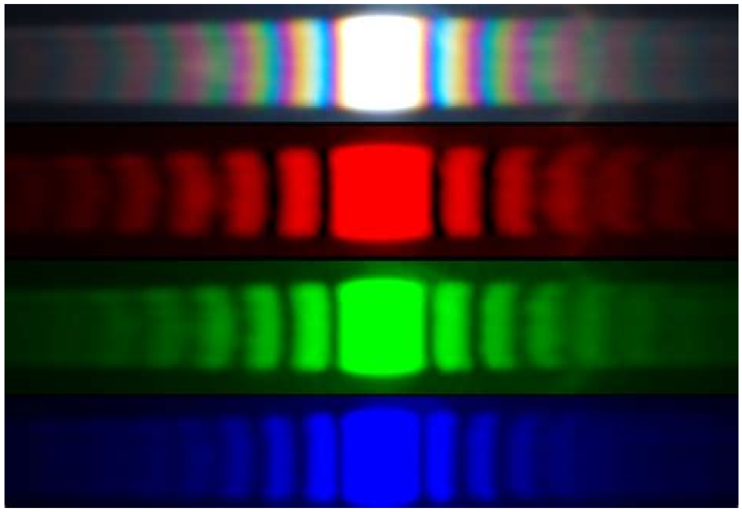

WAVE OPTICS IN COMPUTER GRAPHICS

CARMEN PARK

AUGUST 16, 2025

Abstract

This paper looks at 70 years of wave optics research with a focus on computer graphics. From theoretical physics and mathematics, using Maxwells equations to derive a more practical equation describing electromagnetic discontinuities, to more computable applications and approximations that follow. This is a comprehensive view of wave optics, with a focus on diffraction while also evaluating optimization strategies used.

• Question: Can wave optics be accurately simulated in real-time computer graphics without prohibitive computational costs?
• Scope: looking at hardware vs. algorithmic approaches as well as accessibility challenges for resource-constrained systems.
• Methodology: a comparative analysis of 10 key technical papers (1951-2024) with focuses in mathematical, graphics and optimization findings.

• Analysis: theoretical and approximated findings (Kline, Stam), hybrid approximations developed (Lindsay, Belcour), hardware accelerated sources included (Clausen, Pediredla), questions about accessibility.

• Conclusions: parallel computing has enabled real-time wave optics (30x speedups), thin-film interference models show promise for lower-end hardware, and inquiries in place for future work in algorithmic alternatives to GPU dependence.

Contents

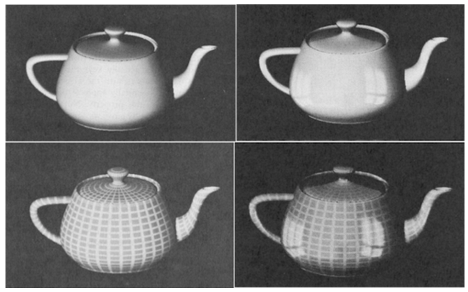

Figure 1: Enter Caption

1 Introduction

1.1 Historical Context

3 dimensional computer graphics emerged in the late 1970s and 1980s with *Utah Tepot* in 1975 as the first widely adopted 3D test model, films like *Tron* (1982) featuring CGI, and studios like Pixar, established in 1985, pioneering techniques for animation. Graphics transitioned from military and academic fields into ubiquitous technologies in entertainment, design, and scientific visualization.

For most of this history, light transport was modeled through classical geometric optics—approximating light as discrete rays interacting with rasterized surfaces. This paradigm, while computationally efficient, fundamentally ignored wave-optical phenomena like diffraction and interference. The Utah Teapot itself, rendered using simple Phong shading, epitomizes this ray-centric approach that dominated graphics pipelines until the 1990s, when Stam's diffraction shaders began introducing wave-optical considerations for specific effects.

This was not for lack of knowledge of light as the propagation of electromagnetic radiation. The wave nature of light as an EM phenomenon had been firmly established in mathematical physics long before computer graphics existed—Maxwell's equations (1865) formalized light's wave behavior, while Fresnel's early 19th-century diffraction studies provided quantitative models. By the time Kline published his asymptotic solutions to Maxwell's equations (1951), the scientific community had already accepted electromagnetic wave theory for nearly a century, with applications ranging from radio to microscopy. These physical principles however, remained largely disjoint from early computer graphics and often still are today, as computational constraints lead to modeling the more simplified geometric optics. The Utah Teapot (1975, Figure 1)—rendered with basic ray tracing, to epitomize this disconnect. While physicists understood light as continuous waves, graphics pioneers modeled it as discrete rays for practicality.

