Sustainable VFX
An Educational Challenge

Volker Helzle
volker.helzle@filmakademie.de
Filmakademie Baden-Württemberg, Animationsinstitut
Ludwigsburg, Germany

Alexander Kreische
alexander.kreische@filmakademie.de
Filmakademie Baden-Württemberg, Animationsinstitut
Ludwigsburg, Germany

Simon Spielmann
simon.spielmann@filmakademie.de
Filmakademie Baden-Württemberg, Animationsinstitut
Ludwigsburg, Germany

Jonas Trottnow
jonas.trottnow"@filmakademie.de
Filmakademie Baden-Württemberg, Animationsinstitut
Ludwigsburg, Germany

Figure 1: Comparison of a key visual with traditional (left) and alternative (right) workflow

ABSTRACT

VFX and animated movie production creation processes are exposed to constant improvement, optimization and innovation, and film schools have to reflect these developments in their education. At the current point in time, almost any visual demand can be realized given the necessary human, technical and time resources. Close to real-time rendering technologies can reduce creation times and therefore increase the number of iterations that can be carried out. Both of this can lead to cost and energy savings but comes with having to deal with unconventional creation processes and quality restrictions. Raising awareness about potential CO_2 emission reductions through innovative technology use is an educational imperative in academic education.

This report compares a typical offline production pipeline for VFX creature shots with a near to real-time workflow using a rasterization engine as commonly used in game engines (real-time). The goal here is to start a, from the viewpoint of the authors, necessary "quality in demand" discussion rather than focusing on brand and feature set comparisons. Required optimizations and shortcomings will be discussed as well as the increased iteration opportunities due to reduction in offline render times. Quality aspects of both approaches are compared in a quantitative study evaluating the overall quality and visual difference presented to a diverse audience.

This work has been created in the context of a film school and the development of future skill-sets. Additionally, in the last years, a sustainability discussion for all aspects of film productions has started. We experience a high demand among young filmmakers to understand what it is that they can actively do to lower the carbon-dioxide footprint and produce more sustainably. Therefore, we will also explain how we address these challenges in our curriculum.

CCS CONCEPTS

• **Applied computing** → *Arts and humanities*; • **Social and professional topics** → **Sustainability**; • **Hardware** → **Impact on the environment**.

SA Educator's Forum '24, December 03-06, 2024, Tokyo, Japan

ACM ISBN 979-8-4007-1136-7/24/12
https://doi.org/10.1145/3680533.3697063

KEYWORDS

real-time rendering, environmental impact, movie production, path tracing, rasterization, vfx, asset creation

ACM Reference Format:

Volker Helzle, Alexander Kreische, Simon Spielmann, and Jonas Trottnow. 2024. Sustainable VFX An Educational Challenge. In *SIGGRAPH Asia 2024 Educator's Forum (SA Educator's Forum '24), December 03-06, 2024.* ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3680533.3697063

1 INTRODUCTION

With a raising awareness about film industry's environmental impact in general and animation/VFX specifically, film schools need to put an emphasis on educating their students about viable options for more sustainable production [ALBERT 2022]. The VFX for a student short film, which recently won a prestigious VFX Award, includes a digital creature that was created by final year VFX, Animation and Technical Directing students. The project duration was 18 months, simulating industry workflows with one goal in mind: creating cinema-grade effects with a believable design, natural movement, photo-realistic rendering and integration. The VFX team settled on a conventional workflow utilizing optimized offline rendering to maximise flexibility but also reasonable use of resources. The research and development department of the film school took this unique opportunity to design and develop a secondary pipeline for a near real-time rendering workflow in order to provide students with an exemplary case study. This was done by relying on rasterized rendering and optimized assets, striving for visually similar appearance in order to evaluate an efficient potential future VFX pipeline. The goal of this was to on the one hand showcase the technical capabilities of near real-time rendering while on the other hand raising awareness for the students about potential energy savings and in general environmental impacts of the industry. We posit that it is especially at the levels of film schools and in general in university education that one can find an environment that allows for experimentation in technological approaches, relatively unrestricted in comparison to the strict budgetary considerations of film industry productions. Moreover, it is at this stage in students' careers that they need to be made aware of the industry's environmental issues in order to themselves be the incubators for change once they find their positions in said industry. The use of close to real-time rendering - whether in regular VFX or through in-camera-VFX with e.g. large-scale LED volumes - can be one element of that change - as has been pointed out[Film Paris Region 2022].

