Plausible Visualization of the Dynamic Digital Factory with Massive Amounts of Lights

Franz Peschel
Daimler AG
franz.peschel@daimler.com

Fabian Scheer
Daimler AG
fabian.scheer@daimler.com

ABSTRACT

In the last years an enormous progress has been made in improving the visual quality of virtual worlds through approximations of former computationally expensive global illumination (GI) effects. Nevertheless, the handling of dynamic or even deformable content in massive data sets illuminated and shadowed by massive light sources still represents some kind of killer application. Besides the gaming industry, these demands are especially made in digital factory planning applications, where moreover preprocessing times and additional manual preparations of the virtual scenarios have to be avoided to keep a short time to production and accordingly to market. Therefore, we present an efficient rendering pipeline and concept for the creation of a plausible illumination for the digital factory, able to handle dynamic content in massive data sets at a large extent, that are plausibly illuminated and shadowed with a lot of light sources and encoded with high dynamic range information at interactive frame rates in a high resolution. The most important aspects for a visually plausible result are analysed on basis of real images of a factory and at last evaluated with the achieved results. Finally, we give a detailed overview and analysis of the performance of the incorporated techniques on modern graphics hardware to identify the main bottlenecks and key points for future research and conclude with an extensive reflection of the benefits of a plausible illumination for digital factory planning applications.

Keywords: Global illumination, massive data, massive light sources, plausible visualization, virtual reality, digital factory

1 INTRODUCTION

Nowadays, the creation of digital mock-ups (DMUs) in early stages of the product or production life cycle is an essential planning instrument to increase the quality and reduce the costs of products or production facilities. Possible bottlenecks or sources of errors can be discussed and analyzed by experts on basis of a holistic visualization. Nevertheless, an entire factory DMU is very complex (e.g. > 25 mio. polygons), consisting of a lot of data contributed by several planning departments or different suppliers. For this reason special rendering methods, like the visibility guided rendering (VGR) described in [12], have to be applied to handle these massive data sets at interactive or realtime frame rates. In this context the visualization plays an important role. In the past, real time or interactive GI effects for such massive amounts of data were hard to achieve, so expensive cluster solutions were applied to obtain a realistic impression. Furthermore, simple flat or phong shading was the only way to achieve the interactivity with the virtual world in real time frame rates. As a consequence stereoscopic techniques in combination with powerwall systems had to be incorporated, to support the perception of depth and the correct assessment of the digital content, like object correlations. This benefit predominantly arised through the binocular disparity depth cue [3], but also neglected an improved effect for the users perception or feeling of being immersed when supporting a better representation for the perception of pictorial depth cues [3, 17], which are evoked by a realistic visualization. Due to the evolved capabilities of modern graphics hardware former computationally expensive GI effects can be approximated at interactive or real time rates. A high visual quality of virtual worlds can be achieved on a single desktop PC, paving the way for desktop virtual reality(VR) and enabling the possibility to reduce the costs considering the provision of expensive hardware.

Against this background we extend former work[19, 23] and present an efficient rendering pipeline and concept for a plausible visualization of the digital factory (DIFA). Dynamic and deformable content of massive data sets illuminated and shadowed by massive amounts of light sources can be handled. Various GI effects are achieved by approximation, but produce a convincing result. The most important and necessary factors of a plausible visualization of the DIFA are analyzed on basis of footages of a real factory and compared with the achieved quality of our results. At the end a detailed overview of the performance of the several techniques to achieve the identified important factors is given, serving as a basis for the identification of the main bottlenecks and key points of future research in this field of application. Our main contributions are the analysis of the key points for an approximated realistic visualization of the digital factory, the specified efficient rendering pipeline for the DIFA and the detailed performance evaluation of the pipeline on modern graphics hardware.

2 RELATED WORK

A lot of progress has been made in the last decades to improve the realism of virtual worlds by global illumination(GI) algorithms at interactive or real time frame rates. [6, 1] present interactive ray tracing to account for massive data sets in virtual reality(VR) applications. [29] and [18] adopt the capabilities of modern graphics hardware and push ray tracing to the next level by considering more complex light effects[29]. Due to the logarithmic complexity, ray tracing is best suited to handle large data models at performant frame rates. Nevertheless, the limits are revealed when dynamic massive data sets illuminated by massive amounts of light sources and the expensive computation of secondary effects, like bounces of indirect lighting, are involved.

