 Single scattering

**Single scattering preset (P+digit):** Besides surface shadows, the demo provides various techniques to integrate volumetric shadows, i.e. to render single scattering. Presets for the available techniques are provided by this setting.

**Scattering technique (C+digit):** Allows you to select particular scattering techniques directly rather than using the presets.

**Scattering quantization (B+digit):** Like *shadow map quantization* but for the resampled shadow maps which are generated specifically for single scattering.

**Scattering moment bias (N+digit):** Like *Moment bias* but only for single scattering techniques based on moment shadow maps. For moment shadow mapping with six moments stored in 64 bits the recommended value is  $4 \cdot 10^{-3}$ .

**Scattering overestimation weight (W+digit):** Single scattering based on moment shadow maps can guarantee that single scattering is always underestimated (too dark) or always overestimated (too bright). Alternatively, you can interpolate linearly between these two extremes using the intermediate options. The best option is usually adaptive overestimation where the interpolation factor varies dependent on the angle between light and view ray.

Some combinations of options are not reasonable or not implemented. In these cases the selected option is marked with  $\#$  and the truly used option is marked

as active. Other combinations can be tested but will not give good results (e.g. moment shadow mapping with an inadequate bias). Therefore, it is recommended to stick to the presets.

### 3 Recommended Experiments

The demo allows you to reproduce all results presented in the paper and more. Here are some hints on what you may do to see these results.

**Soft shadows:** Hold M+2 to go to the submenu for surface shadows. Single scattering should be disabled (P+2) because many scattering techniques require specific kinds of shadow maps. You can try moment soft shadow mapping with interpolation (F+6) or without (F+5, faster but blocky). Recommended environments are the Seaport (E+1), Columns (E+4), Sintel (E+5) and Quadbot (E+6). Try enabling animations (I+4) to get a feeling for aliasing artifacts. For percentage-closer soft shadows (F+3) they will be a lot stronger because multisample antialiasing cannot be used for the shadow map.

**Single scattering:** Hold M+3 to go to the submenu for single scattering. Presets for the novel methods can be enabled by holding P and pressing 4 (lowest quality) to 8 (highest quality). Activating animations (I+4) helps to highlight aliasing artifacts. Recommended environments are the Seaport (E+1), Foliage (E+3) and Escher (E+7).

**Translucent occluders:** There is only one environment with properly sorted, transparent geometry, namely Smoke and Pipes. Note that the particles for the smoke are not automatically oriented towards the camera. Hold E+2 to go to this environment and hold M+2 to go to the submenu for surface shadows. Then enable a shadow mapping technique based on moment shadow maps, e.g. Hamburger 4 moment shadow mapping with 16-bit quantization (S+9). Hold O+2 to enable support for translucent occluders. You should now see partial shadow being cast by the smoke. Hold the left mouse button to change the light direction or enable animations (I+4) to see these shadows interact with shadows of opaque occluders in different configurations. You can also try variance shadow mapping (S+4) or exponential variance shadow mapping (S+7).

**Filtered hard shadows:** Go through the various presets for filtered hard shadows, especially variance shadow mapping (S+4), exponential shadow mapping (S+6), exponential variance shadow mapping (S+7) and moment shadow mapping with 64 bits per texel (S+9) or 128 bits per texel (S+0). A good scene for systematic analysis of the artifacts of these techniques is Thin Wall (E+0). Enable animations (I+2) to see different configurations.

## 4 Notes

**System requirements:** Some features in the demo use compute shaders and cannot function properly if your graphics device does not support them. Direct3D feature level 11.0 is required. If your hardware does not provide this, you can still run the demo but a message informs you about the issue and you cannot use some features. In some cases updating graphics device drivers may help. The demo is provided as Windows executable and should work on Windows 7 SP 1 and above. Other platforms such as OS X or Linux are not supported.

**Exponential variance shadow mapping:** Among other techniques the demo features an implementation of exponential variance shadow maps. This implementation is more or less equivalent to the implementation in Matt Pettineo's shadow sample.

**Variance soft shadow mapping:** The implementation of variance soft shadow mapping is *not* a full implementation of the work by Yang et al. [Yang et al., 2010]. It uses neither kernel subdivision nor hierarchical shadow maps. Adding these features would improve the quality but increase the frame time (in most cases). Our implementation may be viewed as implementation of summed-area variance shadow maps [Lauritzen, 2007] with the addition of the heuristic blocker search proposed by Yang et al. [Yang et al., 2010].