2 ESTABLISHED INDUSTRY WORKFLOW

The original creature asset was created for super close-up shots in 4K. Using industry standard software for modelling, rigging, animation, muscle, fat and skin simulations. Geometry caches, grooming, drool and rigid body dynamics were compiled for final lighting, shading and rendering. The original amount of textures (36 4K UDIM body textures, 4 4K face UDIMs, 9 4K eye UDIMs) was reduced before rendering and optimized shaders were developed to save render time. The final render output consisted of diffuse, specular, sub-surface scattering per light and albedo and ambient occlusion information, while the tech passes included depth information, cryptomattes, position, normal and custom mattes (Figure 2). This workflow provided the needed flexibility for detailed adjustments and control over all passes and light sources. Through the added flexibility via Arbitrary Output Variables (AOVs) per light, lookdev processes moved to compositing. Average computing times can be obtained from table 1.

Figure 2: Selection of render passes for traditional workflow

Table 1: Time and Energy comparison for a single frame and entire production

Pass	Rendering Time Path Tracing	Rendering Time Rasterization
Beauty (incl. groom)	88 min.	2.1 min.
Tech & Mask	80 min.	
Drool & Drool Tech	51 min.	
Time per frame	219 min. (3:39 h)	2.1 min.
Total Energy	14 MWh	0.3 MWh

3 NEAR REAL-TIME WORKFLOW

To convert the high quality asset from an offline renderer to a real-time workflow, several optimizations and conversions were required. Texture resolution had to be reduced for the creature to 10% (from 40 4K UDIMs to 1 8K UDIM). The hair was reduced in number (187k to 51k) and detail (3510k to 461k hair points) and exported as animated cache, then converted from curves to a mesh. The drool FX was also exported as geometry cache with reduced polygon count.

The employed rasterizer, just like GPU versions of offline renderers, features real-time feedback when working on lighting and adjusting shader attributes of static geometry, though it lacks indirect lighting. For an animated scene, it only offers direct lighting, screen space and environment reflections. Compared to the offline solution, the biggest time and resources advantage lies in the rapid final output generation.

While basic render passes (also referred to as AOVs) for the real-time engine like Normal and Depth work out of the box, position, SSS, UVs, motion vectors, and others are more challenging or even impossible. Geometry displacement was not possible and as such the silhouette appears more flat and undefined. Physically correct motion blur for the alembic caches was not possible either. This task was shifted to post-production, applying motion blur as frame interpolation.

Final renders were in 4K and required approximately 130 seconds for beauty and tech passes. Also notice that the near real-time feedback was limited to camera, light and shader interactions. Animation updates required about 10 seconds to process. This relatively long delay was mainly caused by the complexity of the geometry.

Graph 1: "How convincing did you find the staging of the creature?"