More practical approaches for a dynamic and interactive plausible illumination of complex and massive scenes with a lot of light sources can mainly be found by having a look at the gaming industry([25, 5, 16, 15, 26]), especially when desktop VR with standard hardware is desired to avoid expensive hardware solutions. In this context big advancements were made through image space approximations and deferred shading techniques in the last years. [25] uses deferred shading, introduced by [2], to consider huge amounts of lights for shading and relies on visible fragments only. Thus, the computational effort mainly depends on the amount of light sources and the image resolution, while simultaneously being independent of the polygon count and allowing the handling of dynamic or deformable content in real time.

[4] further improves the performance of this method by accumulating the illumination per pixel in a separate light accumulation buffer(LAB) without the consideration of the material color, which is multiplied at the end, thus reducing computational efforts when a lot of light sources are considered. However, occlusion information is missing and the creation of shadow maps for many light sources would still represent the bottleneck. In addition to the consideration of direct light effects([27, 7]), many advancements have been made for approximating GI effects at interactive or real time rates, like screen space ambient occlusion (SSAO)[16], which has become a standard in the gaming industry. In [23] vector based SSAO closer to the original definition of ambient occlusion is presented, especially focusing on the application in digital factory planning scenarios. In [22] SSAO is extended to account for directional occlusion and indirect lighting effects by screen space directional occlusion(SSDO), but is limited to nearly arranged surfaces. In the following sections the indirect lighting effects of SSDO are also mentioned as screen space global illumination(SSGI). [21] present an interesting approach considering the shadow computation of indirect illumination effects approximated by the usage of many secondary light sources. Unfortunately, their approach relies on a point based scene representation, introducing a further preprocessing step we want to avoid to keep minimal production times. In summary, screen space approaches work entirely in screen space without the need for preprocessing steps while simultaneously being independent of the scene complexity.

Recently, [11] introduced cascaded light propagation volumes (CLPVs). In contrast to SSAO or SSGI, CLPVs can compute indirect lighting effects for large areas and are not limited to reveal its effect in cavities or creases only. Three different sized cascaded grids are attached to the camera and indirect illumination is propagated through the grids by using spherical harmonic functions, yielding a performant GI computation. However, small areas in the scene can be missed due to the grid resolution, but fortunately reconsidered with SSGI.

3 SOLVING THE RENDERING EQUATION

According to [9] the rendering equation can expressed as a *Neumann series* using the integral operator notation in the following form[8]:

$$L = E + KL, \tag{1}$$

where K represents the integration over the cosinus modulated BRDF. E denotes the emitted radiance L_e and L the reflected radiance L_r. Equation (1) can be discretized to a simple matrix equation, where E and L are vectors and K is the light transport matrix, characterizing the reflectivity of light in a scene, yielding:

$$L = E + KE + K^2E + K^3E + ... = \sum_{m=0}^{\infty} K^m E \tag{2}$$

E denotes the emitted radiance of a light source, KE the direct illumination of surfaces, K^2E the first bounce indirect illumination, K^3E the second bounce and so on, finally converging against the exact solution of the rendering equation. On basis of this representation the following methods for the computation of single stages of our rendering pipeline and the analysis of important effects in images of a real factory can be broad into a better coherence. Thereby, the operator notation underlines the separability of GI into single computation stages. To achieve an interactive performance of our pipeline we restrict ourself to consider equation (2) just to K^2E, which in most cases is sufficient for a plausible illumination simulation, since the perceptually noticeable differences between the summands strongly decrease from the fourth summand.

4 ANALYZING THE REAL FACTORY

To evaluate the plausibility of an illumination simulation, comparisons to the reality in the factory or at least to images of the real factory have to be done. Furthermore, essential illumination effects for a plausible visualization can be identified by analyzing these images. Several images were used for the analysis under consideration of the summands of the discretized rendering equation (2) in section 3. The four most representative images are presented in fig.2. In addition, various effects influencing the perception of the human visual system are also addressed in the analysis.

4.1 Emission - E

Referring to equation(2), E represents the emitting radiance of light sources. This effect can be noticed in the views of fig.2 at windows with incoming sunlight or at the fluorescent tubes mounted near the ceilings. Since the images were taken with a low dynamic range camera light emitting surfaces appear as bright white areas, thus no further illumination effects can be recognized in this areas. Factories are mainly illuminated by fluorescent tubes and sunlight that shines through windows and in conclusion have to be simulated in a realistic way, since the human perception strongly reacts to these high frequent parts of the illumination.