1.2 Recent Exploration

We've made significant progress since the early days of computer graphics. Researchers can now simulate realistic wave optics effects like diffraction and interference in real-time, thanks to modern GPU power and smarter algorithms. Recent work by Clausen (2024) and Pediredla (2023) shows what's possible with today's hardware - effects that would have been impossible to render interactively just a decade ago.

But there are still challenges. These advanced techniques require powerful graphics cards, as Konyaev's 2011 benchmarks clearly showed. While impressive, this creates a barrier for many potential users. Different applications also have different needs - medical imaging demands extreme accuracy (as seen in Dhillon's 2014 biological models), while entertainment graphics often prioritizes speed over perfect physics.

The field now faces important questions about accessibility and practicality. How can we make these techniques work on less powerful devices? What clever simplifications can maintain visual quality while reducing computation? These aren't just technical problems - they'll determine who actually gets to use these advanced rendering methods in their work.

2 Findings

2.1 Mathematical Findings

"Electromagnetic Theory and Geometrical Optics" (Kline, 1962) published by New York University, is a mathematical paper that provides a thorough history of theories of light and applications of these theories. Kline shows that his work (Kline, 1951), and others are close to approximating the electromagnetic theories of light to geometrical optics, but far from complete models. He develops this argument through reintroducing relevant historical discoveries, as well as relating many discontinuity cases to solvable nonlinear differential equations and partial differential equations like the Schrodinger equation and the eikonal equation. He connects this relevant research to other related fields as well by showing how electromagnetic wave solutions can apply to magnetohydrodynamics, or sonal and shallow wave solutions. Kline mentions some open problems and potential critiques to his proposed work along the way. He closes by stating the next step for diffraction problems is to derive the form of the asymptotic series solution that is valid in diffraction regions.

"An asymptotic solution of Maxwell's Equations" (Kline, 1951) , published by the communications for pure and applied math journal claims that all discontinuities in light (for example diffraction, reflection or refraction) can be described by the pulse solution of Maxwell's equations added to the sum of the infinite series of the “jumps” in discontinuity. This is derived step by step in the first 12-13 pages from Maxwell’s equations and Duhamel’s principle with text to describe each step and eventual findings. The purpose of these findings is to more broadly state that given any discontinuity in light, what follows is one unified way to calculate the effects. Kline also sets out to give the reader tools to access ways of finding the discontinuiy conditions, and ordinary differential equations to solve them. Kline uses an academic tone to address his target audience of mathemeticians, physicists, or individuals interested in the connection between geometric optics and Maxwell’s equations. The article is useful in that the full derivation can more accurately describe lightwave phenomena than most traditional “ray” based optical systems.

2.2 Graphics Findings

“Diffraction Shaders” (Stam, 1999), a computer graphics paper published in SIGGRAPH 99, illustrates using Helmholtz’s wave equation and Fourier transform of spectral densities to create a diffraction shader. Stam similarly describes step by step their mathematical derivation, while also using bidirectional reflection distribution functions (BRDF, a commonly used function for graphics shaders) to frame how their findings can be applied. Using the wave equation, Stam modifies a more traditional ray based BRDF to satisfy a more complex wave based BRDF. Their use of fourier transform separates the respective frequencies of light to achieve the desired results. To show this, they use a CD as an example. How they chose to model the reflective side of a CD is included so that the reader knows what makes up the rendering surface. Stam concludes with a final shader equation, and computer generated images (using Maya) testing their results, in particular, on the CD which successfully displayed a full spectrum of diffractive induced colors.

“Physically Based Real Time Diffraction Using Spherical Harmonics” (Lindsay & Agu, 2006) is another approach by Agu and Lindsay from the Worcester Polytechnic Institute in Massachusetts. They argue that an optimized real-time renderer using spherical coordinates maintains quality without losing frame rate. They reviewed existing methods of wave optics, mentioned their contributions to real-time iridescence modeling, and described the differences and difficulties that come with diffraction, as well as other researchers who have tackled modeling diffraction effects with spherical coordinates. Their choice of spherical harmonics stems from their critique of the point-source calculations at the root of previous attempts. Spherical coordinates account for the entire range of potential lighting, but they also greatly increase the number of calculations needed. As a result, a set of pre-computations of the scene are necessary before real-time rendering is possible. Agu and Lindsay acknowledge this as a drawback of their method. It was also difficult to tell whether or not they were able to account for the spectral nature of diffraction by the lack of color in their supplementary data.