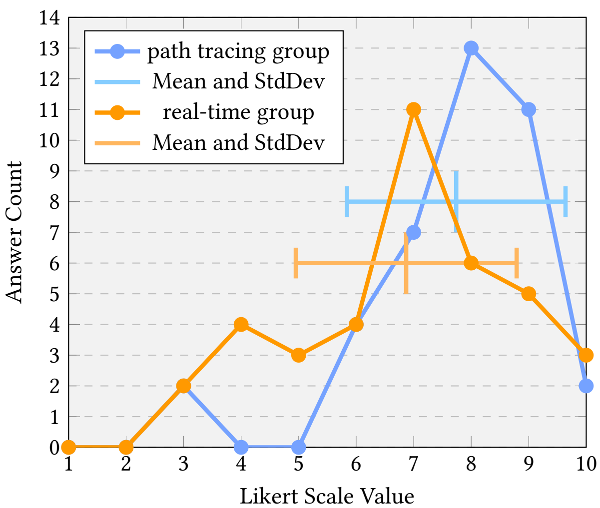

4 COMPARISON

Comparing both versions demands careful consideration. Most importantly, it needs to be emphasized that the near real-time version of the asset is optimized and would also perform better (reduced render time) when used in the offline pipeline. Notice that one goal of this comparison was to find out what an audience response to these versions would be and to identify shortcomings in real-time workflows.

Using a near real-time renderer based on rasterization comes with restrictions and shortcomings. The lack of having similar render passes for a traditional post-production process appears most prominently. Mesh resolution, texture resolution, AOVs and shading need to be limited. Additionally, tasks like grooming, FX and physically correct motion blur remained impossible. Offline rendering setups are still offering additional features including: full AOV output with custom AOVs, ray-traced lighting, fully customizable shading and more - though at the cost of longer render times, limiting fast iterations and near real-time adjustments to the final output.

Similar to professionals, the students were preferring maximum quality imagery and flexibility in post-processing and relighting while successfully targeting optimization. At the same time, relighting in post-production may soon be shifted towards direct adjustments in (real-time) engines, moving work done by a compositor in 2D to the actual 3D scene. Such workflows can lead to enormous time savings as described in the Animation Field Guide[Bousquet 2022].

In future, both worlds (offline and real-time) will most likely evolve towards a parity of feature sets. Advances in graphics hardware and the use of machine learning are enabling more and more offline technologies to become real-time capable[Nicolet et al. 2023]. However, the potential iterations one can perform appear to be most effective with a real-time workflow for now. Hybrid rendering and optimisations towards real-time path tracing are evolving quickly [Ouyang et al. 2021], fulfilling similar time and cost reduction possibilities as rasterized rendering. It should be considered though, that the energy demand of traditional path tracing is significantly higher. A general discussion on quality vs. spent resources (mostly time and energy) seems adequate.

In this comparison scenario additional effort was spent on reducing the complexity of the initial asset for real-time rendering without sacrificing visual fidelity. While this is a big overhead when originally creating the assets for an offline rendering pipeline, this overhead can be reduced drastically when a production targets real-time rendering already at initial asset creation time. On top, assets for real-time use open up opportunities for Cross Media reuse in e.g. games, XR projects, and Virtual Production [Helzle et al. 2022], making the asset more valuable in itself.

Graph 2: "How realistic did you find the creature's surfaces?"

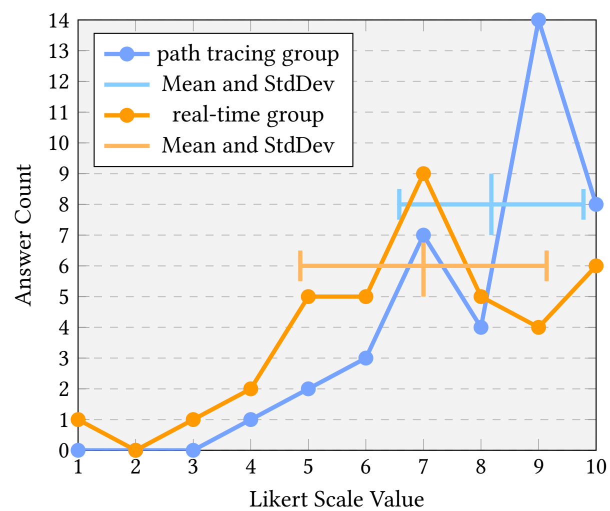

5 EVALUATION AND STUDY

A quantitative study with a heterogeneous audience has been conducted to test the visual impact of creature VFX shots created with a real-time engine compared to a traditional offline path tracer.