4.2 Direct Illumination - KE

In absence of strong sunlight on a sunny day and deep inside the buildings without any windows in range, factories are homogeneously illuminated to achieve ergonomically illuminated workplaces. As a consequence, a lack of hard shadows can be noticed, except for areas under machines and obstacles that are mounted some inches above the ground. The only source for hard shadows represents the sun light, when falling through the windows in a strong intensity. In conclusion, shadow mapping techniques should be applied for the sun light, whereas a diffuse ambient occlusion is sufficient concerning the weak occlusion caused by the fluorescent tubes. Since the tubes are encapsulated by a socket their emitting hemisphere is directed towards the ground.

To achieve a plausible illumination the most remarkable materials also have to be considered to create a visually convincing result, as can be seen by the several specular or glossy effects on the ground and at the pipes in fig.2. The reflecting ground can especially be noticed in view 1 and view 4 and the glossy pipes in view 2 and view 3. Furthermore, strong diffuse reflections can be recognized in view 2 at the walls or at the machinery.

4.3 Indirect Illumination - K^2E

The homogeneous distribution of the light in the scenes as well as the strongly diffuse reflecting surfaces create a smooth ambient illumination and thus it evokes the effect of ambient occlusion in corners and interspaces. Therefore, the ambient brightness and occlusion has to be simulated to obtain a plausible result. Due to indirect bounces of the light in the factory all images of fig.2 seem to be tinted with a green shimmer, which is mainly caused by the green machinery (see view 2 and view 4). As a consequence, the incorporation of indirect illumination effects is essential to achieve a similar appearance concerning indirect bounces and the ambient brightness.

4.4 Glare

Due to the limited capabilities of the used low dynamic range camera used for the images of the factory, the surroundings of light emitting areas are overexposed and glow effects can be recognized at the outer regions. The same effect can be seen at strong reflections on the ground or at metallic surfaces. Those effects are also perceived by the human visual system, even though in a reduced form since the eye can adapt to such illumination situations to a certain degree.

5 DESIGN OF A MASSIVE LIGHTING PIPELINE

In this section a short overview of the massive lighting pipeline (see fig.1) is given, specifying how the several methods for a plausible illumination computation are incorporated into the rendering module of a VR system. For the implementation of the massive lighting pipeline we focused on the visualization of the interior of the factory. As a consequence, hard shadows inside the buildings that merely occur due to strong sun illumination on a sunny day(see section4) are not considered, but can additionally be incorporated as described in our former work[21]. Due to the homogeneous illumination inside the facility convincing occlusion results can be obtained by the SSAO approach of [25]. The main focus lies on the plausible dynamic illumination by thousands of light sources in massive data sets and the simulation of diffuse, glossy and reflective materials. Indirect illumination effects in the nearfield of cavities or creases are achieved by applying SSGI and for larger areas by the application of CLPVs. An ambient term is added by exploiting the capabilities of CLPVs to directly propagate direct light information, which is physically incorrect, but yields visually convincing results, as can be seen in view 2 in fig.2. As a consequence, massive data sets illuminated with massive lights and containing dynamic or deformable content can be handled at interactive rates and in high resolutions without the need for preprocessing stages.

5.1 GBuffer Layout

Since factory planning scenarios with a large spatial extent are considered, the position and normal informa-

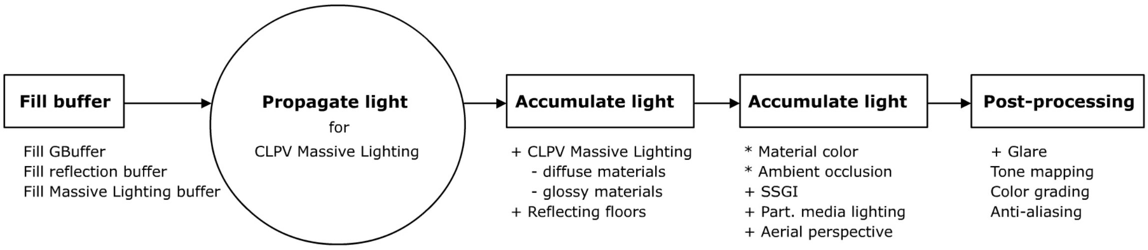

Figure 1: Overview of the massive lighting rendering pipeline

tion has to be stored with 16-bit precision to avoid disturbing z-fighting artifacts as well as gradation artifacts of surface shading. For high screen resolutions this can lead to a large memory requirement. To reduce memory consumption, the world positions are reconstructed from one 16-bit depth value. For the normalized surface normals only the x and y components are needed, since the z component can be reconstructed too. For depth and normal we use the RGB channels of a first render target RT1. Additionally the material color is stored in the RGB channels of a second render target RT2. Further information like specular strength, transparency flags and proxy geometry flags are stored in the alpha channels of RT1 and RT2. In summation two RGBA floating point textures with 16-bit per channel are used for the geometry buffer(GBuffer).