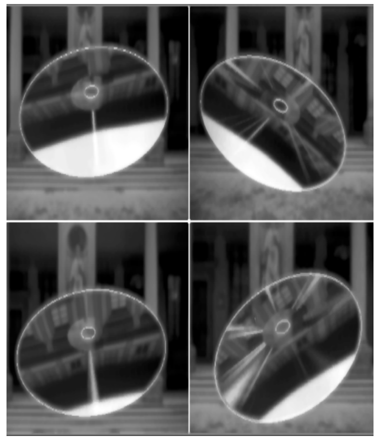

Figure 2: (Lindsay & Agu, 2006)

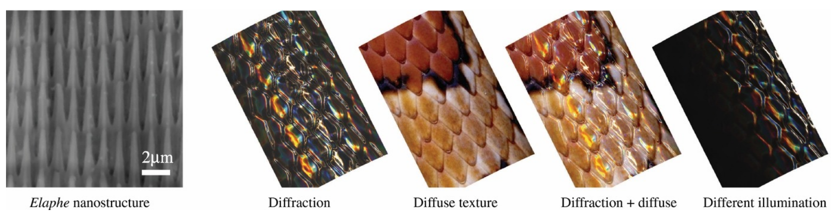

Figure 3: (Dhillon et al., 2014)

"Interactive Diffraction from Biological Nanostructures" (Dhillon et al., 2014), published by the Computer Graphics Forum, extends Stam's research to render diffraction effects of biological nanostructures, they use the example of Elaphe and Xenopeltis snakes, due to their skins spectacular display of structural colors. Dhillon et al. utilize atomic force microscopy (AFM, a high-resolution technique that measures surface topography at the atomic/nanometer scale) to capture highly accurate surface information. Their diffraction rendering differs in that they precompute fourier transform values, and are using OpenGL and 2 (GLSL) fragment shaders, while still using Stam's derived equations for diffraction approach. To show their findings, they set up an experimental setup that sends non-monochromatic light off of Elaphe snake skin. They rendered the same images with light reflecting off the modeled Elaphe skin and using diffraction shaders in OpenGL. Their findings were impressively similar and believe the actual images exhibit more complex diffraction patterns due to the irregularities of biological structures.

"A Practical Extension to Microfacet Theory for the Modeling of Varying Irredescence" (Belcour & Barla, 2017) is a paper and presentation at SIGGRAPH

2017 where the authors acknowledge the complexity of rendering irredesence on surfaces. In the 2017 presentation, Belcour introduces goniochromism, referring to the property of a material to exhibit angle-dependent color shifts when viewed or illuminated from different directions. Their solution to model this type of discontinuity was to apply a thin film interference to a surface. The use of wave optics calculations presented include a spectral BRDF similar to previous sources, with the complex valued taylor series expansion of the electromagentic waves at the film interference and fourier transform separate frequencies. They also utilized Parsevals theorem to approximate from fourier space to a finite sum that is more easily computable. They integrate this with a microfacet model. In order to make this compatible with real-time rendering. This also led to pre-calculations, and preintegrations necessary before rendering.

2.3 Optimization Findings

"A Two-Scale Microfacet Reflectance Model Combining Reflection and Diffraction" (Holzschuch & Pacanowski, 2017), published in ACM Transactions on Graphics, proposes a new reflectance model that accounts for both large-scale surface roughness and fine-scale diffraction effects. The authors argue that traditional microfacet models like Cook–Torrance are insufficient for accurately reproducing measured reflectance when surface features approach the wavelength of visible light. To address this, they introduce a two-scale approach: at the macroscale, they use standard microfacet distributions to model specular reflection, while at the microscale, they integrate a wave-optical diffraction component derived from the convolution of the surface's height-field spectrum and the Fresnel kernel. The model accounts for spectral shifts and angular spreads observed in real-world surfaces, especially metals and coated materials. Holzschuch and Pacanowski validate their model by comparing it to empirical reflectance data, showing significantly improved fits in cases where classical models fall short. They note that their model remains compatible with existing rendering frameworks but requires precomputed data for efficiency. Their work bridges the gap between geometric and wave optics and offers a more complete representation of light behavior on rough surfaces.