We aimed to test the perception of various quality aspects with the end-user audiences. Therefore, we ran the evaluation on 77 participants divided in two groups, each being shown one version and answering the same qualitative questions on a 1-10 likert scale. Among the criteria to score the quality of the VFX were: animation, deformation, digital surfaces and staging. The combined quality aspects were rated high in both groups. On average for the path tracer version across all questions 85% rated the quality criteria 7 and higher. On average for the real-time version across all questions 70% rated the quality criteria 7 and higher.

The evaluation audience were mostly students from another university majoring in media studies in the humanities and in large scale without professional background in VFX. Graph 1 asked for the quality of the staging with peaks for path tracing (Mean=7.74, StdDev=1.53) at Quality 8 and 7 for the real-time version (Mean=6.87, StdDev=1.92).

When asked to score the quality of the creature's surface we noticed weaker responses for the real-time version (Mean=7, StdDev=2.14) in scores 8 to 10 (Graph 2). This indeed makes a lot of sense since displacement was not possible to the extent of the path tracing (Mean=8,18, StdDev=1.6) version resulting in less realistic surface properties and related shading effects like Ambient Occlusion. In Graph 3 we asked for the quality of the creature's movements resulting in somewhat similar results for path tracing (Mean=8.49, StdDev=1.5) and the real-time version (Mean=8.32, StdDev=1.53). This is not surprising as the animation of the creature has not been altered. The response to the question on how natural the deformations of muscle and skin were perceived (Graph 4) resulted in somewhat similar results between path tracing (Mean=8,32, StdDev=1.49) and and the real-time version (Mean=7.95, StdDev=1.84).

Graph 3: "Would you rate the movements of the creature rigid(1) or smooth(10)"

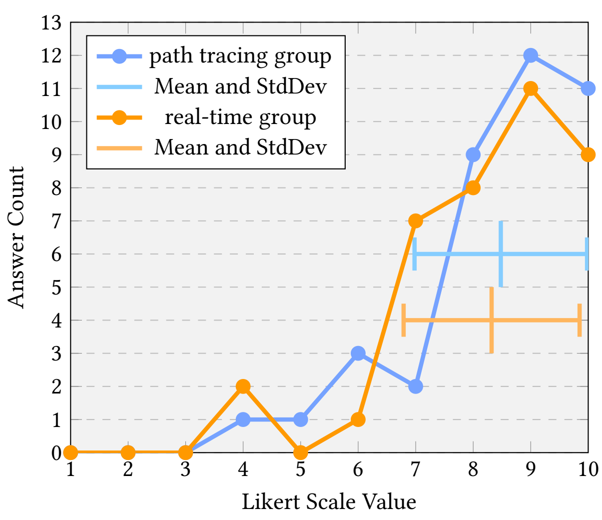

The findings of this study beg the question whether less flexibility and fewer limitations in visual fidelity are worth trading for drastically increased iteration capabilities and reduced render times. Additionally, and most importantly in regards to building awareness among future industry professionals, significant energy savings can be achieved using the real-time approach. The energy budget for the path tracing version (rendering and compositing) was 14 MWh which is within its projected limits for final year projects and emphasises the professional and efficient workflow the students created. However, the energy effort for the real-time version was estimated with just 0.3 MWh and is therefore magnitudes lower by using less than 3% of the path tracing energy budget.

Graph 4: "How natural did you find the deformations (muscles and skin) of the creature?"