5.2 Massive Lighting

[10] describes the possibility to apply CLPVs for a direct diffuse illumination with massive light sources by directly injecting the direct light information into grids. Therefore, a light source buffer, further denoted as massive lighting map(MLM), is created that encapsulates a maximum of $512x512 = 262144$ usable light sources. The MLM contains position, color, orientation and intensity of light sources. Due to the small computational effort of the injection stage of the light sources into the grids, also higher resolutions of the MLM are possible. Hence, area lights like the fluorescent tubes can be properly approximated by many point lights since the amount of light sources will not become a bottleneck. Dividing the amount of all tubes in the factory by the MLM buffer resolution yields the number of possible light sources to approximate a tube. In fact, such a dense sampling of the tubes is necessary, since flickering artifacts can occur when a light source enters or leaves the spatial grid that is attached to the camera due to camera movements. By using a lot of light sources this case can be avoided due to a smooth fading between the grid cells. In the fill buffer stage of our pipeline the lights of the MLM are injected into the CLVP and propagated through the several grids in the propagation stage. In the light accumulation stage the propagated direct illumination can be accessed at the respective position in the scene to fill the LAB with the direct illumination.

5.3 Ground Reflections

To account for visible important details of reflections on the ground, the scene is rendered a second time into another frame buffer from the perspective of a mirrored camera to obtain a mirrored image. In the light accumulation step, the ground plate is automatically identified by a flag in the alpha channel of RT2 in the GBuffer, which is set using the material of the ground plate. Thus, the covered pixels in image space are determined and used as mask for the rendering of the ground reflections. For the simulation of surface roughness the reflections are finally processed with several filtering techniques, e.g. a light streaks approach described in [13] or a simple 5x5 gaussian blur filter.

5.4 Screen Space Ambient Occlusion and Global Illumination

Ambient occlusion simulates the occlusion in cavities, creases and for concave surfaces and is best suited to provide shadow information for uniformly illuminated diffuse or ambient environments[14], typical for factory interiors. Therefore, [23] is applied yielding sufficient results with merely eight to 16 samples in this diffuse environment. By means of the indirect lighting part of the SSDO technique, indirect bounces of light can be computed for nearly arranged surfaces. Since the method is applied to the LAB that already contains the direct illumination of the CLPVs at this stage, indirect bounces of the illumination can be computed as described in [19]. Due to the use of many light sources we do not apply the directional occlusion part of SSDO.

5.5 Participating Media Lighting

On basis of the CLPVs a spatial representation of the illumination is available, that can be further extended by the incorporation of direct light information from the fluorescent tubes to simulate participating media lighting effects. Light sources of the MLM are injected into the CLPV with a higher weighting factor and propagated. By using view aligned volume slicing with screen aligned quads, the light intensities of the CLPV grid voxels can be computed and added to the LAB, simulating the effect of homogeneous participating media lighting. By choosing an appropriate density factor various effects like fog, dust or aerial perspective can be

Figure 2: Top row: Flat shading. Middle row: Representative views for the analysis of a real factory, ranging from view 1(left) to view 4(right). Bottom row: Results of our plausible illumination

simulated. For scenes with a large extent this can also evoke the impression of an aerial perspective and thus increase the depth perception of the scene.

5.6 Tone Mapping, Glare and Color Grading

To account for a more realistic representation of the scene, the illumination computation is based on high dynamic range(HDR) information. Thereby, the capability of the human visual system to adapt to different lighting conditions is simulated by applying the tone mapping operator of [20], allowing a plausible display of the HDR information on low dynamic range output devices. The human perception of bright areas in the factories (see section 4) is also considered by using the approach of [13] to visualize a glare effect in the area of light emitting or strong reflective surfaces. Therefore, we store bright areas above a certain threshold in an additional image which is downsampled and stored multiple times by the factor four. The respective images are filtered with a 5x5 gaussian mask. Afterwards the images are upsampled again and added to the final image composition. Since physically correct indirect bounces of the illumination are neglected, the color temperature of the factory was simulated by a color grading method as it is described in [24] using a 3D Look-Up Table(3D LUT) which simply represents a mapping of a rgb color cube into another rgb color cube and allows us to do the color correction very fast. The color grading can be changed in run time using manipulation methods of the 3D LUT. It is also possible to store or load a 3D LUT.