In the paper "Optimized virtual optical waveguides enhance light throughput in scattering media" (Pediredla et al., 2023) published by Nature Communications, proposes a physically accurate simulator to focus on optimizing virtual optical waveguides. Their goal was to optimize ultrasound frequencies by using gradient-index optics (GRIN) in order to recycle, and maximize the light passing through a medium, like in experiments on a human bladder. Similar to other opimizations, Pedriredla et al. demonstrates that precomputed values, in this instance, focal configurations, can enhance light capture by recycling scattered photons. This study is primarily focused on biomedical optics and photonics imaging, it has strong relevance to graphics from its computational methods and light transport modeling. Its findings also raise questions towards future

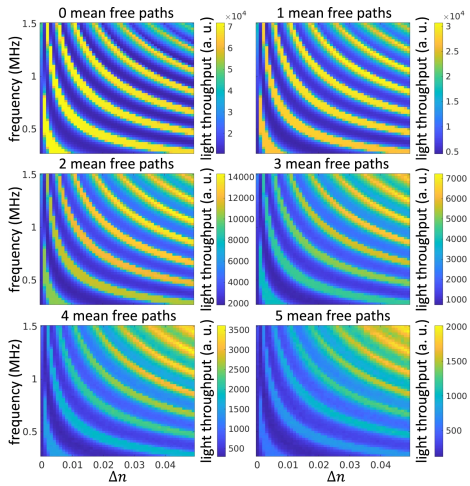

Figure 4: (Pediredla et al., 2023)

optimizations in computer imaging.

“A Practical Real-Time Model for Diffraction on Rough Surfaces” (Clausen et al., 2024), published in JCGT, presents an optimized RGB diffraction model for real-time rendering of wave-optical effects on rough conductive ‘metal’ materials. They build on Clausen et al.’s (2023) spectral work. The authors address two gaps: speckle pattern omission in prior models and computational inefficiency of spectral approaches. They introduce a model wavelength shift function to simulate reddish/bluish diffraction shifts. Procedural speckle patterns using 4D simplex noise are fitted to multivariate normal distributions from macro photos of rough surfaces. Their main innovations include practical real-time performance and VR optimizations. Similar to previous models, there are shader executions that must take place before running this model. Some limitations also include artifacts on very smooth surfaces due to cosine-based covariance approximations. The work bridges wave optics and real-time CG, with nanostructure-aware rendering and interactive diffraction.

“Computer Simulation of Optical Wave Propagation with the Use of Parallel Programming” (Konyaev, Tartakovskii, & Filimonov, 2011), published in Atmospheric and Oceanic Optics, presents two parallel programming strategies for accelerating scalar wave propagation simulations. The authors model optical wave behavior using the parabolic wave equation and solve it using a split-step Fourier method. To address the high computational demand, they implement two approaches: one using OpenMP and Intel’s MKL for CPU multicore acceleration, and another using CUDA for GPU-based processing. Both rely on FFT operations for efficient numerical integration of diffraction terms.

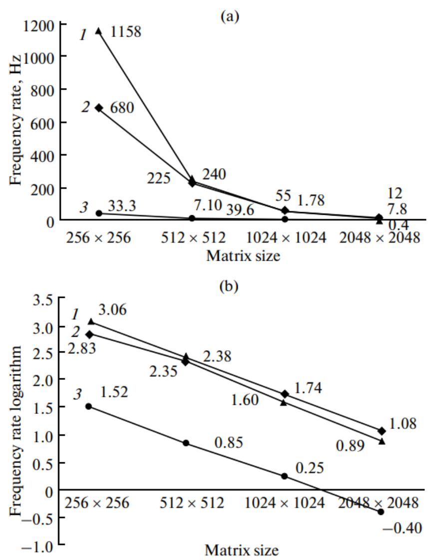

Figure 5: The number of diffraction problem solutions per second (frame rate) as a function of matrix size on a linear scale (a) and logarithmic scale (b): MKL computations (1), CUDA (2), and sequential computations (3)(Konyaev, Tartakovskii, & Filimonov, 2011)

Their results show that OpenMP performs best for grid sizes up to 1024×1024, while CUDA shows superior performance at larger scales. They compare both methods against a baseline using FFTW on a single CPU thread, observing speedups of up to 30× or more, depending on the system configuration. The authors conclude that the choice between CPU or GPU acceleration depends on both the problem size and hardware availability. This work is targeted at optical scientists and engineers aiming to simulate wave propagation over large distances with realistic computation times, and contributes to the broader trend of parallelized wave optics simulation for scalable modeling.