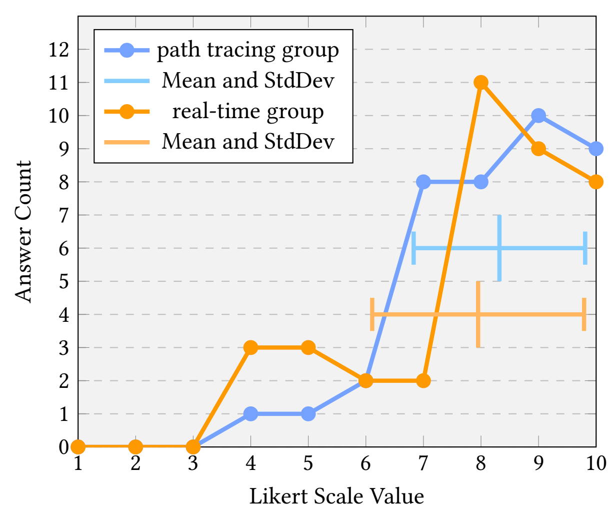

6 CURRICULUM 2.0

The teaching structure at the institution submitting this work is based on "project studies". This means that the students are mostly involved in hands-on projects during the four-and-a-half-year (two-and-a-half for post-graduates) course. This is complemented by lectures mostly from industry experts in all disciplines of media creation with a dedicated focus on Animation, VFX, and XR. Constant feedback to senior lecturers in weekly presentations and one-on-one meetings monitor and improve the progress of the projects. This teaching foundation has proven to provide a good basis for students to apply emerging and new technologies in the creative process of film making. Dedicated workshop seminars like the Set Extension Workshop (Figure 3) allow for early exploration with new technologies and research results and have for several years been focusing on teaching participating students about employing game engines as a means for energy savings[Filmakademie Baden-Württemberg 2023].

In addition to the vivid curriculum structure, the film school also runs its own research and development department. It has engaged in diverse technological developments with a dedicated focus on Digital Humans [Helzle and Goetz 2018] and Virtual Production[Spielmann et al. 2016]. Active involvement in regional and European projects with academic and industry partners has proven to be a constant driver of innovation which can support and enhance the project workflows significantly. This has been particularly evident over the last years with the institution actively partnering in three European Union funded projects with dedicated focus on Extended Reality[EMIL 2024][MAX-R 2024] and the training of associated skill-sets[PANEURAMA 2024].

The dedicated post graduate course "Technical Directing" [Filmakademie Baden-Württemberg 2024] is run by the research and development department and ensures the required technology transfer from the research department into student projects. Students engage in their own research topics mentored by the academic research staff [Maurer 2024]. With their background traditionally in computer science or associated fields they tackle the technologically challenging aspects of film, media and XR productions. Additional lectures and introduction seminars by the research staff support this effort.

Figure 3: Annual Set Extension Workshop

7 CONCLUSION AND OUTLOOK

The presented work should be understood as a starting point for a discussion about quality demands and potential future creation workflows as well as sustainability options in our industry. The study emphasises how close both approaches are in terms of quality and makes room for broader discussion on future pipelines and demands. What was presented to a clearly arranged student project can't easily be applied at large scale - as of now. However, with the increased strive for cost efficient and sustainable solutions, the question arises: Will it be sufficient for future artists and technicians to exclusively target highest visual fidelity or will the focus of creative work also shift towards asset optimization and achieving comparable results with drastically reduced compute power? The authors are aware that studios and streaming providers dictate the technical conditions to their vendors, making it even more difficult to adapt to new creation processes. However, if there is room for change, this work should be understood as a supplemental tool. Various attempts have been made to transfer and convert assets for reuse in other purposes. Most of these solutions are still in an early stage, often developed in-house while demanding for a high level of maintenance [Trottnow et al. 2020]. It might be worth reconsidering asset creation from scratch with the opportunity of close to real-time systems in mind. In any case, the tendency to use near real-time technology in movie production is increasing with high-speed. The proposed near real-time rendering creation pipeline (Blender EEVEE) is close in terms of quality compared to the traditional path tracing workflow (Arnold) and enables faster iteration on lighting and look development, ultimately resulting in faster delivery for the compositing stage. As an educator we see an increased effort in adjusting our curriculum and creation pipelines toward these new opportunities to prepare the next generation of creative talent for a flexible and sustainable future.

ACKNOWLEDGMENTS

The authors would like to thank Tonio Freitag, Maximilian Schubert, Lukas Löffler, Pascal Schober for their ongoing support in realizing this report.