5.7 Anti-aliasing

Aliasing effects caused by an undersampling of high frequent areas can lead to flickering artifacts and jaggy edges, especially during camera movements, reducing the feeling of being immersed in the virtual world. Because of the used deferred shading approach no hardware anti-aliasing can be utilized. We distinguish between a performance and quality version of anti-aliasing. For the performance version an edge anti-aliasing(EdgeAA) is implemented. Edges are detected by building the differences of the surface normals and depth values, followed by a 3x3 gaussian blur applied at the edges. The quality version uses full screen anti-aliasing(FSAA). Therefor the image is computed in doubled resolution. For the super sampling part of FSAA a rotated grid sampling pattern is applied, yielding an optimal result for nearly horizontal or vertical edges, which are often seen in factory environments.

6 RESULTS

To achieve an interactive plausible illumination of massive data sets with massive lights for digital factory planning applications the following main criterias are determined:

- Visually plausible quality
- Interactive or real time frame rates
- Support for massive data sets with dynamic content

- Support for massive dynamic light sources
- Avoidance of preprocessing time

In this section the achieved results are evaluated and reflected considering the listed criterias. The test system consisted of an AMD Athlon 3700+ with 2 GB RAM and a NVIDIA GeForce GTX 285 graphics card. For the performance measurements a typical factory scene with 15.3 million polygons and 4056 fluorescent tubes is used and rendered in a resolution of 1280x1024 pixels. Each of the 4056 tubes of the scenario is approximated by 5 light source, yielding a total of 20280 light sources. To handle the massive amounts of data efficiently a visibility guided renderer[12] is used as the rendering core of the VR-system, filling the buffers for our further illumination computations.

6.1 Visual Quality and Plausibility

To evaluate the plausibility of the proposed method, the images of the real factory and the criterias that were derived by the analysis in section 4 are compared with the achieved results. The results, ranging from view 1 to view 4 in the bottom row of fig.2, show that the light emitting fluorescent tubes with their surrounding glow are plausibly visualized[1]. However, the glow in the images seems a bit more smoothed at the outer edges and is sometimes colorful pigmented against dark backgrounds (see view 1 in fig.2), allowing a further improvement of the simulation. Even though, the direct illumination is spatially approximated by the CLPVs with two band spherical harmonics, the results are visually convincing. Nevertheless, artifacts in form of some isolated bright areas (see view 2 in fig.2) can occur due to the spatial approximation of the grids. These are mainly caused by high intensities in grid cells including direct light sources.

Direct light is propagated in a specific direction in form of cosine-lobes. Thereby, some amount is also send in the opposite direction yielding a scattering of the direct light in the scene. Because of the density of fluorescent tubes and multiple reflections of their light in the factory a similar effect of a diffuse distributed ambient environment can be observed in the images. As a consequence the massive lighting also simulates some kind of physically incorrect indirect illumination in the surroundings of light emitting surfaces and decreases with distance. Thus, a kind of global ambient term is obtained that decreases in distance to a light source in contrast a constant ambient factor, e.g. used in the OpenGL lighting model. However, despite the fact that the light propagation in form of cosine-lobes is directed to the ground, the areas above the fluorescent tubes are also illuminated due to the scattering, thus creating an indirect illumination that appears very similar to the images(see view 2 and view 3 in fig.2). Unfortunately, this effect is colored with direct light information and can not consider the color of a true bounce of indirect illumination. The provided occlusion information obtained by SSAO yields visually plausible results in most of the cases. The limits are revealed in view 2 in fig.2. The shadow beneath the machine in front should be stronger and weaker at the pillar. In view 1 in fig.2 the occlusion on the right side also seems to be to strong, but can be improved with a better parameterization. The indirect illumination effect of SSGI can be clearly noticed in view 2 in fig.2 on the left side on the ground near the green machines, further improving the visual quality towards a plausible impression.
Another important aspect concerning the plausibility is the reflection of the scene on the ground, even though it sometimes appears to be polished due to a lack of the simulation of rough surfaces, which can be incorporated by using additional bump and specular maps[2].
The sparse application of participating media effects in the background of the scenes decreases the contrasts in distance and creates an aerial perspective that further conveys the impression of depth. The effect is especially perceivable in view 2 and view 3 in fig.2.
On basis of the tone mapping operator, the visible contrast and brightness ratios are reproduced and appear very similar to the images of the real factory. The color temperature of the factory could be approximated using the color grading method mentioned in 5.6. In combination with light scattering of the CLPVs a plausible pigmentation of indirect reflected light is obtained and essential for the atmosphere of the visualization as a whole. In conclusion, a plausible and convincing overall quality of the illumination for virtual environments is achieved in comparison to the images of the real factory. Nevertheless, small discrepancies can be recognized in detail, but due to the diffuse ambient environment inside the factory these artifacts are not disturbing, since the human perception is less sensitive to errors in low frequency illumination environments.