3 Analysis

The collection of articles gathered traces the evolution from fundamental physical wave phenomena to the integration of wave optics in computer graphics, followed by ongoing research and optimization efforts. Kline (1951, 1962) distills pure electromagnetic theory into its essential mathematical components, establishing a rigorous foundation for understanding light discontinuities. Stam (1999) then translates these principles into practical computer graphics applications, introducing diffraction shaders based on similar mathematical derivations. Stam's work was so foundational that it is cited in nearly all subsequent papers. Lindsay and Agu (2006) further refine this approach, implementing diffraction in a spherical coordinate system to achieve real-time rendering, though their model leaves spectral accuracy questions unresolved (see fig 1). Later works like Dhillon et al. (2014) validate these models against biological nanostructures, using atomic-force microscopy to bridge theory and nature's complexity, while Belcour and Barla (2017) introduce thin-film interference to handle iridescence. Recent advancements by Clausen et al. (2024) and Pediredla et al. (2023) push to further optimize for real-time constraints and biomedical applications, respectively. This all goes to demonstrate how far the field has evolved from Kline's original asymptotic solutions.

Underlying this progression is a central question, Can wave optics be accurately approximated for practical use, or are sacrifices in precision necessary for performance? The dialogue between these sources reveals both the challenges and innovations in trying to answer this question. SImmilarly, in many of the computing papers, while some include hardware specs used during testing, others do not, some more fundamental questions related to accessable computing fall flat. Kline's asymptotic solutions, while mathematically elegant, are computationally time consuming.They are not suited for real-time applications. Stam's Fourier-based approximations make diffraction tractable in rendering pipelines but simplify the underlying physics, many diffraction efects are not captured with their methods. Lindsay and Agu's spherical harmonics improve efficiency, yet their reliance on precomputation and lack of spectral fidelity highlight the trade-offs inherent in real-time systems.Lindsay and Agu demonstrate that while diffraction computations can be performed rapidly for simple scenes, the calculations become significantly more demanding as scene complexity increases due to amplified light scattering and ray divergence.

The structure of the surfaces being rendered, overlaped with use of diffraction in rendering is at the root of this computational challenge. The most common diffraction seen in reflectance is a result of minute descrepencies from microscopic surface irregularities perturbing incident light waves.

Later works, such as Belcour and Barla (2017) and Holzschuch and Pacanowski (2017), attempt to reconcile these compromises by introducing hybrid models that combine micro-facet theory with wave-optical effects. Their approaches demonstrate that while perfect physical accuracy remains elusive, intelligently structured approximations can deliver visually plausible results without prohibitive computational costs. Meanwhile, Dhillon et al. (2014) push the boundaries by applying these techniques to biological nanostructures, proving that even nature's irregular diffractive surfaces can be simulated, but not without significant data preprocessing. AMT remains largely inaccessible due to its technical requirements and substantial resource demands. While it was adopted in Dhillon et al. it is beyond reach to most potential adopters.

The most recent advancements, like Clausen et al. (2024) and Pediredla et al. (2023), suggest a promising direction by leveraging hardware-aware optimizations. As demonstrated in Figure 5, parallel computing architectures (Konyaev et al., 2011) enable order-of-magnitude speedups for wave-optical simulations—a capability these newer works exploit to narrow the gap between theory and practice. Clausen's procedural speckle patterns and Pediredla's virtual waveguides exemplify how such computational gains permit creative approximations: their approaches remain grounded in physical principles while achieving real-time performance through strategic optimizations, proving wave optics need not be fully abandoned for interactive applications.

Looking ahead, the field appears poised for breakthroughs in adaptive fidelity systems, where real-time rendering dynamically switches between wave and geometric optics based on perceptual or computational demands. Technologies such as GPU-accelerated wave solvers (inspired by Konyaev et al.'s 2011 parallelization methods) and machine learning driven precomputation could further blur the line between accuracy and efficiency. Machine learning however has its own computational costs and ethical questions. Overall however, diffraction and wave optics are on the forefront of new discoveries and realizations.