**Reuse of shadow maps:** No matter which options you activate, the demo will always render the scene to a shadow map at most once per frame. In particular, all single scattering techniques use shadow maps derived from the shadow map created for surface shadows. This means that the selection of shadow mapping techniques is heavily limited when single scattering is enabled. For most techniques a common shadow map is needed and percentage-closer filtering is used. For filtered and prefiltered single scattering with moment shadow maps a four moment shadow map is required.

**Shadow map sampling:** The shadow map is simply fitted to a bounding box around the important parts of the environment. Of course the arising parametrization is subpar and techniques such as trapezoidal shadow maps [Martin and Tan, 2004], parallel-split shadow maps [Zhang et al., 2006] or sample distribution shadow maps [Lauritzen et al., 2011] would provide an advantage. These are not implemented for several reasons:

- The presented work deals with problems that are mostly orthogonal to the chosen parametrization of the shadow map,
- Having a parametrization that is independent of the camera helps to analyze artifacts up close,
- Moment shadow maps can cope with undersampling well.

**Finding shader code:** The demo itself is closed source but its shaders are not. All shaders used by the application are copied to the folder *CreatedShaders*.

They use includes from *ShaderIncludes* and are automatically generated using functions defined in *Shaders*. The files *Shaders/Shadow.fx* and *Shaders/ParticipatingMedia.fx* are of particular interest. The other files often contain unused rudiments. Some functions that compute input quantities on the CPU have been reimplemented in HLSL for reference. The shader code includes doxygen comments and the full documentation is included alongside the demo.

**Graphics debugging:** You may use graphics debugging with Visual Studio or RenderDoc to analyze how the techniques in the demo work. If you do, you should set `enable_shader_debugging = 1` in *ShadowSettings.cfg*.

**Bugs:** The number of available options and settings leads to an immense number of different configurations. Not all of them are supposed to give good results but the demo should not crash. It is recommended to stick to the presets. If you encounter something that should be working but does not, bug reports can be sent to peters (at) cs (dot) uni (minus) bonn (dot) de. Please include *LogFile.txt* in all bug reports.

## 5 Credits

The demo includes models by third parties published under copyleft licenses. Most of them were obtained from [blendswap.com](http://blendswap.com) and we would like to express our gratitude to the operators of this platform and its generous users who provided the models. Here are credits per scene:

**Seaport:** The scene itself has been created by Enrico Steffen. The animated dragon is courtesy of Zoltan Miklosi.

**Foliage:** The whole scene has been created by Eugene Kiver.

**Sintel:** Sintel is the main character in the open movie Sintel. The model is a slightly modified version of the Sintel lite model. It is licensed under Creative Commons Attribution 3.0 Unported.  
© copyright Blender Foundation | [www.sintel.org](http://www.sintel.org)

**Quadbot:** The quadbot comes from the open movie Tears of Steel and was obtained from [blendswap.com](http://blendswap.com). It is licensed under Creative Commons Attribution 3.0 Unported.  
© copyright Blender Foundation | [www.tearsofsteel.org](http://www.tearsofsteel.org)

**Escher:** This scene has been created by Cedrick Münstermann and is inspired by the painting “Convex and Concave” by M. C. Escher.

**Smoke and pipes, columns, planes, thin wall:** Modeled by Christoph Peters in Blender.

## References

- [Lauritzen, 2007] Lauritzen, A. (2007). *GPU Gems 3*, chapter Summed-Area Variance Shadow Maps, pages 157–182. Addison-Wesley.
- [Lauritzen et al., 2011] Lauritzen, A., Salvi, M., and Lefohn, A. (2011). Sample distribution shadow maps. In *Proceedings of the 2011 ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games*, pages 97–102. ACM.
- [Martin and Tan, 2004] Martin, T. and Tan, T.-S. (2004). Anti-aliasing and continuity with trapezoidal shadow maps. In *EGSR04: 15th Eurographics Symposium on Rendering*, pages 153–160. Eurographics Association.
- [Yang et al., 2010] Yang, B., Dong, Z., Feng, J., Seidel, H.-P., and Kautz, J. (2010). Variance soft shadow mapping. In *Computer Graphics Forum*, volume 29, pages 2127–2134.
- [Zhang et al., 2006] Zhang, F., Sun, H., Xu, L., and Lun, L. K. (2006). Parallel-split shadow maps for large-scale virtual environments. In *Proceedings of the 2006 ACM international conference on Virtual reality continuum and its applications*, pages 311–318. ACM.