6.2 Performance

An average performance of 6 frames per second is achieved for the test scene in a resolution of 1280x1024 pixels. Thus, the dynamic digital factory can be interactively explored and manipulated. The computational effort of the several stages of the proposed pipeline for the presented views including a detailed listing of the incorporated methods is depicted in fig. 3.

[1] It should be noted that geometrical discrepancies can occur, due to inconsistencies of the model and the real factory

[2] Applicable in combination with the Virtual Texturing technology of [28]

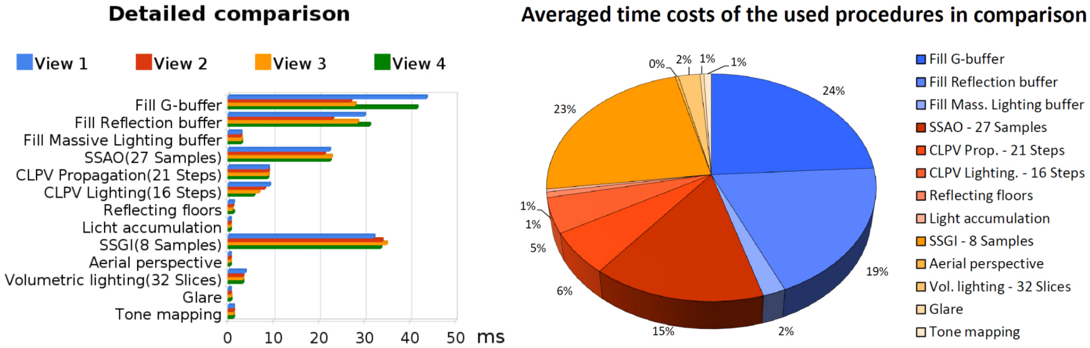

Figure 3: Left: Detailed performance analysis of the several pipeline stages for the considered views. Right: All blue-colored pipeline steps belong to the fill buffer stage, all red-colored to the light accumulation stage and all yellow-colored to the post-processing stage.

Fill Buffer Stage The fill buffer stage with the blue-colored steps in the graph seen on the right side of fig. 3 contains the two time drawing of the scene and therefore is the most expensive stage of the three pipeline stages with an average of 66.06ms. For the filling of the G-Buffer and the Reflection Buffer for the different views strong fluctuations can be noticed in the graph on the left side of fig. 3. These can be explained by the amount of geometry inside the view frustum. In view 2 and view 3 ind fig. 2 a lot of geometry is occluded by the machinery close to the camera, thus the visibility guided renderer can discard a lot of geometry inside the frustum. In average the G-Buffer is filled within 34.97ms, whereas the Reflection Buffer within 28.15ms due to the fact that only a part of the scene covered by the mirror effect has to be rendered. The injection into the CLPVs is independent of the scenes complexity and image resolution and in consequence very performant with merely taking 2.93ms.

Light Accumulation Stage The light accumulation stage with the red-colored steps in the graph is with an average of 40.28ms the second most expensive step, whereas the computation of SSAO with an average of 22.19ms by using 27 samples represents the biggest computational effort. The SSAO performance mainly depends on the image resolution and the amount of used samples. The effort for the light propagation throughout the CLPVs depends on the number of propagation steps and the grid resolution and therefore constantly takes 8.87ms in all views with $21+10+5$ propagations steps and a grid resolution of $48x48x48$. The illumination by the CLPV information depends on the image resolution and the amount of propagation steps in reflection direction through the CLPV to approximate glossy surfaces, which is further denoted as glossy samples. The CLVP performance thus depends on the amount of glossy or reflective areas in the image, which is noticeable on the fluctuations in the timings for the different views, e.g. in view 1 and view 2 more glossy surfaces are visible than in view 3. The average timing is 7.47ms for 16 glossy samples. The last "‘light accumulation"’ step in this stage, which should not be confused with the whole light accumulation stage, is dependent on the image resolution and includes the multiplication of the accumulated light in the LAB with the material color and the ambient occlusion term within an average of 0.6ms.