REFERENCES

ALBERT. 2022. *A Step-Change on Sustainability: BAFTA albert annual review 2022.* Retrieved Oct 8, 2024 from https://shorturl.at/r9AA4

Michele Bousquet. 2022. Animation Field Guide. https://www.unrealengine.com/en-US/animation-field-guide

EMIL. 2024. *EMIL - European Mediea and Immersion Lab.* Retrieved May 15, 2024 from https://emil-xr.eu

Green Screen Interreg Europe Film Paris Region. 2022. *Virtual Production: a study in its environmental impact.* Retrieved August 7, 2024 from https://shorturl.at/PVzyu

Animationsinstitut Filmakademie Baden-Württemberg. 2023. *Research, Innovation, and the Endless Possibilities of Virtual Production.* Retrieved October 8, 2024 from https://shorturl.at/6fZBg

Animationsinstitut Filmakademie Baden-Württemberg. 2024. *Post Graduate Course Technical Directing.* Retrieved August 7, 2024 from https://animationsinstitut.de/en/studies/study-specializations/technical-directing

Volker Helzle and Kai Goetz. 2018. Digital Albert Einstein, a case study. In *ACM SIGGRAPH 2018 Talks* (Vancouver, British Columbia, Canada) *(SIGGRAPH '18)*. Association for Computing Machinery, New York, NY, USA, Article 57, 2 pages. https://doi.org/10.1145/3214745.3214782

Volker Helzle, Simon Spielmann, and Jonas Trottnow. 2022. Green Screens, Green Pixels and Green Shooting. In *ACM SIGGRAPH 2022 Talks* (Vancouver, BC, Canada) *(SIGGRAPH '22)*. Association for Computing Machinery, New York, NY, USA, Article 1, 2 pages. https://doi.org/10.1145/3532836.3536235

Vincent Maurer. 2024. Capturing Light with Robots: A Novel Workflow for Reproducing Realistic Lens Flares. In *SIGGRAPH Asia 2024 Technical Communications* (Tokyo, Japan) *(SA '24)*. Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/3681758.3697995

MAX-R. 2024. *MAX-R, Mixed Augmented and Extended Reality Media Pipeline.* Retrieved August 7, 2024 from https://max-r.eu/

Baptiste Nicolet, Fabrice Rousselle, Jan Novák, Alexander Keller, Wenzel Jakob, and Thomas Müller. 2023. Recursive Control Variates for Inverse Rendering. *Transactions on Graphics (Proceedings of SIGGRAPH)* 42, 4 (Aug. 2023). https://doi.org/10.1145/3592139

Yaobin Ouyang, Shiqiu Liu, Markus Kettunen, Matt Pharr, and Jacopo Pantaleoni. 2021. ReSTIR GI: Path Resampling for Real-Time Path Tracing. *Computer Graphics Forum* (2021). https://doi.org/10.1111/cgf.14378

PANEURAMA. 2024. *PANEURAMA - Innovation Alliance of Animation, Media Arts & Cultural Creative Industriese.* Retrieved August 7, 2024 from https://www.paneurama.eu/

Simon Spielmann, Andreas Schuster, Kai Götz, and Volker Helzle. 2016. VPET: a toolset for collaborative virtual filmmaking. In *SIGGRAPH ASIA 2016 Technical Briefs* (Macau) *(SA '16)*. Association for Computing Machinery, New York, NY, USA, Article 29, 4 pages. https://doi.org/10.1145/3005358.3005370

Jonas Trottnow, William Greenly, Christian Shaw, Sam Hudson, Volker Helzle, Henry Vera, and Dan Ring. 2020. SAUCE: Asset Libraries of the Future. In *The Digital Production Symposium* (Virtual Event, USA) *(DigiPro '20)*. Association for Computing Machinery, New York, NY, USA, Article 3, 5 pages. https://doi.org/10.1145/3403736.3403941