3.1 Conclusion

In the progression of these sources, it is apparent how modern computing advancements have systematically transformed wave optics from theoretical construct to a realistic rendering tool. Where Kline's (1951, 1962) foundational mathematical formulations were computationally unusable for real-time applications, Stam's (1999) Fourier-based approximations first demonstrated how strategic simplifications could make diffraction shaders feasible. A critical turning point happens with parallel computing architectures, seen in Konyaev et al. (2011) in Figure 1, showing 30× speedups through GPU/MKL implementations. Electromagnetic analysis has evolved with technological advancements, from iterative hardware-aware optimizations, to interactive wave-optical effects. This hardware revolution enabled subsequent innovations, yet raises crucial questions about accessibility: Are there algorithmic approaches less dependent on parallel processing that could achieve similar results? Could clever mathematical approximations (like Belcour and Barla's (2017) thin-film interference model) make wave optics rendering usable on less advanced hardware?

Ultimately, the conversation among these sources reflects a broader truth in computer graphics, the best solutions often lie in the balance. Rigorous theory guides innovation, but practical constraints shape its implementation. If Kline (1962) were to witness today's real-time diffraction models, he might critique their simplifications—yet he would likely congradulate the ingenuity that has brought his asymptotic solutions closer to real-world applicability (the asymptotic solutions themselves are still an unsolved question). The path forward will require continued collaboration between physicists, mathematicians, and graphics researchers, ensuring that each step toward optimization remains firmly rooted in the underlying science of light.

References

[1] Belcour, L., & Barla, P. (2017). A practical extension to microfacet theory for the modeling of varying iridescence. *ACM Transactions on Graphics*, *36*(4), Article 65. https://doi.org/10.1145/3072959.3073620

[2] Clausen, O., Misiak, M., Fuhurmann, A., Marroquim, R., & Latoschik, M. E. (2024). A practical real-time model for diffraction on rough surfaces. *Journal of Computer Graphics Techniques*, *13*(1), 1–27. https://jcgt.org/published/0013/01/01/

[3] Dhillon, D. S., Teyssier, J., Single, M., Gaponenko, I., Milinkovitch, M. C., & Zwicker, M. (2014). Interactive diffraction from biological nanostructures. *Computer Graphics Forum*, *33*(8), 177–188. https://doi.org/10.1111/cgf.12425

[4] Garcia, P. (2014). Utah teapots. *Pablo Garcia Studio.* https://www.pablogarcia.org/utah-teapots/

[5] Holzschuch, N., & Pacanowski, R. (2017). A two-scale microfacet reflectance model combining reflection and diffraction. *ACM Transactions on Graphics*, *36*(4), Article 12. https://doi.org/10.1145/3072959.3073621

[6] Kline, M. (1951). An asymptotic solution of Maxwell's equations. *Communications on Pure and Applied Mathematics*, *4*(2–3), 225–262. https://doi.org/10.1002/cpa.3160040203

[7] Kline, M. (1962). Electromagnetic theory and geometrical optics. *Research Report No. EM-171*, Division of Electromagnetic Research. https://archive.org/details/electromagnetict00klin

[8] Konyaev, P. A., Tartakovskii, E. A., & Filimonov, G. A. (2011). Computer simulation of optical wave propagation with the use of parallel programming. *Atmospheric and Oceanic Optics*, *24*(5), 425–431. https://doi.org/10.1134/S1024856011050095

[9] Lindsay, C., & Agu, E. (2006). Physically-based real-time diffraction using spherical harmonics. *Lecture Notes in Computer Science*, 505–517. https://doi.org/10.1007/11919476$_5$1

[10] Pediredla, A., Scopelliti, M. G., Narasimhan, S., Chamanzar, M., & Gkioulekas, I. (2023). Optimized virtual optical waveguides enhance light throughput in scattering media. *Nature Communications*, *14*, Article 5681. https://doi.org/10.1038/s41467-023-40864-z

[11] Stam, J. (1999). Diffraction shaders. *Proceedings of SIGGRAPH 1999*, 101–110. https://www.dgp.toronto.edu/public_user/stam/reality/Research/pdf/diff.pdf