Post-Processing Stage The post-processing stage with the yellow-colored steps in the graph takes an average of 39.3ms. SSGI is included in this stage, since it is applied after the multiplication with the material color. The average timing for SSGI is 33.55ms by using eight samples and a kernel size of eight. SSGI depends on the image resolution and the amount of used samples. The computational effort for the color grading is with 0.5ms very small, as well as the effort for the glare with 0.6ms and the tone mapping with 1.2ms. All three methods depend on the image resolution. Participating media lighting effects take an average of 3.43ms by using 32 slices, whereas the effort mainly depends on the number of the processed pixels and the number of slices to approximate the volume.

7 BENEFITS

The presented method can handle dynamic objects and dynamic light sources in massive data sets. Almost all incorporated techniques are independent of the scene complexity, except the reflections on the ground plate, allowing the interactive exploration of massive virtual worlds with an enhanced visual quality and pave the way for high quality desktop VR. Massive light sources can be handled and no preprocessings[3] are needed for all incorporated methods, keeping a short time to production and market. Our scenarios need no manual preparation times since massive amounts of data can be handled by a visibility guided renderer and light sources

[3] except the data export from the CAD tool to the VGR data base

are automatically set during the data export of the CAD tools by looking up the light source geometry with an identifier.

8 CONCLUSION

In this paper we presented an efficient massive lighting pipeline for the interactive computation of a plausible illumination for massive data sets of the dynamic digital factory illuminated by massive amounts of light sources. No preprocessing steps are necessary, keeping a short time to production and market. The key points for a plausible illumination computation of a factory are identified by a detailed analysis and exemplary shown on basis of images of a real factory. The results of our proposed method are evaluated against the key points and the real images and can confirm the plausibility of the computed illumination. Finally, a detailed performance evaluation of the whole pipeline is presented, including the separate listing of all incorporated techniques to identify the main bottlenecks and connecting factors for future research.

ACKNOWLEDGEMENTS

This work was partially funded by the AVILUS project of the german 'Bildungsministerium für Bildung und Forschung'(BMBF).

REFERENCES

[1] Stephens A., Boulos S., Bigler J., Wald I., and Parker S.G. An Application of Scalable Massive Model Interaction using Shared Memory Systems. In *Proceedings of the 2006 Eurographics Symposium on Parallel Graphics and Visualization*, 2006.

[2] M. Deering, S. Winner, B. Schediwy, C. Duffy, and Ne Hunt. The triangle processor and normal vector shader: a vlsi system for high performance graphics. In *SIGGRAPH '88: Proceedings of the 15th annual conference on Computer graphics and interactive techniques*, 1988.

[3] David Drascic and Paul Milgram. Perceptual issues in augmented reality. In *SPIE Volume 2653: Stereoscopic Displays and Virtual Reality Systems III*, pages 123–134, 1996.

[4] Wolfgang Engel. Wolfgang engels blog, article: Light prepass, http://diaryofagraphicsprogrammer.blogspot.com/2008/03/light-pre-pass-renderer.html, 2010.

[5] Dominic Filion and Rob McNaughton. Effects & techniques. In *SIGGRAPH '08: ACM SIGGRAPH 2008 classes*, pages 133–164, New York, NY, USA, 2008. ACM.

[6] Wald I., Dietrich A., Benthin C., Efremov A., Dahmen T., Günther J., Havran V., Slusallek P., and Seidel H.-P. A Ray Tracing based Virtual Reality Framework for Industrial Design. In *Proceedings of the 2006 IEEE Symposium on Interactive Ray Tracing*, 2006.

[7] J. Isidoro. Shadow mapping: Gpu-based tips and techniques. *Game Developer Conference (GDC) 2006 presentation*, 2006.

[8] Henrik Wann Jensen. *Realistic image synthesis using photon mapping*. A. K. Peters, Ltd., Natick, MA, USA, 2001.

[9] James T. Kajiya. The rendering equation. *SIGGRAPH Comput. Graph.*, 20(4):143–150, 1986.

[10] Anton Kaplanyan. Light propagation volumes in cryengine 3. In *SIGGRAPH '09: ACM SIGGRAPH 2009 Courses*, New York, NY, USA, 2009. ACM.

[11] Anton Kaplanyan and Carsten Dachsbacher. Cascaded light propagation volumes for real-time indirect illumination. In *I3D '10: Proceedings of the 2010 ACM SIGGRAPH symposium on Interactive 3D Graphics and Games*, 2010.

[12] D. Kasik, B. Brüderlin, M. Heyer, and S. Pfützner. Visibility-guided rendering to accelerate 3d graphics hardware performance. In *SIGGRAPH '07: ACM SIGGRAPH 2007 courses*, 2007.

[13] M. Kawase. Frame buffer postprocessing effects in double-s.t.e.a.l (wreckless). *Game Developer Conference (GDC) 2003 lecture, San Jose, CA.*, 2003.

[14] Hayden Landis. Production-ready global illumination. *In SIGGRAPH 2002 Course Note #16: RenderMan in Production*, pages 87–102, 2002.

[15] M. MITTRING. A bit more deferred-cryengine3. *Triangle Game Conference'09*, 2009.

[16] Martin Mittring. Finding next gen: Cryengine 2. In *SIGGRAPH '07: ACM SIGGRAPH 2007 courses*, pages 97–121, New York, NY, USA, 2007. ACM.

[17] John Murray. Some perspectives on visual depth perception. *SIGGRAPH Comput. Graph.*, 28(2):155–157, 1994.

[18] Steven G. Parker, James Bigler, Andreas Dietrich, Heiko Friedrich, Jared Hoberock, David Luebke, David McAllister, Morgan McGuire, Keith Morley, Austin Robison, and Martin Stich. Optix: A general purpose ray tracing engine. *ACM Transactions on Graphics*, August 2010.

[19] Franz Peschel, Fabian Scheer, and Stefan Müller. Interactive plausible illumination for the digital factory. In *Joint Virtual Reality Conference of EuroVR - EGVE - VEC*, 2010.

[20] E. Reinhard, M. Stark, P. Shirley, and J. Ferwerda. Photographic tone reproduction for digital images. In *SIGGRAPH '02: Proceedings of the 29th annual conference on Computer graphics and interactive techniques*. ACM, 2002.

[21] T. Ritschel, T. Grosch, M. H. Kim, H.-P. Seidel, C. Dachsbacher, and J. Kautz. Imperfect shadow maps for efficient computation of indirect illumination. *ACM Trans. Graph.*, 27(5):1–8, 2008.

[22] T. Ritschel, T. Grosch, and H.-P. Seidel. Approximating dynamic global illumination in image space. In *I3D '09: Proceedings of the 2009 symposium on Interactive 3D graphics and games*, pages 75–82, New York, NY, USA, 2009. ACM.

[23] Fabian Scheer and Michael Keutel. Screen space ambient occlusion for virtual and mixed reality factory planning. *Journal of WSCG*, 18(1-3), 2010.

[24] J. Selan. Using lookup tables to accelerate color transformations. In *GPU Gems 3, Chapter 24*, 2007.

[25] Oles Shishkovtsov. Deferred shading in s.t.a.l.k.e.r. In Matt Pharr and Randima Fernando, editors, *GPU Gems 2: Programming Techniques for High-Performance Graphics and General-Purpose Computation*, pages 143–166. Addison-Wesley Professional, 2005.

[26] M. Valient. Deferred rendering in killzone 2, develop conference 2009, http://www.guerrilla-games.com/publications/dr_kz2_rsx_dev07.pdf, 2010.

[27] Michal Valient. Stable rendering of cascaded shadow map. In Wolfgang Engel, editor, *Shader X6, Advanced Rendering Techniques*. Course Technology, CENGAGE Learning, 2008.

[28] J.M.P. van Waveren. id tech 5 challenges, from texture virtualization to massive parallelization. In *SIGGRAPH '09: ACM SIGGRAPH 2009 Courses*, 2009.

[29] Rui Wang, Rui Wang, Kun Zhou, Minghao Pan, and Hujun Bao. An efficient gpu-based approach for interactive global illumination. *ACM Trans. Graph.*, 28(3):1–8, 2009.